

Isilon OneFS

Version 7.2.0

CLI Administration Guide

Copyright © 2013-2015 EMC Corporation. All rights reserved. Published in USA.

Published July, 2015

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Chapter 1	Introduction to this guide	25
	About this guide.....	26
	Isilon scale-out NAS overview.....	26
	Where to go for support.....	26
Chapter 2	Isilon scale-out NAS	27
	OneFS storage architecture.....	28
	Isilon node components.....	28
	Internal and external networks.....	29
	Isilon cluster.....	29
	Cluster administration.....	29
	Quorum.....	29
	Splitting and merging.....	30
	Storage pools.....	31
	IP address pools.....	31
	The OneFS operating system.....	31
	Data-access protocols.....	32
	Identity management and access control.....	32
	Structure of the file system.....	33
	Data layout.....	33
	Writing files.....	34
	Reading files.....	34
	Metadata layout.....	34
	Locks and concurrency.....	35
	Striping.....	35
	Data protection overview.....	35
	N+M data protection.....	36
	Data mirroring.....	37
	The file system journal.....	37
	Virtual hot spare.....	37
	Balancing protection with storage space.....	37
	VMware integration.....	37
	Software modules.....	38
Chapter 3	Introduction to the OneFS command-line interface	39
	OneFS command-line interface overview.....	40
	Syntax diagrams.....	40
	Universal options.....	41
	Command-line interface privileges.....	41
	SmartLock compliance command permissions.....	45
	OneFS time values.....	48
Chapter 4	General cluster administration	49
	General cluster administration overview.....	50
	User interfaces.....	50
	Connecting to the cluster.....	51
	Log in to the web administration interface.....	51

Open an SSH connection to a cluster.....	51
Licensing.....	51
License status.....	52
License configuration.....	54
Activate a license through the command-line interface.....	55
View license information.....	55
Unconfigure a license.....	55
Certificates.....	56
Replace or renew the SSL certificate.....	56
Verify an SSL certificate update.....	57
Self-signed SSL certificate data example.....	58
Cluster identity.....	58
Set the cluster name	58
Cluster contact information.....	59
Specify contact information	59
Cluster date and time.....	60
Set the cluster date and time.....	60
Specify an NTP time server.....	61
SMTP email settings.....	61
Configure SMTP email settings	61
View SMTP email settings.....	62
Configuring the cluster join mode.....	62
Specify the cluster join mode	62
File system settings.....	63
Specify the cluster character encoding.....	63
Enable or disable access time tracking	64
Cluster monitoring.....	64
Monitor the cluster.....	65
View node status.....	65
Monitoring cluster hardware.....	65
View node hardware status.....	65
Chassis and drive states.....	65
Check battery status.....	68
SNMP monitoring.....	68
Events and notifications.....	72
Coalesced events.....	72
Viewing event information.....	74
Responding to events.....	76
Managing event notification settings.....	78
Managing event notification rules.....	80
Cluster maintenance.....	82
Replacing node components.....	82
Upgrading node components.....	82
Managing drive firmware.....	82
Managing cluster nodes.....	87
Upgrading OneFS.....	89
Remote support.....	90
Remote support using SupportIQ.....	90
Remote support using ESRS Gateway.....	93
Cluster administration commands.....	95
isi config.....	96
isi email.....	100
isi email list.....	101
isi exttools.....	101
isi license activate.....	101
isi license status.....	101

isi license unconfigure.....	102
isi perfstat.....	102
isi pkg create.....	102
isi pkg delete.....	102
isi pkg info.....	103
isi pkg install.....	103
isi remotesupport connectemc modify.....	104
isi remotesupport connectemc view.....	105
isi services.....	105
isi set.....	106
isi snmp.....	109
isi snmp list.....	110
isi statistics client.....	110
isi statistics describe.....	116
isi statistics drive.....	116
isi statistics heat.....	118
isi statistics history.....	121
isi statistics list all.....	123
isi statistics list classes.....	124
isi statistics list events.....	124
isi statistics list nodes.....	125
isi statistics list nooutput.....	125
isi statistics list operations.....	126
isi statistics list orderby.....	126
isi statistics list output.....	127
isi statistics list protocols.....	127
isi statistics list stats.....	128
isi statistics list totalby.....	128
isi statistics protocol.....	129
isi statistics pstat.....	134
isi statistics query.....	136
isi statistics system.....	137
isi status.....	139
isi update.....	139
isi version.....	140
isi_for_array.....	141
isi get.....	143
isi_gather_info.....	144
Event commands.....	149
isi events cancel.....	149
isi events list.....	149
isi events notifications create.....	151
isi events notifications modify.....	152
isi events notifications delete.....	154
isi events notifications list.....	154
isi events quiet.....	154
isi events sendtest.....	154
isi events settings list.....	155
isi events settings set.....	155
isi events show.....	155
isi events unquiet.....	156
Hardware commands.....	156
isi batterystatus.....	156
isi devices.....	157
isi servicelight status.....	158
isi servicelight off.....	158

	isi servicelight on	159
	isi drivefirmware status	159
	isi firmware package	160
	isi firmware status	160
	isi firmware update	162
	isi readonly off	163
	isi readonly on	164
	isi readonly show	165
Chapter 5	Access zones	167
	Access zones overview	168
	Access zone base directory rules	168
	Access zones best practices	169
	Access zone limits	169
	Quality of service	170
	Managing access zones	170
	Create an access zone	170
	Associate an IP address pool with an access zone	171
	Modify an access zone	171
	Add an authentication provider to an access zone	171
	Remove an authentication provider from an access zone	172
	Delete an access zone	172
	Access zone commands	172
	isi zone restrictions create	172
	isi zone restrictions delete	173
	isi zone restrictions list	174
	isi zone zones create	174
	isi zone zones delete	177
	isi zone zones list	177
	isi zone zones modify	178
	isi zone zones view	184
Chapter 6	Authentication and access control	185
	Authentication and access control overview	186
	Role-based access	186
	Roles and privileges	186
	Data backup and restore privileges	197
	User permissions utility	198
	Authentication	198
	Supported authentication providers	198
	Authentication provider features	199
	Kerberos authentication	199
	LDAP	200
	Active Directory	200
	NIS	201
	File provider	201
	Local provider	202
	Data access control	202
	Authorization	202
	SMB	203
	NFS	203
	Mixed-permission environments	204
	Managing roles	205
	View roles	205

View privileges.....	205
Create and modify a custom role.....	206
Delete a custom role.....	206
Managing authentication providers.....	206
Managing LDAP providers.....	207
Managing Active Directory providers.....	208
Managing NIS providers.....	209
Managing file providers.....	210
Managing local users and groups.....	214
Managing MIT Kerberos authentication.....	218
Managing access permissions.....	225
View expected user permissions.....	226
Configure access management settings.....	227
Modify ACL policy settings.....	227
Update cluster permissions.....	227
Authentication and access control commands.....	228
isi auth access.....	228
isi auth ads create.....	229
isi auth ads delete.....	232
isi auth ads list.....	233
isi auth ads modify.....	234
isi auth ads spn check.....	237
isi auth ads spn create.....	238
isi auth ads spn delete.....	238
isi auth ads spn list.....	239
isi auth ads trusts controllers list.....	240
isi auth ads trusts list.....	240
isi auth ads trusts view.....	241
isi auth ads view.....	241
isi auth error.....	241
isi auth file create.....	242
isi auth file delete.....	245
isi auth file list.....	245
isi auth file modify.....	246
isi auth file view.....	252
isi auth groups create.....	253
isi auth groups delete.....	254
isi auth groups flush.....	254
isi auth groups list.....	255
isi auth groups modify.....	255
isi auth groups members list.....	257
isi auth groups view.....	258
isi auth id.....	258
isi auth krb5 realm create.....	258
isi auth krb5 realm delete.....	259
isi auth krb5 realm modify.....	259
isi auth krb5 realm list.....	260
isi auth krb5 realm view.....	260
isi auth krb5 create.....	260
isi auth krb5 delete.....	261
isi auth krb5 list.....	261
isi auth krb5 view.....	262
isi auth krb5 domain create.....	262
isi auth krb5 domain delete.....	262
isi auth krb5 domain modify.....	263
isi auth krb5 domain list.....	263

isi auth krb5 domain view.....	263
isi auth krb5 spn create.....	264
isi auth krb5 spn delete.....	264
isi auth krb5 spn check.....	264
isi auth krb5 spn fix.....	265
isi auth krb5 spn import.....	265
isi auth krb5 spn list.....	265
isi auth settings krb5 modify.....	266
isi auth settings krb5 view.....	266
isi auth ldap create.....	266
isi auth ldap delete.....	273
isi auth ldap list.....	274
isi auth ldap modify.....	274
isi auth ldap view.....	284
isi auth local list.....	284
isi auth local view.....	284
isi auth local modify.....	285
isi auth log-level.....	286
isi auth mapping delete.....	287
isi auth mapping dump.....	288
isi auth mapping flush.....	289
isi auth mapping idrange.....	289
isi auth mapping import.....	291
isi auth mapping view.....	291
isi auth mapping modify.....	292
isi auth mapping create.....	293
isi auth mapping token.....	294
isi auth netgroups list.....	294
isi auth nis create.....	295
isi auth nis delete.....	298
isi auth nis list.....	298
isi auth nis modify.....	299
isi auth nis view.....	304
isi auth privileges.....	304
isi auth refresh.....	305
isi auth roles create.....	305
isi auth roles delete.....	305
isi auth roles list.....	306
isi auth roles members list.....	306
isi auth roles modify.....	307
isi auth roles privileges list.....	309
isi auth roles view.....	309
isi auth settings global modify.....	310
isi auth settings global view.....	312
isi auth status.....	312
isi auth users create.....	313
isi auth users delete.....	315
isi auth users flush.....	315
isi auth users list.....	316
isi auth users modify.....	317
isi auth users view.....	319

Chapter 7	Identity management	321
	Identity management overview.....	322
	Identity types.....	322

	Access tokens.....	323
	Access token generation.....	324
	ID mapping.....	324
	User mapping.....	326
	On-disk identity.....	328
	Managing ID mappings.....	329
	Create an identity mapping.....	329
	Modify an identity mapping.....	330
	Delete an identity mapping.....	330
	View an identity mapping.....	330
	Flush the identity mapping cache.....	331
	View a user token.....	331
	Configure identity mapping settings.....	332
	View identity mapping settings.....	332
	Managing user identities.....	332
	View user identity.....	333
	Create a user-mapping rule.....	334
	Merge Windows and UNIX tokens.....	335
	Retrieve the primary group from LDAP.....	335
	Mapping rule options.....	336
	Mapping rule operators.....	337
Chapter 8	Auditing	339
	Auditing overview.....	340
	Syslog.....	340
	Enable syslog.....	340
	Syslog forwarding.....	341
	Protocol audit events.....	341
	Sample config audit log.....	342
	Supported event types.....	343
	Supported audit tools.....	343
	Managing audit settings.....	344
	Enable system configuration auditing.....	344
	Enable protocol access auditing.....	345
	Auditing settings.....	346
	Integrating with the EMC Common Event Enabler.....	347
	Install CEE for Windows.....	347
	Configure CEE for Windows.....	348
	Auditing commands.....	349
	isi audit settings modify.....	349
	isi audit settings view.....	351
	isi audit topics list.....	352
	isi audit topics modify.....	352
	isi audit topics view.....	353
Chapter 9	File sharing	355
	File sharing overview.....	356
	Mixed protocol environments.....	356
	Write caching with SmartCache.....	357
	SMB.....	358
	SMB shares in access zones.....	359
	SMB Multichannel.....	359
	SMB share management through MMC.....	361
	Symbolic links and SMB clients.....	361

	Anonymous access to SMB shares.....	363
	Managing SMB settings.....	363
	Managing SMB shares.....	365
	SMB commands.....	371
NFS.....		395
	NFS exports.....	395
	NFS aliases.....	395
	NFS log files.....	396
	Managing the NFS service.....	396
	Managing NFS exports.....	397
	Managing NFS aliases.....	400
	NFS commands.....	403
FTP.....		444
	View FTP settings.....	444
	Enable FTP file sharing.....	444
	Configure FTP file sharing.....	445
	FTP commands.....	445
HTTP and HTTPS.....		460
	Enable and configure HTTP.....	460
Chapter 10	Home directories	463
	Home directories overview.....	464
	Home directory permissions.....	464
	Authenticating SMB users.....	464
	Home directory creation through SMB.....	464
	Create home directories with expansion variables.....	465
	Create home directories with the --inheritable-path-acl option.....	466
	Create special home directories with the SMB share %U variable... ..	467
	Home directory creation through SSH and FTP.....	468
	Set the SSH or FTP login shell	468
	Set SSH/FTP home directory permissions.....	468
	Set SSH/FTP home directory creation options.....	469
	Provision home directories with dot files.....	470
	Home directory creation in a mixed environment.....	471
	Interactions between ACLs and mode bits.....	471
	Default home directory settings in authentication providers.....	471
	Supported expansion variables.....	472
	Domain variables in home directory provisioning.....	474
Chapter 11	Snapshots	475
	Snapshots overview.....	476
	Data protection with SnapshotIQ.....	476
	Snapshot disk-space usage.....	476
	Snapshot schedules.....	477
	Snapshot aliases.....	477
	File and directory restoration.....	477
	Best practices for creating snapshots.....	478
	Best practices for creating snapshot schedules.....	478
	File clones.....	479
	Shadow-store considerations.....	480
	Snapshot locks.....	480
	Snapshot reserve.....	481
	SnapshotIQ license functionality.....	481
	Creating snapshots with SnapshotIQ.....	481

Create a SnapRevert domain.....	482
Create a snapshot schedule.....	482
Create a snapshot.....	482
Snapshot naming patterns.....	483
Managing snapshots	485
Reducing snapshot disk-space usage.....	485
Delete a snapshot.....	486
Modify snapshot attributes.....	486
Modify a snapshot alias	486
View snapshots.....	487
Snapshot information.....	487
Restoring snapshot data.....	488
Revert a snapshot	488
Restore a file or directory using Windows Explorer.....	488
Restore a file or directory through a UNIX command line.....	489
Clone a file from a snapshot.....	489
Managing snapshot schedules.....	490
Modify a snapshot schedule	490
Delete a snapshot schedule	490
View snapshot schedules	490
Managing snapshot aliases.....	491
Configure a snapshot alias for a snapshot schedule.....	491
Assign a snapshot alias to a snapshot.....	491
Reassign a snapshot alias to the live file system.....	492
View snapshot aliases.....	492
Snapshot alias information.....	492
Managing with snapshot locks.....	493
Create a snapshot lock.....	493
Modify a snapshot lock expiration date.....	493
Delete a snapshot lock.....	494
Snapshot lock information	494
Configure SnapshotIQ settings	494
SnapshotIQ settings	495
Set the snapshot reserve.....	496
Snapshot commands.....	496
Snapshot naming patterns.....	496
isi snapshot schedules create.....	499
isi snapshot schedules modify.....	500
isi snapshot schedules delete.....	502
isi snapshot schedules list.....	502
isi snapshot schedules view.....	504
isi snapshot schedules pending list.....	504
isi snapshot snapshots create.....	505
isi snapshot snapshots modify.....	506
isi snapshot snapshots delete.....	507
isi snapshot snapshots list.....	507
isi snapshot snapshots view.....	509
isi snapshot settings modify.....	509
isi snapshot settings view.....	511
isi snapshot locks create.....	511
isi snapshot locks modify.....	512
isi snapshot locks delete.....	513
isi snapshot locks list.....	514
isi snapshot locks view.....	515
isi snapshot aliases create.....	515
isi snapshot aliases modify.....	515

	isi snapshot aliases delete.....	516
	isi snapshot aliases list.....	516
	isi snapshot aliases view.....	517
Chapter 12	Deduplication with SmartDedupe	519
	Deduplication overview.....	520
	Deduplication jobs.....	520
	Data replication and backup with deduplication.....	521
	Snapshots with deduplication.....	521
	Deduplication considerations.....	521
	Shadow-store considerations.....	522
	SmartDedupe license functionality.....	522
	Managing deduplication.....	522
	Assess deduplication space savings	523
	Specify deduplication settings	523
	View deduplication space savings	524
	View a deduplication report	524
	Deduplication job report information.....	524
	Deduplication information.....	525
	Deduplication commands.....	526
	isi dedupe settings modify.....	526
	isi dedupe settings view.....	527
	isi dedupe stats.....	527
	isi dedupe reports list.....	528
	isi dedupe reports view	529
Chapter 13	Data replication with SyncIQ	531
	SyncIQ backup and recovery overview.....	532
	Replication policies and jobs.....	532
	Automated replication policies.....	533
	Source and target cluster association.....	533
	Full and differential replication.....	534
	Controlling replication job resource consumption.....	534
	Replication reports.....	535
	Replication snapshots.....	535
	Source cluster snapshots.....	535
	Target cluster snapshots.....	536
	Data failover and failback with SyncIQ.....	536
	Data failover.....	537
	Data failback.....	537
	Recovery times and objectives for SyncIQ.....	537
	SyncIQ license functionality.....	538
	Creating replication policies.....	538
	Excluding directories in replication.....	538
	Excluding files in replication.....	539
	File criteria options.....	540
	Configure default replication policy settings	542
	Create a replication policy.....	542
	Create a SyncIQ domain.....	543
	Assess a replication policy	543
	Managing replication to remote clusters.....	543
	Start a replication job.....	544
	Pause a replication job	544
	Resume a replication job	544

Cancel a replication job	544
View active replication jobs	545
Replication job information	545
Initiating data failover and failback with SyncIQ.....	545
Fail over data to a secondary cluster	546
Revert a failover operation.....	546
Fail back data to a primary cluster	547
Performing disaster recovery for SmartLock directories.....	547
Recover SmartLock directories on a target cluster	547
Migrate SmartLock directories	548
Managing replication policies.....	550
Modify a replication policy	550
Delete a replication policy	550
Enable or disable a replication policy	551
View replication policies	551
Replication policy information	552
Managing replication to the local cluster.....	552
Cancel replication to the local cluster	553
Break local target association	553
View replication policies targeting the local cluster.....	553
Remote replication policy information	554
Managing replication performance rules.....	554
Create a network traffic rule	554
Create a file operations rule	554
Modify a performance rule	555
Delete a performance rule	555
Enable or disable a performance rule	555
View performance rules	556
Managing replication reports.....	556
Configure default replication report settings	556
Delete replication reports.....	556
View replication reports	557
Replication report information.....	558
Managing failed replication jobs.....	559
Resolve a replication policy	559
Reset a replication policy	559
Perform a full or differential replication.....	560
Managing changelists.....	560
Create a changelist.....	561
View a changelist.....	561
Changelist information.....	562
Data replication commands.....	564
isi sync policies create.....	564
isi sync policies modify.....	572
isi sync policies delete.....	580
isi sync policies list.....	581
isi sync policies view.....	583
isi sync policies disable.....	584
isi sync policies enable.....	584
isi sync jobs start.....	584
isi sync jobs pause.....	585
isi sync jobs resume.....	585
isi sync jobs cancel.....	586
isi sync jobs list.....	586
isi sync jobs view.....	587
isi sync jobs reports list.....	587

	isi sync jobs reports view.....	588
	isi sync settings modify.....	588
	isi sync settings view.....	589
	isi sync policies resolve.....	589
	isi sync policies reset.....	590
	isi sync target cancel.....	590
	isi sync target list.....	590
	isi sync target view.....	591
	isi sync target break.....	592
	isi sync target reports list.....	592
	isi sync target reports view.....	594
	isi sync target reports subreports list.....	594
	isi sync target reports subreports view.....	596
	isi sync reports list.....	596
	isi sync reports view.....	598
	isi sync reports rotate.....	598
	isi sync reports subreports list.....	598
	isi sync reports subreports view.....	600
	isi sync recovery allow-write.....	600
	isi sync recovery resync-prep.....	601
	isi sync rules create.....	601
	isi sync rules modify.....	602
	isi sync rules delete.....	603
	isi sync rules list.....	604
	isi sync rules view.....	604
	isi_changelist_mod.....	605
Chapter 14	Data layout with FlexProtect	609
	FlexProtect overview.....	610
	File striping.....	610
	Requested data protection.....	610
	FlexProtect data recovery.....	611
	Smartfail.....	611
	Node failures.....	611
	Requesting data protection.....	612
	Requested protection settings.....	612
	Requested protection disk space usage.....	613
Chapter 15	NDMP backup	615
	NDMP backup and recovery overview.....	616
	NDMP two-way backup.....	616
	Snapshot-based incremental backups.....	617
	NDMP protocol support.....	618
	Supported DMAs.....	618
	NDMP hardware support.....	619
	NDMP backup limitations.....	619
	NDMP performance recommendations.....	620
	Excluding files and directories from NDMP backups.....	621
	Configuring basic NDMP backup settings.....	623
	Configure and enable NDMP backup.....	623
	Disable NDMP backup.....	623
	View NDMP backup settings.....	623
	NDMP backup settings.....	623
	Managing NDMP user accounts.....	624

Create an NDMP user account	624
Modify the password of an NDMP user account	624
Delete an NDMP user account	624
View NDMP user accounts	624
Managing NDMP backup devices.....	625
Detect NDMP backup devices	625
Modify an NDMP backup device entry name	625
Delete a device entry for a disconnected NDMP backup device.....	625
View NDMP backup devices	626
Managing NDMP backup ports.....	626
Modify NDMP backup port settings	626
Enable or disable an NDMP backup port.....	626
View NDMP backup ports	627
NDMP backup port settings	627
Managing NDMP backup sessions.....	627
End an NDMP session	627
View NDMP sessions	628
NDMP session information	628
Managing restartable backups.....	629
Configure restartable backups for EMC NetWorker.....	629
View restartable backup contexts.....	630
Delete a restartable backup context.....	630
Configure restartable backup settings.....	630
View restartable backup settings.....	631
Managing file list backups.....	631
Format of a backup file list.....	632
Placement of the file list.....	632
Start a file list backup.....	632
Parallel restore operation.....	633
Specify a serial restore operation.....	633
Sharing tape drives between clusters.....	634
Managing default NDMP settings.....	634
Set default NDMP settings for a directory.....	634
Modify default NDMP settings for a directory.....	634
View default NDMP settings for directories.....	635
NDMP environment variables.....	635
Managing snapshot based incremental backups.....	638
Enable snapshot-based incremental backups for a directory.....	638
Delete snapshots for snapshot-based incremental backups.....	638
View snapshots for snapshot-based incremental backups.....	639
View NDMP backup logs	639
NDMP backup commands.....	639
isi ndmp user create.....	639
isi ndmp user modify.....	640
isi ndmp user delete.....	640
isi ndmp user list.....	640
isi tape rescan.....	641
isi tape rename.....	641
isi tape delete.....	642
isi tape list.....	642
isi fc set.....	643
isi fc disable.....	644
isi fc enable.....	644
isi fc list.....	645
isi ndmp kill.....	646
isi ndmp list.....	646

	isi ndmp probe.....	647
	isi ndmp settings set.....	647
	isi ndmp settings list.....	648
	isi ndmp settings variables create.....	649
	isi ndmp settings variables modify.....	649
	isi ndmp settings variables delete.....	650
	isi ndmp settings variables list.....	651
	isi ndmp dumpdates delete.....	651
	isi ndmp dumpdates list.....	651
	isi ndmp extensions context delete.....	652
	isi ndmp extensions contexts list.....	652
	isi ndmp extensions contexts view.....	653
	isi ndmp extensions settings modify.....	654
	isi ndmp extensions settings view.....	654
Chapter 16	File retention with SmartLock	655
	SmartLock overview.....	656
	Compliance mode.....	656
	SmartLock directories.....	656
	Replication and backup with SmartLock.....	657
	SmartLock replication and backup limitations.....	657
	SmartLock license functionality.....	658
	SmartLock considerations.....	658
	Set the compliance clock.....	659
	View the compliance clock.....	659
	Creating a SmartLock directory.....	659
	Retention periods.....	659
	Autocommit time periods.....	660
	Create a SmartLock directory.....	660
	Managing SmartLock directories.....	661
	Modify a SmartLock directory.....	661
	View SmartLock directory settings.....	661
	SmartLock directory configuration settings.....	662
	Managing files in SmartLock directories.....	665
	Set a retention period through a UNIX command line.....	665
	Set a retention period through Windows Powershell.....	665
	Commit a file to a WORM state through a UNIX command line.....	666
	Commit a file to a WORM state through Windows Explorer.....	666
	Override the retention period for all files in a SmartLock directory..	666
	Delete a file committed to a WORM state	667
	View WORM status of a file.....	667
	SmartLock commands.....	668
	isi worm domains create.....	668
	isi worm domains modify.....	671
	isi worm domains list.....	675
	isi worm domains view	676
	isi worm cdate set.....	676
	isi worm cdate view.....	677
	isi worm files delete.....	677
	isi worm files view.....	677
Chapter 17	Protection domains	679
	Protection domains overview.....	680
	Protection domain considerations.....	680

	Create a protection domain	681
	Delete a protection domain	681
Chapter 18	Data-at-rest-encryption	683
	Data-at-rest encryption overview	684
	Self-encrypting drives	684
	Data security on self-encrypted drives	684
	Data migration to a self-encrypted-drives cluster	685
	Chassis and drive states	685
	Smartfailed drive REPLACE state	688
	Smartfailed drive ERASE state	689
Chapter 19	SmartQuotas	691
	SmartQuotas overview	692
	Quota types	692
	Default quota type	693
	Usage accounting and limits	695
	Disk-usage calculations	696
	Quota notifications	697
	Quota notification rules	697
	Quota reports	698
	Creating quotas	698
	Create an accounting quota	699
	Create an enforcement quota	699
	Managing quotas	700
	Search for quotas	700
	Manage quotas	700
	Export a quota configuration file	701
	Import a quota configuration file	702
	Managing quota notifications	702
	Email quota notification messages	704
	Managing quota reports	706
	Basic quota settings	708
	Advisory limit quota notification rules settings	708
	Soft limit quota notification rules settings	709
	Hard limit quota notification rules settings	710
	Limit notification settings	710
	Quota report settings	711
	Quota commands	712
	isi quota quotas create	712
	isi quota quotas delete	714
	isi quota quotas modify	715
	isi quota quotas list	718
	isi quota quotas view	720
	isi quota quotas notifications clear	721
	isi quota quotas notifications create	722
	isi quota quotas notifications delete	725
	isi quota quotas notifications disable	726
	isi quota quotas notifications list	728
	isi quota quotas notifications modify	730
	isi quota quotas notifications view	733
	isi quota reports create	734
	isi quota reports delete	735
	isi quota reports list	736

- isi quota settings notifications clear736
- isi quota settings notifications create736
- isi quota settings notifications delete 739
- isi quota settings notifications list..... 740
- isi quota settings notifications modify740
- isi quota settings notifications view 742
- isi quota settings reports modify743
- isi quota settings reports view745

Chapter 20 Storage Pools 747

- Storage pools overview 748
- Storage pool functions 748
- Autoprovisioning750
- Node pools 750
 - Node compatibilities750
 - Manual node pools 751
- Virtual hot spare751
- Spillover 752
- Suggested protection 752
- Protection policies 753
- SSD strategies753
- Global namespace acceleration 754
- L3 cache overview755
 - Migration to L3 cache756
 - L3 cache on HD400 node pools756
- Tiers756
- File pool policies756
- Managing node pools through the command-line interface 757
 - Add a compatible node to a node pool758
 - Merge compatible node pools758
 - Delete a compatibility759
 - Create a node pool manually759
 - Add a node to a manually managed node pool760
 - Change the name or protection policy of a node pool760
 - Remove a node from a manually managed node pool760
- Managing L3 cache from the command-line interface761
 - Set L3 cache as the default for new node pools761
 - Enable L3 cache on a specific node pool761
 - Restore SSDs to storage drives for a node pool762
- Managing tiers762
 - Create a tier762
 - Add or move node pools in a tier762
 - Rename a tier763
 - Delete a tier763
- Creating file pool policies763
 - Create a file pool policy764
 - File-matching options for file pool policies764
 - Valid wildcard characters766
 - SmartPools settings766
 - Default file pool requested protection settings769
 - Default file pool I/O optimization settings770
- Managing file pool policies771
 - Modify a file pool policy771
 - Modify default storage pool settings772
 - Configure default file pool policy settings772

	Prioritize a file pool policy.....	773
	Delete a file pool policy.....	773
	Monitoring storage pools.....	774
	Monitor storage pools.....	774
	View the health of storage pools.....	774
	View results of a SmartPools job.....	774
	Storage pool commands.....	775
	isi filepool apply.....	775
	isi filepool default-policy modify.....	777
	isi filepool default-policy view.....	779
	isi filepool policies create.....	779
	isi filepool policies delete.....	782
	isi filepool policies list.....	783
	isi filepool policies modify.....	783
	isi filepool policies view.....	787
	isi filepool templates list.....	787
	isi filepool templates view.....	788
	isi storagepool compatibilities active create.....	788
	isi storagepool compatibilities active delete.....	789
	isi storagepool compatibilities active list.....	789
	isi storagepool compatibilities active view.....	790
	isi storagepool compatibilities available list.....	791
	isi storagepool health.....	792
	isi storagepool list.....	792
	isi storagepool nodepools create.....	792
	isi storagepool nodepools delete.....	793
	isi storagepool nodepools list.....	793
	isi storagepool nodepools modify.....	794
	isi storagepool nodepools view.....	795
	isi storagepool settings modify.....	796
	isi storagepool settings view.....	797
	isi storagepool tiers create.....	797
	isi storagepool tiers delete.....	798
	isi storagepool tiers list.....	798
	isi storagepool tiers modify.....	799
	isi storagepool tiers view.....	799
Chapter 21	System jobs	801
	System jobs overview.....	802
	System jobs library.....	802
	Job operation.....	805
	Job performance impact.....	806
	Job priorities.....	807
	Managing system jobs.....	807
	Start a job.....	808
	Pause a job.....	808
	Modify a job.....	808
	Resume a job.....	809
	Cancel a job.....	809
	Modify job type settings.....	810
	View active jobs.....	810
	View job history.....	810
	Managing impact policies.....	811
	Create an impact policy.....	811
	View impact policy settings.....	812

- Modify an impact policy..... 812
- Delete an impact policy..... 813
- Viewing job reports and statistics..... 813
 - View statistics for a job in progress..... 813
 - View a report for a completed job..... 814
- Job management commands..... 815
 - isi job events list..... 815
 - isi job jobs cancel..... 817
 - isi job jobs list..... 817
 - isi job jobs modify..... 818
 - isi job jobs pause..... 819
 - isi job jobs resume..... 820
 - isi job jobs start..... 820
 - isi job jobs view..... 822
 - isi job policies create..... 823
 - isi job policies delete..... 823
 - isi job policies list..... 824
 - isi job policies modify..... 825
 - isi job policies view..... 826
 - isi job reports list..... 827
 - isi job reports view..... 828
 - isi job statistics list..... 829
 - isi job statistics view..... 830
 - isi job types list..... 831
 - isi job status..... 832
 - isi job types modify..... 833
 - isi job types view..... 834

Chapter 22 Networking 837

- Networking overview..... 838
- Internal network overview..... 838
 - Internal IP address ranges..... 838
 - Internal network failover..... 839
- External client network overview..... 839
 - External network settings..... 839
 - IP address pools..... 840
 - IPv6 support..... 840
 - SmartConnect module..... 841
 - Connection balancing..... 841
 - IP address allocation..... 842
 - IP address failover..... 843
 - IP address rebalancing..... 843
 - SmartConnect DNS service..... 844
 - DNS name resolution..... 844
 - NIC aggregation..... 845
 - Routing options..... 845
 - VLANs..... 847
- Managing internal network settings..... 847
 - Add or remove an internal IP address range..... 847
 - Modify an internal network netmask..... 847
 - Configure and enable internal network failover..... 848
 - Disable internal network failover..... 849
- Managing external network settings..... 849
 - Configure DNS settings..... 849
- Managing external network subnets..... 850

Create a subnet.....	850
Modify a subnet.....	850
Delete a subnet.....	851
View subnets.....	851
Enable or disable VLAN tagging.....	852
Add or remove a DSR address.....	852
Managing IP address pools.....	853
Create an IP address pool.....	853
Modify an IP address pool.....	853
Delete an IP address pool.....	853
View IP address pools.....	854
Add or remove an IP address range.....	854
Configure IP address allocation.....	855
Configure an IP rebalance policy.....	855
Configure an IP failover policy.....	856
Managing SmartConnect Settings.....	857
Configure a SmartConnect zone.....	857
Add or remove a SmartConnect zone alias.....	858
Configure a SmartConnect connection balancing policy.....	858
Configure a SmartConnect service IP address.....	859
Configure a SmartConnect service subnet.....	859
Managing network interface members.....	860
Add or remove a network interface.....	860
Configure NIC Aggregation.....	860
Add or remove a static route.....	862
View network interfaces.....	862
Managing network interface provisioning rules.....	863
Create a network interface provisioning rule.....	863
Modify a network interface provisioning rule.....	863
Delete a network interface provisioning rule.....	864
View network interface provisioning rules.....	864
Managing routing options.....	865
Enable or disable source-based routing.....	865
Add or remove a static route.....	865
Networking commands.....	866
isi networks.....	866
isi networks create pool.....	868
isi networks create rule.....	872
isi networks create subnet.....	873
isi networks delete pool.....	875
isi networks delete rule.....	875
isi networks delete subnet.....	876
isi networks dnscache disable.....	876
isi networks dnscache enable.....	876
isi networks dnscache flush.....	877
isi networks dnscache modify.....	877
isi networks dnscache statistics.....	878
isi networks list interfaces.....	878
isi networks list pools.....	879
isi networks list rules.....	880
isi networks list subnets.....	881
isi networks modify pool.....	881
isi networks modify rule.....	886
isi networks modify subnet.....	887
isi networks sbr enable.....	889
isi networks sbr disable.....	889

Chapter 23	Hadoop	891
	Hadoop overview.....	892
	Hadoop architecture.....	892
	Hadoop compute layer.....	892
	HDFS storage layer.....	892
	How Hadoop is implemented on OneFS.....	893
	Hadoop distributions supported by OneFS.....	893
	WebHDFS.....	894
	Secure impersonation.....	894
	Ambari agent.....	895
	Virtual HDFS racks.....	895
	HDFS implementation considerations.....	896
	HDFS directories and Hadoop user accounts.....	896
	HDFS settings in access zones.....	896
	HDFS and SmartConnect.....	896
	Implementing Hadoop with OneFS.....	897
	Managing the HDFS service.....	897
	Configure HDFS service settings.....	897
	HDFS service settings.....	898
	View HDFS service settings.....	899
	Enable or disable the HDFS service.....	899
	Managing HDFS access zone settings.....	900
	Supported HDFS authentication methods.....	900
	Set the HDFS authentication method in an access zone.....	900
	Configure HDFS authentication properties on the Hadoop client.....	901
	Create a local Hadoop user.....	902
	Set the HDFS root directory in an access zone.....	902
	Enable or disable WebHDFS within an access zone.....	902
	Configure Ambari agent settings.....	903
	Configuring secure impersonation.....	903
	Create a proxy user.....	903
	Modify a proxy user.....	904
	Delete a proxy user.....	904
	List the members of a proxy user.....	904
	View proxy users.....	905
	Managing virtual HDFS racks.....	905
	Create a virtual HDFS rack.....	906
	Modify a virtual HDFS rack.....	906
	Delete a virtual HDFS rack.....	907
	View virtual HDFS racks.....	907
	HDFS commands.....	908
	isi hdfs settings modify.....	908
	isi hdfs settings view.....	909
	isi hdfs proxyusers create.....	909
	isi hdfs proxyusers modify.....	911
	isi hdfs proxyusers delete.....	912
	isi hdfs proxyusers members list.....	913
	isi hdfs proxyusers list.....	914
	isi hdfs proxyusers view.....	914
	isi hdfs racks create.....	915
	isi hdfs racks modify.....	915
	isi hdfs racks delete.....	916
	isi hdfs racks list.....	917
	isi hdfs racks view.....	917

Chapter 24	Antivirus	919
	Antivirus overview.....	920
	On-access scanning.....	920
	Antivirus policy scanning.....	921
	Individual file scanning.....	921
	Antivirus scan reports.....	921
	ICAP servers.....	921
	Supported ICAP servers.....	922
	Antivirus threat responses.....	922
	Configuring global antivirus settings.....	923
	Exclude files from antivirus scans	923
	Configure on-access scanning settings	923
	Configure antivirus threat response settings	924
	Configure antivirus report retention settings.....	924
	Enable or disable antivirus scanning.....	924
	Managing ICAP servers.....	924
	Add and connect to an ICAP server	924
	Test an ICAP server connection.....	925
	Temporarily disconnect from an ICAP server	925
	Reconnect to an ICAP server	925
	Remove an ICAP server	925
	Create an antivirus policy	926
	Managing antivirus policies.....	926
	Modify an antivirus policy	926
	Delete an antivirus policy	926
	Enable or disable an antivirus policy	927
	View antivirus policies	927
	Managing antivirus scans.....	927
	Scan a file.....	927
	Manually run an antivirus policy.....	927
	Stop a running antivirus scan.....	928
	Managing antivirus threats.....	928
	Manually quarantine a file	928
	Remove a file from quarantine	928
	View threats	928
	Antivirus threat information.....	929
	Managing antivirus reports.....	929
	Export an antivirus report.....	929
	View antivirus reports	929
	View antivirus events.....	930
	Antivirus commands.....	930
	isi avscan policy add.....	930
	isi avscan policy edit.....	931
	isi avscan policy delete.....	933
	isi avscan policy.....	933
	isi avscan policy run.....	933
	isi avscan manual.....	934
	isi avscan quarantine.....	934
	isi avscan unquarantine.....	934
	isi avscan report threat.....	935
	isi avscan report scan.....	935
	isi avscan report purge.....	936
	isi avscan settings.....	937
	isi avscan get.....	939

Chapter 25	VMware integration	941
	VMware integration overview	942
	VAAI	942
	VASA	942
	Isilon VASA alarms	942
	VASA storage capabilities	943
	Configuring VASA support	943
	Enable VASA	943
	Download the Isilon vendor provider certificate	944
	Add the Isilon vendor provider	944
	Disable or re-enable VASA	945
	Troubleshooting VASA storage display failures	945

CHAPTER 1

Introduction to this guide

This section contains the following topics:

- [About this guide](#).....26
- [Isilon scale-out NAS overview](#).....26
- [Where to go for support](#).....26

About this guide

This guide describes how the Isilon OneFS command-line interface provides access to cluster configuration, management, and monitoring functionality. This guide also lists and describes all OneFS-specific commands that extend the standard UNIX command set.

We value your feedback. Please let us know how we can improve this document.

- Take the survey at <https://www.research.net/s/isi-docfeedback>.
- Send your comments or suggestions to docfeedback@isilon.com.

Isilon scale-out NAS overview

The EMC Isilon scale-out NAS storage platform combines modular hardware with unified software to harness unstructured data. Powered by the OneFS operating system, an EMC Isilon cluster delivers a scalable pool of storage with a global namespace.

The platform's unified software provides centralized web-based and command-line administration to manage the following features:

- A cluster that runs a distributed file system
- Scale-out nodes that add capacity and performance
- Storage options that manage files, block data, and tiering
- Flexible data protection and high availability
- Software modules that control costs and optimize resources

Where to go for support

You can contact EMC Isilon Technical Support for any questions about EMC Isilon products.

Online Support	Live Chat Create a Service Request
Telephone Support	United States: 1-800-SVC-4EMC (800-782-4362) Canada: 800-543-4782 Worldwide: +1-508-497-7901 For local phone numbers in your country, see EMC Customer Support Centers .
Help with online support	For questions specific to EMC Online Support registration or access, email support@emc.com .

CHAPTER 2

Isilon scale-out NAS

This section contains the following topics:

- [OneFS storage architecture](#)..... 28
- [Isilon node components](#)..... 28
- [Internal and external networks](#)..... 29
- [Isilon cluster](#)..... 29
- [The OneFS operating system](#)..... 31
- [Structure of the file system](#)..... 33
- [Data protection overview](#)..... 35
- [VMware integration](#)..... 37
- [Software modules](#)..... 38

OneFS storage architecture

EMC Isilon takes a scale-out approach to storage by creating a cluster of nodes that runs a distributed file system. OneFS combines the three layers of storage architecture—file system, volume manager, and data protection—into a scale-out NAS cluster.

Each node adds resources to the cluster. Because each node contains globally coherent RAM, as a cluster becomes larger, it becomes faster. Meanwhile, the file system expands dynamically and redistributes content, which eliminates the work of partitioning disks and creating volumes.

Nodes work as peers to spread data across the cluster. Segmenting and distributing data—a process known as striping—not only protects data, but also enables a user connecting to any node to take advantage of the entire cluster's performance.

OneFS uses distributed software to scale data across commodity hardware. Each node helps control data requests, boosts performance, and expands the cluster's capacity. No master device controls the cluster; no slaves invoke dependencies. Instead, each node helps control data requests, boosts performance, and expands the cluster's capacity.

Isilon node components

As a rack-mountable appliance, a storage node includes the following components in a 2U or 4U rack-mountable chassis with an LCD front panel: CPUs, RAM, NVRAM, network interfaces, InfiniBand adapters, disk controllers, and storage media. An Isilon cluster comprises three or more nodes, up to 144.

When you add a node to a cluster, you increase the cluster's aggregate disk, cache, CPU, RAM, and network capacity. OneFS groups RAM into a single coherent cache so that a data request on a node benefits from data that is cached anywhere. NVRAM is grouped to write data with high throughput and to protect write operations from power failures. As the cluster expands, spindles and CPU combine to increase throughput, capacity, and input-output operations per second (IOPS).

EMC Isilon makes several types of nodes, all of which can be added to a cluster to balance capacity and performance with throughput or IOPS:

Node	Use Case
S-Series	IOPS-intensive applications
X-Series	High-concurrency and throughput-driven workflows
NL-Series	Near-primary accessibility, with near-tape value
HD-Series	Maximum capacity

The following EMC Isilon nodes improve performance:

Node	Function
A-Series Performance Accelerator	Independent scaling for high performance
A-Series Backup Accelerator	High-speed and scalable backup-and-restore solution for tape drives over Fibre Channel connections

Internal and external networks

A cluster includes two networks: an internal network to exchange data between nodes and an external network to handle client connections.

Nodes exchange data through the internal network with a proprietary, unicast protocol over InfiniBand. Each node includes redundant InfiniBand ports so you can add a second internal network in case the first one fails.

Clients reach the cluster with 1 GigE or 10 GigE Ethernet. Since every node includes Ethernet ports, the cluster's bandwidth scales with performance and capacity as you add nodes.

Isilon cluster

An Isilon cluster consists of three or more hardware nodes, up to 144. Each node runs the Isilon OneFS operating system, the distributed file-system software that unites the nodes into a cluster. A cluster's storage capacity ranges from a minimum of 18 TB to a maximum of 15.5 PB.

Cluster administration

OneFS centralizes cluster management through a web administration interface and a command-line interface. Both interfaces provide methods to activate licenses, check the status of nodes, configure the cluster, upgrade the system, generate alerts, view client connections, track performance, and change various settings.

In addition, OneFS simplifies administration by automating maintenance with a job engine. You can schedule jobs that scan for viruses, inspect disks for errors, reclaim disk space, and check the integrity of the file system. The engine manages the jobs to minimize impact on the cluster's performance.

With SNMP versions 2c and 3, you can remotely monitor hardware components, CPU usage, switches, and network interfaces. EMC Isilon supplies management information bases (MIBs) and traps for the OneFS operating system.

OneFS also includes a RESTful application programming interface—known as the Platform API—to automate access, configuration, and monitoring. For example, you can retrieve performance statistics, provision users, and tap the file system. The Platform API integrates with OneFS role-based access control to increase security. See the *Isilon Platform API Reference*.

Quorum

An Isilon cluster must have a quorum to work properly. A quorum prevents data conflicts—for example, conflicting versions of the same file—in case two groups of nodes become unsynchronized. If a cluster loses its quorum for read and write requests, you cannot access the OneFS file system.

For a quorum, more than half the nodes must be available over the internal network. A seven-node cluster, for example, requires a four-node quorum. A 10-node cluster requires a six-node quorum. If a node is unreachable over the internal network, OneFS separates the node from the cluster, an action referred to as splitting. After a cluster is split, cluster operations continue as long as enough nodes remain connected to have a quorum.

In a split cluster, the nodes that remain in the cluster are referred to as the majority group. Nodes that are split from the cluster are referred to as the minority group.

When split nodes can reconnect with the cluster and resynchronize with the other nodes, the nodes rejoin the cluster's majority group, an action referred to as merging.

A OneFS cluster contains two quorum properties:

- read quorum (`efs.gmp.has_quorum`)
- write quorum (`efs.gmp.has_super_block_quorum`)

By connecting to a node with SSH and running the `sysctl` command-line tool as root, you can view the status of both types of quorum. Here is an example for a cluster that has a quorum for both read and write operations, as the command's output indicates with a 1, for true:

```
sysctl efs.gmp.has_quorum
efs.gmp.has_quorum: 1
sysctl efs.gmp.has_super_block_quorum
efs.gmp.has_super_block_quorum: 1
```

The degraded states of nodes—such as smartfail, read-only, offline, and so on—affect quorum in different ways. A node in a smartfail or read-only state affects only write quorum. A node in an offline state, however, affects both read and write quorum. In a cluster, the combination of nodes in different degraded states determines whether read requests, write requests, or both work.

A cluster can lose write quorum but keep read quorum. Consider a four-node cluster in which nodes 1 and 2 are working normally. Node 3 is in a read-only state, and node 4 is in a smartfail state. In such a case, read requests to the cluster succeed. Write requests, however, receive an input-output error because the states of nodes 3 and 4 break the write quorum.

A cluster can also lose both its read and write quorum. If nodes 3 and 4 in a four-node cluster are in an offline state, both write requests and read requests receive an input-output error, and you cannot access the file system. When OneFS can reconnect with the nodes, OneFS merges them back into the cluster. Unlike a RAID system, an Isilon node can rejoin the cluster without being rebuilt and reconfigured.

Splitting and merging

Splitting and merging optimize the use of nodes without your intervention.

OneFS monitors every node in a cluster. If a node is unreachable over the internal network, OneFS separates the node from the cluster, an action referred to as splitting. When the cluster can reconnect to the node, OneFS adds the node back into the cluster, an action referred to as merging.

When a node is split from a cluster, it will continue to capture event information locally. You can connect to a split node with SSH and run the `isi events list` command to view the local event log for the node. The local event log can help you troubleshoot the connection issue that resulted in the split. When the split node rejoins the cluster, local events gathered during the split are deleted. You can still view events generated by a split node in the node's event log file located at `/var/log/isi_celog_events.log`.

If a cluster splits during a write operation, OneFS might need to re-allocate blocks for the file on the side with the quorum, which leads allocated blocks on the side without a quorum to become orphans. When the split nodes reconnect with the cluster, the OneFS Collect system job reclaims the orphaned blocks.

Meanwhile, as nodes split and merge with the cluster, the OneFS AutoBalance job redistributes data evenly among the nodes in the cluster, optimizing protection and conserving space.

Storage pools

Storage pools segment nodes and files into logical divisions to simplify the management and storage of data.

A storage pool comprises node pools and tiers. Node pools group equivalent nodes to protect data and ensure reliability. Tiers combine node pools to optimize storage by need, such as a frequently used high-speed tier or a rarely accessed archive.

The SmartPools module groups nodes and files into pools. If you do not activate a SmartPools license, the module provisions node pools and creates one file pool. If you activate the SmartPools license, you receive more features. You can, for example, create multiple file pools and govern them with policies. The policies move files, directories, and file pools among node pools or tiers. You can also define how OneFS handles write operations when a node pool or tier is full. SmartPools reserves a virtual hot spare to reprotect data if a drive fails regardless of whether the SmartPools license is activated.

IP address pools

Within a subnet, you can partition a cluster's external network interfaces into pools of IP address ranges. The pools empower you to customize your storage network to serve different groups of users. Although you must initially configure the default external IP subnet in IPv4 format, you can configure additional subnets in IPv4 or IPv6.

You can associate IP address pools with a node, a group of nodes, or NIC ports. For example, you can set up one subnet for storage nodes and another subnet for accelerator nodes. Similarly, you can allocate ranges of IP addresses on a subnet to different teams, such as engineering and sales. Such options help you create a storage topology that matches the demands of your network.

In addition, network provisioning rules streamline the setup of external connections. After you configure the rules with network settings, you can apply the settings to new nodes.

As a standard feature, the OneFS SmartConnect module balances connections among nodes by using a round-robin policy with static IP addresses and one IP address pool for each subnet. Activating a SmartConnect Advanced license adds features, such as defining IP address pools to support multiple DNS zones.

The OneFS operating system

A distributed operating system based on FreeBSD, OneFS presents an Isilon cluster's file system as a single share or export with a central point of administration.

The OneFS operating system does the following:

- Supports common data-access protocols, such as SMB and NFS.
- Connects to multiple identity management systems, such as Active Directory and LDAP.
- Authenticates users and groups.
- Controls access to directories and files.

Data-access protocols

With the OneFS operating system, you can access data with multiple file-sharing and transfer protocols. As a result, Microsoft Windows, UNIX, Linux, and Mac OS X clients can share the same directories and files.

OneFS supports the following protocols.

SMB

The Server Message Block (SMB) protocol enables Windows users to access the cluster. OneFS works with SMB 1, SMB 2, and SMB 2.1, as well as SMB 3.0 for Multichannel only. With SMB 2.1, OneFS supports client opportunity locks (oplocks) and large (1 MB) MTU sizes. The default file share is `/ifs`.

NFS

The Network File System (NFS) protocol enables UNIX, Linux, and Mac OS X systems to remotely mount any subdirectory, including subdirectories created by Windows users. OneFS works with NFS versions 3 and 4. The default export is `/ifs`.

HDFS

The Hadoop Distributed File System (HDFS) protocol enables a cluster to work with Apache Hadoop, a framework for data-intensive distributed applications. HDFS integration requires you to activate a separate license.

FTP

FTP allows systems with an FTP client to connect to the cluster and exchange files.

HTTP

HTTP gives systems browser-based access to resources. OneFS includes limited support for WebDAV.

Identity management and access control

OneFS works with multiple identity management systems to authenticate users and control access to files. In addition, OneFS features access zones that allow users from different directory services to access different resources based on their IP address. Role-based access control, meanwhile, segments administrative access by role.

OneFS authenticates users with the following identity management systems:

- Microsoft Active Directory (AD)
- Lightweight Directory Access Protocol (LDAP)
- Network Information Service (NIS)
- Local users and local groups
- A file provider for accounts in `/etc/spwd.db` and `/etc/group` files. With the file provider, you can add an authoritative third-party source of user and group information.

You can manage users with different identity management systems; OneFS maps the accounts so that Windows and UNIX identities can coexist. A Windows user account managed in Active Directory, for example, is mapped to a corresponding UNIX account in NIS or LDAP.

To control access, an Isilon cluster works with both the access control lists (ACLs) of Windows systems and the POSIX mode bits of UNIX systems. When OneFS must transform a file's permissions from ACLs to mode bits or from mode bits to ACLs, OneFS merges the permissions to maintain consistent security settings.

OneFS presents protocol-specific views of permissions so that NFS exports display mode bits and SMB shares show ACLs. You can, however, manage not only mode bits but also

ACLs with standard UNIX tools, such as the `chmod` and `chown` commands. In addition, ACL policies enable you to configure how OneFS manages permissions for networks that mix Windows and UNIX systems.

Access zones

OneFS includes an access zones feature. Access zones allow users from different authentication providers, such as two untrusted Active Directory domains, to access different OneFS resources based on an incoming IP address. An access zone can contain multiple authentication providers and SMB namespaces.

RBAC for administration

OneFS includes role-based access control (RBAC) for administration. In place of a root or administrator account, RBAC lets you manage administrative access by role. A role limits privileges to an area of administration. For example, you can create separate administrator roles for security, auditing, storage, and backup.

Structure of the file system

OneFS presents all the nodes in a cluster as a global namespace—that is, as the default file share, `/ifs`.

In the file system, directories are inode number links. An inode contains file metadata and an inode number, which identifies a file's location. OneFS dynamically allocates inodes, and there is no limit on the number of inodes.

To distribute data among nodes, OneFS sends messages with a globally routable block address through the cluster's internal network. The block address identifies the node and the drive storing the block of data.

Note

It is recommended that you do not save data to the root `/ifs` file path but in directories below `/ifs`. The design of your data storage structure should be planned carefully. A well-designed directory optimizes cluster performance and cluster administration.

Data layout

OneFS evenly distributes data among a cluster's nodes with layout algorithms that maximize storage efficiency and performance. The system continuously reallocates data to conserve space.

OneFS breaks data down into smaller sections called blocks, and then the system places the blocks in a stripe unit. By referencing either file data or erasure codes, a stripe unit helps safeguard a file from a hardware failure. The size of a stripe unit depends on the file size, the number of nodes, and the protection setting. After OneFS divides the data into stripe units, OneFS allocates, or stripes, the stripe units across nodes in the cluster.

When a client connects to a node, the client's read and write operations take place on multiple nodes. For example, when a client connects to a node and requests a file, the node retrieves the data from multiple nodes and rebuilds the file. You can optimize how OneFS lays out data to match your dominant access pattern—concurrent, streaming, or random.

Writing files

On a node, the input-output operations of the OneFS software stack split into two functional layers: A top layer, or initiator, and a bottom layer, or participant. In read and write operations, the initiator and the participant play different roles.

When a client writes a file to a node, the initiator on the node manages the layout of the file on the cluster. First, the initiator divides the file into blocks of 8 KB each. Second, the initiator places the blocks in one or more stripe units. At 128 KB, a stripe unit consists of 16 blocks. Third, the initiator spreads the stripe units across the cluster until they span a width of the cluster, creating a stripe. The width of the stripe depends on the number of nodes and the protection setting.

After dividing a file into stripe units, the initiator writes the data first to non-volatile random-access memory (NVRAM) and then to disk. NVRAM retains the information when the power is off.

During the write transaction, NVRAM guards against failed nodes with journaling. If a node fails mid-transaction, the transaction restarts without the failed node. When the node returns, it replays the journal from NVRAM to finish the transaction. The node also runs the AutoBalance job to check the file's on-disk striping. Meanwhile, uncommitted writes waiting in the cache are protected with mirroring. As a result, OneFS eliminates multiple points of failure.

Reading files

In a read operation, a node acts as a manager to gather data from the other nodes and present it to the requesting client.

Because an Isilon cluster's coherent cache spans all the nodes, OneFS can store different data in each node's RAM. By using the internal InfiniBand network, a node can retrieve file data from another node's cache faster than from its own local disk. If a read operation requests data that is cached on any node, OneFS pulls the cached data to serve it quickly.

In addition, for files with an access pattern of concurrent or streaming, OneFS pre-fetches in-demand data into a managing node's local cache to further improve sequential-read performance.

Metadata layout

OneFS protects metadata by spreading it across nodes and drives.

Metadata—which includes information about where a file is stored, how it is protected, and who can access it—is stored in inodes and protected with locks in a B+ tree, a standard structure for organizing data blocks in a file system to provide instant lookups. OneFS replicates file metadata across the cluster so that there is no single point of failure.

Working together as peers, all the nodes help manage metadata access and locking. If a node detects an error in metadata, the node looks up the metadata in an alternate location and then corrects the error.

Locks and concurrency

OneFS includes a distributed lock manager that orchestrates locks on data across all the nodes in a cluster.

The lock manager grants locks for the file system, byte ranges, and protocols, including SMB share-mode locks and NFS advisory locks. OneFS also supports SMB opportunistic locks and NFSv4 delegations.

Because OneFS distributes the lock manager across all the nodes, any node can act as a lock coordinator. When a thread from a node requests a lock, the lock manager's hashing algorithm typically assigns the coordinator role to a different node. The coordinator allocates a shared lock or an exclusive lock, depending on the type of request. A shared lock allows users to share a file simultaneously, typically for read operations. An exclusive lock allows only one user to access a file, typically for write operations.

Striping

In a process known as striping, OneFS segments files into units of data and then distributes the units across nodes in a cluster. Striping protects your data and improves cluster performance.

To distribute a file, OneFS reduces it to blocks of data, arranges the blocks into stripe units, and then allocates the stripe units to nodes over the internal network.

At the same time, OneFS distributes erasure codes that protect the file. The erasure codes encode the file's data in a distributed set of symbols, adding space-efficient redundancy. With only a part of the symbol set, OneFS can recover the original file data.

Taken together, the data and its redundancy form a protection group for a region of file data. OneFS places the protection groups on different drives on different nodes—creating data stripes.

Because OneFS stripes data across nodes that work together as peers, a user connecting to any node can take advantage of the entire cluster's performance.

By default, OneFS optimizes striping for concurrent access. If your dominant access pattern is streaming--that is, lower concurrency, higher single-stream workloads, such as with video--you can change how OneFS lays out data to increase sequential-read performance. To better handle streaming access, OneFS stripes data across more drives. Streaming is most effective on clusters or subpools serving large files.

Data protection overview

An Isilon cluster is designed to serve data even when components fail. By default, OneFS protects data with erasure codes, enabling you to retrieve files when a node or disk fails. As an alternative to erasure codes, you can protect data with two to eight mirrors.

When you create a cluster with five or more nodes, erasure codes deliver as much as 80 percent efficiency. On larger clusters, erasure codes provide as much as four levels of redundancy.

In addition to erasure codes and mirroring, OneFS includes the following features to help protect the integrity, availability, and confidentiality of data:

Feature	Description
Antivirus	OneFS can send files to servers running the Internet Content Adaptation Protocol (ICAP) to scan for viruses and other threats.

Feature	Description
Clones	OneFS enables you to create clones that share blocks with other files to save space.
NDMP backup and restore	OneFS can back up data to tape and other devices through the Network Data Management Protocol. Although OneFS supports both NDMP 3-way and 2-way backup, 2-way backup requires an Isilon Backup Accelerator node.
Protection domains	You can apply protection domains to files and directories to prevent changes.

The following software modules also help protect data, but they require you to activate a separate license:

Licensed Feature	Description
SyncIQ	SyncIQ replicates data on another Isilon cluster and automates failover and failback operations between clusters. If a cluster becomes unusable, you can fail over to another Isilon cluster.
SnapshotIQ	You can protect data with a snapshot—a logical copy of data stored on a cluster.
SmartLock	The SmartLock tool prevents users from modifying and deleting files. You can commit files to a write-once, read-many state: The file can never be modified and cannot be deleted until after a set retention period. SmartLock can help you comply with Securities and Exchange Commission Rule 17a-4.

N+M data protection

OneFS supports N+M erasure code levels of N+1, N+2, N+3, and N+4.

In the N+M data model, N represents the number of nodes, and M represents the number of simultaneous failures of nodes or drives that the cluster can handle without losing data. For example, with N+2 the cluster can lose two drives on different nodes or lose two nodes.

To protect drives and nodes separately, OneFS also supports N+M:B. In the N+M:B notation, M is the number of disk failures, and B is the number of node failures. With N+3:1 protection, for example, the cluster can lose three drives or one node without losing data.

The default protection level for clusters larger than 18 TB is N+2:1. The default for clusters smaller than 18 TB is N+1.

The quorum rule dictates the number of nodes required to support a protection level. For example, N+3 requires at least seven nodes so you can maintain a quorum if three nodes fail.

You can, however, set a protection level that is higher than the cluster can support. In a four-node cluster, for example, you can set the protection level at 5x. OneFS protects the data at 4x until a fifth node is added, after which OneFS automatically reprotects the data at 5x.

Data mirroring

You can protect on-disk data with mirroring, which copies data to multiple locations. OneFS supports two to eight mirrors. You can use mirroring instead of erasure codes, or you can combine erasure codes with mirroring.

Mirroring, however, consumes more space than erasure codes. Mirroring data three times, for example, duplicates the data three times, which requires more space than erasure codes. As a result, mirroring suits transactions that require high performance.

You can also mix erasure codes with mirroring. During a write operation, OneFS divides data into redundant protection groups. For files protected by erasure codes, a protection group consists of data blocks and their erasure codes. For mirrored files, a protection group contains all the mirrors of a set of blocks. OneFS can switch the type of protection group as it writes a file to disk. By changing the protection group dynamically, OneFS can continue writing data despite a node failure that prevents the cluster from applying erasure codes. After the node is restored, OneFS automatically converts the mirrored protection groups to erasure codes.

The file system journal

A journal, which records file-system changes in a battery-backed NVRAM card, recovers the file system after failures, such as a power loss. When a node restarts, the journal replays file transactions to restore the file system.

Virtual hot spare

When a drive fails, OneFS uses space reserved in a subpool instead of a hot spare drive. The reserved space is known as a virtual hot spare.

In contrast to a spare drive, a virtual hot spare automatically resolves drive failures and continues writing data. If a drive fails, OneFS migrates data to the virtual hot spare to protect it. You can reserve as many as four disk drives as a virtual hot spare.

Balancing protection with storage space

You can set protection levels to balance protection requirements with storage space.

Higher protection levels typically consume more space than lower levels because you lose an amount of disk space to storing erasure codes. The overhead for the erasure codes depends on the protection level, the file size, and the number of nodes in the cluster. Since OneFS stripes both data and erasure codes across nodes, the overhead declines as you add nodes.

VMware integration

OneFS integrates with several VMware products, including vSphere, vCenter, and ESXi.

For example, OneFS works with the VMware vSphere API for Storage Awareness (VASA) so that you can view information about an Isilon cluster in vSphere. OneFS also works with the VMware vSphere API for Array Integration (VAAI) to support the following features for block storage: hardware-assisted locking, full copy, and block zeroing. VAAI for NFS requires an ESXi plug-in.

With the Isilon for vCenter plug-in, you can backup and restore virtual machines on an Isilon cluster. With the Isilon Storage Replication Adapter, OneFS integrates with the

VMware vCenter Site Recovery Manager to recover virtual machines that are replicated between Isilon clusters.

Software modules

You can access advanced features by activating licenses for EMC Isilon software modules.

SmartLock

SmartLock protects critical data from malicious, accidental, or premature alteration or deletion to help you comply with SEC 17a-4 regulations. You can automatically commit data to a tamper-proof state and then retain it with a compliance clock.

SyncIQ automated failover and failback

SyncIQ replicates data on another Isilon cluster and automates failover and failback between clusters. If a cluster becomes unusable, you can fail over to another Isilon cluster. Failback restores the original source data after the primary cluster becomes available again.

File clones

OneFS provides provisioning of full read/write copies of files, LUNs, and other clones. OneFS also provides virtual machine linked cloning through VMware API integration.

SnapshotIQ

SnapshotIQ protects data with a snapshot—a logical copy of data stored on a cluster. A snapshot can be restored to its top-level directory.

SmartPools

SmartPools enable you to create multiple file pools governed by file-pool policies. The policies move files and directories among node pools or tiers. You can also define how OneFS handles write operations when a node pool or tier is full.

SmartConnect

If you activate a SmartConnect Advanced license, you can balance policies to evenly distribute CPU usage, client connections, or throughput. You can also define IP address pools to support multiple DNS zones in a subnet. In addition, SmartConnect supports IP failover, also known as NFS failover.

InsightIQ

The InsightIQ virtual appliance monitors and analyzes the performance of your Isilon cluster to help you optimize storage resources and forecast capacity.

Aspera for Isilon

Aspera moves large files over long distances fast. Aspera for Isilon is a cluster-aware version of Aspera technology for non-disruptive, wide-area content delivery.

HDFS

OneFS works with the Hadoop Distributed File System protocol to help clients running Apache Hadoop, a framework for data-intensive distributed applications, analyze big data.

SmartQuotas

The SmartQuotas module tracks disk usage with reports and enforces storage limits with alerts.

CHAPTER 3

Introduction to the OneFS command-line interface

This section contains the following topics:

- [OneFS command-line interface overview](#) 40
- [Syntax diagrams](#) 40
- [Universal options](#) 41
- [Command-line interface privileges](#) 41
- [SmartLock compliance command permissions](#) 45
- [OneFS time values](#) 48

OneFS command-line interface overview

The OneFS command-line interface extends the standard UNIX command set to include commands that enable you to manage an Isilon cluster outside of the web administration interface or LCD panel. You can access the command-line interface by opening a secure shell (SSH) connection to any node in the cluster.

You can run `isi` commands to configure, monitor, and manage Isilon clusters and the individual nodes in a cluster. Brief descriptions, usage information, and examples are provided for each command.

Syntax diagrams

The format of each command is described in a syntax diagram.

The following conventions apply for syntax diagrams:

Element	Description
[]	Square brackets indicate an optional element. If you omit the contents of the square brackets when specifying a command, the command still runs successfully.
< >	Angle brackets indicate a placeholder value. You must replace the contents of the angle brackets with a valid value, otherwise the command fails.
{ }	Braces indicate a group of elements. If the contents of the braces are separated by a vertical bar, the contents are mutually exclusive. If the contents of the braces are not separated by a bar, the contents must be specified together.
	Vertical bars separate mutually exclusive elements within the braces.
...	Ellipses indicate that the preceding element can be repeated more than once. If ellipses follow a brace or bracket, the contents of the braces or brackets can be repeated more than once.

Each `isi` command is broken into three parts: command, required options, and optional options. Required options are positional, meaning that you must specify them in the order that they appear in the syntax diagram. However, you can specify a required option in an alternative order by preceding the text displayed in angle brackets with a double dash. For example, consider `isi snapshot snapshots create`.

```
isi snapshot snapshots create <name> <path>
  [--expires <timestamp>]
  [--alias <string>]
  [--verbose]
```

If the `<name>` and `<path>` options are prefixed with double dashes, the options can be moved around in the command. For example, the following command is valid:

```
isi snapshot snapshots create --verbose --path /ifs/data --alias
newSnap_alias --name newSnap
```

Shortened versions of commands are accepted as long as the command is unambiguous and does not apply to multiple commands. For example, `isi snap snap c newSnap /ifs/data` is equivalent to `isi snapshot snapshots create newSnap /ifs/data` because the root of each word belongs to one command.

exclusively. If a word belongs to more than one command, the command fails. For example, `isi sn snap c newSnap /ifs/data` is not equivalent to `isi snapshot snapshots create newSnap /ifs/data` because the root of `isi sn` could belong to either `isi snapshot` or `isi snmp`.

If you begin typing a word and then press TAB, the rest of the word automatically appears as long as the word is unambiguous and applies to only one command. For example, `isi snap` completes to `isi snapshot` because that is the only valid possibility. However, `isi sn` does not complete, because it is the root of both `isi snapshot` and `isi snmp`.

Universal options

Some options are valid for all commands.

Syntax

```
isi [--timeout <integer>] [--debug] <command> [--help]
```

--timeout <integer>

Specifies the number of seconds before the command times out.

--debug

Displays all calls to the Isilon OneFS Platform API. If a traceback occurs, displays traceback in addition to error message.

--help

Displays a basic description of the command and all valid options for the command.

Examples

The following command causes the `isi sync policies list` command to timeout after 30 seconds:

```
isi --timeout 30 sync policies list
```

The following command displays help output for `isi sync policies list`:

```
isi sync policies list --help
```

Command-line interface privileges

You can perform most tasks granted by a privilege through the command-line interface.

Some OneFS commands require root access; however, if you do not have root access, most of the commands associated with a privilege can be performed through the `sudo` program. The system automatically generates a sudoers file of users based on existing roles.

Prefixing a command with `sudo` allows you to run commands that require root access. For example, if you do not have root access, the following command fails:

```
isi sync policies list
```

However, if you are on the sudoers list, the following command succeeds:

```
sudo isi sync policies list
```

The following tables list all One FS commands available, the associated privilege or root-access requirement, and whether `sudo` is required to run the command.

Note

If you are running in compliance mode, more commands will require `sudo`.

Table 1 Privileges sorted by CLI command

isi command	Privilege	Requires sudo
isi alert	ISI_PRIV_EVENT	x
isi audit	ISI_PRIV_AUDIT	
isi auth - excluding isi auth role	ISI_PRIV_AUTH	
isi auth role	ISI_PRIV_ROLE	
isi avscan	ISI_PRIV_ANTIVIRUS	x
isi batterystatus	ISI_PRIV_STATISTICS	x
isi config	root	
isi dedupe - excluding isi dedupe stats	ISI_PRIV_JOB_ENGINE	
isi dedupe stats	ISI_PRIV_STATISTICS	
isi devices	ISI_PRIV_DEVICES	x
isi domain	root	
isi email	ISI_PRIV_CLUSTER	x
isi events	ISI_PRIV_EVENT	x
isi exttools	root	
isi fc	root	
isi filepool	ISI_PRIV_SMARTPOOLS	
isi firmware	root	
isi ftp	ISI_PRIV_FTP	x
isi get	root	
isi hdfs	root	
isi iscsi	ISI_PRIV_ISCSI	x
isi job	ISI_PRIV_JOB_ENGINE	
isi license	ISI_PRIV_LICENSE	x
isi lun	ISI_PRIV_ISCSI	x
isi ndmp	ISI_PRIV_NDMP	x
isi networks	ISI_PRIV_NETWORK	x
isi nfs	ISI_PRIV_NFS	

Table 1 Privileges sorted by CLI command (continued)

isi command	Privilege	Requires sudo
isi perfstat	ISI_PRIV_STATISTICS	x
isi pkg	root	
isi quota	ISI_PRIV_QUOTA	
isi readonly	root	
isi remotesupport	ISI_PRIV_REMOTE_SUPPORT	
isi servicelight	ISI_PRIV_DEVICES	x
isi services	root	
isi set	root	
isi smartlock	root	
isi smb	ISI_PRIV_SMB	
isi snapshot	ISI_PRIV_SNAPSHOT	
isi snmp	ISI_PRIV_SNMP	x
isi stat	ISI_PRIV_STATISTICS	x
isi statistics	ISI_PRIV_STATISTICS	x
isi status	ISI_PRIV_STATISTICS	x
isi storagepool	ISI_PRIV_SMARTPOOLS	
isi sync	ISI_PRIV_SYNCIQ	
isi tape	ISI_PRIV_NDMP	x
isi target	ISI_PRIV_ISCSI	x
isi update	root	
isi version	ISI_PRIV_CLUSTER	x
isi worm	root	
isi zone	ISI_PRIV_AUTH	

Table 2 CLI commands sorted by privilege

Privilege	isi commands	Requires sudo
ISI_PRIV_ANTIVIRUS	isi avscan	x
ISI_PRIV_AUDIT	isi audit	
ISI_PRIV_AUTH	<ul style="list-style-type: none"> • isi auth - excluding isi auth role • isi zone 	
ISI_PRIV_CLUSTER	<ul style="list-style-type: none"> • isi email • isi version 	x

Table 2 CLI commands sorted by privilege (continued)

Privilege	isi commands	Requires sudo
ISI_PRIV_DEVICES	<ul style="list-style-type: none"> isi devices isi servicelight 	x
ISI_PRIV_EVENT	<ul style="list-style-type: none"> isi alert isi events 	x
ISI_PRIV_FTP	isi ftp	x
ISI_PRIV_ISCSI	<ul style="list-style-type: none"> isi iscsi isi lun isi target 	x
ISI_PRIV_JOB_ENGINE	<ul style="list-style-type: none"> isi job isi dedupe -excluding isi dedupe stats 	
ISI_PRIV_LICENSE	isi license	x
ISI_PRIV_NDMP	<ul style="list-style-type: none"> isi ndmp isi tape 	x
ISI_PRIV_NETWORK	isi networks	x
ISI_PRIV_NFS	isi nfs	
ISI_PRIV_QUOTA	isi quota	
ISI_PRIV_ROLE	isi auth role	
ISI_PRIV_REMOTE_SUPPORT	isi remotesupport	
ISI_PRIV_SMARTPOOLS	<ul style="list-style-type: none"> isi filepool isi storagepool 	
ISI_PRIV_SMB	isi smb	
ISI_PRIV_SNAPSHOT	isi snapshot	
ISI_PRIV_SNMP	isi snmp	x
ISI_PRIV_STATISTICS	<ul style="list-style-type: none"> isi batterystatus isi dedupe stats isi perfstat isi stat isi statistics isi status 	x
ISI_PRIV_SYNCIQ	isi sync	
root	<ul style="list-style-type: none"> isi config isi domain 	

Table 2 CLI commands sorted by privilege (continued)

Privilege	isi commands	Requires sudo
	<ul style="list-style-type: none"> • <code>isi extttools</code> • <code>isi fc</code> • <code>isi firmware</code> • <code>isi get</code> • <code>isi hdfs</code> • <code>isi pkg</code> • <code>isi readonly</code> • <code>isi services</code> • <code>isi set</code> • <code>isi smartlock</code> • <code>isi update</code> • <code>isi worm</code> 	

SmartLock compliance command permissions

If a cluster is running in SmartLock compliance mode, root access is disabled on the cluster. Because of this, if a command requires root access, you can run the command only through the sudo program.

In compliance mode, you can run all `isi` commands that are followed by a space through `sudo`. For example, you can run `isi sync policies create` through `sudo`. In addition, you can also run the following `isi_` commands through `sudo`; these commands are internal and are typically run only by Isilon Technical Support:

- `isi_bootdisk_finish`
- `isi_bootdisk_provider_dev`
- `isi_bootdisk_status`
- `isi_bootdisk_unlock`
- `isi_checkjournal`
- `isi_clean_idmap`
- `isi_client_stats`
- `isi_cpr`
- `isi_cto_update`
- `isi_disk_firmware_reboot`
- `isi_dmi_info`
- `isi_dmilog`
- `isi_dongle_sync`
- `isi_drivenum`
- `isi_dsp_install`

- `isi_dumpjournal`
- `isi_eth_mixer_d`
- `isi_evaluate_provision_drive`
- `isi_fcb_vpd_tool`
- `isi_flexnet_info`
- `isi_flush`
- `isi_for_array`
- `isi_fputil`
- `isi_gather_info`
- `isi_gather_auth_info`
- `isi_gather_cluster_info`
- `isi_gconfig`
- `isi_get_itrace`
- `isi_get_profile`
- `isi_hangdump`
- `isi_hw_check`
- `isi_hw_status`
- `isi_ib_bug_info`
- `isi_ib_fw`
- `isi_ib_info`
- `isi_ilog`
- `isi_imdd_status`
- `isi_inventory_tool`
- `isi_ipmicmc`
- `isi_job_d`
- `isi_kill_busy`
- `isi_km_diag`
- `isi_lid_d`
- `isi_linmap_mod`
- `isi_logstore`
- `isi_lsiexputil`
- `isi_make_abr`
- `isi_mcp`
- `isi_mps_fw_status`
- `isi_netlogger`
- `isi_nodes`
- `isi_ntp_config`
- `isi_ovt_check`
- `isi_patch_d`
- `isi_promptsupport`

- `isi_radish`
- `isi_rbm_ping`
- `isi_repstate_mod`
- `isi_restill`
- `isi_rnvutil`
- `isi_sasphymon`
- `isi_save_itrace`
- `isi_savecore`
- `isi_sed`
- `isi_send_abr`
- `isi_smbios`
- `isi_stats_tool`
- `isi_transform_tool`
- `isi_uftp`
- `isi_umount_ifs`
- `isi_update_cto`
- `isi_update_serialno`
- `isi_vitutil`
- `isi_vol_copy`
- `isi_vol_copy_vnx`

In addition to `isi` commands, you can run the following UNIX commands through `sudo`:

- `date`
- `gcore`
- `ifconfig`
- `kill`
- `killall`
- `nfsstat`
- `ntpdate`
- `nvmecontrol`
- `pciconf`
- `pkill`
- `ps`
- `renice`
- `shutdown`
- `sysctl`
- `tcpdump`
- `top`

OneFS time values

OneFS uses different values for time depending on the application.

You can specify time periods, such as a month, for multiple OneFS applications. However, because some time values have more than one meaning, OneFS defines time values based on the application. The following table describes the time values for OneFS applications:

Module	Month	Year
SnapshotIQ	30 days	365 days (does not account for leap year)
SmartLock	31 days	365 days (does not account for leap year)
SyncIQ	30 days	365 days (does not account for leap year)

CHAPTER 4

General cluster administration

This section contains the following topics:

- [General cluster administration overview](#)..... 50
- [User interfaces](#)..... 50
- [Connecting to the cluster](#)..... 51
- [Licensing](#)..... 51
- [Certificates](#)..... 56
- [Cluster identity](#)..... 58
- [Cluster contact information](#)..... 59
- [Cluster date and time](#)..... 60
- [SMTP email settings](#)..... 61
- [Configuring the cluster join mode](#)..... 62
- [File system settings](#)..... 63
- [Cluster monitoring](#)..... 64
- [Monitoring cluster hardware](#)..... 65
- [Events and notifications](#)..... 72
- [Cluster maintenance](#)..... 82
- [Remote support](#)..... 90
- [Cluster administration commands](#)..... 95
- [Event commands](#)..... 149
- [Hardware commands](#)..... 156

General cluster administration overview

You can manage general OneFS settings and module licenses for the EMC Isilon cluster.

General cluster administration covers several areas. You can manage general settings such as cluster name, date and time, and email. You can monitor the cluster status and performance, including hardware components. You can configure how events and notifications are handled, and you can perform cluster maintenance such as adding, removing, and restarting nodes.

Most management tasks are accomplished through both the web administration or command-line interface; however, you will occasionally encounter a task that can only be managed by one or the other.

User interfaces

OneFS provides several interfaces for managing the EMC Isilon cluster.

Interface	Description	Comment
OneFS web administration interface	The browser-based OneFS web administration interface provides secure access with OneFS-supported browsers. Use this interface to view robust graphical monitoring displays and to perform cluster-management tasks.	The OneFS web administration interface uses port 8080 as its default port.
OneFS command-line interface	Run OneFS <code>isi</code> commands in the command-line interface to configure, monitor, and manage the cluster. Access to the command-line interface is through a secure shell (SSH) connection to any node in the cluster.	The OneFS command-line interface provides an extended standard UNIX command set for managing the cluster.
OneFS API	The OneFS application programming interface (API) is divided into two functional areas: one area enables cluster configuration, management, and monitoring functionality, and the other area enables operations on files and directories on the cluster. You can send requests to the OneFS API through a Representational State Transfer (REST) interface, which is accessed through resource URIs and standard HTTP methods.	You should have a solid understanding of HTTP/1.1 and experience writing HTTP-based client software before you implement client-based software through the OneFS API.
Node front panel	With the exception of accelerator nodes, the front panel of each node contains an LCD screen with five buttons that you can use to monitor node and cluster details.	Node status, events, cluster details, capacity, IP and MAC addresses, throughput, and drive status are available through the node front panel.

Connecting to the cluster

EMC Isilon cluster access is provided through the web administration interface or through SSH. You can use a serial connection to perform cluster-administration tasks through the command-line interface.

You can also access the cluster through the node front panel to accomplish a subset of cluster-management tasks. For information about connecting to the node front panel, see the installation documentation for your node.

Log in to the web administration interface

You can monitor and manage your EMC Isilon cluster from the browser-based web administration interface.

Procedure

1. Open a browser window and type the URL for your cluster in the address field, replacing *<yourNodeIPAddress>* in the following example with the first IP address you provided when you configured ext-1:

```
https://<yourNodeIPAddress>:8080
```

The system displays a message if your security certificates have not been configured. Resolve any certificate configurations and continue to the web site.

2. Log in to OneFS by typing your OneFS credentials in the **Username** and **Password** fields.

After you log into the web administration interface, there is a 4-hour login timeout and a 24-hour session inactivity timeout.

Open an SSH connection to a cluster

You can use any SSH client such as OpenSSH or PuTTY to connect to an EMC Isilon cluster.

Before you begin

You must have valid OneFS credentials to log in to a cluster after the connection is open.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster, using the IP address of the node and port number 22.
2. Log in with your OneFS credentials.

At the OneFS command line prompt, you can use `isi` commands to monitor and manage your cluster.

Licensing

Advanced cluster features are available when you activate licenses for OneFS software modules. Each optional OneFS software module requires you to activate a separate license.

For more information about the following optional software modules, contact your EMC Isilon sales representative.

- HDFS

- InsightIQ
- Isilon for vCenter
- SmartConnect Advanced
- SmartDedupe
- SmartLock
- SmartPools
- SmartQuotas
- SnapshotIQ
- SyncIQ

License status

The status of a OneFS module license indicates whether the functionality provided by a module is available on the cluster.

Licenses exist in one of the following states:

Status	Description
Inactive	The license has not been activated on the cluster. You cannot access the features provided by the corresponding module.
Evaluation	The license has been temporarily activated on the cluster. You can access the features provided by the corresponding module for a limited period of time. After the license expires, the features become unavailable unless the license is reactivated.
Activated	The license has been activated on the cluster. You can access the features provided by the corresponding module.
Expired	The evaluation license has expired on the cluster. You can no longer access the features provided by the corresponding module. The features will remain unavailable unless you reactivate the license.

The following table describes what functionality is available for each license depending on the license's status:

License	Inactive	Evaluation/ Activated	Expired
HDFS	Clients cannot access the cluster through HDFS.	You can configure HDFS settings and clients can access the cluster through HDFS.	You cannot configure HDFS settings. After the HDFS service restarts, clients can no longer access the cluster through HDFS.
InsightIQ	You cannot monitor the cluster with InsightIQ.	You can monitor the cluster with InsightIQ.	InsightIQ stops monitoring the cluster. Data previously collected by InsightIQ is still available on the InsightIQ instance.
Isilon for vCenter	You cannot back up virtual machines that are stored on an	You can back up virtual machines that are stored on an	You cannot create new backups of virtual machines

License	Inactive	Evaluation/ Activated	Expired
	Isilon cluster with Isilon for vCenter.	Isilon cluster with Isilon for vCenter.	that are stored on an Isilon cluster.
SmartPools	All files belong to the default file pool and are governed by the default file pool policy. Virtual hot spare allocation, which reserves space for data repair if a drive fails, is also available.	You can create multiple file pools and file pool policies. You can also manage spillover, which defines how write operations are handled when a storage pool is not writable.	You can no longer manage file pool policies, and the SmartPools job will no longer run. Newly added files will be governed by the default file pool policy, and the SetProtectPlus job will eventually apply the default file pool policy to all files in the cluster. If the SmartPools job is running when the license expires, the job completes before becoming disabled.
SmartConnect Advanced	Client connections are balanced by using a round robin policy. IP address allocation is static. Each external network subnet can be assigned only one IP address pool.	You can access features such as CPU utilization, connection counting, and client connection policies in addition to the round robin policy. You can also configure address pools to support multiple DNS zones within a single subnet, and support IP failover.	You can no longer specify SmartConnect Advanced settings.
SmartDedupe	You cannot deduplicate data with SmartDedupe.	You can deduplicate data with SmartDedupe.	You can no longer deduplicate data. Previously deduplicated data remains deduplicated.
SmartLock	You cannot enforce file retention with SmartLock.	You can enforce file retention with SmartLock.	You cannot create new SmartLock directories or modify SmartLock directory configuration settings for existing directories. You can still commit files to a write once read many (WORM) state, even after the SmartLock license is unconfigured, but you cannot delete WORM-committed files from enterprise directories.

License	Inactive	Evaluation/ Activated	Expired
SnapshotIQ	You can view and manage snapshots generated by OneFS applications. However, you cannot create snapshots or configure SnapshotIQ settings.	You can create, view, and manage snapshots. You can also configure snapshot settings.	You will no longer be able to generate snapshots. Existing snapshot schedules are not deleted; however, the schedules will not generate snapshots. You can still delete snapshots and access snapshot data.
SmartQuotas	You cannot create quotas with SmartQuotas.	You can create quotas with SmartQuotas.	OneFS disables all quotas. Exceeding advisory and soft thresholds does not trigger events. Hard and soft thresholds are not enforced.
SyncIQ	You cannot replicate data with SyncIQ.	You can replicate data with SyncIQ	You will no longer be able to replicate data to remote clusters, and remote clusters will not be able to replicate data to the local cluster. Replication policies will still display a status of enabled; however, future replication jobs created by the policy will fail. If a replication job is in progress when the license expires, the job completes.

License configuration

You can configure or unconfigure some OneFS module licenses.

You can configure a license by performing specific operations through the corresponding module. Not all actions that require you to activate a license will configure the license. Also, not all licenses can be configured. Configuring a license does not add or remove access to any features provided by a module.

You can unconfigure a license only through the `isi license unconfigure` command. You may want to unconfigure a license for a OneFS software module if, for example, you enabled an evaluation version of a module but later decided not to purchase a permanent license. Unconfiguring a module license does not deactivate the license. Unconfiguring a license does not add or remove access to any features provided by a module.

The following table describes both the actions that cause each license to be configured and the results of unconfiguring each license:

License	Cause of configuring	Result of unconfiguring
HDFS	Cannot configure this license.	No system impact.
InsightIQ	Cannot configure this license.	No system impact.

License	Cause of configuring	Result of unconfiguring
Isilon for vCenter	Cannot configure this license.	No system impact.
SmartPools	Create a file pool policy (other than the default file pool policy).	OneFS deletes all file pool policies (except the default file pool policy).
SmartConnect	Configure SmartConnect Advanced settings for at least one IP address pool.	OneFS converts dynamic IP address pools to static IP address pools.
SmartDedupe	Cannot configure this license.	No system impact.
SmartLock	Cannot configure this license.	No system impact.
SnapshotIQ	Create a snapshot schedule.	Deletes all snapshot schedules.
SmartQuotas	Create a quota.	No system impact.
SyncIQ	Create a replication policy.	No system impact.

Activate a license through the command-line interface

You can activate licenses to access optional OneFS modules, which provide advanced cluster features.

Before you begin

Before you can activate a license, you must obtain a valid license key, and you must have root user privileges on your cluster. To obtain a license key, contact your EMC Isilon sales representative.

Procedure

1. Run the `isi license activate` command.

The following command activates a license:

```
isi license activate 1UL9A83P1209Q378C12M0938
```

View license information

You can view information about the current status of any optional Isilon software modules.

Procedure

1. Run the following command:

```
isi license status
```

Unconfigure a license

You can unconfigure a licensed module through the command-line interface.

You must have root user privileges on your Isilon cluster to unconfigure a module license. This procedure is available only through the command-line interface (CLI).

Note

Unconfiguring a license does not deactivate the license.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster.
You must log in as root.
2. Run the `isi license unconfigure` command.

The following command unconfigures the license for SmartConnect:

```
isi license unconfigure -m smartconnect
```

If you do not know the module name, run the `isi license` command for a list of OneFS modules and their status.

OneFS returns a confirmation message similar to the following text: The `SmartConnect` module has been unconfigured. The license is unconfigured, and any processes enabled for the module are disabled.

Certificates

You can renew the Secure Sockets Layer (SSL) certificate for the Isilon web administration interface or replace it with a third-party SSL certificate.

All Platform API communication, which includes communication through the web administration interface, is over SSL. You can replace or renew the self-signed certificate with a certificate that you generate. To replace or renew an SSL certificate, you must be logged in as root.

Replace or renew the SSL certificate

You can replace or renew the Secure Sockets Layer (SSL) certificate, which is used to access the EMC Isilon cluster through a browser.

Before you begin

When you renew or replace a self-signed SSL certificate, you must provide information for your organization in the format that is described in the Self-signed SSL certificate data example.

The following folders are the default locations for the `server.crt` and `server.key` files in OneFS 6.0 and higher.

- SSL certificate: `/usr/local/apache2/conf/ssl.crt/server.crt`
- SSL certificate key: `/usr/local/apache2/conf/ssl.key/server.key`

Procedure

1. Establish an SSH connection to any node in the cluster.
2. At the command prompt, run the following command to create the appropriate directory.

```
mkdir /ifs/local/
```

3. At the command prompt, run the following command to change to the directory.

```
cd /ifs/local/
```


4. Choose the type of certificate you want to install.

Option	Description
Third-party (public or private) CA-issued certificate	<p>a. At the command prompt, run the following command to generate a new Certificate Signing Request (CSR) in addition to a new key, where <i><common_name></i> is the host name, such as <code>isilon.example.com</code>:</p> <pre>openssl req -new -nodes -newkey rsa:1024 -keyout <common_name>.key \ -out <common-name>.csr</pre> <p>b. Send the contents of the <i><common_name>.csr</i> file from the cluster to your Certificate Authority (CA) for signing. When you receive the signed certificate (now a <code>.crt</code> file) from the CA, copy the certificate to <code>/ifs/local/<common-name>.crt</code>.</p>
Self-signed certificate based on the existing (stock) <code>ssl.key</code>	<p>a. At the command prompt, run the following command to create a two-year certificate. Increase or decrease the value for <code>-days</code> to generate a certificate with a different expiration date.</p> <pre>cp /usr/local/apache2/conf/ssl.key/server.key ./openssl req -new \ -days 730 -nodes -x509 -key server.key -out server.crt</pre>

A renewal certificate is created, based on the existing (stock) `ssl.key` file.

5. (Optional) At the command prompt, run the following command to verify the attributes in an SSL certificate.

```
openssl x509 -text -noout -in <common-name>.crt
```

6. Run the following commands to install the certificate and key:

```
isi services -a isi_webui disable
chmod 640 <common_name>.key
isi_for_array -s 'cp /ifs/local/<common-name>.key /usr/local/
apache2/conf/ssl.key/server.key'
isi_for_array -s 'cp /ifs/local/<common-name>.crt /usr/local/
apache2/conf/ssl.crt/server.crt'
isi services -a isi_webui enable
```

7. Run the following command to remove the files in `/ifs/local`.

```
rm /ifs/local/*
```

Verify an SSL certificate update

You can verify the details stored in a Secure Sockets Layer (SSL) certificate.

Procedure

1. Establish an SSH connection to any node in the cluster.

- At the command prompt, run the following command to open and verify the attributes in an SSL certificate.

```
echo QUIT | openssl s_client -connect localhost:8080
```

Self-signed SSL certificate data example

Self-signed SSL certificate renewal or replacement requires you to provide data such as your fully qualified domain name and a contact email address.

When you renew or replace a self-signed SSL certificate, you are asked to provide data in the format shown in the following example. Some fields in the certificate file contain a default value. If you type ' . ', the field is left blank when the certificate is generated.

- Country Name (2 letter code) [XX]:**US**
- State or Province Name (full name) [Some-State]:**Washington**
- Locality Name (for example, city) [default city]:**Seattle**
- Organization Name (for example, company) [Internet Widgits Pty Ltd]:**Isilon**
- Organizational Unit Name (for example, section) []:**Support**
- Common Name (for example, server FQDN or server name) []:**isilon.example.com**
- Email Address []:**support@example.com**

In addition, you should add the following attributes to be sent with your certificate request:

- Challenge password []:**Isilon1**
- Optional company name []:

Cluster identity

You can specify identity attributes for the EMC Isilon cluster.

Cluster name

The cluster name appears on the login page, and it makes the cluster and its nodes more easily recognizable on your network. Each node in the cluster is identified by the cluster name plus the node number. For example, the first node in a cluster named Images may be named Images-1.

Cluster description

The cluster description appears below the cluster name on the login page. The cluster description is useful if your environment has multiple clusters.

Login message

The login message appears as a separate box on the login page. The login message can convey cluster information, login instructions, or warnings that a user should know before logging into the cluster.

Set the cluster name

You can assign a name and add a login message to your EMC Isilon cluster to make the cluster and its nodes more easily recognizable on your network.

Cluster names must begin with a letter and can contain only numbers, letters, and hyphens. The cluster name is added to the node number to identify each node in the cluster. For example, the first node in a cluster named Images may be named Images-1.

Procedure

1. Open the `isi config` command prompt by running the following command:

```
isi config
```

2. Run the `name` command.

The following command sets the name of the cluster to `NewName`:

```
name NewName
```

3. Save your changes by running the following command:

```
commit
```

Cluster contact information

Isilon Technical Support personnel and event notification recipients will communicate with the specified contacts.

You can specify the following contact information for your EMC Isilon cluster:

- Company name and location
- Primary and secondary contact names
- Phone number and email address for each contact

Specify contact information

You can specify contact information so that Isilon Technical Support personnel and event notification recipients can contact you.

SupportIQ is the contact mechanism and must be enabled in order to specify contact information. Please enter company name

Procedure

1. Enable SupportIQ by running the following command:

```
isi_promptsupport -e
```

The system displays the following message:

```
Would you like to enable SupportIQ? [yes]
```

2. Type **yes** and then press ENTER.

The system displays the following message:

```
Please enter company name:
```

3. Type your company name and then press ENTER.

The system displays the following message:

```
Please enter contact name:
```

4. Type your contact name and then press ENTER.

The system displays the following message:

```
Please enter contact phone:
```

5. Type your contact phone and then press ENTER.

The system displays the following message:

```
Please enter contact email:
```

6. Type your contact email address and then press ENTER.

Cluster date and time

The Network Time Protocol (NTP) service is configurable manually, so you can ensure that all nodes in a cluster are synchronized to the same time source.

The NTP method automatically synchronizes cluster date and time settings through an NTP server. Alternatively, you can set the date and time reported by the cluster by manually configuring the service.

Windows domains provide a mechanism to synchronize members of the domain to a master clock running on the domain controllers, so OneFS adjusts the cluster time to that of Active Directory with a service. If there are no external NTP servers configured, OneFS uses the Windows domain controller as the NTP time server. When the cluster and domain time become out of sync by more than 4 minutes, OneFS generates an event notification.

Note

If the cluster and Active Directory become out of sync by more than 5 minutes, authentication will not work.

Set the cluster date and time

You can set the date, time, and time zone that is used by the EMC Isilon cluster.

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem.

2. Specify the current date and time by running the `date` command.

The following command sets the cluster time to 9:47 AM on July 22, 2015:

```
date 2015/07/22 09:47:00
```

3. To verify your time zone setting, run the `timezone` command. The current time zone setting displays. For example:

```
The current time zone is: Pacific Time Zone
```

4. To view a list of valid time zones, run the `help timezone` command. The following options display:

```
Greenwich Mean Time
Eastern Time Zone
Central Time Zone
Mountain Time Zone
Pacific Time Zone
```

```
Arizona
Alaska
Hawaii
Japan
Advanced
```

5. To change the time zone, enter the `timezone` command followed by one of the displayed options.

The following command changes the time zone to `Hawaii`:

```
timezone Hawaii
```

A message confirming the new time zone setting displays. If your desired time zone did not display when you ran the `help timezone` command, enter **timezone Advanced**. After a warning screen, you will proceed to a list of regions. When you select a region, a list of specific time zones for that region appears. Select the desired time zone (you may need to scroll), then enter **OK** or **Cancel** until you return to the `isi config` prompt.

6. Run the `commit` command to save your changes and exit `isi config`.

Specify an NTP time server

You can specify one or more Network Time Protocol (NTP) servers to synchronize the system time on the EMC Isilon cluster. The cluster periodically contacts the NTP servers and sets the date and time based on the information it receives.

Procedure

1. Run the `isi_ntp_config` command.

The following command specifies `ntp.time.server1.com`:

```
isi_ntp_config add server ntp.time.server1.com
```

SMTP email settings

If your network environment requires the use of an SMTP server or if you want to route EMC Isilon cluster event notifications with SMTP through a port, you can configure SMTP email settings.

SMTP settings include the SMTP relay address and port number that email is routed through. You can specify an origination email and subject line for all event notification emails sent from the cluster.

If your SMTP server is configured to support authentication, you can specify a username and password. You can also specify whether to apply encryption to the connection.

Configure SMTP email settings

You can send event notifications through the SMTP mail server. You can also enable SMTP authentication if your SMTP server is configured to use it.

You can configure SMTP email settings if your network environment requires the use of an SMTP server or if you want to route EMC Isilon cluster event notifications with SMTP through a port.

Procedure

1. Run the `isi email` command.

The following example configures SMTP email settings:

```
isi email --mail-relay 10.7.180.45 \
--mail-sender isilon-cluster@company.com \
--mail-subject "Isilon cluster event" --use-smtp-auth yes \
--auth-user SMTPuser --auth-pass Password123 --use-encryption yes
```

View SMTP email settings

You can view SMTP email settings.

Procedure

1. Run the following command:

```
isi email list
```

The system displays information similar to the following example:

```
SMTP relay address:
SMTP relay port:           25
Send email as:
Subject:

Use SMTP AUTH:            No
Username:
Password:                 ****
Use Encryption (TLS):     No
```

Configuring the cluster join mode

The join mode specifies how a node is added to the EMC Isilon cluster and whether authentication is required. OneFS supports manual and secure join modes for adding nodes to the EMC Isilon cluster.

Mode	Description
Manual	Allows you to manually add a node to the cluster without requiring authorization.
Secure	Requires authorization of every node added to the cluster and the node must be added through the web administration interface or through the <code>isi devices -a add -d <unconfigured_node_serial_no></code> command in the command-line interface.
	<p>Note</p> <p>If you specify a secure join mode, you cannot join a node to the cluster through serial console wizard option [2] Join an existing cluster.</p>

Specify the cluster join mode

You can specify the method to use when nodes are added to the EMC Isilon cluster.

Procedure

1. Open the `isi config` command prompt by running the following command:

```
isi config
```

2. Run the `joinmode` command.

The following command prevents nodes from joining the cluster unless the join is initiated by the cluster:

```
joinmode secure
```

3. Save your changes by running the following command:

```
commit
```

File system settings

You can configure global file system settings on an EMC Isilon cluster pertaining to access time tracking and character encoding.

You can enable or disable access time tracking, which monitors the time of access on each file. If necessary, you can also change the default character encoding on the cluster.

Specify the cluster character encoding

You can modify the character encoding set for the EMC Isilon cluster after installation.

Only OneFS-supported character sets are available for selection. UTF-8 is the default character set for OneFS nodes.

Note

If the cluster character encoding is not set to UTF-8, SMB share names are case-sensitive.

You must restart the cluster to apply character encoding changes.

CAUTION

Character encoding is typically established during installation of the cluster. Modifying the character encoding setting after installation may render files unreadable if done incorrectly. Modify settings only if necessary after consultation with Isilon Technical Support

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem.

2. Modify the character encoding by running the `encoding` command.

The following command sets the encoding for the cluster to ISO-8859-1:

```
encoding ISO-8859-1
```

3. Run the `commit` command to save your changes and exit the `isi config` subsystem.
4. Restart the cluster to apply character encoding modifications.

Enable or disable access time tracking

You can enable access time tracking to support features that require it.

By default, the EMC Isilon cluster does not track the timestamp when files are accessed. You can enable this feature to support OneFS features that use it. For example, access-time tracking must be enabled to configure SyncIQ policy criteria that match files based on when they were last accessed.

Note

Enabling access-time tracking may affect cluster performance.

Procedure

1. Enable or disable access time tracking by setting the `atime_enabled` system control.

- To enable access time tracking, run the following command:

```
sysctl efs.bam.atime_enabled=1
```

- To disable access time tracking, run the following command:

```
sysctl efs.bam.atime_enabled=0
```

2. To specify how often to update the last-accessed time, set the `atime_grace_period` system control.

Specify the amount of time as a number of seconds.

The following command configures OneFS to update the last-accessed time every two weeks:

```
sysctl efs.bam.atime_grace_period=1209600
```

Cluster monitoring

You can view health and status information for the EMC Isilon cluster and monitor cluster and node performance.

Run the `isi status` command to review the following information:

- Cluster, node, and drive health
- Storage data such as size and amount used
- IP addresses
- Throughput
- Critical events
- Job status

Additional commands are available to review performance information for the following areas:

- General cluster statistics
- Statistics by protocol or by clients connected to the cluster
- Performance data by drive
- Historical performance data

Advanced performance monitoring and analytics are available through the InsightIQ module, which requires you to activate a separate license. For more information about optional software modules, contact your EMC Isilon sales representative.

Monitor the cluster

You can monitor the health and performance of a cluster with charts and tables.

Procedure

1. Run the following command:

```
isi status
```

View node status

You can view the status of a node.

Procedure

1. (Optional) Run the `isi status` command:

The following command displays information about a node with a logical node number (LNN) of 1:

```
isi status -n 1
```

Monitoring cluster hardware

You can manually check the status of hardware on the EMC Isilon cluster as well as enable SNMP to remotely monitor components.

View node hardware status

You can view the hardware status of a node.

Procedure

1. Click **Dashboard** > **Cluster Overview** > **Cluster Status**.
2. (Optional) In the **Status** area, click the ID number for a node.
3. In the **Chassis and drive status** area, click **Platform**.

Chassis and drive states

You can view chassis and drive state details.

In a cluster, the combination of nodes in different degraded states determines whether read requests, write requests, or both work. A cluster can lose write quorum but keep read quorum. OneFS provides details about the status of chassis and drives in your cluster. The following table describes all the possible states that you may encounter in your cluster.

State	Description	Interface	Error state
HEALTHY	All drives in the node are functioning correctly.	Command-line interface, web	

State	Description	Interface	Error state
		administration interface	
SMARTFAIL or Smartfail or restripe in progress	The drive is in the process of being removed safely from the file system, either because of an I/O error or by user request. Nodes or drives in a smartfail or read-only state affect only write quorum.	Command-line interface, web administration interface	
NOT AVAILABLE	<p>A drive is unavailable for a variety of reasons. You can click the bay to view detailed information about this condition.</p> <hr/> <p>Note</p> <p>In the web administration interface, this state includes the ERASE and SED_ERROR command-line interface states.</p> <hr/>	Command-line interface, web administration interface	X
SUSPENDED	This state indicates that drive activity is temporarily suspended and the drive is not in use. The state is manually initiated and does not occur during normal cluster activity.	Command-line interface, web administration interface	
NOT IN USE	A node in an offline state affects both read and write quorum.	Command-line interface, web administration interface	
REPLACE	The drive was smartfailed successfully and is ready to be replaced.	Command-line interface only	
STALLED	The drive is stalled and undergoing stall evaluation. Stall evaluation is the process of checking drives that are slow or having other issues. Depending on the outcome of the evaluation, the drive may return to service or be smartfailed. This is a transient state.	Command-line interface only	
NEW	The drive is new and blank. This is the state that a drive is in when you run the <code>isi dev</code> command with the <code>-a add</code> option.	Command-line interface only	
USED	The drive was added and contained an Isilon GUID but the drive is not from this node. This drive likely will be formatted into the cluster.	Command-line interface only	

State	Description	Interface	Error state
PREPARING	The drive is undergoing a format operation. The drive state changes to HEALTHY when the format is successful.	Command-line interface only	
EMPTY	No drive is in this bay.	Command-line interface only	
WRONG_TYPE	The drive type is wrong for this node. For example, a non-SED drive in a SED node, SAS instead of the expected SATA drive type.	Command-line interface only	
BOOT_DRIVE	Unique to the A100 drive, which has boot drives in its bays.	Command-line interface only	
SED_ERROR	The drive cannot be acknowledged by the OneFS system. <hr/> Note In the web administration interface, this state is included in Not available.	Command-line interface, web administration interface	X
ERASE	The drive is ready for removal but needs your attention because the data has not been erased. You can erase the drive manually to guarantee that data is removed. <hr/> Note In the web administration interface, this state is included in Not available.	Command-line interface only	
INSECURE	Data on the self-encrypted drive is accessible by unauthorized personnel. Self-encrypting drives should never be used for non-encrypted data purposes. <hr/> Note In the web administration interface, this state is labeled <code>Unencrypted SED</code> .	Command-line interface only	X
UNENCRYPTED SED	Data on the self-encrypted drive is accessible by unauthorized personnel. Self-encrypting drives should never be used for non-encrypted data purposes.	Web administration interface only	X

State	Description	Interface	Error state
	<hr/> <p>Note</p> <p>In the command-line interface, this state is labeled INSECURE.</p> <hr/>		

Check battery status

You can monitor the status of NVRAM batteries and charging systems.

This functionality is available only from the command line on node hardware that supports the command.

Procedure

1. Run the `isi batterystatus` command to view the status of all NVRAM batteries and charging systems on the node.

The system displays output similar to the following example:

```
battery 1 : Good
battery 2 : Good
```

SNMP monitoring

You can use SNMP to remotely monitor the EMC Isilon cluster hardware components, such as fans, hardware sensors, power supplies, and disks. The default Linux SNMP tools or a GUI-based SNMP tool of your choice can be used for this purpose.

You can enable SNMP monitoring on individual nodes on your cluster, and you can also monitor cluster information from any node. Generated SNMP traps are sent to your SNMP network. You can configure an event notification rule that specifies the network station where you want to send SNMP traps for specific events, so that when an event occurs, the cluster sends the trap to that server. OneFS supports SNMP in read-only mode. OneFS supports SNMP version 2c, which is the default value, and SNMP version 3.

Note

OneFS does not support SNMP v1. Although an option for v1/v2c may be displayed, if you select the v1/v2c pair, OneFS will only monitor through SNMP v2c.

You can configure settings for SNMP v3 alone or for both SNMP v2c and v3.

Note

If you configure SNMP v3, OneFS requires the SNMP-specific security level of AuthNoPriv as the default value when querying the cluster. The security level AuthPriv is not supported.

Elements in an SNMP hierarchy are arranged in a tree structure, similar to a directory tree. As with directories, identifiers move from general to specific as the string progresses from left to right. Unlike a file hierarchy, however, each element is not only named, but also numbered.

For example, the SNMP entity `.iso.org.dod.internet.private.enterprises.isilon.oneFSsss`

`sLocalNodeId.0` maps to `.1.3.6.1.4.1.12124.3.2.0`. The part of the name that refers to the OneFS SNMP namespace is the `12124` element. Anything further to the right of that number is related to OneFS-specific monitoring.

Management Information Base (MIB) documents define human-readable names for managed objects and specify their data types and other properties. You can download MIBs that are created for SNMP-monitoring of an Isilon cluster from the web-administration interface or manage them using the command-line interface. MIBs are stored in `/usr/local/share/snmp/mibs/` on a OneFS node. The OneFS ISILON-MIBs serve two purposes:

- Augment the information available in standard MIBs
- Provide OneFS-specific information that is unavailable in standard MIBs

ISILON-MIB is a registered enterprise MIB. Isilon clusters have two separate MIBs:

ISILON-MIB

Defines a group of SNMP agents that respond to queries from a network monitoring system (NMS) called OneFS Statistics Snapshot agents. As the name implies, these agents snapshot the state of the OneFS file system at the time that it receives a request and reports this information back to the NMS.

ISILON-TRAP-MIB

Generates SNMP traps to send to an SNMP monitoring station when the circumstances occur that are defined in the trap protocol data units (PDUs).

The OneFS MIB files map the OneFS-specific object IDs with descriptions. Download or copy MIB files to a directory where your SNMP tool can find them, such as `/usr/share/snmp/mibs/` or `/usr/local/share/snmp/mibs/`, depending on the tool that you use.

To enable Net-SNMP tools to read the MIBs to provide automatic name-to-OID mapping, add `-m All` to the command, as in the following example:

```
snmpwalk -v2c -c public -m All <node IP> isilon
```

If the MIB files are not in the default Net-SNMP MIB directory, you may need to specify the full path, as in the following example. Note that all three lines are a single command.

```
snmpwalk -m /usr/local/share/snmp/mibs/ISILON-MIB.txt:/usr/local\
/share/snmp/mibs/ISILON-TRAP-MIB.txt:/usr/local/share/snmp/mibs \
/ONEFS-TRAP-MIB.txt -v2c -C c -c public <node IP> enterprises.onefs
```

Note

The previous examples are run from the `snmpwalk` command on a cluster. Your SNMP version may require different arguments.

Managing SNMP settings

SNMP can be used to monitor cluster hardware and system information. Settings can be configured through either the web administration interface or the command-line interface.

You can enable SNMP monitoring on individual nodes in the cluster, and you can monitor information cluster-wide from any node when you enable SNMP on each node. When using SNMP on an Isilon cluster, you should use a fixed general username. A password for the general user can be configured in the web administration interface.

You should configure a network monitoring system (NMS) to query each node directly through a static IP address. This approach allows you to confirm that all nodes have

external IP addresses and therefore respond to SNMP queries. Because the SNMP proxy is enabled by default, the SNMP implementation on each node is configured automatically to proxy for all other nodes in the cluster except itself. This proxy configuration allows the Isilon Management Information Base (MIB) and standard MIBs to be exposed seamlessly through the use of context strings for supported SNMP versions. After you download and save the appropriate MIBs, you can configure SNMP monitoring through either the web administration interface or through the command-line interface.

Configure SNMP settings

You can configure SNMP monitoring settings

Note

When SNMP v3 is used, OneFS requires the SNMP-specific security level of AuthNoPriv as the default value when querying the EMC Isilon cluster. The security level AuthPriv is not supported.

Procedure

1. Run the `isi snmp` command.

The following command allows access only through SNMP version 3:

```
isi snmp --protocols "v3"
```

Configure the cluster for SNMP monitoring

You can configure your EMC Isilon cluster to remotely monitor hardware components using SNMP.

Before you begin

When SNMP v3 is used, OneFS requires the SNMP-specific security level of AuthNoPriv as the default value when querying the cluster. The security level AuthPriv is not supported.

You can enable or disable SNMP monitoring, allow SNMP access by version, and configure other settings, some of which are optional. All SNMP access is read-only.

Note

The Isilon cluster does not generate SNMP traps unless you configure an event notification rule to send events.

Procedure

1. Click **Cluster Management** > **General Settings** > **SNMP Monitoring**.
2. In the Service area of the SNMP Monitoring page, enable or disable SNMP monitoring.
 - a. To disable SNMP monitoring, click **Disable**, and then click **Submit**.
 - b. To enable SNMP monitoring, click **Enable**, and then continue with the following steps to configure your settings.
3. In the Downloads area, click **Download** for the MIB file that you want to download. Follow the download process that is specific to your browser.
4. (Optional) If you are using Internet Explorer as your browser, right-click the **Download** link, select **Save As** from the menu, and save the file to your local drive.

You can save the text in the file format that is specific to your Net-SNMP tool.

5. Copy MIB files to a directory where your SNMP tool can find them, such as `/usr/share/snmp/mibs/` or `/usr/local/share/snmp/mibs`, depending on the SNMP tool that you use.

To have Net-SNMP tools read the MIBs to provide automatic name-to-OID mapping, add `-m All` to the command, as in the following example: `snmpwalk -v2c -c public -m All <node IP> isilon`

6. Navigate back to the SNMP Monitoring page and configure General Settings.
 - a. In the Settings area, configure protocol access by selecting the version that you want.

OneFS does not support writable OIDs; therefore, no write-only community string setting is available.
 - b. In the **System location** field, type the system name.

This setting is the value that the node reports when responding to queries. Type a name that helps to identify the location of the node.
 - c. Type the contact email address in the **System contact** field.
7. (Optional) If you selected SNMP v1/v2 as your protocol, locate the SNMP v1/v2c Settings section and type the community name in the **Read-only community** field.

The default community name is `I$ilonpublic`.

Note

OneFS no longer supports SNMP v1. Although an option for v1/v2c may be displayed, if you select the v1/v2c pair, OneFS will only monitor through SNMP v2c.

8. Configure SNMP v3 Settings.
 - a. In the **Read-only user** field, type the SNMP v3 security name to change the name of the user with read-only privileges.

The default read-only user is `general`.
The password must contain at least eight characters and no spaces.
 - b. In the **SNMP v3 password** field, type the new password for the read-only user to set a new SNMP v3 authentication password.

The default password is `password`. We recommend that you change the password to improve security.
 - c. Type the new password in the **Confirm password** field to confirm the new password.
9. Click **Submit**.

View SNMP settings

You can review SNMP monitoring settings.

Procedure

1. Run the following command:

```
isi snmp list
```

Events and notifications

You can monitor the health and performance of your EMC Isilon cluster through OneFS event notifications.

When OneFS identifies an occurrence on your cluster that may require additional attention, an event is generated. OneFS records events related to file system integrity, network connections, hardware, and other vital components of your cluster.

You can select the events that you want to monitor, and you can cancel, quiet, or unquiet events.

In addition, you can configure event notification rules to determine who receives a notification when an event occurs.

Coalesced events

OneFS coalesces related, group events or repeated, duplicate events into a single event.

Coalesced group events

Group events are different types of events that are all related to a single occurrence.

In the following example, a single connection issue might generate the following events:

Event	Description
100010005	A SAS PHY topology problem or change was detected.
100010006	A drive's error log counter indicates there may be a problem.
100010007	A SAS link has exceeded the maximum Bit Error Rate (BER) .
100010008	A SAS link has been disabled for exceeding the maximum Bit Error Rate (BER).

Because the events are triggered by a single occurrence, OneFS creates a group event and combines the related messages under the new group event numbered 24.294. Instead of seeing four events, you will see a single group event alerting you to storage transport issues. You can still view all the grouped events individually if you choose.

To view this coalesced event, run the following command:

```
isi events show 24.924
```

The system displays the following example output of the coalesced group event:

```

ID: 24.924
Type: 199990001
Severity: critical
Value: 0.0
Message: Disk Errors detected (Bay 1)
Node: 21
Lifetime: Sun Jun 17 23:29:29 2012 - Now
Quieted: Not quieted
Specifiers: disk: 35
            val: 0.0
    
```



```

devid: 24
drive_serial: 'XXXXXXXXXXXXXX'
lba: 1953520064L
lnn: 21
drive_type: 'HDD'
device: 'dal'
bay: 1
unit: 805306368
Coalesced by: --
Coalescer Type: Group
Coalesced events:
ID   STARTED   ENDED SEV LNN MESSAGE
24.911 06/17 23:29 -- I 21 Disk stall: Bay 1, Type HDD, LNUM 35.
Disk ...
24.912 06/17 23:29 -- I 21 Sector error: dal block 1953520064
24.913 06/17 23:29 -- I 21 Sector error: dal block 2202232
24.914 06/17 23:29 -- I 21 Sector error: dal block 2202120
24.915 06/17 23:29 -- I 21 Sector error: dal block 2202104
24.916 06/17 23:29 -- I 21 Sector error: dal block 2202616
24.917 06/17 23:29 -- I 21 Sector error: dal block 2202168
24.918 06/17 23:29 -- I 21 Sector error: dal block 2202106
24.919 06/17 23:29 -- I 21 Sector error: dal block 2202105
24.920 06/17 23:29 -- I 21 Sector error: dal block 1048670
24.921 06/17 23:29 -- I 21 Sector error: dal block 223
24.922 06/17 23:29 -- C 21 Disk Repair Initiated: Bay 1, Type
HDD, LNUM...

```

Coalesced duplicate events

Duplicate events are the same message repeated in response to an ongoing issue.

In the following example, a SmartQuotas maximum threshold is repeatedly exceeded, and the system coalesces this sequence of identical but discrete occurrences into one event numbered 1.3035.

To view this coalesced event, run the following command:

```
isi events show 1.3035
```

The system displays the following example output of the coalesced duplicate event:

```

ID: 1.3035
Type: 500010001
Severity: info
Value: 0.0
Message: SmartQuotas threshold violation on quota violated,
domain direc...
Node: All
Lifetime: Thu Jun 14 01:00:00 2012 - Now
Quieted: Not quieted
Specifiers: enforcement: 'advisory'
            domain: 'directory /ifs/quotas'
            name: 'violated'
            val: 0.0
            devid: 0
            lnn: 0
Coalesced by: --
Coalescer Type: Duplicate
Coalesced events:
ID   STARTED   ENDED SEV LNN MESSAGE

```

```

18.621 06/14 01:00 -- I All SmartQuotas threshold violation on quota
vio...
18.630 06/15 01:00 -- I All SmartQuotas threshold violation on quota
vio...
18.638 06/16 01:00 -- I All SmartQuotas threshold violation on quota
vio...
18.647 06/17 01:00 -- I All SmartQuotas threshold violation on quota
vio...
18.655 06/18 01:00 -- I All SmartQuotas threshold violation on quota
vio...

```

Viewing event information

You can view event details, event history, and event logs.

View a list of events

You can view a list of all events.

You can also filter event results by the following criteria:

- Oldest events
- Newest events
- Historical events
- Coalesced events
- Severity level
- Events on a specific node
- Event types

Procedure

1. Run the `isi events list` command.

The system displays output similar to the following example:

```

ID      STARTED   ENDED      SEV LNN MESSAGE
2.2     04/24 03:53 --          C   2   One or more drives (bay(s)
3, 4, 5, 6, ...
1.166  04/24 03:54 --          I   All Software license 'HDFS' will
expire on ...
1.167  04/24 03:54 --          I   All Software license '{license}'
will expir...
1.168  04/24 03:54 --          I   All Software license '{license}'
will expir...
2.57   04/25 17:41 --          I   All Recurring: Software license
'HDFS' will...
2.58   04/25 17:41 --          I   All Recurring: Software license
'{license}'...
1.2     05/02 16:40 --          C   1   One or more drives (bay(s)
3, 4, 5, 6, ...
3.1     05/02 16:50 --          C   3   One or more drives (bay(s)
3, 4, 5, 6, ...
1.227  04/29 20:25 04/29 20:25 C   All Test event sent from CLI
1.228  04/29 20:26 04/29 20:26 C   All Test event sent from CLI
1.229  04/29 22:33 04/29 22:33 C   All Test event sent from CLI
1.259  05/01 00:00 05/01 00:00 I   1   Monthly status
2.59   04/25 17:41 05/02 20:17 I   All Recurring: Software license
'{license}'...

```

The following example command displays a list of events with a severity value of critical:

```
isi events list --severity=critical
```

The system displays output similar to the following example:

```
ID      STARTED      ENDED      SEV LNN MESSAGE
2.2     04/24 03:53 --          C   2   One or more drives (bay(s)
3, 4, 5, 6, ...
1.2     05/02 16:40 --          C   1   One or more drives (bay(s)
3, 4, 5, 6, ...
3.1     05/02 16:50 --          C   3   One or more drives (bay(s)
3, 4, 5, 6, ...
1.227   04/29 20:25 04/29 20:25 C   All Test event sent from CLI
1.228   04/29 20:26 04/29 20:26 C   All Test event sent from CLI
1.229   04/29 22:33 04/29 22:33 C   All Test event sent from CLI
```

View event details

You can view the details of a specific event.

Procedure

1. (Optional) To identify the instance ID of the notification that you want to view, run the following command:

```
isi events list
```

2. To view the details of a specific event, run the `isi events show` command and specify the event instance ID

The following example command displays the details for the event with the instance ID of:

```
isi events show --instanceid=2.57
```

The system displays output similar to the following example:

```
      ID: 2.57
      Type: 400070004
      Severity: info
      Value: 28.0
      Message: Recurring: Software license 'HDFS' will expire on
May 23, 2014
      Node: All
      Lifetime: Fri Apr 25 17:41:34 2014 - Now
      Quieted: Not quieted
      Specifiers: license: 'HDFS'
                  val: 24
                  devid: 0
                  lnn: 0
                  thresh: 0.0
                  exp_date: 'May 23, 2014'
Coalesced by: --
Coalescer Type: Repeat
Coalesced events:
ID      STARTED      ENDED      SEV LNN MESSAGE
1.171   04/25 17:41 --          I   All Software license 'HDFS' will
expire on May 23...
2.43    04/29 00:14 --          I   All Software license 'HDFS' will
expire on May 23...
1.200   04/29 18:34 --          I   All Software license 'HDFS' will
```

```

expire on May 23...
1.232 04/30 17:00 -- I All Software license 'HDFS' will
expire on May 23...
2.72 05/01 03:07 -- I All Software license 'HDFS' will
expire on May 23...
2.75 05/02 16:51 -- I All Software license 'HDFS' will
expire on May 23...
    
```

View the event log

You can log in to a node through the command-line interface and view the contents of the local event log.

Event logs are typically used for support purposes. You can only view the event log using the command-line interface.

Procedure

1. Establish an SSH connection to any node in the EMC Isilon cluster.
2. View the `/var/log/isi_celog_events.log` file.

The log file lists all event activity. Each event row contains one of the following event labels:

Event label	Description
COALESCED: FIRST EVENT	An event was tagged as a possible first event in a series of events that can be coalesced. The first event label is a only a placeholder for a potential parent coalescer event.
COALESCER EVENT: ADDED	A parent coalescer event was created.
COALESCED	An event was added as a child beneath a coalescer event.
CREATOR EV COALID UPDATED	A group was created and the placeholder first event label was updated to include actual group information.
DROPPED	An event did not include any new information and was not stored in the master event database.
FORWARDED_TO_MASTER	An event was forwarded to the master node to be stored in the master event database.
DB: STORED	An event was stored in the master event database.
DB: PURGED	An event was removed from the master event database. The database has a limit of 50,000 entries, and old events are purged when that limit is reached.
INVALID EVENT: DROPPED	An event contained invalid information and was not stored in the master event database.
UPDATE EVENT: DROPPED	A request to update the group information in a parent coalescer event was discontinued.

Responding to events

You can view event details and respond to cluster events.

You can view and manage new events, open events, and recently ended events. You can also view coalesced events and additional, more-detailed information about a specific event. You also can quiet or cancel events.

Quieting, unquieting, and canceling events

You can change an event's state by quieting, unquieting, or canceling an event.

You can select the following actions to change the state of an event:

Quiet

Acknowledges and removes the event from the list of new events and adds the event to a list of quieted events.

Note

If a new event of the same event type is triggered, it is a separate new event and must be quieted.

Unquiet

Returns a quieted event to an unacknowledged state in the list of new events and removes the event from the list of quieted events.

Cancel

Permanently ends an occurrence of an event. The system cancels an event when conditions are met that end its duration, which is bounded by a start time and an end time, or when you cancel the event manually.

Most events are canceled automatically by the system when the event reaches the end of its duration. The event remains in the system until you manually acknowledge or quiet the event. You can acknowledge events through either the web administration interface or the command-line interface.

Quiet an event

You can acknowledge and remove an event from the event list by quieting it.

Procedure

1. (Optional) To identify the instance ID of the event that you want to quiet, run the following command:

```
isi events list
```

2. Quiet the event by running the `isi events quiet` command.

The following example command quiets an event with the instance ID of 1.227:

```
isi events quiet --instanceid=1.227
```

Unquiet an event

You can return a quieted event to an unacknowledged state by unquieting the event.

Procedure

1. (Optional) To identify the instance ID of the event that you want to unquiet, run the following command:

```
isi events list
```

2. Unquiet the event by running the `isi events unquiet` command.

The following example command unquiets an event with the instance ID of 1.227:

```
isi events unquiet --instanceid=1.227
```

Cancel an event

You can end the occurrence of an event by canceling it.

Procedure

1. (Optional) To identify the instance ID of the event that you want to cancel, run the following command:

```
isi events list
```

2. Cancel the event by running the `isi events cancel` command.

The following example command cancels an event with the instance ID of 2.59:

```
isi events cancel --instanceid=2.59
```

Managing event notification settings

You can view and modify event notification settings and configure batch notifications.

Event notification methods

You can define the method by which OneFS delivers notifications.

Email

You can send email messages to distribution lists and apply email templates to notifications. You can also specify SMTP, authorization, and security settings.

SupportIQ

You can deliver notifications to Isilon Technical Support over HTTPS, SMTP, or both.

SNMP trap

You can send SNMP traps to one or more network monitoring stations or trap receivers. Each event can generate one or more SNMP traps. You can download management information base files (MIBs) from the cluster at `/usr/local/share/snmp/mibs/`. The `ISILON-TRAP-MIB.txt` file describes the traps that the cluster can generate, and the `ISILON-MIB.txt` file describes the associated varbinds that accompany the traps.

ESRS

You can receive alerts from the EMC Secure Remote Support (ESRS) Gateway. The ESRS Gateway is a secure, IP-based customer service support system.

The ESRS Gateway is similar to SupportIQ and performs many of the same functions:

- Send alerts regarding the health of your devices.
- Enable support personnel to run the same scripts used by SupportIQ to gather data from your devices.
- Allow support personnel to establish remote access to troubleshoot your cluster.

Event notification settings

You can specify whether you want to receive event notifications as aggregated batches or as individual notifications for each event. Batch notifications are sent every 10 seconds.

The batch options that are described in this table affect both the content and the subject line of notification emails that are sent in response to system events. You can specify event notification batch options when you configure SMTP email settings.

Setting	Option	Description
Notification batch mode	Batch all	Generates a single email for each event notification.
	Batch by severity	Generates an email that contains aggregated notifications for each event of the same severity, regardless of event category.
	Batch by category	Generates an email that contains aggregated notifications for event of the same category, regardless of severity.
	No batching	Generates one email per event.
Custom notification template	No custom notification template is set	Sends the email notification in the default OneFS notification template format.
	Set custom notification template	Sends the email notifications in the format that you defined in your custom template file.

View event notification settings

You can view the current values for event notification settings.

Procedure

1. Run the `isi events settings list` command.

The system displays output similar to the following example:

```
Setting      Value
-----
batch_mode   none
user_template /etc/email/email_template.txt
max_email_size 5242880
```

Modify event notification settings

You can modify settings that affect how event notifications are sent.

You can modify the batch mode, user template, and email size.

Procedure

1. Modify event notifications by running the `isi events settings set` command followed by the name and value of the setting you want to change.

The following example command specifies that event notifications of the same severity must be sent in a batch:

```
isi events settings set --name batch_mode --value severity
```

Managing event notification rules

You can create, modify, or delete event notification rules to determine when and how you receive information about specific system events.

View event notification rules

You can view a list of all notification rules or details for a specific rule.

Procedure

1. To view all notification rules, run the `isi events notifications list` command.

The system displays output similar to the following example:

NAME	TYPE	RECIPIENTS	DESCRIPTION
Isilon support	smtp	1	10 categories, 1 event types
SupportIQ	supportiq	3	10
categories, 1 event types			
CritSupport	smtp	5	10 categories, 2 event types

2. To view the details of a specific notification rule, run the `isi events notifications list` and specify the event name.

The name of the notification rule is case-sensitive.

The following example command displays the details for the SupportIQ notification rule:

```
isi events notifications list --name=SupportIQ
```

The system displays output similar to the following example:

```
Name: SupportIQ
    Type: supportiq
Recipients:
Categories:
LEVELS  EVENT TYPE
    CE   SYS_DISK_EVENTS (100000000)
    CE   NODE_STATUS_EVENTS (200000000)
    CE   REBOOT_EVENTS (300000000)
    CE   SW_EVENTS (400000000)
    CE   QUOTA_EVENTS (500000000)
    CE   SNAP_EVENTS (600000000)
    CE   WINNET_EVENTS (700000000)
    CE   FILESYS_EVENTS (800000000)
    CE   HW_EVENTS (900000000)
    CE   CPOOL_EVENTS (1100000000)
Specific Event types:
LEVELS  EVENT TYPE
    I    SW_CLUSTER_MONTHLY_STATUS (400090001)
```

Create an event notification rule

You can configure event notification rules based on specified events and event types.

You can configure email notification and SNMP trap generation for a specific event.

Procedure

1. Run the `isi events notifications create` command.

The following example command creates a rule called `example rule` that specifies that a notification should be sent to `test@samplecorp.com` when any critical event occurs:

```
isi events notifications create --name=test-rule \
--email=test@samplecorp.com --include-critical=all
```

The name of the notification rule is case-sensitive.

Modify an event notification rule

You can modify event notification rules that you created. System event notification rules cannot be modified.

Procedure

1. (Optional) To identify the name of the notification that you want to cancel, run the following command:

```
isi events notifications list
```

2. Modify an event notification rule by running the `isi events notifications modify` command.

The following example command modifies the notification rule named `test-rule` by adding all emergency level events to list of events that trigger a notification email:

```
isi events notifications modify --name=test-rule --add-
emergency=all
```

The name of the notification rule is case-sensitive.

Send a test event notification

You can generate a test event notification to confirm that event notifications are working as you intend.

Procedure

1. Run the `isi events sendtest` command.

The command returns the instance ID of the test event.

2. (Optional) Run the `isi events list` command to verify that the test event notification was sent and that the returned instance ID is displayed in the list.

Delete an event notification rule

You can delete event notification rules that you created. System event notification rules cannot be deleted.

Procedure

1. (Optional) To identify the name of the notification that you want to cancel, run the following command:

```
isi events notifications list
```

2. Delete an event notification rule by running the `isi events notifications delete` command.

The following example command deletes the notification rule named test-rule:

```
isi events notifications delete --name=test-rule
```

The name of the notification rule is case-sensitive.

Cluster maintenance

Trained service personnel can replace or upgrade components in Isilon nodes.

Isilon Technical Support can assist you with replacing node components or upgrading components to increase performance.

Replacing node components

If a node component fails, Isilon Technical Support will work with you to quickly replace the component and return the node to a healthy status.

Trained service personnel can replace the following field replaceable units (FRUs):

- battery
- boot flash drive
- SATA/SAS Drive
- memory (DIMM)
- fan
- front panel
- intrusion switch
- network interface card (NIC)
- InfiniBand card
- NVRAM card
- SAS controller
- power supply

If you configure your cluster to send alerts to Isilon, Isilon Technical Support will contact you if a component needs to be replaced. If you do not configure your cluster to send alerts to Isilon, you must initiate a service request.

Upgrading node components

You can upgrade node components to gain additional capacity or performance.

Trained service personnel can upgrade the following components in the field:

- drive
- memory (DIMM)
- network interface card (NIC)

If you want to upgrade components in your nodes, contact Isilon Technical Support.

Managing drive firmware

If the firmware of any drive in a cluster becomes obsolete, the cluster performance or hardware reliability might get affected. To ensure overall data integrity, you may update

the drive firmware to the latest revision by installing the drive support package or the drive firmware package.

You can determine whether the drive firmware on your cluster is of the latest revision by viewing the status of the drive firmware.

Note

We recommend that you contact EMC Isilon Technical Support before updating the drive firmware.

Drive firmware update overview

You can update the drive firmware through drive support packages or drive firmware packages.

Download and install either of these packages from <http://support.emc.com> depending on the OneFS version running on your cluster and the type of drives on the nodes.

Drive Support Package

For clusters running OneFS 7.1.1 and later, install a drive support package to update the drive firmware. You do not need to reboot the affected nodes to complete the firmware update. A drive support package provides the following additional capabilities:

- Updates the following drive configuration information:
 - List of supported drives
 - Drive firmware metadata
 - SSD wear monitoring data
 - SAS and SATA settings and attributes
- Automatically updates the drive firmware for new and replacement drives to the latest revision before those drives are formatted and used in a cluster. This is applicable only for clusters running OneFS 7.2 and later.

Note

Firmware of drives in use cannot be updated automatically.

Drive Firmware Package

For clusters running OneFS versions earlier than 7.1.1, or for clusters with non-bootflash nodes, install a cluster-wide drive firmware package to update the drive firmware. You must reboot the affected nodes to complete the firmware update.

Install a drive support package

For clusters running OneFS 7.1.1 and later, install the drive support package to update your drive firmware to the latest supported revision.

Procedure

1. Go to the [EMC Support](#) page that lists all the available versions of the drive support package.
2. Click the latest version of the drive support package and download the file.
3. Open a secure shell (SSH) connection to any node in the cluster and log in.
4. Create or check for the availability of the directory structure `/ifs/data/Isilon_Support/dsp`.
5. Copy the downloaded file to the `dsp` directory through SCP, FTP, SMB, NFS, or any other supported data-access protocols.

6. Unpack the file by running the `tar` command.

For example, unpack drive support package version 1.4 as follows:

```
tar -zxvf Drive_Support_v1.4.tgz
```

7. Install the package by running the `isi_dsp_install` command.

For example, install drive support package version 1.4 as follows:

```
isi_dsp_install Drive_Support_v1.4.tar
```

Note

- You must run the `isi_dsp_install` command to install the drive support package. Do not use the `isi pkg` command.
- The installation process takes care of installing all the necessary files from the drive support package followed by the uninstallation of the package. You do not need to delete the package after its installation or prior to installing a later version.

View drive firmware status

You can view the status of the drive firmware on the cluster to determine whether you need to update the firmware.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Perform one of the following tasks:
 - To view the drive firmware status of all the nodes, run the following command:

```
isi drivefirmware status
```

- To view the drive firmware status of drives on a specific node, run the `isi devices` command with the `-a fwstatus` option. Run the following command to view the drive firmware status of each drive on node 1:

```
isi devices -a fwstatus -d 1
```

The output of the previous command is shown in the following example:

Node 1	Model	FW	Desired FW
Bay 1	HGST HUS724030ALA640	MF80AAC0	
Bay 2	HGST HUS724030ALA640	MF80AAC0	
Bay 3	HGST HUS724030ALA640	MF80AAC0	

Run the following command to view the drive firmware status on node 1 and disk 12:

```
isi devices -a fwstatus -d 1:12
```

If a drive firmware update is not required, the `Desired FW` column is empty.

Update the drive firmware

Determine the type of drive in the node, and then execute the update. The node restarts automatically after the loading process is complete.

Before you begin

Install the drive firmware package.

Note

Do not restart or power off the node before the update is complete. When the update process completes successfully, the node restarts automatically.

Procedure

1. Log in to the node through either a serial console port or an internal SSH connection between nodes.
-

Note

External network interfaces are disabled as part of the reboot process. The command in the next step will fail if you run the command from an external network interface.

2. Determine whether a drive firmware update is required by typing the following command:

```
isi_disk_firmware_reboot
```

3. Determine whether any of the drives on the node that need to be updated are Western Digital drives by typing the following command:

```
isi_radish -q
```

CAUTION

If you are performing a drive firmware update on Western Digital drives, you must update the drives sequentially to avoid significant drive damage.

4. Type one of the following commands depending on whether the node contains any Western Digital drives that require a drive firmware update:
 - To update the drive firmware of a node with any Western Digital drives, type the following command to perform a sequential update:


```
isi_disk_firmware_reboot -sv
```
 - To update the drive firmware of a node without any Western Digital drives, type the following command to perform a parallel update:


```
isi_disk_firmware_reboot -p
```

A drive firmware update takes 20–60 seconds, depending on the drive model. A node containing Western Digital drives takes approximately fifteen minutes to complete the entire process.

After the update process is complete, the node reboots automatically.

If the update is unsuccessful, the LED display on the front panel of the node will indicate an error and the node will not reboot. Wait for a few minutes and then run the `reboot` command to reboot the node manually. If this process is unsuccessful, contact EMC Isilon Technical Support.

Verify a drive firmware update

After you update the drive firmware in a node, confirm that the firmware is updated properly and that the affected drives are operating correctly.

Procedure

1. Ensure that no drive firmware updates are currently in progress by running the following command:

```
isi devices
```

If a drive is currently being updated, [FW_UPDATE] appears in the status column.

2. Verify that all drives have been updated by running the following command:

```
isi drivefirmware status
```

If all drives have been updated, the `Desired FW` column is empty.

3. Verify that all affected drives are operating in a healthy state by running the following command:

```
isi devices
```

If a drive is operating in a healthy state, [HEALTHY] appears in the status column.

Drive firmware status information

You can view information about the status of the drive firmware through the OneFS command-line interface.

The following example shows the output of the `isi drivefirmware status` command:

Model	FW	Desired FW	Count	Nodes
HGST HUS724030ALA640	MF80AAC0	30	1	

Where:

Model

Displays the name of the drive model.

FW

Displays the version number of the firmware currently running on the drives.

Desired FW

If the drive firmware should be upgraded, displays the version number of the drive firmware that the firmware should be updated to.

Count

Displays the number of drives of this model that are currently running the specified drive firmware.

Nodes

Displays the LNNs of nodes that the specified drives exist in.

The following example shows the output of the `isi devices` command with the `-a fwstatus` option:

Node 1	Model	FW	Desired FW

Bay 1	HGST	HUS724030ALA640	MF80AAC0
Bay 2	HGST	HUS724030ALA640	MF80AAC0
Bay 3	HGST	HUS724030ALA640	MF80AAC0

Where:

Drive

Displays the number of the bay that the drive is in.

Note

This column is not labeled in the output. The information appears under the node number.

Model

Displays the name of the drive model.

FW

Displays the version number of the firmware currently running on the drive.

Desired FW

Displays the version number of the drive firmware that the drive should be updated to. If a drive firmware update is not required, the `Desired FW` column is empty.

Automatic update of drive firmware

For clusters running OneFS 7.2 or later, install the latest drive support package on a node to automatically update the firmware for a new or replacement drive.

The information within the drive support package determines whether the firmware of a drive must be updated before the drive is formatted and used. If an update is available, the drive is automatically updated with the latest firmware.

Note

New and replacement drives added to a cluster are formatted regardless of the status of their firmware revision. You can identify a firmware update failure by viewing the firmware status for the drives on a specific node. In case of a failure, run the `isi devices` command with the `fwupdate` action on the node to update the firmware manually. For example, run the following command to manually update the firmware on node 1:

```
isi devices -a fwupdate -d 1
```

Managing cluster nodes

You can add and remove nodes from a cluster. You can also shut down or restart the entire cluster.

Add a node to a cluster

You can add a new node to an existing EMC Isilon cluster.

Before you begin

Before you add a node to a cluster, verify that an internal IP address is available. Add IP addresses as necessary before you add a new node.

If a new node is running a different version of OneFS than a cluster, the system changes the node version of OneFS to match the cluster.

Note

For specific information about version compatibility between OneFS and EMC Isilon hardware, refer to the *Isilon Supportability and Compatibility Guide*.

Procedure

1. Identify the serial number of the node to be added by running the following command:

```
isi devices --action discover
```

2. Join the node to the cluster by running the following command, where *<serial_num>* is the serial number of the node:

```
isi devices --action add --device <serial_num>
```

Remove a node from the cluster

You can remove a node from a cluster. When you remove a node, the system smartfails the node to ensure that data on the node is transferred to other nodes in the cluster.

Removing a storage node from a cluster deletes the data from that node. Before the system deletes the data, the FlexProtect job safely redistributes data across the nodes remaining in the cluster.

Procedure

1. Run the `isi devices` command.

The following command removes a node with a logical node number (LNN) of 2 from the cluster:

```
isi devices --action smartfail --device 2
```

Modify the LNN of a node

You can modify the logical node number (LNN) of a node. This procedure is available only through the command-line interface (CLI).

The nodes within your cluster can be renamed to any name/integer between 1 and 144. By changing the name of your node, you are resetting the LNN.

Note

Although you can specify any integer as an LNN, we recommend that you do not specify an integer greater than 144. Specifying LNNs above 144 can result in significant performance degradation.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Open the `isi config` command prompt by running the following command:

```
isi config
```

3. Run the `lnnset` command.

The following command switches the LNN of a node from 12 to 73:

```
lnnset 12 73
```


4. Enter `commit` .

Results

You might need to reconnect to your SSH session before the new node name is automatically changed.

Restart or shut down the cluster

You can restart or shut down the EMC Isilon cluster.

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem

2. Restart or shutdown nodes on the cluster.

- To restart a single node or all nodes on the cluster, run the `reboot` command.

The following command restarts a single node by specifying the logical node number (lnn):

```
reboot 7
```

- To shut down a single node or all nodes on the cluster, run the `shutdown` command.

The following command shuts down all nodes on the cluster:

```
shutdown all
```

Upgrading OneFS

Two options are available for upgrading the OneFS operating system: a rolling upgrade or a simultaneous upgrade. Before upgrading OneFS software, a pre-upgrade check must be performed.

A rolling upgrade individually upgrades and restarts each node in the EMC Isilon cluster sequentially. During a rolling upgrade, the cluster remains online and continues serving clients with no interruption in service, although some connection resets may occur on SMB clients. Rolling upgrades are performed sequentially by node number, so a rolling upgrade takes longer to complete than a simultaneous upgrade. The final node in the upgrade process is the node that you used to start the upgrade process.

Note

Rolling upgrades are not available for all clusters. For instructions on how to plan an upgrade, prepare the cluster for upgrade, and perform an upgrade of the operating system, see the *OneFS Upgrade Planning and Process Guide*.

A simultaneous upgrade installs the new operating system and restarts all nodes in the cluster at the same time. Simultaneous upgrades are faster than rolling upgrades but require a temporary interruption of service during the upgrade process. Your data is inaccessible during the time that it takes to complete the upgrade process.

Before beginning either a simultaneous or rolling upgrade, OneFS compares the current cluster and operating system with the new version to ensure that the cluster meets certain criteria, such as configuration compatibility (SMB, LDAP, SmartPools), disk availability, and the absence of critical cluster events. If upgrading puts the cluster at risk, OneFS warns you, provides information about the risks, and prompts you to confirm whether to continue the upgrade.

If the cluster does not meet the pre-upgrade criteria, the upgrade does not proceed, and the unsupported statuses are listed.

Remote support

Isilon Technical Support personnel can remotely manage your Isilon cluster to troubleshoot an open support case with your permission.

You can enable remote customer service support through SupportIQ or the EMC Secure Remote Support (ESRS) Gateway.

Remote support using SupportIQ

Isilon Technical Support personnel can remotely manage your Isilon cluster to troubleshoot an open support case with your permission. The Isilon SupportIQ module allows Isilon Technical Support personnel to gather diagnostic data about the cluster.

Isilon Technical Support representatives run scripts that gather data about cluster settings and operations. The SupportIQ agent then uploads the information to a secure Isilon FTP site so it is available for Isilon Technical Support personnel to review. These scripts do not affect cluster services or data availability.

Note

The SupportIQ scripts are based on the Isilon `isi_gather_info` log-gathering tool.

The SupportIQ module is included with the OneFS operating system and does not require you to activate a separate license. You must enable and configure the SupportIQ module before SupportIQ can run scripts to gather data. The feature may have been enabled when the cluster was first set up, but you can enable or disable SupportIQ through the Isilon web administration interface.

In addition to enabling the SupportIQ module to allow the SupportIQ agent to run scripts, you can enable remote access, which allows Isilon Technical Support personnel to monitor cluster events and remotely manage your cluster using SSH or the web administration interface. Remote access helps Isilon Technical Support to quickly identify and troubleshoot cluster issues. Other diagnostic tools are available for you to use in conjunction with Isilon Technical Support to gather and upload information such as packet capture metrics.

Note

If you enable remote access, you must also share cluster login credentials with Isilon Technical Support personnel. Isilon Technical Support personnel remotely access your cluster only in the context of an open support case and only after receiving your permission.

Configuring SupportIQ

OneFS logs contain data that Isilon Technical Support personnel can securely upload, with your permission, and then analyze to troubleshoot cluster problems. The SupportIQ technology must be enabled and configured for this process.

When SupportIQ is enabled, Isilon Technical Support personnel can request logs through scripts that gather cluster data and then upload the data to a secure location. You must enable and configure the SupportIQ module before SupportIQ can run scripts to gather data. The feature may have been enabled when the cluster was first set up.

You can also enable remote access, which allows Isilon Technical Support personnel to troubleshoot your cluster remotely and run additional data-gathering scripts. Remote access is disabled by default. To enable remote SSH access to your cluster, you must provide the cluster password to a Technical Support engineer.

Enable and configure SupportIQ

You can enable and configure SupportIQ to allow the SupportIQ agent to run scripts that gather and upload information about your cluster to Isilon Technical Support personnel.

Procedure

1. Run the following command:

```
isi_promptsupport -e
```

2. Follow the command prompts to enter the following information:

- Company name
- Contact name
- Contact phone
- Contact email

3. Log into the web administration interface and click **Cluster Management** > **General Settings** > **SupportIQ**.

SupportIQ alerts, HTTPS proxy, and remote access must be configured through the web administration interface.

4. In the **SupportIQ Settings** area, select the **Enable SupportIQ** check box.
5. For **SupportIQ alerts**, select an option.
 - **Send alerts via SupportIQ agent (HTTPS) and by email (SMTP)** – SupportIQ delivers notifications to Isilon through the SupportIQ agent over HTTPS and by email over SMTP.
 - **Send alerts via SupportIQ agent (HTTPS)** – SupportIQ delivers notifications to Isilon only through the SupportIQ agent over HTTPS.
6. (Optional) Enable HTTPS proxy support for SupportIQ.
 - a. Select the **HTTPS proxy for SupportIQ** check box.
 - b. In the **Proxy host** field, type the IPv4 address, IPv6 address, or fully qualified domain name (FQDN) of the HTTP proxy server.
 - c. In the **Proxy port** field, type the port number on which the HTTP proxy server receives requests.
 - d. (Optional) In the **Username** field, type the user name for the proxy server.
 - e. (Optional) In the **Password** field, type the password for the proxy server.
7. (Optional) Enable remote access to the cluster.
 - a. Select the **Enable remote access to cluster via SSH and web interface** check box.
 - b. Review the remote-access end user license agreement (EULA) and, if you agree to the terms and conditions, select the **I have read and agree to...** check box.
8. Click **Submit**.

A successful configuration is indicated by a message similar to `SupportIQ settings have been updated.`

Disable SupportIQ

You can disable SupportIQ so the SupportIQ agent does not run scripts to gather and upload data about your Isilon cluster.

Procedure

1. Run the following command:

```
isi_promptsupport -e
```

2. Follow the command prompts.

SupportIQ scripts

When SupportIQ is enabled, Isilon Technical Support personnel can request logs with scripts that gather cluster data and then upload the data. The SupportIQ scripts are located in the `/usr/local/SupportIQ/Scripts/` directory on each node.

The following table lists the data-gathering activities that SupportIQ scripts perform. These scripts can be run automatically, at the request of an Isilon Technical Support representative, to collect information about your cluster's configuration settings and operations. The SupportIQ agent then uploads the information to a secure Isilon FTP site, so that it is available for Isilon Technical Support personnel to analyze. The SupportIQ scripts do not affect cluster services or the availability of your data.

Action	Description
Clean watch folder	Clears the contents of <code>/var/crash</code> .
Get application data	Collects and uploads information about OneFS application programs.
Generate dashboard file daily	Generates daily dashboard information.
Generate dashboard file sequence	Generates dashboard information in the sequence that it occurred.
Get ABR data (as built record)	Collects as-built information about hardware.
Get ATA control and GMirror status	Collects system output and invokes a script when it receives an event that corresponds to a predetermined <code>eventid</code> .
Get cluster data	Collects and uploads information about overall cluster configuration and operations.
Get cluster events	Gets the output of existing critical events and uploads the information.
Get cluster status	Collects and uploads cluster status details.
Get contact info	Extracts contact information and uploads a text file that contains it.
Get contents (var/crash)	Uploads the contents of <code>/var/crash</code> .
Get job status	Collects and uploads details on a job that is being monitored.
Get domain data	Collects and uploads information about the cluster's Active Directory Services (ADS) domain membership.

Action	Description
Get file system data	Collects and uploads information about the state and health of the OneFS <code>/ifs/</code> file system.
Get IB data	Collects and uploads information about the configuration and operation of the InfiniBand back-end network.
Get logs data	Collects and uploads only the most recent cluster log information.
Get messages	Collects and uploads active <code>/var/log/messages</code> files.
Get network data	Collects and uploads information about cluster-wide and node-specific network configuration settings and operations.
Get NFS clients	Runs a command to check if nodes are being used as NFS clients.
Get node data	Collects and uploads node-specific configuration, status, and operational information.
Get protocol data	Collects and uploads network status information and configuration settings for the NFS, SMB, FTP, and HTTP protocols.
Get Pcap client stats	Collects and uploads client statistics.
Get readonly status	Warns if the chassis is open and uploads a text file of the event information.
Get usage data	Collects and uploads current and historical information about node performance and resource usage.
<code>isi_gather_info</code>	Collects and uploads all recent cluster log information.
<code>isi_gather_info --incremental</code>	Collects and uploads changes to cluster log information that have occurred since the most recent full operation.
<code>isi_gather_info --incremental single node</code>	Collects and uploads details for a single node. Prompts you for the node number.
<code>isi_gather_info single node</code>	Collects and uploads changes to cluster log information that have occurred since the most recent full operation. Prompts you for the node number.
Upload the dashboard file	Uploads dashboard information to the secure Isilon Technical Support FTP site.

Remote support using ESRS Gateway

EMC Isilon clusters support enablement of the ESRS Gateway.

The EMC Secure Remote Support (ESRS) Gateway is a secure, IP-based customer service support system. The EMC ESRS Gateway features include 24x7 remote monitoring and secure authentication with AES 256-bit encryption and RSA digital certificates. You can select monitoring on a node-by node basis, allow or deny remote support sessions, and review remote customer service activities.

The ESRS Gateway is similar to SupportIQ and performs many of the same functions:

- Send alerts regarding the health of your devices.
- Enable support personnel to run the same scripts used by SupportIQ to gather data from your devices.

- Allow support personnel to establish remote access to troubleshoot your cluster.

An important difference between SupportIQ and the ESRS Gateway is that SupportIQ management is cluster-wide; SupportIQ manages all nodes. The ESRS Gateway manages nodes individually; you select which nodes should be managed.

You can only enable one remote support system on your Isilon cluster. The EMC products you use and your type of environment determine which system is most appropriate for your Isilon cluster:

- If your environment comprises one or more EMC products that can be monitored, use the ESRS Gateway.
- If ESRS is currently implemented in your environment, use the ESRS Gateway.
- If your use of ESRS requires the ESRS Client, use SupportIQ. Isilon nodes do not support ESRS Client connectivity.
- If you have a high-security environment, use the ESRS Gateway.
- If the only EMC products in your environment are Isilon nodes, use SupportIQ.

See the most recent version of the document titled *EMC Secure Remote Support Technical Description* for a complete description of EMC Secure Remote Support features and functionality.

Additional documentation on ESRS can be found on the EMC Online Support site.

Configuring ESRS Gateway support

You can configure support for the ESRS Gateway on your Isilon cluster.

Before configuring ESRS Gateway support on your Isilon cluster, at least one ESRS Gateway server must be installed and configured. The server acts as the single point of entry and exit for IP-based remote support activities and monitoring notifications. You can also set up a secondary gateway server as a failover, specify whether to use SMTP if ESRS transmission fails, and specify whether an email should be sent upon transmission failure.

ESRS Gateway support also requires you to designate a subnet as a point for remote access by support personnel. We recommend that you designate a subnet that is dedicated to remote connections through the ESRS Gateway, and that the subnet contains a static IP address pool in the System access zone. If you cannot dedicate a subnet for remote connections, ensure that the first IP address pool in the designated subnet is configured to use static IP addresses and is assigned to the System access zone.

When you enable support for the ESRS Gateway on a cluster, the serial number and IP address of each node is sent to the ESRS Gateway server. Once node information is received, you can:

- Select which nodes you want managed through the ESRS Gateway with the ESRS Configuration Tool.
- Create rules for remote support connection to Isilon nodes with the ESRS Policy Manager.

See the most recent version of the document titled *EMC Secure Remote Site Planning Guide* for a complete description of ESRS Gateway server requirements, installation, and configuration.

See the most recent version of the document titled *EMC Secure Remote Support Gateway for Windows Operations Guide* for a complete description of the ESRS Configuration Tool.

See the most recent version of the document titled *EMC Secure Remote Support Policy Manager Operations Guide* for a complete description of the ESRS Policy Manager.

Additional documentation on ESRS can be found on the EMC Online Support site.

Enable and configure ESRS Gateway support

You can enable support for the ESRS Gateway on an Isilon cluster.

Before you begin

An ESRS Gateway server must be installed and configured before you can enable ESRS Gateway support on an Isilon cluster. SupportIQ must be disabled.

Procedure

1. Run the `isi remotesupport connectemc modify` command to enable and configure ESRS Gateway support.

The following command enables ESRS Gateway support, specifies a primary gateway, a remote support subnet, and that an SMTP failover should be used.

```
isi remotesupport connectemc modify --enabled yes \  
-- primary-esrs-gateway gw-serv-esrs1 --use-smtp-failover yes \  
--remote-support-subnet subnet0
```

Disable ESRS Gateway support

You can disable ESRS Gateway support on the Isilon cluster.

You can disable support for the ESRS Gateway in order to use SupportIQ. Isilon clusters only allow one remote support system to be enabled at a time.

Procedure

1. Disable ESRS Gateway support on an Isilon cluster by running the following command:

```
isi remotesupport connectemc modify --enabled no
```

View ESRS Gateway settings

You can view ESRS Gateway settings specified on an EMC Isilon cluster.

Procedure

1. Run the `isi remotesupport connectemc view` command.

The system displays output similar to the following example:

```
Enabled: yes  
Primary Esrs Gateway: gw-serv-esrs1  
Secondary Esrs Gateway: gw-serv-esrs2  
Use Smtip Failover: yes  
Email Customer On Failure: no  
Remote Support Subnet: subnet0
```

Cluster administration commands

You can configure, monitor, and manage the Isilon cluster using cluster administration commands.

isi config

Opens a new prompt where node and cluster settings can be altered.

The command-line prompt changes to indicate that you are in the `isi config` subsystem. While you are in the `isi config` subsystem, other OneFS commands are unavailable and only `isi config` commands are valid.

Syntax

```
isi config
```

Note

- The following commands are not recognized unless you are currently at the `isi config` command prompt.
- Changes are not applied until you run the `commit` command.
- Some commands require you to restart the cluster.

Commands

`changes`

Displays a list of changes to the configuration that have not been committed.

`commit`

Commits configuration settings and then exits `isi config`.

`date <time-and-date>`

Displays or sets the current date and time on the cluster.

<time-and-date>

Sets cluster time to the time specified.

Specify *<time-and-date>* in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

Specify *<time>* as one of the following values.

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

h

Specifies hours

s

Specifies seconds

`deliprange [<interface-name> [<ip-range>]]`

Displays a list of internal network IP addresses that can be assigned to nodes or removes specified addresses from the list.

<interface-name>

Specifies the name of the interface as one of the following values:

int-a

int-b

failover

<ip-range>

Specifies the range of IP addresses that can no longer be assigned to nodes. Specify in the form *<lowest-ip><highest-ip>*.

encoding [list] [*<encoding>*]

Sets the default encoding character set for the cluster.



Character encoding is typically established during installation of the cluster. Incorrectly modifying character encoding settings may render files unreadable. Modify settings only if necessary and after consultation with Isilon Technical Support.

list

Displays the list of supported character sets.

exit

Exits the `isi config` subsystem.

help

Displays a list of all `isi config` commands. For information about specific commands, the syntax is `help [<command>]`.

interface *<interface-name>*{enable | disable}

The interface command displays the IP ranges, netmask, and MTU and enables or disables internal interfaces. When issued with no argument, this command displays IP ranges, netmask, and MTU of all interfaces. When issued with an *<interface-name>* argument, this command displays IP ranges, netmask, and MTU for only the specified interface.

{enable | disable}

Enables or disables the specified interface.

<interface-name>

Specifies the name of the interface as `int-a` or `int-b`.

iprange [*<interface-name>*] [*<lowest-ip><highest-ip>*]

Displays a list of internal IP addresses that can be assigned to nodes, or adds addresses to the list.

<interface-name>

Specifies the name of the interface as `int-a`, `int-b`, or `failover`.

<lowest-ip><highest-ip>

Specifies the range of IP addresses that can be assigned to nodes.

ipset

Obsolete. Use `l3nset` to renumber cluster nodes. The IP address cannot be set manually.

`joinmode` [*<mode>*]

Displays the setting for how nodes are added to the current cluster. Options *<mode>* specifies the cluster add node setting as one of the following values.

`manual`

Configures the cluster so that joins can be initiated by either the node or the cluster.

`secure`

Configures the cluster so that joins can be initiated by only the cluster.

`l3nset` [*<old-l3n>* *<new-l3n>*]

Displays a table of logical node number (LNN), device ID, and internal IP address for each node in the cluster when run without arguments. Changes the LNN when specified.

<old l3n>

Specifies the old LNN that is to be changed.

<new l3n>

Specifies the new LNN that is replacing the old LNN value for that node.

Note

The new LNN must not be currently assigned to another node. Users logged in to the shell or web administration interface of a node whose LNN is changed must log in again to view the new LNN.

`migrate` [*<interface-name>*] [*<old-ip-range>*] {*<new-ip-range>* [-n *<netmask>*]}

Displays a list of IP address ranges that can be assigned to nodes or both adds and removes IP ranges from that list.

<interface-name>

Specifies the name of the interface as `int-a`, `int-b`, and `failover`.

<old-ip-range>

Specifies the range of IP addresses that can no longer be assigned to nodes. If unspecified, all existing IP ranges are removed before the new IP range is added. Specify in the form of *<lowest-ip>* *<highest-ip>*.

<new-ip-range>

Specifies the range of IP addresses that can be assigned to nodes. Specify in the form of *<lowest-ip>* *<highest-ip>*.

-n *<netmask>*

Specifies a new netmask for the interface.

Note

If more than one node is given a new IP address, the cluster reboots when the change is committed. If only one node is given a new IP address, only that node is rebooted.

`mtu` [*<value>*]

Displays the size of the maximum transmission unit (MTU) that the cluster uses for internal network communications when run with no arguments. Sets a new size of the MTU value, when specified. This command is for the internal network only.

Note

This command is not valid for clusters with an InfiniBand back end.

<value>

Specifies the new size of the MTU value. Any value is valid, but not all values may be compatible with your network. The most common settings are 1500 for standard frames and 9000 for jumbo frames.

name [*<new_name>*]

Displays the names currently assigned to clusters when run with no arguments. Assigns new names to clusters, as specified.

<new_name>

Specifies a new name for the cluster.

netmask [*<interface-name>* [*<ip-mask>*]]

Displays the subnet IP mask that the cluster is currently using or sets new subnet IP masks, as specified. Specifies the interface name as *int-a* and *int-b*.

<interface-name>

Specifies the name of the interface. Valid names are *int-a* and *int-b*.

<ip-mask>

Specifies the new IP mask for the interface.

quit

Exits the `isi config` subsystem.

reboot [{*<node_lnn>* | *all*}]

Reboots one or more nodes, specified by LNN. If no nodes are specified, reboots the node from which the command is run. To reboot the cluster, specify *all*.

Note

If run on an unconfigured node, this command does not accept any arguments.

remove

Deprecated. Instead, run the `isi devices -a smartfail` command.

shutdown [{*<node_lnn>* | *all*}]

Shuts down one or more nodes, specified by LNN. If no nodes are specified, shuts down the node from which the command is run. To shut down the cluster, specify *all*.

Note

If run on an unconfigured node, this command does not accept any arguments.

status [*advanced*]

Displays current information on the status of the cluster. To display additional information, including device health, specify *advanced*.

timezone [*<timezone identifier>*]

Displays the current time zone or specifies new time zones. Specifies the new timezone for the cluster as one of the following values:

<timezone identifier>

Specifies the new time zone for the cluster as one of the following values:

Greenwich Mean Time

Eastern Time Zone

Central Time Zone

Mountain Time Zone

Pacific Time Zone

Arizona

Alaska

Hawaii

Japan

Advanced. Opens a prompt with more time zone options.

version

Displays information about the current OneFS version.

wizard

Activates a wizard on unconfigured nodes and reactivates the wizard if you exit it during the initial node configuration process. The wizard prompts you through the node-configuration steps.

isi email

Configures email settings for the cluster.

Syntax

```
isi email
  [--mail-relay] <string>
  [--smtp-port <integer>]
  [--mail-sender <string>]
  [--mail-subject <string>]
  [--use-smtp-auth {yes | no}]
  [--auth-user <string>]
  [--use-encryption {yes | no}]
```

Options

--mail-relay *<string>*

Specifies the SMTP relay address.

--smtp-port *<integer>*

Specifies the SMTP relay port. The default value is 25.

--mail-sender *<string>*

Specifies the originator email address.

--mail-subject *<string>*

Specifies the prefix string for the email subject.

--use-smtp-auth {yes | no}

Specifies whether using SMTP authentication. **Yes** enables SMTP authentication.

```
{--auth-user | -u} <string>
```

Specifies the username for SMTP authentication.

```
{--auth-pass | -p} <string>
```

Specifies the password for SMTP authentication.

```
--use-encryption {yes | no}
```

Specifies whether using encryption (TLS) for SMTP authentication. *Yes* enables encryption.

isi email list

Displays email settings for the cluster

Syntax

```
isi email list
```

isi exttools

Provides subcommands for interacting with supported third-party tools.

Nagios is the only third-party tool that is supported in this release. Multiple Isilon clusters can be monitored with the configuration file that is generated when this command is used.

Syntax

```
isi exttools nagios_config
```

isi license activate

Activates a OneFS module license.

Syntax

```
isi license activate --key <key>
```

Options

```
{--key | -k} <key>
```

Activates the specified license key. Multiple keys can be specified by separating them with either spaces or commas.

isi license status

Displays the license status of all OneFS modules.

Syntax

```
isi license status
```

isi license unconfigure

Unconfigures a OneFS module license.

Unconfiguring a license disables any recurring jobs or scheduled operations that you have activated for that module, but it does not deactivate the license.

Syntax

```
isi license unconfigure --module <module>
```

Options

`--module <module>`

Specifies the name of the license that you want to unconfigure.

isi perfstat

Displays a realtime snapshot of network and file I/O performance.

Syntax

```
isi perfstat
```

isi pkg create

Creates OneFS patches. The `isi pkg create` command is intended solely for use by EMC Isilon Engineering personnel to create new patches for the OneFS operating system. As such, it does not function as intended in a customer environment. This description is information-only.

Syntax

```
isi pkg create <patch-spec-file>
```

Options

`<patch-spec-file>`

Provide the description and parameters for the patch that you are creating.

isi pkg delete

Uninstalls a patch.

Syntax

```
isi pkg delete <patch-name>
```

Options

`<patch-name>`

Required. Uninstalls a patch from the cluster. The patch name must be the name of an installed patch.

Examples

To uninstall a package named `patch-example.tar` from the cluster, run the following command.

```
isi pkg delete patch-example.tar
```

Use the `isi pkg info` command to verify that the patch was successfully uninstalled.

isi pkg info

Displays information about patches that are installed on the cluster.

Syntax

```
isi pkg info <patch-name>
```

Options

<patch-name>

Displays information about only the specified patch. *<patch-name>* can be the path to a tar archive or the URL of a patch on an HTTP or FTP site. If you omit this option, the system displays all installed patches.

Examples

When you examine a specific patch, more information is returned than when you run `isi pkg info` without arguments to get information on all patches.

To get information for a patch named `patch-example.tar`, run the following command.

```
isi pkg info <path> patch-example.tar
```

The system displays the package name and date of installation, similar to the following output.

```
Information for patch-example:
Description:
Package Name : patch-example - 2009-10-11
```

If the patch is not installed, the system displays the following output.

```
patch-example.tar It is not installed.
```

isi pkg install

Installs a patch from a tar archive or an HTTP or FTP site

Syntax

```
isi pkg install <patch-name>
```

Options

<patch-name>

Required. Installs the specified patch on the cluster. *<patch-name>* can be either a path to a tar archive or a URL for a patch on an HTTP or FTP site.

Examples

To install a patch named `patch-example.tar` on the cluster, run the following command.

```
isi pkg install patch-example.tar
```

The system displays output similar to the following example.

```
Preparing to install the package... Installing
the package... Committing the installation...
Package successfully installed
```

If necessary, use the `isi pkg info` command to verify that the patch was successfully installed.

isi remotesupport connectemc modify

Enables or disables support for EMC Secure Remote Support (ESRS) Gateway on an Isilon node.

Syntax

```
isi remotesupport connectemc modify
[--enabled {yes|no}]
[--primary-esrs-gateway <string>]
[--secondary-esrs-gateway <string>]
[--use-smtp-failover {yes|no}]
[--email-customer-on-failure {yes|no}]
[--remote-support-subnet <string>]
```

Options

`--enabled {yes|no}`

Specifies whether support for ESRS Gateway is enabled on the Isilon cluster.

`--primary-esrs-gateway <string>`

Specifies the primary ESRS Gateway server. The gateway server acts as the single point of entry and exit for IP-based remote support activities and monitoring notifications. You can specify the gateway as an IP address or as the gateway name.

`--secondary-esrs-gateway <string>`

Specifies an optional secondary ESRS Gateway server that acts as a failover server. You can specify the gateway as an IP address or as the gateway name.

`--use-smtp-failover {yes|no}`

Specifies whether to send event notifications to a failover SMTP address upon ESRS transmission failure. The SMTP email address is specified using the `isi email` command.

`--email-customer-on-failure {yes|no}`

Specifies whether to send an alert to a customer email address upon failure of other notification methods.

`--remote-support-subnet <string>`

Specifies the subnet on the Isilon cluster to be used for remote support connections.

Note

We recommend that you designate a subnet that is dedicated to remote connections through the ESRS Gateway, and that the subnet contains a static IP address pool in the System access zone. If you cannot dedicate a subnet for remote connections, ensure that the first IP address pool in the designated subnet is configured to use static IP addresses and is assigned to the System access zone.

Examples

The following command enables ESRS Gateway support, specifies an IP address as the primary gateway, specifies subnet3 as the subnet that handles remote support, and directs OneFS to email the customer contact if all transmission methods fail:

```
isi remotesupport connectemc modify --enabled=yes \
-- primary-esrs-gateway=198.51.100.24 \
--email-customer-on-failure=yes --remote-support-subnet=subnet3
```

isi remotesupport connectemc view

Displays EMC Secure Remote Support (ESRS) Gateway settings on an Isilon node.

Syntax

```
isi remotesupport connectemc view
```

Options

This command has no options.

isi services

Displays a list of available services. The `-l` and `-a` options can be used separately or together.

Syntax

```
isi services
[-l | -a]
[<service> [{enable | disable}]]
```

Options

`-l`

Lists all available services and the current status of each. This is the default value for this command.

`- a`

Lists all services, including hidden services, and the current status of each.

`<service> {enable | disable}`

Enables or disables the specified service.

Examples

The following example shows the command to enable a specified hidden service.

```
isi services -a <hidden-service> enable
```

isi set

Works similar to `chmod`, providing a mechanism to adjust OneFS-specific file attributes, such as the requested protection, or to explicitly restripe files. Files can be specified by path or LIN.

Syntax

```
isi set
  [-f -F -L -n -v -r
  -R]
  [-p <policy>]
  [-w <width>]
  [-c {on | off}]
  [-g <restripe_goal>]
  [-e <encoding>]
  [-d <@r drives>]
  [-a {<default> | <streaming> | <random> | <custom{1..5}>}]
  [-l {<concurrency> | <streaming> | <random>}]
  [--diskpool {<id> | <name>}]
  [-A {on | off}]
  [-P {on | off}]
  [{--strategy | -s} {<avoid> | <metadata> | <metadata-write> |
  <data>}
  [<file> {<path> | <lin>}]
```

Options

-f

Suppresses warnings on failures to change a file.

-F

Includes the `/ifs/.ifsvar` directory content and any of its subdirectories. Without `-F`, the `/ifs/.ifsvar` directory content and any of its subdirectories are skipped. This setting allows the specification of potentially dangerous, unsupported protection policies.

-L

Specifies file arguments by LIN instead of path.

-n

Displays the list of files that would be changed without taking any action.

-v

Displays each file as it is reached.

-r

Runs a restripe.

-R

Sets protection recursively on files.

-p <policy>

Specifies protection policies in the following forms:

+M

Where **M** is the number of node failures that can be tolerated without loss of data. **+M** must be a number from, where numbers 1 through 4 are valid.

+D:M

Where **D** indicates the number of drive failures and **M** indicates number of node failures that can be tolerated without loss of data. **D** must be a number from 1 through 4 and **M** must be any value that divides into **D** evenly. For example, +2:2 and +4:2 are valid, but +1:2 and +3:2 are not.

Nx

Where **N** is the number of independent mirrored copies of the data that will be stored. **N** must be a number, with 1 through 8 being valid choices.

-w <width>

Specifies the number of nodes across which a file is striped. Typically, $w = N + M$, but width can also mean the total of the number of nodes that are used.

You can set a maximum width policy of 32, but the actual protection is still subject to the limitations on **N** and **M**.

-c {on | off}

Specifies whether write-coalescing is turned on.

-g <restripe goal>

Specifies the restripe goal. The following values are valid:

```
repair
reprotect
rebalance
retune
```

-e <encoding>

Specifies the encoding of the filename. The following values are valid:

```
EUC-JP
EUC-JP-MS
EUC-KR
ISO-8859-1
ISO-8859-10
ISO-8859-13
ISO-8859-14
ISO-8859-15
ISO-8859-160
ISO-8859-2
ISO-8859-3
ISO-8859-4
ISO-8859-5
ISO-8859-6
ISO-8859-7
ISO-8859-8
ISO-8859-9
UTF-8
```

UTF-8-MAC

Windows-1252

Windows-949

Windows-SJIS

-d *<@r drives>*

Specifies the minimum number of drives that the file is spread across.

-a *<value>*

Specifies the file access pattern optimization setting. The following values are valid:

default

streaming

random

custom1

custom2

custom3

custom4

custom5

-l *<value>*

Specifies the file layout optimization setting. This is equivalent to setting both the **-a** and **-d** flags.

concurrency

streaming

random

--diskpool *<id| name>*

Sets the preferred diskpool for a file.

-A {on | off}

Specifies whether file access and protections settings should be managed manually.

-P {on | off}

Specifies whether the file inherits values from the applicable file pool policy.

{--strategy | -s} *<value>*

Sets the SSD strategy for a file. The following values are valid:

If the value is `metadata-write`, all copies of the file's metadata are laid out on SSD storage if possible, and user data still avoids SSDs. If the value is `data`, Both the file's meta- data and user data (one copy if using mirrored protection, all blocks if FEC) are laid out on SSD storage if possible.

avoid

Writes all associated file data and metadata to HDDs only. The data and metadata of the file are stored so that SSD storage is avoided, unless doing so would result in an out-of-space condition.

metadata

Writes both file data and metadata to HDDs. One mirror of the metadata for the file is on SSD storage if possible, but the strategy for data is to avoid SSD storage.

metadata-write

Writes file data to HDDs and metadata to SSDs, when available. All copies of metadata for the file are on SSD storage if possible, and the strategy for data is to avoid SSD storage.

data

Uses SSD node pools for both data and metadata. Both the metadata for the file and user data, one copy if using mirrored protection and all blocks if FEC, are on SSD storage if possible.

***<file>*{*<path>*|*<lin>*}**

Specifies a file by path or LIN.

isi snmp

Manages SNMP settings.

When SNMP v3 is used, OneFS requires AuthNoPriv as the default value. AuthPriv is not supported.

Syntax

```
isi snmp
  [--syslocation <string>]
  [--syscontact <string>]
  [--protocols <value>]
  [--rocommunity <string>]
  [--v3-rouser <string>]
  [--v3-password <string>]
```

Options

--syslocation *<string>*

Specifies the SNMP network. read-only field that the SNMP implementation on the cluster can report back to a user when queried. It's purely informational for the user. This just sets the value of the standard system location OID.

--syscontact *<string>*

Sets the SNMP network contact address.

--protocols *<value>*

Specifies SNMP protocols. The following values are valid:

v1/v2c

v3

all

Note

v1 and v2 are controlled together and must be specified together, as shown.

{ **--rocommunity** | **-c** } *<string>*

Specifies the read-only community name.

{ **--v3-rouser** | **-u** } *<string>*

Specifies the SNMP v3 read-only user security name.

{ **--v3-password** | **-p** } *<string>*

Specifies the SNMP v3 auth password.

isi snmp list

Displays SNMP settings.

Syntax

```
isi snmp list
```

isi statistics client

Displays the most active, by throughput, clients accessing the cluster for each supported protocol. You can specify options to track access by user, for example, more than one user on the same client host access the cluster.

Syntax

```
isi statistics client
[--csv]
[--csvstring <string>]
[--noconversion]
[--noheader]
[--top]
[--interval <integer>]
[--repeat <integer>]
[--degraded]
[--timeout <integer>]
[--nodes <value>]
[--protocols <value>]
[--classes <string>]
[--orderby <column>]
[--total]
[--totalby <column>]
[--nooutput <column>]
[--output <column>]
[--long]
[--zero]
[--local_addr <integer>]
[--local_names <string>]
[--remote_addr <integer>]
[--remote_names <string>]
[--user_ids <integer>]
[--user_names <string>]
[--numeric]
[--wide]
```

Options

```
{--csv | -c}
```

Displays data as comma-separated values.

Note

Disables top-style display and dynamic unit conversion.

```
{--csvstring | -C} <string>
```

Displays data as a csv-style separated list, with the specified string as separator.

Note

If specified, `--csvstring` overrides `--csv` and disables top-style display and dynamic unit conversion.

`--noconversion`

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

`--noheader`

Displays data without column headings.

`{--top | -t}`

Displays results in top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -i} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <integer>`

Specifies how many times to run the report before quitting.

Note

To run the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Causes the report to continue if some nodes do not respond.

`{--timeout | -o} <integer>`

Specifies the number of seconds before remote commands time out.

`--nodes`

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list, for example, `--nodes=1-2, 5-9`. The default value is `all`. The following values are valid:

- `all`
- `<int>`
- `<int>-<int>`

`--protocols <value>`

Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list, for example `--protocols=http,papi`. The following values are valid:

- `all`
- `external`
- `ftp`
- `hdfs`
- `http`

- internal
- irp
- iscsi
- jobd
- lsass_in
- lsass_out
- nlm
- nfs3
- nfs4
- papi
- siq
- smb1
- smb2

`--classes <string>`

Specify which operation classes to report statistics on. The default setting is all classes. The following values are valid:

all

All classes

read

File and stream reading

write

File and stream writing

create

File link node stream and directory creation

delete

File link node stream and directory deletion

namespace_read

Attribute stat and ACL reads; lookup directory reading

namespace_write

Renames; attribute setting; permission time and ACL writes

file_state

Open close; locking: acquire release break check; notification

session_state

Negotiation inquiry or manipulation of protocol connection or session state

other

File-system information for other uncategorized operations

`--orderby <column>`

Specifies how rows are ordered. The following values are valid:

- Class
- In
- InAvg
- InMax

- InMin
- LocalAddr
- LocalName
- Node
- NumOps
- Ops
- Out
- OutAvg
- OutMax
- OutMin
- Proto
- RemoteAddr
- RemoteName
- TimeAvg
- TimeMax
- TimeMin
- TimeStamp
- UserID
- UserName

`--total`

Groups and aggregates results as implied by filtering options.

`--totalby <column>`

Aggregates results according to specified fields. The following values are valid:

- Class
- Group
- LocalAddr
- LocalName
- Node
- Proto
- RemoteAddr
- RemoteName
- UserId
- UserName

Note

`--totalby` overrides `--total` option, when specified.

`--nooutput <column>`

Specifies which columns are not displayed. Columns are excluded from the list of currently active columns specified by the `--output` option or `--long` options or from the default column list if it is not overridden by other output options.

`--output <column>`

Specifies which columns to display. The following values are valid:

TimeStamp

Displays the time at which the isi statistics tool last gathered data. Displayed in POSIX time (number of seconds elapsed since January 1, 1970).

NumOps

Displays the number of times an operation has been performed.

Ops

Displays the rate at which an operation has been performed. Displayed in operations per second.

InMax

Displays the maximum input (received) bytes for an operation.

InMin

Displays the minimum input (received) bytes for an operation.

In

Displays the rate of input for an operation since the last time isi statistics collected the data. Displayed in bytes per second.

InAvg

Displays the average input (received) bytes for an operation.

OutMax

Displays the maximum output (sent) bytes for an operation.

OutMin

Displays the minimum output (sent) bytes for an operation.

Out

Displays the rate of output for an operation since the last time isi statistics collected the data. Displayed in bytes per second.

OutAvg

Displays the average output (sent) bytes for an operation.

TimeMax

Displays the maximum elapsed time taken to complete an operation. Displayed in microseconds.

TimeMin

Displays the minimum elapsed time taken to complete an operation. Displayed in microseconds.

TimeAvg

Displays the average elapsed time taken to complete an operation. Displayed in microseconds.

Node

Displays the node on which the operation was performed.

Proto

Displays the protocol of the operation.

Class

Displays the class of the operation.

UserID

Displays the numeric UID of the user issuing the operation request or the unique logical unit number (LUN) identifier in the case of the iSCSI protocol.

UserName

Displays the resolved text name of the UserID, or the target and LUN in the case of the iSCSI protocol. In either case, if resolution cannot be performed, UNKNOWN is displayed.

LocalAddr

Displays the IP address of the host receiving the operation request. Displayed in dotted-quad form.

LocalName

Displays the resolved text name of the LocalAddr, if resolution can be performed.

RemoteAddr

Displays the IP address of the host sending the operation request. Displayed in dotted-quad form.

RemoteName

Displays the resolved text name of the RemoteAddr, if resolution can be performed.

`--long`

Displays all possible columns.

`--zero`

Shows table entries with no values.

`--local_addrs <integer>`

Specifies local IP addresses for which statistics will be reported.

`--local-names <string>`

Specifies local host names for which statistics will be reported.

`--remote_addrs <integer>`

Specifies remote IP addresses for which statistics will be reported.

`--remote_names <string>`

Specifies remote client names for which statistics will be reported.

`--user_ids <integer>`

Specifies user ids for which statistics will be reported. The default setting is all users.

`--user_names <string>`

Specifies user names for which statistics will be reported. The default setting is all users.

`--numeric`

If text identifiers of local hosts, remote clients, or users are in the list of columns to display (the default setting is for them to be displayed), display the unresolved numeric equivalent of these columns.

`--wide`

Displays resolved names with a wide column width.

isi statistics describe

Displays documentation on given statistics.

Syntax

```
isi statistics describe --stats <statistics-key-string>
```

Options

```
{--stats | -s} <statistics-key-string>
```

Displays documentation on specified statistics. For a complete list of statistics, run `isi statistics list stats`.

isi statistics drive

Displays performance information by drive.

Syntax

```
isi statistics drive
  [--csv]
  [--csvstring <string>]
  [--noconversion]
  [--noheader]
  [--top]
  [--interval <integer>]
  [--repeat <integer>]
  [--degraded]
  [--timeout <integer>]
  [--long]
  [--nodes <value>]
  [--timestamp]
  [--type <value>]
  [--orderby <column>]
```

Options

```
{--csv | -c}
```

Displays data as a comma-separated list.

Note

Disables top-style display and dynamic unit conversion.

```
{--csvstring | -C} <string>
```

Display data as a csv-style separated list with the specified string as separator.

Note

Overrides `--csv`, if specified, and disables top-style display and dynamic unit conversion.

```
--noconversion
```

Displays all data in base quantities, without dynamic conversion. If set, this parameter also disables the display of units within the data table.

```
--noheader
```

Displays data without column headings.

`{--top | -t}`

Displays the results in a top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -I} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <integer>`

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Sets the report to continue running if some nodes do not respond.

`{--timeout | -o} <integer>`

Specifies the number of seconds before remote commands timeout.

`--long`

Displays all possible columns.

`--nodes`

Specifies which nodes to report statistics on. The default value is `all`. The following values are valid:

- `all`
- `<int>`
- `<int><int>`
- `*`

`--orderby <column>`

Specifies how the rows are ordered. The following values are valid:

`Busy`

`BytesIn`

`BytesOut`

`Drive`

`Inodes`

`OpsIn`

`OpsOut`

`Queued`

`SizeIn`

`SizeOut`

`Slow`

`TimeAvg`

`TimeInQ`

`Type`

Used

`--timestamp`

Displays the time at which the `isi statistics` tool last gathered data. Time is displayed in Epoch seconds.

`--type`

Specifies the drive types for which statistics will be reported. The default setting is all drives. The following values are valid:

sata

sas

ssd

isi statistics heat

Displays the most active `/ifs` paths for various metrics.

Syntax

```
isi statistics heat
  [--csv]
  [--csvstring <string>]
  [--noconversion]
  [--noheader]
  [--top]
  [--interval <integer>]
  [--repeat <integer>]
  [--degraded]
  [--timeout <integer>]
  [--nodes <value>]
  [--events <string>]
  [--classes <string>]
  [--orderby <column>]
  [--totalby <column>]
  [--maxpath <integer>]
  [--pathdepth <integer>]
  [--limit <integer>]
  [--output <column>]
  [--long]
```

Options

`{--csv | -c}`

Displays data as a comma-separated list.

Note

Disables the top-style display and dynamic unit conversion.

`{--csvstring | -C} <string>`

Displays data as a csv-style separated list with specified string as separator.

Note

Overrides `--csv`, if specified, and disables top-style display and dynamic unit conversion.

`--noconversion`

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

`--noheader`

Displays data without column headings.

`{--top | -t}`

Displays the results in top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -I} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <integer>`

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Sets the report to continue running if some nodes do not respond.

`{--timeout | -o} <integer>`

Specifies the number of seconds before remote commands timeout.

`--nodes`

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list—for example, `--nodes=1-2,5-9`. The default value is `all`.

The following values are valid:

- `all`
- `<int>`
- `<int><int>`

`--events <string>`

Specifies which event types for the specified information are reported. The following values are valid:

blocked

Access to the LIN was blocked waiting for a resource to be released by another operation. Class is `other`.

contended

A LIN is experiencing cross-node contention; it is being accessed simultaneously through multiple nodes. Class is `other`.

deadlocked

The attempt to lock the LIN resulted in deadlock. Class is `other`.

link

The LIN has been linked into the file system; the LIN associated with this event is the parent directory and not the linked LIN. Class is `namespace_write`.

lock

The LIN is locked. Class is `other`.

lookup

A name is looked up in a directory; the LIN for the directory searched is the one associated with the event. Class is `namespace_read`.

read

A read was performed. Class is `read`.

rename

A file or directory was renamed. The LIN associated with this event is the directory where the rename took place for either the source directory or the destination directory, if they differ. Class is `namespace_write`.

getattr

A file or directory attribute has been read. Class is `namespace_read`.

setattr

A file or directory attribute has been added, modified, or deleted. Class is `namespace_write`.

unlink

A file or directory has been unlinked from the file system, the LIN associated with this event is the parent directory of the removed item. Class is `namespace_write`.

write

A write was performed. Class is `write`.

`--classes <string>`

Specifies which classes for the specified information will be reported. The default setting is all classes. The following values are valid:

read

File and stream reading

write

File and stream writing

create

File, link, node, stream, and directory creation

delete

File, link, node, stream, and directory deletion

namespace_read

Attribute, stat, and ACL reads; lookup, directory reading

namespace_write

Renames; attribute setting; permission, time, and ACL writes

file_state

Open, close; locking: acquire, release, break, check; notification

session_state

Negotiation, inquiry, or manipulation of protocol connection or session state

other

File-system information

`--orderby <column>`

Specifies how rows are ordered. The following values are valid:

- Class
- Event
- LIN
- Node
- Ops

- Path
- TimeStamp

`--totalby <column>`
 Aggregates results according to specified fields. The following values are valid:

- Class
- Event
- LIN
- Node
- Ops
- Path
- TimeStamp

`--maxpath <integer>`
 Specifies the maximum path length to look up in the file system.

`-pathdepth <integer>`
 Reduces paths to the specified depth.

`--limit <integer>`
 Displays only the specified number of entries after totaling and ordering.

`--output <column>`
 Specifies the columns to display. The following values are valid:

timestamp
 Displays the time at which the isi statistics tool last gathered data. Displayed in POSIX time (number of seconds elapsed since January 1, 1970).

Ops
 Displays the rate at which an operation has been performed. Displayed in operations per second.

Node
 Displays the node on which the operation was performed.

Event
 Displays the name of the event.

Class
 Displays the class of the operation.

LIN
 Displays the LIN for the file or directory associated with the event.

Path
 Displays the path associated with the event LIN.

`--long`
 Displays all possible columns.

isi statistics history

Displays historical statistics.

Syntax

```
isi_statistics_history
[--csv]
```

```

[--csvstring <string>]
[--noconversion]
[--noheader]
[--top]
[--interval <number>]
[--repeat <number>]
[--degraded]
[--timeout <number>]
[--nodes <value>]
[--stats <string>]
[--onecolumn]
[--formattime]
[--begin <number>]
[--end <number>]
[--resolution <number>]
[--zero]

```

Options

`{--csv | -c}`

Displays data as a comma-separated list.

Note

Disables top-style display and dynamic unit conversion.

`{--csvstring | -C} <string>`

Display data as a csv-style separated list with specified string as separator.

Note

Overrides `--csv`, if specified, and disables top-style display and dynamic unit conversion.

`--noconversion`

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

`--noheader`

Displays data without column headings.

`{--top | -t}`

Displays results in a top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -I} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <number>`

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Sets the report to continue running if some nodes do not respond.

`{--timeout | -o} <number>`

Specifies number of seconds before remote commands time out.

`--nodes`

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list—for example, `--nodes=1-2,5-9`. The default value is `all`. The following values are valid:

- `all`
- `<int>`
- `<int>-<int>`

`--stats <string>`

Specifies which statistics should be reported for requested nodes, where the value for `<string>` is a statistics key. Use the `isi statistics list stats` command for a complete listing of statistics keys.

`{--onecolumn | -1}`

Displays output one series at a time in one column rather than in a grid format.

`{--formattime | -F}`

Formats series times rather than using UNIX Epoch timestamp format.

`--begin <number>`

Specifies begin time in UNIX Epoch timestamp format.

`--end <number>`

Specifies end time in UNIX Epoch timestamp format.

`--resolution <number>`

Specifies the minimum interval between series data points in seconds.

`--zero`

Displays grid rows with no valid series points.

isi statistics list all

Displays a list of valid list-mode arguments.

Syntax

```
isi statistics list all
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list classes

Displays a list of valid arguments for the `--classes` option.

Syntax

```
isi statistics list classes
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list events

Displays a list of valid arguments for the `--events` option.

Syntax

```
isi statistics list events
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--protocol`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list nodes

Displays a list of valid arguments for the `--nodes` option.

Syntax

```
isi statistics list nodes
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list nooutput

Displays a list of valid arguments for the `--nooutput` option.

Syntax

```
isi statistics list nooutput
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list operations

Displays a list of valid arguments for the `--operations` option.

Syntax

```
isi statistics list operations
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list orderby

Displays a list of valid arguments for the `--orderby` option.

Syntax

```
isi statistics list orderby
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list output

Displays a list of valid arguments for the `--output` option.

Syntax

```
isi statistics list output
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list protocols

Displays a list of valid arguments for the `--protocols` option.

Syntax

```
isi statistics list protocols
  [--client]
  [--protocol]
  [--heat]
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list stats

Displays a list of valid arguments for the `--stats` option.

Syntax

```
isi statistics list stats  
  [--client]  
  [--protocol]  
  [--heat]  
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics list totalby

Displays a list of valid arguments for the `--totalby` option.

Syntax

```
isi statistics list totalby  
  [--client]  
  [--protocol]  
  [--heat]  
  [--drive]
```

Options

`--client`

Displays valid option values for client mode.

`--protocol`

Displays valid option values for protocol mode.

`--heat`

Displays valid option values for heat mode.

`--drive`

Displays valid option values for drive mode.

isi statistics protocol

Displays statistics by protocol, such as NFSv3 and HTTP.

Syntax

```
isi statistics protocol
  [--csv]
  [--csvstring <string>]
  [--noconversion]
  [--noheader]
  [--top]
  [--interval <integer>]
  [--repeat <integer>]
  [--degraded]
  [--timeout <integer>]
  [--long]
  [--nodes {all | <int> | <int>--<int>}]
  [--protocols <protocol>...]
  [--classes <class>...]
  [--orderby <column>...]
  [--total]
  [--totalby <column>...]
  [--nooutput <column>...]
  [--output <column>...]
  [--long]
  [--zero]
  [--operations <operation>...]
```

Options

`{--csv | -c}`

Displays data as comma-separated values.

Note

Disables top-style display and dynamic unit conversion.

`{--csvstring | -C} <string>`

Displays data as a csv-style separated list, with the specified string as a separator.

Note

If specified, `--csvstring` overrides `--csv` and disables top-style display and dynamic unit conversion.

`--noconversion`

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units in the data table.

`--noheader`

Displays data without column headings.

`{--top | -t}`

Displays results in top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -I} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <integer>`

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Causes the report to continue running if some nodes do not respond.

`{--timeout | -o} <integer>`

Specifies the number of seconds before remote commands timeout.

`--nodes`

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list—for example, `--nodes=1-2,5-9`. The default value is `all`. The following values are valid:

- `all`
- `<int>`
- `<int><int>`

`--protocols <value>`

Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list—for example, `--protocols=http,papi`. The following values are valid:

- `all`
- `external`
- `ftp`
- `hdfs`
- `http`
- `internal`
- `irp`
- `iscsi`
- `jobd`
- `lsass_in`
- `lsass_out`
- `nlm`
- `nfs3`
- `nfs4`
- `papi`
- `siq`
- `smb1`
- `smb2`

`--classes <class>`

Specifies which operation classes to report statistics on. The default setting is `all`. The following values are valid:

all

All classes

read

File and stream reading

write

File and stream writing

create

File link node stream and directory creation

delete

File link node stream and directory deletion

namespace_read

Attribute stat and ACL reading; lookup directory reading

namespace_write

Renames; attribute setting; permission time and ACL writes

file_state

Open, close; locking: acquire, release, break, check; notification

session_state

Negotiation inquiry or manipulation of protocol connection or session state

other

File-system information. Multiple values can be specified in a comma-separated list.

--orderby <column>

Specifies how rows are ordered. The following values are valid:

- Class
- In
- InAvg
- InMax
- InMin
- InStdDev
- Node
- NumOps
- Ops
- Out
- OutAvg
- OutMax
- OutMin
- OutStdDev
- Proto
- TimeAvg
- TimeMax
- TimeMin

- TimeStamp
- TimeStdDev

--total

Groups and aggregates the results according to the filtering options.

--totalby <column>

Aggregates results according to specified fields. The following values are valid:

- Class
- Node
- Proto
- Op

Note

--totalby overrides --total, when specified.

--nooutput <column>

Specifies which columns are not displayed. The columns are excluded from the list of the active columns that are specified by --output or --long or from the default column list if it is not overridden by other output options. Multiple values can be specified in a comma-separated list. The following values are valid:

- Class
- In
- InAvg
- InMax
- InMin
- InStdDev
- Node
- NumOps
- Ops
- Out
- OutAvg
- OutMax
- OutMin
- OutStdDev
- Proto
- TimeAvg
- TimeMax
- TimeMin
- TimeStamp
- TimeStdDev

--output <column>

Specifies which columns to display. The following values are valid:

timestamp

Displays the time at which the `isi statistics` tool last gathered data. Displayed in POSIX time (number of seconds elapsed since January 1, 1970). Specify *<time-and-date>* in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

Specify *<time>* as one of the following values.

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

h

Specifies hours

s

Specifies seconds

NumOps

Displays the number of times an operation has been performed. Multiple values can be specified in a comma-separated list.

Ops

Displays the rate at which an operation has been performed. Displayed in operations per second.

InMax

Displays the maximum input (received) bytes for an operation.

InMin

Displays the minimum input (received) bytes for an operation.

In

Displays the rate of input for an operation since the last time `isi statistics` collected the data. Displayed in bytes per second.

InAvg

Displays the average input (received) bytes for an operation.

OutMax

Displays the maximum output (sent) bytes for an operation.

OutMin

Displays the minimum output (sent) bytes for an operation.

Out

Displays the rate of output for an operation since the last time `isi statistics` collected the data. Displayed in bytes per second.

OutAvg

Displays the average output (sent) bytes for an operation.

TimeMax

Displays the maximum elapsed time taken to complete an operation. Displayed in microseconds.

TimeMin

Displays the minimum elapsed time taken to complete an operation. Displayed in microseconds.

TimeAvg

Displays the average elapsed time taken to complete an operation. Displayed in microseconds.

Node

Displays the node on which the operation was performed.

Proto

Displays the protocol of the operation.

Class

Displays the class of the operation.

InStdDev

Displays the standard deviation of the input (received) bytes for an operation. Displayed in bytes.

OutStdDev

Displays the standard deviation of the output (sent) bytes for an operation. Displayed in bytes.

Op

Displays the name of the operation

`--long`

Displays all possible columns.

`--zero`

Shows table entries with no values.

`--operations <operation>`

Specifies the operations on which statistics are reported. To view a list of valid values, run the `isi statistics list operations` command. Multiple values can be specified in a comma-separated list.

isi statistics pstat

Displays a selection of cluster-wide and protocol data.

Syntax

```
isi statistics pstat
  [--csv]
  [--csvstring <string>]
  [--noconversion]
  [--noheader]
  [--top]
  [--interval <integer>]
  [--repeat <integer>]
  [--degraded]
  [--timeout <integer>]
  [--protocol <protocol>]
```

Options

`{--csv | -c}`

Displays data as comma-separated values.

Note

Disables top-style display and dynamic unit conversion.

`{--csvstring | -C} <string>`

Displays data as a csv-style separated list, with the specified string as separator.

Note

If specified, `--csvstring` overrides `--csv` and disables top-style display and dynamic unit conversion.

`--noconversion`

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

`--noheader`

Displays data without column headings.

`{--top | -t}`

Displays results in top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -i} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <integer>`

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Sets the report to continue running if some nodes do not respond.

`{--timeout | -o} <integer>`

Specifies number of seconds before remote commands time out.

`--protocol <protocol>`

Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list—for example, `--protocols=http,papi`. The following values are valid:

- ftp
- hdfs
- http
- irp
- iscsi
- jobd
- lsass_in

- lsass_out
- nlm
- nfs3
- nfs4
- papi
- siq
- smb1
- smb2

isi statistics query

Displays highly customizable information on any statistic in the cluster statistics library.

Syntax

```
isi statistics query
  [--csv]
  [--csvstring <string>]
  [--noconversion]
  [--noheader]
  [--top]
  [--interval <integer>]
  [--repeat <integer>]
  [--degraded]
  [--timeout <integer>]
  [--nofooter]
  [--nodes {all | <int>[-<int>]}]
  [--stats <string>]
```

Options

{--csv | -c}

Displays data as comma-separated values.

Note

Disables top-style display and dynamic unit conversion.

{--csvstring | -C} <string>

Displays data as a csv-style separated list, with the specified string as a separator.

Note

If specified, --csvstring overrides --csv and disables top-style display and dynamic unit conversion.

--noconversion

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

--noheader

Displays data without column headings.

{--top | -t}

Displays results in top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

```
{--interval | -I} <integer>
```

Reports data at the interval specified in seconds.

```
{--repeat | -r} <integer>
```

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

```
{--degraded | -d}
```

Causes the report to continue if some nodes do not respond.

```
{--timeout | -o} <integer>
```

Specifies the number of seconds before remote commands time out.

```
--nofooter
```

Suppresses display of the footer row that contains aggregation data.

```
--nodes {all | <int>[-<int>]}
```

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list—for example, `--nodes=1-2,5-9`. The default value is `all`.

```
{--stats | statistics key} <key>
```

Specifies which statistics should be reported for requested nodes. Run the `isi statistics list stats` command for a complete list of statistics keys.

isi statistics system

Displays general cluster statistics, including op rates for all supported protocols and network and disk traffic.

Syntax

```
isi statistics system
  [--csv]
  [--csvstring <string>]
  [--noconversion]
  [--noheader]
  [--top]
  [--interval <integer>]
  [--repeat <integer>]
  [--degraded]
  [--timeout <integer>]
  [--running <integer>]
  [--nodes]
  [--timestamp]
  [--oprates]
```

Options

```
{--csv | -c}
```

Displays data as a comma-separated list.

Note

Disables top-style display and dynamic unit conversion.

`{--csvstring | -C} <string>`

Display data as a csv-style separated list with specified string as separator.

Note

Overrides `--csv`, if specified, and disables top-style display and dynamic unit conversion.

`--noconversion`

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

`--noheader`

Displays data without column headings.

`{--top | -t}`

Displays results in a top-style display, where data is continuously overwritten in a single table.

Note

If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--interval | -I} <integer>`

Reports data at the interval specified in seconds.

`{--repeat | -r} <integer>`

Specifies how many times to run the report before quitting.

Note

To set the report to run indefinitely, specify `-1`.

`{--degraded | -d}`

Sets the report to continue running if some nodes do not respond.

`{--timeout | -o} <integer>`

Specifies number of seconds before remote commands time out.

`--running <integer>`

Displays statistics with aggregation of cluster statistics over a given interval in the number of iterations of the tool repetition interval. If running averages are requested, they appear on the row labeled **Avg**.

`--nodes`

Displays information on individual nodes.

`--timestamp`

Displays the time at which the `isi statistics` tool last gathered data. ;The timestamp is displayed in Epoch seconds.

`--oprates`

Displays the protocol operation rate statistics instead of the default throughput statistics.

isi status

Displays information about the current status of the cluster, events, and jobs.

Syntax

```
isi status
  [-q]
  [-r]
  [-w]
  [-D]
  [-d <storage-pool-name>]
  [-n <id>]
```

Options

-q

Omits event and protection operations and displays only information on the status of the cluster.

-r

Displays the raw size.

-w

Displays results without truncations.

-D

Displays more detailed information on running protection operations, including a list of worker processes. Also displays more information on failed protection operations, including a list of errors.

-d <storage-pool-name>

Displays a node pool or tier view of the file system instead of a cluster view. If a storage pool name such as a tier or a node pool is specified, only information for that pool is reported.

-n <id>

Displays the same information for an individual node, specified by logical node number (LNN), in addition to statistics for each disk in the node.

isi update

Updates a cluster to a newer version of OneFS.

You are prompted to specify where the image to use to update the cluster is located. After the image is loaded, you are prompted to reboot the cluster.

Syntax

```
isi update
  [--rolling]
  [--manual]
  [--drain-time <duration>]
  [--check-only]
```

Options

--rolling

Performs a rolling update, allowing the cluster to remain available during the update. When a rolling update is interrupted, the same update command can be issued to restart the rolling update. The update then attempts to continue where the previous update was interrupted. Rolling updates are not supported for all versions. Contact your Isilon representative for information about which versions support this option.

`--manual`

Causes rolling update process to pause and wait for user input before rebooting each node.

`--drain-time <duration>`

Sets the update process to suspend a node from its SmartConnect pool. The process then waits for clients to disconnect or for the specified *<duration>* to elapse before rebooting the node. The default *<duration>* units are in seconds. You can specify different time units by adding a letter to the end of the time, however. The following values are valid:

m
Minutes

h
Hours

d
Days

w
Weeks

`--check-only`

Provides information about potential failures across the cluster but does not initiate the upgrade process.

isi version

Displays detailed information about the Isilon cluster software properties.

Syntax

```
isi version
  [<os-info>]
```

Options

<os-info>

Optional variable that limits the output to specified pieces of information. If you do not include an *<os-info>* value, the system displays all information. Only the following values for *<os-info>* are acceptable.

osversion

Displays the name, build, release date, and current operating system version.

osbuild

Displays build information.

osrelease

Displays the version string for the software.

ostype

Displays the name of the operating system.

osrevision

Displays the revision number as a base-10 number.

copyright

Displays the current copyright information for the software.

Examples

The following command displays the name of the operating system only.

```
isi version ostype
```

isi_for_array

Runs commands on multiple nodes in an array, either in parallel or in serial.

When options conflict, the one specified last takes precedence.

Note

The `-k`, `-u`, `-p`, and `-q` options are valid only for SSH transport.

Syntax

```
isi_for_array
  [--array-name <array>]
  [--array-file <filename>]
  [--directory <directory>]
  [--diskless]
  [--known-hosts-file <filename>]
  [--user <user>]
  [--nodes <nodes>]
  [--password <password>]
  [--pre-command <command>]
  [--query-password]
  [--quiet]
  [--serial]
  [--storage]
  [--transport <transport-type>]
  [--throttle <setting>]
  [--exclude-nodes <nodes>]
  [--exclude-down-nodes]
```

Options

`{--array-name | -a} <array>`

Uses `<array>`.

`{--array-file | -A} <filename>`

Reads array information from `<filename>`. The default looks first for `$HOME/.array.xml`, then for `/etc/ifs/array.xml`.

`{--directory | -d} <directory>`

Runs commands from the specified directory on remote computers. The current working directory is the default directory. An empty `<directory>` results in commands being run in the user's home directory on the remote computer.

`{--diskless | -D}`

Runs commands from diskless nodes.

- `{--known-hosts-file | -k} <filename>`
 Uses *<filename>* for SSH known hosts file instead of the default `/dev/null` directory.
- `{--user | -u | -l} <user>`
 Logs in as *<user>* instead of as the default root user.
- `{--nodes | -n} <nodes>`
 Runs commands on the specified nodes, which can be specified multiple times. Must be a list of either node names or ranges of node IDs; for example, `1, 3-5, node8, 10`. If no nodes are explicitly listed, the whole array is used.
- `{--password | -p | --pw} <password>`
 Uses the specified password instead of the default password.
- `{--pre-command | -P} <command>`
 Runs the specified command before any other commands. This is useful for setting up the environment and it can be specified multiple times. You can specify `-` to reset the list of pre-commands.
- `{--query-password | -q}`
 Prompts the user for a password.
- `{--quiet | -Q}`
 Suppresses printing of the host prefix for each output line.
- `{--serial | -s}`
 Runs commands in serial instead of parallel.
- `{--storage | -S}`
 Run commands from storage nodes.
- `{--transport | -t} <transport-type>`
 Specifies the network transport type. The default value is `rpc`. Valid transports values are `rpc` or `ssh`.
- `{--throttle | -T} <setting>`
 Adjusts throttling. To disable throttling, specify `0`. The default value is `24`.
- `{--exclude-nodes | -x} <nodes>`
 Excludes specified nodes from the command. This argument is specified in the same manner as the `-n` option.
- `{--exclude-down-nodes | -X}`
 Excludes offline nodes from the command. This command is limited to cluster local use only.

Examples

In SmartLock compliance mode, to run `isi_for_array` for a command that requires root privileges, you must specify `sudo` twice. For example, the following command runs `isi stat` on each node in a compliance cluster.

```
sudo isi_for_array -u compadmin sudo isi stat
```

isi get

Displays information about a set of files, including the requested protection, current actual protection, and whether write-coalescing is enabled.

Requested protection appears in one of three colors: green, yellow, or red. Green indicates full protection. Yellow indicates degraded protection under a mirroring policy. Red indicates a loss of one or more data blocks under a parity policy.

Syntax

```
isi get {[[-a] [-d] [-g] [-s] [{-D | -DD | -DDC}] [-R] <path>}
| {[[-g] [-s] [{-D | -DD | -DDC}] [-R] -L <lin>}}
```

Options

-a

Displays the hidden "." and ".." entries of each directory.

-d

Displays the attributes of a directory instead of the contents.

-g

Displays detailed information, including snapshot governance lists.

-s

Displays the protection status using words instead of colors.

-D

Displays more detailed information.

-DD

Includes information about protection groups and security descriptor owners and groups.

-DDC

Includes cyclic redundancy check (CRC) information.

-R

Displays information about the subdirectories and files of the specified directories.

<path>

Displays information about the specified file or directory.

Specify as a file or directory path.

-L <lin>

Displays information about the specified file or directory.

Specify as a file or directory LIN.

Examples

The following command displays information on `ifs/home/` and all of its subdirectories:

```
isi get -R /ifs/home
```

The system displays output similar to the following example:

```

POLICY  LEVEL PERFORMANCE COAL  FILE
default 4x/2 concurrency on   ./
default 8x/3 concurrency on   ../
default 4x/2 concurrency on   admin/
default 4x/2 concurrency on   ftp/
default 4x/2 concurrency on   newUser1/
default 4x/2 concurrency on   newUser2/

/ifs/home/admin:
default 4+2/2 concurrency on   .zshrc

/ifs/home/ftp:
default 4x/2 concurrency on   incoming/
default 4x/2 concurrency on   pub/

/ifs/home/ftp/incoming:

/ifs/home/ftp/pub:

/ifs/home/newUser1:
default 4+2/2 concurrency on   .cshrc
default 4+2/2 concurrency on   .login
default 4+2/2 concurrency on   .login_conf
default 4+2/2 concurrency on   .mail_aliases
default 4+2/2 concurrency on   .mailrc
default 4+2/2 concurrency on   .profile
default 4+2/2 concurrency on   .rhosts
default 4+2/2 concurrency on   .shrc
default 4+2/2 concurrency on   .zshrc

/ifs/home/newUser2:
default 4+2/2 concurrency on   .cshrc
default 4+2/2 concurrency on   .login
default 4+2/2 concurrency on   .login_conf
default 4+2/2 concurrency on   .mail_aliases
default 4+2/2 concurrency on   .mailrc
default 4+2/2 concurrency on   .profile
default 4+2/2 concurrency on   .rhosts
default 4+2/2 concurrency on   .shrc
default 4+2/2 concurrency on   .zshrc

```

isi_gather_info

Collects and uploads the most recent cluster log information to SupportIQ.

Multiple instances of `-i`, `-f`, `-s`, `-S`, and `-l` are allowed.

`gather_expr` and `analysis_expr` can be quoted.

The default temporary directory is `/ifs/data/Isilon_Support/` (change with `-L` or `-T`).

Syntax

```

isi_gather_info
  [-h]
  [-v]
  [-u <user>]
  [-p <password>]
  [-i]
  [--incremental]
  [-l]
  [-f <filename>]

```



```

[-n <nodes>]
[--local-only]
[--skip-node-check]
[-s gather-script]
[-S gather-expr]
[-l gather-expr]
[-a analysis-script]
[-A analysis-expr]
[-t <tarfile>]
[-x exclude_tool]
[-I]
[-L]
[-T <temp-dir>]
[--tardir <dir>]
[--symlinkdir <dir>]
[--varlog_recent]
[--varlog_all]
[--nologs]
[--group <name>]
[--noconfig]
[--save-only]
[--save]
[--upload]
[--noupload]
[--re-upload <filename>]
[--verify-upload]
[--http]
[--nohttp]
[--http-host <host>]
[--http-path <dir>]
[--http-proxy <host>]
[--http-proxy-port <port>]
[--ftp]
[--noftp]
[--ftp-user <user>]
[--ftp-pass <password>]
[--ftp-host <host>]
[--ftp-path <dir>]
[--ftp-port <alt-port>]
[--ftp-proxy <host>]
[--ftp-proxy-port <port>]
[--ftp-mode <mode>]
[--esrs]
[--email]
[--noemail]
[--email-addresses]
[--email-subject]
[--email-body]
[--skip-size-check]

```

Options

- h
Prints this message and exits.
- v
Prints version info and exits.
- u <user>
Specifies the login as <user> instead of as the default root user.
- p <password>
Uses <password>.
- i

Includes only the listed utility. See also the `-l` option for a list of utilities to include. The special value `all` may be used to include every known utility.

`--incremental`
Gathers only those logs that changed since last log upload.

`-l`
Lists utilities and groups that can be included. See `-i` and `--group`.

`-f <filename>`
Gathers `<filename>` from each node. The value must be an absolute path.

`-n <nodes>`
Gathers information from only the specified nodes. Nodes must be a list or range of LNNs, for example, `1, 4-10, 12, 14`. If no nodes are specified, the whole array is used. Note that nodes are automatically excluded if they are down.

`--local-only`
Gathers information only from only the local node. Run this option when gathering files from the `/ifs` filesystem.

`--skip-node-check`
Skips the check for node availability.

`-s gather-script`
Runs `<gather-script>` on every node.

`-S gather-expr`
Runs `<gather-expr>` on every node.

`-l gather-expr`
Runs `<gather-expr>` on the local node.

`-a analysis-script`
Runs `<analysis-script>` on results.

`-A analysis-expr`
Runs `<analysis-expr>` on every node.

`-t <tarfile>`
Saves all results to the specified `<tarfile>` rather than to the default tar file.

`-x exclude_tool`
Excludes the specified tool or tools from being gathered from each node. Multiple tools can be listed as comma-separated values.

`-I`
Saves results to `/ifs`. This is the default setting.

`-L`
Save all results to local storage `/var/crash/support/`.

`-T <temp-dir>`
Saves all results to `<temp-dir>` instead of the default directory. `-T` overrides `-L` and `-l`.

`--tardir <dir>`
Places the final package directly into the specified directory.

`--symlinkdir <dir>`

Creates a symlink to the final package in the specified directory.

`--varlog_recent`
Gathers all logs in `/var/log`, with the exception of the compressed and rotated old logs. The default setting is all logs.

`--varlog_all`
Gathers all logs in `/var/log`, including compressed and rotated old logs. This is the default setting.

`--nologs`
Does not gather the required minimum number of logs.

`--group <name>`
Adds a specific group of utilities to the tar file.

`--noconfig`
Uses built-in default values and bypasses the configuration file.

`--save-only`
Saves the CLI-specified configuration to file and exits.

`--save`
Saves the CLI-specified configuration to file and runs it.

`--upload`
Uploads logs to Isilon Technical Support automatically. This is the default setting.

`--noupload`
Specifies no automatic upload to Isilon Technical Support.

`--re-upload <filename>`
Re-uploads the specified `<filename>`.

`--verify-upload`
Creates a tar file and uploads to test connectivity.

`--http`
Attempts HTTP upload. This is the default setting.

`--nohttp`
Specifies no HTTP upload attempt.

`--http-host <host>`
Specifies an alternate HTTP site for upload.

`--http-path <dir>`
Specifies an alternate HTTP upload directory.

`--http-proxy <host>`
Specifies the proxy server to use.

`--http-proxy-port <port>`
Specifies the proxy port to use.

`--ftp`
Attempts FTP upload. This setting is the default value.

`--noftp`
Specifies no FTP upload attempt.

`--ftp-user <user>`

Specifies an alternate user for FTP (default: anonymous).

`--ftp-pass <password>`
Specifies an alternate password for FTP.

`--ftp-host <host>`
Specifies an alternate FTP site for upload.

`--ftp-path DIR`
Specifies an alternate FTP upload directory.

`--ftp-port <alt-port>`
Specifies an alternate FTP port for upload.

`--ftp-proxy <host>`
Specifies the proxy server to use.

`--ftp-proxy-port <port>`
Specifies the proxy port to use.

`--ftp-mode <mode>`
Specifies the mode of FTP file transfer. The following values are valid: `both`, `active`, `passive`. The default value is `both`.

`--esrs`
Attempts ESRS upload.

`--email`
Attempts SMTP upload. If set, SMTP is tried first.

`--noemail`
Specifies no SMTP upload attempt. This is the default value.

`--email-addresses`
Specifies email addresses as comma-separated strings.

`--email-subject`
Specifies an alternative email subject.

`--email-body`
Specifies alternative email text shown on head of body.

`--skip-size-check`
Does not check the size of the gathered file.

Event commands

You can access and configure OneFS events and notification rules settings using the event commands. Running `isi events` without subcommands is equivalent to running `isi events list`.

isi events cancel

Cancels events.

Syntax

```
isi events cancel --instanceids <id>
```

Options

```
{--instanceids | -i} {<id> | <id,id,id...>}
```

Specifies the instance ID of the event that you want to cancel.

Specifies multiple event instances in a comma-separated list.

You can specify all event instances with `all`.

isi events list

Displays a list of system events.

Syntax

```
isi events list
[--oldest {-<rel-time> | <spec-time>}]
[--newest {-<rel-time> | <spec-time>}]
[--history]
[--coalesced]
[--severity <value>]
[--limit <integer>]
[--localdb]
[--nodes <node>]
[--types <integer>]
[--columns <column>]
[--sort-by <column>]
[--csv]
[--wide]
```

Options

```
{--oldest | -o} {-<rel-time> | <spec-time>}
```

Displays only events that have a start time after a specified date and time.

Specify `-<rel-time>` in the following format, where `d`, `h`, and `m` specify days, hours and minutes:

```
<integer>{d | h | m}
```

Specify `<spec-time>` in the following format, where `<mm>/<dd>/YYYY <HH>:<MM>` are the numerical month, day, year, hour and minute:

```
<mm>/<dd>/]YYYY <HH>:<MM>
```

```
{--newest | -n} {<rel-time> | <spec-time>
```

Displays only events that have a start time after a specified date and time

Specify *<rel-time>* in the following format, where d, h, and m specify days, hours and minutes:

```
<integer>{d | h | m}
```

Specify *<spec-time>* in the following format, where <MM>/<DD>/]YYYY <hh>:<mm> are the numerical month, day, year, hour and minute:

```
<MM>/<DD>/]YYYY <hh>:<mm>
```

```
{--history | -s}
```

Retrieves only historical events, which are those that have an end date or are quieted events.

```
{--coalesced | -C}
```

Includes coalesced events in results.

```
{--severity} | -v} <value>
```

Retrieves events for a specified level or levels . Multiple levels can be specified in a comma-separated list. the following values are valid:

```
info
```

```
warn
```

```
critical
```

```
emergency
```

```
{--limit | -L} <integer>
```

Limits results to the specified number of events.

```
{--localdb | -l}
```

Uses the local DB rather than the master.

```
--nodes <node>
```

Specifies which nodes to report statistics on. Default is all. The following values are valid:

- all
- <int>
- <int><int>

```
{--types | -i} <integer>
```

Retrieves all instances of the listed event types. Multiple types can be specified in a comma-separated list.

```
{--columns | -c} <column>
```

Specifies event list columns in a comma-separated list. The following values are valid:

```
type
```

```
id
```

```
coalesce_id
```

```
start_time
```

```
end_time
```

```

    lnn
    severity
    value
    quiet
    message
  {--sort-by | -b} <column>
    Specifies which column to sort the rows by. The default sort column is start_time.
  --csv
    Displays rows in CSV format and suppresses headers.
  {--wide | -w}
    Displays table in wide mode without truncations.

```

isi events notifications create

Creates a new event notification rule. Notifications rule parameters must be created in order. For example, an `--exclude` parameter that is specified after an `--include` parameter is not the same as specifying `--include` before `--exclude`.

Syntax

```

isi events notifications create --name <name>
  [--email <email-address>]
  [--snmp
  <SNMP-community>@<<SNMP host>]
  [--include-all <id>[,<id>]]
  [--include-info <id>[,<id>]]
  [--include-warn <id>[,<id>]]
  [--include-critical <id>[,<id>]]
  [--include-emergency <id>[,<id>]]
  [--exclude-all <id>[,<id>]]
  [--exclude-info <id>[,<id>]]
  [--exclude-warn <id>[,<id>]]
  [--exclude-critical <id>[,<id>]]
  [--exclude-emergency <id>[,<id>]]

```

Options

```

--name <name>
    Specifies the name of the notification rule being created.

--email <email-address>
    Specifies the email address to send an SNMP event. Multiple email address can be
    delimited with commas.

--snmp <SNMP-community>@<SNMP host>
    Specifies the SNMP community and hostname to send snmp event. Community and
    hostname are connected by an @ symbol. Multiple entries can be specified in a
    comma-separated list.

--include-all <id>[,<id>]
    Configures specified events for all severities (info, warn, critical, emergency). --
    include=all configures all events for all severities.

--include-info <id>[,<id>]

```

Configures specified events for info severity. `--include-info=all` configures all events for info.

`--include-warn <id>[,<id>]`

Configures specified events for warn severity. `--include-warn=all` configures all events for warn.

`--include-critical <id>[,<id>]`

Configures specified events for critical severity. `--include-critical=all` configures all events for critical.

`--include-emergency <id>[,<id>]`

Configures specified events for emergency severity. `--include-critical=all` configures all events for emergency.

`--exclude-all <id>[,<id>]`

Excludes specified events for all severities (info, warn, critical, emergency). `--exclude-all=all` results in no configured events.

`--exclude-info <id>[,<id>]`

Excludes specified events for info severity. `--exclude-info=all` excludes all info events.

`--exclude-warn <id>[,<id>]`

Excludes specified events for warn severity. `--exclude-warn=all` excludes all warn events.

`--exclude-critical <id>[,<id>]`

Excludes specified events for critical severity. `--exclude-critical=all` excludes all critical events.

`--exclude-emergency <id>[,<id>]`

Excludes specified events for emergency severity. `--exclude-emergency=all` excludes all emergency events.

isi events notifications modify

Modifies an event notification rule. Notifications rule parameters must be created in order.

Syntax

```
isi events notifications modify <name>
  [--email <email-address>]
  [--snmp
  <SNMP-community>@<<SNMP host>>]
  [--add-all <id>[,<id>]]
  [--add-info <id>[,<id>]]
  [--add-warn <id>[,<id>]]
  [--add-critical <id>[,<id>]]
  [--add-emergency <id>[,<id>]]
  [--delete-all <id>[,<id>]]
  [--delete-info <id>[,<id>]]
  [--delete-warn <id>[,<id>]]
  [--delete-critical <id>[,<id>]]
  [--delete-emergency <id>[,<id>]]
```

Options

<name>

Specifies the name of the notification rule being modified.

`--target-name <name>`

Specifies a new name for the notification rule.

`--email <email-address>`

Specifies the email address to send an SNMP event. Multiple email address can be delimited with commas.

`--snmp <SNMP-community>@<SNMP host>`

Specifies the SNMP community and hostname to send snmp event. Community and hostname are connected by an @ symbol. Multiple entries can be specified in a comma-separated list.

`--add-all <id>[,<id>]`

Configures specified events for all severities (info, warn, critical, emergency). `--add=all` configures all events for all severities.

`--add-info <id>[,<id>]`

Configures specified events for info severity. `--add-info=all` configures all events for info.

`--add-warn <id>[,<id>]`

Configures specified events for warn severity. `--add-warn=all` configures all events for warn.

`--add-critical <id>[,<id>]`

Configures specified events for critical severity. `--add-critical=all` configures all events for critical.

`--add-emergency <id>[,<id>]`

Configures specified events for emergency severity. `--add-critical=all` configures all events for emergency.

`--delete-all <id>[,<id>]`

Excludes specified events for all severities (info, warn, critical, emergency). `--deletes-all=all` results in no configured events.

`--delete-info <id>[,<id>]`

Excludes specified events for info severity. `--deletes-info=all` deletes all info events.

`--delete-warn <id>[,<id>]`

Excludes specified events for warn severity. `--deletes-warn=all` deletes all warn events.

`--delete-critical <id>[,<id>]`

Excludes specified events for critical severity. `--deletes-critical=all` deletes all critical events.

`--delete-emergency <id>[,<id>]`

Excludes specified events for emergency severity. `--deletes-emergency=all` deletes all emergency events.

isi events notifications delete

Deletes a notification rule.

Syntax

```
isi events notifications delete --name <name>
```

Options

```
{--name | -n} <name>
```

Specifies the name of the notification rule to delete.

isi events notifications list

Displays a list of settings for a notification rule.

Syntax

```
isi events notifications list
[--name <name>]
```

Options

```
{--name | -n} <name>
```

Specifies the name of the event notification rule.

isi events quiet

Marks an event as quieted. A quieted event is acknowledged when it is marked. The event is not removed or canceled.

Syntax

```
isi events quiet --instanceids <id>
```

Options

```
{--instanceids | -i} [<id>| <id,id,id...>]
```

Specifies the instance ID of the event that you want to quiet.

Specifies multiple event instances in a comma-separated list.

You can specify all event instances with `all`.

isi events sendtest

Sends a test event notification to verify event notification settings.

Syntax

```
isi events sendtest
--wait
```

Options

```
{--wait | -w}
```

Specifies a wait for an event existence in the master database.

isi events settings list

Displays a list of global settings and values.

Syntax

```
isi events settings list --name <name>
```

Options

```
{--name | -n} <name>
```

Specifies the name of the setting to display.

isi events settings set

Changes the values of global settings.

Syntax

```
isi events settings set --name <name> --value <value>
```

Options

```
{--name | -n} <name>
```

Specifies the name of the setting to be changed.

```
{--value | -v} <value>
```

Specifies the new value for the specified setting.

isi events show

Displays information for an event.

Syntax

```
isi events show --instanceid <id>
  [--wide]
  [--localdb]
```

Options

```
{--instanceid | -i} <id>
```

Specifies the ID of the event to view.

```
{--wide | -w}
```

Displays the event information in wide mode.

```
{--localdb | -l}
```

Uses localdb instead of the master.

isi events unquiet

Returns a quieted event to an unacknowledged state.

Syntax

```
isi events unquiet --instanceid <id>
```

Options

```
{--instanceid | -i} {<id> | <id,id,id...>}
```

Instance ID of the event that you want to unquiet.

Specify multiple event instances in a comma-separated list.

Specify all event instances with `all`.

Hardware commands

You can check the status of your cluster hardware, including specific node components, through the hardware commands.

isi batterystatus

Displays the current state of NVRAM batteries and charging systems on node hardware that supports this feature.

Syntax

```
isi batterystatus
```

Options

There are no options for this command.

Examples

To view the current state of NVRAM batteries and charging systems, run the following command:

```
isi batterystatus
```

The system displays output similar to the following example:

```
battery 1 : Good (10)
battery 2 : Good (10)
```

If the node hardware is not compatible, the system displays output similar to the following:

```
Battery status not supported on this hardware.
```

isi devices

Displays information about devices in the cluster and changes their status.

Syntax

```
isi devices
  [--action <action>]
  [--device {<LNN>:<drive> | <node-serial>}]
  [--grid]
  [--log <syslog-tag>]
  [--timeout <timeout>]
```

Options

If no options are specified with this command, the current status of the local node is displayed.

`--action <action>`

Designates the action to perform on a target device.
The following actions are available:

status

Displays the status of the given device.

smartfail

SmartFails the given node or drive.

stopfail

Ends the SmartFail operation on the given node. If the node is still attached to the cluster, the node is returned to a healthy state. If the node has been removed from the cluster, the node enters a suspended state.

add

Adds the given node or drive to the cluster.

format

Formats the specified drive.

fwstatus

Displays the firmware status of drives.

fwupdate

Updates the firmware of drives.

discover

Scans the internal cluster network for nodes that have not been added to the cluster.

confirm

Displays the join status of the node.

queue

Queues the specified node to be added to the cluster. Once the node is connected to the internal cluster network, the node will be automatically added to the cluster. Specify `--device` as a node serial number.

unqueue

Removes the specified node from the queue of nodes to be added to the cluster. If the node is connected to the internal cluster network, the node will not be automatically added to the cluster. Specify `--device` as a node serial number.

queuelist

Displays the list of nodes that are queued to be added to the cluster.

```
{--device | -d} {<LNN>:<drive> | <node-serial>}
```

Sets the target device on which to perform an action. If only <LNN> is specified, the action is performed on the entire node. If only <drive> is specified, the action is performed on the specified drive in the local node.

The following values are valid for <drive>:

<N>	Where <N> is a valid bay number.
bay<N>	Where <N> is a valid bay number.
Inum<N>	Where <N> is a valid Inum number.

Specify <node-serial> as a node serial number.

```
{--grid | -g}
```

Formats the requested system output to display in a grid to represent the physical layout of the drive bays in the node chassis.

```
{--log | -L} <syslog-tag>
```

If the command succeeds, tags the command output with the specified tag and logs it in /var/log/messages. This option is valid only if one of the following actions is being performed: smartfail, stopfail, add, format, purpose, fwupdate, queue, unqueue, queuelist.

```
{--timeout | -t} <timeout>
```

Establishes a timeout limit for the cluster information gathering period.

isi servicelight status

Indicates whether the LED service light on the back panel of a node is on or off.

Syntax

```
isi servicelight status
```

Options

There are no options for this command.

Examples

To display the status of the service light, run the following command.

```
isi servicelight status
```

The system displays output similar to the following example.

```
The service LED is off
```

isi servicelight off

Turns off the LED service light on the back panel of a node.

Syntax

```
isi servicelight off
```

Options

There are no options for this command.

Examples

To turn off the LED service light on the back panel of a node, run the following command.

```
isi servicelight off
```

isi servicelight on

Turns on the LED service light on the back panel of a node.

Syntax

```
isi servicelight on
```

Options

There are no options for this command.

Examples

To turn on the LED service light on the back panel of a node, run the following command.

```
isi servicelight on
```

isi drivefirmware status

Displays the status of drive firmware on the cluster.

Syntax

```
isi drivefirmware status
  [--action <action>]
  [--device <node>:<drive>]
  [--grid]
  [--log <syslog-tag>]
  [--timeout <timeout>]
```

Options

If no options are specified with this command, the current drive firmware status of the local node is displayed.

{--local | -L}

Displays information from the local node only.

{--diskless | -D}

Displays information only from diskless nodes such as accelerators.

{--storage | -S}

Displays information only from storage nodes.

{--include-nodes | -n} <nodes>

Displays information only from the specified nodes.

{--exclude-nodes | -x} <nodes>

Displays information from all nodes except those that are specified.

--verbose

Displays more detailed information.

isi firmware package

Displays information related to the installed firmware package.

Syntax

```
isi firmware package
  [--local]
  [--diskless]
  [--storage]
  [--include-nodes <nodes>]
  [--exclude-nodes <nodes>]
```

Options

{--local | -L}

Displays information from the local node only.

{--diskless | -D}

Displays information only from diskless nodes such as accelerators.

{--storage | -S}

Displays information only from storage nodes.

{--include-nodes | -n} <nodes>

Displays information only from the specified nodes.

{--exclude-nodes | -x} <nodes>

Displays information from all nodes except those that are specified.

Examples

To display the status of all firmware, run the following command:

```
isi firmware status
```

To display firmware package information from all storage nodes in the cluster, run the following command:

```
isi firmware package --storage
```

isi firmware status

Displays information on firmware types and versions.

Syntax

```
isi firmware status
  [--local]
  [--diskless]
  [--storage]
  [--include-nodes <nodes>]
  [--exclude-nodes <nodes>]
  [--include-device <device>]
  [--exclude-device <device>]
  [--include-type <device-type>]
  [--exclude-type <device-type>]
  [--save]
  [--verbose]
```


Options`{--local | -L}`

Displays information from the local node only.

`{--diskless | -D}`

Displays information only from diskless nodes such as accelerators.

`{--storage | -S}`

Displays information only from storage nodes.

`{--include-nodes | -n} <nodes>`

Displays information only from the specified nodes.

`{--exclude-nodes | -x} <nodes>`

Displays information from all nodes except those that are specified.

`{--include-device | -d} <device>`

Displays information only from the specified device.

`--exclude-device <device>`

Displays information from all devices except those that are specified.

`{--include-type | -t} <device-type>`

Displays information only from the specified device type.

`--exclude-type <device-type>`

Displays information from all device types except those that are specified.

`--save`Save the output of the status to `/etc/ifs/firmware_versions`.`{--verbose | -v}`

Displays more detailed information.

Examples

To display firmware package information from nodes two and three, run the following command:

```
isi firmware status --include-nodes 2,3
```

The system displays output similar to the following example.

Device	Type	Firmware	Nodes
IsilonIB	Network	4.8.930+205-0002-05_A	2-3
Lsi	DiskCtrl	6.28.00.00+01.28.02.00+1.17+0.99c	2-3

To display firmware package information for the network device type, run the following command:

```
isi firmware status --include-type network
```

The system displays output similar to the following example.

Device	Type	Firmware	Nodes
IsilonIB	Network	4.8.930+205-0002-05_A	1-6

isi firmware update

Updates firmware on devices within the cluster to match those in the installed firmware package.

Each node is updated one at a time. Each time a node is updated, the node is restarted. The system does not begin to update the next node until the previous node has rejoined the cluster.

Syntax

```
isi firmware update
  [--local]
  [--diskless]
  [--storage]
  [--include-nodes <nodes>]
  [--exclude-nodes <nodes>]
  [--include-device <device>]
  [--exclude-device <device>]
  [--include-type <device-type>]
  [--exclude-type <device-type>]
  [--force]
  [--verbose]
```

Options

{--local | -L}

Updates the local node only.

{--diskless | -D}

Updates diskless nodes such as accelerators only.

{--storage | -S}

Updates storage nodes only.

{--include-nodes | -n} <nodes>

Updates the specified nodes only.

{--exclude-nodes | -x} <nodes>

Updates all nodes except those that are specified.

{--include-device | -d} <device>

Updates the specified device only.

--exclude-device <device>

Updates all devices except those that are specified.

{--include-type | -t} <device-type>

Updates the specified device type only.

--exclude-type <device-type>

Updates all device types except those that are specified.

--force

Forces the update.

{--verbose | -v}

Displays more detailed information.

Examples

To update the firmware on nodes two and three, run the following command:

```
isi firmware update --include-nodes 2,3
```

To update the firmware for the network device type only, run the following command:

```
isi firmware update --include-type network
```

isi readonly off

Sets nodes to read-write mode.

This command only clears any user-specified requests for read-only mode. If the node has been placed into read-only mode by the system, it will remain in read-only mode until the conditions which triggered read-only mode have cleared.

Syntax

```
isi readonly off
  [--nodes <nodes>]
  [--verbose]
```

Options

If no options are specified, the local node is set to read-write mode.

`--nodes <nodes>`

Specifies the nodes to apply read-write settings to. The following values for `<nodes>` are valid:

- all
- *
- `<int>`
- `<int>-<int>`

Multiple values can be specified in a comma-separated list.

`{--verbose | -v}`

Displays more detailed information.

Examples

To apply read-write settings to every node in the cluster, run the following command.

```
isi readonly off --nodes all
```

The system displays output similar to the following example.

```
Read-only changes committed successfully
```

Use the `isi readonly show` command to confirm the read-write settings of the cluster. The system displays output similar to the following example.

node	mode	status
1	read/write	
2	read/write	
3	read/write	
4	read/write	

```
5 read/write
6 read/write
```

isi readonly on

Sets nodes to read-only mode.

If read-only mode is currently disallowed for this node, it will remain read/write until read-only mode is allowed again.

Syntax

```
isi readonly on
  [--nodes <nodes>]
  [--verbose]
```

Options

If no options are specified, the local node is set to read-only mode.

`--nodes <nodes>`

Specifies the nodes to apply read-only settings to. The following values for *<nodes>* are valid:

- all
- *
- *<int>*
- *<int>-<int>*

Multiple values can be specified in a comma-separated list.

`{--verbose | -v}`

Displays more detailed information.

Examples

To apply read-only settings to every node in the cluster, run the following command.

```
isi readonly on --nodes all
```

The system displays output similar to the following example.

```
Read-only changes committed successfully
```

Use the `isi readonly show` command to confirm the read-only settings of the cluster. The system displays output similar to the following example.

node	mode	status
1	read-only	user-ui
2	read-only	user-ui
3	read-only	user-ui
4	read-only	user-ui
5	read-only	user-ui
6	read-only	user-ui

isi readonly show

Displays a list of read-only settings for the cluster.

Syntax

```
isi readonly show
```

Options

There are no options for this command.

Examples

To display the read-only settings for the cluster, run the following command.

```
isi readonly show
```

The system displays output similar to the following example.

node	mode	status
1	read/write	
2	read/write	
3	read/write	
4	read/write	
5	read/write	
6	read/write	

CHAPTER 5

Access zones

This section contains the following topics:

- [Access zones overview](#) 168
- [Access zone base directory rules](#) 168
- [Access zones best practices](#) 169
- [Access zone limits](#) 169
- [Quality of service](#) 170
- [Managing access zones](#) 170
- [Access zone commands](#) 172

Access zones overview

Although the default view of an EMC Isilon cluster is that of one physical machine, you can partition a cluster into multiple virtual containers called access zones. Access zones allow you to isolate data and control who can access data in each zone.

Access zones support all configuration settings for authentication and identity management services on a cluster, so you can configure authentication providers and provision SMB shares and NFS exports on a zone-by-zone basis. When you create an access zone, a local provider is created automatically, which allows you to configure each access zone with a list of local users and groups. You can also authenticate through a different authentication provider in each access zone.

To control data access, you can direct incoming connections to the access zone through a specific IP address pool. Associating an access zone with an IP address pool restricts authentication to the associated access zone and reduces the number of available and accessible SMB shares and NFS exports. Another advantage to multiple access zones is the ability to configure audit protocol access for individual access zones. You can modify the default list of successful and failed protocol audit events and then generate reports through a third-party tool for an individual access zone.

A cluster includes a built-in access zone named System, where you manage all aspects of a cluster and other access zones. By default, all cluster IP addresses connect to the System zone. Even if you create additional access zones, you configure all access zones in the System zone. Role-based access, which primarily allows configuration actions, is available through only the System zone. All administrators, including those given privileges by a role, must connect to the System zone to configure a cluster.

Configuration management of a non-System access zone is not permitted through SSH, the OneFS Platform API, or the web administration interface. However, you can create and delete SMB shares in an access zone through the Microsoft Management Console (MMC).

Access zone base directory rules

You must assign a base directory to each access zone. A base directory defines the file system tree exposed by an access zone and isolates data contained in the directory to the access zone.

A base directory path is unique for each access zone and cannot overlap or be nested inside base directories of other access zones.

Base directories restrict configuration options for several features such as SMB share paths, NFS exports, the HDFS root directory, and the local provider home directory template. You must observe the following rules when specifying a base directory:

- The base directory cannot be identical to the base directory of any other access zone, except the System zone. For example, you cannot specify `/ifs/data/hr` to both the zone2 and zone3 access zones.
- Cannot overlap with the file system tree of a base directory in any other access zone, except the System zone. For example, if `/ifs/data/hr` is assigned to zone2, you cannot assign `/ifs/data/hr/personnel` to zone3.
- The base directory of the default System access zone is `/ifs` and cannot be modified.

Note

Assigning a base directory that is identical to or overlaps with the System zone is allowed, but only recommended as a temporary base directory when modifying the base directory path and migrating data to the new directory.

Access zones best practices

You can avoid configuration problems on the EMC Isilon cluster when creating access zones by following best practices guidelines.

Best practice	Details
Create unique base directories.	The base directory path of each access zone must be unique and cannot overlap or be nested inside the base directory of another access zone.
Separate the function of the System zone from other access zones.	If you choose to create any additional access zones, do not allow data access in both the System zone and created zones. Reserve the System zone for configuration access, and create additional zones for data access. Move current data out of the System zone and into a new access zone.
Create access zones to isolate data access for different clients or users.	Do not create access zones if a workflow requires data sharing between different classes of clients or users.
Assign only one authentication provider of each type to each access zone.	An access zone is limited to a single Active Directory provider; however, OneFS allows multiple LDAP, NIS, and file authentication providers in each access zone. It is recommended that you assign only one type of each provider per access zone in order to simplify administration.
Avoid overlapping UID or GID ranges for authentication providers in the same access zone.	The potential for zone access conflicts is slight but possible if overlapping UIDs/GIDs are present in the same access zone.
Configure a single DNS server for all access zones.	OneFS does not support one DNS server per access zone. It is recommended that all access zones point to a single DNS server.

Access zone limits

You can follow access zone limits guidelines to help size the workloads on the OneFS system.

If you configure multiple access zones on an EMC Isilon cluster, limits guidelines are recommended for optimal system performance. The limits described in the *EMC Isilon Guidelines for Large Workloads* publication are recommended for heavy enterprise workflows on a cluster, treating each access zone as a separate physical machine.

Quality of service

You can set upper bounds on quality of service by assigning specific physical resources to each access zone.

Quality of service addresses physical hardware performance characteristics that can be measured, improved, and sometimes guaranteed. Characteristics measured for quality of service include but are not limited to throughput rates, CPU usage, and disk capacity. When you share physical hardware in an EMC Isilon cluster across multiple virtual instances, competition exists for the following services:

- CPU
- Memory
- Network bandwidth
- Disk I/O
- Disk capacity

Access zones do not provide logical quality of service guarantees to these resources, but you can partition these resources between access zones on a single cluster. The following table describes a few ways to partition resources to improve quality of service:

Use	Notes
NICs	You can assign specific NICs on specific nodes to an IP address pool that is associated with an access zone. By assigning these NICs, you can determine the nodes and interfaces that are associated with an access zone. This enables the separation of CPU, memory, and network bandwidth.
SmartPools	SmartPools are separated by node hardware equivalence classes, usually into multiple tiers of high, medium, and low performance. The data written to a SmartPool is written only to the disks in the nodes of that pool. Associating an IP address pool with only the nodes of a single SmartPool enables partitioning of disk I/O resources.
SmartQuotas	Through SmartQuotas, you can limit disk capacity by a user or a group or in a directory. By applying a quota to an access zone's base directory, you can limit disk capacity used in that access zone.

Managing access zones

You can create access zones on the EMC Isilon cluster, view and modify access zone settings, and delete access zones.

Create an access zone

You can create an access zone to isolate data and restrict which users can access the data.

Procedure

1. Run the `isi zone zones create` command.

The following command creates an access zone named zone3 and sets the base directory to `/ifs/hr/data`:

```
isi zone zones create zone3 /ifs/hr/data
```

The following command creates an access zone named zone3, sets the base directory to `/ifs/hr/data` and creates the directory on the EMC Isilon cluster if it does not already exist:

```
isi zone zones create zone3 /ifs/hr/data --create-path
```

Associate an IP address pool with an access zone

You can associate an IP address pool with an access zone to ensure that clients can only connect to the access zone through IP addresses in the pool.

Procedure

1. Run the `isi networks modify pool` command.

You must specify the name of the IP address pool in the following format: `<subnet-name>:<pool-name>`.

The following command associates zone3 with pool1 on subnet1:

```
isi networks modify pool --name=subnet1:pool1 --access-zone=zone3
```

Modify an access zone

You can modify the properties of any access zone except the name of the built-in System zone.

Procedure

1. Run the `isi zone zones modify` command.

The following command renames the zone3 access zone to zone5 and removes all current authentication providers from the access zone:

```
isi zone zones modify zone3 --name=zone5 --clear-auth-providers
```

Add an authentication provider to an access zone

You can add an authentication provider to an existing access zone.

Procedure

1. Run the `isi zone zones modify` command with the `--add-auth-providers` option.

You must specify the name of the authentication provider in the following format: `<provider-type>:<provider-name>`.

The following command adds a file authentication provider named HR-Users to the zone3 access zone:

```
isi zone zones modify zone3 --add-auth-providers=file:hr-users
```

Remove an authentication provider from an access zone

You can remove an authentication provider from an access zone. This does not remove the authentication provider from the EMC Isilon cluster; the provider remains available for future use.

Procedure

1. Run the `isi zone zones modify` command with the `--remove-auth-providers` option.

You must specify the name of the authentication provider in the following format:
<provider-type>.<provider-name>.

The following command removes the file authentication provider named HR-Users from the zone3 access zone:

```
isi zone zones modify zone3 --remove-auth-providers=file:hr-users
```

The following command removes all authentication providers from the zone3 access zone:

```
isi zone zones modify zone3 --clear-auth-providers
```

Delete an access zone

You can delete any access zone except the built-in System zone. When you delete an access zone, all associated authentication providers remain available to other access zones, but IP addresses are not reassigned to other access zones. SMB shares, NFS exports, and HDFS data paths are deleted when you delete an access zone; however, the directories and data still exist, and you can map new shares, exports, or paths in another access zone.

Procedure

1. Run the `isi zone zones delete` command.

The following command deletes the zone3 access zone :

```
isi zone zones delete zone3
```

Access zone commands

You can configure and manage access zones through access zone commands.

isi zone restrictions create

Prohibits user or group access to the `/ifs` directory. Attempts to read or write files by restricted users or groups return `ACCESS DENIED` errors.

Syntax

```
isi zone restrictions create <zone> {<user> | --uid <integer>
| --group <string> | --gid <integer> | --sid <string>
| --wellknown <string>}
[--verbose]
```

Options

<zone>

Specifies an access zone by name.

`<user>`
Specifies a user by name.

`--uid <integer>`
Specifies a user by UID.

`--group <string>`
Specifies a group by name.

`--gid <integer>`
Specifies a group by GID.

`--sid <string>`
Specifies an object by user or group SID.

`--wellknown <name>`
Specifies a well-known user, group, machine, or account name.

`{--verbose | -v}`
Returns a success or fail message after running the command.

isi zone restrictions delete

Removes a restriction that prohibits user or group access to the `/ifs` directory.

Syntax

```
isi zone restrictions delete <zone> {<user> | --uid <integer>
| --group <string> | --gid <integer> | --sid <string>
| --wellknown <string>}
[--force]
[--verbose]
```

Options

`<zone>`
Specifies an access zone by name.

`<user>`
Specifies a user by name.

`--uid <integer>`
Specifies a user by UID.

`--group <string>`
Specifies a group by name.

`--gid <integer>`
Specifies a group by GID.

`--sid <string>`
Specifies an object by user or group SID.

`--wellknown <string>`
Specifies an object by well-known SID.

`{--force | -f}`
Suppresses command-line prompts and messages.

```
{--verbose | -v}
```

Returns a success or fail message after running the command.

isi zone restrictions list

Displays a list of users or groups that are prohibited from accessing the `/ifs` directory.

Syntax

```
isi zone restrictions list <zone>
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<zone>

Specifies an access zone by name.

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

Examples

To display a list of restricted users for the built-in System zone, run the following command:

```
isi zone restrictions list system
```

isi zone zones create

Creates an access zone.

Syntax

```
isi zone zones create <name> <path>
  [--cache-size <size>]
  [--map-untrusted <workgroup>]
  [--auth-providers <provider-type>:<provider-name>]
  [--netbios-name <string>]
  [--all-auth-providers {yes | no}]
  [--user-mapping-rules <string>]
  [--home-directory-umask <integer>]
  [--skeleton-directory <string>]
  [--audit-success <operations>]
  [--audit-failure <operations>]
  [--hdfs-authentication {all | simple_only | kerberos_only}]
```

```
[--hdfs-root-directory <path>]
[--webhdfs-enabled {yes | no}]
[--hdfs-ambari-server <string>]
[--hdfs-ambari-namenode <string>]
[--syslog-forwarding-enabled {yes | no}]
[--syslog-audit-events <operations>]
[--create-path]
[--verbose]
```

Options

<name>

Specifies the name of the access zone.

<path>

Specifies the base directory path for the zone. Paths for zones must not overlap, meaning that you cannot create nested zones.

--cache-size <size>

Specifies the maximum size of the zone's in-memory identity cache in bytes. Valid values are integers in the range 1000000 - 50000000. The default value is 5000000.

--map-untrusted <workgroup>

Maps untrusted domains to the specified NetBIOS workgroup during authentication.

--auth-providers <provider-type>:<provider-name>

Specifies one or more authentication providers, separated by commas, for authentication to the access zone. Authentication providers are checked in the order specified. You must specify the name of the authentication provider in the following format: <provider-type>:<provider-name>.

--netbios-name <string>

Specifies the NetBIOS name.

--all-auth-providers {yes | no}

Specifies whether to authenticate through all available authentication providers. If no, authentication is through the list of providers specified by the --auth-providers option.

--user-mapping-rules <string>

Specifies one or more user mapping rules, separated by commas, for the access zone.

--home-directory-umask <integer>

Specifies the permissions to set on auto-created user home directories.

--skeleton-directory <string>

Sets the skeleton directory for user home directories.

--audit-success <operations>

Specifies one or more filters, separated by commas, for auditing protocol operations that succeeded. The following operations are valid:

- close
- create
- delete
- get_security

- logoff
 - logon
 - read
 - rename
 - set_security
 - tree_connect
 - write
- `--audit-failure <operations>`
 Specifies one or more filters, separated by commas, for auditing protocol operations that failed. The following operations are valid:
- close
 - create
 - delete
 - get_security
 - logoff
 - logon
 - read
 - rename
 - set_security
 - tree_connect
 - write
 - all
- `--hdfs-authentication <operations>`
 Specifies the allowed authentication type for the HDFS protocol. Valid values are `all`, `simple_only`, or `kerberos_only`
- `--hdfs-root-directory <path>`
 Specifies the root directory for the HDFS protocol.
- `--webhdfs-enabled {yes | no}`
 Enables or disables WebHDFS on the zone.
- `--hdfs-ambari-server <string>`
 Specifies the Ambari server that receives communication from an Ambari agent. The value must be a resolvable hostname, FQDN, or IP address.
- `--hdfs-ambari-namenode <string>`
 Specifies a point of contact in the access zone that Hadoop services managed through the Ambari interface should connect through. The value must be a resolvable IP address or a SmartConnect zone name.
- `--syslog-forwarding-enabled {yes | no}`
 Enables or disables syslog forwarding of zone audit events.
- `--syslog-audit-events <operations>`
 Sets the filter for the auditing protocol operations to forward to syslog. You must specify the `--syslog-audit-events` parameter for each additional filter. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

The `all` option specifies all valid filter operations.

`--create-path`

Specifies that the value entered as the access zone path is to be created if it doesn't already exist.

`{--verbose | -v}`

Displays the results of running the command.

isi zone zones delete

Deletes an access zone. All authentication providers that are associated with the access zone remain available to other zones, but IP addresses are not reassigned. You cannot delete the built-in System zone.

Syntax

```
isi zone zones delete <zone>
  [--force]
  [--verbose]
```

Options

`<zone>`

Specifies the name of the access zone to delete.

`{--force | -f}`

Suppresses command-line prompts and messages.

`{--verbose | -v}`

Displays the results of running the command.

isi zone zones list

Displays a list of access zones in the cluster.

Syntax

```
isi zone zones list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
```

```
[--no-footer]
[--verbose]
```

Options

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

Examples

To view a list of all access zones in the cluster, run the following command:

```
isi zone zones list
```

isi zone zones modify

Modifies an access zone.

Syntax

```
isi zone zones modify <zone>
  [--name <string>]
  [--path <path>]
  [--cache-size <size>]
  [--map-untrusted <string>]
  [--auth-providers <provider-type>:<provider-name>]
  [--clear-auth-providers]
  [--add-auth-providers <provider-type>:<provider-name>]
  [--remove-auth-providers <provider-type>:<provider-name>]
  [--netbios-name <string>]
  [--all-auth-providers {yes | no}]
  [--user-mapping-rules <string>]
  [--clear-user-mapping-rules]
  [--add-user-mapping-rules <string>]
  [--remove-user-mapping-rules <string>]
  [--home-directory-umask <integer>]
  [--skeleton-directory <string>]
  [--audit-success <operations>]
  [--clear-audit-success]
  [--add-audit-success <operations>]
  [--remove-audit-success <operations>]
  [--audit-failure <operations>]
  [--clear-audit-failure]
  [--add-audit-failure <operations>]
  [--remove-audit-failure <operations>]
  [--hdfs-authentication {all | simple_only | kerberos_only}]
  [--hdfs-root-directory <path>]
  [--webhdfs-enabled {yes | no}]
  [--hdfs-ambari-server <string>]
  [--hdfs-ambari-namenode <string>]
  [--syslog-forwarding-enabled {yes | no}]
```

```
[--syslog-audit-events <operations>]
[--clear-syslog-audit-events <operations>]
[--add-syslog-audit-events <operations>]
[--remove-syslog-audit-events <string>]
[--create-path]
[--verbose]
```

Options

<zone>

Specifies the name of the access zone to modify.

`--name <string>`

Specifies a new name for the access zone. You cannot change the name of the built-in System access zone.

`--path <path>`

Specifies the base directory path for the zone. Paths for zones must not overlap, meaning that you cannot create nested zones.

`--cache-size <size>`

Specifies the maximum size of the zone's in-memory cache in bytes. Valid values are integers in the range 1000000 - 50000000. The default value is 50000000.

`--map-untrusted <string>`

Specifies the NetBIOS workgroup to map untrusted domains to during authentication.

`--auth-providers <provider-type>:<provider-name>`

Specifies one or more authentication providers, separated by commas, for authentication to the access zone. This option overwrites any existing entries in the authentication providers list. To add or remove providers without affecting the current entries, configure settings for `--add-auth-providers` or `--remove-auth-providers`.

`--clear-auth-providers`

Removes all authentication providers from the access zone.

`--add-auth-providers <provider-type>:<provider-name>`

Adds one or more authentication providers, separated by commas, to the access zone.

`--remove-auth-providers <provider-type>:<provider-name>`

Removes one or more authentication providers, separated by commas, from the access zone.

`--netbios-name <string>`

Specifies the NetBIOS name.

`--all-auth-providers {yes | no}`

Specifies whether to authenticate through all available authentication providers. If `no`, authentication is through the list of providers specified by the `--auth-providers` option.

`--user-mapping-rules <string>`

Specifies one or more user mapping rules, separated by commas, for the access zone. This option overwrites all entries in the user mapping rules list. To add or remove mapping rules without overwriting the current entries, configure settings with `--add-user-mapping-rules` or `--remove-user-mapping-rules`.

```
--clear-user-mapping-rules
    Removes all user mapping rules from the access zone.
--add-user-mapping-rules <string>
    Adds one or more user mapping rules, separated by commas, to the access zone.
--remove-user-mapping-rules <string>
    Removes one or more user mapping rules, separated by commas, from the access
    zone.
--home-directory-umask <integer>
    Specifies the permissions to set on auto-created user home directories.
--skeleton-directory <string>
    Sets the skeleton directory for user home directories.
--audit-success <operations>
    Specifies one or more filters, separated by commas, for auditing protocol operations
    that succeeded. This option overwrites the current list of filter operations. The
    following operations are valid:
    • close
    • create
    • delete
    • get_security
    • logoff
    • logon
    • read
    • rename
    • set_security
    • tree_connect
    • write
    • all
    To add or remove filters without affecting the current list, configure settings with --
    add-audit-success or --remove-audit-success.
--clear-audit-success
    Clears all filters for auditing protocol operations that succeeded.
--add-audit-success <operations>
    Adds one or more filters, separated by commas, for auditing protocol operations that
    succeeded. The following operations are valid:
    • close
    • create
    • delete
    • get_security
    • logoff
    • logon
    • read
```

- rename
- set_security
- tree_connect
- write
- all

`--remove-audit-success <operations>`

Removes one or more filters, separated by commas, for auditing protocol operations that succeeded. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

`--audit-failure <operations>`

Specifies one or more filters, separated by commas, for auditing protocol operations that failed. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

This option overwrites the current list of filter operations. To add or remove filters without affecting the current list, configure settings with `--add-audit-failure` or `--remove-audit-failure`.

`--clear-audit-failure`

Clears all filters for auditing protocol operations that failed.

`--add-audit-failure <operations>`

Adds one or more filters, separated by commas, for auditing protocol operations that failed. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

`--remove-audit-failure <operations>`

Removes one or more filters, separated by commas, for auditing protocol operations that failed. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

`--hdfs-authentication <operations>`

Specifies the allowed authentication type for the HDFS protocol. Valid values are `all`, `simple_only`, or `kerberos_only`.

`--hdfs-root-directory <path>`

Specifies the root directory for the HDFS protocol.

`--webhdfs-enabled {yes | no}`

Enables or disables WebHDFS on the zone.

`--hdfs-ambari-server <string>`

Specifies the Ambari server that receives communication from an Ambari agent. The value must be a resolvable hostname, FQDN, or IP address.

`--hdfs-ambari-namenode <string>`

Specifies a point of contact in the access zone that Hadoop services managed through the Ambari interface should connect through. The value must be a resolvable IP address or a SmartConnect zone name.

`--syslog-forwarding-enabled {yes | no}`

Enables or disables syslog forwarding of zone audit events.

`--syslog-audit-events <operations>`

Sets the filter for the auditing protocol operations to forward to syslog. You must specify the `--syslog-audit-events` parameter for each additional filter. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

The `all` option specifies all valid filter operations.

`--clear-syslog-audit-events <operations>`

Clears the filter setting for the auditing protocol operations that are forwarded to syslog. You can specify `all` to clear all valid filter operations.

`--add-syslog-audit-events <operations>`

Adds a filter for the auditing protocol operations to forward to syslog. You must specify the `--syslog-audit-events` parameter for each additional filter. The following operations are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

`--remove-syslog-audit-events <string>`

Removes a filter for the auditing protocol operations to forward to syslog. You must specify the `--syslog-audit-events` parameter for each additional filter. The `all` option specifies all valid filter operations. Specify `--remove-syslog-audit-events` for each filter setting that you want to add.

`--create-path`

Specifies that the zone path is to be created if it doesn't already exist.

`{--verbose | -v}`

Displays the results of running the command.

isi zone zones view

Displays the properties of an access zone.

Syntax

```
isi zone zones view <zone>
```

Options

<zone>

Specifies the name of the access zone to view.

CHAPTER 6

Authentication and access control

This section contains the following topics:

- [Authentication and access control overview](#) 186
- [Role-based access](#) 186
- [Authentication](#) 198
- [Data access control](#) 202
- [Authorization](#) 202
- [Managing roles](#) 205
- [Managing authentication providers](#) 206
- [Managing access permissions](#) 225
- [Authentication and access control commands](#) 228

Authentication and access control overview

OneFS supports several methods for ensuring that your cluster remains secure, including UNIX- and Windows-style permissions for data-level access control, access zones for data isolation, and role-based administration control access to system configuration settings.

OneFS is designed for a mixed environment that allows you to configure both Access Control Lists (ACLs) and standard UNIX permissions on the cluster file system.

Note

In most situations, the default settings are sufficient. You can configure additional access zones, custom roles, and permissions policies as necessary for your particular environment.

Role-based access

You can assign role-based access to delegate administrative tasks to selected users.

Role based access control (RBAC) allows the right to perform particular administrative actions to be granted to any user who can authenticate to a cluster. Roles are created by a Security Administrator, assigned privileges, and then assigned members. All administrators, including those given privileges by a role, must connect to the System zone to configure the cluster. When these members log in to the cluster through a configuration interface, they have these privileges. All administrators can configure settings for access zones, and they always have control over all access zones on the cluster.

Roles also give you the ability to assign privileges to member users and groups. By default, only the root user and the admin user can log in to the web administration interface through HTTP or the command-line interface through SSH. Using roles, the root and admin users can assign others to built-in or customer roles that have login and administrative privileges to perform specific administrative tasks.

Note

As a best practice, assign users to roles that contain the minimum set of necessary privileges. For most purposes, the default permission policy settings, system access zone, and built-in roles are sufficient. You can create role-based access management policies as necessary for your particular environment.

Roles and privileges

In addition to controlling access to files and directories through ACLs and POSIX mode bits, OneFS controls configuration-level access through administrator roles. A role is a collection of OneFS privileges that are usually associated with a configuration subsystem. Those privileges are granted to members of that role as they log in to the cluster through the Platform API, command-line interface, or web administration interface

Roles

You can permit and limit access to administrative areas of your EMC Isilon cluster on a per-user basis through roles.

OneFS includes built-in administrator roles with predefined sets of privileges that cannot be modified. The following list describes what you can and cannot do through roles:

- You can assign privileges through role membership.
- You can add any user to a role as long as the user can authenticate to the cluster.
- You can create custom roles and assign privileges to those roles.
- You can add users singly or as groups, including well-known groups.
- You can assign a user as a member of more than one role.
- You can add a group to a role, which grants to all users who are members of that group all of the privileges associated with the role.
- You cannot assign privileges directly to users or groups.

Note

When OneFS is first installed, only users with root- or admin-level can log in and assign users to roles.

Built-in roles

Built-in roles include privileges to perform a set of administrative functions.

The following tables describe each of the built-in roles from most powerful to least powerful. The tables include the privileges and read/write access levels, if applicable, that are assigned to each role. You can assign users and groups to built-in roles and to roles that you create.

Table 3 SecurityAdmin role

Description	Privileges	Read/write access
Administer security configuration on the cluster, including authentication providers, local users and groups, and role membership.	ISI_PRIV_LOGIN_CONSOLE	N/A
	ISI_PRIV_LOGIN_PAPI	N/A
	ISI_PRIV_LOGIN_SSH	N/A
	ISI_PRIV_AUTH	Read/write
	ISI_PRIV_ROLE	Read/write

Table 4 SystemAdmin role

Description	Privileges	Read/write access
Administer all aspects of cluster configuration that are not specifically handled by the SecurityAdmin role.	ISI_PRIV_LOGIN_CONSOLE	N/A
	ISI_PRIV_LOGIN_PAPI	N/A
	ISI_PRIV_LOGIN_SSH	N/A
	ISI_PRIV_SYS_SHUTDOWN	N/A
	ISI_PRIV_SYS_SUPPORT	N/A
	ISI_PRIV_SYS_TIME	N/A
	ISI_PRIV_ANTIVIRUS	Read/write
	ISI_PRIV_AUDIT	Read/write
	ISI_PRIV_CLUSTER	Read/write

Table 4 SystemAdmin role (continued)

Description	Privileges	Read/write access
	ISI_PRIV_DEVICES	Read/write
	ISI_PRIV_EVENT	Read/write
	ISI_PRIV_FTP	Read/write
	ISI_PRIV_HDFS	Read/write
	ISI_PRIV_HTTP	Read/write
	ISI_PRIV_ISCSI	Read/write
	ISI_PRIV_JOB_ENGINE	Read/write
	ISI_PRIV_LICENSE	Read/write
	ISI_PRIV_NDMP	Read/write
	ISI_PRIV_NETWORK	Read/write
	ISI_PRIV_NFS	Read/write
	ISI_PRIV_NTP	Read/write
	ISI_PRIV_QUOTA	Read/write
	ISI_PRIV_REMOTE_SUPPORT	Read/write
	ISI_PRIV_SMARTPOOLS	Read/write
	ISI_PRIV_SMB	Read/write
	ISI_PRIV_SNAPSHOT	Read/write
	ISI_PRIV_STATISTICS	Read/write
	ISI_PRIV_SYNCIQ	Read/write
	ISI_PRIV_VCENTER	Read/write
ISI_PRIV_WORM	Read/write	
ISI_PRIV_NS_TRAVERSE	N/A	
ISI_PRIV_NS_IFS_ACCESS	N/A	

Table 5 AuditAdmin role

Description	Privileges	Read/write access
View all system configuration settings.	ISI_PRIV_LOGIN_CONSOLE	N/A
	ISI_PRIV_LOGIN_PAPI	N/A
	ISI_PRIV_LOGIN_SSH	N/A
	ISI_PRIV_ANTIVIRUS	Read-only
	ISI_PRIV_AUDIT	Read-only
	ISI_PRIV_CLUSTER	Read-only

Table 5 AuditAdmin role (continued)

Description	Privileges	Read/write access
	ISI_PRIV_DEVICES	Read-only
	ISI_PRIV_EVENT	Read-only
	ISI_PRIV_FTP	Read-only
	ISI_PRIV_HDFS	Read-only
	ISI_PRIV_HTTP	Read-only
	ISI_PRIV_ISCSI	Read-only
	ISI_PRIV_JOB_ENGINE	Read-only
	ISI_PRIV_LICENSE	Read-only
	SI_PRIV_NDMP	Read-only
	ISI_PRIV_NETWORK	Read-only
	ISI_PRIV_NFS	Read-only
	ISI_PRIV_NTP	Read-only
	ISI_PRIV_QUOTA	Read-only
	ISI_PRIV_REMOTE_SUPPORT	Read-only
	ISI_PRIV_SMARTPOOLS	Read-only
	ISI_PRIV_SMB	Read-only
	ISI_PRIV_SNAPSHOT	Read-only
	ISI_PRIV_STATISTICS	Read-only
	ISI_PRIV_SYNCIQ	Read-only
	ISI_PRIV_VCENTER	Read-only
	ISI_PRIV_WORM	Read-only

Table 6 VMwareAdmin role

Description	Privileges	Read/write access
Administers remotely all aspects of storage needed by VMware vCenter.	ISI_PRIV_LOGIN_PAPI	N/A
	ISI_PRIV_ISCSI	Read/write
	ISI_PRIV_NETWORK	Read/write
	ISI_PRIV_SMARTPOOLS	Read/write
	ISI_PRIV_SNAPSHOT	Read/write
	ISI_PRIV_SYNCIQ	Read/write
	ISI_PRIV_VCENTER	Read/write
	ISI_PRIV_NS_TRAVERSE	N/A

Table 6 VMwareAdmin role (continued)

Description	Privileges	Read/write access
	ISI_PRIV_NS_IFS_ACCESS	N/A

Table 7 BackupAdmin role

Description	Privileges	Read/write access
Allows backup and restore of files from <code>/ifs</code>	ISI_PRIV_IFS_BACKUP	Read-only
	ISI_PRIV_IFS_RESTORE	Read-only

Custom roles

Custom roles supplement built-in roles.

You can create custom roles and assign privileges mapped to administrative areas in your EMC Isilon cluster environment. For example, you can create separate administrator roles for security, auditing, storage provisioning, and backup.

You can designate certain privileges as read-only or read/write when adding the privilege to a role. You can modify this option at any time.

You can add or remove privileges as user responsibilities grow and change.

Privileges

Privileges permit users to complete tasks on an EMC Isilon cluster.

Privileges are associated with an area of cluster administration such as Job Engine, SMB, or statistics.

Privileges have one of two forms:

Action

Allows a user to perform a specific action on a cluster. For example, the `ISI_PRIV_LOGIN_SSH` privilege allows a user to log in to a cluster through an SSH client.

Read/Write

Allows a user to view or modify a configuration subsystem such as statistics, snapshots, or quotas. For example, the `ISI_PRIV_SNAPSHOT` privilege allows an administrator to create and delete snapshots and snapshot schedules. A read/write privilege can grant either read-only or read/write access. Read-only access allows a user to view configuration settings; read/write access allows a user to view and modify configuration settings.

Privileges are granted to the user on login to a cluster through the OneFS API, the web administration interface, SSH, or a console session. A token is generated for the user, which includes a list of all privileges granted to the user. Each URI, web-administration interface page, and command requires a specific privilege to view or modify the information available through any of these interfaces.

Note

Privileges are not granted to users that do not connect to the System Zone during login or to users that connect through the deprecated Telnet service, even if they are members of a role.

OneFS privileges

Privileges in OneFS are assigned through role membership; privileges cannot be assigned directly to users and groups.

Table 8 Login privileges

OneFS privilege	User right	Privilege type
ISI_PRIV_LOGIN_CONSOLE	Log in from the console	Action
ISI_PRIV_LOGIN_PAPI	Log in to the Platform API and the web administration interface	Action
ISI_PRIV_LOGIN_SSH	Log in through SSH	Action

Table 9 System privileges

OneFS privilege	User right	Privilege type
ISI_PRIV_SYS_SHUTDOWN	Shut down the system	Action
ISI_PRIV_SYS_SUPPORT	Run cluster diagnostic tools	Action
ISI_PRIV_SYS_TIME	Change the system time	Action

Table 10 Security privileges

OneFS privilege	User right	Privilege type
ISI_PRIV_AUTH	Configure external authentication providers	Read/write
ISI_PRIV_ROLE	Create new roles and assign privileges	Read/write

Table 11 Configuration privileges

OneFS privilege	User right	Privilege type
ISI_PRIV_ANTIVIRUS	Configure antivirus scanning	Read/write
IS_PRIV_AUDIT	Configure audit capabilities	Read/write
ISI_PRIV_CLUSTER	Configure cluster identity and general settings	Read/write

Table 11 Configuration privileges (continued)

OneFS privilege	User right	Privilege type
ISI_PRIV_DEVICES	Create new roles and assign privileges	Read/write
ISI_PRIV_EVENT	View and modify system events	Read/write
ISI_PRIV_FTP	Configure FTP server	Read/write
ISI_PRIV_HDFS	Configure HDFS server	Read/write
ISI_PRIV_HTTP	Configure HTTP server	Read/write
ISI_PRIV_ISCSI	Configure iSCSI server	Read/write
ISI_PRIV_JOB_ENGINE	Schedule cluster-wide jobs	Read/write
ISI_PRIV_LICENSE	Activate OneFS software licenses	Read/write
ISI_PRIV_NDMP	Configure NDMP server	Read/write
ISI_PRIV_NETWORK	Configure network interfaces	Read/write
ISI_PRIV_NFS	Configure the NFS server	Read/write
ISI_PRIV_NTP	Configure NTP	Read/write
ISI_PRIV_QUOTA	Configure file system quotas	Read/write
ISI_PRIV_REMOTE_SUPPORT	Configure remote support	Read/write
ISI_PRIV_SMARTPOOLS	Configure storage pools	Read/write
ISI_PRIV_SMB	Configure the SMB server	Read/write
ISI_PRIV_SNAPSHOT	Schedule, take, and view snapshots	Read/write
ISI_PRIV_SNMP	Configure SNMP server	Read/write
ISI_PRIV_STATISTICS	View file system performance statistics	Read/write
ISI_PRIV_SYNCIQ	Configure SyncIQ	Read/write
ISI_PRIV_VCENTER	Configure VMware for vCenter	Read/write
ISI_PRIV_WORM	Configure SmartLock directories	Read/write

Table 12 Platform API-only privileges

OneFS privilege	User right	Privilege type
ISI_PRIV_EVENT	View and modify system events	Read/write
ISI_PRIV_LICENSE	Activate OneFS software licenses	Read/write
ISI_PRIV_STATISTICS	View file system performance statistics	Read/write

Table 13 File access privileges

OneFS privilege	User right	Privilege type
ISI_PRIV_IFS_BACKUP	Back up files from <code>/ifs</code> . <hr/> Note This privilege circumvents traditional file access checks, such as mode bits or NTFS ACLs.	Action
ISI_PRIV_IFS_RESTORE	Restore files from <code>/ifs</code> . <hr/> Note This privilege circumvents traditional file access checks, such as mode bits or NTFS ACLs.	Action

Command-line interface privileges

You can perform most tasks granted by a privilege through the command-line interface.

Some OneFS commands require root access, but if you do not have root access, you can perform most of the commands associated with a privilege through the `sudo` program. The system automatically generates a `sudoers` file of users based on existing roles. Prefixing a command with `sudo` allows you to run commands that require root access. For example, if you do not have root access, the following command fails:

```
isi alert list
```

If you are on the `sudoers` list because you are a member of a role that has the `ISI_PRIV_EVENT` privilege, the following command succeeds:

```
sudo isi alert list
```

The following tables list all One FS commands available, the associated privilege or root-access requirement, and whether `sudo` is required to run the command.

Note

If you are running in compliance mode, additional sudo commands are available.

Table 14 Privileges sorted by CLI command

isi command	Privilege	Requires sudo
isi alert	ISI_PRIV_EVENT	x
isi audit	ISI_PRIV_AUDIT	
isi auth - excluding isi auth role	ISI_PRIV_AUTH	
isi auth role	ISI_PRIV_ROLE	
isi avscan	ISI_PRIV_ANTIVIRUS	x
isi batterystatus	ISI_PRIV_STATISTICS	x
isi config	root	
isi dedupe - excluding isi dedupe stats	ISI_PRIV_JOB_ENGINE	
isi dedupe stats	ISI_PRIV_STATISTICS	
isi devices	ISI_PRIV_DEVICES	x
isi drivefirmware	root	
isi domain	root	
isi email	ISI_PRIV_CLUSTER	x
isi events	ISI_PRIV_EVENT	x
isi exttools	root	
isi fc	root	
isi filepool	ISI_PRIV_SMARTPOOLS	
isi firmware	root	
isi ftp	ISI_PRIV_FTP	x
isi get	root	
isi hdfs	ISI_PRIV_HDFS	
isi iscsi	ISI_PRIV_ISCSI	x
isi job	ISI_PRIV_JOB_ENGINE	
isi license	ISI_PRIV_LICENSE	x
isi lun	ISI_PRIV_ISCSI	x
isi ndmp	ISI_PRIV_NDMP	x
isi networks	ISI_PRIV_NETWORK	x
isi nfs	ISI_PRIV_NFS	
isi perfstat	ISI_PRIV_STATISTICS	x

Table 14 Privileges sorted by CLI command (continued)

isi command	Privilege	Requires sudo
isi pkg	root	
isi quota	ISI_PRIV_QUOTA	
isi readonly	root	
isi remotesupport	ISI_PRIV_REMOTE_SUPPORT	
isi servicelight	ISI_PRIV_DEVICES	x
isi services	root	
isi set	root	
isi smb	ISI_PRIV_SMB	
isi snapshot	ISI_PRIV_SNAPSHOT	
isi snmp	ISI_PRIV_SNMP	x
isi stat	ISI_PRIV_STATISTICS	x
isi statistics	ISI_PRIV_STATISTICS	x
isi status	ISI_PRIV_STATISTICS	x
isi storagepool	ISI_PRIV_SMARTPOOLS	
isi sync	ISI_PRIV_SYNCIQ	
isi tape	ISI_PRIV_NDMP	x
isi target	ISI_PRIV_ISCSI	x
isi update	root	
isi version	ISI_PRIV_CLUSTER	x
isi worm	ISI_PRIV_WORM	
isi zone	ISI_PRIV_AUTH	

Table 15 CLI commands sorted by privilege

Privilege	isi commands	Requires sudo
ISI_PRIV_ANTIVIRUS	<ul style="list-style-type: none"> isi avscan 	x
ISI_PRIV_AUDIT	<ul style="list-style-type: none"> isi audit 	
ISI_PRIV_AUTH	<ul style="list-style-type: none"> isi auth - excluding isi auth role isi zone 	
ISI_PRIV_IFS_BACKUP	N/A	N/A
ISI_PRIV_CLUSTER	<ul style="list-style-type: none"> isi email isi version 	x

Table 15 CLI commands sorted by privilege (continued)

Privilege	isi commands	Requires sudo
ISI_PRIV_DEVICES	<ul style="list-style-type: none"> isi devices isi servicelight 	x
ISI_PRIV_EVENT	<ul style="list-style-type: none"> isi alert isi events 	x
ISI_PRIV_FTP	<ul style="list-style-type: none"> isi ftp 	x
ISI_PRIV_HDFS	<ul style="list-style-type: none"> isi hdfs 	
ISI_PRIV_ISCSI	<ul style="list-style-type: none"> isi iscsi isi lun isi target 	x
ISI_PRIV_JOB_ENGINE	<ul style="list-style-type: none"> isi job isi dedupe - excluding isi dedupe stats 	
ISI_PRIV_LICENSE	<ul style="list-style-type: none"> isi license 	x
ISI_PRIV_NDMP	<ul style="list-style-type: none"> isi ndmp isi tape 	x
ISI_PRIV_NETWORK	<ul style="list-style-type: none"> isi networks 	x
ISI_PRIV_NFS	<ul style="list-style-type: none"> isi nfs 	
ISI_PRIV_QUOTA	<ul style="list-style-type: none"> isi quota 	
ISI_PRIV_ROLE	<ul style="list-style-type: none"> isi auth role 	
ISI_PRIV_REMOTE_SUPPORT	<ul style="list-style-type: none"> isi remotesupport 	
ISI_PRIV_IFS_RESTORE	N/A	N/A
ISI_PRIV_SMARTPOOLS	<ul style="list-style-type: none"> isi filepool isi storagepool 	
ISI_PRIV_SMB	<ul style="list-style-type: none"> isi smb 	
ISI_PRIV_SNAPSHOT	<ul style="list-style-type: none"> isi snapshot 	
ISI_PRIV_SNMP	<ul style="list-style-type: none"> isi snmp 	x
ISI_PRIV_STATISTICS	<ul style="list-style-type: none"> isi batterystatus isi dedupe stats isi perfstat isi stat 	x

Table 15 CLI commands sorted by privilege (continued)

Privilege	isi commands	Requires sudo
	<ul style="list-style-type: none"> isi statistics isi status 	
ISI_PRIV_SYNCIQ	<ul style="list-style-type: none"> isi sync 	
ISI_PRIV_WORM	<ul style="list-style-type: none"> isi worm 	
root	<ul style="list-style-type: none"> isi config isi domain isi drivefirmware isi exttools isi fc isi firmware isi get isi pkg isi readonly isi services isi set isi update 	

Data backup and restore privileges

You can assign privileges to a user that are explicitly for cluster data backup and restore actions.

Two privileges allow a user to backup and restore cluster data over supported client-side protocols: ISI_PRIV_IFS_BACKUP and ISI_PRIV_IFS_RESTORE.

CAUTION

These privileges circumvent traditional file access checks, such as mode bits or NTFS ACLs.

Most cluster privileges allow changes to cluster configuration in some manner. The backup and restore privileges allow access to cluster data from the System zone, the traversing of all directories, and reading of all file data and metadata regardless of file permissions.

Users assigned these privileges use the protocol as a backup protocol to another machine without generating access-denied errors and without connecting as the root user. These two privileges are supported over the following client-side protocols:

- SMB
- RAN API
- FTP
- SSH

Over SMB, the `ISI_PRIV_IFS_BACKUP` and `ISI_PRIV_IFS_RESTORE` privileges emulate the Windows privileges `SE_BACKUP_NAME` and `SE_RESTORE_NAME`. The emulation means that normal file-open procedures are protected by file system permissions. To enable the backup and restore privileges over the SMB protocol, you must open files with the `FILE_OPEN_FOR_BACKUP_INTENT` option, which occurs automatically through Windows backup software such as Robocopy. Application of the option is not automatic when files are opened through general file browsing software such as Windows File Explorer.

Both `ISI_PRIV_IFS_BACKUP` and `ISI_PRIV_IFS_RESTORE` privileges primarily support Windows backup tools such as Robocopy. A user must be a member of the BackupAdmin built-in role to access all Robocopy features, which includes copying file DACL and SACL metadata.

User permissions utility

You can view expected user or group permissions to a given file or directory with the expected user permissions utility.

The command-line interface expected user permissions utility provides quick discovery of user and group permissions, displaying ACL or mode bits permissions. The utility does not display privileges or SMB share permissions, however.

Note

You must be a member of a role that has `ISI_PRIV_LOGIN_SSH` and `ISI_PRIV_AUTH` privileges to run this utility.

For information about viewing group or user permissions, see the View expected user permissions topic.

Authentication

OneFS supports local and remote authentication providers to verify that users attempting to access an EMC Isilon cluster are who they claim to be. Anonymous access, which does not require authentication, is supported for protocols that allow it.

OneFS supports concurrent multiple authentication provider types, which are analogous to directory services. For example, OneFS is often configured to authenticate Windows clients with Active Directory and to authenticate UNIX clients with LDAP. You can also configure NIS, designed by Sun Microsystems, to authenticate users and groups when they access a cluster.

Note

OneFS is RFC 2307-compliant.

Supported authentication providers

You can configure local and remote authentication providers to authenticate or deny user access to an EMC Isilon cluster.

The following table compares features that are available with each of the authentication providers that OneFS supports. In the following table, an x indicates that a feature is fully supported by a provider; an asterisk (*) indicates that additional configuration or support from another provider is required.

Authentication provider	NTLM	Kerberos	User/group management	Netgroups	UNIX properties (RFC 2307)	Windows properties
Active Directory	x	x			*	x
LDAP	*	x		x	x	*
NIS				x	x	
Local	x		x		x	x
File	x			x	x	
MIT Kerberos		x		*	*	*

Authentication provider features

You can configure authentication providers for your environment.

Authentication providers support a mix of the features described in the following table.

Feature	Description
Authentication	All authentication providers support plain-text authentication. You can configure some providers to support NTLM or Kerberos authentication also.
Users and groups	OneFS provides the ability to manage users and groups directly on the cluster.
Netgroups	Specific to NFS, netgroups restrict access to NFS exports.
UNIX-centric user and group properties	Login shell, home directory, UID, and GID. Missing information is supplemented by configuration templates or additional authentication providers.
Windows-centric user and group properties	NetBIOS domain and SID. Missing information is supplemented by configuration templates.

Kerberos authentication

Kerberos is a network authentication provider that negotiates encryption tickets for securing a connection. OneFS supports Active Directory Kerberos and MIT Kerberos authentication providers on an EMC Isilon cluster. If you configure an Active Directory provider, Kerberos authentication is provided automatically. MIT Kerberos works independently of Active Directory.

For MIT Kerberos authentication, you define an administrative domain known as a realm. Within this realm, an authentication server has the authority to authenticate a user, host, or service. You can optionally define a Kerberos domain to allow additional domain extensions to be associated with a realm.

The authentication server in a Kerberos environment is called the Key Distribution Center (KDC) and distributes encrypted tickets. When a user authenticates with an MIT Kerberos provider within a realm, an encrypted ticket with the user's service principal name (SPN) is created and validated to securely pass the user's identification for the requested service.

You can include an MIT Kerberos provider in specific access zones for authentication. Each access zone may include at most one MIT Kerberos provider. You can discontinue authentication through an MIT Kerberos provider by removing the provider from all the referenced access zones.

Keytabs and SPNs overview

A Key Distribution Center (KDC) is an authentication server that stores accounts and keytabs for users connecting to a network service within an EMC Isilon cluster. A keytab is a key table that stores keys to validate and encrypt Kerberos tickets.

One of the fields in a keytab entry is a service principal name (SPN). An SPN identifies a unique service instance within a cluster. Each SPN is associated with a specific key in the KDC. Users can use the SPN and its associated keys to obtain Kerberos tickets that enable access to various services on the cluster. A member of the SecurityAdmin role can create new keys for the SPNs and modify them later as necessary. An SPN for a service typically appears as `<service>/<fqdn>@<realm>`.

Note

SPNs must match the SmartConnect zone name and the FQDN hostname of the cluster. If the SmartConnect zone settings are changed, you must update the SPNs on the cluster to match the changes.

MIT Kerberos protocol support

MIT Kerberos supports certain standard network communication protocols such as HTTP, HDFS, and NFS. MIT Kerberos does not support SMB, SSH, and FTP protocols.

For the NFS protocol support, MIT Kerberos must be enabled for an export and also a Kerberos provider must be included within the access zone.

LDAP

The Lightweight Directory Access Protocol (LDAP) is a networking protocol that enables you to define, query, and modify directory services and resources.

OneFS can authenticate users and groups against an LDAP repository in order to grant them access to the cluster. OneFS supports Kerberos authentication for an LDAP provider.

The LDAP service supports the following features:

- Users, groups, and netgroups.
- Configurable LDAP schemas. For example, the `ldapsam` schema allows NTLM authentication over the SMB protocol for users with Windows-like attributes.
- Simple bind authentication, with and without SSL.
- Redundancy and load balancing across servers with identical directory data.
- Multiple LDAP provider instances for accessing servers with different user data.
- Encrypted passwords.

Active Directory

The Active Directory directory service is a Microsoft implementation of Lightweight Directory Access Protocol (LDAP), Kerberos, and DNS technologies that can store information about network resources. Active Directory can serve many functions, but the primary reason for joining the cluster to an Active Directory domain is to perform user and group authentication.

When the cluster joins an Active Directory domain, a single Active Directory machine account is created. The machine account establishes a trust relationship with the domain

and enables the cluster to authenticate and authorize users in the Active Directory forest. By default, the machine account is named the same as the cluster. If the cluster name is more than 15 characters long, the name is hashed and displayed after joining the domain.

Note

If you configure an Active Directory provider, Kerberos authentication is provided automatically.

Whenever possible, observe the following guidelines when you configure Active Directory providers on a cluster:

- Configure a single Active Directory instance if all domains have a trust relationship.
- Configure multiple Active Directory instances only to grant access to multiple sets of mutually-untrusted domains.

NIS

The Network Information Service (NIS) provides authentication and identity uniformity across local area networks. OneFS includes an NIS authentication provider that enables you to integrate the cluster with your NIS infrastructure.

NIS, designed by Sun Microsystems, can authenticate users and groups when they access the cluster. The NIS provider exposes the `passwd`, `group`, and `netgroup` maps from an NIS server. Hostname lookups are also supported. You can specify multiple servers for redundancy and load balancing.

Note

NIS is different from NIS+, which OneFS does not support.

File provider

A file provider enables you to supply an authoritative third-party source of user and group information to an EMC Isilon cluster. A third-party source is useful in UNIX and Linux environments that synchronize `/etc/passwd`, `/etc/group`, and `etc/netgroup` files across multiple servers.

Standard BSD `/etc/spwd.db` and `/etc/group` database files serve as the file provider backing store on a cluster. You generate the `spwd.db` file by running the `pwd_mkdb` command in the OneFS command-line interface (CLI). You can script updates to the database files.

On an Isilon cluster, a file provider hashes passwords with `libcrypt`. For the best security, we recommend that you use the Modular Crypt Format in the source `/etc/passwd` file to determine the hashing algorithm. OneFS supports the following algorithms for the Modular Crypt Format:

- MD5
- NT-Hash
- SHA-256
- SHA-512

For information about other available password formats, run the `man 3 crypt` command in the CLI to view the `crypt` man pages.

Note

The built-in System file provider includes services to list, manage, and authenticate against system accounts such as root, admin, and nobody. We recommended that you do not modify the System file provider.

Local provider

The local provider provides authentication and lookup facilities for user accounts added by an administrator.

Local authentication is useful when Active Directory, LDAP, or NIS directory services are not configured or when a specific user or application needs access to the cluster. Local groups can include built-in groups and Active Directory groups as members.

In addition to configuring network-based authentication sources, you can manage local users and groups by configuring a local password policy for each node in the cluster. OneFS settings specify password complexity, password age and re-use, and password-attempt lockout policies.

Data access control

You can configure an EMC Isilon cluster so that both UNIX and Windows users have access to content over NFS and SMB, regardless of the protocol that stored the data.

The OneFS operating system supports multiprotocol data access over Server Message Block (SMB) and Network File System (NFS) with a unified security model. For NFS, the default export on a cluster, `/ifs`, enables Linux and UNIX clients to remotely mount any subdirectory, including subdirectories created by Windows users. Linux and UNIX clients also can mount ACL-protected subdirectories created by a OneFS administrator.

Conversely, for SMB the default file share on a cluster, `/ifs`, provides Windows users access to file system resources over the network that includes resources stored by UNIX and Linux systems. The same access model applies to directories and files.

By default, OneFS maintains the same file permissions regardless of the client's operating system, the user's identity management system, or the file sharing protocol. When OneFS must transform a file's permissions from ACLs to mode bits or vice versa, it merges the permissions into an optimal representation that uniquely balances user expectations and file security.

Authorization

OneFS supports two types of authorization data on a file: Windows-style access control lists (ACLs) and POSIX mode bits (UNIX permissions). Authorization type is based on the ACL policies that are set and on the file-creation method.

Access to a file or directory is governed by either a Windows access control list (ACL) or UNIX mode bits. Regardless of the security model, OneFS enforces access rights consistently across access protocols. A user is granted or denied the same rights to a file when using SMB for Windows file sharing as when using NFS for UNIX file sharing.

An EMC Isilon cluster includes global policy settings that enable you to customize the default ACL and UNIX permissions to best support your environment. Generally, files that are created over SMB or in a directory that has an ACL receive an ACL; otherwise, OneFS relies on the POSIX mode bits that define UNIX permissions. In either case, the owner is represented by a UNIX identifier (UID or GID) or by its Windows identifier (SID). The

primary group is represented by a GID or SID. Although mode bits are present when a file has an ACL, the mode bits are provided for only protocol compatibility, not for access checks.

Note

Although you can configure ACL policies to optimize a cluster for UNIX or Windows, you should do so only if you understand how ACL and UNIX permissions interact.

The OneFS file system installs with UNIX permissions as the default. By using Windows Explorer or OneFS administrative tools, you can give a file or directory an ACL. In addition to Windows domain users and groups, ACLs in OneFS can include local, NIS, and LDAP users and groups. After you give a file an ACL, OneFS stops enforcing the file's mode bits, which remain only as an estimate of the effective permissions.

SMB

You can configure SMB shares to provide Windows clients network access to file system resources on the cluster. You can grant permissions to users and groups to carry out operations such as reading, writing, and setting access permissions on SMB shares.

ACLs

In Windows environments, file and directory permissions, referred to as access rights, are defined in access control lists (ACLs). Although ACLs are more complex than mode bits, ACLs can express much more granular sets of access rules. OneFS checks the ACL processing rules commonly associated with Windows ACLs.

A Windows ACL contains zero or more access control entries (ACEs), each of which represents the security identifier (SID) of a user or a group as a trustee. In OneFS, an ACL can contain ACEs with a UID, GID, or SID as the trustee. Each ACE contains a set of rights that allow or deny access to a file or folder. An ACE can optionally contain an inheritance flag to specify whether the ACE should be inherited by child folders and files.

Note

Instead of the standard three permissions available for mode bits, ACLs have 32 bits of fine-grained access rights. Of these, the upper 16 bits are general and apply to all object types. The lower 16 bits vary between files and directories but are defined in a way that allows most applications to apply the same bits for files and directories.

Rights grant or deny access for a given trustee. You can block user access explicitly through a deny ACE or implicitly by ensuring that a user does not directly, or indirectly through a group, appear in an ACE that grants the right.

NFS

You can configure NFS exports to provide UNIX clients network access to file system resources on the cluster.

UNIX permissions

In a UNIX environment, file and directory access is controlled by POSIX mode bits, which grant read, write, or execute permissions to the owning user, the owning group, and everyone else.

OneFS supports the standard UNIX tools for viewing and changing permissions, `ls`, `chmod`, and `chown`. For more information, run the `man ls`, `man chmod`, and `man chown` commands.

All files contain 16 permission bits, which provide information about the file or directory type and the permissions. The lower 9 bits are grouped as three 3-bit sets, called triples, which contain the read, write, and execute (rwx) permissions for each class of users—owner, group, and other. You can set permissions flags to grant permissions to each of these classes.

Unless the user is root, OneFS checks the class to determine whether to grant or deny access to the file. The classes are not cumulative: The first class matched is applied. It is therefore common to grant permissions in decreasing order.

Mixed-permission environments

When a file operation requests an object's authorization data, for example, with the `ls -l` command over NFS or with the **Security** tab of the **Properties** dialog box in Windows Explorer over SMB, OneFS attempts to provide that data in the requested format. In an environment that mixes UNIX and Windows systems, some translation may be required when performing create file, set security, get security, or access operations.

NFS access of Windows-created files

If a file contains an owning user or group that is a SID, the system attempts to map it to a corresponding UID or GID before returning it to the caller.

In UNIX, authorization data is retrieved by calling `stat(2)` on a file and examining the owner, group, and mode bits. Over NFSv3, the `GETATTR` command functions similarly. The system approximates the mode bits and sets them on the file whenever its ACL changes. Mode bit approximations need to be retrieved only to service these calls.

Note

SID-to-UID and SID-to-GID mappings are cached in both the OneFS ID mapper and the `stat` cache. If a mapping has recently changed, the file might report inaccurate information until the file is updated or the cache is flushed.

SMB access of UNIX-created files

No UID-to-SID or GID-to-SID mappings are performed when creating an ACL for a file; all UIDs and GIDs are converted to SIDs or principals when the ACL is returned.

OneFS initiates a two-step process for returning a security descriptor, which contains SIDs for the owner and primary group of an object:

1. The current security descriptor is retrieved from the file. If the file does not have a discretionary access control list (DACL), a synthetic ACL is constructed from the file's lower 9 mode bits, which are separated into three sets of permission triples—one each for owner, group, and everyone. For details about mode bits, see the UNIX permissions topic.
2. Two access control entries (ACEs) are created for each triple: the allow ACE contains the corresponding rights that are granted according to the permissions; the deny ACE

contains the corresponding rights that are denied. In both cases, the trustee of the ACE corresponds to the file owner, group, or everyone. After all of the ACEs are generated, any that are not needed are removed before the synthetic ACL is returned.

Managing roles

You can view, add, or remove members of any role. Except for built-in roles, whose privileges you cannot modify, you can add or remove OneFS privileges on a role-by-role basis.

Note

Roles take both users and groups as members. If a group is added to a role, all users who are members of that group are assigned the privileges associated with the role. Similarly, members of multiple roles are assigned the combined privileges of each role.

View roles

You can view information about built-in and custom roles.

Procedure

1. Run one of the following commands to view roles.

- To view a basic list of all roles on the cluster, run the following command:

```
isi auth roles list
```

- To view detailed information about each role on the cluster, including member and privilege lists, run the following command:

```
isi auth roles list --verbose
```

- To view detailed information about a single role, run the following command, where *<role>* is the name of the role:

```
isi auth roles view <role>
```

View privileges

You can view user privileges.

This procedure must be performed through the command-line interface (CLI). You can view a list of your privileges or the privileges of another user using the following commands:

Procedure

1. Establish an SSH connection to any node in the cluster.
2. To view privileges, run one of the following commands.

- To view a list of all privileges, run the following command:

```
isi auth privileges --verbose
```

- To view a list of your privileges, run the following command:

```
isi auth id
```

- To view a list of privileges for another user, run the following command, where *<user>* is a placeholder for another user by name:

```
isi auth mapping token <user>
```

Create and modify a custom role

You can create an empty custom role and then add users and privileges to the role.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. Run the following command to create a role, where *<name>* is the name that you want to assign to the role and *<string>* specifies an optional description:

```
isi auth roles create <name> [--description <string>]
```

3. Run the following command to add a user to the role, where *<role>* is the name of the role and *<string>* is the name of the user:

```
isi auth roles modify <role> [--add-user <string>]
```

Note

You can also modify the list of users assigned to a built-in role.

4. Run the following command to add a privilege with read/write access to the role, where *<role>* is the name of the role and *<string>* is the name of the privilege:

```
isi auth roles modify <role> [--add-priv <string>]
```

5. Run the following command to add a privilege with read-only access to the role, where *<role>* is the name of the role and *<string>* is the name of the privilege:

```
isi auth roles modify <role> [--add-priv-ro <string>]
```

Delete a custom role

Deleting a role does not affect the privileges or users that are assigned to it. Built-in roles cannot be deleted.

Procedure

1. Run the following command to delete a custom role, where *<name>* is the name of the role that you want to delete:

```
isi auth roles delete <name>
```

Managing authentication providers

You can configure one or more LDAP, Active Directory, NIS, file, and Kerberos providers. A local provider is created automatically when you create an access zone, which allows you to create a configuration for each access zone so it has its own list of local users that can authenticate to it. You also can create a password policy for each local provider to enforce password complexity.

Managing LDAP providers

You can view, configure, modify, and delete LDAP providers. You can discontinue authentication through an LDAP provider by removing it from all access zones that are using it.

Configure an LDAP provider

By default, when you configure an LDAP provider, it is automatically added to the System access zone.

Procedure

1. If the LDAP server allows anonymous queries, you can create an LDAP provider by running the `isi auth ldap create` command with the following parameters, where variables in angle brackets are placeholders for values specific to your environment:

```
isi auth ldap create <name> --base-dn=<base-distinguished-name> \
--server-uris=<uri>
```

The following command joins the user `test` to the LDAP server `test-ldap.example.com`:

```
isi auth ldap create test \
--base-dn="dc=test-ldap,dc=example,dc=com" \
--server-uris="ldap://test-ldap.example.com"
```

Note

You can specify multiple servers by repeating the `--server-uris` parameter with the URI value or with a comma-separated list, such as `--server-uris="ldap://a.example.com,ldap://b.example.com"`.

2. If the LDAP server does not allow anonymous queries, you can create an LDAP provider by running the `isi auth ldap create` command, where variables in angle brackets are placeholders for values specific to your environment:

```
isi auth ldap create <name> --bind-dn=<distinguished-name> \
--bind-password=<password> --server-uris=<uri>
```

The following command joins the LDAP server `test-ldap.example.com` and binds to user `test` in the organizational unit `users`:

```
isi auth ldap create test-ldap \
--bind-dn="cn=test,ou=users,dc=test-ldap,dc=example,dc=com" \
--bind-password="mypasswd" \
--server-uris="ldap://test-ldap.example.com"
```

Note

The bind DN must have the proper permissions set.

Modify an LDAP provider

You can modify any setting for an LDAP provider except its name. You must specify at least one server for the provider to be enabled.

Procedure

1. Run the following command to modify an LDAP provider, where *<provider-name>* is a placeholder for the name of the provider that you want to modify:

```
isi auth ldap modify <provider-name>
```

Delete an LDAP provider

When you delete an LDAP provider, it is removed from all access zones. As an alternative, you can stop using an LDAP provider by removing it from each access zone that contains it so that the provider remains available for future use.

For information about the parameters and options that are available for this procedure, run the `isi auth ldap delete --help` command.

Procedure

1. Run the following command to delete an LDAP provider, where *<name>* is a placeholder for the name of the LDAP provider that you want to delete.

```
isi auth ldap delete <name>
```

Managing Active Directory providers

You can view, configure, modify, and delete Active Directory providers. OneFS includes a Kerberos configuration file for Active Directory in addition to the global Kerberos configuration file, both of which you can configure through the command-line interface. You can discontinue authentication through an Active Directory provider by removing it from all access zones that are using it.

Configure an Active Directory provider

You can configure one or more Active Directory providers, each of which must be joined to a separate Active Directory domain. By default, when you configure an Active Directory provider, it is automatically added to the System access zone.

Note

Consider the following information when you configure an Active Directory provider:

- When you join Active Directory from OneFS, cluster time is updated from the Active Directory server, as long as an NTP server has not been configured for the cluster.
- If you migrate users to a new or different Active Directory domain, you must re-set the ACL domain information after you configure the new provider. You can reset the domain information with third-party tools, such as Microsoft SubInACL.

Procedure

1. Run the following command to configure an Active Directory provider, where *<name>* is a placeholder for the fully qualified Active Directory name and *<user>* is a placeholder for a user name with permission to join machines to the given domain.

```
isi auth ads create <name> <user>
```


Modify an Active Directory provider

You can modify the advanced settings for an Active Directory provider.

Procedure

1. Run the following command to modify an Active Directory provider, where *<provider-name>* is a placeholder for the name of the provider that you want to modify.

```
isi auth ads modify <provider-name>
```

Delete an Active Directory provider

When you delete an Active Directory provider, you disconnect the cluster from the Active Directory domain that is associated with the provider, disrupting service for users who are accessing it. After you leave an Active Directory domain, users can no longer access the domain from the cluster.

Procedure

1. Run the following command to Delete an Active Directory provider, where *<name>* is a placeholder for the Active Directory name that you want to delete.

```
isi auth ads delete <name>
```

Managing NIS providers

You can view, configure, and modify NIS providers or delete providers that are no longer needed. You can discontinue authentication through an NIS provider by removing it from all access zones that are using it.

Configure an NIS provider

You can configure multiple NIS providers, each with its own settings, and add them to one or more access zones. By default, when you configure an NIS provider, it is automatically added to the System access zone.

Procedure

1. Configure an NIS provider by running the `isi auth nis create` command with the following syntax:

```
isi auth nis create <name> --servers=<server> \
--nis-domain=<domain>
```

The following example joins the NIS server `nistest.example.com` to `nistest`:

```
isi auth nis create nistest --servers="nistest.example.com" \
--nis-domain="example.com"
```

Note

You can specify multiple servers by repeating the `--servers` parameter for each server or with a comma-separated list, such as `--server="a.example.com,b.example.com"`.

Modify an NIS provider

You can modify any setting for an NIS provider except its name. You must specify at least one server for the provider to be enabled.

Procedure

1. Run the following command to modify an NIS provider, where *<provider-name>* is a placeholder for provider that you want to modify.

```
isi auth nis modify <provider-name>
```

Delete an NIS provider

When you delete an NIS provider, it is removed from all access zones. As an alternative, you can stop using an NIS provider by removing it from each access zone that contains it, so that the provider remains available for future use.

Procedure

1. Run the following command to delete an NIS provider, where *<name>* is a placeholder for the name of the NIS provider that you want to delete.

```
isi auth nis delete <name>
```

Managing file providers

You can configure one or more file providers, each with its own combination of replacement files, for each access zone. Password database files, which are also called user database files, must be in binary format.

Each file provider pulls directly from up to three replacement database files: a group file that has the same format as `/etc/group`; a netgroups file; and a binary password file, `spwd.db`, which provides fast access to the data in a file that has the `/etc/master.passwd` format. You must copy the replacement files to the cluster and reference them by their directory path.

Note

If the replacement files are located outside the `/ifs` directory tree, you must distribute them manually to every node in the cluster. Changes that are made to the system provider's files are automatically distributed across the cluster.

Configure a file provider

You can specify replacement files for any combination of users, groups, and netgroups.

Procedure

1. Run the following command to configure a file provider, where *<name>* is your name for the file provider.

```
isi auth file create <name>
```

Generate a password file

Password database files, which are also called user database files, must be in binary format.

This procedure must be performed through the command-line interface (CLI). For command-usage guidelines, run the `man pwd_mkdb` command.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. Run the `pwd_mkdb <file>` command, where `<file>` is the location of the source password file.

Note

By default, the binary password file, `spwd.db`, is created in the `/etc` directory. You can override the location to store the `spwd.db` file by specifying the `-d` option with a different target directory.

The following command generates an `spwd.db` file in the `/etc` directory from a password file that is located at `/ifs/test.passwd`:

```
pwd_mkdb /ifs/test.passwd
```

The following command generates an `spwd.db` file in the `/ifs` directory from a password file that is located at `/ifs/test.passwd`:

```
pwd_mkdb -d /ifs /ifs/test.passwd
```

Modify a file provider

You can modify any setting for a file provider, including its name.

Note

Although you can rename a file provider, there are two caveats: you can rename a file provider through only the web administration interface and you cannot rename the System file provider.

Procedure

1. Run the following command to modify a file provider, where `<provider-name>` is a placeholder for the name that you supplied for the provider.

```
isi auth file modify <provider-name>
```

Delete a file provider

To stop using a file provider, you can clear all of its replacement file settings or you can permanently delete the provider.

Note

You cannot delete the System file provider.

Procedure

1. Run the following command to delete a file provider, where `<name>` is a placeholder for the name of the provider that you want to delete.

```
isi auth file delete <name>
```

Password file format

The file provider uses a binary password database file, `spwd.db`. You can generate a binary password file from a `master.passwd`-formatted file by running the `pwd_mkdb` command.

The `master.passwd` file contains ten colon-separated fields, as shown in the following example:

```
admin:*:10:10::0:0:Web UI Administrator:/ifs/home/admin:/bin/zsh
```

The fields are defined below in the order in which they appear in the file.

Note

UNIX systems often define the `passwd` format as a subset of these fields, omitting the Class, Change, and Expiry fields. To convert a file from `passwd` to `master.passwd` format, add `:0:0:` between the GID field and the Gecos field.

Username

The user name. This field is case-sensitive. OneFS does not limit the length; many applications truncate the name to 16 characters, however.

Password

The user's encrypted password. If authentication is not required for the user, you can substitute an asterisk (*) for a password. The asterisk character is guaranteed to not match any password.

UID

The UNIX user identifier. This value must be a number in the range 0-4294967294 that is not reserved or already assigned to a user. Compatibility issues occur if this value conflicts with an existing account's UID.

GID

The group identifier of the user's primary group. All users are a member of at least one group, which is used for access checks and can also be used when creating files.

Class

This field is not supported by OneFS and should be left empty.

Change

OneFS does not support changing the passwords of users in the file provider. This field is ignored.

Expiry

OneFS does not support the expiration of user accounts in the file provider. This field is ignored.

Gecos

This field can store a variety of information but is usually used to store the user's full name.

Home

The absolute path to the user's home directory, beginning at `/ifs`.

Shell

The absolute path to the user's shell. If this field is set to `/sbin/nologin`, the user is denied command-line access.

Group file format

The file provider uses a group file in the format of the `/etc/group` file that exists on most UNIX systems.

The `group` file consists of one or more lines containing four colon-separated fields, as shown in the following example:

```
admin:*:10:root,admin
```

The fields are defined below in the order in which they appear in the file.

Group name

The name of the group. This field is case-sensitive. Although OneFS does not limit the length of the group name, many applications truncate the name to 16 characters.

Password

This field is not supported by OneFS and should contain an asterisk (*).

GID

The UNIX group identifier. Valid values are any number in the range 0-4294967294 that is not reserved or already assigned to a group. Compatibility issues occur if this value conflicts with an existing group's GID.

Group members

A comma-delimited list of user names.

Netgroup file format

A netgroup file consists of one or more netgroups, each of which can contain members. Hosts, users, or domains, which are members of a netgroup, are specified in a member triple. A netgroup can also contain another netgroup.

Each entry in a `netgroup` file consists of the netgroup name, followed by a space-delimited set of member triples and nested netgroup names. If you specify a nested netgroup, it must be defined on a separate line in the file.

A member triple takes the following form:

```
(<host>, <user>, <domain>)
```

Where `<host>` is a placeholder for a machine name, `<user>` is a placeholder for a user name, and `<domain>` is a placeholder for a domain name. Any combination is valid except an empty triple: `(,,)`.

The following sample file contains two netgroups. The `rootgrp` netgroup contains four hosts: two hosts are defined in member triples and two hosts are contained in the nested `othergrp` netgroup, which is defined on the second line.

```
rootgrp (myserver, root, somedomain.com) (otherserver, root,
somedomain.com) othergrp
othergrp (other-win,, somedomain.com) (other-linux,, somedomain.com)
```

Note

A new line signifies a new netgroup. You can continue a long netgroup entry to the next line by typing a backslash character (\) in the right-most position of the first line.

Managing local users and groups

When you create an access zone, each zone includes a local provider that allows you to create and manage local users and groups. Although you can view the users and groups of any authentication provider, you can create, modify, and delete users and groups in the local provider only.

View a list of users and groups by provider

You can view users and groups by a provider type.

Procedure

1. Run the following command to view a list of users and groups for a specified provider, where *<provider-type>* is a placeholder for your provider-type string and *<provider-name>* is a placeholder for the name that you assigned the specific provider:

```
isi auth users list --provider="<provider-type>:<provider-name>"
```

2. To list users and groups for an LDAP provider type that is named Unix LDAP, run a command similar to the following example:

```
isi auth users list --provider="lsa-ldap-provider:Unix LDAP"
```

Create a local user

Each access zone includes a local provider that allows you to create and manage local users and groups. When creating a local user account, you can configure its name, password, home directory, UNIX user identifier (UID), UNIX login shell, and group memberships.

Procedure

1. Run the following command to create a local user, where *<name>* is your name for the user, *<provider-name>* specifies the provider for this user, and *<string>* is the password for this user.

```
isi auth users create <name> --provider="local:<provider-name>" \
--password="<string>"
```

Note

A user account is disabled if no password is specified. If you do not create a password when you create the user account, you can add a password later by running the `isi auth users modify` command, specifying the appropriate user by username, UID, or SID.

Create a local group

In the local provider of an access zone, you can create groups and assign members to them.

Procedure

1. Run the following command to create a local group, where *<name>* and *<provider-name>* are values that you provide to define the group.

```
isi auth groups create <name> --provider "local:<provider-name>"
```

Naming rules for local users and groups

Local user and group names must follow naming rules in order to ensure proper authentication and access to the EMC Isilon cluster.

You must adhere to the following naming rules when creating and modifying local users and groups:

- The maximum name length is 104 characters. It is recommended that names do not exceed 64 characters.
- Names cannot contain the following invalid characters:
"/ \ [] : ; | = , + * ? < >
- Names can contain any special character that is not in the list of invalid characters. It is recommended that names do not contain spaces.
- Names are not case sensitive.

Configure or modify a local password policy

You can configure and modify a local password policy for a local provider.

This procedure must be performed through the command-line interface (CLI).

Note

Separate password policies are configured for each access zone. Each access zone in the cluster contains a separate instance of the local provider, which allows each access zone to have its own list of local users who can authenticate. Password complexity is configured for each local provider, not for each user.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. (Optional) Run the following command to view the current password settings:

```
isi auth local view system
```

3. Run the `isi auth local modify system` command, choosing from the parameters described in Local password policy default settings.

The `--password-complexity` parameter must be specified for each setting.

```
isi auth local modify system --password-complexity=lowercase \  
--password-complexity=uppercase --password-complexity=numeric \  
--password-complexity=symbol
```

The following command configures a local password policy for a local provider:

```
isi auth local modify <provider-name> \
--min-password-length=20 \
--lockout-duration=20m \
--lockout-window=5m \
--lockout-threshold=5 \
--add-password-complexity=uppercase \
--add-password-complexity=numeric
```

Local password policy settings

You can configure local password policy settings and specify the default for each setting through the `isi auth local modify` command. Password complexity increases the number of possible passwords that an attacker must check before the correct password is guessed.

Setting	Description	Comments
<code>min-password-length</code>	Minimum password length in characters.	Long passwords are best. The minimum length should not be so long that users have a difficult time entering or remembering the password.
<code>password-complexity</code>	A list of cases that a new password must contain. By default, the list is empty.	You can specify as many as four cases. The following cases are valid: <ul style="list-style-type: none"> • uppercase • lowercase • numeric • symbol (excluding # and @)
<code>min-password-age</code>	The minimum password age. You can set this value using characters for units; for example, 4W for 4 weeks, 2d for 2 Days.	A minimum password age ensures that a user cannot enter a temporary password and then immediately change it to the previous password. Attempts to check or set a password before the time expires are denied.
<code>max-password-age</code>	The maximum password age. You can set this value using characters for units; for example, 4W for 4 weeks, 2d for 2 Days.	Attempts to login after a password expires forces a password change. If a password change dialog cannot be presented, the user is not allowed to login.
<code>password-history-length</code>	The number of historical passwords to keep. New passwords are checked against this list and rejected if the password is already present. The max history length is 24.	To avoid recycling of passwords, you can specify the number of previous passwords to remember. If a new password matches a remembered previous password, it is rejected.
<code>lockout-duration</code>	The length of time in seconds that an account is locked after a configurable	After an account is locked, it is unavailable from all sources until it is unlocked. OneFS provides two configurable options to avoid

Setting	Description	Comments
	number of bad passwords are entered.	administrator interaction for every locked account: <ul style="list-style-type: none"> Specify how much time must elapse before the account is unlocked. Automatically reset the incorrect-password counter after a specified time, in seconds.
lockout-threshold	The number of incorrect password attempts before an account is locked. A value of zero disables account lockout.	After an account is locked, it is unavailable from all sources until it is unlocked.
lockout-window	The time that elapses before the incorrect password attempts count is reset.	If the configured number of incorrect password attempts is reached, the account is locked and lockout-duration determines the length of time that the account is locked. A value of zero disables the window.

Modify a local user

You can modify any setting for a local user account except the user name.

Procedure

1. Run the following command to modify a local group, where *<name>* or *<gid>* or *<sid>* are placeholders for the user identifiers and *<provider-name>* is a placeholder for the name of the local provider associated with the user:

```
isi auth users modify (<name> or --gid <gid> or --sid <sid>) \
--provider "local:<provider-name>"
```

Modify a local group

You can add or remove members from a local group.

Procedure

1. Run the following command to modify a local group, where *<name>* or *<gid>* or *<sid>* are placeholders for the group identifiers and *<provider-name>* is a placeholder for the name of the local provider associated with the group:

```
isi auth groups modify (<name> or --gid <gid> or --sid <sid>) \
--provider "local:<provider-name>"
```

Delete a local user

A deleted user can no longer access the cluster through the command-line interface, web administration interface, or file access protocol. When you delete a local user account, its home directory remains in place.

Procedure

1. Run the following command to delete a local user, where *<uid>* and *<sid>* are placeholders for the UID and SID of the user that you want to delete, and *<provider-name>* is a placeholder for the local provider associated with the user.

```
isi auth users delete <name> --uid <uid> --sid <sid> \
--provider "local:<provider-name>"
```

Delete a local group

You can delete a local group even if members are assigned to it. Deleting a group does not affect the members of that group.

Procedure

1. Run the following command to delete a local group, where *<group>* is a placeholder for the name of the group that you want to delete:

```
isi auth groups delete <group>
```

Note

You can run the command with *<gid>* or *<sid>* instead of *<group>*.

Managing MIT Kerberos authentication

You can configure an MIT Kerberos provider for authentication without Active Directory. Configuring an MIT Kerberos provider involves creating an MIT Kerberos realm, creating a provider, and joining a predefined realm. Optionally, you can configure an MIT Kerberos domain for the provider. You can also update the encryption keys if there are any configuration changes to the Kerberos provider. You can include the provider in one or more access zones.

Managing MIT Kerberos realms

An MIT Kerberos realm is an administrative domain that defines the boundaries within which an authentication server has the authority to authenticate a user or service. You can create, view, edit, or delete a realm. As a best practice, specify a realm name using uppercase characters.

Create an MIT Kerberos realm

You can create an MIT Kerberos realm by defining a Key Distribution Center (KDC) and an administrative server.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to create an MIT Kerberos realm.

Procedure

1. Run the `isi auth krb5 realm create` command to create an MIT Kerberos realm.

For example, run the following command to create a realm by specifying a KDC and an administrative server:

```
isi auth krb5 realm create <realm> --kdc <kdc> --admin-server
<admin-server>
```

Modify an MIT Kerberos realm

You can modify an MIT Kerberos realm by modifying the Key Distribution Center (KDC), the domain (optional), and the administrative server settings for that realm.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to delete an MIT Kerberos provider.

Procedure

1. Run the `isi auth krb5 realm modify` command to modify an MIT Kerberos realm.

For example, run the following command to modify an MIT Kerberos realm by specifying an alternate KDC and an administrative server:

```
isi auth krb5 realm modify <realm> --is-default-realm true --kdc
<kdc> --admin-server <admin-server>
```

View an MIT Kerberos realm

You can view details related to the name, Key Distribution Centers (KDCs), and the administrative server associated with an MIT Kerberos realm.

Procedure

1. Run the `isi auth krb5 realm view` command to view details for an MIT Kerberos realm.

For example, run the following command to view the details for a realm:

```
isi auth krb5 realm view <realm>
```

Delete an MIT Kerberos realm

You can delete one or more MIT Kerberos realms and all the associated MIT Kerberos domains.

Before you begin

Kerberos realms are referenced by Kerberos providers. Before you can delete a realm for which you have created a provider, you must first delete that provider.

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to delete an MIT Kerberos realm.

Procedure

1. Run the `isi auth krb5 realm delete` command to delete an MIT Kerberos realm.

For example, run the following command to delete a realm:

```
isi auth krb5 realm delete <realm>
```

Managing MIT Kerberos providers

You can create view, delete, or modify an MIT Kerberos provider. You can also configure the Kerberos provider settings.

Creating an MIT Kerberos provider

You can create an MIT Kerberos provider by obtaining the credentials for accessing a cluster through the Key Distribution Center (KDC) of the Kerberos realm. This process is also known as joining a realm. Thus when you create a Kerberos provider you also join a realm that has been previously defined.

Depending on how OneFS manages your Kerberos environment, you can create a Kerberos provider through one of the following methods:

- Accessing the Kerberos administration server and creating keys for services on the OneFS cluster.
- Manually transferring the Kerberos key information in the form of keytabs.

Create an MIT Kerberos provider and join a realm with administrator credentials

You can create an MIT Kerberos provider and join an MIT Kerberos realm using the credentials authorized to access the Kerberos administration server. You can then create keys for the various services on the EMC Isilon cluster. This is the recommended method for creating a Kerberos provider and joining a Kerberos realm.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to access the Kerberos administration server.

Procedure

1. Run the following command to create a Kerberos provider and join a Kerberos realm, where `<realm>` is the name of the Kerberos realm which already exists or is created if it does not exist:

```
isi auth krb5 create <realm> <user> --<kdc>=<string>
```

The following example joins a user with admin credentials to the *cluster-name.company.com* realm:

```
isi auth krb5 create cluster-name.company.com aima/admin --  
kdc=<kdc-name>.domain.company.com
```

Create an MIT Kerberos provider and join a realm with a keytab file

You can create an MIT Kerberos provider and join an MIT Kerberos realm through a keytab file. Follow this method only if your Kerberos environment is managed by manually transferring the Kerberos key information through the keytab files.

Before you begin

Make sure that the following prerequisites are met:

- You must create and copy a keytab file to a node on the cluster.
- You must be a member of a role that has `ISI_PRIV_AUTH` privileges to access the Kerberos administration server.

Procedure

1. Run the following command to create a Kerberos provider and join a Kerberos realm, where *<realm>* is the placeholder for the realm name which should exist as a realm object already:

```
isi auth krb5 create <realm> --keytab-file=<string>
```

For example, run the following command to create a Kerberos provider and join the *cluster-name.company.com* realm using a keytab file:

```
isi auth krb5 create cluster-name.company.com --keytab-file=/tmp/krb5.keytab
```

View an MIT Kerberos provider

You can view the properties of an MIT Kerberos provider after creating it.

Procedure

1. Run the following command to view the properties of a Kerberos provider:

```
isi auth krb5 view <provider-name>
```

List the MIT Kerberos providers

You can list one or more MIT Kerberos providers and display the list in a specific format. You can also specify a limit for the number of providers to be listed.

Procedure

1. Run the `isi auth krb5 list` command to list one or more Kerberos providers.

For example, run the following command to list the first five Kerberos providers in a tabular format without any headers or footers:

```
isi auth krb5 list -l 5 --format table --no-header --no-footer
```

Delete an MIT Kerberos provider

You can delete an MIT Kerberos provider and remove it from all the referenced access zones. When you delete a provider, you also leave an MIT Kerberos realm.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to delete a Kerberos provider.

Procedure

1. Run the `isi auth krb5 delete` command as follows to delete a Kerberos provider.

```
isi auth krb5 delete <provider-name>
```

Configure MIT Kerberos provider settings

You can configure the settings of a Kerberos provider to allow the DNS records to locate the Key Distribution Center (KDC), Kerberos realms, and the authentication servers associated with a Kerberos realm. These settings are global to all Kerberos users across

all nodes, services, and zones. Some settings are applicable only to client-side Kerberos and are independent of the Kerberos provider.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to view or modify the settings of a Kerberos provider.

Procedure

1. Run the `isi auth settings krb5` command with the `view` or `modify` subcommand.
2. Specify the settings to modify.

Managing MIT Kerberos domains

You can optionally define MIT Kerberos domains to allow additional domain extensions to be associated with an MIT Kerberos realm. You can always specify a default domain for a realm.

You can create, modify, delete, and view an MIT Kerberos domain. A Kerberos domain name is a DNS suffix that you specify typically using lowercase characters.

Add an MIT Kerberos domain to a realm

You can optionally add an MIT Kerberos domain to an MIT Kerberos realm to enable additional Kerberos domain extensions to be associated with a Kerberos realm.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to associate a Kerberos domain with a Kerberos realm.

Procedure

1. Add a Kerberos domain by running the `isi auth krb5 domain create` command.

For example, run the following command to add a Kerberos domain to a Kerberos realm:

```
isi auth krb5 domain create <domain>
```

Modify an MIT Kerberos domain

You can modify an MIT Kerberos domain by modifying the realm settings.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to modify an MIT Kerberos domain.

Procedure

1. Run the `isi auth krb5 domain modify` command to modify a Kerberos domain.

For example, run the following command to modify a Kerberos domain by specifying an alternate Kerberos realm:

```
isi auth krb5 domain modify <domain> --realm <realm>
```

View an MIT Kerberos domain mapping

You can view the properties of an MIT Kerberos domain mapping.

Procedure

1. Run the `isi auth krb5 domain view` command with a value specified for the `<domain>` variable to view the properties of a Kerberos domain mapping:

```
isi auth krb5 domain view <domain>
```

List MIT Kerberos domains

You can list one or more MIT Kerberos domains and display the list in a tabular, JSON, CSV, or list format. You can also specify a limit for the number of domains to be listed.

Procedure

1. Run the `isi auth krb5 domain list` command to list one or more MIT Kerberos domains.

For example, run the following command to list the first ten MIT Kerberos domains in a tabular format without any headers or footers:

```
isi auth krb5 domain list -l=10 --format=table --no-header --no-footer
```

Delete an MIT Kerberos domain mapping

You can delete one or more MIT Kerberos domain mappings.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to delete an MIT Kerberos domain mapping.

Procedure

1. Run the `isi auth krb5 domain delete` command to delete an MIT Kerberos domain mapping.

For example, run the following command to delete a domain mapping:

```
isi auth krb5 domain delete <domain>
```

Managing SPNs and keys

A service principal name (SPN) is the name referenced by a client to identify an instance of a service on an EMC Isilon cluster. An MIT Kerberos provider authenticates services on a cluster through SPNs.

You can perform the following operations on SPNs and their associated keys:

- Update the SPNs if there are any changes to the SmartConnect zone settings that are based on those SPNs
- List the registered SPNs to compare them against a list of discovered SPNs
- Update keys associated with the SPNs either manually or automatically
- Import keys from a keytab table
- Delete specific key versions or delete all the keys associated with an SPN

View SPNs and keys

You can view the service principal names (SPNs) and their associated keys that are registered for an MIT Kerberos provider. Clients obtain Kerberos tickets and access services on EMC Isilon clusters through SPNs and their associated keys.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to view SPNs and keys.

Procedure

1. Run the `isi auth krb5 spn list` command to list one or more SPNs and their associated keys and the Key version numbers (Kvno).

For example, run the following command to list the first five SPNs for an MIT Kerberos provider in a tabular format without any headers or footers:

```
isi auth krb5 list <provider-name> -l 5 --format table --no-header
--no-footer <spn-list>
```

Delete keys

You can delete specific key versions or all the keys associated with a service principal name (SPN).

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to delete keys.

After creating new keys due to security reasons or to match configuration changes, follow this procedure to delete older version of the keys so that the keytab table is not populated with redundant keys.

Procedure

1. Run the `isi auth krb5 spn delete` command to delete all keys for a specified SPN or a specific version of a key.

For example, run the following command to delete all the keys associated with an SPN for an MIT Kerberos provider:

```
isi auth krb5 spn delete <provider-name> <spn> --all
```

The *<provider-name>* is the name of the MIT Kerberos provider. You can delete a specific version of the key by specifying a key version number value for the `kvno` argument and including that value in the command syntax.

Manually add or update a key for an SPN

You can manually add or update keys for a service principal name (SPN). This process creates a new key for the specified SPN.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to add or update a key for an SPN.

Procedure

1. Run the `isi auth krb5 spn create` command to add or update keys for an SPN.

For example, run the following command to add or update a key for an SPN by specifying the *<provider-name>*, *<user>*, and *<spn>* positional arguments:

```
isi auth krb5 spn create <provider-name> <user> <spn>
```

Automatically update an SPN

You can automatically update or add a service principal name (SPN) if it is registered with an MIT Kerberos provider but does not appear in the list of discovered SPNs.

Before you begin

You must be a member of a role that has `ISI_PRIV_AUTH` privileges to automatically update an SPN.

Procedure

1. Run the `isi auth krb5 spn check` command to compare the list of registered SPNs against the list of discovered SPNs.

Proceed to the next step if the comparison does not show similar results.

2. Run the `isi auth krb5 spn fix` command to fix the missing SPNs.

For example, run the following command to add missing SPNs for an MIT Kerberos service provider:

```
isi auth krb5 spn fix <provider-name> <user>
```

You can optionally specify a password for *<user>* which is the placeholder for a user who has the permission to join clients to the given domain.

Import a keytab file

An MIT Kerberos provider joined through a legacy keytab file might not have the ability to manage keys through the Kerberos admin credentials. In such a case, import a new keytab file and then add the keytab file keys to the provider.

Before you begin

Make sure that the following pre-requisites are met before you import a keytab file:

- You must create and copy a keytab file to a node on the cluster where you will perform this procedure.
- You must be a member of a role that has `ISI_PRIV_AUTH` privileges to import a keytab file.

Procedure

1. Import the keys of a keytab file by running the `isi auth krb5 spn import` command.

For example, run the following command to import the keys of the *<keytab-file>* to the provider referenced as *<provider-name>*:

```
isi auth krb5 spn import <provider-name> <keytab-file>
```

Managing access permissions

The internal representation of identities and permissions can contain information from UNIX sources, Windows sources, or both. Because access protocols can process the

information from only one of these sources, the system may need to make approximations to present the information in a format the protocol can process.

View expected user permissions

You can view the expected permissions for user access to a file or directory.

This procedure must be performed through the command-line interface (CLI).

Procedure

1. Establish an SSH connection to any node in the cluster.
2. View expected user permissions by running the `isi auth access` command.

The following command displays permissions in `/ifs/` for the user that you specify in place of `<username>`:

```
isi auth access <username> /ifs/
```

The system displays output similar to the following example:

```
User
  Name : <username>
  UID  : 2018
  SID  :
SID:S-1-5-21-2141457107-1514332578-1691322784-1018
File
  Owner : user:root
  Group : group:wheel
  Mode  : drwxrwxrwx
  Relevant Mode : d---rwx---
Permissions
  Expected : user:<username> \
  allow
dir_gen_read,dir_gen_write,dir_gen_execute,delete_child
```

3. View mode-bits permissions for a user by running the `isi auth access` command.

The following command displays verbose-mode file permissions information in `/ifs/` for the user that you specify in place of `<username>`:

```
isi auth access <username> /ifs/ -v
```

The system displays output similar to the following example:

```
User Name : <username> UID \
: 2018 SID : SID:S-1-5-21-2141457107-1514332578-1691322784-1018
File Owner : user:root Group : group:wheel Mode : drwxrwxrwx
Relevant Mode : d---rwx--- Permissions Expected : user:<username>
allow dir_gen_read,dir_gen_write,dir_gen_execute,delete_child
```

4. View expected ACL user permissions on a file for a user by running the `isi auth access` command.

The following command displays verbose-mode ACL file permissions for the file `file_with_acl.tx` in `/ifs/data/` for the user that you specify in place of `<username>`:

```
isi auth access <username> /ifs/data/file_with_acl.tx -v
```

The system displays output similar to the following example:

```
User Name : <username> \
UID : 2097 SID :
SID:S-1-7-21-2141457107-1614332578-1691322789-1018
File Owner : user:<username> Group : group:wheel
Permissions Expected : user:<username>
allow file_gen_read,file_gen_write,std_write_dac
Relevant Acl: group:<group-name> Users_allow file_gen_read
user:<username> allow std_write_dac,file_write,
append,file_write_ext_attr,file_write_attr
group:wheel allow file_gen_read,file_gen_write
```

Configure access management settings

Default access settings include whether to send NTLMv2 responses for SMB connections, the identity type to store on disk, the Windows workgroup name for running in local mode, and character substitution for spaces encountered in user and group names.

Procedure

1. Configure access management settings by running the `isi auth settings global modify` command.

The following command modifies global settings for a workgroup:

```
isi auth settings global modify \
--send-ntlmv2=false --on-disk-identity=native \
--space-replacement="_" --workgroup=WORKGROUP
```

Modify ACL policy settings

You can modify ACL policy settings but the default ACL policy settings are sufficient for most cluster deployments.

CAUTION

Because ACL policies change the behavior of permissions throughout the system, they should be modified only as necessary by experienced administrators with advanced knowledge of Windows ACLs. This is especially true for the advanced settings, which are applied regardless of the cluster's environment.

For UNIX, Windows, or balanced environments, the optimal permission policy settings are selected and cannot be modified. You can choose to manually configure the cluster's default permission settings if necessary to support your particular environment, however.

Procedure

1. Run the following command to modify ACL policy settings, where *<provider-name>* specifies the name of the provider:

```
isi auth ads modify <provider-name>
```

Update cluster permissions

You can update file permissions or ownership by running the PermissionRepair job.

To prevent permissions issues that can occur after changing the on-disk identity, run this authentication and access control job with `convert` mode specified to ensure that the changes are fully propagated throughout the cluster.

Procedure

1. Update cluster permissions by running the `isi job jobs start` command with the following syntax.

The following command updates cluster permissions, where `permissionrepair` specifies the job type, where variables in angle brackets are placeholders for values specific to your environment:

```
isi job start permissionrepair --priority <1-10> \
--policy <policy> --mode <clone | inherit | convert > \
--mapping-type=<system | sid | unix | native> --zone <zone-name>
```

Note

You cannot combine the `--template` parameter with the `convert` mode option, but you can combine the parameter with the `clone` and `inherit` mode options. Conversely, you cannot combine the `--mapping-type` and `--zone` parameters with the `clone` and `inherit` mode options, but you can combine the parameters with the `convert` mode option.

Example 1 Examples

The following example updates cluster permissions, where `permissionrepair` specifies the job type, the priority is 3, the chosen mode is `convert`, and the mapping type is `unix`:

```
isi job jobs start permissionrepair --priority=3 \
--policy myPolicy --mode=convert --mapping-type=unix \
--template <isi path> --path </ifs directory> --zone zone2
```

Authentication and access control commands

You can control access to your cluster through the authentication and access control commands.

isi auth access

Lists that permissions that a user has to access a given file.

Note

This command does not display SMB share permissions or privileges.

Syntax

```
isi auth access {<name> | --uid <integer> | --sid <string>} <path>
[--zone <string>]
[--numeric]
[--verbose]
```

Options

`<name>`

Specifies the user name.

`--sid <string>`
 Specifies the user by SID.

`--uid <integer>`
 Specifies the user by UID

`<path>`
 Specifies the file path in `/ifs`.

`--zone <string>`
 Specifies the access zone for the user.

`{--numeric | -n}`
 Displays the numeric identifier of the user.

`{--verbose | -v}`
 Displays more detailed information.

isi auth ads create

Configures an Active Directory provider and joins an Active Directory domain.

Syntax

```
isi auth ads create <name> <user>
  [--password <string>]
  [--account <string>]
  [--organizational-unit <string>]
  [--kerberos-nfs-spn {yes | no} ]
  [--dns-domain <dns-domain>]
  [--allocate-gids {yes | no}]
  [--allocate-uids {yes | no}]
  [--cache-entry-expiry <duration>]
  [--assume-default-domain {yes | no}]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--domain-offline-alerts {yes | no}]
  [--home-directory-template <path>]
  [--ignore-all-trusts {yes | no}]
  [--ignored-trusted-domains <dns-domain>]
  [--include-trusted-domains <dns-domain>]
  [--ldap-sign-and-seal {yes | no}]
  [--node-dc-affinity <string>]
  [--node-dc-affinity-timeout <timestamp>]
  [--login-shell <path>]
  [--lookup-domains <dns-domain>]
  [--lookup-groups {yes | no}]
  [--lookup-normalize-groups {yes | no}]
  [--lookup-normalize-users {yes | no}]
  [--lookup-users {yes | no}]
  [--machine-password-changes {yes | no}]
  [--machine-password-lifespan <duration>]
  [--nss-enumeration {yes | no}]
  [--sfu-support {none | rfc2307}]
  [--store-sfu-mappings {yes | no}]
  [--verbose]
```

Options

`<name>`
 Specifies a fully qualified Active Directory domain name, which will also be used as the provider name.

`<user>`

Specifies the user name of an account that has permission to join machine accounts to the Active Directory domain.

`--password <string>`

Specifies the password of the provided user account. If you omit this option, you will be prompted to supply a password.

`--account <string>`

Specifies the machine account name to use in Active Directory. By default, the cluster name is used.

`--organizational-unit <string>`

Specifies the name of the organizational unit (OU) to connect to on the Active Directory server. Specify the OU in the form **OuName** or **OuName1/SubName2**.

`--kerberos-nfs-spn {yes | no}`

Specifies whether to add SPNs for using Kerberized NFS.

`--dns-domain <dns-domain>`

Specifies a DNS search domain to use instead of the domain that is specified in the `--name` setting.

`--allocate-gids {yes | no}`

Enables or disables GID allocation for unmapped Active Directory groups. Active Directory groups without GIDs can be proactively assigned a GID by the ID mapper. If this option is disabled, GIDs are not proactively assigned, but when a user's primary group does not include a GID, the system may allocate one.

`--allocate-uids {yes | no}`

Enables or disables UID allocation for unmapped Active Directory users. Active Directory users without UIDs can be proactively assigned a UID by the ID mapper. If this option is disabled, UIDs are not proactively assigned, but when a user's identity does not include a UID, the system may allocate one.

`--cache-entry-expiry <duration>`

Specifies how long to cache a user or group, in the format *<integer>*{Y|M|W|D|H|m|s}.

`--assume-default-domain {yes | no}`

Specifies whether to look up unqualified user names in the primary domain. If this option is set to `no`, the primary domain must be specified for each authentication operation.

`--check-online-interval <duration>`

Specifies the time between provider online checks, in the format *<integer>*{Y|M|W|D|H|m|s}.

`--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time that a user logs in, if a home directory does not already exist for the user.

`--domain-offline-alerts {yes | no}`

Specifies whether to send an alert if the domain goes offline. If this option is set to `yes`, notifications are sent as specified in the global notification rules. The default value is `no`.

`--home-directory-template <path>`

Specifies the template path to use when creating home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically

replaced with strings at home directory creation time that represent specific variables. For example, %U, %D, and %Z are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

`--ignore-all-trusts {yes | no}`

Specifies whether to ignore all trusted domains.

`--ignored-trusted-domains <dns-domain>`

Specifies a list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

`--include-trusted-domains <dns-domain>`

Specifies a list of trusted domain to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--ldap-sign-and-seal {yes | no}`

Specifies whether to use encryption and signing on LDAP requests to a DC.

`{--node-dc-affinity | -x} <string>`

Specifies the domain controller that the node should exclusively communicate with (affinitize to). This option should be used with a timeout value, which is configured using the `--node-dc-affinity-timeout` option. Otherwise, the default timeout value of 30 minutes is assigned.

Note

This setting is for debugging purposes and should be left unconfigured during normal operation. To disable this feature, use a timeout value of 0.

`{--node-dc-affinity-timeout} <timestamp>`

Specifies the timeout setting for the local node affinity to a domain controller, using the date format `<YYYY> <MM> <DD>` or the date/time format `<YYYY> <MM> <DD>T<hh>:<mm>[:<ss>]`.

Note

A value of 0 disables the affinity. When affinitization is disabled, communication with the specified domain controller may not end immediately. It may persist until another domain controller can be chosen.

`--login-shell <path>`

Specifies the full path to the login shell to use if the Active Directory server does not provide login-shell information. This setting applies only to users who access the file system through SSH.

`--lookup-domains <string>`

Specifies a list of domains to which user and group lookups are to be limited. Repeat this option to specify multiple list items.

`--lookup-groups {yes | no}`

Specifies whether to look up Active Directory groups in other providers before allocating a GID.

`--lookup-normalize-groups {yes | no}`

Specifies whether to normalize Active Directory group names to lowercase before looking them up.

```
--lookup-normalize-users {yes | no}
    Specifies whether to normalize Active Directory user names to lowercase before
    looking them up.
--lookup-users {yes | no}
    Specifies whether to look up Active Directory users in other providers before
    allocating a UID.
--machine-password-password {yes | no}
    Specifies whether to enable periodic changes of the machine account password for
    security purposes.
--machine-password-lifespan <duration>
    Sets the maximum age of the machine account password, in the format <integer>{Y|M|
    W|D|H|m|s}.
--nss-enumeration {yes | no}
    Specifies whether to allow the Active Directory provider to respond to getpwent and
    getgrent requests.
--sfu-support {none | rfc2307}
    Specifies whether to support RFC 2307 attributes for Windows domain controllers.
    RFC 2307 is required for Windows UNIX Integration and for Services For UNIX (SFU)
    technologies.
--store-sfu-mappings {yes | no}
    Specifies whether to store SFU mappings permanently in the ID mapper.
{--verbose | -v}
    Displays the results of running the command.
```

isi auth ads delete

Deletes an Active Directory provider, which includes leaving the Active Directory domain that the provider is joined to. Leaving an Active Directory domain disrupts service for users who are accessing the domain. After you leave an Active Directory domain, users can no longer access the domain from the cluster.

Syntax

```
isi auth ads delete <provider-name>
    [--force]
    [--verbose]
```

Options

```
<provider-name>
    Specifies the name of the provider to delete.
{--force | -f}
    Suppresses command-line prompts and messages.
{--verbose | -v}
    Displays the results of running the command.
```


Examples

To leave an Active Directory domain named `some.domain.org` and delete the authentication provider that is associated with it, run the following command:

```
isi auth ads delete some.domain.org
```

At the confirmation prompt, type `y`.

isi auth ads list

Displays a list of Active Directory providers.

Syntax

```
isi auth ads list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

Examples

To view a list of all the Active Directory providers that the cluster is joined to, run the following command:

```
isi auth ads list
```

The system displays output similar to the following example:

Name	Authentication	Status	DC Name	Site
AD.EAST.EMC.COM	Yes	online	-	BOS
AD.NORTH.EMC.COM	Yes	online	-	VAN
AD.SOUTH.EMC.COM	No	online	-	TIJ
AD.WEST.EMC.COM	Yes	online	-	SEA

Total: 4				

isi auth ads modify

Modifies an Active Directory authentication provider.

Syntax

```
isi auth ads modify <provider-name>
  [--reset-schannel {yes | no}]
  [--domain-controller <string>]
  [--allocate-gids {yes | no}]
  [--allocate-uids {yes | no}]
  [--cache-entry-expiry <duration>]
  [--assume-default-domain {yes | no}]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--domain-offline-alerts {yes | no}]
  [--home-directory-template <path>]
  [--ignore-all-trusts {yes | no}]
  [--ignored-trusted-domains <dns-domain>]
  [--clear-ignored-trusted-domains]
  [--add-ignored-trusted-domains <dns-domain>]
  [--remove-ignored-trusted-domains <dns-domain>]
  [--include-trusted-domains <dns-domain>]
  [--clear-include-trusted-domains]
  [--add-include-trusted-domains <dns-domain>]
  [--remove-include-trusted-domains <dns-domain>]
  [--ldap-sign-and-seal {yes | no}]
  [--node-dc-affinity <string>]
  [--node-dc-affinity-timeout <timestamp>]
  [--login-shell <path>]
  [--lookup-domains <dns-domain>]
  [--clear-lookup-domains]
  [--add-lookup-domains <dns-domain>]
  [--remove-lookup-domains <dns-domain>]
  [--lookup-groups {yes | no}]
  [--lookup-normalize-groups {yes | no}]
  [--lookup-normalize-users {yes | no}]
  [--lookup-users {yes | no}]
  [--machine-password-changes {yes | no}]
  [--machine-password-lifespan <duration>]
  [--nss-enumeration {yes | no}]
  [--sfu-support {none | rfc2307}]
  [--store-sfu-mappings {yes | no}]
  [--verbose]
```

Options

<provider-name>

Specifies the domain name that the Active Directory provider is joined to, which is also the Active Directory provider name.

`--reset-schannel {yes | no}`

Resets the secure channel to the primary domain.

`--domain-controller <dns-domain>`

Specifies a domain controller.

`--allocate-gids {yes | no}`

Enables or disables GID allocation for unmapped Active Directory groups. Active Directory groups without GIDs can be proactively assigned a GID by the ID mapper. If this option is disabled, GIDs are not assigned proactively, but when a user's primary group does not include a GID, the system may allocate one.

`--allocate-uids {yes | no}`

Enables or disables UID allocation for unmapped Active Directory users. Active Directory users without UIDs can be proactively assigned a UID by the ID mapper. If this option is disabled, UIDs are not assigned proactively, but when a user's identity does not include a UID, the system may allocate one.

`--cache-entry-expiry <duration>`

Specifies how long to cache a user or group, in the format *<integer>*{Y|M|W|D|H|m|s}.

`--assume-default-domain {yes | no}`

Specifies whether to look up unqualified user names in the primary domain. If this option is set to `no`, the primary domain must be specified for each authentication operation.

`--check-online-interval <duration>`

Specifies the time between provider online checks, in the format *<integer>*{Y|M|W|D|H|m|s}.

`--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

`--domain-offline-alerts {yes | no}`

Specifies whether to send an alert if the domain goes offline. If this option is set to `yes`, notifications are sent as specified in the global notification rules. The default value is `no`.

`--home-directory-template <path>`

Specifies the template path to use when creating home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

`--ignore-all-trusts {yes | no}`

Specifies whether to ignore all trusted domains.

`--ignored-trusted-domains <dns-domain>`

Specifies a list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

`--clear-ignored-trusted-domains`

Clears the list of ignored trusted domains if `--ignore-all-trusts` is disabled.

`--add-ignored-trusted-domains <dns-domain>`

Adds a domain to the list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

`--remove-ignored-trusted-domains <dns-domain>`

Removes a specified domain from the list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

`--include-trusted-domains <dns-domain>`

Specifies a list of trusted domains to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--clear-include-trusted-domains`

Clears the list of trusted domains to include if `--ignore-all-trusts` is enabled.

`--add-include-trusted-domains <dns-domain>`

Adds a domain to the list of trusted domains to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--remove-include-trusted-domains <dns-domain>`

Removes a specified domain from the list of trusted domains to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--ldap-sign-and-seal {yes | no}`

Specifies whether to use encryption and signing on LDAP requests to a domain controller.

`{--node-dc-affinity [-x] <string>`

Specifies the domain controller that the node should exclusively communicate with (affinitize). This option should be used with a timeout value, which is configured using the `--node-dc-affinity-timeout` option. Otherwise, the default timeout value of 30 minutes is assigned.

Note

This setting is for debugging purposes and should be left unconfigured during normal operation. To disable this feature, use a timeout value of 0.

`{--node-dc-affinity-timeout} <timestamp>`

Specifies the timeout setting for the local node affinity to a domain controller, using the date format `<YYYY>-<MM>-<DD>` or the date/time format `<YYYY>-<MM>-<DD>T<hh>:<mm>[:<ss>]`.

Note

A value of 0 disables the affinity. When affinitization is disabled, communication with the specified domain controller may not end immediately. It may persist until another domain controller can be chosen.

`--login-shell <path>`

Specifies the path to the login shell to use if the Active Directory server does not provide login-shell information. This setting applies only to users who access the file system through SSH.

`--lookup-domains <string>`

Specifies a list of domains to which user and group lookups are to be limited. Repeat this option to specify multiple list items.

`--clear-lookup-domains`

Clears the list of restricted domains for user and group lookups.

`--add-lookup-domains <string>`

Adds an entry to the restricted list of domains to use for user and group lookups. Repeat this option to specify multiple list items.

`--remove-lookup-domains <string>`

Removes an entry from the list of domains to use for user and group lookups. Repeat this option to specify multiple list items.

`--lookup-groups {yes | no}`

Specifies whether to look up Active Directory groups in other providers before allocating a GID.

`--lookup-normalize-groups {yes | no}`
Specifies whether to normalize Active Directory group names to lowercase before looking them up.

`--lookup-normalize-users {yes | no}`
Specifies whether to normalize Active Directory user names to lowercase before looking them up.

`--lookup-users {yes | no}`
Specifies whether to look up Active Directory users in other providers before allocating a UID.

`--machine-password-password {yes | no}`
Specifies whether to enable periodic changes of the machine account password for security purposes.

`--machine-password-lifespan <duration>`
Sets the maximum age of the machine account password, in the format *<integer>*{Y|M|W|D|H|m|s}.

`--nss-enumeration {yes | no}`
Specifies whether to allow the Active Directory provider to respond to `getpwent` and `getgrent` requests.

`--sfu-support {none | rfc2307}`
Specifies whether to support RFC 2307 attributes for domain controllers. RFC 2307 is required for Windows UNIX Integration and for Services For UNIX (SFU) technologies.

`--store-sfu-mappings {yes | no}`
Specifies whether to store SFU mappings permanently in the ID mapper.

`{--verbose | -v}`
Displays the results of running the command.

isi auth ads spn check

Checks valid service principal names (SPNs).

Syntax

```
isi auth ads spn check --domain <string>
  [--machinecreds]
  [--user <string> [--password <string>]]
  [--repair]
```

Options

`{--domain | -D} <string>`
Specifies the DNS domain name for the user or group that is attempting to connect to the cluster.

`--machinecreds`
Directs the system to use machine credentials when connecting to the cluster.

`{--user | -U} <string>`
Specifies an administrative user account to connect to the cluster, if required.

```
{--password | -P} <string>
```

Specifies the administrative user account password.

```
{--repair | -r}
```

Repairs missing SPNs.

isi auth ads spn create

Adds one or more service principal names (SPNs) for a machine account. SPNs must be propagated to all domain controllers to make them available to clients.

Syntax

```
isi auth ads spn create --spn <string>
  [--domain <string>]
  [--account <string>]
  [--machinecreds]
  [--user <string> [--password <string>]]
```

Options

```
{--spn | -s} <string>
```

Specifies an SPN to register. Repeat this option to specify multiple list items.

```
{--domain | -D} <string>
```

Specifies the DNS domain name for the user or group that is attempting to connect to the cluster.

```
{--account | -a} <string>
```

Specifies the address of the machine account. If no account is specified, the machine account of the cluster is used.

```
{--user | -U} <string>
```

Specifies an administrative user account to connect to the cluster, if required.

```
{--password | -P} <string>
```

Specifies the administrative user account password.

```
--machinecreds
```

Directs the system to use machine credentials when connecting to the cluster.

isi auth ads spn delete

Deletes one or more SPNs that are registered against a machine account.

Syntax

```
isi auth ads spn delete --spn <string>
  [--domain <string>]
  [--account <string>]
  [--machinecreds]
  [--user <string> [--password <string>]]
```

Options

```
{--spn | -s} <string>
```

Specifies an SPN to delete. Repeat this option to specify multiple list items.

```
{--domain | -D} <string>
```

Specifies the DNS domain name for the user or group that is attempting to connect to the cluster.

{--account | -a} *<string>*

Specifies the address of the machine account. If no account is specified, the machine account of the cluster is used.

--machinecreds

Directs the system to use machine credentials when connecting to the cluster.

{--user | -U} *<string>*

Specifies an administrative user account to connect to the cluster, if required.

{--password | -P} *<string>*

Specifies the administrative user account password.

isi auth ads spn list

Displays a list of service principal names (SPNs) that are registered against a machine account.

Syntax

```
isi auth ads spn list --domain <string>
  [--account <string>
  [--machinecreds]
  [--user <string> [--password <string>]]
```

Options

{--domain | -D} *<string>*

Specifies the DNS domain name for the user or group that is attempting to connect to the cluster.

{--account | -a} *<string>*

Specifies the address of the machine account. If no account is specified, the machine account of the cluster is used.

--machinecreds

Directs the system to use machine credentials when connecting to the cluster.

{--user | -U} *<string>*

Specifies an administrative user account to connect to the cluster, if required.

{--password | -P} *<string>*

Specifies the administrative user account password.

Examples

Run the following command to view a list of SPNs that are currently registered against the machine account:

```
isi auth ads spn list
```

The system displays output similar to the following example:

```
HOST/test
HOST/test.sample.isilon.com
```

isi auth ads trusts controllers list

Displays a list of domain controllers for a trusted domain.

Syntax

```
isi auth ads trusts controllers list <provider>
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<provider>

Specifies an Active Directory provider.

{--limit | -l} *<integer>*

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

The following command displays a list of trusted domains in an Active Directory provider named ad.isilon.com:

```
isi auth ads trusts controllers list ad.isilon.com
```

isi auth ads trusts list

Displays a list of trusted domains.

Syntax

```
isi auth ads trusts list <provider>
```

Options

<provider>

Specifies an Active Directory provider.

isi auth ads trusts view

Displays the properties of a trusted domain.

Syntax

```
isi auth ads trusts view <provider> <domain>
```

Options

<provider>

Specifies an Active Directory provider.

<domain>

Specifies the trusted domain to view.

isi auth ads view

Displays the properties of an Active Directory provider.

Syntax

```
isi auth ads view <provider-name>
  [--verbose]
```

Options

<provider-name>

Specifies the name of the provider to view.

{--verbose | -v}

Displays more detailed information.

isi auth error

Displays error code definitions from the authentication log files.

Syntax

```
isi auth error <error-code>
```

Options

<error-code>

Specifies the error code to convert.

Examples

To view the definition of error code 4, run the following command:

```
isi auth error 4
```

The system displays output similar to the following example:

```
4 = ERROR_TOO_MANY_OPEN_FILES
```

isi auth file create

Creates a file provider.

Syntax

```
isi auth file create <name>
  [--password-file <path>]
  [--group-file <path>]
  [--authentication {yes | no}]
  [--cache-entry-expiry <duration>]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--findable-users <string>]
  [--group-domain <string>]
  [--home-directory-template <path>]
  [--listable-groups <string>]
  [--listable-users <string>]
  [--login-shell <path>]
  [--modifiable-groups <string>]
  [--modifiable-users <string>]
  [--netgroup-file <path>]
  [--normalize-groups {yes | no}]
  [--normalize-users {yes | no}]
  [--ntlm-support {all | v2only | none}]
  [--provider-domain <string>]
  [--restrict-findable {yes | no}]
  [--restrict-listable {yes | no}]
  [--restrict-modifiable {yes | no}]
  [--unfindable-groups <string>]
  [--unfindable-users <string>]
  [--unlistable-groups <string>]
  [--unlistable-users <string>]
  [--unmodifiable-groups <string>]
  [--unmodifiable-users <string>]
  [--user-domain <string>]
  [--verbose]
```

Options

<name>

Sets the file provider name.

`--password-file <path>`

Specifies the path to a `passwd.db` replacement file.

`--group-file <path>`

Specifies the path to a `group` replacement file.

`--authentication {yes | no}`

Enables or disables the use of the provider for authentication as well as identity. The default value is `yes`.

`--cache-entry-expiry <duration>`

Specifies the length of time after which the cache entry will expire, in the format `<integer>{[Y | M | W | D | H | m | s]}`. To turn off cache expiration, set this value to `off`.

`--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

- `--enabled {yes | no}`
Enables or disables the provider.
- `--findable-groups <string>`
Specifies a list of groups that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable group. If populated, groups that are not included in this list cannot be resolved.
- `--findable-users <string>`
Specifies a list of users that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable user. If populated, users that are not included in this list cannot be resolved.
- `--group-domain <string>`
Specifies the domain that this provider will use to qualify groups. The default group domain is `FILE_GROUPS`.
- `--home-directory-template <path>`
Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.
- `--listable-groups <string>`
Specifies a group that can be listed if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be listed.
- `--listable-users <string>`
Specifies a user that can be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be listed.
- `--login-shell <path>`
Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.
- `--modifiable-groups <string>`
Specifies a group that can be modified in this provider if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be modified.
- `--modifiable-users <string>`
Specifies a user that can be modified in this provider if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be modified.
- `--netgroup-file <path>`
Specifies the path to a `netgroup` replacement file.
- `--normalize-groups {yes | no}`
Normalizes group names to lowercase before lookup.
- `--normalize-users {yes | no}`
Normalizes user names to lowercase before lookup.

```
--ntlm-support {all | v2only | none}
```

For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM.

```
--provider-domain <string>
```

Specifies the domain that the provider will use to qualify user and group names.

```
--restrict-findable {yes | no}
```

Specifies whether to check the provider for filtered lists of findable and unfindable users and groups.

```
--restrict-listable {yes | no}
```

Specifies whether to check the provider for filtered lists of listable and unlistable users and groups.

```
--restrict-modifiable {yes | no}
```

Specifies whether to check the provider for filtered lists of modifiable and unmodifiable users and groups.

```
--unfindable-groups <string>
```

If `--restrict-findable` is enabled and the findable groups list is empty, specifies a group that cannot be resolved by this provider. Repeat this option to specify multiple list items.

```
--unfindable-users <string>
```

If `--restrict-findable` is enabled and the findable users list is empty, specifies a user that cannot be resolved by this provider. Repeat this option to specify multiple list items.

```
--unlistable-groups <string>
```

If `--restrict-listable` is enabled and the listable groups list is empty, specifies a group that cannot be listed by this provider. Repeat this option to specify multiple list items.

```
--unlistable-users <string>
```

If `--restrict-listable` is enabled and the listable users list is empty, specifies a user that cannot be listed by this provider. Repeat this option to specify multiple list items.

```
--unmodifiable-groups <string>
```

If `--restrict-modifiable` is enabled and the modifiable groups list is empty, specifies a group that cannot be modified. Repeat this option to specify multiple list items.

```
--unmodifiable-users <string>
```

If `--restrict-modifiable` is enabled and the modifiable users list is empty, specifies a user that cannot be modified. Repeat this option to specify multiple list items.

```
--user-domain <string>
```

Specifies the domain that this provider will use to qualify users. The default user domain is `FILE_USERS`.

```
{--verbose | -v}
```

Displays more detailed information.

isi auth file delete

Deletes a file provider.

Syntax

```
isi auth file delete <provider-name>
  [--force]
  [--verbose]
```

Options

<provider-name>

Specifies the name of the provider to delete.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

isi auth file list

Displays a list of file providers.

Syntax

```
isi auth file list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} *<integer>*

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth file modify

Modifies a file provider.

Syntax

```
isi auth file modify <provider-name>
  [--provider <string>]
  [--password-file <path>]
  [--group-file <path>]
  [--authentication {yes | no}]
  [--cache-entry-expiry <duration>]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--clear-findable-groups]
  [--add-findable-groups <string>]
  [--remove-findable-groups <string>]
  [--findable-users <string>]
  [--clear-findable-users]
  [--add-findable-users <string>]
  [--remove-findable-users <string>]
  [--group-domain <string>]
  [--home-directory-template <path>]
  [--listable-groups <string>]
  [--clear-listable-groups]
  [--add-listable-groups <string>]
  [--remove-listable-groups <string>]
  [--listable-users <string>]
  [--clear-listable-users]
  [--add-listable-users <string>]
  [--remove-listable-users <string>]
  [--login-shell <path>]
  [--modifiable-groups <string>]
  [--clear-modifiable-groups]
  [--add-modifiable-groups <string>]
  [--remove-modifiable-groups <string>]
  [--modifiable-users <string>]
  [--clear-modifiable-users]
  [--add-modifiable-users <string>]
  [--remove-modifiable-users <string>]
  [--netgroup-file <path>]
  [--normalize-groups {yes | no}]
  [--normalize-users {yes | no}]
  [--ntlm-support {all | v2only | none}]
  [--provider-domain <string>]
  [--restrict-findable {yes | no}]
  [--restrict-listable {yes | no}]
  [--restrict-modifiable {yes | no}]
  [--unfindable-groups <string>]
  [--clear-unfindable-groups]
  [--add-unfindable-groups <string>]
  [--remove-unfindable-groups <string>]
  [--unfindable-users <string>]
  [--clear-unfindable-users]
  [--add-unfindable-users <string>]
  [--remove-unfindable-users <string>]
  [--unlistable-groups <string>]
  [--clear-unlistable-groups]
  [--add-unlistable-groups <string>]
  [--remove-unlistable-groups <string>]
  [--unlistable-users <string>]
  [--clear-unlistable-users]
  [--add-unlistable-users <string>]
  [--remove-unlistable-users <string>]
```

```

[--unmodifiable-groups <string>]
[--clear-unmodifiable-groups]
[--add-unmodifiable-groups <string>]
[--remove-unmodifiable-groups <string>]
[--unmodifiable-users <string>]
[--clear-unmodifiable-users]
[--add-unmodifiable-users <string>]
[--remove-unmodifiable-users <string>]
[--user-domain <string>]
[--verbose]

```

Options

<provider-name>

Specifies the name of the file provider to modify. This setting cannot be modified.

`--provider <string>`

Specifies an authentication provider of the format *<type>:<instance>*. Valid provider types are `ads`, `ldap`, `nis`, `file`, and `local`. For example, an LDAP provider named `auth1` can be specified as `ldap:auth1`.

`--password-file <path>`

Specifies the path to a `passwd.db` replacement file.

`--group-file <path>`

Specifies the path to a `group` replacement file.

`--authentication {yes | no}`

Enables or disables the use of the provider for authentication as well as identity. The default value is `yes`.

`--cache-entry-expiry <duration>`

Specifies the length of time after which the cache entry will expire, in the format *<integer>{[Y | M | W | D | H | m | s]}*. To turn off cache expiration, set this value to `off`.

`--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

`--enabled {yes | no}`

Enables or disables the provider.

`--enumerate-groups {yes | no}`

Specifies whether to allow the provider to enumerate groups.

`--enumerate-users {yes | no}`

Specifies whether to allow the provider to enumerate users.

`--findable-groups <string>`

Specifies a group that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be resolved. This option overwrites any existing entries in the findable groups list; to add or remove groups without affecting current entries, use `--add-findable-groups` or `--remove-findable-groups`.

`--clear-findable-groups`

Removes all entries from the list of findable groups.

```
--add-findable-groups <string>
    Adds an entry to the list of findable groups that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--remove-findable-groups <string>
    Removes an entry from the list of findable groups that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--findable-users <string>
    Specifies a user that can be found in the provider if --restrict-findable is
    enabled. Repeat this option to specify multiple list items. If populated, any users that
    are not included in this list cannot be resolved. This option overwrites any existing
    entries in the findable users list; to add or remove users without affecting current
    entries, use --add-findable-users or --remove-findable-users.
--clear-findable-users
    Removes all entries from the list of findable users.
--add-findable-users <string>
    Adds an entry to the list of findable users that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--remove-findable-users <string>
    Removes an entry from the list of findable users that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--group-domain <string>
    Specifies the domain that the provider will use to qualify groups. The default group
    domain is FILE_GROUPS.
--group-file <path>
    Specifies the path to a group replacement file.
--home-directory-template <path>
    Specifies the path to use as a template for naming home directories. The path must
    begin with /ifs and can include special character sequences that are dynamically
    replaced with strings at home directory creation time that represent specific
    variables. For example, %U, %D, and %Z are replaced with the user name, provider
    domain name, and zone name, respectively. For more information, see the Home
    directories section.
--listable-groups <string>
    Specifies a group that can be viewed in this provider if --restrict-listable is
    enabled. Repeat this option to specify multiple list items. If populated, any groups
    that are not included in this list cannot be viewed. This option overwrites any existing
    entries in the listable groups list; to add or remove groups without affecting current
    entries, use --add-listable-groups or --remove-listable-groups.
--clear-listable-groups
    Removes all entries from the list of viewable groups.
--add-listable-groups <string>
    Adds an entry to the list of viewable groups that is checked if --restrict-
    listable is enabled. Repeat this option to specify multiple list items.
--remove-listable-groups <string>
```


Removes an entry from the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--listable-users <string>`

Specifies a user that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be viewed. This option overwrites any existing entries in the listable users list; to add or remove users without affecting current entries, use `--add-listable-users` or `--remove-listable-users`.

`--clear-listable-users`

Removes all entries from the list of viewable users.

`--add-listable-users <string>`

Adds an entry to the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-listable-users <string>`

Removes an entry from the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--login-shell <path>`

Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.

`--modifiable-groups <string>`

Specifies a group that can be modified if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be modified. This option overwrites any existing entries in the modifiable groups list; to add or remove groups without affecting current entries, use `--add-modifiable-groups` or `--remove-modifiable-groups`.

`--clear-modifiable-groups`

Removes all entries from the list of modifiable groups.

`--add-modifiable-groups <string>`

Adds an entry to the list of modifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--remove-modifiable-groups <string>`

Removes an entry from the list of modifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--modifiable-users <string>`

Specifies a user that can be modified if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be modified. This option overwrites any existing entries in the modifiable users list; to add or remove users without affecting current entries, use `--add-modifiable-users` or `--remove-modifiable-users`.

`--clear-modifiable-users`

Removes all entries from the list of modifiable users.

`--add-modifiable-users <string>`

Adds an entry to the list of modifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--remove-modifiable-users <string>`
 Removes an entry from the list of modifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--netgroup-file <path>`
 Specifies the path to a `netgroup` replacement file.

`--normalize-groups {yes | no}`
 Normalizes group names to lowercase before lookup.

`--normalize-users {yes | no}`
 Normalizes user names to lowercase before lookup.

`--ntlm-support {all | v2only | none}`
 For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM and is recommended.

`--password-file <path>`
 Specifies the path to a `passwd.db` replacement file.

`--provider-domain <string>`
 Specifies the domain that this provider will use to qualify user and group names.

`--restrict-findable {yes | no}`
 Specifies whether to check this provider for filtered lists of findable and unfindable users and groups.

`--restrict-listable {yes | no}`
 Specifies whether to check this provider for filtered lists of viewable and unviewable users and groups.

`--restrict-modifiable {yes | no}`
 Specifies whether to check this provider for filtered lists of modifiable and unmodifiable users and groups.

`--unfindable-groups <string>`
 If `--restrict-findable` is enabled and the findable groups list is empty, specifies a group that cannot be resolved by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unfindable groups list; to add or remove groups without affecting current entries, use `--add-unfindable-groups` or `--remove-unfindable-groups`.

`--clear-unfindable-groups`
 Removes all entries from the list of unfindable groups.

`--add-unfindable-groups <string>`
 Adds an entry to the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-unfindable-groups <string>`
 Removes an entry from the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--unfindable-users <string>`
 If `--restrict-findable` is enabled and the findable users list is empty, specifies a user that cannot be resolved by this provider. Repeat this option to specify multiple

list items. This option overwrites any existing entries in the unfindable users list; to add or remove users without affecting current entries, use `--add-unfindable-users` or `--remove-unfindable-users`.

`--clear-unfindable-users`

Removes all entries from the list of unfindable groups.

`--add-unfindable-users <string>`

Adds an entry to the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-unfindable-users <string>`

Removes an entry from the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--unlistable-groups <string>`

If `--restrict-listable` is enabled and the viewable groups list is empty, specifies a group that cannot be listed by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unlistable groups list; to add or remove groups without affecting current entries, use `--add-unlistable-groups` or `--remove-unlistable-groups`.

`--clear-unlistable-groups`

Removes all entries from the list of unviewable groups.

`--add-unlistable-groups <string>`

Adds an entry to the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-unlistable-groups <string>`

Removes an entry from the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--unlistable-users <string>`

If `--restrict-listable` is enabled and the viewable users list is empty, specifies a user that cannot be listed by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unlistable users list; to add or remove users without affecting current entries, use `--add-unlistable-users` or `--remove-unlistable-users`.

`--clear-unlistable-users`

Removes all entries from the list of unviewable users.

`--add-unlistable-users <string>`

Adds an entry to the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-unlistable-users <string>`

Removes an entry from the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--unmodifiable-groups <string>`

If `--restrict-modifiable` is enabled and the modifiable groups list is empty, specifies a group that cannot be modified. Repeat this option to specify multiple list items. This option overwrites any existing entries in this provider's unmodifiable

groups list; to add or remove groups without affecting current entries, use `--add-unmodifiable-groups` or `--remove-unmodifiable-groups`.

`--clear-unmodifiable-groups`
Removes all entries from the list of unmodifiable groups.

`--add-unmodifiable-groups <string>`
Adds an entry to the list of unmodifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--remove-unmodifiable-groups <string>`
Removes an entry from the list of unmodifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--unmodifiable-users <string>`
If `--restrict-modifiable` is enabled and the modifiable users list is empty, specifies a user that cannot be modified. Repeat this option to specify multiple list items. This option overwrites any existing entries in this provider's unmodifiable users list; to add or remove users without affecting current entries, use `--add-unmodifiable-users` or `--remove-unmodifiable-users`.

`--clear-unmodifiable-users`
Removes all entries from the list of unmodifiable users.

`--add-unmodifiable-users <string>`
Adds an entry to the list of unmodifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--remove-unmodifiable-users <string>`
Removes an entry from the list of unmodifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

`--user-domain <string>`
Specifies the domain that this provider will use to qualify users. The default user domain is `FILE_USERS`.

`{--verbose | -v}`
Displays the results of running the command.

isi auth file view

Displays the properties of a file provider.

Syntax

```
isi auth file view <provider-name>
```

Options

`<provider-name>`

Specifies the name of the provider to view.

isi auth groups create

Creates a local group.

Syntax

```
isi auth groups create <name>
  [--gid <integer>]
  [--add-user <name>]
  [--add-uid <integer>]
  [--add-sid <string>]
  [--add-wellknown <name>]
  [--sid <string>]
  [--zone <string>]
  [--provider <string>]
  [--verbose]
  [--force]
```

Options

<name>

Specifies the group name.

`--gid <integer>`

Overrides automatic allocation of the UNIX group identifier (GID) with the specified value. Setting this option is not recommended.

`--add-user <name>`

Specifies the name of the user to add to the group. Repeat this option to specify multiple users.

`--add-uid <integer>`

Specifies the UID of the user to add to the group. Repeat this option to specify multiple users.

`--add-sid <string>`

Specifies the SID of the user to add to the group. Repeat this option to specify multiple users.

`--add-wellknown <name>`

Specifies a wellknown persona name to add to the group. Repeat this option to specify multiple personas.

`--sid <string>`

Sets the Windows security identifier (SID) for the group, for example S-1-5-21-13.

`--zone <string>`

Specifies the access zone in which to create the group.

`--provider <string>`

Specifies a local authentication provider in the specified access zone.

`{--verbose | -v}`

Displays more detailed information.

`{--force | -f}`

Suppresses command-line prompts and messages.

isi auth groups delete

Removes a local group from the system. Members of a group are removed before the group is deleted.

Syntax

```
isi auth groups delete {<group> | --gid <integer> | --sid <string>}
  [--zone <string>]
  [--provider <string>]
  [--force]
  [--verbose]
```

Options

This command requires *<group>*, *--gid <integer>*, or *--sid <string>*.

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GID.

<group>

--sid <string>

Specifies the group by SID.

--zone <string>

Specifies the name of the access zone that contains the group.

--provider <string>

Specifies the group's authentication provider.

{*--force* | *-f*}

Suppresses command-line prompts and messages.

{*--verbose* | *-v*}

Displays the results of running the command.

isi auth groups flush

Flushes cached group information.

Syntax

```
isi auth groups flush
```

Options

There are no options for this command.

Examples

To flush all cached group information, run the following command:

```
isi auth groups flush
```

isi auth groups list

Displays a list of groups.

Syntax

```
isi auth groups list
  [--domain <string>]
  [--zone <string>]
  [--provider <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--domain <string>`

Specifies the provider domain.

`--zone <string>`

Specifies an access zone.

`--provider <string>`

Specifies an authentication provider.

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi auth groups modify

Modifies a local group.

Syntax

```
isi auth groups modify {<group> | --gid <integer> | --sid <string>}
  [--new-gid <integer>]
  [--add-uid <integer>]
  [--remove-uid <integer>]
  [--add-user <name>]
  [--remove-user <name>]
  [--add-sid <string>]
  [--remove-sid <string>]
  [--add-wellknown <name>]
  [--remove-wellknown <name>]
  [--zone <string>]
  [--provider <string>]
```

```
[--verbose]
[--force]
```

Options

This command requires *<group>*, `--gid <integer>`, or `--sid <string>`.

<group>

Specifies the group by name.

`--gid <integer>`

Specifies the group by GID.

`--sid <string>`

Specifies the group by SID.

`--new-gid <integer>`

Specifies a new GID for the group. Setting this option is not recommended.

`--add-uid <integer>`

Specifies the UID of a user to add to the group. Repeat this option to specify multiple list items.

`--remove-uid <integer>`

Specifies the UID of a user to remove from the group. Repeat this option to specify multiple list items.

`--add-user <name>`

Specifies the name of a user to add to the group. Repeat this option to specify multiple list items.

`--remove-user <name>`

Specifies the name of a user to remove from the group. Repeat this option to specify multiple list items.

`--add-sid <string>`

Specifies the SID of an object to add to the group, for example S-1-5-21-13. Repeat this option to specify multiple list items.

`--remove-sid <string>`

Specifies the SID of an object to remove from the group. Repeat this option to specify multiple list items.

`--add-wellknown <name>`

Specifies a well-known SID to add to the group. Repeat this option to specify multiple list items.

`--remove-wellknown <name>`

Specifies a well-known SID to remove from the group. Repeat this option to specify multiple list items.

`--zone <string>`

Specifies the group's access zone.

`--provider <string>`

Specifies the group's authentication provider.

`{--verbose | -v}`

Displays more detailed information.


```
{--force | -f}
```

Suppresses command-line prompts and messages.

isi auth groups members list

Displays a list of members that are associated with a group.

Syntax

```
isi auth groups members list {<group> | --gid <integer> | --sid
<string>}
  [--zone <string>]
  [--provider <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

This command requires *<group>*, *--gid <integer>*, or *--sid <string>*.

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GUID.

--sid <string>

Specifies the group by SID.

--zone <string>

Specifies an access zone.

--provider <string>

Specifies an authentication provider.

{--limit | -l} *<integer>*

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth groups view

Displays the properties of a group.

Syntax

```
isi auth groups view {<group> | --gid <integer> | --sid <string>}
  [--zone <string>]
  [--show-groups]
  [--provider <string>]
```

Options

This command requires *<group>*, *--gid <integer>*, or *--sid <string>*.

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GUID.

--sid <string>

Specifies the group by SID.

--zone <string>

Specifies an access zone.

--show-groups

Displays groups that include this group as a member.

--provider <string>

Specifies an authentication provider.

isi auth id

Displays your access token.

Syntax

```
isi auth id
```

Options

There are no options for this command.

isi auth krb5 realm create

Creates an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm create <realm>
  [--is-default-realm <boolean>]
  [--kdc <string>]
  [--admin-server <string>]
  [--default-domain <string>]
```

Options

<realm>

Specifies the name of the realm.

`--is-default-realm <boolean>`

Specifies whether realm will be the default.

`--kdc <string>`

Specifies the hostname or IP address of a Key Distribution Center (KDC). Specify `--kdc` for each additional hostname or IP address of a KDC.

`--admin-server <string>`

Specifies the hostname or IP address of the administrative server (master KDC).

`--default-domain <string>`

Specifies the default domain for the realm used for translating the v4 principal names.

isi auth krb5 realm delete

Deletes an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm delete <realm>
    [--force]
```

Options

`<realm>`

Specifies the Kerberos realm name.

`{--force | -f}`

Specifies not to ask for a confirmation.

isi auth krb5 realm modify

Modifies an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm modify <realm>
    [--is-default-realm <boolean>]
    [--kdc <string>]
    [--admin-server <string>]
    [--default-domain <string>]
```

Options

`<realm>`

Specifies the Kerberos realm name.

`--is-default-realm <boolean>`

Specifies whether the Kerberos realm will be the default.

`--kdc <string>`

Specifies the hostname or IP address of the Key Distribution Center (KDC). Specify `--kdc` for each additional hostname or IP address of the KDC.

`--admin-server <string>`

Specifies the hostname or IP address of the administrative server (master KDC).

`--default-domain <string>`

Specifies the default domain for the Kerberos realm used for translating the v4 principal names.

isi auth krb5 realm list

Displays a list of MIT Kerberos realms.

Syntax

```
isi auth krb5 realm list
    [--limit <integer>]
    [--format {table | json | csv | list}]
    [--no-header]
    [--no-footer]
```

Options

{--limit | -l} <integer>

Specifies the number of Kerberos realms to display.

--format {table | json | csv | list}

Specifies whether to display the Kerberos realms in a tabular, JSON, CSV, or list format.

{--no-header | -a}

Specifies not to display the headers in the CSV or tabular formats.

{--no-footer | -z}

Specifies not to display the table summary footer information.

isi auth krb5 realm view

Displays the properties of an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm view <realm>
```

Options

<realm>

Specifies the Kerberos realm name.

isi auth krb5 create

Creates an MIT Kerberos provider and joins a user to an MIT Kerberos realm.

Syntax

```
isi auth krb5 create <realm> {<user> | --keytab-file <string> }
    [--password <string>]
    [--spn<string>]
    [--is-default-realm <boolean>]
    [--kdc <string>]
    [--admin-server <string>]
    [--default-domain <string>]
```

Options

<realm>

Specifies the Kerberos realm name.

<user>

Specifies a user name with permissions to create the service principal names (SPNs) in the given Kerberos realm.

`--keytab-file <string>`

Specifies the keytab file to import.

`--password <string>`

Specifies the password used for joining a Kerberos realm.

`--spn <string>`

Specifies the SPNs to register. Specify `--spn` for each additional SPN that you want to register.

`--is-default-realm <boolean>`

Specifies whether the Kerberos realm will be the default.

`--kdc <string>`

Specifies the hostname or IP address of the Key Distribution Center (KDC). Specify `--kdc` for each additional hostname or IP address of the KDC.

`--admin-server <string>`

Specifies the hostname or IP address of the administrative server (master KDC).

`--default-domain <string>`

Specifies the default Kerberos domain for the Kerberos realm used for translating v4 principal names.

isi auth krb5 delete

Deletes an MIT Kerberos authentication provider and removes the user from an MIT Kerberos realm.

Syntax

```
isi auth krb5 delete <provider-name>
                    [--force]
```

Options

<provider-name>

Specifies the Kerberos provider name.

`{--force | -f}`

Specifies not to ask for a confirmation.

isi auth krb5 list

Displays a list of MIT Kerberos authentication providers.

Syntax

```
isi auth krb5 list
                [--limit <integer>]
                [--format {table | json | csv | list}]
                [--no-header]
                [--no-footer]
```

Options

```
{--limit | -l} <integer>
```

Specifies the number of Kerberos providers to display.

```
--format {table | json | csv | list}
```

Specifies to display the Kerberos providers in a tabular, JSON, CSV, or list format.

```
{--no-header | -a}
```

Specifies not to display the headers in the CSV or tabular formats.

```
{--no-footer | -z}
```

Specifies not to display the table summary footer information.

isi auth krb5 view

Displays the properties of an MIT Kerberos authentication provider.

Syntax

```
isi auth krb5 view <provider-name>
```

Options

```
<provider-name>
```

Specifies the Kerberos provider name.

isi auth krb5 domain create

Creates an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain create <domain>
  [--realm <string>]
```

Options

```
<domain>
```

Specifies the name of the Kerberos domain.

```
--realm <string>
```

Specifies the name of the Kerberos realm.

isi auth krb5 domain delete

Deletes an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain delete <domain>
  [--force]
```

Options

```
<domain>
```

Specifies the name of the Kerberos domain.

```
{--force | -f}
```

Specifies not to ask for a confirmation.

isi auth krb5 domain modify

Modifies an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain modify <domain>
    [--realm <string>]
```

Options

<domain>

Specifies the Kerberos domain name.

`--realm <string>`

Specifies the Kerberos realm name.

isi auth krb5 domain list

Displays a list of MIT Kerberos domain mappings.

Syntax

```
isi auth krb5 domain list
    [--limit <integer>]
    [--format {table | json | csv | list}]
    [--no-header]
    [--no-footer]
```

Options

`{--limit | -l} <integer>`

Specifies the number of Kerberos domain mappings to display.

`--format {table | json | csv | list}`

Specifies whether to display the Kerberos domain mappings in a tabular, JSON, CSV, or list formats.

`{--no-header | -a}`

Specifies not to display the headers in the CSV or tabular formats.

`{--no-footer | -z}`

Specifies not to display the table summary footer information.

isi auth krb5 domain view

Displays the properties of an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain view <domain>
```

Options

<domain>

Specifies the Kerberos domain name.

isi auth krb5 spn create

Creates or updates keys for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn create <provider-name> <user> <spn>
[--password <string>]
```

Options

<provider-name>

Specifies the Kerberos provider name.

<user>

Specifies a user name with permissions to create the service principal names (SPNs) in the Kerberos realm.

<spn>

Specifies the SPN.

`--password <string>`

Specifies the password used during the modification of a Kerberos realm.

isi auth krb5 spn delete

Deletes keys from an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn delete <provider-name> <spn> {<kvno> | --all}
```

Options

<provider-name>

Specifies the Kerberos provider name.

<spn>

Specifies the service principal name (SPN).

<kvno>

Specifies the key version number.

`--all`

Deletes all the key versions.

isi auth krb5 spn check

Checks for missing service principal names (SPNs) for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn check <provider-name>
```

Options

<provider-name>

Specifies the Kerberos provider name.

isi auth krb5 spn fix

Adds the missing service principal names (SPNs) for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn fix <provider-name> <user>
    [--password <string>]
    [--force]
```

Options

<provider-name>

Specifies the Kerberos provider name.

<user>

Specifies a user name with permissions to join clients to the given Kerberos domain.

`--password <string>`

Specifies the password that was used when modifying the Kerberos realm.

`{--force | -f}`

Specifies not to ask for a confirmation.

isi auth krb5 spn import

Imports keys from a keytab file for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn import <provider-name> <keytab-file>
```

Options

<provider-name>

Specifies the Kerberos provider name.

<keytab-file>

Specifies the keytab file to import.

isi auth krb5 spn list

Lists the service principal names (SPNs) and keys registered for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn list <provider-name>
    [--limit <integer>]
    [--format {table | json | csv | list}]
    [--no-header]
    [--no-footer]
```

Options

<provider-name>

Specifies the Kerberos provider name.

`{--limit | -l} <integer>`

Specifies the number of SPNs and keys to display.

`--format {table | json | csv | list}`

Specifies to display the SPNs and keys in a tabular, JSON, CSV, or list format.

```
{--no-header | -a}
```

Specifies not to display the headers in the CSV or tabular formats.

```
{--no-footer | -z}
```

Specifies not to display the table summary footer information.

isi auth settings krb5 modify

Modifies the global settings of an MIT Kerberos authentication provider.

Syntax

```
isi auth settings krb5 modify
  [--always-send-preauth <boolean> | --revert-always-send-
preauth]
  [--default-realm <string>]
  [--dns-lookup-kdc <boolean> | --revert-dns-lookup-kdc]
  [--dns-lookup-realm <boolean> | --revert-dns-lookup-realm]
```

Options

`--always-send-preauth <boolean>`

Specifies whether to send preauth.

`--revert-always-send-preauth`

Sets the value of `--always-send-preauth` to the system default.

`--default-realm <string>`

Specifies the default Kerberos realm name.

`--dns-lookup-kdc <boolean>`

Allows DNS to find Key Distribution Centers (KDCs).

`--revert-dns-lookup-kdc`

Sets the value of `--dns-lookup-kdc` to the system default.

`--dns-lookup-realm <boolean>`

Allows DNS to find the Kerberos realm names.

`--revert-dns-lookup-realm`

Sets the value of `--dns-lookup-realm` to the system default.

isi auth settings krb5 view

Displays MIT Kerberos provider authentication settings.

Syntax

```
isi auth settings krb5 view
```

isi auth ldap create

Creates an LDAP provider.

Syntax

```
isi auth ldap create <name>
  [--base-dn <string>]
```

```

[--server-uris <string>]
[--alternate-security-identities-attribute <string>]
[--authentication {yes | no}]
[--balance-servers {yes | no}]
[--bind-dn <string>]
[--bind-timeout <integer>]
[--cache-entry-expiry <duration>]
[--certificate-authority-file <string>]
[--check-online-interval <duration>]
[--cn-attribute <string>]
[--create-home-directory {yes | no}]
[--crypt-password-attribute <string>]
[--email-attribute <string>]
[--enabled {yes | no}]
[--enumerate-groups {yes | no}]
[--enumerate-users {yes | no}]
[--findable-groups <string>]
[--findable-users <string>]
[--gecos-attribute <string>]
[--gid-attribute <string>]
[--group-base-dn <string>]
[--group-domain <string>]
[--group-filter <string>]
[--group-members-attribute <string>]
[--group-search-scope <scope>]
[--home-directory-template <string>]
[--homedir-attribute <string>]
[--ignore-tls-errors {yes | no}]
[--listable-groups <string>]
[--listable-users <string>]
[--login-shell <string>]
[--member-of-attribute <string>]
[--name-attribute <string>]
[--netgroup-base-dn <string>]
[--netgroup-filter <string>]
[--netgroup-members-attribute <string>]
[--netgroup-search-scope <scope>]
[--netgroup-triple-attribute <string>]
[--normalize-groups {yes | no}]
[--normalize-users {yes | no}]
[--nt-password-attribute <string>]
[--ntlm-support {all | v2only | none}]
[--provider-domain <string>]
[--require-secure-connection {yes | no}]
[--restrict-findable {yes | no}]
[--restrict-listable {yes | no}]
[--search-scope <scope>]
[--search-timeout <integer>]
[--shell-attribute <string>]
[--uid-attribute <string>]
[--unfindable-groups <string>]
[--unfindable-users <string>]
[--unique-group-members-attribute <string>]
[--unlistable-groups <string>]
[--unlistable-users <string>]
[--user-base-dn <string>]
[--user-domain <string>]
[--user-filter <string>]
[--user-search-scope <scope>]
[--bind-password <string>]
[--set-bind-password]
[--verbose]

```

Options

<name>

Sets the LDAP provider name.

--base-dn <string>

- Sets the root of the tree in which to search for identities. For example, CN=myuser, CN=Users, DC=mycompany, DC=com.
- `--server-uris <string>`
Specifies a list of LDAP server URIs to be used when accessing the server. Repeat this option to specify multiple items.
- `--alternate-security-identities-attribute <string>`
Specifies the name to be used when searching for alternate security identities. This name is used when OneFS attempts to resolve a kerberos principal to a user.
- `--authentication {yes | no}`
Enables or disables the use of the provider for authentication as well as identity. The default value is yes.
- `--balance-servers {yes | no}`
Makes the provider connect to a random server on each request.
- `--bind-dn <string>`
Specifies the distinguished name to use when binding to the LDAP server. For example, CN=myuser, CN=Users, DC=mycompany, DC=com.
- `--bind-timeout <integer>`
Specifies the timeout in seconds when binding to the LDAP server.
- `--cache-entry-expiry <duration>`
Specifies the duration of time to cache a user or group, in the format `<integer>[Y | M | W | D | H | m | s]`.
- `--certificate-authority-file <path>`
Specifies the path to the root certificates file.
- `--check-online-interval <duration>`
Specifies the time between provider online checks, in the format `<integer>[Y | M | W | D | H | m | s]`.
- `--cn-attribute <string>`
Specifies the LDAP attribute that contains common names. The default value is cn.
- `--create-home-directory {yes | no}`
Specifies whether to automatically create a home directory the first time a user logs in, if a home directory does not already exist for the user.
- `--crypt-password-attribute <string>`
Specifies the LDAP attribute that contains UNIX passwords. This setting has no default value.
- `--email-attribute <string>`
Specifies the LDAP attribute that contains email addresses. The default value is mail.
- `--enabled {yes | no}`
Enables or disables the provider.
- `--enumerate-groups {yes | no}`
Specifies whether to allow the provider to enumerate groups.
- `--enumerate-users {yes | no}`

Specifies whether to allow the provider to enumerate users.

`--findable-groups <string>`

Specifies a list of groups that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable group. If populated, groups that are not included in this list cannot be resolved.

`--findable-users <string>`

Specifies a list of users that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable user. If populated, users that are not included in this list cannot be resolved.

`--gecos-attribute <string>`

Specifies the LDAP attribute that contains GECOS fields. The default value is `gecos`.

`--gid-attribute <string>`

Specifies the LDAP attribute that contains GIDs. The default value is `gidNumber`.

`--group-base-dn <string>`

Specifies the distinguished name of the entry at which to start LDAP searches for groups.

`--group-domain <string>`

Specifies the domain that the provider will use to qualify groups. The default group domain is `LDAP_GROUPS`.

`--group-filter <string>`

Sets the LDAP filter for group objects.

`--group-members-attribute <string>`

Specifies the LDAP attribute that contains group members. The default value is `memberUid`.

`--group-search-scope <scope>`

Defines the default depth from the base distinguished name (DN) to perform LDAP searches for groups.

The following values are valid:

default

Applies the setting in `--search-scope`.

Note

You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree

Searches the base DN and all entries below it.

children

Searches all entries below the base DN, excluding the base DN.

- `--home-directory-template <path>`
 Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information about home directory variables, see Home directories.
- `--homedir-attribute <string>`
 Specifies the LDAP attribute that contains home directories. The default value is `homeDirectory`.
- `--ignore-tls-errors {yes | no}`
 Continues over a secure connection even if identity checks fail.
- `--listable-groups <string>`
 Specifies a list of groups that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed.
- `--listable-users <string>`
 Specifies a list of users that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed.
- `--login-shell <path>`
 Specifies the pathname of the user's login shell for users who access the file system through SSH.
- `--member-of-attribute <string>`
 Sets the attribute to be used when searching LDAP for reverse memberships. This LDAP value should be an attribute of the user type `posixAccount` that describes the groups in which the POSIX user is a member.
- `--name-attribute <string>`
 Specifies the LDAP attribute that contains UIDs, which are used as login names. The default value is `uid`.
- `--netgroup-base-dn <string>`
 Specifies the distinguished name of the entry at which to start LDAP searches for netgroups.
- `--netgroup-filter <string>`
 Sets the LDAP filter for netgroup objects.
- `--netgroup-members-attribute <string>`
 Specifies the LDAP attribute that contains netgroup members. The default value is `memberNisNetgroup`.
- `--netgroup-search-scope <scope>`
 Defines the depth from the base distinguished name (DN) to perform LDAP searches for netgroups.
 The following values are valid:

default

Applies the setting in `--search-scope`.

Note

You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree

Searches the base DN and all entries below it.

children

Searches all entries below the base DN, excluding the base DN.

`--netgroup-triple-attribute <string>`

Specifies the LDAP attribute that contains netgroup triples. The default value is `nisNetgroupTriple`.

`--normalize-groups {yes | no}`

Normalizes group names to lowercase before lookup.

`--normalize-users {yes | no}`

Normalizes user names to lowercase before lookup.

`--nt-password-attribute <string>`

Specifies the LDAP attribute that contains Windows passwords. The default value is `ntpasswdhash`.

`--ntlm-support {all | v2only | none}`

For users with NTLM-compatible credentials, specifies which NTLM versions to support.

`--provider-domain <string>`

Specifies the domain that the provider will use to qualify user and group names.

`--require-secure-connection {yes | no}`

Specifies whether to require a TLS connection.

`--restrict-findable {yes | no}`

Specifies whether to check the provider for filtered lists of findable and unfindable users and groups.

`--restrict-listable {yes | no}`

Specifies whether to check the provider for filtered lists of listable and unlistable users and groups.

`--search-scope <scope>`

Defines the default depth from the base distinguished name (DN) to perform LDAP searches.

The following values are valid:

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree

Searches the base DN and all entries below it.

children

Searches all entries below the base DN, excluding the base DN itself.

`--search-timeout <integer>`

Specifies the number of seconds after which to stop retrying and fail a search. The default value is 100.

`--shell-attribute <string>`

Specifies the LDAP attribute that contains a user's UNIX login shell. The default value is `loginShell`.

`--uid-attribute <string>`

Specifies the LDAP attribute that contains UID numbers. The default value is `uidNumber`.

`--unfindable-groups <string>`

If `--restrict-findable` is enabled and the findable groups list is empty, specifies a list of groups that cannot be resolved by this provider. Repeat this option to specify multiple list items.

`--unfindable-users <string>`

If `--restrict-findable` is enabled and the findable users list is empty, specifies a list of users that cannot be resolved by this provider. Repeat this option to specify multiple list items.

`--unique-group-members-attribute <string>`

Specifies the LDAP attribute that contains unique group members. This attribute is used to determine which groups a user belongs to if the LDAP server is queried by the user's DN instead of the user's name. This setting has no default value.

`--unlistable-groups <string>`

If `--restrict-listable` is enabled and the listable groups list is empty, specifies a list of groups that cannot be listed by this provider that cannot be viewed. Repeat this option to specify multiple list items.

`--unlistable-users <string>`

If `--restrict-listable` is enabled and the listable users list is empty, specifies a list of users that cannot be listed by this provider that cannot be viewed. Repeat this option to specify multiple list items.

`--user-base-dn <string>`

Specifies the distinguished name of the entry at which to start LDAP searches for users.

`--user-domain <string>`

Specifies the domain that the provider will use to qualify users. The default user domain is `LDAP_USERS`.

`--user-filter <string>`

Sets the LDAP filter for user objects.

`--user-search-scope <scope>`

Defines the depth from the base distinguished name (DN) to perform LDAP searches for users.

The following values are valid:

default

Applies the search scope that is defined in the default query settings.

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree

Searches the base DN and all entries below it.

children

Searches all entries below the base DN, excluding the base DN itself.

`--bind-password <string>`

Sets the password for the distinguished name that is used when binding to the LDAP server. To set the password interactively, use the `--set-bind-password` option instead.

`--set-bind-password`

Interactively sets the password for the distinguished name that is used when binding to the LDAP server. This option cannot be used with `--bind-password`.

`{--verbose | -v}`

Displays the results of running the command.

isi auth ldap delete

Deletes an LDAP provider.

Syntax

```
isi auth ldap delete <provider-name>
  [--force]
  [--verbose]
```

Options

`<provider-name>`

Specifies the name of the provider to delete.

`{--force | -f}`

Suppresses command-line prompts and messages.

`<provider-name>`

Specifies the name of the provider to delete.

`{--verbose | -v}`

Displays more detailed information.

isi auth ldap list

Displays a list of LDAP providers.

Syntax

```
isi auth ldap list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth ldap modify

Modifies an LDAP provider.

Syntax

```
isi auth ldap modify <provider-name>
  [--base-dn <string>]
  [--server-uris <string>]
  [--clear-server-uris]
  [--add-server-uris <string>]
  [--remove-server-uris <string>]
  [--alternate-security-identities-attribute <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--bind-dn <string>]
  [--bind-timeout <integer>]
  [--cache-entry-expiry <duration>]
  [--certificate-authority-file <string>]
  [--check-online-interval <duration>]
  [--cn-attribute <string>]
  [--create-home-directory {yes | no}]
  [--crypt-password-attribute <string>]
  [--email-attribute <string>]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--clear-findable-groups]
  [--add-findable-groups <string>]
  [--remove-findable-groups <string>]
```

```

[--findable-users <string>]
[--clear-findable-users]
[--add-findable-users <string>]
[--remove-findable-users <string>]
[--gecos-attribute <string>]
[--gid-attribute <string>]
[--group-base-dn <string>]
[--group-domain <string>]
[--group-filter <string>]
[--group-members-attribute <string>]
[--group-search-scope <scope>]
[--homedir-attribute <string>]
[--home-directory-template <string>]
[--ignore-tls-errors {yes | no}]
[--listable-groups <string>]
[--clear-listable-groups]
[--add-listable-groups <string>]
[--remove-listable-groups <string>]
[--listable-users <string>]
[--clear-listable-users]
[--add-listable-users <string>]
[--remove-listable-users <string>]
[--login-shell <string>]
[--member-of-attribute <string>]
[--name-attribute <string>]
[--netgroup-base-dn <string>]
[--netgroup-filter <string>]
[--netgroup-members-attribute <string>]
[--netgroup-search-scope <scope>]
[--netgroup-triple-attribute <string>]
[--normalize-groups {yes | no}]
[--normalize-users {yes | no}]
[--nt-password-attribute <string>]
[--ntlm-support {all | v2only | none}]
[--provider-domain <string>]
[--require-secure-connection {yes | no}]
[--restrict-findable {yes | no}]
[--restrict-listable {yes | no}]
[--search-scope <scope>]
[--search-timeout <integer>]
[--shell-attribute <string>]
[--uid-attribute <string>]
[--unfindable-groups <string>]
[--clear-unfindable-groups]
[--add-unfindable-groups <string>]
[--remove-unfindable-groups <string>]
[--unfindable-users <string>]
[--clear-unfindable-users]
[--add-unfindable-users <string>]
[--remove-unfindable-users <string>]
[--unique-group-members-attribute <string>]
[--unlistable-groups <string>]
[--clear-unlistable-groups]
[--add-unlistable-groups <string>]
[--remove-unlistable-groups <string>]
[--unlistable-users <string>]
[--clear-unlistable-users]
[--add-unlistable-users <string>]
[--remove-unlistable-users <string>]
[--user-base-dn <string>]
[--user-domain <string>]
[--user-filter <string>]
[--user-search-scope <scope>]
[--bind-password <string>]
[--set-bind-password]
[--verbose]

```

Options

- <provider-name>*
Specifies the name of the LDAP provider to modify.
- `--base-dn <string>`
Sets the root of the tree in which to search for identities. For example, `CN=myuser, CN=Users, DC=mycompany, DC=com`.
- `--server-uris <string>`
Specifies a list of LDAP server URIs to be used when accessing the server. Repeat this option to specify multiple items.
- `--clear-server-uris`
Removes all entries from the list of server URIs.
- `--add-server-uris <string>`
Adds an entry to the list of server URIs. Repeat this option to specify multiple list items.
- `--remove-server-uris <string>`
Removes an entry from the list of server URIs. Repeat this option to specify multiple list items.
- `--alternate-security-identities-attribute <string>`
Specifies the name to be used when searching for alternate security identities. This name is used when OneFS attempts to resolve a kerberos principal to a user.
- `--authentication {yes | no}`
Enables or disables the use of this provider for authentication as well as identity. The default value is `yes`.
- `--balance-servers {yes | no}`
Makes this provider connect to a random server on each request.
- `--bind-dn <string>`
Specifies the distinguished name to use when binding to the LDAP server. For example, `CN=myuser, CN=Users, DC=mycompany, DC=com`.
- `--bind-timeout <integer>`
Specifies the timeout in seconds when binding to the LDAP server.
- `--cache-entry-expiry <duration>`
Specifies the amount of time to cache a user or group, in the format `<integer>[Y | M | W | D | H | m | s]`.
- `--certificate-authority-file <path>`
Specifies the path to the root certificates file.
- `--check-online-interval <duration>`
Specifies the time between provider online checks, in the format `<integer>[Y | M | W | D | H | m | s]`.
- `--cn-attribute <string>`
Specifies the LDAP attribute that contains common names. The default value is `cn`.
- `--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user. The directory path is specified in the path template through the `--home-directory-template` command.

`--crypt-password-attribute <string>`

Specifies the LDAP attribute that contains UNIX passwords. This setting has no default value.

`--email-attribute <string>`

Specifies the LDAP attribute that contains email addresses. The default value is `mail`.

`--enabled {yes | no}`

Enables or disables this provider.

`--enumerate-groups {yes | no}`

Specifies whether to allow the provider to enumerate groups.

`--enumerate-users {yes | no}`

Specifies whether to allow the provider to enumerate users.

`--findable-groups <string>`

Specifies a list of groups that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be resolved in this provider. This option overwrites the entries in the findable groups list; to add or remove groups without affecting current entries, use `--add-findable-groups` or `--remove-findable-groups`.

`--clear-findable-groups`

Removes the list of findable groups.

`--add-findable-groups <string>`

Adds an entry to the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-findable-groups <string>`

Removes an entry from the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--findable-users <string>`

Specifies a list of users that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be resolved in this provider. This option overwrites the entries in the findable users list; to add or remove users without affecting current entries, use `--add-findable-users` or `--remove-findable-users`.

`--clear-findable-users`

Removes the list of findable users.

`--add-findable-users <string>`

Adds an entry to the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-findable-users <string>`

Removes an entry from the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--gecos-attribute <string>`

Specifies the LDAP attribute that contains GECOS fields. The default value is `gecos`.

`--gid-attribute <string>`

Specifies the LDAP attribute that contains GIDs. The default value is `gidNumber`.

`--group-base-dn <string>`

Specifies the distinguished name of the entry at which to start LDAP searches for groups.

`--group-domain <string>`

Specifies the domain that this provider will use to qualify groups. The default group domain is `LDAP_GROUPS`.

`--group-filter <string>`

Sets the LDAP filter for group objects.

`--group-members-attribute <string>`

Specifies the LDAP attribute that contains group members. The default value is `memberUid`.

`--group-search-scope <scope>`

Defines the default depth from the base distinguished name (DN) to perform LDAP searches for groups.

The following values are valid:

default

Applies the setting in `--search-scope`.

Note

You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree

Searches the base DN and all entries below it.

children

Searches all entries below the base DN, excluding the base DN.

`--home-directory-template <path>`

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

`--homedir-attribute <string>`

Specifies the LDAP attribute that is used when searching for the home directory. The default value is `homeDirectory`.

`--ignore-tls-errors {yes | no}`

Continues over a secure connection even if identity checks fail.

`--listable-groups <string>`

Specifies a list of groups that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed in this provider. This option overwrites the entries in the listable groups list; to add or remove groups without affecting current entries, use `--add-listable-groups` or `--remove-listable-groups`.

`--clear-listable-groups`

Removes all entries from the list of viewable groups.

`--add-listable-groups <string>`

Adds an entry to the list of listable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-listable-groups <string>`

Removes an entry from the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--listable-users <string>`

Specifies a list of users that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed in this provider. This option overwrites the entries in the listable users list; to add or remove users without affecting current entries, use `--add-listable-users` or `--remove-listable-users`.

`--clear-listable-users`

Removes all entries from the list of viewable users.

`--add-listable-users <string>`

Adds an entry to the list of listable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-listable-users <string>`

Removes an entry from the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--login-shell <path>`

Specifies the pathname to the user's login shell, for users who access the file system through SSH.

`--member-of-attribute <string>`

Sets the attribute to be used when searching LDAP for reverse memberships. This LDAP value should be an attribute of the user type `posixAccount` that describes the groups in which the POSIX user is a member.

`--name-attribute <string>`

Specifies the LDAP attribute that contains UIDs, which are used as login names. The default value is `uid`.

`--netgroup-base-dn <string>`
 Specifies the distinguished name of the entry at which to start LDAP searches for netgroups.

`--netgroup-filter <string>`
 Sets the LDAP filter for netgroup objects.

`--netgroup-members-attribute <string>`
 Specifies the LDAP attribute that contains netgroup members. The default value is `memberNisNetgroup`.

`--netgroup-search-scope <scope>`
 Defines the depth from the base distinguished name (DN) to perform LDAP searches for netgroups.
 The following values are valid:

default
 Applies the setting in `--search-scope`.

Note

You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base
 Searches only the entry at the base DN.

onelevel
 Searches all entries exactly one level below the base DN.

subtree
 Searches the base DN and all entries below it.

children
 Searches all entries below the base DN, excluding the base DN.

`--netgroup-triple-attribute <string>`
 Specifies the LDAP attribute that contains netgroup triples. The default value is `nisNetgroupTriple`.

`--normalize-groups {yes | no}`
 Normalizes group names to lowercase before lookup.

`--normalize-users {yes | no}`
 Normalizes user names to lowercase before lookup.

`--nt-password-attribute <string>`
 Specifies the LDAP attribute that contains Windows passwords. The default value is `ntpasswdhash`.

`--ntlm-support {all | v2only | none}`
 For users with NTLM-compatible credentials, specifies which NTLM versions to support.
 The following values are valid:

`all`

`v2only`


```

    none
--provider-domain <string>
    Specifies the domain that this provider will use to qualify user and group names.
--require-secure-connection {yes | no}
    Specifies whether to require a TLS connection.
--restrict-findable {yes | no}
    Specifies whether to check this provider for filtered lists of findable and unfindable
    users and groups.
--restrict-listable {yes | no}
    Specifies whether to check this provider for filtered lists of viewable and unviewable
    users and groups.
--search-scope <scope>
    Defines the default depth from the base distinguished name (DN) to perform LDAP
    searches.
    The following values are valid:
base
    Searches only the entry at the base DN.
onelevel
    Searches all entries exactly one level below the base DN.
subtree
    Searches the base DN and all entries below it.
children
    Searches all entries below the base DN, excluding the base DN itself.
--search-timeout <integer>
    Specifies the number of seconds after which to stop retrying and fail a search. The
    default value is 100.
--shell-attribute <string>
    Specifies the LDAP attribute that is used when searching for a user's UNIX login shell.
    The default value is loginShell.
--uid-attribute <string>
    Specifies the LDAP attribute that contains UID numbers. The default value is
    uidNumber.
--unfindable-groups <string>
    Specifies a group that cannot be found in this provider if --restrict-findable
    is enabled. Repeat this option to specify multiple list items. This option overwrites
    the entries in the unfindable groups list; to add or remove groups without affecting
    current entries, use --add-unfindable-groups or --remove-unfindable-
    groups.
--clear-unfindable-groups
    Removes all entries from the list of unfindable groups.
--add-unfindable-groups <string>
    Adds an entry to the list of unfindable groups that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--remove-unfindable-groups <string>

```

Removes an entry from the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--unfindable-users <string>`

Specifies a user that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable users list; to add or remove users without affecting current entries, use `--add-unfindable-users` or `--remove-unfindable-users`.

`--clear-unfindable-users`

Removes all entries from the list of unfindable groups.

`--add-unfindable-users <string>`

Adds an entry to the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-unfindable-users <string>`

Removes an entry from the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--unique-group-members-attribute <string>`

Specifies the LDAP attribute that contains unique group members. This attribute is used to determine which groups a user belongs to if the LDAP server is queried by the user's DN instead of the user's name. This setting has no default value.

`--unlistable-groups <string>`

Specifies a group that cannot be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable groups list; to add or remove groups without affecting current entries, use `--add-unlistable-groups` or `--remove-unlistable-groups`.

`--clear-unlistable-groups`

Removes all entries from the list of unviewable groups.

`--add-unlistable-groups <string>`

Adds an entry to the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-unlistable-groups <string>`

Removes an entry from the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--unlistable-users <string>`

Specifies a user that cannot be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable users list; to add or remove users without affecting current entries, use `--add-unlistable-users` or `--remove-unlistable-users`.

`--clear-unlistable-users`

Removes all entries from the list of unviewable users.

`--add-unlistable-users <string>`

Adds an entry to the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-unlistable-users <string>`
 Removes an entry from the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--user-base-dn <string>`
 Specifies the distinguished name of the entry at which to start LDAP searches for users.

`--user-domain <string>`
 Specifies the domain that this provider will use to qualify users. The default user domain is `LDAP_USERS`.

`--user-filter <string>`
 Sets the LDAP filter for user objects.

`--user-search-scope <scope>`
 Defines the depth from the base distinguished name (DN) to perform LDAP searches for users. The valid values are as follows:
 The following values are valid:

default
 Applies the setting in `--search-scope`.

Note

You cannot specify `--search-scope=default`. For example, if you specify `--user-search-scope=default`, the search scope is set to the value of `--search-scope`.

base
 Searches only the entry at the base DN.

onelevel
 Searches all entries exactly one level below the base DN.

subtree
 Searches the base DN and all entries below it.

children
 Searches all entries below the base DN, excluding the base DN.

`--bind-password <string>`
 Sets the password for the distinguished name that is used when binding to the LDAP server. To set the password interactively, use the `--set-bind-password` option instead.

`--set-bind-password`
 Interactively sets the password for the distinguished name that is used when binding to the LDAP server. This option cannot be used with `--bind-password`.

`{--verbose | -v}`
 Displays the results of running the command.

isi auth ldap view

Displays the properties of an LDAP provider.

Syntax

```
isi auth ldap view <provider-name>
```

Options

<provider-name>

Specifies the name of the provider to view.

isi auth local list

Displays a list of local providers.

Syntax

```
isi auth local list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} *<integer>*

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth local view

Displays the properties of a local provider.

Syntax

```
isi auth local view <provider-name>
```

Options

<provider-name>

Specifies the name of the provider to view.

isi auth local modify

Modifies a local provider.

Syntax

```
isi auth local modify <provider-name>
  [--authentication {yes | no}]
  [--create-home-directory {yes | no}]
  [--home-directory-template <string>]
  [--lockout-duration <duration>]
  [--lockout-threshold <integer>]
  [--lockout-window <duration>]
  [--login-shell <string>]
  [--machine-name <string>]
  [--min-password-age <duration>]
  [--max-password-age <duration>]
  [--min-password-length <integer>]
  [--password-prompt-time <duration>]
  [--password-complexity{lowercase | uppercase |
    numeric | symbol}]
  [--clear-password-complexity]
  [--add-password-complexity {lowercase | uppercase |
    numeric | symbol}]
  [--remove-password-complexity <string>]
  [--password-history-length <integer>]
  [--verbose]
```

Options

<provider-name>

Specifies the name of the local provider to modify.

`--authentication {yes | no}`

Uses the provider for authentication as well as identity. The default setting is `yes`.

`--create-home-directory {yes | no}`

Creates a home directory the first time a user logs in.

`--home-directory-template <string>`

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

`--lockout-duration <duration>`

Sets the length of time that an account will be inaccessible after multiple failed login attempts.

`--lockout-threshold <integer>`

Specifies the number of failed login attempts after which an account will be locked out.

`--lockout-window <duration>`

Sets the time in which the number of failed attempts specified by the `--lockout-threshold` option must be made for an account to be locked out. Duration is specified in the format `<integer>{[Y | M | W | D | H | m | s]}`.

`--login-shell <string>`

Specifies the path to the UNIX login shell.

`--machine-name <string>`
Specifies the domain to use to qualify user and group names for the provider.

`--min-password-age <duration>`
Sets the minimum password age, in the format `<integer>{Y | M | W | D | H | m | s}`.

`--max-password-age <duration>`
Sets the maximum password age, in the format `<integer>{Y | M | W | D | H | m | s}`.

`--min-password-length <integer>`
Sets the minimum password length.

`--password-prompt-time <duration>`
Sets the remaining time until a user is prompted for a password change, in the format `<integer>{Y | M | W | D | H | m | s}`.

`--password-complexity {lowercase | uppercase | numeric | symbol}`
Specifies the conditions that a password is required to meet. A password must contain at least one character from each specified option to be valid. For example, if `lowercase` and `numeric` are specified, a password must contain at least one lowercase character and one digit to be valid. Symbols are valid, excluding `#` and `@`.

`--clear-password-complexity`
Clears the list of parameters against which to validate new passwords.

`--add-password-complexity {lowercase | uppercase | numeric | symbol}`
Adds items to the list of parameters against which to validate new passwords. Repeat this command to specify additional password-complexity options.

`--remove-password-complexity <string>`
Removes items from the list of parameters against which to validate new passwords. Repeat this command to specify each password-complexity option that you want to remove.

`--password-history-length <integer>`
Specifies the number of previous passwords to store to prevent reuse of a previous password. The max password history length is 24.

`{--verbose | -v}`
Displays more detailed information.

isi auth log-level

Displays or modifies the run-time log level of the authentication service.

Syntax

```
isi auth log-level [--set <string>]
```

Options

`{--set | -s} <string>`

Sets the log level for the current node. The log level determines how much information is logged.

The following values are valid and are organized from least to most information:

- always
- error
- warning
- info
- verbose
- debug
- trace

Note

Levels **verbose**, **debug**, and **trace** may cause performance issues. Levels **debug** and **trace** log information that likely will be useful only when consulting EMC Isilon Technical Support.

Examples

To set the log level to **debug**, run the following command:

```
isi auth log-level --set=debug
```

isi auth mapping delete

Deletes one or more identity mappings.

Syntax

```
isi auth mapping delete {<source>| --source-uid <integer>
| --source-gid <integer> | --source-sid <string> | --all}
[{{--only-generated | --only-external | --2way | --target <string>
| --target-uid <integer> | --target-gid <integer> | --target-sid
<string>}}]
[--zone<string>]
```

Options

<source>

Specifies the mapping source by identity type, in the format *<type>.<value>*—for example, **UID:2002**.

--source-uid *<integer>*

Specifies the mapping source by UID.

--source-gid *<integer>*

Specifies the mapping source by GID.

--source-sid *<string>*

Specifies the mapping source by SID.

--all

Deletes all identity mappings in the specified access zone. Can be used in conjunction with **--only-generated** and **--only-external** for additional filtering.

--only-generated

Only deletes identity mappings that were created automatically and that include a generated UID or GID from the internal range of user and group IDs. Must be used in conjunction with **--all**.

--only-external

Only deletes identity mappings that were created automatically and that include a UID or GID from an external authentication source. Must be used in conjunction with `--all`.

`--2way`

Specifies or deletes a two-way, or reverse, mapping.

`--target <string>`

Specifies the mapping target by identity type, in the format `<type>:<value>`—for example, `UID:2002`.

`--target-uid <integer>`

Specifies the mapping target by UID.

`--target-gid <integer>`

Specifies the mapping target by GID.

`--target-sid <string>`

Specifies the mapping target by SID.

`--zone <string>`

Deletes identity mappings in the specified access zone. If no access zone is specified, mappings are deleted from the default System zone.

isi auth mapping dump

Displays or prints the kernel mapping database.

Syntax

```
isi auth mapping dump
  [--file <path>]
  [--zone <string>]
```

Options

If no option is specified, the full kernel mapping database is displayed.

`{--file | -f} <path>`

Prints the database to the specified output file.

`--zone <string>`

Displays the database from the specified access zone. If no access zone is specified, displays all mappings.

Examples

To view the kernel mapping database, run the following command:

```
isi auth mapping dump
```

The system displays output similar to the following example:

```
["ZID:1", "UID:6299", [{"SID:S-1-5-21-1195855716-1407", 128}]]
["ZID:1", "GID:1000000", [{"SID:S-1-5-21-1195855716-513", 48}]]
["ZID:1", "SID:S-1-5-21-1195855716-1407", [{"UID:6299", 144}]]
["ZID:1", "SID:S-1-5-21-1195855716-513", [{"GID:1000000", 32}]]
```


isi auth mapping flush

Flushes the cache for one or all identity mappings. Flushing the cache might be useful if the ID mapping rules have been modified.

Syntax

```
isi auth mapping flush {--all | --source <string>
  | --source-uid <integer> | --source-gid <integer>
  | --source-sid <string>}
[--zone<string>]
```

Options

You must specify either `--all` or one of the source options.

`--all`

Flushes all identity mappings on the EMC Isilon cluster.

`--source <string>`

Specifies the mapping source by identity type, in the format `<type>:<value>`—for example, `UID:2002`.

`--source-uid <integer>`

Specifies the source identity by UID.

`--source-gid <integer>`

Specifies the source identity by GID.

`--source-sid <string>`

Specifies the source identity by SID.

`--zone<string>`

Specifies the access zone of the source identity. If no access zone is specified, any mapping for the specified source identity is flushed from the default System zone.

isi auth mapping idrange

Displays or modifies the range that UIDs and GIDs are generated from.

Syntax

```
isi auth mapping idrange {--set-uid-low <integer>
  | --set-uid-high <integer> | --set-uid-hwm <integer>
  | --set-gid-low <integer> | --set-gid-high <integer>
  | --set-gid-hwm <integer> | --get-uid-range | --get-gid-range}...
```

Options

Note

When modifying a UID or GID range, make sure that your settings meet the following requirements:

- Existing IDs are not included
 - A mapping does not overlap with another range that might be used by other IDs on the cluster
 - The mapping is large enough to avoid running out of unused IDs; if all IDs in the range are in use, ID allocation will fail.
-

`--set-uid-low <integer>`

Sets the lowest UID value in the range.

`--set-uid-high <integer>`

Sets the highest UID value in the range.

`--set-uid-hwm <integer>`

Specifies the next UID that will be allocated (the high water mark).

Note

- If the high water mark is set to more than the high UID value, UID allocation will fail.
 - The high water mark cannot be set to less than the lowest UID value in the range. If the specified `<integer>` value is less than the low UID value, the high water mark is set to the low UID value.
-

`--set-gid-low <integer>`

Sets the lowest GID value in the range.

`--set-gid-high <integer>`

Sets the highest GID value in the range.

`--set-gid-hwm <integer>`

Specifies the next GID that will be used (the high water mark).

Note

- If the high water mark is set to more than the high GID value, GID allocation will fail.
 - High water mark cannot be set to less than the lowest GID value in the range. If the specified `<integer>` value is less than the low GID value, high water mark is set to the low GID value.
-

`--get-uid-range`

Displays the current UID range.

`--get-gid-range`

Displays the current GID range.

isi auth mapping import

Imports mappings from a source file to the ID mapping database.

Syntax

```
isi auth mapping import --file <path>
  [--overwrite]
```

Options

{--file | -f} <path>

Specifies the full path to the file to import. File content must be in the same format as the output that is displayed by running the `isi auth mapping dump` command.

{--overwrite | -o}

Overwrites existing entries in the mapping database file.

isi auth mapping view

Displays mappings for an identity.

Syntax

```
isi auth mapping view {<id>| --uid <integer>
  | --gid <integer> | --sid <string>}
  [--nocreate]
  [--zone <string>]
```

Options

<id>

Specifies the ID of the source identity type in the format `<type>:<value>`—for example, `UID:2002`.

--uid <integer>

Specifies the mapping source by UID.

--gid <integer>

Specifies the mapping source by GID.

--sid <string>

Specifies the mapping source by SID.

--nocreate

Specifies that nonexistent mappings should not be created.

--zone

Specifies the access zone of the source identity. If no access zone is specified, OneFS displays mappings from the default System zone.

Examples

The following command displays mappings for a user whose UID is 2002 in the zone3 access zone:

```
isi auth mapping view uid:2002 --zone=zone3
```

The system displays output similar to the following example:

Type	Mapping
Name	test1
On-disk	UID:2002
Unix UID	2002
Unix GID	None
SMB	S-1-5-21-1776575851-2890035977-2418728619-1004
NFSv4	test1

isi auth mapping modify

Sets or modifies a mapping between two identities.

Syntax

```
isi auth mapping modify {<source>| --source-uid <integer>
| --source-gid <integer> | --source-sid <string> | --target
<string>
| --target-uid <integer> | --target-gid <string> | --
target-sid <string>}
[--on-disk]
[--2way]
[--zone<string>]
```

Options

<source>

Specifies the mapping source by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--source-uid <integer>

Specifies the mapping source by UID.

--source-gid <integer>

Specifies the mapping source by GID.

--source-sid <string>

Specifies the mapping source by SID.

--target <string>

Specifies the mapping target by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--target-uid <integer>

Specifies the mapping target by UID.

--target-gid <integer>

Specifies the mapping target by GID.

--target-sid <string>

Specifies the mapping target by SID.

--on-disk

Specifies that the source on-disk identity should be represented by the target identity.

--2way

Specifies a two-way, or reverse, mapping.

--zone<string>

Specifies the access zone that the ID mapping is applied to. If no access zone is specified, the mapping is applied to the default System zone.

isi auth mapping create

Creates a manual mapping between a source identity and target identity or automatically generates a mapping for a source identity.

Syntax

```
isi auth mapping create {<source>| --source-uid <integer>
  | --source-gid <integer> | --source-sid <string>}
[--uid | --gid | --sid]
[--on-disk]
[--2way]
[--target <string> | --target-uid <integer>
  | --target-gid <string> | --target-sid <string>}]
[--zone<string>]
```

Options

<source>

Specifies the mapping source by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--source-uid <integer>

Specifies the mapping source by UID.

--source-gid <integer>

Specifies the mapping source by GID.

--source-sid <string>

Specifies the mapping source by SID.

--uid

Generates a mapping if one does not exist for the identity; otherwise, retrieves the mapped UID.

--gid

Generates a mapping if one does not exist for the identity; otherwise, retrieves the mapped GID.

--sid

Generates a mapping if one does not exist for the identity; otherwise, retrieves the mapped SID.

--on-disk

Specifies that the source on-disk identity should be represented by the target identity.

--2way

Specifies a two-way, or reverse, mapping.

--target <string>

Specifies the mapping target by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--target-uid <integer>

Specifies the mapping target by UID.

--target-gid <integer>

Specifies the mapping target by GID.

`--target-sid <string>`

Specifies the mapping target by SID.

`--zone <string>`

Specifies the access zone that the ID mapping is applied to. If no access zone is specified, the mapping is applied to the default System zone.

isi auth mapping token

Displays the access token that is calculated for a user during authentication.

Syntax

```
isi auth mapping token {<user> | --uid <integer>
  | --kerberos-principal <string>}
  [--zone <string>]
  [--primary-gid <integer>]
  [--gid <integer>]
```

Options

This command requires `<user>` or `--uid <integer>` or `--kerberos-principal <string>`.

`<user>`

Specifies the user by name.

`--uid <integer>`

Specifies the user by UID.

`--kerberos-principal <string>`

Specifies the Kerberos principal by name. For example, `user@realm.com`.

`--zone <string>`

Specifies the name of the access zone that contains the mapping.

`--primary-gid <integer>`

Specifies the primary GID.

`--gid <integer>`

Specifies a token GID. Repeat this option to specify multiple GIDs.

isi auth netgroups list

Displays information about a netgroup.

Syntax

```
isi auth netgroups list --netgroup <string>
  [--recursive]
  [--ignore]
  [--raw]
```

Options

`--netgroup <string>`

Specifies the name of a netgroup.

`--recursive`

Recursively resolves nested netgroups.

`--ignore`
Ignores errors and unresolvable netgroups.

`--raw`
Displays raw netgroup information.

isi auth nis create

Creates an NIS provider.

Syntax

```
isi auth nis create <name>
  [--nis-domain <string>]
  [--servers <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--cache-entry-expiry <duration>]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--findable-users <string>]
  [--group-domain <string>]
  [--home-directory-template <path>]
  [--hostname-lookup {yes | no}]
  [--listable-groups <string>]
  [--listable-users <string>]
  [--login-shell <path>]
  [--normalize-groups {yes | no}]
  [--normalize-users {yes | no}]
  [--provider-domain <string>]
  [--ntlm-support {all | v2only | none}]
  [--request-timeout <integer>]
  [--restrict-findable {yes | no}]
  [--restrict-listable {yes | no}]
  [--retry-time <integer>]
  [--unfindable-groups <string>]
  [--unfindable-users <string>]
  [--unlistable-groups <string>]
  [--unlistable-users <string>]
  [--user-domain <string>]
  [--ypmatch-using-tcp {yes | no}]
  [--verbose]
```

Options

`<name>`

Sets the name of the NIS provider.

`--nis-domain <string>`

Specifies the NIS domain name.

`--servers <string>`

Specifies a list of NIS servers to be used by this provider. Repeat this option to specify multiple list items.

`--authentication {yes | no}`

Enables or disables the use of the provider for authentication as well as identity. The default value is `yes`.

`--balance-servers {yes | no}`

Makes the provider connect to a random server on each request.

`--cache-entry-expiry <duration>`

Specifies amount of time to cache a user or group, in the format *<integer>{Y | M | W | D | H | m | s}*.

`--check-online-interval <duration>`

Specifies the time between provider online checks, in the format *<integer>{Y | M | W | D | H | m | s}*.

`--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

`--enabled {yes | no}`

Enables or disables the provider.

`--enumerate-groups {yes | no}`

Specifies whether to allow the provider to enumerate groups.

`--enumerate-users {yes | no}`

Specifies whether to allow the provider to enumerate users.

`--findable-groups <string>`

Specifies a group that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be resolved.

`--findable-users <string>`

Specifies a user that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be resolved.

`--group-domain <string>`

Specifies the domain that this provider will use to qualify groups. The default group domain is `NIS_GROUPS`.

`--home-directory-template <path>`

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

`--hostname-lookup {yes | no}`

Enables or disables host name lookups.

`--listable-groups <string>`

Specifies a group that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed.

`--listable-users <string>`

Specifies a user that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed.


```
--login-shell <path>
    Specifies the path to the user's login shell. This setting applies only to users who
    access the file system through SSH.
--normalize-groups {yes | no}
    Normalizes group name to lowercase before lookup.
--normalize-users {yes | no}
    Normalizes user name to lowercase before lookup.
--provider-domain <string>
    Specifies the domain that this provider will use to qualify user and group names.
--ntlm-support {all | v2only | none}
    For users with NTLM-compatible credentials, specifies which NTLM versions to
    support. Valid values are all, v2only, and none. NTLMv2 provides additional
    security over NTLM.
--request-timeout <integer>
    Specifies the request timeout interval in seconds.
--restrict-findable {yes | no}
    Specifies whether to check this provider for filtered lists of findable and unfindable
    users and groups.
--restrict-listable {yes | no}
    Specifies whether to check this provider for filtered lists of viewable and unviewable
    users and groups.
--retry-time <integer>
    Sets the timeout period in seconds after which a request will be retried.
--unfindable-groups <string>
    If --restrict-findable is enabled and the findable groups list is empty,
    specifies a group that cannot be resolved by this provider. Repeat this option to
    specify multiple list items.
--unfindable-users <string>
    If --restrict-findable is enabled and the findable users list is empty, specifies
    a user that cannot be resolved by this provider. Repeat this option to specify multiple
    list items.
--unlistable-groups <string>
    If --restrict-listable is enabled and the listable groups list is empty,
    specifies a group that cannot be viewed by this provider. Repeat this option to specify
    multiple list items.
--unlistable-users <string>
    If --restrict-listable is enabled and the listable users list is empty, specifies
    a user that cannot be viewed by this provider. Repeat this option to specify multiple
    list items.
--user-domain <string>
    Specifies the domain that this provider will use to qualify users. The default user
    domain is NIS_USERS.
--ypmatch-using-tcp {yes | no}
    Uses TCP for YP Match operations.
```

```
{--verbose | -v}
```

Displays the results of running the command.

isi auth nis delete

Deletes an NIS provider.

Syntax

```
isi auth nis delete <provider-name>
  [--force]
  [--verbose]
```

Options

<provider-name>

Specifies the name of the provider to delete.

```
{--force | -f}
```

Suppresses command-line prompts and messages.

```
{--verbose | -v}
```

Returns a success or fail message after running the command.

isi auth nis list

Displays a list of NIS providers and indicates whether a provider is functioning correctly.

Syntax

```
isi auth nis list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi auth nis modify

Modifies an NIS provider.

Syntax

```
isi auth nis modify <provider-name>
  [--nis-domain <string>]
  [--servers <string>]
  [--clear-servers]
  [--add-servers <string>]
  [--remove-servers <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--cache-entry-expiry <duration>]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--clear-findable-groups]
  [--add-findable-groups <string>]
  [--remove-findable-groups <string>]
  [--findable-users <string>]
  [--clear-findable-users]
  [--add-findable-users <string>]
  [--remove-findable-users <string>]
  [--group-domain <string>]
  [--home-directory-template <string>]
  [--hostname-lookup {yes | no}]
  [--listable-groups <string>]
  [--clear-listable-groups]
  [--add-listable-groups <string>]
  [--remove-listable-groups <string>]
  [--listable-users <string>]
  [--clear-listable-users]
  [--add-listable-users <string>]
  [--remove-listable-users <string>]
  [--login-shell <string>]
  [--normalize-groups {yes | no}]
  [--normalize-users {yes | no}]
  [--provider-domain <string>]
  [--ntlm-support {all | v2only | none}]
  [--request-timeout <integer>]
  [--restrict-findable {yes | no}]
  [--restrict-listable {yes | no}]
  [--retry-time <integer>]
  [--unfindable-groups <string>]
  [--clear-unfindable-groups]
  [--add-unfindable-groups <string>]
  [--remove-unfindable-groups <string>]
  [--unfindable-users <string>]
  [--clear-unfindable-users]
  [--add-unfindable-users <string>]
  [--remove-unfindable-users <string>]
  [--unlistable-groups <string>]
  [--clear-unlistable-groups]
  [--add-unlistable-groups <string>]
  [--remove-unlistable-groups <string>]
  [--unlistable-users <string>]
  [--clear-unlistable-users]
  [--add-unlistable-users <string>]
  [--remove-unlistable-users <string>]
  [--user-domain <string>]
  [--ypmatch-using-tcp {yes | no}]
  [--verbose]
```

Options*<provider-name>*

Specifies the name of the NIS provider to modify.

`--nis-domain <string>`

Specifies the NIS domain name.

`--servers <string>`Specifies a list of NIS server to be used by this provider. Repeat this option to specify multiple list items. This option overwrites the entries in the NIS servers list; to add or remove servers without affecting current entries, use `--add-servers` or `--remove-servers`.`--clear-servers`

Removes all entries from the list of NIS servers.

`--add-servers <string>`

Adds an entry to the list of NIS servers. Repeat this option to specify multiple items.

`--remove-servers <string>`

Removes an entry from the list of NIS servers. Repeat this option to specify multiple items.

`--authentication {yes | no}`Enables or disables the use of this provider for authentication as well as identity. The default value is `yes`.`--balance-servers {yes | no}`

Makes this provider connect to a random server on each request.

`--cache-entry-expiry <duration>`Specifies amount of time to cache a user or group, in the format `<integer>[{Y | M | W | D | H | m | s}]`.`--check-online-interval <duration>`Specifies the time between provider online checks, in the format `<integer>[{Y | M | W | D | H | m | s}]`.`--create-home-directory {yes | no}`

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

`--enabled {yes | no}`

Enables or disables this provider.

`--enumerate-groups {yes | no}`

Specifies whether to allow this provider to enumerate groups.

`--enumerate-users {yes | no}`

Specifies whether to allow this provider to enumerate users.

`--findable-groups <string>`Specifies a group that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be resolved. This option overwrites the entries in the findable groups list; to add or remove groups without affecting current entries, use `--add-findable-groups` or `--remove-findable-groups`.

```

--clear-findable-groups
    Removes all entries from the list of findable groups.
--add-findable-groups <string>
    Adds an entry to the list of findable groups that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--remove-findable-groups <string>
    Removes an entry from the list of findable groups that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--findable-users <string>
    Specifies a user that can be found in this provider if --restrict-findable is
    enabled. Repeat this option to specify multiple list items. If populated, users that are
    not included in this list cannot be resolved. This option overwrites the entries in the
    findable users list; to add or remove users without affecting current entries, use --
    add-findable-users or --remove-findable-users.
--clear-findable-users
    Removes all entries from the list of findable users.
--add-findable-users <string>
    Adds an entry to the list of findable users that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--remove-findable-users <string>
    Removes an entry from the list of findable users that is checked if --restrict-
    findable is enabled. Repeat this option to specify multiple list items.
--group-domain <string>
    Specifies the domain that this provider will use to qualify groups. The default group
    domain is NIS_GROUPS.
--home-directory-template <path>
    Specifies the path to use as a template for naming home directories. The path must
    begin with /ifs and can include special character sequences that are dynamically
    replaced with strings at home directory creation time that represent specific
    variables. For example, %U, %D, and %Z are replaced with the user name, provider
    domain name, and zone name, respectively. For more information, see the Home
    directories section.
--hostname-lookup {yes | no}
    Enables or disables host name lookups.
--listable-groups <string>
    Specifies a group that can be viewed in this provider if --restrict-listable is
    enabled. Repeat this option to specify multiple list items. If populated, groups that
    are not included in this list cannot be viewed. This option overwrites the entries in the
    listable groups list; to add or remove groups without affecting current entries, use --
    add-listable-groups or --remove-listable-groups.
--clear-listable-groups
    Removes all entries from the list of viewable groups.
--add-listable-groups <string>

```

Adds an entry to the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-listable-groups <string>`
Removes an entry from the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--listable-users <string>`
Specifies a user that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed. This option overwrites the entries in the listable users list; to add or remove users without affecting current entries, use `--add-listable-users` or `--remove-listable-users`.

`--clear-listable-users`
Removes all entries from the list of viewable users.

`--add-listable-users <string>`
Adds an entry to the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-listable-users <string>`
Removes an entry from the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--login-shell <path>`
Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.

`--normalize-groups {yes | no}`
Normalizes group names to lowercase before lookup.

`--normalize-users {yes | no}`
Normalizes user names to lowercase before lookup.

`--provider-domain <string>`
Specifies the domain that this provider will use to qualify user and group names.

`--ntlm-support {all | v2only | none}`
For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM.

`--request-timeout <integer>`
Specifies the request timeout interval in seconds.

`--restrict-findable {yes | no}`
Specifies whether to check this provider for filtered lists of findable and unfindable users and groups.

`--restrict-listable {yes | no}`
Specifies whether to check this provider for filtered lists of viewable and unviewable users and groups.

`--retry-time <integer>`
Sets the timeout period in seconds after which a request will be retried.

`--unfindable-groups <string>`

Specifies a group that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable groups list; to add or remove groups without affecting current entries, use `--add-unfindable-groups` or `--remove-unfindable-groups`.

`--clear-unfindable-groups`

Removes all entries from the list of unfindable groups.

`--add-unfindable-groups <string>`

Adds an entry to the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-unfindable-groups <string>`

Removes an entry from the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--unfindable-users <string>`

Specifies a user that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable users list; to add or remove users without affecting current entries, use `--add-unfindable-users` or `--remove-unfindable-users`.

`--clear-unfindable-users`

Removes all entries from the list of unfindable groups.

`--add-unfindable-users <string>`

Adds an entry to the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--remove-unfindable-users <string>`

Removes an entry from the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

`--unlistable-groups <string>`

Specifies a group that cannot be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable groups list; to add or remove groups without affecting current entries, use `--add-unlistable-groups` or `--remove-unlistable-groups`.

`--clear-unlistable-groups`

Removes all entries from the list of unlistable groups.

`--add-unlistable-groups <string>`

Adds an entry to the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-unlistable-groups <string>`

Removes an entry from the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--unlistable-users <string>`

Specifies a user that cannot be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the

entries in the unlistable users list; to add or remove users without affecting current entries, use `--add-unlistable-users` or `--remove-unlistable-users`.

`--clear-unlistable-users`
Removes all entries from the list of unviewable users.

`--add-unlistable-users <string>`
Adds an entry to the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--remove-unlistable-users <string>`
Removes an entry from the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

`--user-domain <string>`
Specifies the domain that this provider will use to qualify users. The default user domain is `NIS_USERS`.

`--ypmatch-using-tcp {yes | no}`
Uses TCP for YP Match operations.

`{--verbose | -v}`
Displays the results of running the command.

isi auth nis view

Displays the properties of an NIS provider.

Syntax

```
isi auth nis view <provider-name>
```

Options

`<provider-name>`

Specifies the name of the provider to view.

isi auth privileges

Displays a list of system privileges.

Syntax

```
isi auth privileges
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.


```
{--verbose | -v}
```

Displays more detailed information.

Note

When using the `--verbose` option, the output `Read Write: No` means that the privileges are read-only.

isi auth refresh

Refreshes authentication system configuration settings.

Syntax

```
isi auth refresh
```

Options

There are no options for this command.

isi auth roles create

Creates a custom role.

This command creates an empty role. To assign privileges and add members to the role, run the `isi auth roles modify` command.

Syntax

```
isi auth roles create <name>
  [--description <string>]
  [--verbose]
```

Options

<name>

Specifies the name of the role.

`--description <string>`

Specifies a description of the role.

```
{--verbose | -v}
```

Displays the results of running the command.

isi auth roles delete

Deletes a role.

Syntax

```
isi auth roles delete <role>
  [--force]
  [--verbose]
```

Options

<role>

Specifies the name of the role to delete.

```
{--force | -f}
```

Suppresses command-line prompts and messages.

```
{--verbose | -v}
```

Displays more detailed information.

isi auth roles list

Displays a list of roles.

Syntax

```
isi auth roles list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi auth roles members list

Displays a list of the members of a role.

Syntax

```
isi auth roles members list <role>
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
<role>
```

Specifies a role by name.

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

Examples

To view the members of the SystemAdmin role, run the following command:

```
isi auth roles members list systemadmin
```

In the following sample output, the SystemAdmin role currently contains one member, a user named admin:

```
Type Name
-----
user admin
-----
Total: 1
```

isi auth roles modify

Modifies a role.

Syntax

```
isi auth roles modify <role>
  [--name <string>]
  [--description <string>]
  [--add-group <string>]
  [--remove-group <string>]
  [--add-gid <integer>]
  [--remove-gid <integer>]
  [--add-uid <integer>]
  [--remove-uid <integer>]
  [--add-user <string>]
  [--remove-user <string>]
  [--add-sid <string>]
  [--remove-sid <string>]
  [--add-wellknown <string>]
  [--remove-wellknown <string>]
  [--add-priv <string>]
  [--add-priv-ro <string>]
  [--remove-priv <string>]
  [--verbose]
```

Options

<role>

Specifies the name of the role to modify.

--name <string>

Specifies a new name for the role. Applies to custom roles only.

--description <string>

Specifies a description of the role.

--add-group <string>

Adds a group with the specified name to the role. Repeat this option for each additional item.

--remove-group <string>

Removes a group with the specified name from the role. Repeat this option for each additional item.

`--add-gid <integer>`
 Adds a group with the specified GID to the role. Repeat this option for each additional item.

`--remove-gid <integer>`
 Removes a group with the specified GID from the role. Repeat this option for each additional item.

`--add-uid <integer>`
 Adds a user with the specified UID to the role. Repeat this option for each additional item.

`--remove-uid <integer>`
 Removes a user with the specified UID from the role. Repeat this option for each additional item.

`--add-user <string>`
 Adds a user with the specified name to the role. Repeat this option for each additional item.

`--remove-user <string>`
 Removes a user with the specified name from the role. Repeat this option for each additional item.

`--add-sid <string>`
 Adds a user or group with the specified SID to the role. Repeat this option for each additional item.

`--remove-sid <string>`
 Removes a user or group with the specified SID from the role. Repeat this option for each additional item.

`--add-wellknown <string>`
 Adds a well-known SID with the specified name—for example, Everyone—to the role. Repeat this option for each additional item.

`--remove-wellknown <string>`
 Removes a well-known SID with the specified name from the role. Repeat this option for each additional item.

`--add-priv <string>`
 Adds a read/write privilege to the role. Applies to custom roles only. Repeat this option for each additional item.

`--add-priv-ro <string>`
 Adds a read-only privilege to the role. Applies to custom roles only. Repeat this option for each additional item.

`--remove-priv <string>`
 Removes a privilege from the role. Applies to custom roles only. Repeat this option for each additional item.

`{--verbose | -v}`
 Displays the results of running the command.

isi auth roles privileges list

Displays a list of privileges that are associated with a role.

Syntax

```
isi auth roles privileges list <role>
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<role>

Specifies a role by name.

`--limit [-l] <integer>`

Displays no more than the specified number of items.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`--no-header [-a]`

Displays table and CSV output without headers.

`--no-footer [-z]`

Displays table output without footers.

`--verbose [-v]`

Displays more detailed information.

Examples

To list the privileges that are associated with the built-in SecurityAdmin role, run the following command:

```
isi auth roles privileges list securityadmin
```

The system displays output similar to the following example:

```
ID
-----
ISI_PRIV_LOGIN_CONSOLE
ISI_PRIV_LOGIN_PAPI
ISI_PRIV_LOGIN_SSH
ISI_PRIV_AUTH
ISI_PRIV_ROLE
-----
Total: 5
```

isi auth roles view

Displays the properties of a role.

Syntax

```
isi auth roles view <role>
```

Options*<role>*

Specifies the name of the role to view.

isi auth settings global modify

Modifies the global authentication settings.

Syntax

```
isi auth settings global modify
[--send-ntlmv2 {yes | no} | --revert-send-ntlmv2]
[--space-replacement <character> | --revert-space-replacement]
[--workgroup <string> | --revert-workgroup]
[--provider-hostname-lookup <string>]
[--cache-cred-lifetime <duration> | --revert-cache-cred-lifetime]
[--cache-id-lifetime <duration> | --revert-cache-id-lifetime]
[--on-disk-identity {native | unix | sid}
 | --revert-on-disk-identity]
[--rpc-max-requests <integer> | --revert-rpc-max-requests]
[--unknown-gid <integer> | --revert-unknown-gid]
[--unknown-uid <integer> | --revert-unknown-uid]
[--verbose]
```

Options*--send-ntlmv2* {yes | no}

Specifies whether to send only NTLMv2 responses to an SMB client. The default value is no. Valid values are yes, no. The default value is no.

*--revert-send-ntlmv2*Reverts the *--send-ntlmv2* setting to the system default value.*--space-replacement* *<character>*

For clients that have difficulty parsing spaces in user and group names, specifies a substitute character. Be careful to choose a character that is not in use.

*--revert-space-replacement*Reverts the *--space-replacement* setting to the system default value.*--workgroup* *<string>*

Specifies the NetBIOS workgroup. The default value is WORKGROUP.

*--revert-workgroup*Reverts the *--workgroup* setting to the system default value.*--provider-hostname-lookup* *<string>*

Allows hostname lookup through authentication providers. Applies to NIS only.

--alloc-retries *<integer>*

Specifies the number of times to retry an ID allocation before failing.

*--revert-alloc-retries*Reverts the *--alloc-retries* setting to the system default value.*--cache-cred-lifetime* *<duration>*Specifies the length of time to cache credential responses from the ID mapper, in the format *<integer>*{Y | M | W | D | H | m | s}.*--revert-cache-cred-lifetime*Reverts the *--cache-cred-lifetime* setting to the system default value.

`--cache-id-lifetime <duration>`
 Specifies the length of time to cache ID responses from the ID mapper, in the format `<integer>[Y|M|W|D|H|m|s]`.

`--revert-cache-id-lifetime`
 Reverts the `--cache-id-lifetime` setting to the system default value.

`--on-disk-identity <string>`
 Controls the preferred identity to store on disk. If OneFS is unable to convert an identity to the preferred format, it is stored as is. This setting does not affect identities that are already stored on disk. The accepted values are listed below.

native
 Allows OneFS to determine the identity to store on disk. This is the recommended setting.

unix
 Always stores incoming UNIX identifiers (UIDs and GIDs) on disk.

sid
 Stores incoming Windows security identifiers (SIDs) on disk unless the SID was generated from a UNIX identifier. If the SID was generated from a UNIX identifier, OneFS converts it back to the UNIX identifier and stores it on disk.

Note

To prevent permission errors after changing the on-disk identity, run `isi job jobs start PermissionRepair` with the `convert` mode specified.

`--revert-on-disk-identity`
 Sets the `--on-disk-identity` setting to the system default value.

`--rpc-max-requests <integer>`
 Specifies the maximum number of simultaneous ID mapper requests allowed. The default value is 64.

`--revert-rpc-max-requests`
 Sets the `--rpc-max-requests` setting to the system default value.

`--unknown-gid <integer>`
 Specifies the GID to use for the unknown (anonymous) group.

`--revert-unknown-gid`
 Sets the `--unknown-gid` setting to the system default value.

`--unknown-uid <integer>`
 Specifies the UID to use for the unknown (anonymous) user.

`--revert-unknown-uid`
 Sets the `--unknown-uid` setting to the system default value.

`{--verbose | -v}`
 Displays more detailed information.

isi auth settings global view

Displays global authentication settings.

Syntax

```
isi auth settings global view
```

Options

There are no options for this command.

Examples

To view the current authentication settings on the cluster, run the following command:

```
isi auth settings global view
```

The system displays output similar to the following example:

```

        Send NTLMv2: No
        Space Replacement:
          Workgroup: WORKGROUP
Provider Hostname Lookup: disabled
        Alloc Retries: 5
        Cache Cred Lifetime: 15m
        Cache ID Lifetime: 15m
        On Disk Identity: native
          RPC Block Time: 5s
          RPC Max Requests: 16
          RPC Timeout: 30s
System GID Threshold: 80
System UID Threshold: 80
  GID Range Enabled: Yes
    GID Range Min: 1000000
    GID Range Max: 2000000
  UID Range Enabled: Yes
    UID Range Min: 1000000
    UID Range Max: 2000000
  Min Mapped Rid: 2147483648
    Group UID: 4294967292
      Null GID: 4294967293
      Null UID: 4294967293
    Unknown GID: 4294967294
    Unknown UID: 4294967294

```

isi auth status

Displays provider status, including available authentication providers and which providers are function correctly.

Syntax

```
isi auth status [<zone><string>]
  [--limit [-l | <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<zone><string>

Specifies an access zone by name.

`--limit [-l] <integer>`

Specifies the number of providers to display.

`--format {table | json | csv | list}`

Displays providers in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi auth users create

Creates a user account.

Syntax

```
isi auth users create <name>
  [--enabled {yes | no}]
  [--expiry <timestamp>]
  [--email <string>]
  [--gecos <string>]
  [--home-directory <path>]
  [--password <string>]
  [--password-expires {yes | no}]
  [{--primary-group <name> | --primary-group-gid <integer>
   | --primary-group-sid <string>}]
  [--prompt-password-change {yes | no}]
  [--shell <path>]
  [--uid <integer>]
  [--zone <string>]
  [--provider <string>]
  [--set-password]
  [--verbose]
  [--force]
```

Options

`<name>`

Specifies the user name.

`--enabled {yes | no}`

Enables or disables the user.

`{--expiry | -x} <timestamp>`

Specifies the time at which the user account will expire, using the date format `<YYYY>-<MM>-<DD>` or the date/time format `<YYYY>-<MM>-<DD>T<hh>:<mm>[:<ss>]`.

`--email <string>`

Specifies the email address of the user.

`--gecos <string>`

Specifies the values for the following Gecos field entries in the user's password file:

```
Full Name:
Office Location:
```

```
Office Phone:
Home Phone:
Other information:
```

Values must be entered as a comma-separated list, and values that contain spaces must be enclosed in quotation marks. For example, the `--gecos="Jane Doe",Seattle,555-5555,, "Temporary worker"` option with these values results in the following entries:

```
Full Name: Jane Doe
Office Location: Seattle
Office Phone: 555-5555
Home Phone:
Other information: Temporary worker
```

- `--home-directory <path>`
Specifies the path to the user's home directory.
- `--password <string>`
Sets the user's password to the specified value. This option cannot be used with the `--set-password` option.
- `--password-expires {yes | no}`
Specifies whether to allow the password to expire.
- `--primary-group <name>`
Specifies the user's primary group by name.
- `--primary-group-gid <integer>`
Specifies the user's primary group by GID.
- `--primary-group-sid <string>`
Specifies the user's primary group by SID.
- `--prompt-password-change {yes | no}`
Prompts the user to change the password during the next login.
- `--shell <path>`
Specifies the path to the UNIX login shell.
- `--uid <integer>`
Overrides automatic allocation of the UNIX user identifier (UID) with the specified value. Setting this option is not recommended.
- `--zone <string>`
Specifies the access zone in which to create the user.
- `--provider <string>`
Specifies a local authentication provider in the specified access zone.
- `--set-password`
Sets the password interactively. This option cannot be used with the `--password` option.
- `{--verbose | -v}`
Displays the results of running the command.
- `{--force | -f}`
Suppresses command-line prompts and messages.

isi auth users delete

Deletes a local user from the system.

Syntax

```
isi auth users delete {<user> | --uid <integer> | --sid <string>}
  [--zone <string>]
  [--provider <string>]
  [--force]
  [--verbose]
```

Options

This command requires *<user>*, *--uid <integer>*, or *--sid <string>*.

<user>

Specifies the user by name.

--uid <integer>

Specifies the user by UID.

--sid <string>

Specifies the user by SID.

--zone <string>

Specifies the name of the access zone that contains the user.

--provider <string>

Specifies the name of the authentication provider that contains the user.

{*--force* | *-f*}

Suppresses command-line prompts and messages.

{*--verbose* | *-v*}

Displays the results of running the command.

isi auth users flush

Flushes cached user information.

Syntax

```
isi auth users flush
```

Options

There are no options for this command.

Examples

To flush all cached user information, run the following command:

```
isi auth user flush
```

isi auth users list

Displays a list of users. If no options are specified, all users in the System access zone are displayed.

Note

The `--domain` option must be specified to list Active Directory users.

Syntax

```
isi auth users list
  [--domain <string>]
  [--zone <string>]
  [--provider <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--domain <string>`

Displays only the users in the specified provider domain.

`--zone <string>`

Specifies the access zone whose users you want to list. The default access zone is System.

`--provider <string>`

Displays only the users in the specified authentication provider. The syntax for specifying providers is `<provider-type>:<provider-name>`, being certain to use the colon separator; for example, `isi auth users list --provider="lsa-ldap-provider:Unix LDAP"`.

`{--limit | -l} <integer>`.

Displays no more than the specified number of items.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi auth users modify

Modifies a local user.

Syntax

```
isi auth users modify {<user> | --uid <integer> | --sid <string>}
  [--enabled {yes | no}]
  [--expiry <timestamp>]
  [--unlock]
  [--email <string>]
  [--gecos <string>]
  [--home-directory <path>]
  [--password <string>]
  [--password-expires {yes | no}]
  [{--primary-group <string> | --primary-group-gid <integer>
  | --primary-group-sid <string>}]
  [--prompt-password-change {yes | no}]
  [--shell <path>]
  [--new-uid <integer>]
  [--zone <string>]
  [--add-group <name>]
  [--add-gid <id>]
  [--remove-group <name>]
  [--remove-gid <id>]
  [--provider <string>]
  [--set-password]
  [--verbose]
  [--force]
```

Options

This command requires *<user>*, *--uid <integer>*, or *--sid <string>*.

<user>

Specifies the user by name.

--uid <integer>

Specifies the user by UID.

--sid <string>

Specifies the user by SID.

--enabled {yes | no}

Enables or disables the user.

{--expiry | -x} <timestamp>

Specifies the time at which the user account will expire, using the date format *<YYYY>*
<MM><DD> or the date/time format *<YYYY><MM><DD>[T<hh>:<mm>]:<ss>*.

--unlock

Unlocks the user account if locked.

--email <string>

Specifies the email address of the user.

--gecos <string>

Specifies the values for the following Gecos field entries in the user's password file:

```
Full Name:
Office Location:
Office Phone:
```

```
Home Phone:
Other information:
```

Values must be entered as a comma-separated list, and values that contain spaces must be enclosed in quotation marks. For example, the `--gecos= "Jane Doe", Seattle, 555-5555, , "Temporary worker"` option with these values results in the following entries:

```
Full Name: Jane Doe
Office Location: Seattle
Office Phone: 555-5555
Home Phone:
Other information: Temporary worker
```

- `--home-directory <path>`
Specifies the path to the user's home directory.
- `--password <string>`
Sets the user's password to the specified value. This option cannot be used with the `--set-password` option.
- `--password-expires {yes | no}`
Specifies whether to allow the password to expire.
- `--primary-group <name>`
Specifies the user's primary group by name.
- `--primary-group-gid <integer>`
Specifies the user's primary group by GID.
- `--primary-group-sid <string>`
Specifies the user's primary group by SID.
- `--prompt-password-change {yes | no}`
Prompts the user to change the password during the next login.
- `--shell <path>`
Specifies the path to the UNIX login shell.
- `--new-uid <integer>`
Specifies a new UID for the user. Setting this option is not recommended.
- `--zone <string>`
Specifies the name of the access zone that contains the user.
- `--add-group <name>`
Specifies the name of a group to add the user to. Repeat this option to specify multiple list items.
- `--add-gid <integer>`
Specifies the GID of a group to add the user to. Repeat this option to specify multiple list items.
- `--remove-group <name>`
Specifies the name of a group to remove the user from. Repeat this option to specify multiple list items.
- `--remove-gid <integer>`

Specifies the GID of a group to remove the user from. Repeat this option to specify multiple list items.

`--provider <string>`

Specifies an authentication provider of the format `<type>:<instance>`. Valid provider types are `ads`, `ldap`, `nis`, `file`, and `local`. For example, an LDAP provider named `auth1` can be specified as `ldap:auth1`.

`--set-password`

Sets the password interactively. This option cannot be used with the `--password` option.

`{--verbose | -v}`

Displays the results of running the command.

`{--force | -f}`

Suppresses command-line prompts and messages.

isi auth users view

Displays the properties of a user.

Syntax

```
isi auth users view {<user> | --uid <integer> | --sid <string>}
  [--cached]
  [--show-groups]
  [--resolve-names]
  [--zone <string>]
  [--provider <string>]
```

Options

This command requires `<user>`, `--uid <integer>`, or `--sid <string>`.

`<user>`

Specifies the user by name.

`--uid <integer>`

Specifies the user by UID.

`--sid <string>`

Specifies the user by SID.

`--cached`

Returns only cached information.

`--show-groups`

Displays groups that include the user as a member.

`--resolve-names`

Resolves the names of all related groups and related identities.

`--zone <string>`

Specifies the name of the access zone that contains the user.

`--provider <string>`

Specifies the name of the authentication provider that contains the user in the format `<type>:<instance>`. Valid values for type are `ads`, `ldap`, `nis`, `file`, and `local`. For example an LDAP provider named `auth1` can be specified as `ldap:auth1`, or an Active Directory provider can be specified as `ads:YORK.east.com`.

CHAPTER 7

Identity management

This section contains the following topics:

- [Identity management overview](#) 322
- [Identity types](#) 322
- [Access tokens](#) 323
- [Access token generation](#) 324
- [Managing ID mappings](#) 329
- [Managing user identities](#) 332

Identity management overview

In environments with several different types of directory services, OneFS maps the users and groups from the separate services to provide a single unified identity on an EMC Isilon cluster and uniform access control to files and directories, regardless of the incoming protocol. This process is called identity mapping.

Isilon clusters are frequently deployed in multiprotocol environments with multiple types of directory services, such as Active Directory and LDAP. When a user with accounts in multiple directory services logs in to a cluster, OneFS combines the user's identities and privileges from all the directory services into a native access token.

You can configure OneFS settings to include a list of rules for access token manipulation to control user identity and privileges. For example, you can set a user mapping rule to merge an Active Directory identity and an LDAP identity into a single token that works for access to files stored over both SMB and NFS. The token can include groups from Active Directory and LDAP. The mapping rules that you create can solve identity problems by manipulating access tokens in many ways, including the following examples:

- Authenticate a user with Active Directory but give the user a UNIX identity.
- Select a primary group from competing choices in Active Directory or LDAP.
- Disallow login of users that do not exist in both Active Directory and LDAP.

For more information about identity management, see the white paper *Managing identities with the Isilon OneFS user mapping service* at EMC Online Support.

Identity types

OneFS supports three primary identity types, each of which you can store directly on the file system. Identity types are user identifier and group identifier for UNIX, and security identifier for Windows.

When you log on to an EMC Isilon cluster, the user mapper expands your identity to include your other identities from all the directory services, including Active Directory, LDAP, and NIS. After OneFS maps your identities across the directory services, it generates an access token that includes the identity information associated with your accounts. A token includes the following identifiers:

- A UNIX user identifier (UID) and a group identifier (GID). A UID or GID is a 32-bit number with a maximum value of 4,294,967,295.
- A security identifier (SID) for a Windows user account. A SID is a series of authorities and sub-authorities ending with a 32-bit relative identifier (RID). Most SIDs have the form S-1-5-21-*A*-*B*-*C*-*RID*, where *A*, *B*, and *C* are specific to a domain or computer and *RID* denotes the object in the domain.
- A primary group SID for a Windows group account.
- A list of supplemental identities, including all groups in which the user is a member.

The token also contains privileges that stem from administrative role-based access control.

On an Isilon cluster, a file contains permissions, which appear as an access control list (ACL). The ACL controls access to directories, files, and other securable system objects.

When a user tries to access a file, OneFS compares the identities in the user's access token with the file's ACL. OneFS grants access when the file's ACL includes an access control entry (ACE) that allows the identity in the token to access the file and that does

not include an ACE that denies the identity access. OneFS compares the access token of a user with the ACL of a file.

Note

For more information about access control lists, including a description of the permissions and how they correspond to POSIX mode bits, see the white paper titled *EMC Isilon multiprotocol data access with a unified security model* on the EMC Online Support web site.

When a name is provided as an identifier, it is converted into the corresponding user or group object and the correct identity type. You can enter or display a name in various ways:

- UNIX assumes unique case-sensitive namespaces for users and groups. For example, Name and name represent different objects.
- Windows provides a single, case-insensitive namespace for all objects and also specifies a prefix to target an Active Directory domain; for example, domain\name.
- Kerberos and NFSv4 define principals, which require names to be formatted the same way as email addresses; for example, name@domain.com.

Multiple names can reference the same object. For example, given the name support and the domain example.com, support, EXAMPLE\support and support@example.com are all names for a single object in Active Directory.

Access tokens

An access token is created when the user first makes a request for access.

Access tokens represent who a user is when performing actions on the cluster and supply the primary owner and group identities during file creation. Access tokens are also compared against the ACL or mode bits during authorization checks.

During user authorization, OneFS compares the access token, which is generated during the initial connection, with the authorization data on the file. All user and identity mapping occurs during token generation; no mapping takes place during permissions evaluation.

An access token includes all UIDs, GIDs, and SIDs for an identity, in addition to all OneFS privileges. OneFS reads the information in the token to determine whether a user has access to a resource. It is important that the token contains the correct list of UIDs, GIDs, and SIDs. An access token is created from one of the following sources:

Source	Authentication
Username	<ul style="list-style-type: none"> • SMB impersonate user • Kerberized NFSv3 • Kerberized NFSv4 • NFS export user mapping • HTTP • FTP • HDFS
Privilege Attribute Certificate (PAC)	<ul style="list-style-type: none"> • SMB NTLM

Source	Authentication
	<ul style="list-style-type: none"> Active Directory Kerberos
User identifier (UID)	<ul style="list-style-type: none"> NFS AUTH_SYS mapping

Access token generation

For most protocols, the access token is generated from the username or from the authorization data that is retrieved during authentication.

The following steps present a simplified overview of the complex process through which an access token is generated:

Step	Process	Description
1	User identity lookup	<p>Using the initial identity, the user is looked up in all configured authentication providers in the access zone, in the order in which they are listed, until a match is found. The user identity and group list are retrieved from the authenticating provider. Any SIDs, UIDs, or GIDs are added to the initial token.</p> <hr/> <p>Note</p> <p>An exception to this behavior occurs if the AD provider is configured to call other providers, such as LDAP or NIS.</p>
2	ID mapping	The user's identifiers are associated across directory services. All SIDs are converted to their equivalent UID/GID and vice versa. These ID mappings are also added to the access token.
3	User mapping	Access tokens from other directory services are combined. If the username matches any user mapping rules, the rules are processed in order and the token is updated accordingly.
4	On-disk identity calculation	The default on-disk identity is calculated from the final token and the global setting. These identities are used for newly created files.

ID mapping

The Identity (ID) mapping service maintains relationship information between mapped Windows and UNIX identifiers to provide consistent access control across file sharing protocols within an access zone.

Note

ID mapping and user mapping are different services, despite the similarity in names.

During authentication, the authentication daemon requests identity mappings from the ID mapping service in order to create access tokens. Upon request, the ID mapping service returns Windows identifiers mapped to UNIX identifiers or UNIX identifiers mapped to Windows identifiers. When a user authenticates to a cluster over NFS with a UID or GID, the ID mapping service returns the mapped Windows SID, allowing access to files that another user stored over SMB. When a user authenticates to the cluster over

SMB with a SID, the ID mapping service returns the mapped UNIX UID and GID, allowing access to files that a UNIX client stored over NFS.

Mappings between UIDs or GIDs and SIDs are stored according to access zone in a cluster-distributed database called the ID map. Each mapping in the ID map is stored as a one-way relationship from the source to the target identity type. Two-way mappings are stored as complementary one-way mappings.

Mapping Windows IDs to UNIX IDs

When a Windows user authenticates with an SID, the authentication daemon searches the external Active Directory provider to look up the user or group associated with the SID. If the user or group has only an SID in the Active Directory, the authentication daemon requests a mapping from the ID mapping service.

Note

User and group lookups may be disabled or limited, depending on the Active Directory settings. You enable user and group lookup settings through the `isi auth ads modify` command.

If the ID mapping service does not locate and return a mapped UID or GID in the ID map, the authentication daemon searches other external authentication providers configured in the same access zone for a user that matches the same name as the Active Directory user.

If a matching user name is found in another external provider, the authentication daemon adds the matching user's UID or GID to the access token for the Active Directory user, and the ID mapping service creates a mapping between the UID or GID and the Active Directory user's SID in the ID map. This is referred to as an *external mapping*.

Note

When an external mapping is stored in the ID map, the UID is specified as the on-disk identity for that user. When the ID mapping service stores a generated mapping, the SID is specified as the on-disk identity.

If a matching user name is not found in another external provider, the authentication daemon assigns a UID or GID from the ID mapping range to the Active Directory user's SID, and the ID mapping service stores the mapping in the ID map. This is referred to as a *generated mapping*. The ID mapping range is a pool of UIDs and GIDs allocated in the mapping settings.

After a mapping has been created for a user, the authentication daemon retrieves the UID or GID stored in the ID map upon subsequent lookups for the user.

Mapping UNIX IDs to Windows IDs

The ID mapping service creates temporary UID-to-SID and GID-to-SID mappings only if a mapping does not already exist. The UNIX SIDs that result from these mappings are never stored on disk.

UIDs and GIDs have a set of predefined mappings to and from SIDs.

If a UID-to-SID or GID-to-SID mapping is requested during authentication, the ID mapping service generates a temporary UNIX SID in the format `S-1-22-1-⟨UID⟩` or `S-1-22-2-⟨GID⟩` by applying the following rules:

- For UIDs, the ID mapping service generates a UNIX SID with a domain of `S-1-22-1` and a resource ID (RID) matching the UID. For example, the UNIX SID for UID 600 is `S-1-22-1-600`.

- For GIDs, the ID mapping service generates a UNIX SID with a domain of S-1-22-2 and an RID matching the GID. For example, the UNIX SID for GID 800 is S-1-22-2-800.

ID mapping ranges

In access zones with multiple external authentication providers, such as Active Directory and LDAP, it is important that the UIDs and GIDs from different providers that are configured in the same access zone do not overlap. Overlapping UIDs and GIDs between providers within an access zone might result in some users gaining access to other users' directories and files.

The range of UIDs and GIDs that can be allocated for generated mappings is configurable in each access zone through the `isi auth settings mappings modify` command. The default range for both UIDs and GIDs is 1000000–2000000 in each access zone.

Do not include commonly used UIDs and GIDs in your ID ranges. For example, UIDs and GIDs below 1000 are reserved for system accounts and should not be assigned to users or groups.

User mapping

User mapping provides a way to control permissions by specifying a user's security identifiers, user identifiers, and group identifiers. OneFS uses the identifiers to check file or group ownership.

With the user-mapping feature, you can apply rules to modify which user identity OneFS uses, add supplemental user identities, and modify a user's group membership. The user-mapping service combines a user's identities from different directory services into a single access token and then modifies it according to the rules that you create.

Note

You can configure mapping rules on a per-zone basis. Mapping rules must be configured separately in each access zone that uses them. OneFS maps users only during login or protocol access.

Default user mappings

Default user mappings determine access if explicit user-mapping rules are not created.

If you do not configure rules, a user who authenticates with one directory service receives the identity information in other directory services when the account names are the same. For example, a user who authenticates with an Active Directory domain as Desktop\jane automatically receives identities in the final access token for the corresponding UNIX user account for jane from LDAP or NIS.

In the most common scenario, OneFS is connected to two directory services, Active Directory and LDAP. In such a case, the default mapping provides a user with the following identity attributes:

- A UID from LDAP
- The user SID from Active Directory
- An SID from the default group in Active Directory

The user's groups come from Active Directory and LDAP, with the LDAP groups and the autogenerated group GID added to the list. To pull groups from LDAP, the mapping service queries the `memberUid` attribute. The user's home directory, `gecos`, and shell come from Active Directory.

Elements of user-mapping rules

You combine operators with user names to create a user-mapping rule.

The following elements affect how the user mapper applies a rule:

- The operator, which determines the operation that a rule performs
- Fields for usernames
- Options
- A parameter
- Wildcards

User-mapping best practices

You can follow best practices to simplify user mapping.

Use Active Directory with RFC 2307 and Windows Services for UNIX

Use Microsoft Active Directory with Windows Services for UNIX and RFC 2307 attributes to manage Linux, UNIX, and Windows systems. Integrating UNIX and Linux systems with Active Directory centralizes identity management and eases interoperability, reducing the need for user-mapping rules. Make sure your domain controllers are running Windows Server 2003 or later.

Employ a consistent username strategy

The simplest configurations name users consistently, so that each UNIX user corresponds to a similarly named Windows user. Such a convention allows rules with wildcard characters to match names and map them without explicitly specifying each pair of accounts.

Do not use overlapping ID ranges

In networks with multiple identity sources, such as LDAP and Active Directory with RFC 2307 attributes, you should ensure that UID and GID ranges do not overlap. It is also important that the range from which OneFS automatically allocates UIDs and GIDs does not overlap with any other ID range. OneFS automatically allocates UIDs and GIDs from the range 1,000,000-2,000,000. If UIDs and GIDs overlap multiple directory services, some users might gain access to other users' directories and files.

Avoid common UIDs and GIDs

Do not include commonly used UIDs and GIDs in your ID ranges. For example, UIDs and GIDs below 1000 are reserved for system accounts; do not assign them to users or groups.

Do not use UPNs in mapping rules

You cannot use a user principal name (UPN) in a user mapping rule. A UPN is an Active Directory domain and username that are combined into an Internet-style name with an @ symbol, such as an email address: jane@example. If you include a UPN in a rule, the mapping service ignores it and may return an error. Instead, specify names in the format DOMAIN\user.com.

Group rules by type and order them

The system processes every mapping rule by default, which can present problems when you apply a deny-all rule—for example, to deny access to all unknown users. In addition, replacement rules might interact with rules that contain wildcard characters. To minimize complexity, it is recommended that you group rules by type and organize them in the following order:

1. Replacement rules: Specify all rules that replace an identity first to ensure that OneFS replaces all instances of the identity.
2. Join, add, and insert rules: After the names are set by any replacement operations, specify join, add, and insert rules to add extra identifiers.
3. Allow and deny rules: Specify rules that allow or deny access last.

Note

Stop all processing before applying a default deny rule. To do so, create a rule that matches allowed users but does nothing, such as an add operator with no field options, and has the break option. After enumerating the allowed users, you can place a catchall deny at the end to replace anybody unmatched with an empty user.

To prevent explicit rules from being skipped, in each group of rules, order explicit rules before rules that contain wildcard characters.

Add the LDAP or NIS primary group to the supplemental groups

When an Isilon cluster is connected to Active Directory and LDAP, a best practice is to add the LDAP primary group to the list of supplemental groups. This lets OneFS honor group permissions on files created over NFS or migrated from other UNIX storage systems. The same practice is advised when an Isilon cluster is connected to both Active Directory and NIS.

On-disk identity

After the user mapper resolves a user's identities, OneFS determines an authoritative identifier for it, which is the preferred on-disk identity.

OneFS stores either UNIX or Windows identities in file metadata on disk. On-disk identity types are UNIX, SID, and native. Identities are set when a file is created or a file's access control data is modified. Almost all protocols require some level of mapping to operate correctly, so choosing the preferred identity to store on disk is important. You can configure OneFS to store either the UNIX or the Windows identity, or you can allow OneFS to determine the optimal identity to store.

On-disk identity types are UNIX, SID, and native. Although you can change the type of on-disk identity, the native identity is best for a network with UNIX and Windows systems. In native on-disk identity mode, setting the UID as the on-disk identity improves NFS performance.

Note

The SID on-disk identity is for a homogeneous network of Windows systems managed only with Active Directory. When you upgrade from a version earlier than OneFS 6.5, the on-disk identity is set to UNIX. When you upgrade from OneFS 6.5 or later, the on-disk identity setting is preserved. On new installations, the on-disk identity is set to native.

The native on-disk identity type allows the OneFS authentication daemon to select the correct identity to store on disk by checking for the identity mapping types in the following order:

Order	Mapping type	Description
1	Algorithmic mapping	An SID that matches S-1-22-1-UID or S-1-22-2-GID in the internal ID mapping database is converted back to the corresponding UNIX identity, and the UID and GID are set as the on-disk identity.
2	External mapping	A user with an explicit UID and GID defined in a directory service (such as Active Directory with RFC 2307 attributes, LDAP, NIS, or the OneFS file provider or local provider) has the UNIX identity set as the on-disk identity.
3	Persistent mapping	Mappings are stored persistently in the identity mapper database. An identity with a persistent mapping in the identity mapper database uses the destination of that mapping as the on-disk identity, which occurs primarily with manual ID mappings. For example, if there is an ID mapping of GID:10000 to S-1-5-32-545, a request for the on-disk storage of GID:10000 returns S-1-5-32-545.
4	No mapping	If a user lacks a UID or GID even after querying the other directory services and identity databases, its SID is set as the on-disk identity. In addition, to make sure a user can access files over NFS, OneFS allocates a UID and GID from a preset range of 1,000,000 to 2,000,000. In native on-disk identity mode, a UID or GID that OneFS generates is never set as the on-disk identity.

Note

If you change the on-disk identity type, you should run the PermissionRepair job in convert mode to make sure that the disk representation of all files is consistent with the changed setting.

Managing ID mappings

You can create, modify, and delete identity mappings and configure ID mapping settings.

Create an identity mapping

You can create a manual identity mapping between source and target identities or automatically generate a mapping for a source identity.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth mapping create` command.

The following command specifies IDs of source and target identities in the zone3 access zone to create a two-way mapping between the identities:

```
isi auth mapping create --2way --source-sid=S-1-5-21-12345 \
--target-uid=5211 --zone=zone3
```

Modify an identity mapping

You can modify the configuration of an identity mapping.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth mapping modify` command.

The following command modifies the mapping of the user with UID 4236 in the zone3 access zone to include a reverse, 2-way mapping between the source and target identities:

```
isi auth mapping modify --source-uid=4236 \
--target-sid=S-1-5-21-12345 --zone=zone3 --2way
```

Delete an identity mapping

You can delete one or more identity mappings.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth mapping delete` command.

The following command deletes all identity mappings in the zone3 access zone:

```
isi auth mapping delete --all --zone=zone3
```

The following command deletes all identity mappings in the zone3 access zone that were both created automatically and include a UID or GID from an external authentication source:

```
isi auth mapping delete --all --only-external --zone=zone3
```

The following command deletes the identity mapping of the user with UID 4236 in the zone3 access zone:

```
isi auth mapping delete --source-uid=4236 --zone=zone3
```

View an identity mapping

You can display mapping information for a specific identity.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth mapping view` command.

The following command displays mappings for the user with UID 4236 in the zone3 access zone:

```
isi auth mapping view --uid=4236 --zone=zone3
```

The system displays output similar to the following example:

```
Name: user_36
On-disk: UID: 4236
Unix uid: 4236
Unix gid: -100000
SMB: S-1-22-1-4236
```

Flush the identity mapping cache

You can flush the ID map cache to remove in-memory copies of all or specific identity mappings.

Modifications to ID mappings may cause the cache to become out-of-sync and users might experience slowness or stalls when authenticating. You can flush the cache to synchronize the mappings.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth mapping flush` command.

The following command flushes all identity mappings on the EMC Isilon cluster:

```
isi auth mapping flush --all
```

The following command flushes the mapping of the user with UID 4236 in the zone3 access zone:

```
isi auth mapping flush --source-uid=4236 --zone=zone3
```

View a user token

You can view the contents of an access token generated for a user during authentication.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth mapping token` command.

The following command displays the access token of a user with UID 4236 in the zone3 access zone:

```
isi auth mapping token --uid=4236 --zone=zone3
```

The system displays output similar to the following example:

```
User
Name: user_36
UID: 4236
SID: S-1-22-1-4236
On Disk: 4236
ZID: 3
Zone: zone3
Privileges: -
Primary Group
Name: user_36
GID: 4236
```

```
SID: S-1-22-2-4236
On Disk: 4236
```

Configure identity mapping settings

You can enable or disable automatic allocation of UIDs and GIDS and customize the range of ID values in each access zone. The default range is 1000000–2000000.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth settings mapping modify` command.

The following command enables automatic allocation of both UIDs and GIDs in the zone3 access zone and sets their allocation ranges to 25000–50000:

```
isi auth settings mapping modify --gid-range-enabled=yes \
--gid-range-min=25000 --gid-range-max=50000 --uid-range-
enabled=yes \
--uid-range-min=25000 --uid-range-max=50000 --zone=zone3
```

View identity mapping settings

You can view the current configuration of identity mapping settings in each zone.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi auth settings mapping view` command.

The following command displays the current settings in the zone3 access zone:

```
isi auth settings mapping view --zone=zone3
```

The system displays output similar to the following example:

```
GID Range Enabled: Yes
GID Range Min: 25000
GID Range Max: 50000
UID Range Enabled: Yes
UID Range Min: 25000
UID Range Max: 50000
```

Managing user identities

You can manage user identities by creating user-mapping rules.

When you create user-mapping rules, it is important to remember the following information:

- You can only create user-mapping rules if you are connected to the EMC Isilon cluster through the System zone; however, you can apply user-mapping rules to specific access zones. If you create a user-mapping rule for a specific access zone, the rule applies only in the context of its zone.
- When you change user-mapping on one node, OneFS propagates the change to the other nodes.

- After you make a user-mapping change, the OneFS authentication service reloads the configuration.

View user identity

You can view the identities and group membership that a specified user has within the Active Directory and LDAP directory services.

This procedure must be performed through the command-line interface (CLI).

Note

The OneFS user access token contains a combination of identities from Active Directory and LDAP if both directory services are configured. You can run the following commands to discover the identities that are within each specific directory service.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. View a user identity from Active Directory only by running the `isi auth users view` command.

The following command displays the identity of a user named `stand` in the Active Directory domain named `YORK`:

```
isi auth users view --user=YORK\stand --show-groups
```

The system displays output similar to the following example:

```
Name: YORK\stand
DN:
CN=stand,CN=Users,DC=york,DC=hull,DC=example,DC=com
DNS Domain: york.hull.example.com
Domain: YORK
Provider: lsa-activedirectory-provider:YORK.HULL.EXAMPLE.COM
Sam Account Name: stand
UID: 4326
SID: S-1-5-21-1195855716-1269722693-1240286574-591111
Primary Group
  ID : GID:1000000
  Name : YORK\york_sh_udg
Additional Groups: YORK\sd-york space group
                  YORK\york_sh_udg
                  YORK\sd-york-group
                  YORK\sd-group
                  YORK\domain users
```

3. View a user identity from LDAP only by running the `isi auth users view` command.

The following command displays the identity of an LDAP user named `stand`:

```
isi auth user view --user=stand --show-groups
```

The system displays output similar to the following example:

```
Name: stand
DN:
uid=stand,ou=People,dc=colorado4,dc=hull,dc=example,dc=com
DNS Domain: -
Domain: LDAP_USERS
Provider: lsa-ldap-provider:Unix LDAP
Sam Account Name: stand
```

```

        UID: 4326
        SID: S-1-22-1-4326
    Primary Group
        ID : GID:7222
        Name : stand
    Additional Groups: stand
                    sd-group
                    sd-group2

```

Create a user-mapping rule

You can create user-mapping rules to manage user identities on the cluster.

You can create the first mapping rule with the `--user-mapping-rules` option for the `isi zone zones modify System` command. If you try to add a second rule with the command above, however, it replaces the existing rule rather than adding the new rule to the list of rules. To add more rules to the list of rules, you must use the `--add-user-mapping-rules` option with the `isi zone zones modify System` command.

Note

If you do not specify an access zone, user-mapping rules are created in the System zone.

Procedure

1. To create a rule to merge the Active Directory user with a user from LDAP, run the following command, where `<user-a>` and `<user-b>` are placeholders for the identities to be merged; for example, `user_9440` and `lduser_010`, respectively:

```
isi zone zones modify System --add-user-mapping-rules \
"<DOMAIN> <user-a> &= <user-b>"
```

Run the following command to view the rule:

```
isi zone zones view System
```

If the command runs successfully, the system displays the mapping rule, which is visible in the User Mapping Rules line of the output:

```

        Name: System
        Cache Size: 4.77M
    Map Untrusted:
        SMB Shares: -
    Auth Providers: -
    Local Provider: Yes
    NetBIOS Name:
    All SMB Shares: Yes
    All Auth Providers: Yes
    User Mapping Rules: <DOMAIN>\<user_a> &= <user_b>
    Home Directory Umask: 0077
    Skeleton Directory: /usr/share/skel
    Zone ID: 1

```

2. To verify the changes to the token, run a command similar to the following example:

```
isi auth mapping token <DOMAIN>\\<user-a>
```

If the command runs successfully, the system displays output similar to the following example:

```

User
Name : <DOMAIN>\<user-a>

```

```

        UID : 1000201
        SID : S-1-5-21-1195855716-1269722693-1240286574-11547
    ZID: 1
    Zone: System
    Privileges: -
Primary Group
    Name : <DOMAIN>\domain users
    GID : 1000000
    SID : S-1-5-21-1195855716-1269722693-1240286574-513
Supplemental Identities
    Name : Users
    GID : 1545
    SID : S-1-5-32-545

    Name : lduser_010
    UID : 10010
    SID : S-1-22-1-10010

    Name : example
    GID : 10000
    SID : S-1-22-2-10000

    Name : ldgroup_20user
    GID : 10026
    SID : S-1-22-2-10026

```

Merge Windows and UNIX tokens

You can use either the join or append operator to merge two user names into a single token.

When Windows and UNIX user names do not match across directory services, you can write user-mapping rules that use either the join or the append operator to merge two user names into a single token. For example, if a user's Windows username is win_bob and the user's UNIX username is UNIX_bob, you can join or append them.

When you append an account to another account, the append operator adds information from one identity to another. OneFS appends the fields that the options specify from the source identity to the target identity. OneFS appends the identifiers to the additional group list.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. Write a rule similar to the following example to join the Windows and UNIX user names, where *<win-username>* and *<UNIX-username>* are placeholders for the user's Windows and UNIX accounts:

```
MYDOMAIN\<win-username> &= <UNIX-username> []
```

3. Write a rule similar to the following example to append the UNIX account to the Windows account with the groups option:

```
MYDOMAIN\<win-username> ++ <UNIX-username> [groups]
```

Retrieve the primary group from LDAP

You can create a user-mapping rule to insert or append primary group information from LDAP into a user's access token.

By default, the user-mapping service combines information from AD and LDAP but gives precedence to the information from AD. Mapping rules control how OneFS combines the information. You can retrieve the primary group information from LDAP instead of AD.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. Write a rule similar to the following example to insert information from LDAP into a user's access token:

```
*\* += * [group]
```

3. Write a rule similar to the following example to append other information from LDAP to a user's access token:

```
*\* ++ * [user,groups]
```

Mapping rule options

Mapping rules can contain options that target the fields of an access token.

A field represents an aspect of a cross-domain access token, such as the primary UID and primary user SID from a user that you select. You can see some of the fields in the OneFS web administration interface. **User** in the web administration interface is the same as username. You can also see fields in an access token by running the command `isi auth mapping token`.

When you create a rule, you can add an option to manipulate how OneFS combines aspects of two identities into a single token. For example, an option can force OneFS to append the supplement groups to a token.

A token includes the following fields that you can manipulate with user mapping rules:

- username
- unix_name
- primary_uid
- primary_user_sid
- primary_gid
- primary_group_sid
- additional_ids (includes supplemental groups)

Options control how a rule combines identity information in a token. The break option is the exception: It stops OneFS from processing additional rules.

Although several options can apply to a rule, not all options apply to all operators. The following table describes the effect of each option and the operators that they work with.

Option	Operator	Description
user	insert, append	Copies the primary UID and primary user SID, if they exist, to the token.
groups	insert, append	Copies the primary GID and primary group SID, if they exist, to the token.
groups	insert, append	Copies all the additional identifiers to the token. The additional identifiers exclude the primary UID, the primary GID, the primary user SID, and the primary group SID.

Option	Operator	Description
default_user	all operators except remove groups	If the mapping service fails to find the second user in a rule, the service tries to find the username of the default user. The name of the default user cannot include wildcards. When you set the option for the default user in a rule with the command-line interface, you must set it with an underscore: default_user.
break	all operators	Stops the mapping service from applying rules that follow the insertion point of the break option. The mapping service generates the final token at the point of the break.

Mapping rule operators

The operator determines what a mapping rule does.

You can create user-mapping rules through either the web-administration interface, where the operators are spelled out in a list, or from the command-line interface.

When you create a mapping rule with the OneFS command-line interface (CLI), you must specify an operator with a symbol. The operator affects the direction in which the mapping service processes a rule. For more information about creating a mapping rule, see the white paper *Managing identities with the Isilon OneFS user mapping service*. The following table describes the operators that you can use in a mapping rule.

A mapping rule can contain only one operator.

Operator	Web interface	CLI	Direction	Description
append	Append fields from a user	++	Left-to-right	Modifies an access token by adding fields to it. The mapping service appends the fields that are specified in the list of options (user, group, groups) to the first identity in the rule. The fields are copied from the second identity in the rule. All appended identifiers become members of the additional groups list. An append rule without an option performs only a lookup operation; you must include an option to alter a token.
insert	Insert fields from a user	+=	Left-to-right	Modifies an existing access token by adding fields to it. Fields specified in the options list (user, group, groups) are copied from the new identity and inserted into the identity in the token. When the rule inserts a primary user or primary group, it become the new primary user and primary group in the token. The previous primary user and primary group move to the additional identifiers list. Modifying the primary user leaves the token's username unchanged. When inserting the additional groups from an identity, the service adds the new groups to the existing groups.
replace	Replace one user with a different user	=>	Left-to-right	Removes the token and replaces it with the new token that is identified by the second username. If the second username is empty,

Operator	Web interface	CLI	Direction	Description
				the mapping service removes the first username in the token, leaving no username. If a token contains no username, OneFS denies access with a <code>no such user</code> error.
remove groups	Remove supplemental groups from a user	--	Unary	Modifies a token by removing the supplemental groups.
join	Join two users together	&=	Bidirectional	Inserts the new identity into the token. If the new identity is the second user, the mapping service inserts it after the existing identity; otherwise, the service inserts it before the existing identity. The location of the insertion point is relevant when the existing identity is already the first in the list because OneFS uses the first identity to determine the ownership of new file system objects.

CHAPTER 8

Auditing

This section contains the following topics:

- [Auditing overview](#)340
- [Syslog](#) 340
- [Protocol audit events](#) 341
- [Supported event types](#) 343
- [Supported audit tools](#) 343
- [Managing audit settings](#)344
- [Integrating with the EMC Common Event Enabler](#)347
- [Auditing commands](#) 349

Auditing overview

You can audit system configuration changes and SMB and NFS protocol activity on an EMC Isilon cluster. All audit data is stored and protected in the cluster file system and organized by audit topics.

When you enable system configuration auditing, no additional configuration is required; all configuration events that are handled by the application programming interface (API) through the command-line interface (CLI) are tracked and recorded in the `config` audit topic directories.

Auditing can detect many potential sources of data loss, including fraudulent activities, inappropriate entitlements, and unauthorized access attempts. Customers in industries such as financial services, health care, life sciences, and media and entertainment, as well as in governmental agencies, must meet stringent regulatory requirements developed to protect against these sources of data loss.

You can enable and configure protocol auditing for one or more access zones in a cluster. If you enable protocol auditing for an access zone, file-access events through the SMB and NFS protocol are recorded in the protocol audit topic directories. The `protocol` audit log file is consumable by auditing applications that support the EMC Common Event Enabler (CEE). You can specify which events to log in each access zone. For example, you might want to audit the default set of `protocol` events in the System access zone but audit only successful attempts to delete files in a different access zone.

The audit events are logged on the individual nodes where the SMB or NFS client initiated the activity. The events are then stored in a binary file under `/ifs/.ifsvar/audit/logs`. The logs automatically roll over to a new file after the size reaches 1 GB.

Syslog

Syslog is a protocol that is used to convey certain event notification messages. The root user can configure an Isilon cluster to log audit events and forward them to syslog by using the syslog forwarder.

By default, all protocol events that occur on a particular node are forwarded to the `/var/log/audit_protocol.log` file, regardless of the access zone the event originated from.

Syslog is configured with an identity of `audit_protocol`, a facility of `syslog`, and a priority level of `info`.

Enable syslog

By default, audit event forwarding to syslog is not enabled when auditing is enabled. To enable this feature, you must configure audit syslog settings through the command line interface for zones.

Before you begin

Note

To enable audit event forwarding, you must configure audit syslog settings for each access zone. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a Secure Shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi zone zones modify` command with the `--syslog-forwarding-enabled` option to enable or disable audit syslog.

The following command enables audit syslog for an access zone named UserZone:

```
isi zone zones modify UserZone --syslog-forwarding-enabled=yes
```

The following command disables audit syslog for the UserZone access zone:

```
isi zone zones modify UserZone --syslog-forwarding-enabled=no
```

3. To view the audit settings, run the following command:

```
isi audit settings view
```

4. To view the audit configuration settings, run the following command:

```
isi zone zones view <zone>
```

Note

The syslog forwarder forwards only the zone's audit events to syslog that are set by the `--syslog-audit-events` parameter. This parameter is set to a list of comma-separated audit event types or "all" to set all the audit events. Only audit events that are defined by running the `isi zone zones modify` command with the `--audit-success` and `--audit-failure` options are eligible for forwarding to syslog.

Syslog forwarding

The syslog forwarder is a daemon that, when enabled, retrieves configuration changes and protocol audit events in an access zone and forwards the events to syslog. Only user-defined audit success and failure events are eligible for being forwarded to syslog.

On each node there is an audit syslog forwarder daemon running that will log audit events to the same node's syslog daemon.

Protocol audit events

By default, audited access zones track only certain events on the EMC Isilon cluster, including successful and failed attempts to access files and directories.

The default tracked events are create, close, delete, rename, and set_security.

The names of generated events are loosely based on the Windows I/O request packet (IRP) model in which all operations begin with a create event to obtain a file handle. A create event is required before all I/O operations, including the following: close, create, delete, get_security, read, rename, set_security, and write. A close event marks when the client is finished with the file handle that was produced by a create event.

These internally stored events are translated to events that are forwarded through CEE to the auditing application. The CEE export facilities on OneFS perform this mapping. CEE can be used to connect to any third party application that supports CEE.

Different SMB and NFS clients issue different requests, and one particular version of a platform such as Windows or Mac OS X using SMB might differ from another. Similarly,

different versions of an application such as Microsoft Word or Windows Explorer might make different protocol requests. For example, a client with a Windows Explorer window open might generate many events if an automatic or manual refresh of that window occurs. Applications issue requests with the logged-in user's credentials, but you should not assume that all requests are purposeful user actions.

When enabled, OneFS audit will track all changes that are made to the files and directories in SMB shares and NFS exports.

Sample config audit log

You can view both configuration audit and protocol audit logs by running the `isi_audit_viewer` command on any node in the Isilon cluster.

The `isi_audit_viewer -t config` command produces output similar to the following example:

```
[0: Fri Jan 23 16:17:03 2015] {"id":"524e0928-a35e-11e4-9d0c-005056302134","timestamp":1422058623106323,"payload":"PAPI config logging started."}

[1: Fri Jan 23 16:17:03 2015] {"id":"5249b99d-a35e-11e4-9d0c-005056302134","timestamp":1422058623112992,"payload":{"user":{"token":{"UID":0, "GID":0, "SID": "SID:S-1-22-1-0", "GSID": "SID:S-1-22-2-0", "GROUPS": ["SID:S-1-5-11", "GID:5", "GID:20", "GID:70", "GID:10"]}, "protocol": 17, "zone id": 1, "client": "10.7.220.97", "local": "10.7.177.176" }}, "uri": "/1/protocols/smb/shares", "method": "POST", "args": "", "body": {"path": "/ifs/data", "name": "Test" }}}

[2: Fri Jan 23 16:17:05 2015] {"id":"5249b99d-a35e-11e4-9d0c-005056302134","timestamp":1422058625144567,"payload":{"status":201,"statusmsg":"Created","body":{"id":"Test" }}}

[3: Fri Jan 23 16:17:39 2015] {"id":"67e7ca62-a35e-11e4-9d0c-005056302134","timestamp":1422058659345539,"payload":{"user":{"token":{"UID":0, "GID":0, "SID": "SID:S-1-22-1-0", "GSID": "SID:S-1-22-2-0", "GROUPS": ["SID:S-1-5-11", "GID:5", "GID:20", "GID:70", "GID:10"]}, "protocol": 17, "zone id": 1, "client": "10.7.220.97", "local": "10.7.177.176" }}, "uri": "/1/audit/settings", "method": "PUT", "args": "", "body": {"config_syslog_enabled": true }}}

[4: Fri Jan 23 16:17:39 2015] {"id":"67e7ca62-a35e-11e4-9d0c-005056302134","timestamp":1422058659387928,"payload":{"status":204,"statusmsg":"No Content","body":{}}}
```

Events come in pairs; a *pre event* is logged before the command is carried out and a *post event* is logged after the event is triggered. These events can be correlated by matching the `id` field. In the above logs, events 1 and 2 are paired, and events 3 and 4 are paired.

The pre event always comes first, and contains user token information, the PAPI path, and whatever arguments were passed to the PAPI call. In event 1, a POST request was made to `/1/protocols/smb/shares` with arguments `path=/ifs/data` and `name=Test`. The post event contains the HTTP return status and any output returned from the server.

Supported event types

You can view or modify the event types that are audited in an access zone.

For the most current list of supported auditing tools, see the Isilon [Third-Party Software & Hardware Compatibility Guide](#).

The following event types are configured by default on each audited access zone:

Event name	Example protocol activity
create	<ul style="list-style-type: none"> • Create a file or directory • Open a file, directory, or share • Mount a share • Delete a file
close	<ul style="list-style-type: none"> • Close a directory • Close a modified or unmodified file
rename	Rename a file or directory
delete	Delete a file or directory
set_security	Attempt to modify file or directory permissions

The following event types are available for exporting through CEE:

Event name	Example protocol activity
read	The first read request on an open file handle
write	The first write request on an open file handle
close	The client is finished with an open file handle
get_security	The client reads security information for an open file handle

The following protocol audit events can not be exported through CEE:

Event name	Example protocol activity
logon	SMB session create request by a client
logoff	SMB session logoff
tree_connect	SMB first attempt to access a share

Supported audit tools

You can configure OneFS to send protocol auditing logs to servers that support the EMC Common Event Enabler (CEE).

CEE has been tested and verified to work on several third-party software vendors. For the most current list of supported auditing tools, see the Isilon [Third-Party Software & Hardware Compatibility Guide](#).

Note

We recommend that you install and configure third-party auditing applications before you enable the OneFS auditing feature. Otherwise, the large backlog consumed by this feature may cause results to not be updated for a considerable amount of time.

Managing audit settings

You can enable and disable system configuration and protocol access audit settings, in addition to configuring integration with the EMC Common Event Enabler.

Enable system configuration auditing

OneFS can audit system configuration events on your Isilon cluster. When you enable or disable system configuration auditing, no additional configuration is required. If you enable configuration auditing, all configuration events that are handled by the platform API including writes, modifications, and deletions are tracked and recorded in the config audit topic directories.

To start collecting auditing information, enable configuration change auditing in either the OneFS web administration interface or the OneFS command-line interface (CLI). After being enabled, the audit function will track all configuration changes made over the WebUI or CLI, including the date and time the change occurred, what user made the change, and what the change was.

Configuration events are logged to `/var/log/audit_config.log`. Configuration events are not forwarded to the Common Event Enabler (CEE).

You can generate a `config` log by enabling configuration change auditing and making a modification to the cluster via PAPI. Configuration change logs are populated in the config topic in the audit back-end store under `/ifs/.ifsvvar/audit`.

Procedure

1. Run the `isi audit settings modify` command.

The following command enables system configuration auditing on the cluster:

```
isi audit settings modify --config-auditing-enabled=yes
```

Sample config audit log

You can view both configuration audit and protocol audit logs by running the `isi_audit_viewer` command on any node in the Isilon cluster.

The `isi_audit_viewer -t config` command produces output similar to the following example:

```
[0: Fri Jan 23 16:17:03 2015] {"id":"524e0928-a35e-11e4-9d0c-005056302134","timestamp":1422058623106323,"payload":"PAPI config logging started."}

[1: Fri Jan 23 16:17:03 2015] {"id":"5249b99d-a35e-11e4-9d0c-005056302134","timestamp":1422058623112992,"payload":{"user":{"token":{"UID":0, "GID":0, "SID": "SID:S-1-22-1-0", "GSID": "SID:S-1-22-2-0", "GROUPS": ["SID:S-1-5-11", "GID:5", "GID:20", "GID:70", "GID:10"]}, "protocol": 17, "zone id": 1, "client":
```



```

"10.7.220.97", "local": "10.7.177.176" }}, "uri": "/1/protocols/smb/
shares", "method": "POST", "args": "", "body": {"path": "/ifs/data",
"name": "Test"}}}

[2: Fri Jan 23 16:17:05 2015] {"id": "5249b99d-
a35e-11e4-9d0c-005056302134", "timestamp": 1422058625144567, "payload":
{"status": 201, "statusmsg": "Created", "body": {"id": "Test"}}}

[3: Fri Jan 23 16:17:39 2015] {"id": "67e7ca62-
a35e-11e4-9d0c-005056302134", "timestamp": 1422058659345539, "payload":
{"user": {"token": {"UID": 0, "GID": 0, "SID": "SID:S-1-22-1-0", "GSID":
"SID:S-1-22-2-0", "GROUPS": ["SID:S-1-5-11", "GID:5", "GID:20", "GID:
70", "GID:10"], "protocol": 17, "zone id": 1, "client":
"10.7.220.97", "local": "10.7.177.176" }}, "uri": "/1/audit/
settings", "method": "PUT", "args": "", "body": {"config_syslog_enabled":
true}}}

[4: Fri Jan 23 16:17:39 2015] {"id": "67e7ca62-
a35e-11e4-9d0c-005056302134", "timestamp": 1422058659387928, "payload":
{"status": 204, "statusmsg": "No Content", "body": {}}}

```

Events come in pairs; a *pre event* is logged before the command is carried out and a *post event* is logged after the event is triggered. These events can be correlated by matching the `id` field. In the above logs, events 1 and 2 are paired, and events 3 and 4 are paired.

The pre event always comes first, and contains user token information, the PAPI path, and whatever arguments were passed to the PAPI call. In event 1, a POST request was made to `/1/protocols/smb/shares` with arguments `path=/ifs/data` and `name=Test`. The post event contains the HTTP return status and any output returned from the server.

Enable protocol access auditing

You can audit SMB and NFS protocol access to generate events on a per-access zone basis and forward the events to the EMC Common Event Enabler (CEE) for export to third-party products.

The following protocol events are collected for audited access zones by default: `create`, `delete`, `rename`, and `set_security`. You can modify the set of events that are audited in an access zone by running the `isi zone zones modify` command.

Note

Because each audited event consumes system resources, we recommend that you only configure zones for events that are needed by your auditing application. In addition, we recommend that you install and configure third-party auditing applications before you enable the OneFS auditing feature. Otherwise, the large backlog performed by this feature may cause results to not be updated for a considerable amount of time. Additionally, you can manually configure the time that you want audit events to be forwarded by running the `isi audit settings modify --cee log-time` command.

Procedure

1. Run the `isi audit settings modify` command.

The following command enables SMB and NFS protocol access auditing in the System access zone, and forwards logged events to a CEE server:

```
isi audit settings modify --protocol-auditing-enabled=yes \
--cee-server-uris=http://sample.com:12228/cee \
--hostname=cluster.domain.com --audited-zones=System
```

Protocol events are written to the `/var/log/audit_protocol.log` file. After the auditing event has been logged, a CEE forwarder service handles forwarding the event to CEE. The event is forwarded through an HTTP PUT operation. At this point, CEE will forward the audit event to a defined endpoint.

Auditing settings

Basic settings for audit configuration are available through the `isi audit settings modify` command. When you audit protocol events for an access zone, a default set of audit events are logged. You can modify the list of audit events to log by running the `isi zone zones modify <zone>` command, where `<zone>` is the name of an audited access zone—for example, the System zone.

Options

`--protocol-auditing-enabled {yes | no}`

Enables or disables the auditing of I/O events.

`--audited-zones <zones>`

Specifies one or more access zones, separated by commas, that will be audited if protocol auditing is enabled. This option overwrites all entries in the list of access zones; to add or remove access zones without affecting current entries, use `--add-audited-zones` or `--remove-audited-zones`.

`--clear-audited-zones`

Clears the list of access zones to audit.

`--add-audited-zones <zones>`

Adds one or more access zones, separated by commas, to the list of zones that will be audited if protocol auditing is enabled.

`--remove-audited-zones <zones>`

Removes one or more access zones, separated by commas, which will be audited if protocol auditing is enabled.

`--cee-server-uris <uris>`

Specifies one or more CEE server URIs, separated by commas, where audit logs will be forwarded if protocol auditing is enabled. This option overwrites all entries in the list of CEE server URIs. To add or remove URIs without affecting current entries, use `--add-cee-server-uris` or `--remove-cee-server-uris`.

`--clear-cee-server-uris`

Clears the list of CEE server URIs to which audit logs are forwarded.

`--add-cee-server-uris <uris>`

Adds one or more CEE server URIs, separated by commas, where audit logs are forwarded if protocol auditing is enabled.

`--remove-cee-server-uris <uris>`

Removes one or more CEE server URIs, separated by commas, from the list of URIs where audit logs are forwarded if protocol auditing is enabled.

`--cee-log-time <date>`

Specifies a date after which the audit CEE forwarder will forward logs. To forward SMB or NFS traffic logs, specify `topic`. Specify `<date>` in the following format:

```
[protocol]@<YYYY>-<MM>-<DD> <HH>:<MM>:<SS>
```

`--syslog-log-time <date>`

Specifies a date after which the audit syslog forwarder will forward logs. To forward SMB or NFS traffic logs, specify `topic`. To forward configuration change logs, specify `config`. Specify `<date>` in the following format:

```
[protocol|config]@<YYYY>-<MM>-<DD> <HH>:<MM>:<SS>
```

`--hostname <string>`

Specifies the hostname of this cluster for reporting protocol events to CEE servers. This is typically the SmartConnect zone name. The hostname is used to construct the UNC path of audited files and directories—for example, `\\hostname\ifs\data\file.txt`.

`--config-auditing-enabled {yes | no}`

Enables or disables the auditing of requests to modify application programming interface (API) configuration settings.

`{--verbose | -v}`

Displays the results of running the command.

Integrating with the EMC Common Event Enabler

OneFS integration with the EMC Common Event Enabler (CEE) enables third-party auditing applications to collect and analyze SMB and NFS protocol auditing logs.

For the most current list of supported auditing tools, see the Isilon [Third-Party Software & Hardware Compatibility Guide](#).

OneFS supports the Common Event Publishing Agent (CEPA) component of CEE for Windows. For integration with OneFS, you must install and configure CEE for Windows on a supported Windows client.

Note

We recommend that you install and configure third-party auditing applications before you enable the OneFS auditing feature. Otherwise, the large backlog performed by this feature may cause results to not be updated for a considerable time.

Install CEE for Windows

To integrate CEE with OneFS, you must first install CEE on a computer that is running the Windows operating system.

Before you begin

Be prepared to extract files from the `.iso` file, as described in the following steps. If you are not familiar with the process, consider choosing one of the following methods:

1. Install WinRAR or another suitable archival program that can open `.iso` files as an archive, and copy the files.

2. Burn the image to a CD-ROM, and then copy the files.
3. Install SlySoft Virtual CloneDrive, which allows you to mount an ISO image as a drive that you can copy files from.

Note

You should install a minimum of two servers. We recommend that you install CEE 6.6.0 or later.

Procedure

1. Download the CEE framework software from EMC Online Support:
 - a. In a web browser, go to <https://support.emc.com/search/>.
 - b. In the search field, type **Common Event Enabler for Windows**, and then click the **Search** icon.
 - c. Click **Common Event Enabler <Version> for Windows**, where <Version> is 6.2 or later, and then follow the instructions to open or save the .iso file.

2. From the .iso file, extract the 32-bit or 64-bit EMC_CEE_Pack executable file that you need.

After the extraction completes, the **EMC Common Event Enabler** installation wizard opens.

3. Click **Next** to proceed to the **License Agreement** page.
4. Select the **I accept...** option to accept the terms of the license agreement, and then click **Next**.
5. On the **Customer Information** page, type your user name and organization, select your installation preference, and then click **Next**.
6. On the **Setup Type** page, select **Complete**, and then click **Next**.
7. Click **Install** to begin the installation.

The **Installing EMC Common Event Enabler** page displays the progress of the installation. When the installation is complete, the **InstallShield Wizard Completed** page appears.

8. Click **Finish** to exit the wizard.
9. Restart the system.

Configure CEE for Windows

After you install CEE for Windows on a client computer, you must configure additional settings through the Windows Registry Editor (`regedit.exe`).

Procedure

1. Open the Windows Registry Editor.
2. Configure the following registry keys, if supported by your audit application:

Setting	Registry location	Key	Value
CEE HTTP listen port	[HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\Configuration]	HttpPort	12228
Enable audit	[HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\CEPP\Audit\Configuration]	Enabled	1

Setting	Registry location	Key	Value
remote endpoints			
Audit remote endpoints	[HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\CEPP\Audit\Configuration]	EndPoint	<i>⟨EndPoint⟩</i>

Note

- The HttpPort value must match the port in the CEE URIs that you specify during OneFS protocol audit configuration.
 - The EndPoint value must be in the format *⟨EndPoint_Name⟩@⟨IP_Address⟩*. You can specify multiple endpoints by separating each value with a semicolon (;).
-

The following key specifies a single remote endpoint:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\CEPP\Audit\Configuration] EndPoint = AuditApplication@10.7.1.2
```

The following key specifies multiple remote endpoints:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\CEPP\Audit\Configuration] EndPoint = AuditApplication@192.168.22.3;AuditApplication@192.168.33.2
```

3. Close the Windows Registry Editor.

Auditing commands

You can audit system configuration events and SMB and NFS protocol access events on the EMC Isilon cluster. All audit data is stored in files called audit topics, which collect log information that you can process further with auditing tools for Windows.

isi audit settings modify

Enables or disables auditing of system configuration changes and protocol access, and configures additional protocol-auditing settings.

To enable auditing of system configuration changes, you must set the `--config-auditing-enabled` option to `yes`. No other settings are available.

To enable auditing of protocol access, you must set the `--protocol-auditing-enabled` option to `yes`, and you must also specify which access zones to audit by setting the `--audited-zones` option.

Note

If you are integrating with a third-party auditing application, it is recommended that you install and configure third-party auditing applications before you enable the OneFS auditing feature. Otherwise, the large backlog performed by this feature may cause results to not be updated for a considerable time.

Syntax

```
isi audit settings modify
  [--protocol-auditing-enabled {yes | no} ]
  [--audited-zones <zones> | --clear-audited-zones]
  [--add-audited-zones <zones>]
  [--remove-audited-zones <zones>]
  [--cee-server-uris <uris> | --clear-cee-server-uris]
  [--add-cee-server-uris <uris>]
  [--remove-cee-server-uris <uris>]
  [--hostname <string>]
  [--config-auditing-enabled {yes | no}]
  [--config-syslog-enabled {yes | no}]
  [--verbose]
```

Options

`--protocol-auditing-enabled {yes | no}`

Enables or disables the auditing of data-access requests through the SMB and NFS protocol.

`--audited-zones <zones>`

Specifies one or more access zones, separated by commas, which will be audited if protocol auditing is enabled. This option overwrites all entries in the list of access zones; to add or remove access zones without affecting current entries, use `--add-audited-zones` or `--remove-audited-zones`.

`--clear-audited-zones`

Clears the list of access zones to audit.

`--add-audited-zones <zones>`

Adds one or more access zones, separated by commas, to the list of zones that will be audited if protocol auditing is enabled.

`--remove-audited-zones <zones>`

Removes one or more access zones, separated by commas, which will be audited if protocol auditing is enabled.

`--cee-server-uris <uris>`

Specifies one or more CEE server URIs, separated by commas, where audit logs will be forwarded if protocol auditing is enabled. The OneFS CEE export service uses round robin load-balancing when exporting events to multiple CEE servers. This option overwrites all entries in the list of CEE server URIs. To add or remove URIs without affecting current entries, use `--add-cee-server-uris` or `--remove-cee-server-uris`.

`--clear-cee-server-uris`

Clears the list of CEE server URIs to which audit logs are forwarded.

`--add-cee-server-uris <uris>`

Adds one or more CEE server URIs, separated by commas, to the list of URIs where audit logs are forwarded.

`--remove-cee-server-uris <uris>`

Removes one or more CEE server URIs, separated by commas, from the list of URIs where audit logs are forwarded.

`--cee-log-time <date>`

Specifies a date after which the audit CEE forwarder will forward logs. To forward SMB or NFS traffic logs, specify `topic`. Specify `<date>` in the following format:

```
[protocol]@<YYYY>-<MM>-<DD> <HH>:<MM>:<SS>
```

`--syslog-log-time <date>`

Specifies a date after which the audit syslog forwarder will forward logs. To forward SMB or NFS traffic logs, specify `topic`. To forward configuration change logs, specify `config`. Specify `<date>` in the following format:

```
[protocol|config]@<YYYY>-<MM>-<DD> <HH>:<MM>:<SS>
```

`--hostname <string>`

Specifies the name of the storage cluster to use when forwarding protocol events—typically, the SmartConnect zone name. When SmartConnect is not implemented, the value must match the hostname of the cluster as your third-party audit application recognizes it. If the field is left blank, events from each node are filled with the node name (clustername + lnn). This setting is required only if needed by your third-party audit application.

`--config-auditing-enabled {yes | no}`

Enables or disables the auditing of requests made through the API for system configuration changes.

`--config-syslog-enabled {yes | no}`

Enables or disables the forwarding of system configuration changes to syslog.

`{--verbose | -v}`

Displays the results of running the command.

Note

OneFS collects the following protocol events for audited access zones by default: `create`, `close`, `delete`, `rename`, and `set_security`. You can specify the successful and failed events that are audited in an access zone by running the `isi zone zones modify` command. Because each audited event consumes system resources, you should only log events that are supported by your auditing application.

isi audit settings view

Displays audit configuration settings.

Syntax

```
isi audit settings view
```

Options

There are no options for this command.

Examples

To view current audit settings, run the following command:

```
isi audit settings view
```

The system displays output similar to the following text:

```
Protocol Auditing Enabled: Yes
    Audited Zones: System, zoneA
    CEE Server URIs: http://example.com:12228/cee
    Hostname: mycluster
Config Auditing Enabled: Yes
Config Syslog Enabled: Yes
```

isi audit topics list

Displays a list of configured audit topics, which are internal collections of audit data.

Syntax

```
isi audit topics list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi audit topics modify

Modifies the properties of an audit topic.

Syntax

```
isi audit topics modify <name>
  [--max-cached-messages <integer>]
  [--verbose]
```

Options

`<name>`

Specifies the name of the audit topic to modify. Valid values are `protocol` and `config`.

`--max-cached-messages <integer>`

Specifies the maximum number of audit messages to cache before writing them to a persistent store. The larger the number, the more efficiently audit events can be processed. If you specify 0, each audit event is sent synchronously.

{--verbose | -v}

Displays the results of running the command.

isi audit topics view

Displays the properties of an audit topic.

Syntax

```
isi audit topics view <name>
```

Options

<name>

Specifies the name of the audit topic whose properties you want to view. Valid values are `protocol` and `config`.

CHAPTER 9

File sharing

This section contains the following topics:

- [File sharing overview](#) 356
- [SMB](#) 358
- [NFS](#) 395
- [FTP](#) 444
- [HTTP and HTTPS](#) 460

File sharing overview

Multi-protocol support in OneFS enables files and directories on the Isilon cluster to be accessed through SMB for Windows file sharing, NFS for UNIX file sharing, secure shell (SSH), FTP, and HTTP. By default, only the SMB and NFS protocols are enabled.

OneFS creates the `/ifs` directory, which is the root directory for all file system data on the cluster. The `/ifs` directory is configured as an SMB share and an NFS export by default. You can create additional shares and exports within the `/ifs` directory tree.

Note

We recommend that you do not save data to the root `/ifs` file path but in directories below `/ifs`. The design of your data storage structure should be planned carefully. A well-designed directory structure optimizes cluster performance and administration.

You can set Windows- and UNIX-based permissions on OneFS files and directories. Users who have the required permissions and administrative privileges can create, modify, and read data on the cluster through one or more of the supported file sharing protocols.

- SMB. Allows Microsoft Windows and Mac OS X clients to access files that are stored on the cluster.
- NFS. Allows Linux and UNIX clients that adhere to the RFC1813 (NFSv3) and RFC3530 (NFSv4) specifications to access files that are stored on the cluster.
- HTTP and HTTPS (with optional DAV). Allows clients to access files that are stored on the cluster through a web browser.
- FTP. Allows any client that is equipped with an FTP client program to access files that are stored on the cluster through the FTP protocol.

Mixed protocol environments

The `/ifs` directory is the root directory for all file system data in the cluster, serving as an SMB share, an NFS export, and a document root directory. You can create additional shares and exports within the `/ifs` directory tree. You can configure your OneFS cluster to use SMB or NFS exclusively. You can also enable HTTP, FTP, and SSH.

Access rights are consistently enforced across access protocols on all security models. A user is granted or denied the same rights to a file whether using SMB or NFS. Clusters running OneFS support a set of global policy settings that enable you to customize the default access control list (ACL) and UNIX permissions settings.

OneFS is configured with standard UNIX permissions on the file tree. Through Windows Explorer or OneFS administrative tools, you can give any file or directory an ACL. In addition to Windows domain users and groups, ACLs in OneFS can include local, NIS, and LDAP users and groups. After a file is given an ACL, the mode bits are no longer enforced and exist only as an estimate of the effective permissions.

Note

We recommend that you configure ACL and UNIX permissions only if you fully understand how they interact with one another.

Write caching with SmartCache

Write caching accelerates the process of writing data to the cluster. OneFS includes a write-caching feature called SmartCache, which is enabled by default for all files and directories.

If write caching is enabled, OneFS writes data to a write-back cache instead of immediately writing the data to disk. OneFS can write the data to disk at a time that is more convenient.

Note

We recommend that you keep write caching enabled. You should also enable write caching for all file pool policies.

OneFS interprets writes to the cluster as either synchronous or asynchronous, depending on a client's specifications. The impacts and risks of write caching depend on what protocols clients use to write to the cluster, and whether the writes are interpreted as synchronous or asynchronous. If you disable write caching, client specifications are ignored and all writes are performed synchronously.

The following table explains how clients' specifications are interpreted, according to the protocol.

Protocol	Synchronous	Asynchronous
NFS	The stable field is set to <code>data_sync</code> or <code>file_sync</code> .	The stable field is set to <code>unstable</code> .
SMB	The <code>write-through</code> flag has been applied.	The <code>write-through</code> flag has not been applied.

Write caching for asynchronous writes

Writing to the cluster asynchronously with write caching is the fastest method of writing data to your cluster.

Write caching for asynchronous writes requires fewer cluster resources than write caching for synchronous writes, and will improve overall cluster performance for most workflows. However, there is some risk of data loss with asynchronous writes.

The following table describes the risk of data loss for each protocol when write caching for asynchronous writes is enabled:

Protocol	Risk
NFS	If a node fails, no data will be lost except in the unlikely event that a client of that node also crashes before it can reconnect to the cluster. In that situation, asynchronous writes that have not been committed to disk will be lost.
SMB	If a node fails, asynchronous writes that have not been committed to disk will be lost.

We recommend that you do not disable write caching, regardless of the protocol that you are writing with. If you are writing to the cluster with asynchronous writes, and you decide that the risks of data loss are too great, we recommend that you configure your clients to use synchronous writes, rather than disable write caching.

Write caching for synchronous writes

Write caching for synchronous writes costs cluster resources, including a negligible amount of storage space. Although it is not as fast as write caching with asynchronous writes, unless cluster resources are extremely limited, write caching with synchronous writes is faster than writing to the cluster without write caching.

Write caching does not affect the integrity of synchronous writes; if a cluster or a node fails, none of the data in the write-back cache for synchronous writes is lost.

SMB

OneFS includes a configurable SMB service to create and manage SMB shares. SMB shares provide Windows clients network access to file system resources on the cluster. You can grant permissions to users and groups to carry out operations such as reading, writing, and setting access permissions on SMB shares.

The `/ifs` directory is configured as an SMB share and is enabled by default. OneFS supports both user and anonymous security modes. If the user security mode is enabled, users who connect to a share from an SMB client must provide a valid user name with proper credentials.

SMB shares act as checkpoints, and users must have access to a share in order to access objects in a file system on a share. If a user has access granted to a file system, but not to the share on which it resides, that user will not be able to access the file system regardless of privileges. For example, assume a share named `ABCDOCS` contains a file named `file1.txt` in a path such as: `/ifs/data/ABCDOCS/file1.txt`. If a user attempting to access `file1.txt` does not have share privileges on `ABCDOCS`, that user cannot access the file even if originally granted read and/or write privileges to the file.

The SMB protocol uses security identifiers (SIDs) for authorization data. All identities are converted to SIDs during retrieval and are converted back to their on-disk representation before they are stored on the cluster.

When a file or directory is created, OneFS checks the access control list (ACL) of its parent directory. If the ACL contains any inheritable access control entries (ACEs), a new ACL is generated from those ACEs. Otherwise, OneFS creates an ACL from the combined file and directory create mask and create mode settings.

OneFS supports the following SMB clients:

SMB version	Supported operating systems
1.0	Windows 2000 or later Windows XP or later Mac OS X 10.5 or later
2.0	Windows Vista or later Windows Server 2008 or later Mac OS X 10.9 or later
2.1	Windows 7 or later Windows Server 2008 R2 or later
3.0 - Multichannel only	Windows 8 or later Windows Server 2012 or later

SMB shares in access zones

You can create and manage SMB shares within access zones.

You can create access zones that partition storage on the EMC Isilon cluster into multiple virtual containers. Access zones support all configuration settings for authentication and identity management services on the cluster, so you can configure authentication providers and provision SMB shares on a zone-by-zone basis. When you create an access zone, a local provider is created automatically, which allows you to configure each access zone with a list of local users and groups. You can also authenticate through a different Active Directory provider in each access zone, and you can control data access by directing incoming connections to the access zone from a specific IP address in a pool. Associating an access zone with an IP address pool restricts authentication to the associated access zone and reduces the number of available and accessible SMB shares.

Here are a few ways to simplify SMB management with access zones:

- Migrate multiple SMB servers, such as Windows file servers or NetApp filers, to a single Isilon cluster, and then configure a separate access zone for each SMB server.
- Configure each access zone with a unique set of SMB share names that do not conflict with share names in other access zones, and then join each access zone to a different Active Directory domain.
- Reduce the number of available and accessible shares to manage by associating an IP address pool with an access zone to restrict authentication to the zone.
- Configure default SMB share settings that apply to all shares in an access zone.

The Isilon cluster includes a built-in access zone named System, where you manage all aspects of the cluster and other access zones. If you don't specify an access zone when managing SMB shares, OneFS will default to the System zone.

SMB Multichannel

SMB Multichannel supports establishing a single SMB session over multiple network connections.

SMB Multichannel is a feature of the SMB 3.0 protocol that provides the following capabilities:

Increased throughput

OneFS can transmit more data to a client through multiple connections over high speed network adapters or over multiple network adapters.

Connection failure tolerance

When an SMB Multichannel session is established over multiple network connections, the session is not lost if one of the connections has a network fault, which enables the client to continue to work.

Automatic discovery

SMB Multichannel automatically discovers supported hardware configurations on the client that have multiple available network paths and then negotiates and establishes a session over multiple network connections. You are not required to install components, roles, role services, or features.

SMB Multichannel requirements

You must meet software and NIC configuration requirements to support SMB Multichannel on the EMC Isilon cluster.

OneFS can only support SMB Multichannel when the following software requirements are met:

- Windows Server 2012, 2012r2 or Windows 8, 8.1 clients
- SMB Multichannel must be enabled on both the EMC Isilon cluster and the Windows client computer. It is enabled on the Isilon cluster by default.

SMB Multichannel establishes a single SMB session over multiple network connections only on supported network interface card (NIC) configurations. SMB Multichannel requires at least one of the following NIC configurations on the client computer:

- Two or more network interface cards.
- One or more network interface cards that support Receive Side Scaling (RSS).
- One or more network interface cards configured with link aggregation. Link aggregation enables you to combine the bandwidth of multiple NICs on a node into a single logical interface.

Client-side NIC configurations supported by SMB Multichannel

SMB Multichannel automatically discovers supported hardware configurations on the client that have multiple available network paths.

Each node on the EMC Isilon cluster has at least one RSS-capable network interface card (NIC). Your client-side NIC configuration determines how SMB Multichannel establishes simultaneous network connections per SMB session.

Client-side NIC Configuration	Description
Single RSS-capable NIC	SMB Multichannel establishes a maximum of four network connections to the Isilon cluster over the NIC. The connections are more likely to be spread across multiple CPU cores, which reduces the likelihood of performance bottleneck issues and achieves the maximum speed capability of the NIC.
Multiple NICs	<p>If the NICs are RSS-capable, SMB Multichannel establishes a maximum of four network connections to the Isilon cluster over each NIC. If the NICs on the client are not RSS-capable, SMB Multichannel establishes a single network connection to the Isilon cluster over each NIC. Both configurations allow SMB Multichannel to leverage the combined bandwidth of multiple NICs and provides connection fault tolerance if a connection or a NIC fails.</p> <hr/> <p>Note</p> <p>SMB Multichannel cannot establish more than eight simultaneous network connections per session. In a multiple NIC configuration, this might limit the number connections allowed per NIC. For example, if the configuration contains three RSS-capable NICs, SMB Multichannel might establish three connections over the first NIC, three connections over the second NIC and two connections over the third NIC.</p>
Aggregated NICs	<p>SMB Multichannel establishes multiple network connections to the Isilon cluster over aggregated NICs, which results in balanced connections across CPU cores, effective consumption of combined bandwidth, and connection fault tolerance.</p> <hr/> <p>Note</p> <p>The aggregated NIC configuration inherently provides NIC fault tolerance that is not dependent upon SMB.</p>

SMB share management through MMC

OneFS supports the Shared Folders snap-in for the Microsoft Management Console (MMC), which allows SMB shares on the EMC Isilon cluster to be managed using the MMC tool.

Typically, you connect to the global System zone through the web administration interface or the command line interface to manage and configure shares. If you configure access zones, you can connect to a zone through the MMC Shared Folders snap-in to directly manage all shares in that zone.

You can establish a connection through the MMC Shared Folders snap-in to an Isilon node and perform the following SMB share management tasks:

- Create and delete shared folders
- Configure access permission to an SMB share
- View a list of active SMB sessions
- Close open SMB sessions
- View a list of open files
- Close open files

When you connect to a zone through the MMC Shared Folders snap-in, you can view and manage all SMB shares assigned to that zone; however, you can only view active SMB sessions and open files on the specific node that you are connected to in that zone. Changes you make to shares through the MMC Shared Folders snap-in are propagated across the cluster.

MMC connection requirements

You can connect to an EMC Isilon cluster through the MMC Shared Folders snap-in if you meet access requirements.

The following conditions are required to establish a connection through the MMC Shared Folders snap-in:

- You must run the Microsoft Management Console (MMC) from a Windows workstation that is joined to the domain of an Active Directory (AD) provider configured on the cluster.
- You must be a member of the local `<cluster>\Administrators` group.

Note

Role-based access control (RBAC) privileges do not apply to the MMC. A role with SMB privileges is not sufficient to gain access.

- You must log in to a Windows workstation as an Active Directory user that is a member of the local `<cluster>\Administrators` group.

Symbolic links and SMB clients

OneFS enables SMB2 clients to access symbolic links in a seamless manner. Many administrators deploy symbolic links to virtually reorder file system hierarchies, especially when crucial files or directories are scattered around an environment.

In an SMB share, a symbolic link (also known as a symlink or a soft link) is a type of file that contains a path to a target file or directory. Symbolic links are transparent to applications running on SMB clients, and they function as typical files and directories.

Support for relative and absolute links is enabled by the SMB client. The specific configuration depends on the client type and version.

A symbolic link that points to a network file or directory that is not in the path of the active SMB session is referred to as an absolute (or remote) link. Absolute links always point to the same location on a file system, regardless of the present working directory, and usually contain the root directory as part of the path. Conversely, a relative link is a symbolic link that points directly to a user's or application's working directory, so you do not have to specify the full absolute path when creating the link.

OneFS exposes symbolic links through the SMB2 protocol, enabling SMB2 clients to resolve the links instead of relying on OneFS to resolve the links on behalf of the clients. To transverse a relative or absolute link, the SMB client must be authenticated to the SMB shares that the link can be followed through. However, if the SMB client does not have permission to access the share, access to the target is denied and Windows will not prompt the user for credentials.

SMB2 and NFS links are interoperable for relative links only. For maximum compatibility, create these links from a POSIX client.

Note

SMB1 clients (such as Windows XP or 2002) may still use relative links, but they are traversed on the server side and referred to as "shortcut files." Absolute links do not work in these environments.

Enabling symbolic links

Before you can fully use symbolic links in an SMB environment, you must enable them.

For Windows SMB clients to traverse each type of symbolic link, you must enable them on the client. Windows supports the following link types:

- local to local
- remote to remote
- local to remote
- remote to local

You must run the following Windows command to enable all four link types:

```
fsutil behavior set SymlinkEvaluation L2L:1 R2R:1 L2R:1 R2L:1
```

For POSIX clients using Samba, you must set the following options in the `[global]` section of your Samba configuration file (`smb.conf`) to enable Samba clients to traverse relative and absolute links:

```
follow symlinks=yes
wide links=yes
```

In this case, "wide links" in the `smb.conf` file refers to absolute links. The default setting in this file is `no`.

Managing symbolic links

After enabling symbolic links, you can create or delete them from the Windows command prompt or a POSIX command line.

Create symbolic links using the Windows `mklink` command on an SMB2 client or the `ln` command from a POSIX command-line interface. For example, an administrator may want

to give a user named `User1` access to a file named `File1.doc` in the `/ifs/data/` directory without giving specific access to that directory by creating a link named `Link1`:

```
mklink \ifs\home\users\User1\Link1 \ifs\data\Share1\File1.doc
```

When you create a symbolic link, it is designated as a file link or directory link. Once the link is set, the designation cannot be changed. You can format symbolic link paths as either relative or absolute.

To delete symbolic links, use the `del` command in Windows, or the `rm` command in a POSIX environment.

Keep in mind that when you delete a symbolic link, the target file or directory still exists. However, when you delete a target file or directory, a symbolic link continues to exist and still points to the old target, thus becoming a broken link.

Anonymous access to SMB shares

You can configure anonymous access to SMB shares by enabling the local Guest user and allowing impersonation of the guest user.

For example, if you store files such as browser executables or other data that is public on the internet, anonymous access allows any user to access the SMB share without authenticating.

Managing SMB settings

You can enable or disable the SMB service, configure global settings for the SMB service, and configure default SMB share settings that are specific to each access zone.

View global SMB settings

You can view the global SMB settings that are applied to all nodes on the EMC Isilon cluster.

Procedure

1. Run the `isi smb settings global view` command.

The system displays output similar to the following example:

```
Access Based Share Enum: No
Dot Snap Accessible Child: Yes
Dot Snap Accessible Root: Yes
Dot Snap Visible Child: No
Dot Snap Visible Root: Yes
Enable Security Signatures: No
    Guest User: nobody
    Ignore Eas: No
Onefs Cpu Multiplier: 4
Onefs Num Workers: 0
Require Security Signatures: No
    Server String: Isilon Server
Srv Cpu Multiplier: 4
Srv Num Workers: 0
Support Multichannel: Yes
Support NetBIOS: No
Support Smb2: Yes
```

Configure global SMB settings

You can configure global settings for SMB file sharing.

CAUTION

Modifying global SMB file sharing settings could result in operational failures. Be aware of the potential consequences before modifying these settings.

Procedure

1. Run the `isi smb settings global modify` command.

The following example command specifies that read-only cannot be deleted from SMB shares:

```
isi smb settings global modify --allow-delete-readonly=no
```

Enable or disable the SMB service

The SMB service is enabled by default.

Note

You can determine whether the service is enabled or disabled by running the `isi services -l` command.

Procedure

1. Run the `isi services` command.

The following command disables the SMB service:

```
isi services smb disable
```

The following command enables the SMB service:

```
isi services smb enable
```

Enable or disable SMB Multichannel

SMB Multichannel is required for multiple, concurrent SMB sessions from a Windows client computer to a node in an EMC Isilon cluster. SMB Multichannel is enabled in the Isilon cluster by default.

You can enable or disable SMB Multichannel only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi smb settings global modify` command.

The following command enables SMB Multichannel on the EMC Isilon cluster:

```
isi smb settings global modify --support-multichannel=yes
```

The following command disables SMB Multichannel on the EMC Isilon cluster:

```
isi smb settings global modify --support-multichannel=no
```

View default SMB share settings

You can view the default SMB share settings specific to an access zone.

Procedure

- Run the `isi smb settings shares view` command.

The following example command displays the default SMB share settings configured for zone5 :

```
isi smb settings shares view --zone=zone5
```

The system displays output similar to the following example:

```

Access Based Enumeration: No
Access Based Enumeration Root Only: No
  Allow Delete Readonly: No
  Allow Execute Always: No
    Change Notify: norecuse
  Create Permissions: default acl
Directory Create Mask: 0700
Directory Create Mode: 0000
  File Create Mask: 0700
  File Create Mode: 0100
    Hide Dot Files: No
      Host ACL: -
  Impersonate Guest: never
  Impersonate User:
Mangle Byte Start: 0xED00
  Mangle Map: 0x01-0x1F:-1
  Ntfs ACL Support: Yes
    Oplocks: Yes
  Strict Flush: Yes
  Strict Locking: No

```

Configure default SMB share settings

You can configure SMB share settings specific to each access zone.

The default settings are applied to all new shares that are added to the access zone.

CAUTION

If you modify the default settings, the changes are applied to all existing shares in the access zone.

Procedure

1. Run the `isi smb settings shares modify` command.

The following command specifies that guests are never allowed access to shares in zone5:

```
isi smb settings global modify --zone=zone5 --impersonate-guest=never
```

Managing SMB shares

You can configure the rules and other settings that govern the interaction between your Windows network and individual SMB shares on the cluster.

OneFS supports %U, %D, %Z, %L, %0, %1, %2, and %3 variable expansion and automatic provisioning of user home directories.

You can configure the users and groups that are associated with an SMB share, and view or modify their share-level permissions.

Note

We recommend that you configure advanced SMB share settings only if you have a solid understanding of the SMB protocol.

Create an SMB share

When you create an SMB share, you can override the default permissions, performance, and access settings. You can configure SMB home directory provisioning by including expansion variables in the share path to automatically create and redirect users to their own home directories.

Before you begin

You must specify a path to use as the SMB share. Shares are specific to access zones and the share path must exist under the zone path. You can specify an existing path or create the path at the time you create the share. Create access zones before you create SMB shares.

You can specify one or more expansion variables in the directory path but you must set the flags to true for both the `--allow-variable-expansion` and `--auto-create-directory` parameters. If you do not specify these settings, the variable expansion string is interpreted literally by the system.

Procedure

1. Run the `isi smb shares create` command.

The following commands creates a directory at `/ifs/zone5/data/share1`, creates a share named `share1` using that path, and adds the share to the existing access zone named `zone5`:

```
mkdir /ifs/data/share1
isi smb shares create \
--name=share1 --path=/ifs/data/share1 --zone=zone5 \
--browsable=true --description="Example Share 1"
```

Note

Share names can contain up to 80 characters, and can only contain alphanumeric characters, hyphens, and spaces. Also, if the cluster character encoding is not set to UTF-8, SMB share names are case-sensitive.

The following command creates a directory at `/ifs/data/share2`, converts it to an SMB share, and adds the share to the default System zone because no zone is specified:

```
isi smb shares create share2 --path=/ifs/data/share2 \
--create-path --browsable=true --description="Example Share 2"
```

The following command creates a directory at `/ifs/data/share3` and converts it to an SMB share. The command also applies an ACL to the share:

```
isi smb shares create share3 --path=/ifs/data/share3 \
--create-path --browsable=true --description="Example Share 3" \
--inheritable-path-acl=true --create-permissions="default acl"
```

Note

If no default ACL is configured and the parent directory does not have an inheritable ACL, an ACL is created for the share with the `directory-create-mask` and `directory-create-mode` settings.

The following command creates the directory `/ifs/data/share4` and converts it to a non-browsable SMB share. The command also configures the use of mode bits for permissions control:

```
isi smb shares create --name=share4 --path=/ifs/data/share4 \
--create-path --browsable=false --description="Example Share 4" \
--inheritable-path-acl=true --create-permissions="use create mask
\
and mode"
```

2. The following command creates home directories for each user that connects to the share, based on the user's NetBIOS domain and user name.

In this example, if a user is in a domain named `DOMAIN` and has a username of `user_1`, the path `/ifs/home/%D/%U` expands to `/ifs/home/DOMAIN/user_1`.

```
isi smb shares modify HOMEDIR --path=/ifs/home/%D/%U \
--allow-variable-expansion=yes --auto-create-directory=yes
```

The following command creates a share named `HOMEDIR` with the existing path `/ifs/share/home`:

```
isi smb shares create HOMEDIR /ifs/share/home
```

3. Run the `isi smb shares permission modify` command to enable access to the share.

The following command allows the well-known user `Everyone` full permissions to the `HOMEDIR` share:

```
isi smb shares permission modify HOMEDIR --wellknown Everyone \
--permission-type allow --permission full
```

Modify an SMB share

You can modify the settings of individual SMB shares.

Before you begin

SMB shares are zone-specific. When you modify a share, you must identify the access zone that the share belongs to. If you do not identify the access zone, OneFS defaults to the System zone. If the share you want to modify has the same name as a share in the System zone, the share in the System zone is modified.

Procedure

1. Run the `isi smb shares modify` command.

In the following example, the file path for `share1` in `zone5` points to `/ifs/zone5/data`. The following commands modifies the file path of `share1` to `/ifs/zone5/etc`, which is another directory in the `zone5` path:

```
isi smb shares modify share1 --zone=zone5 \
--path=/ifs/zone5/etc
```

Note

If the cluster character encoding is not set to UTF-8, SMB share names are case-sensitive.

Delete an SMB share

You can delete SMB shares that are no longer needed.

Before you begin

SMB shares are zone-specific. When you delete a share, you must identify the access zone that the share belongs to. If you do not identify the access zone, OneFS defaults to the System zone. If the share you want to delete has the same name as a share in the System zone, the share in the System zone is deleted.

If you delete an SMB share, the share path is deleted but the directory it referenced still exists. If you create a new share with the same path as the share that was deleted, the directory that the previous share referenced will be accessible again through the new share.

Procedure

1. Run the `isi smb shares delete` command.

The following command deletes a share named `Share1` from the access zone named `zone-5`:

```
isi smb shares delete Share1 --zone=zone-5
```

2. Type **yes** at the confirmation prompt.

Limit access to /ifs share for the Everyone account

By default, the `/ifs` root directory is configured as an SMB share in the System access zone. It is recommended that you restrict the Everyone account of this share to read-only access.

Procedure

1. Run the `isi smb shares permission modify` command.

The following example changes the Everyone account permissions to read-only on the SMB share configured for the `/ifs` directory:

```
isi smb shares permission modify ifs --wellknown=Everyone \  
-d allow -p read
```

2. (Optional) Verify the change by running the following command to list permissions on the share:

```
isi smb shares permission list ifs
```

Configure anonymous access to a single SMB share

You can configure anonymous access to data stored on a single share through Guest user impersonation.

Procedure

1. Enable the Guest user account in the access zone that contains the share you want by running the `isi auth users modify` command.

The following command enables the guest user in the access zone named zone3:

```
isi auth users modify Guest --enabled=yes --zone=zone3
```

2. Set guest impersonation on the share you want to allow anonymous access to by running the `isi smb share modify` command.

The following command configures guest impersonation on a share named share1 in zone3:

```
isi smb share modify share1 --zone=zone3 \
--impersonate-guest=always
```

3. Verify that the Guest user account has permission to access the share by running the `isi smb share permission list` command.

The following command list the permissions for share1 in zone3:

```
isi smb share permission list share1 --zone=zone3
```

The system displays output similar to the following example

Account	Account Type	Run as Root	Permission Type	Permission
Everyone	wellknown	False	allow	read
Guest	user	False	allow	full

Configure anonymous access to all SMB shares in an access zone

You can configure anonymous access to data stored in an access zone through Guest user impersonation.

Procedure

1. Enable the Guest user account in the access zone that contains the share you want by running the `isi auth users modify` command.

The following command enables the guest user in the access zone named zone3:

```
isi auth users modify Guest --enabled=yes --zone=zone3
```

2. Set guest impersonation as the default value for all shares in the access zone by running the `isi smb settings share modify` command.

The following command configures guest impersonation for all shares in zone3:

```
isi smb settings share modify --zone=zone3 \
--impersonate-guest=always
```

3. Verify that the Guest user account has permission to each share in the access zone by running the `isi smb share permission list` command.

The following command list the permissions for share1 in zone3:

```
isi smb share permission list share1 --zone=zone3
```

The system displays output similar to the following example

Account	Account Type	Run as Root	Permission Type	Permission
Everyone	wellknown	False	allow	read
Guest	user	False	allow	full

Configure multi-protocol home directory access

For users who will access this share through FTP or SSH, you can make sure that their home directory path is the same whether they connect through SMB or they log in through FTP or SSH.

This command directs the SMB share to use the home directory template that is specified in the user's authentication provider. This procedure is available only through the command-line interface.

Procedure

1. Establish an SSH connection to any node in the cluster.
2. Run the following command, where *<homedir_share>* is the name of the SMB share:

```
isi smb share modify <homedir_share> --path=""
```

Supported expansion variables

You can include expansion variables in an SMB share path or in an authentication provider's home directory template.

OneFS supports the following expansion variables. You can improve performance and reduce the number of shares to be managed when you configure shares with expansion variables. For example, you can include the %U variable for a share rather than create a share for each user. When a %U is included in the name so that each user's path is different, security is still ensured because each user can view and access only his or her home directory.

Note

When you create an SMB share through the web administration interface, you must select the **Allow Variable Expansion** check box or the string is interpreted literally by the system.

Variable	Value	Description
%U	User name (for example, user_001)	Expands to the user name to allow different users to use different home directories. This variable is typically included at the end of the path. For example, for a user named user1, the path <code>/ifs/home/%U</code> is mapped to <code>/ifs/home/user1</code> .
%D	NetBIOS domain name (for example, YORK for YORK.EAST.EXAMPLE.COM)	Expands to the user's domain name, based on the authentication provider: <ul style="list-style-type: none"> • For Active Directory users, %D expands to the Active Directory NetBIOS name. • For local users, %D expands to the cluster name in uppercase characters. For example, for a cluster named cluster1, %D expands to CLUSTER1. • For users in the System file provider, %D expands to UNIX_USERS. • For users in other file providers, %D expands to FILE_USERS. • For LDAP users, %D expands to LDAP_USERS.

Variable	Value	Description
		<ul style="list-style-type: none"> For NIS users, %D expands to NIS_USERS.
%Z	Zone name (for example, ZoneABC)	Expands to the access zone name. If multiple zones are activated, this variable is useful for differentiating users in separate zones. For example, for a user named user1 in the System zone, the path <code>/ifs/home/%Z/%U</code> is mapped to <code>/ifs/home/System/user1</code> .
%L	Host name (cluster host name in lowercase)	Expands to the host name of the cluster, normalized to lowercase. Limited use.
%0	First character of the user name	Expands to the first character of the user name.
%1	Second character of the user name	Expands to the second character of the user name.
%2	Third character of the user name	Expands to the third character of the user name.

Note

If the user name includes fewer than three characters, the %0, %1, and %2 variables wrap around. For example, for a user named ab, the variables maps to a, b, and a, respectively. For a user named a, all three variables map to a.

SMB commands

You can access and configure the SMB file sharing service through the SMB commands.

isi smb log-level

Configures the log level settings for SMB shares on the node.

Syntax

```
isi smb log-level
  [--set | -s] <string>
  [--last-packets | -l] <integer>
```

Options

`{--set | -s} <string>`

Specifies the level of data logged for SMB shares on the node. Specify one of the following valid options:

- always
- error
- warning
- info
- verbose
- debug
- trace

`{--last-packets | -l} <integer>`

Specifies the last number of packets to be logged when an SMB session on the node closes.

isi smb log-level add

Specifies a log-level filter of SMB share information on the node.

Syntax

```
isi smb log-level add
  [--level | -l] <string>
  [--ip | -i] <ip-address>
  [--smb-op | -o] <string>
```

Options

`{--level | -l} <string>`

Sets the level information to be logged. The following values are valid:

- always
- error
- warning
- info
- verbose
- debug
- trace

`{--ip | -i} <ip-address>`

Specifies IP addresses filters. Valid format is `x.x.x.x` in a comma-separated list, or `x:x:x:x:x:x:x:x`.

`{--smb-op | -o} <string>`

Specifies operations filters. The following values are valid:

- read
- write
- session-setup
- logoff
- flush
- notify
- tree-connect
- tree-disconnect
- create
- delete
- oplock
- locking
- set-info
- query
- close
- create-directory

- delete-directory

isi smb log-level delete

Removes a log-level filter of SMB share information from the node.

Syntax

```
isi smb log-level delete
  [--id <string>]
```

Options

```
{--id | -i} <string>
```

Specifies the ID of the log-level filter to be deleted from the node.

isi smb log-level list

Displays the current log level and log information.

Syntax

```
isi smb log-level list
```

Options

There are no options for this command.

isi smb openfiles close

Closes an open file.

Note

To view a list of open files, run the `isi smb openfiles list` command.

Syntax

```
isi smb openfiles close <id>
  [--force]
```

Options

```
<id>
```

Specifies the ID of the open file to close.

```
{--force | -f}
```

Suppresses command-line prompts and messages.

Examples

The following command closes a file with an ID of 32:

```
isi smb openfiles close 32
```

isi smb openfiles list

Displays a list of files that are open in SMB shares.

Syntax

```
isi smb openfiles list
  [--limit <integer>]
```

```
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} *<integer>*

Displays no more than the specified number of smb openfiles.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi smb sessions delete

Deletes SMB sessions, filtered first by computer and then optionally by user.

Note

Any open files are automatically closed before an SMB session is deleted.

Syntax

```
isi smb sessions delete <computer-name>
  [{--user <name> | --uid <id> | --sid <sid>}]
  [--force]
  [--verbose]
```

Options

<computer-name>

Required. Specifies the computer name. If a --user, --uid, or --sid option is not specified, the system deletes all SMB sessions associated with this computer.

--user *<string>*

Specifies the name of the user. Deletes only those SMB sessions to the computer that are associated with the specified user.

--uid *<id>*

Specifies a numeric user identifier. Deletes only those SMB sessions to the computer that are associated with the specified user identifier.

--sid *<sid>*

Specifies a security identifier. Deletes only those SMB sessions to the computer that are associated with the security identifier.

{--force | -f}

Specifies that the command execute without prompting for confirmation.

Examples

The following command deletes all SMB sessions associated with a computer named computer1:

```
isi smb sessions delete computer1
```

The following command deletes all SMB sessions associated with a computer named computer1 and a user named user1:

```
isi smb sessions delete computer1 --user=user1
```

isi smb sessions delete-user

Deletes SMB sessions, filtered first by user then optionally by computer.

Note

Any open files are automatically closed before an SMB session is deleted.

Syntax

```
isi smb sessions delete-user {<user> | --uid <id> | --sid <sid> }
  [--computer-name <string>]
  [--force]
  [--verbose]
```

Options

<user>

Required. Specifies the user name. If the `--computer-name` option is omitted, the system deletes all SMB sessions associated with this user.

{--computer-name | -C} <string>

Deletes only the user's SMB sessions that are associated with the specified computer.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

Examples

The following command deletes all SMB sessions associated with a user called user1:

```
isi smb sessions delete-user user1
```

The following command deletes all SMB sessions associated with a user called user1 and a computer called computer1:

```
isi smb sessions delete-user user1 \
--computer-name=computer1
```

isi smb sessions list

Displays a list of open SMB sessions.

Syntax

```
isi smb sessions list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Specifies the maximum number of SMB sessions to list.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi smb settings global modify

Modifies global SMB settings.

Syntax

```
isi smb settings global modify
  [--access-based-share-enum {yes | no}]
  [--revert-access-based-share-enum]
  [--dot-snap-accessible-child {yes | no}]
  [--revert-dot-snap-accessible-child]
  [--dot-snap-accessible-root]
  [--revert-dot-snap-accessible-root]
  [--dot-snap-visible-child {yes | no}]
  [--revert-dot-snap-visible-child]
  [--dot-snap-visible-root {yes | no}]
  [--revert-dot-snap-visible-root]
  [--enable-security-signatures {yes | no}]
  [--revert-enable-security-signatures]
  [--guest-user <string>]
  [--revert-guest-user]
  [--ignore-eas {yes | no}]
  [--revert-ignore-eas]
  [--onefs-cpu-multiplier <integer>]
  [--revert-onefs-cpu-multiplier]
  [--onefs-num-workers <integer>]
  [--revert-onefs-num-workers]
  [--require-security-signatures {yes | no}]
  [--revert-require-security-signatures]
  [--server-string <string>]
  [--revert-server-string]
  [--srv-cpu-multiplier <integer>]
```



```
[--revert_srv-cpu-multiplier]
[--srv-num-workers <integer>]
[--revert-srv-num-workers]
[--support-multichannel {yes | no}]
[--revert-support-multichannel]
[--support-netbios {yes | no}]
[--revert-support-netbios]
[--support-smb2 {yes | no}]
[--revert-support-smb2]
[--verbose]
```

Options

- `--access-based-share-enum {yes | no}`
Enumerates only the files and folders that the requesting user has access to.
- `--revert-access-based-share-enum`
Sets the value to the system default for `--access-based-share-enum`.
- `--dot-snap-accessible-child {yes | no}`
Specifies whether to make the `/ifs/.snapshot` directory visible in subdirectories of the share root. The default setting is `no`.
- `--revert-dot-snap-accessible-child`
Sets the value to the system default for `--dot-snap-accessible-child`.
- `--dot-snap-accessible-root {yes | no}`
Specifies whether to make the `/ifs/.snapshot` directory accessible at the share root. The default setting is `yes`.
- `--revert-dot-snap-accessible-root`
Sets the value to the system default for `--dot-snap-accessible-root`.
- `--dot-snap-visible-child {yes | no}`
Specifies whether to make the `/ifs/.snapshot` directory visible in subdirectories of the share root. The default setting is `no`.
- `--revert-dot-snap-visible-child`
Sets the value to the system default for `--dot-snap-visible-child`.
- `--dot-snap-visible-root {yes | no}`
Specifies whether to make the `/ifs/.snapshot` directory visible at the root of the share. The default setting is `no`.
- `--revert-dot-snap-visible-root`
Sets the value to the system default for `--dot-snap-visible-root`.
- `--enable-security-signatures {yes | no}`
Indicates whether the server supports signed SMB packets.
- `--revert-enable-security-signatures`
Sets the value to the system default for `--enable-security-signatures`.
- `--guest-user <integer>`
Specifies the fully qualified user to use for guest access.
- `--revert-guest-user`
Sets the value to the system default for `--guest-user`.
- `--ignore-eas {yes | no}`

Specifies whether to ignore EAs on files.

`--revert-ignore-eas`
Sets the value to the system default for `--ignore-eas`.

`--onefs-cpu-multiplier` *<integer>*
Specifies the number of OneFS worker threads to configure based on the number of CPUs. Valid numbers are **1-4**.

`--revert-onefs-cpu-multiplier`
Sets the value to the system default for `--onefs-cpu-multiplier`.

`--onefs-num-workers` *<integer>*
Specifies the number of OneFS worker threads that are allowed to be configured. Valid numbers are **0-1024**. If set to **0**, the number of SRV workers will equal the value specified by `--onefs-cpu-multiplier` times the number of CPUs.

`--revert-onefs-num-workers`
Sets the value to the system default for `--onefs-num-workers`.

`--require-security-signatures` {yes | no}
Specifies whether packet signing is required. If set to **yes**, signing is always required. If set to **no**, signing is not required but clients requesting signing will be allowed to connect if the `--enable-security-signatures` option is set to **yes**.

`--revert-require-security-signatures`
Sets the value to the system default for `--require-security-signatures`.

`--server-string` *<string>*
Provides a description of the server.

`--revert-server-string`
Sets the value to the system default for `--revert-server-string`.

`--srv-cpu-multiplier` *<integer>*
Specifies the number of SRV worker threads to configure per CPU. Valid numbers are **1-8**.

`--revert_srv-cpu-multiplier`
Sets the value to the system default for `--srv-cpu-multiplier`.

`--srv-num-workers` *<integer>*
Specifies the number of OneFS worker threads that are allowed to be configured. Valid numbers are **0-1024**. If set to **0**, the number of SRV workers will equal the value specified by `--srv-cpu-multiplier` times the number of CPUs.

`--revert-srv-num-workers`
Sets the value to the system default for `--revert-srv-num-workers`.

`--support-multichannel` {yes | no}
Specifies whether Multichannel for SMB 3.0 is enabled on the cluster. SMB Multichannel is enabled by default.

`--revert-support-multichannel`
Set the value of `--support-multichannel` back to the default system value.

`--support-netbios` {yes | no}
Specifies whether to support the NetBIOS protocol.

```
--revert-support-netbios
```

Sets the value to the system default for `--support-netbios`.

```
--support-smb2 {yes | no}
```

Specifies whether to support the SMB 2.0 protocol. The default setting is `yes`.

```
--revert-support-smb2
```

Sets the value to the system default for `--support-smb2`.

isi smb settings global view

Displays the default SMB configuration settings.

Syntax

```
isi smb settings global view
```

Options

There are no options for this command.

isi smb settings shares modify

Modifies default settings for SMB shares.

Syntax

```
isi smb settings shares modify
[--access-based-enumeration {yes | no}]
[--revert-access-based-enumeration]
[--access-based-enumeration-root-only {yes | no}]
[--revert-access-based-enumeration-root-only]
[--allow-delete-readonly {yes | no}]
[--revert-allow-delete-readonly]
[--allow-execute-always {yes | no}]
[--revert-allow-execute-always]
[--change-notify {all | norecurse | none}]
[--revert-change-notify]
[--create-permissions {"default acl" | "inherit mode bits" | "use
create mask and mode"}]
[--revert-create-permissions]
[--directory-create-mask <integer>]
[--revert-directory-create-mask]
[--directory-create-mode <integer>]
[--revert-directory-create-mode]
[--file-create-mask <integer>]
[--revert-file-create-mask]
[--file-create-mode <integer>]
[--revert-file-create-mode]
[--hide-dot-files {yes | no}]
[--revert-hide-dot-files]
[--host-acl <host-acl>]
[--revert-host-acl]
[--clear-host-acl]
[--add-host-acl <string>]
[--remove-host-acl <string>]
[--impersonate-guest {always | "bad user" | never}]
[--revert-impersonate-guest]
[--impersonate-user <string>]
[--revert-impersonate-user]
[--mangle-byte-start <integer>]
[--revert-mangle-byte-start]
[--mangle-map <mangle-map>]
[--revert-mangle-map]
[--clear-mangle-map]
```

```
[--add-mangle-map <string>]
[--remove-mangle-map <string>]
[--ntfs-acl-support {yes | no}]
[--revert-ntfs-acl-support]
[--oplocks {yes | no}]\
[--revert-oplocks]
[--strict-flush {yes | no}]
[--revert-strict-flush]
[--strict-locking {yes | no}]
[--revert-strict-locking]
[--zone <string>]
```

Options

`--access-based-enumeration {yes | no}`
 Specifies whether access-based enumeration is enabled.

`--revert-access-based-enumeration`
 Sets the value to the system default for `--access-based-enumeration`.

`--access-based-enumeration-root-only {yes | no}`
 Specifies whether access-based enumeration is only enabled on the root directory of the share.

`--revert-access-based-enumeration-root-only`
 Sets the value to the system default for `--access-based-enumeration-root-only`.

`--allow-delete-readonly {yes | no}`
 Specifies whether read-only files can be deleted.

`--revert-allow-delete-readonly`
 Sets the value to the system default for `--allow-delete-readonly`.

`--allow-execute-always {yes | no}`
 Specifies whether a user with read access to a file can also execute the file.

`--revert-allow-execute-always`
 Sets the value to the system default for `--allow-execute-always`.

`--change-notify {norecurse | all | none}`
 Defines the change notify setting. The acceptable values are `norecurse`, `all`, and `none`.

`--revert-change-notify`
 Sets the value to the system default for `--change-notify`.

`--create-permissions {"default acl"|"inherit mode bits"|"use create mask and mode"}`
 Sets the default permissions to apply when a file or directory is created.

`--revert-create-permissions`
 Sets the value to the system default for `--create-permissions`.

`--directory-create-mask <integer>`
 Defines which mask bits are applied when a directory is created.

`--revert-directory-create-mask`
 Sets the value to the system default for `--directory-create-mask`.

`--directory-create-mode <integer>`
 Defines which mode bits are applied when a directory is created.

`--revert-directory-create-mode`
 Sets the value to the system default for `--directory-create-mode`.

`--file-create-mask <integer>`
 Defines which mask bits are applied when a file is created.

`--revert-file-create-mask`
 Sets the value to the system default for `--file-create-mask`.

`--file-create-mode <integer>`
 Defines which mode bits are applied when a file is created.

`--revert-file-create-mode`
 Sets the value to the system default for `--file-create-mode`.

`--hide-dot-files {yes | no}`
 Specifies whether to hide files that begin with a period—for example, UNIX configuration files.

`--revert-hide-dot-files`
 Sets the value to the system default for `--hide-dot-files`.

`--host-acl <string>`
 Specifies which hosts are allowed access. Specify `--host-acl` for each additional host ACL clause. This will replace any existing ACL.

`--revert-host-acl`
 Sets the value to the system default for `--host-acl`.

`--clear-host-acl <string>`
 Clears the value for an ACL expressing which hosts are allowed access.

`--add-host-acl <string>`
 Adds an ACE to the already-existing host ACL. Specify `--add-host-acl` for each additional host ACL clause to be added.

`--remove-host-acl <string>`
 Removes an ACE from the already-existing host ACL. Specify `--remove-host-acl` for each additional host ACL clause to be removed.

`--impersonate-guest {always | "bad user" | never}`
 Allows guest access to the share. The acceptable values are `always`, `"bad user"`, and `never`.

`--revert-impersonate-guest`
 Sets the value to the system default for `--impersonate-guest`.

`--impersonate-user <string>`
 Allows all file access to be performed as a specific user. This must be a fully qualified user name.

`--revert-impersonate-user`
 Sets the value to the system default for `--impersonate-user`.

`--mangle-byte-start <string>`
 Specifies the `wchar_t` starting point for automatic invalid byte mangling.

```

--revert-mangle-byte-start
    Sets the value to the system default for --mangle-byte-start.
--mangle-map <string>
    Maps characters that are valid in OneFS but are not valid in SMB names.
--revert-mangle-map
    Sets the value to the system default for --mangle-map.
--clear-mangle-map <string>
    Clears the values for character mangle map.
--add-mangle-map <string>
    Adds a character mangle map. Specify --add-mangle-map for each additional Add
    character mangle map.
--remove-mangle-map <string>
    Removes a character mangle map. Specify --remove-mangle-map for each
    additional Remove character mangle map.
--ntfs-acl-support {yes | no}
    Specifies whether ACLs can be stored and edited from SMB clients.
--revert-ntfs-acl-support
    Sets the value to the system default for --ntfs-acl-support.
--oplocks {yes | no}
    Specifies whether to allow oplock requests.
--revert-oplocks
    Sets the value to the system default for --oplocks.
--strict-flush {yes | no}
    Specifies whether to always honor flush requests.
--revert-strict-flush
    Sets the value to the system default for --strict-flush.
--strict-locking {yes | no}
    Specifies whether the server will check for and enforce file locks.
--revert-strict-locking
    Sets the value to the system default for --strict-locking.
--zone <string>
    Specifies the name of the access zone.

```

isi smb settings shares view

Displays default settings for all SMB shares or for SMB shares in a specified access zone.

Syntax

```

isi smb settings shares view
  [--zone <string>]

```

Options

```

--zone <string>

```

Specifies the name of the access zone. Displays only the settings for shares in the specified zone.

isi smb shares create

Creates an SMB share.

Syntax

```
isi smb shares create <name> <path>
  [--zone <string>]
  [--inheritable-path-acl {yes | no}]
  [--create-path]
  [--host-acl <string>]
  [--description <string>]
  [--csc-policy {none | documents | manual | programs}]
  [--allow-variable-expansion {yes | no}]
  [--auto-create-directory {yes | no}]
  [--browsable {yes | no}]
  [--allow-execute-always {yes | no}]
  [--directory-create-mask <integer>]
  [--strict-locking {yes | no}]
  [--hide-dot-files {yes | no}]
  [--impersonate-guest {always | "bad user" | never}]
  [--strict-flush {yes | no}]
  [--access-based-enumeration {yes | no}]
  [--access-based-enumeration-root-only {yes | no}]
  [--mangle-byte-start <string>]
  [--file-create-mask <integer>]
  [--create-permissions {"default acl" | "inherit mode bits"
    | "use create mask and mode"}]
  [--mangle-map <string>]
  [--impersonate-user <string>]
  [--change-notify <string>]
  [--oplocks {yes | no}]
  [--allow-delete-readonly {yes | no}]
  [--directory-create-mode <integer>]
  [--ntfs-acl-support {yes | no}]
  [--file-create-mode <integer>]
```

Options

<name>

Required. Specifies the name for the new SMB share.

<path>

Required. Specifies the full path of the SMB share to create, beginning at /ifs.

--zone <string>

Specifies the access zone the new SMB share is assigned to. If no access zone is specified, the new SMB share is assigned to the default `System` zone.

{--inheritable-path-acl | -i}{yes|no}

If set to `yes`, if the parent directory has an inheritable access control list (ACL), its ACL will be inherited on the share path. The default setting is `no`.

--create-path

Creates the SMB-share path if one doesn't exist.

--host-acl <string>

Specifies the ACL that defines host access. Specify `--host-acl` for each additional host ACL clause.

--description <string>

Specifies a description for the SMB share.

```
--csc-policy {none | documents | manual | programs}, -C {none |
documents | manual | programs}
```

Sets the client-side caching policy for the share. Valid values are `none`, `documents`, `manual`, and `programs`.

```
--allow-variable-expansion {yes | no}
```

Specifies automatic expansion of variables for home directories.

```
--directory-create-mask <integer>
```

Creates home directories automatically.

```
--browsable {yes | no}, -b {yes | no}
```

If set to `yes`, makes the share visible in net view and the browse list. The default setting is `yes`.

```
--allow-execute-always {yes | no}
```

If set to `yes`, allows a user with read access to a file to also execute the file. The default setting is `no`.

```
--directory-create-mask <integer>
```

Defines which mask bits are applied when a directory is created.

```
--strict-locking {yes | no}
```

If set to `yes`, directs the server to check for and enforce file locks. The default setting is `no`.

```
--hide-dot-files {yes | no}
```

If set to `yes`, hides files that begin with a decimal—for example, UNIX configuration files. The default setting is `no`.

```
--impersonate-guest {always | "bad user" | never}
```

Allows guest access to the share. The acceptable values are `always`, `"bad user"`, and `never`.

```
--strict-flush {yes | no}
```

If set to `yes`, flush requests are always honored. The default setting is `yes`.

```
--access-based-enumeration {yes | no}
```

If set to `yes`, enables access-based enumeration only on the files and folders that the requesting user can access. The default setting is `no`.

```
--access-based-enumeration-root-only {yes | no}
```

If set to `yes`, enables access-based enumeration only on the root directory of the SMB share. The default setting is `no`.

```
--mangle-byte-start <string>
```

Specifies the `wchar_t` starting point for automatic invalid byte mangling.

```
--file-create-mask <integer>
```

Defines which mask bits are applied when a file is created.

```
--create-permissions {"default acl" | "inherit mode bits" |
"use create mask and mode"}
```


Sets the default permissions to apply when a file or directory is created. Valid values are "default acl", "inherit mode bits", and "use create mask and mode"

`--mangle-map <string>`

Maps characters that are valid in OneFS but are not valid in SMB names.

`--impersonate-user <string>`

Allows all file access to be performed as a specific user. This value must be a fully qualified user name.

`--change-notify {norecurse | all | none}`

Defines the change notify setting. The acceptable values are `norecurse`, `all`, or `none`.

`--oplocks {yes | no}`

If set to `yes`, allows oplock requests. The default setting is `yes`.

`--allow-delete-readonly {yes | no}`

If set to `yes`, allows read-only files to be deleted. The default setting is `no`.

`--directory-create-mode <integer>`

Defines which mode bits are applied when a directory is created.

`--ntfs-acl-support {yes | no}`

If set to `yes`, allows ACLs to be stored and edited from SMB clients. The default setting is `yes`.

`--file-create-mode <integer>`

Defines which mode bits are applied when a file is created.

Example 2 Examples

The following sample command specifies that the subnet 10.7.215.0/24 is allowed access to the share, but all other subnets are denied access:

```
--host-acl=allow:10.7.216.0/24 --host-acl=deny:ALL
```

isi smb shares delete

Deletes an SMB share.

Syntax

```
isi smb shares delete <share>
  [--zone <string>]
  [--force]
```

Options

`<share>`

Specifies the name of the SMB share to delete.

`--zone <string>`

Specifies the access zone the SMB share is assigned to. If no access zone is specified, the system deletes the SMB share with the specified name assigned to the default `System` zone, if found.

```
{--force | -f}
```

Suppresses command-line prompts and messages.

Examples

The following command deletes a share named "test-smb" in the "example-zone" access zone without displaying a warning prompt:

```
isi smb shares delete test-smb --zone example-zone --force
```

isi smb shares list

Displays a list of SMB shares.

Syntax

```
isi smb shares list
  [--zone <string>]
  [--limit <integer>]
  [--sort {name | path | description}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
--zone <string>
```

Specifies the access zone. Displays all SMB shares in the specified zone. If no access zone is specified, the system displays all SMB shares in the default `System` zone.

```
{--limit | -l} <integer>
```

Specifies the maximum number of items to list.

```
--sort {name | path | description}
```

Specifies the field to sort items by.

```
{--descending | -d}
```

Sorts the data in descending order.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
--verbose | -v
```

Displays more detailed information.

isi smb shares modify

Modifies an SMB share's settings.

Syntax

```
isi smb shares modify <share>
  [--name <string>]
```

```

[--path <path>]
[--zone <string>]
[--new-zone <string>]
[--host-acl <host-acl>]
[--revert-host-acl]
[--clear-host-acl]
[--add-host-acl <string>]
[--remove-host-acl <string>]
[--description <string>]
[--csc-policy {manual | documents | programs | none}]
[--revert-csc-policy]
[--allow-variable-expansion {yes | no}]
[--revert-allow-variable-expansion]
[--auto-create-directory {yes | no}]
[--revert-auto-create-directory {yes | no}]
[--browsable {yes | no}]
[--revert-browsable]
[--allow-execute-always {yes | no}]
[--revert-allow-execute-always]
[--directory-create-mask <integer>]
[--revert-directory-create-mask]
[--strict-locking {yes | no}]
[--revert-strict-locking]
[--hide-dot-files {yes | no}]
[--revert-hide-dot-files]
[--impersonate-guest {always | "bad user" | never}]
[--revert-impersonate-guest]
[--strict-flush {yes | no}]
[--revert-strict-flush]
[--access-based-enumeration {yes | no}]
[--revert-access-based-enumeration]
[--access-based-enumeration-root-only {yes | no}]
[--revert-access-based-enumeration-root-only]
[--mangle-byte-start <integer>]
[--revert-mangle-byte-start]
[--file-create-mask <integer>]
[--revert-file-create-mask]
[--create-permissions {"default acl" | "inherit mode bits"
 | "use create mask and mode"}]
[--revert-create-permissions]
[--mangle-map <mangle-map>]
[--revert-mangle-map]
[--clear-mangle-map]
[--add-mangle-map <string>]
[--remove-mangle-map <string>]
[--impersonate-user <string>]
[--revert-impersonate-user]
[--change-notify {all | norecurse | none}]
[--revert-change-notify]
[--oplocks {yes | no}]
[--revert-oplocks]
[--allow-delete-readonly {yes | no}]
[--revert-allow-delete-readonly]
[--directory-create-mode <integer>]
[--revert-directory-create-mode]
[--ntfs-acl-support {yes | no}]
[--revert-ntfs-acl-support]
[--file-create-mode <integer>]
[--revert-file-create-mode]
[--verbose]

```

Options

<share>

Required. Specifies the name of the SMB share to modify.

--name *<name>*

Specifies the name for the SMB share.

`--path <path>`
 Specifies a new path for the SMB share, starting in `/ifs`.

`--zone <string>`
 Specifies the access zone that the SMB share is assigned to. If no access zone is specified, the system modifies the SMB share with the specified name assigned to the default `System` zone, if found.

`--new-zone <string>`
 Specifies the new access zone that SMB share will be reassigned to.

`--host-acl <host-acl>`
 An ACL expressing which hosts are allowed access. Specify `--host-acl` for each additional host ACL clause.

`--revert-host-acl`
 Sets the value to the system default for `--host-acl`.

`--clear-host-acl`
 Clears the value of an ACL that expresses which hosts are allowed access.

`--add-host-acl <string>`
 Adds an ACL expressing which hosts are allowed access. Specify `--add-host-acl` for each additional host ACL clause to add.

`--remove-host-acl <string>`
 Removes an ACL expressing which hosts are allowed access. Specify `--remove-host-acl` for each additional host ACL clause to remove.

`--description <string>`
 The description for this SMB share.

`--csc-policy, -C {manual | documents | programs | none}`
 Specifies the client-side caching policy for the shares.

`--revert-csc-policy`
 Sets the value to the system default for `--csc-policy`.

`{--allow-variable-expansion | -a} {yes | no}`
 Allows the automatic expansion of variables for home directories.

`--revert-allow-variable-expansion`
 Sets the value to the system default for `--allow-variable-expansion`.

`{--auto-create-directory | -d} {yes | no}`
 Automatically creates home directories.

`--revert-auto-create-directory`
 Sets the value to the system default for `--auto-create-directory`.

`{--browsable | -b} {yes | no}`
 The share is visible in the net view and the browse list.

`--revert-browsable`
 Sets the value to the system default for `--browsable`.

`--allow-execute-always {yes | no}`
 Allows users to execute files they have read rights for.

```

--revert-allow-execute-always
    Sets the value to the system default for --allow-execute-always.
--directory-create-mask <integer>
    Specifies the directory create mask bits.
--revert-directory-create-mask
    Sets the value to the system default for --directory-create-mask.
--strict-locking {yes | no}
    Specifies whether byte range locks contend against the SMB I/O.
--revert-strict-locking
    Sets the value to the system default for --strict-locking.
--hide-dot-files {yes | no}
    Hides files and directories that begin with a period ".".
--revert-hide-dot-files
    Sets the value to the system default for --hide-dot-files.
--impersonate-guest {always | "bad user" | never}
    Specifies the condition in which user access is done as the guest account.
--revert-impersonate-guest
    Sets the value to the system default for --impersonate-guest.
--strict-flush {yes | no}
    Handles the SMB flush operations.
--revert-strict-flush
    Sets the value to system default for --strict-flush.
--access-based-enumeration {yes | no}
    Specifies to only enumerate files and folders that the requesting user has access to.
--revert-access-based-enumeration
    Sets the value to the system default for --access-based-enumeration.
--access-based-enumeration-root-only {yes | no}
    Specifies access-based enumeration on only the root directory of the share.
--revert-access-based-enumeration-root-only
    Sets the value to the system default for --access-based-enumeration-root-only.
--mangle-byte-start <integer>
    Specifies the wchar_t starting point for automatic byte mangling.
--revert-mangle-byte-start
    Sets the value to the system default for --mangle-byte-start.
--file-create-mask <integer>
    Specifies the file create mask bits.
--revert-file-create-mask
    Sets the value to the system default for --file-create-mask.
--create-permissions {"default acl"|"inherit mode bits"|"use
create mask and mode"}

```

Sets the create permissions for new files and directories in a share.

`--revert-create-permissions`
Sets the value to the system default for `--create-permissions`.

`--mangle-map <mangle-map>`
The character mangle map. Specify `--mangle-map` for each additional character mangle map.

`--revert-mangle-map`
Sets the value to the system default for `--mangle-map`.

`--clear-mangle-map`
Clears the value for character mangle map.

`--add-mangle-map <string>`
Adds a character mangle map. Specify `--add-mangle-map` for each additional Add character mangle map.

`--remove-mangle-map <string>`
Removes a character mangle map. Specify `--remove-mangle-map` for each additional Remove character mangle map.

`--impersonate-user <string>`
The user account to be used as a guest account.

`--revert-impersonate-user`
Sets the value to the system default for `--impersonate-user`.

`--change-notify {all | norecurse | none}`
Specifies the level of change notification alerts on a share.

`--revert-change-notify`
Sets the value to the system default for `--change-notify`.

`--oplocks {yes | no}`
Supports oplocks.

`--revert-oplocks`
Sets the value for the system default of `--oplocks`.

`--allow-delete-readonly {yes | no}`
Allows the deletion of read-only files in the share.

`--revert-allow-delete-readonly`
Sets the value for the system default of `--allow-delete-readonly`.

`--directory-create-mode <integer>`
Specifies the directory create mode bits.

`--revert-directory-create-mode`
Sets the value for the system default of `--directory-create-mode`.

`--ntfs-acl-support {yes | no}`
Supports NTFS ACLs on files and directories.

`--revert-ntfs-acl-support`
Sets the value for the system default of `--ntfs-acl-support`.

`--file-create-mode <integer>`

Specifies the file create mode bits.

`--revert-file-create-mode`

Sets the value for the system default of `--file-create-mode`.

isi smb shares permission create

Creates permissions for an SMB share.

Syntax

```
isi smb shares permission create <share> {<user> | --group <name>
  | --gid <id> | --uid <id> | --sid <string> | --wellknown <string>}
  {--run-as-root | --permission-type {allow | deny}}
  --permission {full | change | read}}
[--zone <zone>]
[--verbose]
```

Options

<share>

Specifies the name of the SMB share.

<user>

Specifies a user by name.

`--group <name>`

Specifies a group by name.

`--gid <id>`

Specifies a group by UNIX group identifier.

`--uid <id>`

Specifies a user by UNIX user identifier.

`--sid <string>`

Specifies an object by its Windows security identifier.

`--wellknown <string>`

Specifies a well-known user, group, machine, or account name.

`{--permission-type | -d} {deny | allow}`

Specifies whether to allow or deny a permission.

`{--permission | -p} {read | full | change}`

Specifies the level of control to allow or deny.

`--run-as-root {yes | no}`

If set to `yes`, allows the account to run as root. The default setting is `no`.

`--zone <zone>`

Specifies an access zone.

`{--verbose | -v}`

Displays more detailed information.

isi smb shares permission delete

Deletes user or group permissions for an SMB share.

Syntax

```
isi smb shares permission delete <share> {<user> | --group <name>
  |--gid <id> | --uid <id> | --sid <string> | --wellknown <string>}
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<share>

Required. Specifies the SMB share name.

<user>

Specifies a user by name.

--group *<name>*

Specifies a group by name.

--gid *<id>*

Specifies a group by UNIX group identifier.

--uid *<id>*

Specifies a user by UNIX user identifier.

--sid *<string>*

Specifies an object by its Windows security identifier.

--wellknown *<string>*

Specifies a well-known user, group, machine, or account name.

--zone *<string>*

Specifies an access zone.

{--force | -f}

Specifies that you want the command to execute without prompting for confirmation.

{--verbose | -v}

Displays more detailed information.

isi smb shares permission list

Displays a list of permissions for an SMB share.

Syntax

```
isi smb shares permission list <share>
  [--zone <zone>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

<share>

Specifies the name of the SMB share to display.

--zone *<zone>*

Specifies the access zone to display.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

isi smb shares permission modify

Modifies permissions for an SMB share.

Syntax

```
isi smb shares permission modify <share> {<user> | --group <name>
| --gid <id> | --uid <id> | --sid <string> | --wellknown <string>}
{--run-as-root | --permission-type {allow | deny}
--permission {full | change | read}}
[--zone <zone>]
[--verbose]
```

Options

<share>

Specifies the name of the SMB share.

<user>

Specifies a user by name.

`--group <name>`

Specifies a group by name.

`--gid <id>`

Specifies a group by UNIX group identifier.

`--uid <id>`

Specifies a user by UNIX user identifier.

`--sid <string>`

Specifies an object by its Windows security identifier.

`--wellknown <string>`

Specifies a well-known user, group, machine, or account name.

`{--permission-type | -d} {deny | allow}`

Specifies whether to allow or deny a permission.

`{--permission | -p} {read | full | change}`

Specifies the level of control to allow or deny.

`--run-as-root {yes | no}`

If set to `yes`, allows the account to run as root. The default setting is `no`.

`--zone <zone>`

Specifies an access zone.

`{--verbose | -v}`

Displays more detailed information.

isi smb shares permission view

Displays a single permission for an SMB share.

Syntax

```
isi smb shares permission view <share> {<user> |
--group <name> | --gid <integer>
| --uid <integer> | --sid <string>
| --wellknown <string>}
[--zone <string>]
```

Options

<share>

Specifies the name of the SMB share.

<user>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <integer>

Specifies a numeric group identifier.

--uid <integer>

Specifies a numeric user identifier.

--sid <string>

Specifies a security identifier.

--wellknown <string>

Specifies a well-known user, group, machine, or account name.

--zone <string>

Specifies an access zone.

isi smb shares view

Displays information about an SMB share.

Syntax

```
isi smb shares view <share>
[--zone <string>]
```

Options

<share>

Specifies the name of the SMB share to view.

--zone <string>

Specifies the access zone that the SMB share is assigned to. If no access zone is specified, the system displays the SMB share with the specified name assigned to the default `System` zone, if found.

NFS

OneFS provides an NFS server so you can share files on your cluster with NFS clients that adhere to the RFC1813 (NFSv3) and RFC3530 (NFSv4) specifications.

In OneFS, the NFS server is fully optimized as a multi-threaded service running in user space instead of the kernel. This architecture load balances the NFS service across all nodes of the cluster, providing the stability and scalability necessary to manage up to thousands of connections across multiple NFS clients.

NFS mounts execute and refresh quickly, and the server constantly monitors fluctuating demands on NFS services and makes adjustments across all nodes to ensure continuous, reliable performance. Using a built-in process scheduler, OneFS helps ensure fair allocation of node resources so that no client can seize more than its fair share of NFS services.

The NFS server also supports access zones defined in OneFS, so that clients can access only the exports appropriate to their zone. For example, if NFS exports are specified for Zone 2, only clients assigned to Zone 2 can access these exports.

To simplify client connections, especially for exports with large path names, the NFS server also supports aliases, which are shortcuts to mount points that clients can specify directly.

For secure NFS file sharing, OneFS supports NIS and LDAP authentication providers.

NFS exports

You can manage individual NFS export rules that define mount-points (paths) available to NFS clients and how the server should perform with these clients.

In OneFS, you can create, delete, list, view, modify, and reload NFS exports.

NFS export rules are zone-aware. Each export is associated with a zone, can only be mounted by clients on that zone, and can only expose paths below the zone root. By default, any export command applies to the client's current zone.

Each rule must have at least one path (mount-point), and can include additional paths. You can also specify that all subdirectories of the given path or paths are mountable. Otherwise, only the specified paths are exported, and child directories are not mountable.

An export rule can specify a particular set of clients, enabling you to restrict access to certain mount-points or to apply a unique set of options to these clients. If the rule does not specify any clients, then the rule applies to all clients that connect to the server. If the rule does specify clients, then that rule is applied only to those clients.

NFS aliases

You can create and manage aliases as shortcuts for directory path names in OneFS. If those path names are defined as NFS exports, NFS clients can specify the aliases as NFS mount points.

NFS aliases are designed to give functional parity with SMB share names within the context of NFS. Each alias maps a unique name to a path on the file system. NFS clients can then use the alias name in place of the path when mounting.

Aliases must be formed as top-level Unix path names, having a single forward slash followed by name. For example, you could create an alias named `/q4` that maps

to `/ifs/data/finance/accounting/winter2015` (a path in OneFS). An NFS client could mount that directory through either of:

```
mount cluster_ip:/q4
```

```
mount cluster_ip:/ifs/data/finance/accounting/winter2015
```

Aliases and exports are completely independent. You can create an alias without associating it with an NFS export. Similarly, an NFS export does not require an alias.

Each alias must point to a valid path on the file system. While this path is absolute, it must point to a location beneath the zone root (`/ifs` on the System zone). If the alias points to a path that does not exist on the file system, any client trying to mount the alias would be denied in the same way as attempting to mount an invalid full pathname.

NFS aliases are zone-aware. By default, an alias applies to the client's current access zone. To change this, you can specify an alternative access zone as part of creating or modifying an alias.

Each alias can only be used by clients on that zone, and can only apply to paths below the zone root. Alias names are unique per zone, but the same name can be used in different zones—for example, `/home`.

When you create an alias in the web administration interface, the alias list displays the status of the alias. Similarly, using the `--check` option of the `isi nfs aliases` command, you can check the status of an NFS alias (status can be: good, illegal path, name conflict, not exported, or path not found).

NFS log files

OneFS writes log messages associated with NFS events to a set of files in `/var/log`.

With the log level option, you can now specify the detail at which log messages are output to log files. The following table describes the log files associated with NFS.

Log file	Description
<code>nfs.log</code>	Primary NFS server functionality (v3, v4, mount)
<code>rpc_lockd.log</code>	NFS v3 locking events through the NLM protocol
<code>rpc_statd.log</code>	NFS v3 reboot detection through the NSM protocol
<code>isi_netgroup_d.log</code>	Netgroup resolution and caching

Managing the NFS service

You can enable or disable the NFS service and specify the NFS versions to support, including NFSv3 and NFSv4. NFS settings are applied across all nodes in the cluster.

View NFS settings

You can view the global NFS settings that are applied to all nodes in the cluster.

Procedure

1. Run the `isi nfs settings global view` command.

The system displays output similar to the following example:

```
Lock Protection Level: 2
    NFSv3 Enabled: Yes
    NFSv4 Enabled: No
    NFS Service Enabled: Yes
```

Configure NFS file sharing

You can set the lock protection level and NFS version support.

These settings are applied across all nodes in the cluster.

Procedure

1. Run the `isi nfs settings global modify` command.

The following command sets the lock protection level to 4:

```
isi nfs settings global modify --lock-protection=4
```

The following command enables NFSv4 support:

```
isi nfs settings global modify --nfsv4-enabled=yes
```

Enable or disable the NFS service

In OneFS, the NFSv3 service is enabled by default. You can also enable NFSv4.

Note

You can determine whether NFS services are enabled or disabled by running the `isi nfs settings global view` command.

Procedure

1. Run the `isi nfs settings global modify` command.

The following command disables the NFSv3 service:

```
isi nfs settings global modify --nfsv3-enabled no
```

The following command enables the NFSv4 service:

```
isi nfs settings global modify --nfsv4-enabled yes
```

Managing NFS exports

You can create NFS exports, view and modify export settings, and delete exports that are no longer needed.

The `/ifs` directory is the top-level directory for data storage in OneFS, and is also the path defined in the default export. By default, the `/ifs` export disallows root access, but other enables UNIX clients to mount this directory and any subdirectories beneath it.

Note

We recommend that you modify the default export to limit access only to trusted clients, or to restrict access completely. To help ensure that sensitive data is not compromised, other exports that you create should be lower in the OneFS file hierarchy, and can be protected by access zones or limited to specific clients with either root, read-write, or read-only access, as appropriate.

Configure default NFS export settings

The default NFS export settings are applied to new NFS exports. You can override these settings when you create or modify an export.

You can view the current default export settings by running the `isi nfs settings export view` command.

⚠ CAUTION

We recommend that you not modify default export settings unless you are sure of the result.

Procedure

- Run the `isi nfs settings export modify` command.

The following command specifies a maximum export file size of one terabyte:

```
isi nfs settings export modify --max-file-size 1099511627776
```

The following command restores the maximum export file size to the system default:

```
isi nfs settings export modify --revert-max-file-size
```

Create a root-squashing rule for an export

By default, the NFS service implements a root-squashing rule for the default NFS export. This prevents root users on NFS clients from exercising root privileges on the NFS server.

In OneFS, the default NFS export is `/ifs`, the top-level directory where cluster data is stored.

Procedure

1. Use the `isi nfs exports view` command to view the current settings of the default export.

The following command displays the settings of the default export:

```
isi nfs exports view 1
```

2. Confirm the following default values for these settings, which show that root is mapped to nobody, thereby restricting root access:

```
Map Root
  Enabled: True
  User: Nobody
  Primary Group: -
  Secondary Groups: -
```

3. If the root-squashing rule, for some reason, is not in effect, you can implement it for the default NFS export by running the `isi nfs export modify` command, as follows:

```
isi nfs exports modify 1 --map-root-enabled true --map-root nobody
```

Results

With these settings, regardless of the users' credentials on the NFS client, they would not be able to gain root privileges on the NFS server.

Create an NFS export

You can create the NFS exports necessary to service the needs of your organization.

Before you begin

Each directory path that you designate for an export must already exist in the `/ifs` directory tree.

A directory path can be used by more than one export, provided those exports do not have any of the same explicit clients.

Procedure

1. Run the `isi nfs exports create` command.

The following command creates an export supporting client access to multiple paths and their subdirectories:

```
isi nfs exports create /ifs/data/projects,/ifs/home --all-dirs=yes
```

2. (Optional) To view the export ID, which is required for modifying or deleting the export, run the `isi nfs exports list` command.

Check NFS exports for errors

You can check for errors in NFS exports, such as conflicting export rules, invalid paths, and unresolvable hostnames and netgroups.

This procedure is available only through the CLI.

Procedure

1. Run the `isi nfs exports check` command.

In the following example output, no errors were found:

```
ID Message
-----
-----
Total: 0
```

In the following example output, export 1 contains a directory path that does not currently exist:

```
ID Message
-----
1  '/ifs/test' does not exist
-----
Total: 1
```

Modify an NFS export

You can modify the settings for an individual NFS export.

Procedure

1. Run the `isi nfs exports modify` command.

The following command adds a client with read-write access to NFS export 2.

```
isi nfs exports modify 2 --add-read-write-clients 10.1.249.137
```

This command would override the export's access-restriction setting if there was a conflict. For example, if the export was created with read-write access disabled, the client, 10.1.249.137, would still have read-write permissions on the export.

Delete an NFS export

You can delete NFS exports that are no longer needed.

Before you begin

You need the export ID number to delete the export. To display a list of exports and their ID numbers, you can run the `isi nfs exports list` command.

Procedure

1. Run the `isi nfs exports delete` command.

The following command deletes an export whose ID is 2:

```
isi nfs exports delete 2
```

The following command deletes an export whose ID is 3 without displaying a confirmation prompt. Be careful when using the `--force` option.

```
isi nfs exports delete 3 --force
```

2. If you did not specify the `--force` option, type **yes** at the confirmation prompt.

Managing NFS aliases

You can create NFS aliases to simplify exports that clients connect to. An NFS alias maps an absolute directory path to a simple directory path.

For example, suppose you created an NFS export to `/ifs/data/hq/home/archive/first-quarter/finance`. You could create the alias `/finance1` to map to that directory path.

NFS aliases can be created in any access zone, including the System zone.

Create an NFS alias

You can create an NFS alias to map a long directory path to a simple pathname.

Aliases must be formed as a simple Unix-style directory path, for example, `/home`.

Procedure

1. Run the `isi nfs aliases create` command.

The following command creates an alias to a full pathname in OneFS in an access zone named `hq-home`:

```
isi nfs aliases create /home /ifs/data/offices/hq/home --zone hq-home
```

When you create an NFS alias, OneFS performs a health check. If, for example, the full path that you specify is not a valid path, OneFS issues a warning:

```
Warning: health check on alias '/home' returned 'path not found'
```

Nonetheless, the alias is created, and you can create the directory that the alias points to at a later time.

Modify an NFS alias

You can modify an NFS alias, for example, if an export directory path has changed.

Aliases must be formed as a simple Unix-style directory path, for example, `/home`.

Procedure

1. Run the `isi nfs aliases modify` command.

The following command changes the name of an alias in the access zone `hq-home`:

```
isi nfs aliases modify /home --zone hq-home --name /home1
```

When you modify an NFS alias, OneFS performs a health check. If, for example, the path to the alias is not valid, OneFS issues a warning:

```
Warning: health check on alias '/home' returned 'not exported'
```

Nonetheless, the alias is modified, and you can create the export at a later time.

Delete an NFS alias

You can delete an NFS alias.

If an NFS alias is mapped to an NFS export, deleting the alias can disconnect clients that used the alias to connect to the export.

Procedure

1. Run the `isi nfs aliases delete` command.

The following command deletes the alias `/home` in an access zone named `hq-home`:

```
isi nfs aliases delete /home --zone hq-home
```

When you delete an NFS alias, OneFS asks you to confirm the operation:

```
Are you sure you want to delete NFS alias /home? (yes/[no])
```

2. Type **yes**, and then press ENTER.

The alias is deleted, unless an error condition was found, for example, you typed the name of the alias incorrectly.

List NFS aliases

You can view a list of NFS aliases that have already been defined for a particular zone. Aliases in the system zone are listed by default.

Procedure

1. Run the `isi nfs aliases list` command.

The following command lists aliases that have been created in the system zone (the default):

```
isi nfs aliases list
```

This command lists aliases that have been created in an access zone named `hq-home`:

```
isi nfs aliases list --zone hq-home
```

Output from the command looks similar to the following example:

Zone	Name	Path
hq-home	/home	/ifs/data/offices/newyork
hq-home	/root_alias	/ifs/data/offices
hq-home	/project	/ifs/data/offices/project

Total: 3		

View an NFS alias

You can view the settings of an NFS alias in the specified access zone.

Procedure

1. Run the `isi nfs aliases view` command.

The following command provides information on an alias in the access zone, `hq-home`, including the health of the alias:

```
isi nfs aliases view /projects --zone hq-home --check
```

Output from the command looks similar to the following example:

Zone	Name	Path	Health
hq-home	/projects	/ifs/data/offices/project	good

Total: 1			

NFS commands

You can access and configure the NFS file sharing service through the NFS commands.

isi nfs aliases create

Creates an NFS alias, which is a shortcut to a directory path under `/ifs`. An alias' target directory can be made into an NFS export, in which case NFS clients can mount the alias name.

Syntax

```
isi nfs aliases create <name> <path>
  [--zone <string>]
  [--force]
```

Options

<name>

The name of the alias. Alias names must be formed as Unix root directory with a single forward slash followed by the name. For example, `/home`.

<path>

The OneFS directory pathname the alias links to. The pathname must be an absolute path below the access zone root. For example, `/ifs/data/ugroup1/home`.

`--zone`

The access zone in which the alias is active.

`--force`

Forces creation of the alias without requiring confirmation.

Example

The following command creates an alias in a zone named `ugroup1`:

```
isi nfs aliases create /home /ifs/data/ugroup1/home
--zone ugroup1
```

isi nfs aliases delete

Deletes a previously created alias.

Syntax

```
isi nfs aliases delete <name>
  [--zone <string>]
  [--force]
```

Options

<name>

The name of the alias to be deleted.

`--zone <string>`

The access zone in which the alias is active.

`--force`

Forces the alias to be deleted without having to confirm the operation.

Example

The following command deletes an alias from a zone named `ugroup1`.

```
isi nfs aliases delete /projects --zone ugroup1
```

If you do not use the `--force` option in the command, OneFS asks you to confirm the deletion. Type **yes** to confirm or **no** to decline the operation, then press ENTER.

isi nfs aliases list

List NFS aliases available in the current access zone.

Syntax

```
isi nfs aliases list
  [--check]
  [--zone <string>]
  [--limit <integer>]
  [--sort {zone | name | path | health}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

`--check`

For the current zone, displays a list of aliases and their health status.

`--zone <string>`

The access zone in which the alias is active.

`--limit <integer>`

Displays no more than the specified number of NFS aliases.

`--sort {zone | name | path | health}`

Specifies the field to sort by.

`--descending`

Specifies to sort the data in descending order.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`--no-header`

Displays table and CSV output without headers.

`--no-footer`

Displays table output without footers.

Example

The following command displays a table of the aliases in a zone named `ugroup1` including their health status.

```
isi nfs aliases list --zone ugroup1 --check
```

Output from the command is similar to the following example:

```

Zone      Name          Path          Health
-----
ugroup1  /home         /ifs/data/offices/newyork good
ugroup1  /root_alias   /ifs/data/offices      good
ugroup1  /project      /ifs/data/offices/project good
-----
Total: 3

```

isi nfs aliases modify

Modify the name, zone, or absolute path of an alias.

Syntax

```

isi nfs aliases modify <alias>
  [--zone <string>]
  [--new-zone <string>]
  [--name <string>]
  [--path <string>]

```

Options

<alias>

The current name of the alias, for example, /home.

--zone <string>

The access zone in which the alias is currently active.

--new-zone <string>

The new access zone in which the alias is to be active.

--name <string>

A new name for the alias.

--path <string>

The new OneFS directory pathname the alias should link to. The pathname must be an absolute path below the access zone root. For example, /ifs/data/ugroup2/home.

Example

The following command modifies the zone, name, and path of an existing alias:

```

isi nfs aliases modify /home --name /users --zone ugroup1 --new-zone
ugroup2
--path /ifs/data/ugroup2/users

```

isi nfs aliases view

View information about an alias in the current zone.

Syntax

```

isi nfs aliases view <name>
  [--zone <string>]
  [--check]

```

Options

<name>

The name of the alias.

--zone <string>

The access zone in which the alias is active.

`--check`

Include the health status of the alias.

Example

The following command displays a table of information, including the health status, of an alias named `/projects` in the current zone.

```
isi nfs aliases view /projects --check
```

isi nfs exports check

Checks for and lists configuration errors for NFS exports, including conflicting export rules, bad paths, unresolvable host names, and unresolvable net groups.

Syntax

```
isi nfs exports check
  [--limit <integer>]
  [--zone <string>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--limit <integer>`

Displays no more than the specified number of NFS exports.

`--zone <string>`

Specifies the access zone in which the export was created.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`--no-header`

Displays table and CSV output without headers.

`--no-footer`

Displays table output without footers.

`--verbose`

Displays more detailed information.

Examples

The following command checks the exports in a zone named `zone-1`:

```
isi nfs exports check --zone zone-1
```

If the check finds no problems, it returns an empty table. If, however, the check finds a problem, it returns a display similar to the following:

```
ID      Message
-----
3      '/ifs/data/project' does not exist
-----
Total: 1
```

isi nfs exports create

Creates an NFS export.

Note

To view the default NFS export settings that will be applied when creating an export, run the `isi nfs settings export view` command.

Syntax

```
isi nfs exports create <paths>
  [--block-size <size>]
  [--can-set-time {yes | no}]
  [--case-insensitive {yes | no}]
  [--case-preserving {yes | no}]
  [--chown-restricted {yes | no}]
  [--directory-transfer-size <size>]
  [--link-max <integer>]
  [--max-file-size <size>]
  [--name-max-size <integer>]
  [--no-truncate {yes | no}]
  [--return-32bit-file-ids {yes | no}]
  [--symlinks {yes | no}]
  [--zone <string>]
  [--clients <string>]
  [--description <string>]
  [--root-clients <string>]
  [--read-write-clients <string>]
  [--read-only-clients <string>]
  [--all-dirs {yes | no}]
  [--encoding <string>]
  [--security-flavors {unix | krb5 | krb5i | krb5p}]
  [--snapshot <snapshot>]
  [--map-lookup-uid {yes | no}]
  [--map-retry {yes | no}]
  [--map-root-enabled {yes | no}]
  [--map-non-root-enabled {yes | no}]
  [--map-failure-enabled {yes | no}]
  [--map-all <identity>]
  [--map-root <identity>]
  [--map-non-root <identity>]
  [--map-failure <identity>]
  [--map-full {yes | no}]
  [--commit-asynchronous {yes | no}]
  [--read-only {yes | no}]
  [--readdirplus {yes | no}]
  [--read-transfer-max-size <size>]
  [--read-transfer-multiple <integer>]
  [--read-transfer-size <size>]
  [--setattr-asynchronous {yes | no}]
  [--time-delta <time delta>]
  [--write-datasync-action {datasync | filesync | unstable}]
  [--write-datasync-reply {datasync | filesync}]
  [--write-filesync-action {datasync | filesync | unstable}]
  [--write-filesync-reply filesync]
  [--write-unstable-action {datasync | filesync | unstable}]
  [--write-unstable-reply {datasync | filesync | unstable}]
  [--write-transfer-max-size <size>]
  [--write-transfer-multiple <integer>]
  [--write-transfer-size <size>]
  [--force]
  [--verbose]
```

Options

`--paths <paths>...`

Required. Specifies the path to be exported, starting at `/ifs`. This option can be repeated to specify multiple paths.

`--block-size <size>`

Specifies the block size, in bytes.

`--can-set-time {yes | no}`

If set to `yes`, enables the export to set time. The default setting is `no`.

`--case-insensitive {yes | no}`

If set to `yes`, the server will report that it ignores case for file names. The default setting is `no`.

`--case-preserving {yes | no}`

If set to `yes`, the server will report that it always preserves case for file names. The default setting is `no`.

`--chown-restricted {yes | no}`

If set to `yes`, the server will report that only the superuser can change file ownership. The default setting is `no`.

`--directory-transfer-size <size>`

Specifies the preferred directory transfer size. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 128K.

`--link-max <integer>`

The reported maximum number of links to a file.

`--max-file-size <size>`

Specifies the maximum allowed file size on the server (in bytes). If a file is larger than the specified value, an error is returned.

`--name-max-size <integer>`

The reported maximum length of characters in a filename.

`--no-truncate {yes | no}`

If set to `yes`, too-long file names will result in an error rather than be truncated.

`--return-32bit-file-ids {yes | no}`

Applies to NFSv3 and NFSv4. If set to `yes`, limits the size of file identifiers returned from `readdir` to 32-bit values. The default value is `no`.

Note

This setting is provided for backward compatibility with older NFS clients, and should not be enabled unless necessary.

`--symlinks {yes | no}`

If set to `yes`, advertises support for symlinks. The default setting is `no`.

`--zone <string>`

Access zone in which the export should apply. The default zone is `system`.

`--clients <string>`

Specifies a client to be allowed access through this export. Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple clients by repeating this option.

Note

This option replaces the entire list of clients. To add or remove a client from the list, specify `--add-clients` or `--remove-clients`.

`--description <string>`

The description for this NFS export.

`--root-clients <string>`

Allows the root user of the specified client to execute operations as the root user of the cluster. This option overrides the `--map-all` and `--map-root` option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

`--read-write-clients <string>`

Grants read/write privileges to the specified client for this export. This option overrides the `--read-only` option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

`--read-only-clients <string>`

Makes the specified client read-only for this export. This option overrides the `--read-only` option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

`--all-dirs {yes | no}`

If set to `yes`, this export will cover all directories. The default setting is `no`.

`--encoding <string>`

Specifies the character encoding of clients connecting through this NFS export.

Valid values and their corresponding character encodings are provided in the following table. These values are taken from the node's `/etc/encodings.xml` file, and are not case sensitive.

Value	Encoding
cp932	Windows-SJIS
cp949	Windows-949
cp1252	Windows-1252
euc-kr	EUC-KR
euc-jp	EUC-JP
euc-jp-ms	EUC-JP-MS
utf-8-mac	UTF-8-MAC
utf-8	UTF-8
iso-8859-1	ISO-8859-1 (Latin-1)

Value	Encoding
iso-8859-2	ISO-8859-2 (Latin-2)
iso-8859-3	ISO-8859-3 (Latin-3)
iso-8859-4	ISO-8859-4 (Latin-4)
iso-8859-5	ISO-8859-5 (Cyrillic)
iso-8859-6	ISO-8859-6 (Arabic)
iso-8859-7	ISO-8859-7 (Greek)
iso-8859-8	ISO-8859-8 (Hebrew)
iso-8859-9	ISO-8859-9 (Latin-5)
iso-8859-10	ISO-8859-10 (Latin-6)
iso-8859-13	ISO-8859-13 (Latin-7)
iso-8859-14	ISO-8859-14 (Latin-8)
iso-8859-15	ISO-8859-15 (Latin-9)
iso-8859-16	ISO-8859-16 (Latin-10)

`--security-flavors {unix | krb5 | krb5i | krb5p} ...`

Specifies a security flavor to support. To support multiple security flavors, repeat this option for each additional entry. The following values are valid:

sys

Sys or UNIX authentication.

krb5

Kerberos V5 authentication.

krb5i

Kerberos V5 authentication with integrity.

krb5p

Kerberos V5 authentication with privacy.

`--snapshot {<snapshot> | <snapshot-alias>}`

Specifies the ID of a snapshot or snapshot alias to export. If you specify this option, directories will be exported in the state captured in either the specified snapshot or the snapshot referenced by the specified snapshot alias. If the snapshot does not capture the exported path, the export will be inaccessible to users.

If you specify a snapshot alias, and the alias is later modified to reference a new snapshot, the new snapshot will be automatically applied to the export.

Because snapshots are read-only, clients will not be able to modify data through the export unless you specify the ID of a snapshot alias that references the live version of the file system.

Specify *<snapshot>* or *<snapshot-alias>* as the ID or name of a snapshot or snapshot alias.

`--map-lookup-uid {yes | no}`

If set to *yes*, incoming UNIX user identifiers (UIDs) will be looked up locally. The default setting is *no*.

`--map-retry {yes | no}`

If set to `yes`, the system retries failed user-mapping lookups. The default setting is `no`.

`--map-root-enabled {yes | no}`

Enable/disable mapping incoming root users to a specific account.

`--map-non-root-enabled {yes | no}`

Enable/disable mapping incoming non-root users to a specific account.

`--map-failure-enabled {yes | no}`

Enable/disable mapping users to a specific account after failing an auth lookup.

`--map-all <identity>`

Specifies the default identity that operations by any user will execute as. If this option is not set to `root`, you can allow the root user of a specific client to execute operations as the root user of the cluster by including the client in the `--root-clients` list.

`--map-root <identity>`

Map incoming root users to a specific user and/or group ID.

`--map-non-root <identity>`

Map non-root users to a specific user and/or group ID.

`--map-failure <identity>`

Map users to a specific user and/or group ID after a failed auth attempt.

`--map-full {yes | no}`

Determines how user mapping is accomplished if a user is specified in an export option such as `--map-root` or `--map-all`. When enabled, a user mapping queries the OneFS user database and retrieves users from the applicable authentication subsystem, such as local authentication or Active Directory. When disabled, only local authentication is queried.

The default setting is `yes`.

`--commit-asynchronous {yes | no}`

If set to `yes`, enables commit data operations to be performed asynchronously. The default setting is `no`

`--read-only {yes | no}`

Determines the default privileges for all clients accessing the export.

If set to `yes`, you can grant read/write privileges to a specific client by including the client in the `--read-write-clients` list.

If set to `no`, you can make a specific client read-only by including the client in the `--read-only-clients` list. The default setting is `no`.

`--readdirplus {yes | no}`

Applies to NFSv3 only. If set to `yes`, enables processing of readdir-plus requests. The default setting is `no`.

`--read-transfer-max-size <size>`

Specifies the maximum read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

`--read-transfer-multiple <integer>`
 Specifies the suggested multiple read size to report to NFSv3 and NFSv4 clients. Valid values are 0–4294967295. The initial default value is 512.

`--read-transfer-size <size>`
 Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 128K.

`--setattr-asynchronous {yes | no}`
 If set to `yes`, performs set-attributes operations asynchronously. The default setting is `no`.

`--time-delta <float>`
 Specifies server time granularity, in seconds.

`--write-datasync-action {datasync | filesync | unstable}`
 Applies to NFSv3 and NFSv4 only. Specifies an alternate `datasync` write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `datasync`, which performs the request as specified.

`--write-datasync-reply {datasync | filesync}`
 Applies to NFSv3 and NFSv4 only. Specifies an alternate `datasync` reply method. The following values are valid:

- `datasync`
- `filesync`

The default value is `datasync` (does not respond differently).

`--write-filesync-action {datasync | filesync | unstable}`
 Applies to NFSv3 and NFSv4 only. Specifies an alternate `filesync` write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `filesync`, which performs the request as specified.

`--write-filesync-reply {filesync}`
 Applies to NFSv3 and NFSv4 only. Specifies an alternate `filesync` reply method. The only valid value is `filesync` (does not respond differently).

`--write-unstable-action {datasync | filesync | unstable}`
 Specifies an alternate `unstable`-write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable`, which performs the request as specified.

`--write-unstable-reply {datasync | filesync | unstable}`

Specifies an alternate unstable-reply method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable` (does not respond differently).

`--write-transfer-max-size <size>`

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

`--write-transfer-multiple <integer>`

Specifies the suggested write transfer multiplier to report to NFSv3 and NFSv4 clients. Valid values are `0-4294967295`. The initial default value is `512`.

`--write-transfer-size <size>`

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

`{--force | -f}`

If set to `no` (default), a confirmation prompt displays when the command runs. If set to `yes`, the command executes without prompting for confirmation.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command creates an NFS export for a particular zone and set of clients:

```
isi nfs exports create /ifs/data/ugroup1/home
--description 'Access to home dirs for user group 1'
--zone ugroup1 --clients 10.1.28.1 --clients 10.1.28.2
```

The following command creates an NFS export with multiple directory paths and a custom security type (Kerberos 5):

```
isi nfs exports create /ifs/data/projects /ifs/data/templates
--security-flavors krb5
```

isi nfs exports delete

Deletes an NFS export.

Syntax

```
isi nfs exports delete <id>
[--zone <string>]
[--force]
[--verbose]
```

Options

`<id>`

Specifies the ID of the NFS export to delete. You can use the `isi nfs exports list` command to view a list of exports and their IDs in the current zone.

`--zone <string>`

Specifies the access zone in which the export was created. Without this switch, the command defaults to the current zone.

`--force`

Suppresses command-line prompts and messages.

`--verbose`

Displays more detailed information.

isi nfs exports list

Displays a list of NFS exports.

Syntax

```
isi nfs exports list
  [--zone <string>]
  [--limit <integer>]
  [--sort <field>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--zone <string>`

Specifies the name of the access zone in which the export was created.

`{--limit <integer>`

Displays no more than the specified number of NFS exports.

`--sort <field>`

Specifies the field to sort by. Valid values are as follows:

- id
- zone
- paths
- description
- clients
- root_clients
- read_only_clients
- read_write_clients
- unresolved_clients
- all_dirs
- block_size
- can_set_time
- commit_asynchronous
- directory_transfer_size

- encoding
 - map_lookup_uid
 - map_retry
 - map_all
 - map_root
 - map_full
 - max_file_size
 - read_only
 - readdirplus
 - readdirplus_prefetch
 - return_32bit_file_ids
 - read_transfer_max_size
 - read_transfer_multiple
 - read_transfer_size
 - security_flavors
 - setattr_asynchronous
 - symlinks
 - time_delta
 - write_datasync_action
 - write_datasync_reply
 - write_filesync_action
 - write_filesync_reply
 - write_unstable_action
 - write_unstable_reply
 - write_transfer_max_size
 - write_transfer_multiple
 - write_transfer_size
- descending
Specifies to sort the data in descending order.
- format {table | json | csv | list}
Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- no-header
Displays table and CSV output without headers.
- no-footer
Displays table output without footers.
- verbose
Displays more detailed information.

Examples

The following command lists NFS exports, by default in the current zone:

```
isi nfs exports list
```

The following command lists NFS exports in a specific zone:

```
isi nfs exports list --zone hq-home
```

isi nfs exports modify

Modifies an NFS export.

Note

You can use the `isi nfs settings export view` command to see the full list of default settings for exports.

Syntax

```
isi nfs exports modify <ID>
  [--block-size <size>]
  [--revert-block-size]
  [--can-set-time {yes | no}]
  [--revert-can-set-time]
  [--case-insensitive {yes | no}]
  [--revert-case-insensitive]
  [--case-preserving {yes | no}]
  [--revert-case-preserving]
  [--chown-restricted {yes | no}]
  [--revert-chown-restricted]
  [--directory-transfer-size <size>]
  [--revert-directory-transfer-size]
  [--link-max <integer>]
  [--revert-link-max]
  [--max-file-size <size>]
  [--revert-max-file-size]
  [--name-max-size <integer>]
  [--revert-name-max-size]
  [--no-truncate {yes | no}]
  [--revert-no-truncate]
  [--return-32bit-file-ids {yes | no}]
  [--revert-return-32bit-file-ids]
  [--symlinks {yes | no}]
  [--revert-symlinks]
  [--new-zone <string>]
  [--description <string>]
  [--paths <path>]
  [--clear-paths]
  [--add-paths <string>]
  [--remove-paths <string>]
  [--clients <string>]
  [--clear-clients]
  [--add-clients <string>]
  [--remove-clients <string>]
  [--root-clients <string>]
  [--clear-root-clients]
  [--add-root-clients <string>]
  [--remove-root-clients <string>]
  [--read-write-clients <string>]
  [--clear-read-write-clients]
  [--add-read-write-clients <string>]
  [--remove-read-write-clients <string>]
  [--read-only-clients <string>]
  [--clear-read-only-clients]
  [--add-read-only-clients <string>]
  [--remove-read-only-clients <string>]
  [--all-dirs {yes | no}]
  [--revert-all-dirs]
  [--encoding <string>]
  [--revert-encoding]
```



```

[--security-flavors {unix | krb5 | krb5i | krb5p}]
[--revert-security-flavors]
[--clear-security-flavors]
[--add-security-flavors {unix | krb5 | krb5i | krb5p}]
[--remove-security-flavors <string>]
[--snapshot <snapshot>]
[--revert-snapshot]
[--map-lookup-uid {yes | no}]
[--revert-map-lookup-uid]
[--map-retry {yes | no}]
[--revert-map-retry]
[--map-root-enabled {yes | no}]
[--revert-map-root-enabled]
[--map-non-root-enabled {yes | no}]
[--revert-map-non-root-enabled]
[--map-failure-enabled {yes | no}]
[--revert-map-failure-enabled]
[--map-all <identity>]
[--revert-map-all]
[--map-root <identity>]
[--revert-map-root]
[--map-non-root <identity>]
[--revert-map-non-root]
[--map-failure <identity>]
[--revert-map-failure]
[--map-full {yes | no}]
[--revert-map-full]
[--commit-asynchronous {yes | no}]
[--revert-commit-asynchronous]
[--read-only {yes | no}]
[--revert-read-only]
[--readdirplus {yes | no}]
[--revert-readdirplus]
[--read-transfer-max-size <size>]
[--revert-read-transfer-max-size]
[--read-transfer-multiple <integer>]
[--revert-read-transfer-multiple]
[--read-transfer-size <size>]
[--revert-read-transfer-size]
[--setattr-asynchronous {yes | no}]
[--revert-setattr-asynchronous]
[--time-delta <time delta>]
[--revert-time-delta]
[--write-datasync-action {datasync | filesync |unstable}]
[--revert-write-datasync-action]
[--write-datasync-reply {datasync | filesync}]
[--revert-write-datasync-reply]
[--write-filesync-action {datasync | filesync |unstable}]
[--revert-write-filesync-action]
[--write-filesync-reply filesync]
[--write-unstable-action {datasync | filesync |unstable}]
[--revert-write-unstable-action]
[--write-unstable-reply {datasync | filesync |unstable}]
[--revert-write-unstable-reply]
[--write-transfer-max-size <size>]
[--revert-write-transfer-max-size]
[--write-transfer-multiple <integer>]
[--revert-write-transfer-multiple]
[--write-transfer-size <size>]
[--revert-write-transfer-size]
[--zone <string>]
[--force]
[--verbose]

```

Options

<ID>

The export ID number. You can use the `isi nfs exports list` command to view all the exports and their ID numbers in the current access zone.

`--block-size <size>`
Specifies the block size, in bytes.

`--revert-block-size`
Restores the setting to the system default.

`--can-set-time {yes | no}`
If set to `yes`, enables the export to set time. The default setting is `no`.

`--revert-can-set-time`
Restores the setting to the system default.

`--case-insensitive {yes | no}`
If set to `yes`, the server will report that it ignores case for file names. The default setting is `no`.

`--revert-case-insensitive`
Restores the setting to the system default.

`--case-preserving {yes | no}`
If set to `yes`, the server will report that it always preserves case for file names. The default setting is `no`.

`--revert-case-preserving`
Restores the setting to the system default.

`--chown-restricted {yes | no}`
If set to `yes`, the server will report that only the superuser can change file ownership. The default setting is `no`.

`--revert-chown-restricted`
Restores the setting to the system default.

`--directory-transfer-size <size>`
Specifies the preferred directory transfer size. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `128K`.

`--revert-directory-transfer-size`
Restores the setting to the system default.

`--link-max <integer>`
The reported maximum number of links to a file.

`--revert-link-max`
Restores the setting to the system default.

`--max-file-size <size>`
Specifies the maximum allowed file size on the server (in bytes). If a file is larger than the specified value, an error is returned.

`--revert-max-file-size`
Restores the setting to the system default.

`--name-max-size <integer>`

The reported maximum length of characters in a filename.

`--revert-name-max-size`

Restores the setting to the system default.

`--no-truncate {yes | no}`

If set to `yes`, too-long file names will result in an error rather than be truncated.

`--revert-no-truncate`

Restores the setting to the system default.

`--return-32bit-file-ids {yes | no}`

Applies to NFSv3 and later. If set to `yes`, limits the size of file identifiers returned from `readdir` to 32-bit values. The default value is `no`.

Note

This setting is provided for backward compatibility with older NFS clients, and should not be enabled unless necessary.

`--revert-return-32bit-file-ids`

Restores the setting to the system default.

`--symlinks {yes | no}`

If set to `yes`, advertises support for symlinks. The default setting is `no`.

`--revert-symlinks`

Restores the setting to the system default.

`--new-zone <string>`

Specifies a new access zone in which the export should apply. The default zone is `system`.

`--description <string>`

The description for this NFS export.

`--paths <paths>...`

Required. Specifies the path to be exported, starting at `/ifs`. This option can be repeated to specify multiple paths.

`--clear-paths`

Clear any of the paths originally specified for the export. The path must be within the `/ifs` directory.

`--add-paths <paths>...`

Add to the paths originally specified for the export. The path must be within `/ifs`. This option can be repeated to specify multiple paths.

`--remove-paths <paths>...`

Remove a path from the paths originally specified for the export. The path must be within `/ifs`. This option can be repeated to specify multiple paths to be removed.

`--clients <string>`

Specifies a client to be allowed access through this export. Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple clients by repeating this option.

`--clear-clients`

Clear the full list of clients originally allowed access through this export.

- `--add-clients <string>`
 Specifies a client to be added to the list of clients with access through this export. Specify clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple clients by repeating this option.
- `--remove-clients <string>`
 Specifies a client to be removed from the list of clients with access through this export. Specify clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple clients by repeating this option.
- `--root-clients <string>`
 Allows the root user of the specified client to execute operations as the root user of the cluster. This option overrides the `--map-all` and `--map-root` option for the specified client.
 Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.
- `--clear-root-clients`
 Clear the full list of root clients originally allowed access through this export.
- `--add-root-clients <string>`
 Specifies a root client to be added to the list of root clients with access through this export. Specify root clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple root clients by repeating this option.
- `--remove-root-clients <string>`
 Specifies a root client to be removed from the list of root clients with access through this export. Specify root clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple root clients by repeating this option.
- `--read-write-clients <string>`
 Grants read/write privileges to the specified client for this export. This option overrides the `--read-only` option for the specified client.
 Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.
- `--clear-read-write-clients`
 Clear the full list of read-write clients originally allowed access through this export.
- `--add-read-write-clients <string>`
 Specifies a read-write client to be added to the list of read-write clients with access through this export. Specify read-write clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple read-write clients by repeating this option.
- `--remove-read-write-clients <string>`
 Specifies a read-write client to be removed from the list of read-write clients with access through this export. Specify read-write clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple read-write clients by repeating this option.
- `--read-only-clients <string>`
 Makes the specified client read-only for this export. This option overrides the `--read-only` option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

`--clear-read-only-clients`

Clear the full list of read-only clients originally allowed access through this export.

`--add-read-only-clients <string>`

Specifies a read-only client to be added to the list of read-only clients with access through this export. Specify read-only clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple read-only clients by repeating this option.

`--remove-read-only-clients <string>`

Specifies a read-only client to be removed from the list of read-only clients with access through this export. Specify read-only clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple read-only clients by repeating this option.

`--all-dirs {yes | no}`

If set to `yes`, this export will cover all directories. The default setting is `no`.

`--revert-all-dirs`

Restores the setting to the system default.

`--encoding <string>`

Specifies the character encoding of clients connecting through this NFS export.

Valid values and their corresponding character encodings are provided in the following table. These values are taken from the node's `/etc/encodings.xml` file, and are not case sensitive.

Value	Encoding
cp932	Windows-SJIS
cp949	Windows-949
cp1252	Windows-1252
euc-kr	EUC-KR
euc-jp	EUC-JP
euc-jp-ms	EUC-JP-MS
utf-8-mac	UTF-8-MAC
utf-8	UTF-8
iso-8859-1	ISO-8859-1 (Latin-1)
iso-8859-2	ISO-8859-2 (Latin-2)
iso-8859-3	ISO-8859-3 (Latin-3)
iso-8859-4	ISO-8859-4 (Latin-4)
iso-8859-5	ISO-8859-5 (Cyrillic)
iso-8859-6	ISO-8859-6 (Arabic)
iso-8859-7	ISO-8859-7 (Greek)
iso-8859-8	ISO-8859-8 (Hebrew)

Value	Encoding
iso-8859-9	ISO-8859-9 (Latin-5)
iso-8859-10	ISO-8859-10 (Latin-6)
iso-8859-13	ISO-8859-13 (Latin-7)
iso-8859-14	ISO-8859-14 (Latin-8)
iso-8859-15	ISO-8859-15 (Latin-9)
iso-8859-16	ISO-8859-16 (Latin-10)

`--revert-encoding`

Restores the setting to the system default.

`--security-flavors {unix | krb5 | krb5i | krb5p} ...`

Specifies a security flavor to support. To support multiple security flavors, repeat this option for each additional entry. The following values are valid:

sys

Sys or UNIX authentication.

krb5

Kerberos V5 authentication.

krb5i

Kerberos V5 authentication with integrity.

krb5p

Kerberos V5 authentication with privacy.

`--revert-security-flavors`

Restores the setting to the system default.

`--snapshot {<snapshot> | <snapshot-alias>}`

Specifies the ID of a snapshot or snapshot alias to export. If you specify this option, directories will be exported in the state captured in either the specified snapshot or the snapshot referenced by the specified snapshot alias. If the snapshot does not capture the exported path, the export will be inaccessible to users.

If you specify a snapshot alias, and the alias is later modified to reference a new snapshot, the new snapshot will be automatically applied to the export.

Because snapshots are read-only, clients will not be able to modify data through the export unless you specify the ID of a snapshot alias that references the live version of the file system.

Specify *<snapshot>* or *<snapshot-alias>* as the ID or name of a snapshot or snapshot alias.

`--revert-snapshot`

Restores the setting to the system default.

`--map-lookup-uid {yes | no}`

If set to *yes*, incoming UNIX user identifiers (UIDs) will be looked up locally. The default setting is *no*.

`--revert-map-lookup-uid`

Restores the setting to the system default.

`--map-retry {yes | no}`

If set to `yes`, the system will retry failed user-mapping lookups. The default setting is `no`.

```
--revert-map-retry
```

Restores the setting to the system default.

```
--map-root-enabled {yes | no}
```

Enable/disable mapping incoming root users to a specific account.

```
--revert-map-root-enabled
```

Restores the setting to the system default.

```
--map-non-root-enabled {yes | no}
```

Enable/disable mapping incoming non-root users to a specific account.

```
--revert-map-non-root-enabled
```

Restores the setting to the system default.

```
--map-failure-enabled {yes | no}
```

Enable/disable mapping users to a specific account after failing an auth lookup.

```
--revert-map-failure-enabled
```

Restores the setting to the system default.

```
--map-all <identity>
```

Specifies the default identity that operations by any user will execute as. If this option is not set to `root`, you can allow the root user of a specific client to execute operations as the root user of the cluster by including the client in the `--root-clients` list.

```
--revert-map-all
```

Restores the setting to the system default.

```
--map-root <identity>
```

Map incoming root users to a specific user and/or group ID.

```
--revert-map-root
```

Restores the setting to the system default.

```
--map-non-root <identity>
```

Map non-root users to a specific user and/or group ID.

```
--revert-map-non-root
```

Restores the setting to the system default.

```
--map-failure <identity>
```

Map users to a specific user and/or group ID after a failed auth attempt.

```
--revert-map-failure
```

Restores the setting to the system default.

```
--map-full {yes | no}
```

Determines how user mapping is accomplished if a user is specified in an export option such as `--map-root` or `--map-all`. When enabled, a user mapping queries the OneFS user database and retrieves users from the applicable authentication subsystem, such as local authentication or Active Directory. When disabled, only local authentication is queried.

The default setting is `yes`.

```
--revert-map-full
```

Restores the `--map-full` setting to the system default, `yes`.

`--commit-asynchronous {yes | no}`
 If set to `yes`, enables commit data operations to be performed asynchronously. The default setting is `no`.

`--revert-commit-asynchronous`
 Restores the setting to the system default.

`--read-only {yes | no}`
 Determines the default privileges for all clients accessing the export.
 If set to `yes`, you can grant read/write privileges to a specific client by including the client in the `--read-write-clients` list.
 If set to `no`, you can make a specific client read-only by including the client in the `--read-only-clients` list. The default setting is `no`.

`--revert-read-only`
 Restores the setting to the system default.

`--readdirplus {yes | no}`
 Applies to NFSv3 only. If set to `yes`, enables processing of readdir-plus requests. The default setting is `no`.

`--revert-readdirplus`
 Restores the setting to the system default.

`--read-transfer-max-size <size>`
 Specifies the maximum read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

`--revert-read-transfer-max-size`
 Restores the setting to the system default.

`--read-transfer-multiple <integer>`
 Specifies the suggested multiple read size to report to NFSv3 and NFSv4 clients. Valid values are `0-4294967295`. The initial default value is `512`.

`--revert-read-transfer-multiple`
 Restores the setting to the system default.

`--read-transfer-size <size>`
 Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `128K`.

`--revert-read-transfer-size`
 Restores the setting to the system default.

`--setattr-asynchronous {yes | no}`
 If set to `yes`, performs set-attributes operations asynchronously. The default setting is `no`.

`--revert-setattr-asynchronous`
 Restores the setting to the system default.


```
--time-delta <float>
    Specifies server time granularity, in seconds.
--revert-time-delta
    Restores the setting to the system default.
--write-datasync-action {datasync | filesync | unstable}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync write method. The
    following values are valid:
    • datasync
    • filesync
    • unstable
    The default value is datasync, which performs the request as specified.
--revert-write-datasync-action
    Restores the setting to the system default.
--write-datasync-reply {datasync | filesync}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync reply method. The
    following values are valid:
    • datasync
    • filesync
    The default value is datasync (does not respond differently).
--revert-write-datasync-reply
    Restores the setting to the system default.
--write-filesync-action {datasync | filesync | unstable}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync write method. The
    following values are valid:
    • datasync
    • filesync
    • unstable
    The default value is filesync, which performs the request as specified.
--revert-write-filesync-action
    Restores the setting to the system default.
--write-filesync-reply {filesync}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync reply method. The
    only valid value is filesync (does not respond differently).
--write-unstable-action {datasync | filesync | unstable}
    Specifies an alternate unstable-write method. The following values are valid:
    • datasync
    • filesync
    • unstable
    The default value is unstable, which performs the request as specified.
--revert-write-unstable-action
    Restores the setting to the system default.
--write-unstable-reply {datasync | filesync | unstable}
```

Specifies an alternate unstable-reply method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable` (does not respond differently).

`--revert-write-unstable-reply`

Restores the setting to the system default.

`--write-transfer-max-size <size>`

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

`--revert-write-transfer-max-size`

Restores the setting to the system default.

`--write-transfer-multiple <integer>`

Specifies the suggested write transfer multiplier to report to NFSv3 and NFSv4 clients. Valid values are `0-4294967295`. The initial default value is `512`.

`--revert-write-transfer-multiple`

Restores the setting to the system default.

`--write-transfer-size <size>`

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

`--revert-write-transfer-size`

Restores the setting to the system default.

`--zone`

Access zone in which the export was originally created.

`{--force | -f}`

If set to `no` (default), a confirmation prompt displays when the command runs. If set to `yes`, the command executes without prompting for confirmation.

`{--verbose | -v}`

Displays more detailed information.

isi nfs exports reload

Reloads the NFS exports configuration.

Syntax

```
isi nfs exports reload
```

Options

There are no options for this command.

isi nfs exports view

View an NFS export.

Syntax

```
isi nfs exports view <id>
  [--zone <string>]
```

Options

<id>

Specifies the ID of the NFS export to display. If you do not know the ID, use the `isi nfs exports list` command to view a list of exports and their associated IDs.

`--zone <string>`

Specifies the name of the access zone in which the export was created.

isi nfs log-level

Configure NFS logging in OneFS.

To use this command, you must have `ISI_PRIV_NFS` or higher administrative privileges in OneFS.

Syntax

```
isi nfs log-level
  --set {always | error | warning | info | verbose | debug | trace}
  --reset
```

Options

None

With no optional parameter specified, displays the currently set log level.

`--set {always | error | warning | info | verbose | debug | trace}`

Specifies the desired NFS log level for this node, as described in the following table.

Log level	Description
always	Specifies that all NFS events are logged in NFS log files.
error	Specifies that only NFS error conditions are logged in NFS log files.
warning	Specifies that only NFS warning conditions are logged in NFS log files.
info	Specifies that only NFS information conditions are logged in NFS log files.
verbose	Specifies verbose logging.
debug	Adds information that EMC Isilon can use to troubleshoot issues
trace	Adds tracing information that EMC Isilon can use to pinpoint issues

`--reset`

Restores the log level to the system default.

Examples

The following command sets the NFS log level to `error` only:

```
isi nfs log-level --set error
```

The following command resets the NFS log level to the system default:

```
isi nfs log-level --reset
```

isi nfs netgroup bgwrite

Gets and sets the time between writes to the backup drive.

Syntax

```
isi nfs netgroup bgwrite
  [--minutes <number>]
```

Options

```
{--minutes | -m} <number>
```

Specifies the time in minutes between writes to the backup drive.

isi nfs netgroup check

Updates all cached netgroups.

Syntax

```
isi nfs netgroup check
  [--node <string>]
```

Options

```
{--node | -n} <string>
```

Specifies the node to send the update command.

isi nfs netgroup expiration

Sets the time between each netgroup expiration.

Syntax

```
isi nfs netgroup expiration
  [--minutes <number>]
```

Options

```
{--minutes | -m} <number>
```

Specifies the time in minutes between each netgroup expiration.

isi nfs netgroup flush

Specifies the NFS netgroup to flush.

Syntax

```
isi nfs netgroup flush
  [--node <string>]
```

Options

```
{--node | -n} <string>
```

Specifies the node to send the flush command to.

isi nfs netgroup lifetime

Gets and sets the time in minutes before stale cache entries are wiped.

Syntax

```
isi nfs netgroup lifetime
  [--minutes <number>]
```

Options

```
{--minutes | -m} <number>
```

Specifies the time in minutes before stale cache entries are wiped.

isi nfs netgroup retry

Gets and sets the retry interval.

Syntax

```
isi nfs netgroup retry
  [--seconds <integer>]
```

Options

```
{--seconds | -s} <integer>
```

Specifies the time in seconds between attempts to retry an NFS netgroup.

isi nfs nlm locks list

Applies to NFSv3 only. Displays a list of NFS Network Lock Manager (NLM) advisory locks.

Syntax

```
isi nfs nlm locks list
  [--limit <integer>]
  [--sort {client | path | lock_type | range | created}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
--limit <integer>
```

Displays no more than the specified number of NFS nlm locks.

```
--sort {client | path | lock_type | range | created}
```

Specifies the field to sort by.

```
--descending
```

Specifies to sort the data in descending order.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

- `--no-header`
Displays table and CSV output without headers.
- `--no-footer`
Displays table output without footers.
- `--verbose`
Displays more detailed information.

Examples

To view a detailed list of all current NLM locks, run the following command:

```
isi nfs nlm locks list --verbose
```

In the following sample output, there are currently three locks: one on `/ifs/home/test1/file.txt` and two on `/ifs/home/test2/file.txt`.

Client	Path	Lock Type	Range
machineName/10.72.134.119	/ifs/home/test1/file.txt	exclusive	[0, 2]
machineName/10.59.166.125	/ifs/home/test2/file.txt	shared	[10, 20]
machineName/10.63.119.205	/ifs/home/test2/file.txt	shared	[10, 20]

isi nfs nlm locks waiters

Displays a list of clients that are waiting to place a Network Lock Manager (NLM) lock on a currently locked file. This command applies to NFSv3 only.

Syntax

```
isi nfs nlm locks waiters
  [--limit <integer>]
  [--sort {client | path | lock_type | range | created}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

- `--limit <integer>`
Displays no more than the specified number of NLM locks.
- `--sort {client | path | lock_type | range | created}`
Specifies the field to sort by.
- `--descending`
Specifies to sort the data in descending order.
- `--format {table | json | csv | list}`
Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- `--no-header`
Displays table and CSV output without headers.
- `--no-footer`
Displays table output without footers.
- `--verbose`
Displays more detailed information.

Examples

The following command displays a detailed list of clients waiting to lock a currently-locked file:

```
isi nfs nlm locks waiters --verbose
```

The system displays output similar to the following example:

Client	Path	Lock Type	Range
machineName/1.2.34.5	/ifs/home/test1/file.txt	exclusive	[0, 2]

isi nfs nlm sessions disconnect

Removes an NFS client, cancels any pending Network Lock Manager (NLM) locks, and unlocks all of the client's current locks. This command applies to NFSv3 only.

Syntax

```
isi nfs nlm sessions disconnect <client>
[--force]
```

Options

<client>

Required. Specifies the client to disconnect, in the form *<machine_name>/<server_ip>*.

`--force`

Suppresses command-line prompts and messages.

Example

The following command disconnects a client named *client1* from a server whose IP address is 192.168.7.143:

```
isi nfs nlm sessions disconnect client1/192.168.7.143
```

isi nfs nlm sessions list

Displays a list of clients holding NFS Network Lock Manager (NLM) locks. This command applies to NFSv3 only.

Syntax

```
isi nfs nlm sessions list
[--limit <integer>]
[--sort {client | path | lock_type | range | created}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

`--limit` *<integer>*

Displays no more than the specified number of NFS nlm sessions.

`--sort` {client | path | lock_type | range | created}

Specifies the field to sort by.

`--descending`

Specifies to sort the data in descending order.

- `--format {table | json | csv | list}`
Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- `--no-header`
Displays table and CSV output without headers.
- `--no-footer`
Displays table output without footers.
- `--verbose`
Displays more detailed information.

Example

To view a list of active NLM sessions, run the following command:

```
isi nfs nlm sessions list
```

isi nfs settings export modify

Modifies the default settings that are applied when creating NFS exports.

Note

You can view the currently configured default NFS export settings by running the `isi nfs settings export view` command.

Syntax

```
isi nfs exports modify <ID>
  [--block-size <size>]
  [--revert-block-size]
  [--can-set-time {yes|no}]
  [--revert-can-set-time]
  [--case-insensitive {yes|no}]
  [--revert-case-insensitive]
  [--case-preserving {yes|no}]
  [--revert-case-preserving]
  [--chown-restricted {yes|no}]
  [--revert-chown-restricted]
  [--directory-transfer-size <size>]
  [--revert-directory-transfer-size]
  [--link-max <integer>]
  [--revert-link-max]
  [--max-file-size <size>]
  [--revert-max-file-size]
  [--name-max-size <integer>]
  [--revert-name-max-size]
  [--no-truncate {yes|no}]
  [--revert-no-truncate]
  [--return-32bit-file-ids {yes|no}]
  [--revert-return-32bit-file-ids]
  [--symlinks {yes|no}]
  [--revert-symlinks]
  [--all-dirs {yes|no}]
  [--revert-all-dirs]
  [--encoding <string>]
  [--revert-encoding]
  [--security-flavors {unix|krb5|krb5i|krb5p}]
  [--revert-security-flavors]
  [--clear-security-flavors]
  [--add-security-flavors {unix|krb5|krb5i|krb5p}]
  [--remove-security-flavors <string>]
  [--snapshot <snapshot>]
```



```

[--revert-snapshot]
[--map-lookup-uid {yes|no}]
[--revert-map-lookup-uid]
[--map-retry {yes|no}]
[--revert-map-retry]
[--map-root-enabled {yes|no}]
[--revert-map-root-enabled]
[--map-non-root-enabled {yes|no}]
[--revert-map-non-root-enabled]
[--map-failure-enabled {yes|no}]
[--revert-map-failure-enabled]
[--map-all <identity>]
[--revert-map-all]
[--map-root <identity>]
[--revert-map-root]
[--map-non-root <identity>]
[--revert-map-non-root]
[--map-failure <identity>]
[--revert-map-failure]
[--map-full {yes|no}]
[--revert-map-full]
[--commit-asynchronous {yes|no}]
[--revert-commit-asynchronous]
[--read-only {yes|no}]
[--revert-read-only]
[--readdirplus {yes|no}]
[--revert-readdirplus]
[--read-transfer-max-size <size>]
[--revert-read-transfer-max-size]
[--read-transfer-multiple <integer>]
[--revert-read-transfer-multiple]
[--read-transfer-size <size>]
[--revert-read-transfer-size]
[--setattr-asynchronous {yes|no}]
[--revert-setattr-asynchronous]
[--time-delta <integer>]
[--revert-time-delta]
[--write-datasync-action {datasync|filesync|unstable}]
[--revert-write-datasync-action]
[--write-datasync-reply {datasync|filesync}]
[--revert-write-datasync-reply]
[--write-filesync-action {datasync|filesync|unstable}]
[--revert-write-filesync-action]
[--write-filesync-reply filesync]
[--write-unstable-action {datasync|filesync|unstable}]
[--revert-write-unstable-action]
[--write-unstable-reply {datasync|filesync|unstable}]
[--revert-write-unstable-reply]
[--write-transfer-max-size <size>]
[--revert-write-transfer-max-size]
[--write-transfer-multiple <integer>]
[--revert-write-transfer-multiple]
[--write-transfer-size <size>]
[--revert-write-transfer-size]
[--zone <string>]
[--force]
[--verbose]

```

Options

`--block-size <size>`

Specifies the block size, in bytes.

`--revert-block-size`

Restores the setting to the system default.

`--can-set-time {yes|no}`

If set to `yes`, enables the export to set time. The default setting is `no`.

`--revert-can-set-time`
Restores the setting to the system default.

`--case-insensitive {yes|no}`
If set to `yes`, the server will report that it ignores case for file names. The default setting is `no`.

`--revert-case-insensitive`
Restores the setting to the system default.

`--case-preserving {yes|no}`
If set to `yes`, the server will report that it always preserves case for file names. The default setting is `no`.

`--revert-case-preserving`
Restores the setting to the system default.

`--chown-restricted {yes|no}`
If set to `yes`, the server will report that only the superuser can change file ownership. The default setting is `no`.

`--revert-chown-restricted`
Restores the setting to the system default.

`--directory-transfer-size <size>`
Specifies the preferred directory transfer size. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 128K.

`--revert-directory-transfer-size`
Restores the setting to the system default.

`--link-max <integer>`
The reported maximum number of links to a file.

`--revert-link-max`
Restores the setting to the system default.

`--max-file-size <size>`
Specifies the maximum allowed file size on the server (in bytes). If a file is larger than the specified value, an error is returned.

`--revert-max-file-size`
Restores the setting to the system default.

`--name-max-size <integer>`
The reported maximum length of characters in a filename.

`--revert-name-max-size`
Restores the setting to the system default.

`--no-truncate {yes|no}`
If set to `yes`, too-long file names will result in an error rather than be truncated.

`--revert-no-truncate`
Restores the setting to the system default.

`--return-32bit-file-ids {yes|no}`

Applies to NFSv3 and later. If set to `yes`, limits the size of file identifiers returned from `readdir` to 32-bit values. The default value is `no`.

Note

This setting is provided for backward compatibility with older NFS clients, and should not be enabled unless necessary.

`--revert-return-32bit-file-ids`

Restores the setting to the system default.

`--symlinks {yes|no}`

If set to `yes`, advertises support for symlinks. The default setting is `no`.

`--revert-symlinks`

Restores the setting to the system default.

`--new-zone <string>`

Specifies a new access zone in which the export should apply. The default zone is `system`.

`--all-dirs {yes|yesno}`

If set to `yes`, this export will cover all directories. The default setting is `no`.

`--revert-all-dirs`

Restores the setting to the system default.

`--encoding <string>`

Specifies the character encoding of clients connecting through this NFS export.

Valid values and their corresponding character encodings are provided in the following table. These values are taken from the node's `/etc/encodings.xml` file, and are not case sensitive.

Value	Encoding
cp932	Windows-SJIS
cp949	Windows-949
cp1252	Windows-1252
euc-kr	EUC-KR
euc-jp	EUC-JP
euc-jp-ms	EUC-JP-MS
utf-8-mac	UTF-8-MAC
utf-8	UTF-8
iso-8859-1	ISO-8859-1 (Latin-1)
iso-8859-2	ISO-8859-2 (Latin-2)
iso-8859-3	ISO-8859-3 (Latin-3)
iso-8859-4	ISO-8859-4 (Latin-4)
iso-8859-5	ISO-8859-5 (Cyrillic)
iso-8859-6	ISO-8859-6 (Arabic)

Value	Encoding
iso-8859-7	ISO-8859-7 (Greek)
iso-8859-8	ISO-8859-8 (Hebrew)
iso-8859-9	ISO-8859-9 (Latin-5)
iso-8859-10	ISO-8859-10 (Latin-6)
iso-8859-13	ISO-8859-13 (Latin-7)
iso-8859-14	ISO-8859-14 (Latin-8)
iso-8859-15	ISO-8859-15 (Latin-9)
iso-8859-16	ISO-8859-16 (Latin-10)

`--revert-encoding`

Restores the setting to the system default.

`--security-flavors {unix|krb5|krb5i|krb5p} ...`

Specifies a security flavor to support. To support multiple security flavors, repeat this option for each additional entry. The following values are valid:

sys

Sys or UNIX authentication.

krb5

Kerberos V5 authentication.

krb5i

Kerberos V5 authentication with integrity.

krb5p

Kerberos V5 authentication with privacy.

`--revert-security-flavors`

Restores the setting to the system default.

`--snapshot {<snapshot>|<snapshot-alias>}`

Specifies the ID of a snapshot or snapshot alias to export. If you specify this option, directories will be exported in the state captured in either the specified snapshot or the snapshot referenced by the specified snapshot alias. If the snapshot does not capture the exported path, the export will be inaccessible to users.

If you specify a snapshot alias, and the alias is later modified to reference a new snapshot, the new snapshot will be automatically applied to the export.

Because snapshots are read-only, clients will not be able to modify data through the export unless you specify the ID of a snapshot alias that references the live version of the file system.

Specify *<snapshot>* or *<snapshot-alias>* as the ID or name of a snapshot or snapshot alias.

`--revert-snapshot`

Restores the setting to the system default.

`--map-lookup-uid {yes|no}`

If set to *yes*, incoming UNIX user identifiers (UIDs) will be looked up locally. The default setting is *no*.

`--revert-map-lookup-uid`

Restores the setting to the system default.

`--map-retry {yes|no}`
 If set to `yes`, the system will retry failed user-mapping lookups. The default setting is `no`.

`--revert-map-retry`
 Restores the setting to the system default.

`--map-root-enabled {yes|no}`
 Enable/disable mapping incoming root users to a specific account.

`--revert-map-root-enabled`
 Restores the setting to the system default.

`--map-non-root-enabled {yes|no}`
 Enable/disable mapping incoming non-root users to a specific account.

`--revert-map-non-root-enabled`
 Restores the setting to the system default.

`--map-failure-enabled {yes|no}`
 Enable/disable mapping users to a specific account after failing an auth lookup.

`--revert-map-failure-enabled`
 Restores the setting to the system default.

`--map-all <identity>`
 Specifies the default identity that operations by any user will run as. If this option is not set to `root`, you can allow the root user of a specific client to run operations as the root user of the cluster by including the client in the `--root-clients` list.

`--revert-map-all`
 Restores the setting to the system default.

`--map-root <identity>`
 Map incoming root users to a specific user and/or group ID.

`--revert-map-root`
 Restores the setting to the system default.

`--map-non-root <identity>`
 Map non-root users to a specific user and/or group ID.

`--revert-map-non-root`
 Restores the setting to the system default.

`--map-failure <identity>`
 Map users to a specific user and/or group ID after a failed auth attempt.

`--revert-map-failure`
 Restores the setting to the system default.

`--map-full {yes|no}`
 Determines how user mapping is accomplished if a user is specified in an export option such as `--map-root` or `--map-all`. When enabled, a user mapping queries the OneFS user database and retrieves users from the applicable authentication subsystem, such as local authentication or Active Directory. When disabled, only local authentication is queried.

The default setting is *yes*.

`--revert-map-full`
Restores the `--map-full` setting to the system default, *yes*.

`--commit-asynchronous {yes|no}`
If set to *yes*, enables commit data operations to be performed asynchronously. The default setting is *no*.

`--revert-commit-asynchronous`
Restores the setting to the system default.

`--read-only {yes|no}`
Determines the default privileges for all clients accessing the export.
If set to *yes*, you can grant read/write privileges to a specific client by including the client in the `--read-write-clients` list.
If set to *no*, you can make a specific client read-only by including the client in the `--read-only-clients` list. The default setting is *no*.

`--revert-read-only`
Restores the setting to the system default.

`--readdirplus {yes|no}`
Applies to NFSv3 only. If set to *yes*, enables processing of readdir-plus requests. The default setting is *no*.

`--revert-readdirplus`
Restores the setting to the system default.

`--read-transfer-max-size <size>`
Specifies the maximum read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: *b* for bytes; *K* for KB; *M* for MB; or *G* for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295*b*. The initial default value is 512*K*.

`--revert-read-transfer-max-size`
Restores the setting to the system default.

`--read-transfer-multiple <integer>`
Specifies the suggested multiple read size to report to NFSv3 and NFSv4 clients. Valid values are 0-4294967295. The initial default value is 512.

`--revert-read-transfer-multiple`
Restores the setting to the system default.

`--read-transfer-size <size>`
Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: *b* for bytes; *K* for KB; *M* for MB; or *G* for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295*b*. The initial default value is 128*K*.

`--revert-read-transfer-size`
Restores the setting to the system default.

`--setattr-asynchronous {yes|no}`
If set to *yes*, performs set-attributes operations asynchronously. The default setting is *no*.

```

--revert-setattr-asynchronous
    Restores the setting to the system default.
--time-delta <integer>
    Specifies server time granularity, in seconds.
--revert-time-delta
    Restores the setting to the system default.
--write-datasync-action {datasync|filesync|unstable}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync write method. The
    following values are valid:
    • datasync
    • filesync
    • unstable
    The default value is datasync, which performs the request as specified.
--revert-write-datasync-action
    Restores the setting to the system default.
--write-datasync-reply {datasync|filesync}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync reply method. The
    following values are valid:
    • datasync
    • filesync
    The default value is datasync (does not respond differently).
--revert-write-datasync-reply
    Restores the setting to the system default.
--write-filesync-action {datasync|filesync|unstable}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync write method. The
    following values are valid:
    • datasync
    • filesync
    • unstable
    The default value is filesync, which performs the request as specified.
--revert-write-filesync-action
    Restores the setting to the system default.
--write-filesync-reply {filesync}
    Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync reply method. The
    only valid value is filesync (does not respond differently).
--write-unstable-action {datasync|filesync|unstable}
    Specifies an alternate unstable-write method. The following values are valid:
    • datasync
    • filesync
    • unstable
    The default value is unstable, which performs the request as specified.
--revert-write-unstable-action

```

Restores the setting to the system default.

```
--write-unstable-reply {datasync|filesync|unstable}
```

Specifies an alternate unstable-reply method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable` (does not respond differently).

```
--revert-write-unstable-reply
```

Restores the setting to the system default.

```
--write-transfer-max-size <size>
```

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

```
--revert-write-transfer-max-size
```

Restores the setting to the system default.

```
--write-transfer-multiple <integer>
```

Specifies the suggested write transfer multiplier to report to NFSv3 and NFSv4 clients. Valid values are `0-4294967295`. The initial default value is `512`.

```
--revert-write-transfer-multiple
```

Restores the setting to the system default.

```
--write-transfer-size <size>
```

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `512K`.

```
--revert-write-transfer-size
```

Restores the setting to the system default.

```
--zone
```

Access zone in which the export was originally created.

```
--force
```

If set to `no` (default), a confirmation prompt displays when the command runs. If set to `yes`, the command runs without prompting for confirmation.

```
--verbose
```

Displays more detailed information.

isi nfs settings export view

Displays default NFS export settings.

Syntax

```
isi nfs settings export view
[--zone <string>]
```

Options

```
--zone <string>
```


Specifies the access zone in which the default settings apply.

Example

To view the currently-configured default export settings, run the following command:

```
isi nfs settings export view
```

The system displays output similar to the following example:

```
Read Write Clients: -
Unresolved Clients: -
  All Dirs: No
  Block Size: 8.0K
  Can Set Time: Yes
  Case Insensitive: No
  Case Preserving: Yes
  Chown Restricted: No
  Commit Asynchronous: No
Directory Transfer Size: 128.0K
  Encoding: DEFAULT
  Link Max: 32767
  Map Lookup UID: No
  Map Retry: Yes
  Map Root
    Enabled: True
    User: nobody
  Primary Group: -
  Secondary Groups: -
  Map Non Root
    Enabled: False
    User: nobody
  Primary Group: -
  Secondary Groups: -
  Map Failure
    Enabled: False
    User: nobody
  Primary Group: -
  Secondary Groups: -
  Map Full: Yes
  Max File Size: 8192.00000P
  Name Max Size: 255
  No Truncate: No
  Read Only: No
  Readdirplus: Yes
Return 32Bit File Ids: No
Read Transfer Max Size: 1.00M
Read Transfer Multiple: 512
  Read Transfer Size: 128.0K
  Security Type: unix
Setattr Asynchronous: No
  Snapshot: -
  Symlinks: Yes
  Time Delta: 1.0 ns
Write Datasync Action: datasync
  Write Datasync Reply: datasync
Write Filesync Action: filesync
  Write Filesync Reply: filesync
Write Unstable Action: unstable
  Write Unstable Reply: unstable
Write Transfer Max Size: 1.00M
Write Transfer Multiple: 512
  Write Transfer Size: 512.0K
```

isi nfs settings global modify

Modifies the default NFS global settings.

Syntax

```
isi nfs settings global modify
  [--lock-protection <integer>]
  [--nfsv3-enabled {yes | no}]
  [--nfsv4-enabled {yes | no}]
  [--force]
```

Options

`--lock-protection <integer>`

Specifies the number of nodes failures that can happen before a lock might be lost.

`--nfsv3-enabled {yes | no}`

Specifies that NFSv3 is enabled.

`--nfsv4-enabled {yes | no}`

Specifies that NFSv4 is enabled.

`{--force}`

Causes the command to be executed without your confirmation.

isi nfs settings global view

Displays the global options for NFS settings.

Syntax

```
isi nfs settings global view
```

Options

There are no options for this command.

isi nfs settings zone modify

Modifies the default NFS zone settings for the NFSv4 ID mapper.

Syntax

```
isi nfs settings zone modify
  [--nfsv4-domain <string>]
  [--revert-nfsv4-domain]
  [--nfsv4-replace-domain {yes | no}]
  [--revert-nfsv4-replace-domain]
  [--nfsv4-no-domain {yes | no}]
  [--revert-nfsv4-no-domain]
  [--nfsv4-no-domain-uids {yes | no}]
  [--revert-nfsv4-no-domain-uids]
  [--nfsv4-no-names {yes | no}]
  [--revert-nfsv4-no-names]
  [--nfsv4-allow-numeric-ids {yes | no}]
  [--revert-nfsv4-allow-numeric-ids]
  [--zone <string>]
  [--verbose]
```

Options

`--nfsv4-domain <string>`

Specifies the NFSv4 domain name.

```

--revert-nfsv4-domain
    Return the setting to the system default. Default is localdomain.
--nfsv4-replace-domain {yes | no}
    Replace owner/group domain with the domainname. For NFSv4.
--revert-nfsv4-replace-domain
    Returns setting to the system default. Default is yes.
--nfsv4-no-domain {yes | no}
    Sends owners/groups without domainname. For NFSv4.
--revert-nfsv4-no-domain
    Returns setting to the system default. Default is no.
--nfsv4-no-domain-uids {yes | no}
    Send UIDs/GIDs without domainname. For NFSv4.
--revert-nfsv4-no-domain-uids
    Returns setting to the system default. Default is yes.
--nfsv4-no-names {yes | no}
    Always send owners/groups as UIDs/GIDs. For NFSv4.
--revert-nfsv4-no-names
    Returns setting to the system default. Default is no.
--nfsv4-allow-numeric-ids {yes | no}
    Sends owners/groups as UIDs/GIDs when look-ups fail or if no-names=1. For NFSv4.
--revert-nfsv4-allow-numeric-ids
    Returns setting to the system default. Default is yes.
--zone <string>
    Specifies the access zone.
--verbose
    Displays more detailed information.

```

Example

The following command specifies that the NFS server would accept UIDs/GIDs in place of user names:

```
isi nfs settings zone modify --nfsv4-no-names yes
```

isi nfs settings zone view

Displays the default NFSv4-related access zone settings.

Syntax

```
isi nfs settings zone view
[--zone <string>]
```

Options

```
--zone <string>
```

Specifies the access zone for which you want to view NFSv4-related settings.

Example

The following command specifies that you want to examine NFSv4-related settings for an access zone named Zone1:

```
isi nfs settings zone view --zone=Zone1
```

FTP

OneFS includes a secure FTP service called vsftpd, which stands for Very Secure FTP Daemon, that you can configure for standard FTP and FTPS file transfers.

View FTP settings

You can view a list of current FTP configuration settings.

Procedure

1. Run the `isi ftp list` command.

The system displays output similar to the following example:

```
accept-timeout          60
allow-anon-access       NO
allow-anon-upload       YES
allow-dirlists          YES
allow-downloads         YES
allow-local-access      YES
allow-writes            YES
always-chdir-homedir   YES
anon-chown-username     root
anon-root-path          /ifs/home/ftp
anon-umask              077
ascii-mode              off
connect-timeout         60
data-timeout            300
dirlist-localtime      NO
dirlist-names          hide
file-create-perm        0666
local-root-path         local user home directory
local-umask             077
server-to-server        NO
session-support         YES
session-timeout         300
user-config-dir         (none)
denied-user-list        (none)
limit-anon-passwords   NO
anon-password-list     (disabled)
chroot-local-mode      No local users chrooted; exception list
inactive
chroot-exception-list  (none)
```

Enable FTP file sharing

The FTP service, vsftpd, is disabled by default.

Note

You can determine whether the service is enabled or disabled by running the `isi services -l` command.

Procedure

1. Run the following command:

```
isi services vsftpd enable
```

The system displays the following confirmation message:

```
The service 'vsftpd' has been enabled.
```

After you finish

You can configure FTP settings by running the `isi ftp` command.

Configure FTP file sharing

You can set the FTP service to allow any node in the cluster to respond to FTP requests through a standard user account.

Before you begin

You must enable the FTP service before you can use it.

You can enable the transfer of files between remote FTP servers and enable anonymous FTP service on the root by creating a local user named `anonymous` or `ftp`.

When configuring FTP access, make sure that the specified FTP root is the home directory of the user who logs in. For example, the FTP root for local user `jsmith` should be `ifs/home/jsmith`.

Procedure

1. Run the `isi ftp <action>` command.

You must run this command separately for each action.

The following command enables server-to-server transfers:

```
isi ftp server-to-server=yes
```

The following command disables anonymous uploads:

```
isi ftp allow-anon-upload=no
```

FTP commands

You can access and configure the FTP service through the FTP commands.

isi ftp accept-timeout

Sets and displays data connection timeout

Syntax

```
isi ftp accept-timeout <value>
```

Options

<value>

Specifies the time, in seconds, that a remote client has to establish a PASV style data connection before timeout. All integers between 30 and 600 are valid values. If no options are specified, the current timeout is displayed. The default value is 60.

Examples

To set the data connection timeout to 5 minutes, run the following command:

```
isi ftp accept-timeout 300
```

isi ftp allow-anon-access

Sets and displays whether anonymous access is permitted.

Syntax

```
isi ftp allow-anon-access <value>
```

Options

<value>

Controls whether anonymous logins are permitted or not. If enabled, both the usernames ftp and anonymous are recognized as anonymous logins. Valid values are **YES** and **NO**. If no options are specified, displays whether or not anonymous logins are permitted. The default value is **NO**.

Examples

To allow anonymous access, run the following command:

```
isi ftp allow-anon-access YES
```

isi ftp allow-anon-upload

Sets and displays whether anonymous users are permitted to upload files.

Syntax

```
isi ftp allow-anon-upload <value>
```

Options

<value>

Controls whether anonymous users are able to upload files under certain conditions. Valid values are **YES** and **NO**. For anonymous users to be able to upload, the `isi ftp allow-writes` command must be set to **YES**, and the anonymous user must have write permission on the desired upload location. If no options are specified, displays whether anonymous users are permitted to upload files. The default value is **YES**.

Examples

To disable anonymous users from uploading files, run the following command:

```
isi ftp allow-anon-upload NO
```

isi ftp allow-dirlists

Sets and displays whether directory list commands are permitted.

Syntax

```
isi ftp allow-dirlists <value>
```

Options

<value>

Controls whether directory list commands are enabled. Valid values are **YES** and **NO**. If no options are specified, displays whether directory list commands are permitted. The default value is **YES**.

Examples

To disable directory list commands, run the following command:

```
isi ftp allow-dirlists NO
```

isi ftp allow-downloads

Sets and displays whether downloads are permitted.

Syntax

```
isi ftp allow-downloads <value>
```

Options

<value>

Controls whether files can be downloaded. Valid values are **YES** and **NO**. If no options are specified, displays whether downloads are permitted. The default value is **YES**.

Examples

To disable downloads from being permitted, run the following command:

```
isi ftp allow-downloads NO
```

isi ftp allow-local-access

Displays and controls whether local logins are permitted.

Syntax

```
isi ftp allow-local-access <value>
```

Options

<value>

Valid values are **YES** and **NO**. If set to **YES**, normal user accounts can be used to log in. If no options are specified, displays whether commands that change the file system are permitted. The default value is **YES**.

Examples

To deny local login permission, run the following command:

```
isi ftp allow-local-access NO
```

isi ftp allow-writes

Sets and displays whether commands that change the filesystem are permitted.

Syntax

```
isi ftp allow-writes <value>
```

Options

<value>

If no options are specified, displays whether commands that change the file system are permitted. Controls whether any of the following commands are allowed:

- **STOR**
- **DELE**
- **RNFR**
- **RNTO**
- **MKD**
- **RMD**
- **APPE**
- **SITE**

Valid values are **YES** and **NO**. The default value is **YES**.

Examples

To disable commands that change the file system, run the following command:

```
isi ftp allow-writes NO
```

isi ftp always-chdir-homedir

Specifies whether FTP always changes directories to the home directory of the user.

Syntax

```
isi ftp always-chdir-homedir <value>
```

Options

<value>

Controls whether FTP always initially changes directories to the home directory of the user. The default value is **YES**. If set to **NO**, you can set up a `chroot` area in FTP without having a home directory for the user. If no options are specified, displays the current setting. Valid values are **YES** and **NO**.

isi ftp anon-chown-username

Displays and specifies the owner of anonymously uploaded files.

Syntax

```
isi ftp anon-chown-username <value>
```

Options

<value>

Gives ownership of anonymously uploaded files to the specified user. The value must be a local username. If no options are specified, displays the owner of anonymously uploaded files. The default value is `root`.

Examples

The following command sets the owner of anonymously uploaded files to be "user1":

```
isi ftp anon-chown-username user1
```


isi ftp anon-password-list

Displays the list of anonymous user passwords.

Syntax

```
isi ftp anon-password-list
```

Options

There are no options for this command.

Examples

To display a list of anonymous user passwords, run the following command:

```
isi ftp anon-password-list
```

The system displays output similar to the following example:

```
anon-password-list: 1234
                    password
```

isi ftp anon-password-list add

Adds passwords to the anonymous password list.

Syntax

```
isi ftp anon-password-list add <value>
```

Options

<value>

Specifies the password being added to the anonymous password list.

Examples

The following command adds "1234" to the anonymous password list:

```
isi ftp anon-password-list add 1234
```

The system displays output similar to the following example:

```
anon-password-list: added password '1234' to list
```

isi ftp anon-password-list remove

Removes passwords from the anonymous password list.

Syntax

```
isi ftp anon-password-list remove <value>
```

Options

<value>

Specifies which password to remove from the anonymous password list.

Examples

The following command removes "1234" from the anonymous password list:

```
isi ftp anon-password-list remove 1234
```

The system displays output similar to the following example:

```
anon-password-list: removed password '1234' from list
```

isi ftp anon-root-path

Displays and specifies the root path for anonymous users.

Syntax

```
isi ftp anon-root-path --value <ifs-directory>
[--reset]
```

Options

```
{--value | -v} <ifs-directory>
```

Represents a directory in */ifs* that the Very Secure FTP Daemon (VSFTPD) will try to change to after an anonymous login. Valid values are paths in */ifs*. If no options are specified, displays the root path for anonymous users. The default value is */ifs/home/ftp*.

```
--reset
```

Resets the value to */ifs/home/ftp*.

Examples

The following command sets the root path for anonymous users to */ifs/home/newUser/*:

```
isi ftp anon-root-path --value /ifs/home/newUser/
```

The system displays output similar to the following example:

```
anon-root-path: /ifs/home/ftp -> /ifs/home/newUser/
```

isi ftp anon-umask

Displays and specifies the anonymous user file creation umask.

Syntax

```
isi ftp anon-umask <value>
```

Options

```
<value>
```

Specifies the umask for file creation by anonymous users. Valid values are octal umask numbers. If no options are specified, displays the current anonymous user file creation umask. The default value is **077**.

Note

The value must contain the '0' prefix; otherwise, the value will be treated as a base 10 integer

.

Examples

The following command sets the umask for file creation by anonymous users to 066:

```
isi ftp anon-umask 066
```

The system displays output similar to the following example:

```
anon-umask: 077 -> 066
```

isi ftp ascii-mode

Sets and displays whether ASCII downloads and uploads are permitted.

Syntax

```
isi ftp ascii-mode <value>
```

Options

<value>

Determines whether ASCII downloads and uploads are enabled. The following values are valid:

- **both** ASCII mode data transfers are honored on both downloads and uploads.
- **download** ASCII mode data transfers are honored on downloads.
- **off** ASCII mode data transfers will not be honored.
- **upload** ASCII mode data transfers are honored on uploads.

If no options are specified, displays whether ASCII downloads and uploads are permitted. The default value is **off**.

Examples

To allow both ASCII downloads and uploads, run the following command:

```
isi ftp ascii-mode both
```

The system displays output similar to the following example:

```
ascii-mode: off -> both
```

isi ftp chroot-exception-list

Displays the list of local user chroot exceptions.

Syntax

```
isi ftp chroot-exception-list
```

Options

There are no options for this command.

Examples

To view a list of local user chroot exceptions, run the following command:

```
isi ftp chroot-exception-list
```

The system displays output similar to the following example:

```
chroot-exception-list:
user1
user2
user3
```

isi ftp chroot-exception-list add

Adds users to the `chroot` exception list.

Syntax

```
isi ftp chroot-exception-list add <value>
```

Options

<value>

Specifies the user being added to the `chroot` exception list.

Examples

The following command adds `newUser` to the `chroot` exception list:

```
isi ftp chroot-exception-list add newUser
```

The system displays output similar to the following example:

```
chroot-exception-list: added user 'newUser' to list
```

isi ftp chroot-exception-list remove

Removes users from the `chroot` exception list.

Syntax

```
isi ftp chroot-exception-list remove <value>
```

Options

<value>

Required. Specifies the user being removed from the `chroot` exception list.

Examples

The following command removes `newUser` from the `chroot` exception list:

```
isi ftp chroot-exception-list remove newUser
```

The system displays output similar to the following example:

```
chroot-exception-list: removed user 'newUser' from list
```

isi ftp chroot-list-enable

Displays or specifies the lookup of `chroot` exceptions.

Syntax

```
isi ftp chroot-list-enable <value>
```

Options

<value>

Enables support for using `chroot-exception-list` entries as exceptions to `chroot` mode when set to **YES**. Valid values are **YES** and **NO**. If no options are specified, displays whether the lookup is enabled. The default value is **NO**.

isi ftp chroot-local-mode

Specifies which users are placed in a chroot jail in their home directory after they login.

Syntax

```
isi ftp chroot-local-mode <value>
```

Options

<value>

Specifies which users are placed in a chroot jail in their home directory after they login. The following values are valid:

- **all** All local users are placed in a chroot jail in their home directory after they login.
- **all-with-exceptions** All local users except those in the chroot exception list are placed in a chroot jail in their home directory after they login.
- **none** No local users are placed in a chroot jail in their home directory after they login.
- **none-with-exceptions** Only users in the chroot exception list are placed in a chroot jail in their home directory after they login.

If no options are specified, displays the current setting. The default value is none.

Examples

To place users who are not on the chroot exception list in a chroot jail in their home directory after they login, run the following command:

```
isi ftp chroot-local-mode --value=all-with-exceptions
```

The system displays output similar to the following example:

```
chroot-local-mode: all -> all-with-exceptions \  
chroot-exception-list is active.
```

To place only users in the chroot exception list in a chroot jail in their home directory after they login, run the following command:

```
isi ftp chroot-local-mode --value=none-with-exceptions
```

The system displays output similar to the following example:

```
chroot-local-mode: none -> none-with-exceptions \  
chroot-exception-list is active.
```

isi ftp connect-timeout

Specifies and displays the data connection response timeout.

Syntax

```
isi ftp connect-timeout <value>
```

Options

<value>

Specifies the timeout (in seconds) for a remote client to respond to our PORT style data connection. Valid values are integers between 30 and 600. If no options are

specified, displays the current data connection response timeout. The default value is 60.

Examples

To set the timeout to two minutes, run the following command:

```
isi ftp connect-timeout 120
```

The system displays output similar to the following example:

```
connect-timeout: 60 -> 120
```

isi ftp data-timeout

Specifies the data connection stall timeout.

Syntax

```
isi ftp data-timeout <value>
```

Options

<value>

Specifies the maximum time (in seconds) data transfers are allowed to stall with no progress before the remote client is removed. Valid values are integers between 30 and 600. The default value is 300.

Examples

To set the timeout to 1 minute, run the following command:

```
isi ftp data-timeout 60
```

The system displays output similar to the following example:

```
connect-timeout: 60 -> 120
```

isi ftp denied-user-list

Displays the list of denied users.

Syntax

```
isi ftp denied-user-list
```

Options

There are no options for this command.

Examples

To view the list of denied users, run the following command:

```
isi ftp denied-user-list
```

The system displays output similar to the following example:

```
denied-user-list:
unwelcomeUser1
unwelcomeUser1
unwelcomeUser2
unwelcomeUser4
```

isi ftp denied-user-list add

Adds users to the list of denied users.

Syntax

```
isi ftp denied-user-list add <value>
```

Options

<value>

Specifies the name of the user being added to the denied user list.

Examples

The following command adds *unwelcomeUser* to the list of denied users:

```
isi ftp denied-user-list add unwelcomeUser
```

The system displays output similar to the following example:

```
denied-user-list: added user 'unwelcomeUser' to list
```

isi ftp denied-user-list delete

Removes users from the list of denied users.

Syntax

```
isi ftp denied-user-list delete <value>
```

Options

<value>

Specifies the name of the user being removed from the denied user list.

Examples

The following command removes *approvedUser* from the list of denied users:

```
isi ftp denied-user-list remove approvedUser
```

The system displays output similar to the following example:

```
denied-user-list: removed user 'unwelcomUser' from list
```

isi ftp dirlist-localtime

Specifies and displays whether the time displayed in directory listings is in your local time zone.

Syntax

```
isi ftp dirlist-localtime <value>
```

Options

<value>

Specifies whether the time displayed in directory listings is in your local time zone. Valid values are **YES** and **NO**. If **NO**, time displays on GMT. If **YES** the time displays in your local time zone. If no options are specified, the current setting is displayed. The default value is **NO**.

The last-modified times returned by commands issued inside of the FTP shell are also affected by this parameter.

Examples

To set the time displayed in directory listings to your local time zone, run the following command:

```
isi ftp dirlist-localtime YES
```

The system displays output similar to the following example:

```
dirlist-localtime: NO -> YES
```

isi ftp dirlist-names

Displays and controls what information is displayed about users and groups in directory listings.

Syntax

```
isi ftp dirlist-names <value>
```

Options

<value>

Determines what information is displayed about users and groups in directory listings. The following values are valid:

- **hide** All user and group information in directory listings is displayed as **ftp**.
- **numeric** Numeric IDs are shown in the user and group fields of directory listings.
- **textual** Textual names are shown in the user and group fields of directory listings.

If no options are specified, displays the current setting. The default value is **hide**.

Examples

To show numeric IDs of users and groups in directory listings, run the following command:

```
isi ftp dirlist-names numeric
```

System displays output similar to the following example:

```
dirlist-names: hide -> numeric
```

isi ftp file-create-perm

Specifies and displays file creation permissions

Syntax

```
isi ftp file-create-perm <value>
```

Options

<value>

Specifies the permissions with which uploaded files are created. Valid values are octal permission numbers. If no options are specified, this command displays the current file creation permission setting. The default value is **0666**.

Note

To uploaded files to be executable, consider changing the permissions to 0777.

Examples

To set the octal permission number to 0777, run the following command:

```
isi ftp file-create-perm 0777
```

The system displays output similar to the following example:

```
file-create-perm: 0666 -> 0777
```

isi ftp local-root-path

Displays and specified the root path for local users. VSFTPD attempts to change into the directory specified by the root path after a logical login.

Syntax

```
isi ftp local-root-path --value <ifs-directory>
[--reset]
```

Options

```
{--value | -v} <ifs-directory>
```

Specifies a directory in `/ifs` that VSFTPD attempts to change into after a local login. Valid values are paths in `/ifs`. If no options are specified, the current root path for local users is displayed. The default value is the local user home directory.

--reset

Resets to use the local user home directory.

Examples

The following command sets the root path for local users to `/ifs/home/newUser`:

```
isi ftp local-root-path --value=/ifs/home/newUser
```

The system displays output similar to the following example:

```
local-root-path: local user home directory -> /ifs/home/newUser
```

To set the root path for local users back to the local user home directory, run the following command:

```
isi ftp local-root-path --reset
```

The system displays output similar to the following example:

```
local-root-path: /ifs/home/newUser1 -> local user home directory
```

isi ftp local-umask

Displays and specifies the local user file creation umask.

Syntax

```
isi ftp local-umask <value>
```

Options

<value>

Specifies the umask for file creation by local users. Valid values are octal umask numbers. If no options are specified, displays the current local user file creation umask. The default value is 077.

Note

Value must contain the '0' prefix, otherwise the value will be treated as a base 10 integer.

Examples

The following command sets the local user file creation umask to 066:

```
isi ftp local-umask 066
```

The system displays output similar to the following example:

```
local-umask: 077 -> 066
```

isi ftp server-to-server

Sets and displays whether server-to-server (FXP) transfers are permitted.

Syntax

```
isi ftp server-to-server <value>
```

Options

<value>

Specifies whether or not to allow FXP transfers. Valid values are **YES** and **NO**. If no options are specified, displays current setting. The default value is **NO**.

Examples

To allow FXP transfers, run the following command:

```
isi ftp server-to-server YES
```

The system displays output similar to the following example:

```
server-to-server: NO -> YES
```

isi ftp session-support

Enables or disables FTP session support.

Syntax

```
isi ftp session-support <value>
```

Options

<value>

Valid values are **YES** and **NO**. If set to **YES**, the command maintains login sessions for each user through Pluggable Authentication Modules (PAM). If set to **NO**, the command prevents automatic home directory creation if that functionality is

otherwise available. If no options are specified, displays whether FTP session support is enabled. The default value is **YES**.

isi ftp session-timeout

Specifies and displays the idle system timeout.

Syntax

```
isi ftp session-timeout <value>
```

Options

<value>

Specifies the maximum time (in seconds) that a remote client may spend between FTP commands before the remote client is kicked off. Valid values are integers between 30 and 600. If no options are specified, displays the current idle system timeout. The default value is 300.

Examples

To set the timeout to one minute, run the following command:

```
isi ftp session-timeout 60
```

The system displays output similar to the following example:

```
session-timeout: 300 -> 60
```

isi ftp user-config-dir

Displays and specifies the user configuration directory.

Syntax

```
isi ftp user-config-dir --value<directory>
[--reset]
```

Options

If no options are specified, displays the current user configuration directory path.

{--value | -v} <directory>

Specifies the directory where user-specific configurations that override global configurations can be found. The default value is the local user home directory.

--reset

Reset to use the local user home directory.

Examples

The following command sets the user configuration directory to /ifs/home/User/directory:

```
isi ftp user-config-dir --value=/ifs/home/User/directory
```

To set the user configuration directory back to the local user home directory, run the following command:

```
isi ftp user-config-dir --reset
```

HTTP and HTTPS

OneFS includes a configurable HTTP service, which is used to request files that are stored on the cluster and to interact with the web administration interface.

OneFS supports both HTTP and its secure variant, HTTPS. Each node in the cluster runs an instance of the Apache HTTP Server to provide HTTP access. You can configure the HTTP service to run in different modes.

Both HTTP and HTTPS are supported for file transfer, but only HTTPS is supported for Platform API calls. The HTTPS-only requirement includes the web administration interface. In addition, OneFS supports a form of the web-based DAV (WebDAV) protocol that enables users to modify and manage files on remote web servers. OneFS performs distributed authoring, but does not support versioning and does not perform security checks. You can enable DAV in the web administration interface.

Enable and configure HTTP

You can configure HTTP and DAV to enable users to edit and manage files collaboratively across remote web servers.

This procedure is available only through the web administration interface.

Procedure

1. Click **Protocols > HTTP Settings**.
2. From the **Service** options, select one of the following settings:

Option	Description
Enable HTTP	Allows HTTP access for cluster administration and browsing content on the cluster.
Disable HTTP and redirect to the web interface	Allows only administrative access to the web administration interface. This is the default setting.
Disable HTTP entirely	Closes the HTTP port used for file access. Users can continue to access the web administration interface by specifying the port number in the URL. The default port is 8080.

3. In the **Document root directory** field, type or click **Browse** to navigate to an existing directory in `/ifs`, or click **File System Explorer** to create a new directory and set its permissions.

Note

The HTTP server runs as the daemon user and group. To properly enforce access controls, you must grant the daemon user or group read access to all files under the document root, and allow the HTTP server to traverse the document root.

4. In the **Server hostname** field, type the HTTP server name. The server hostname must be a fully-qualified, SmartConnect zone name and valid DNS name. The name must begin with a letter and contain only letters, numbers, and hyphens (-).
5. In the **Administrator email address** field, type an email address to display as the primary contact for issues that occur while serving files.

6. From the **Authentication** list, select an authentication setting:

Option	Description
Off	Disables HTTP authentication.
Basic Authentication Only	Enables HTTP basic authentication. User credentials are sent in plain text.
Integrated Authentication Only	Enables HTTP authentication via NTLM, Kerberos, or both.
Integrated and Basic Authentication	Enables both basic and integrated authentication.
Basic Authentication with Access Controls	Enables HTTP basic authentication and enables the Apache web server to perform access checks.
Integrated Authentication with Access Controls	Enables HTTP integrated authentication via NTLM and Kerberos, and enables the Apache web server to perform access checks.
Integrated and Basic Auth with Access Controls	Enables HTTP basic authentication and integrated authentication, and enables the Apache web server to perform access checks.

7. Click the **Enable DAV** check box. This allows multiple users to manage and modify files collaboratively across remote web servers.
8. Click the **Disable access logging** check box.
9. Click **Submit**.

CHAPTER 10

Home directories

This section contains the following topics:

- [Home directories overview](#) 464
- [Home directory permissions](#) 464
- [Authenticating SMB users](#) 464
- [Home directory creation through SMB](#) 464
- [Home directory creation through SSH and FTP](#) 468
- [Home directory creation in a mixed environment](#) 471
- [Interactions between ACLs and mode bits](#) 471
- [Default home directory settings in authentication providers](#) 471
- [Supported expansion variables](#) 472
- [Domain variables in home directory provisioning](#) 474

Home directories overview

When you create a local user, OneFS automatically creates a home directory for the user. OneFS also supports dynamic home directory provisioning for users who access the cluster by connecting to an SMB share or by logging in through FTP or SSH. Regardless of the method by which a home directory was created, you can configure access to the home directory through a combination of SMB, SSH, and FTP.

Home directory permissions

You can set up a user's home directory with a Windows ACL or with POSIX mode bits, which are then converted into a synthetic ACL. The method by which a home directory is created determines the initial permissions that are set on the home directory.

When you create a local user, the user's home directory is created with mode bits by default.

For users who authenticate against external sources, you can specify settings to create home directories dynamically at login time. If a home directory is created during a login through SSH or FTP, it is set up with mode bits; if a home directory is created during an SMB connection, it receives either mode bits or an ACL. For example, if an LDAP user first logs in through SSH or FTP, the user's home directory is created with mode bits. If the same user first connects through an SMB share, the home directory is created with the permissions indicated by the configured SMB settings. If the `--inheritable-path-acl` option is enabled, an ACL is generated; otherwise, mode bits are used.

Authenticating SMB users

You can authenticate SMB users from authentication providers that can handle NT hashes.

SMB sends an NT password hash to authenticate SMB users, so only users from authentication providers that can handle NT hashes can log in over SMB. The following OneFS-supported authentication providers can handle NT hashes:

- Active Directory
- Local
- LDAPSAM (LDAP with Samba extensions enabled)

Home directory creation through SMB

You can create SMB shares by including expansion variables in the share path. Expansion variables give users to access their home directories by connecting to the share. You can also enable dynamic provisioning of home directories that do not exist at SMB connection time.

Note

Share permissions are checked when files are accessed, before the underlying file system permissions are checked. Either of these permissions can prevent access to the file or directory.

Create home directories with expansion variables

You can configure settings with expansion variables to create SMB share home directories.

When users access the EMC Isilon cluster over SMB, home directory access is through SMB shares. You can configure settings with a path that uses a variable expansion syntax, allowing a user to connect to their home directory share.

Note

Home directory share paths must begin with `/ifs/` and must be in the root path of the access zone in which the home directory SMB share is created.

In the following commands, the `--allow-variable-expansion` option is enabled to indicate that `%U` should be expanded to the user name, which is `user411` in this example. The `--auto-create-directory` option is enabled to create the directory if it does not exist:

```
isi smb shares create HOMEDIR --path=/ifs/home/%U \
  --allow-variable-expansion=yes --auto-create-directory=yes
isi smb shares permission modify HOMEDIR --wellknown Everyone \
  --permission-type allow --permission full
isi smb shares view HOMEDIR
```

The system displays output similar to the following example:

```
Share Name: HOMEDIR
Path: /ifs/home/%U
Description:
Client-side Caching Policy: manual
Automatically expand user names or domain names: True
Automatically create home directories for users: True
Browsable: True
Permissions:
Account Account Type Run as Root Permission Type Permission
-----
Everyone wellknown False allow full
-----
Total: 1
...
```

When `user411` connects to the share with the `net use` command, the user's home directory is created at `/ifs/home/user411`. On `user411`'s Windows client, the `net use m: command` connects `/ifs/home/user411` through the HOMEDIR share:

```
net use m: \\cluster.company.com\HOMEDIR /u:user411
```

Procedure

1. Run the following commands on the cluster with the `--allow-variable-expansion` option enabled. The `%U` expansion variable expands to the user name, and the `--auto-create-directory` option is enabled to create the directory if it does not exist:

```
isi smb shares create HOMEDIR --path=/ifs/home/%U \
  --allow-variable-expansion=yes --auto-create-directory=yes
isi smb shares permission modify HOMEDIR --wellknown Everyone \
  --permission-type allow --permission full
```

2. Run the following command to view the home directory settings:

```
isi smb shares view HOMEDIR
```

The system displays output similar to the following example:

```

Share Name: HOMEDIR
Path: /ifs/home/%U
Description:
Client-side Caching Policy: manual
Automatically expand user names or domain names: True
Automatically create home directories for users: True
Browsable: True
Permissions:
Account Account Type Run as Root Permission Type Permission
-----
Everyone wellknown False allow full
-----
Total: 1
...

```

If user411 connects to the share with the `net use` command, user411's home directory is created at `/ifs/home/user411`. On user411's Windows client, the `net use m: command` connects `/ifs/home/user411` through the HOMEDIR share, mapping the connection similar to the following example:

```
net use m: \\cluster.company.com\HOMEDIR /u:user411
```

Create home directories with the `--inheritable-path-acl` option

You can enable the `--inheritable-path-acl` option on a share to specify that it is to be inherited on the share path if the parent directory has an inheritable ACL.

Before you begin

To perform most configuration tasks, you must log on as a member of the SecurityAdmin role.

By default, an SMB share's directory path is created with a synthetic ACL based on mode bits. You can enable the `--inheritable-path-acl` option to use the inheritable ACL on all directories that are created, either at share creation time or for those dynamically provisioned when connecting to that share.

Procedure

1. Run commands similar to the following examples to enable the `--inheritable-path-acl` option on the cluster to dynamically provision a user home directory at first connection to a share on the cluster:

```
isi smb shares create HOMEDIR_ACL --path=/ifs/home/%U \
--allow-variable-expansion=yes --auto-create-directory=yes \
--inheritable-path-acl=yes
```

```
isi smb shares permission modify HOMEDIR_ACL \
--wellknown Everyone \
--permission-type allow --permission full
```

2. Run a `net use` command, similar to the following example, on a Windows client to map the home directory for user411:

```
net use q: \\cluster.company.com\HOMEDIR_ACL /u:user411
```

- Run a command similar to the following example on the cluster to view the inherited ACL permissions for the user411 share:

```
cd /ifs/home/user411
ls -lde .
```

The system displays output similar to the following example:

```
drwx-----+ 2 user411 Isilon Users 0 Oct 19 16:23 ./
OWNER: user:user411
GROUP: group:Isilon Users
CONTROL:dacl_auto_inherited,dacl_protected
0: user:user411 allow dir_gen_all,object_inherit,container_inherit
```

Create special home directories with the SMB share %U variable

The special SMB share name %U enables you to create a home-directory SMB share that appears the same as a user's user name.

You typically set up a %U SMB share with a share path that includes the %U expansion variable. If a user attempts to connect to a share matching the login name and it does not exist, the user connects to the %U share instead and is directed to the expanded path for the %U share.

Note

If another SMB share exists that matches the user's name, the user connects to the explicitly named share rather than to the %U share.

Procedure

- Run the following command to create a share that matches the authenticated user login name when the user connects to the share:

```
isi smb share create %U /ifs/home/%U \
  --allow-variable-expansion=yes --auto-create-directory=yes \
  --zone=System
```

After running this command, user Zachary will see a share named 'zachary' rather than '%U', and when Zachary tries to connect to the share named 'zachary', he will be directed to /ifs/home/zachary. On a Windows client, if Zachary runs the following commands, he sees the contents of his /ifs/home/zachary directory:

```
net use m: \\cluster.ip\zachary /u:zachary
cd m:
dir
```

Similarly, if user Claudia runs the following commands on a Windows client, she sees the directory contents of /ifs/home/claudia:

```
net use m: \\cluster.ip\claudia /u:claudia
cd m:
dir
```

Zachary and Claudia cannot access one another's home directory because only the share 'zachary' exists for Zachary and only the share 'claudia' exists for Claudia.

Home directory creation through SSH and FTP

You can configure home directory support for users who access the cluster through SSH or FTP by modifying authentication provider settings.

Set the SSH or FTP login shell

You can use the `--login-shell` option to set the default login shell for the user.

By default, the `--login-shell` option, if specified, overrides any login-shell information provided by the authentication provider, except with Active Directory. If the `--login-shell` option is specified with Active Directory, it simply represents the default login shell if the Active Directory server does not provide login-shell information.

Procedure

1. Run the following command to set the login shell for all local users to `/bin/bash`:

```
isi auth local modify System --login-shell /bin/bash
```

2. Run the following command to set the default login shell for all Active Directory users in your domain to `/bin/bash`:

```
isi auth ads modify YOUR.DOMAIN.NAME.COM --login-shell /bin/bash
```

Set SSH/FTP home directory permissions

You can specify home directory permissions for a home directory that is accessed through SSH or FTP by setting a `umask` value.

Before you begin

To perform most configuration tasks, you must log on as a member of the SecurityAdmin role.

When a user's home directory is created at login through SSH or FTP, it is created using POSIX mode bits. The permissions setting on a user's home directory is set to `0755`, then masked according to the `umask` setting of the user's access zone to further limit permissions. You can modify the `umask` setting for a zone with the `--home-directory-umask` option, specifying an octal number as the `umask` value.

Procedure

1. Run the following command to view `umask` setting:

```
isi zone zones view System
```

The system displays output similar to the following example:

```
Name: System
Path: /ifs
Cache Size: 4.77M
Map Untrusted:
Auth Providers: -
NetBIOS Name:
All Auth Providers: Yes
```

```

User Mapping Rules: -
Home Directory Umask: 0077
  Skeleton Directory: /usr/share/skel
    Audit Success: create, delete, rename, set_security, close
    Audit Failure: create, delete, rename, set_security, close
HDFS Authentication: all
  HDFS Keytab: /etc/hdfs.keytab
HDFS Root Directory: /ifs
  WebHDFS Enabled: Yes
Syslog Forwarding Enabled: No
  Syslog Audit Events: create, delete, rename, set_security
  Zone ID: 1

```

In the command result, you can see the default setting for Home Directory Umask for the created home directory is 0700, which is equivalent to (0755 & ~(077)). You can modify the Home Directory Umask setting for a zone with the `--home-directory-umask` option, specifying an octal number as the umask value. This value indicates the permissions that are to be disabled, so larger mask values indicate fewer permissions. For example, a umask value of 000 or 022 yields created home directory permissions of 0755, whereas a umask value of 077 yields created home directory permissions of 0700.

2. Run a command similar to the following example to allow a group/others write/execute permission in a home directory:

```
isi zone zones modify System --home-directory-umask=022
```

In this example, user home directories will be created with mode bits 0755 masked by the umask field, set to the value of 022. Therefore, user home directories will be created with mode bits 0755, which is equivalent to (0755 & ~(022)).

Set SSH/FTP home directory creation options

You can configure home directory support for a user who accesses the cluster through SSH or FTP by specifying authentication provider options.

Procedure

1. Run the following command to view settings for an Active Directory authentication provider on the cluster:

```
isi auth ads list
```

The system displays output similar to the following example:

```

Name                Authentication Status DC Name Site
-----
YOUR.DOMAIN.NAME.COM Yes                online -     SEA
-----
Total: 1

```

2. Run the `isi auth ads modify` command with the `--home-directory-template` and `--create-home-directory` options.

```

isi auth ads modify YOUR.DOMAIN.NAME.COM \
--home-directory-template=/ifs/home/ADS/%D/%U \
--create-home-directory=yes

```

3. Run the `isi auth ads view` command with the `--verbose` option.

The system displays output similar to the following example:

```
Name: YOUR.DOMAIN.NAME.COM
NetBIOS Domain: YOUR
...
Create Home Directory: Yes
Home Directory Template: /ifs/home/ADS/%D/%U
Login Shell: /bin/sh
```

4. Run the `id` command.

The system displays output similar to the following example:

```
uid=1000008(<your-domain>\user_100) gid=1000000(<your-domain>
\domain users)
groups=1000000(<your-domain>\domain users),1000024(<your-domain>
\c1t),1545(Users)
```

5. (Optional) To verify this information from an external UNIX node, run the `ssh` command from an external UNIX node.

For example, the following command would create `/ifs/home/ADS/<your-domain>/user_100` if it did not previously exist:

```
ssh <your-domain>\\user_100@cluster.isilon.com
```

Provision home directories with dot files

You can provision home directories with dot files.

Before you begin

To perform most configuration tasks, you must log on as a member of the SecurityAdmin role.

The skeleton directory, which is located at `/usr/share/skel` by default, contains a set of files that are copied to the user's home directory when a local user is created or when a user home directory is dynamically created during login. Files in the skeleton directory that begin with `dot.` are renamed to remove the `dot` prefix when they are copied to the user's home directory. For example, `dot.cshrc` is copied to the user's home directory as `.cshrc`. This format enables dot files in the skeleton directory to be viewable through the command-line interface without requiring the `ls -a` command.

For SMB shares that might use home directories that were provisioned with dot files, you can set an option to prevent users who connect to the share through SMB from viewing the dot files.

Procedure

1. Run the following command to display the default skeleton directory in the System access zone:

```
isi zone zones view System
```

The system displays output similar to the following example:

```
Name: System
...
Skeleton Directory: /usr/share/skel
```

2. Run the `isi zone zones modify` command to modify the default skeleton directory.

The following command modifies the default skeleton directory, `/usr/share/skel`, in an access zone, where `System` is the value for the `<zone>` option and `/usr/share/skel2` is the value for the `<path>` option:

```
isi zone zones modify System --skeleton-directory=/usr/share/skel2
```

Home directory creation in a mixed environment

If a user logs in through both SMB and SSH, it is recommended that you configure home directory settings so the path template is the same for the SMB share and each authentication provider against which the user is authenticating through SSH.

Interactions between ACLs and mode bits

Home directory setup is determined by several factors, including how users authenticate and the options that specify home directory creation.

A user's home directory may be set up with either ACLs or POSIX mode bits, which are converted into a synthetic ACL. The directory of a local user is created when the local user is created, and the directory is set up with POSIX mode bits by default. Directories can be dynamically provisioned at log in for users who authenticate against external sources, and in some cases for users who authenticate against the File provider. In this situation, the user home directory is created according to how the user first logs in.

For example, if an LDAP user first logs in through SSH or FTP and the user home directory is created, it is created with POSIX mode bits. If that same user first connects through an SMB home directory share, the home directory is created as specified by the SMB option settings. If the `--inherited-path-acl` option is enabled, ACLs are generated.

Otherwise, POSIX mode bits are used.

Default home directory settings in authentication providers

The default settings that affect how home directories are set up differ, based on the authentication provider that the user authenticates against.

Authentication provider	Home directory	Home directory creation	UNIX login shell
Local	<ul style="list-style-type: none"> <code>--home-directory-template=/ifs/home/%U</code> <code>--create-home-directory=yes</code> <code>--login-shell=/bin/sh</code> 	Enabled	/bin/sh
File	<ul style="list-style-type: none"> <code>--home-directory-template=""</code> 	Disabled	None

Authentication provider	Home directory	Home directory creation	UNIX login shell
	<ul style="list-style-type: none"> • <code>--create-home-directory=no</code> 		
Active Directory	<ul style="list-style-type: none"> • <code>--home-directory-template=/ifs/home/%D/%U</code> • <code>--create-home-directory=no</code> • <code>--login-shell=/bin/sh</code> <hr/> <p>Note If available, provider information overrides this value.</p> <hr/>	Disabled	/bin/sh
LDAP	<ul style="list-style-type: none"> • <code>--home-directory-template=""</code> • <code>--create-home-directory=no</code> 	Disabled	None
NIS	<ul style="list-style-type: none"> • <code>--home-directory-template=""</code> • <code>--create-home-directory=no</code> 	Disabled	None

Supported expansion variables

You can include expansion variables in an SMB share path or in an authentication provider's home directory template.

OneFS supports the following expansion variables. You can improve performance and reduce the number of shares to be managed when you configure shares with expansion variables. For example, you can include the %U variable for a share rather than create a share for each user. When a %U is included in the name so that each user's path is different, security is still ensured because each user can view and access only his or her home directory.

Note

When you create an SMB share through the web administration interface, you must select the **Allow Variable Expansion** check box or the string is interpreted literally by the system.

Variable	Value	Description
%U	User name (for example, user_001)	Expands to the user name to allow different users to use different home directories. This variable is typically included at the end of the path. For example, for a user named user1, the path <code>/ifs/home/%U</code> is mapped to <code>/ifs/home/user1</code> .
%D	NetBIOS domain name (for example, YORK for YORK.EAST.EXAMPLE.COM)	Expands to the user's domain name, based on the authentication provider: <ul style="list-style-type: none"> For Active Directory users, %D expands to the Active Directory NetBIOS name. For local users, %D expands to the cluster name in uppercase characters. For example, for a cluster named cluster1, %D expands to CLUSTER1. For users in the System file provider, %D expands to UNIX_USERS. For users in other file providers, %D expands to FILE_USERS. For LDAP users, %D expands to LDAP_USERS. For NIS users, %D expands to NIS_USERS.
%Z	Zone name (for example, ZoneABC)	Expands to the access zone name. If multiple zones are activated, this variable is useful for differentiating users in separate zones. For example, for a user named user1 in the System zone, the path <code>/ifs/home/%Z/%U</code> is mapped to <code>/ifs/home/System/user1</code> .
%L	Host name (cluster host name in lowercase)	Expands to the host name of the cluster, normalized to lowercase. Limited use.
%0	First character of the user name	Expands to the first character of the user name.
%1	Second character of the user name	Expands to the second character of the user name.
%2	Third character of the user name	Expands to the third character of the user name.

Note

If the user name includes fewer than three characters, the %0, %1, and %2 variables wrap around. For example, for a user named ab, the variables maps to a, b, and a, respectively. For a user named a, all three variables map to a.

Domain variables in home directory provisioning

You can use domain variables to specify authentication providers when provisioning home directories.

The domain variable (%D) is typically used for Active Directory users, but it has a value set that can be used for other authentication providers. %D expands as described in the following table for the various authentication providers.

Authenticated user	%D expansion
Active Directory user	Active Directory NetBIOS name—for example, YORK for provider YORK.EAST.EXAMPLE.COM.
Local user	The cluster name in all-uppercase characters—for example, if the cluster is named MyCluster, %D expands to MYCLUSTER.
File user	<ul style="list-style-type: none"> • UNIX_USERS (for System file provider) • FILE_USERS (for all other file providers)
LDAP user	LDAP_USERS (for all LDAP authentication providers)
NIS user	NIS_USERS (for all NIS authentication providers)

CHAPTER 11

Snapshots

This section contains the following topics:

- [Snapshots overview](#)..... 476
- [Data protection with SnapshotIQ](#).....476
- [Snapshot disk-space usage](#)..... 476
- [Snapshot schedules](#).....477
- [Snapshot aliases](#)..... 477
- [File and directory restoration](#).....477
- [Best practices for creating snapshots](#)..... 478
- [Best practices for creating snapshot schedules](#)..... 478
- [File clones](#).....479
- [Snapshot locks](#)..... 480
- [Snapshot reserve](#)..... 481
- [SnapshotIQ license functionality](#).....481
- [Creating snapshots with SnapshotIQ](#).....481
- [Managing snapshots](#) 485
- [Restoring snapshot data](#)..... 488
- [Managing snapshot schedules](#).....490
- [Managing snapshot aliases](#).....491
- [Managing with snapshot locks](#)..... 493
- [Configure SnapshotIQ settings](#) 494
- [Set the snapshot reserve](#).....496
- [Snapshot commands](#)..... 496

Snapshots overview

A OneFS snapshot is a logical pointer to data that is stored on a cluster at a specific point in time.

A snapshot references a directory on a cluster, including all data stored in the directory and its subdirectories. If the data referenced by a snapshot is modified, the snapshot stores a physical copy of the data that was modified. Snapshots are created according to user specifications or are automatically generated by OneFS to facilitate system operations.

To create and manage snapshots, you must activate a SnapshotIQ license on the cluster. Some applications must generate snapshots to function but do not require you to activate a SnapshotIQ license; by default, these snapshots are automatically deleted when OneFS no longer needs them. However, if you activate a SnapshotIQ license, you can retain these snapshots. You can view snapshots generated by other modules without activating a SnapshotIQ license.

You can identify and locate snapshots by name or ID. A snapshot name is specified by a user and assigned to the virtual directory that contains the snapshot. A snapshot ID is a numerical identifier that OneFS automatically assigns to a snapshot.

Data protection with SnapshotIQ

You can create snapshots to protect data with the SnapshotIQ software module. Snapshots protect data against accidental deletion and modification by enabling you to restore deleted and modified files. To use SnapshotIQ, you must activate a SnapshotIQ license on the cluster.

Snapshots are less costly than backing up your data on a separate physical storage device in terms of both time and storage consumption. The time required to move data to another physical device depends on the amount of data being moved, whereas snapshots are always created almost instantaneously regardless of the amount of data referenced by the snapshot. Also, because snapshots are available locally, end-users can often restore their data without requiring assistance from a system administrator. Snapshots require less space than a remote backup because unaltered data is referenced rather than recreated.

Snapshots do not protect against hardware or file-system issues. Snapshots reference data that is stored on a cluster, so if the data on the cluster becomes unavailable, the snapshots will also be unavailable. Because of this, it is recommended that you back up your data to separate physical devices in addition to creating snapshots.

Snapshot disk-space usage

The amount of disk space that a snapshot consumes depends on both the amount of data stored by the snapshot and the amount of data the snapshot references from other snapshots.

Immediately after OneFS creates a snapshot, the snapshot consumes a negligible amount of disk space. The snapshot does not consume additional disk space unless the data referenced by the snapshot is modified. If the data that a snapshot references is modified, the snapshot stores read-only copies of the original data. A snapshot consumes only the space that is necessary to restore the contents a directory to the state it was in when the snapshot was taken.

To reduce disk-space usage, snapshots that reference the same directory reference each other, with older snapshots referencing newer snapshots. If a file is deleted, and several snapshots reference the file, a single snapshot stores a copy of the file, and the other snapshots reference the file from the snapshot that stored the copy. The reported size of a snapshot reflects only the amount of data stored by the snapshot and does not include the amount of data referenced by the snapshot.

Because snapshots do not consume a set amount of storage space, there is no available-space requirement for creating a snapshot. The size of a snapshot grows according to how the data referenced by the snapshot is modified. A cluster cannot contain more than 20,000 snapshots.

Snapshot schedules

You can automatically generate snapshots according to a snapshot schedule.

With snapshot schedules, you can periodically generate snapshots of a directory without having to manually create a snapshot every time. You can also assign an expiration period that determines when SnapshotIQ deletes each automatically generated snapshot.

Snapshot aliases

A snapshot alias is a logical pointer to a snapshot. If you specify an alias for a snapshot schedule, the alias will always point to the most recent snapshot generated by that schedule. Assigning a snapshot alias allows you to quickly identify and access the most recent snapshot generated according to a snapshot schedule.

If you allow clients to access snapshots through an alias, you can reassign the alias to redirect clients to other snapshots. In addition to assigning snapshot aliases to snapshots, you can also assign snapshot aliases to the live version of the file system. This can be useful if clients are accessing snapshots through a snapshot alias, and you want to redirect the clients to the live version of the file system.

File and directory restoration

You can restore the files and directories that are referenced by a snapshot by copying data from the snapshot, cloning a file from the snapshot, or reverting the entire snapshot.

Copying a file from a snapshot duplicates the file, which roughly doubles the amount of storage space consumed. Even if you delete the original file from the non-snapshot directory, the copy of the file remains in the snapshot.

Cloning a file from a snapshot also duplicates the file. However, unlike a copy, which immediately consumes additional space on the cluster, a clone does not consume any additional space on the cluster unless the clone or cloned file is modified.

Reverting a snapshot replaces the contents of a directory with the data stored in the snapshot. Before a snapshot is reverted, SnapshotIQ creates a snapshot of the directory that is being replaced, which enables you to undo the snapshot revert later. Reverting a snapshot can be useful if you want to undo a large number of changes that you made to files and directories. If new files or directories have been created in a directory since a snapshot of the directory was created, those files and directories are deleted when the snapshot is reverted.

Note

If you move a directory, you cannot revert snapshots of the directory that were taken before the directory was moved.

Best practices for creating snapshots

Consider the following snapshot best practices when working with a large number of snapshots.

It is recommended that you do not create more than 1,000 snapshots of a single directory to avoid performance degradation. If you create a snapshot of a root directory, that snapshot counts towards the total number of snapshots for any subdirectories of the root directory. For example, if you create 500 snapshots of `/ifs/data` and 500 snapshots of `/ifs/data/media`, you have created 1000 snapshots of `/ifs/data/media`. Avoid creating snapshots of directories that are already referenced by other snapshots.

It is recommended that you do not create more than 1000 hard links per file in a snapshot to avoid performance degradation. Always attempt to keep directory paths as shallow as possible. The deeper the depth of directories referenced by snapshots, the greater the performance degradation.

Creating snapshots of directories higher on a directory tree will increase the amount of time it takes to modify the data referenced by the snapshot and require more cluster resources to manage the snapshot and the directory. However, creating snapshots of directories lower on directories trees will require more snapshot schedules, which can be difficult to manage. It is recommended that you do not create snapshots of `/ifs` or `/ifs/data`.

You can create up to 20,000 snapshots on a cluster at a time. If you create a large number of snapshots, you might not be able to manage snapshots through the OneFS web administration interface. However, you can manage any number of snapshots through the OneFS command-line interface.

Note

It is recommended that you do not disable the snapshot delete job. Disabling the snapshot delete job prevents unused disk space from being freed and can also cause performance degradation.

If the system clock is set to a time zone other than Coordinated Universal Time (UTC), SnapshotIQ modifies snapshot duration periods to match Daylight Savings Time (DST). Upon entering DST, snapshot durations are increased by an hour to adhere to DST; when exiting DST, snapshot durations are decreased by an hour to adhere to standard time.

Best practices for creating snapshot schedules

Snapshot schedule configurations can be categorized by how they delete snapshots: ordered deletions and unordered deletions.

An ordered deletion is the deletion of the oldest snapshot of a directory. An unordered deletion is the deletion of a snapshot that is not the oldest snapshot of a directory. Unordered deletions take approximately twice as long to complete and consume more cluster resources than ordered deletions. However, unordered deletions can save space by retaining a smaller total number of snapshots.

The benefits of unordered deletions versus ordered deletions depend on how often the data referenced by the snapshots is modified. If the data is modified frequently, unordered deletions will save space. However, if data remains unmodified, unordered deletions will most likely not save space, and it is recommended that you perform ordered deletions to free cluster resources.

To implement ordered deletions, assign the same duration period for all snapshots of a directory. The snapshots can be created by one or multiple snapshot schedules. Always ensure that no more than 1000 snapshots of a directory are created.

To implement unordered snapshot deletions, create several snapshot schedules for a single directory, and then assign different snapshot duration periods for each schedule. Ensure that all snapshots are created at the same time when possible.

The following table describes snapshot schedules that follow snapshot best practices:

Table 16 Snapshot schedule configurations

Deletion type	Snapshot frequency	Snapshot time	Snapshot expiration	Max snapshots retained
Ordered deletion (for mostly static data)	Every hour	Beginning at 12:00 AM Ending at 11:59 AM	1 month	720
Unordered deletion (for frequently modified data)	Every other hour	Beginning at 12:00 AM Ending at 11:59 PM	1 day	27
	Every day	At 12:00 AM	1 week	
	Every week	Saturday at 12:00 AM	1 month	
	Every month	The first Saturday of the month at 12:00 AM	3 months	

File clones

SnapshotIQ enables you to create file clones that share blocks with existing files in order to save space on the cluster. A file clone usually consumes less space and takes less time to create than a file copy. Although you can clone files from snapshots, clones are primarily used internally by OneFS.

The blocks that are shared between a clone and cloned file are contained in a hidden file called a shadow store. Immediately after a clone is created, all data originally contained in the cloned file is transferred to a shadow store. Because both files reference all blocks from the shadow store, the two files consume no more space than the original file; the clone does not take up any additional space on the cluster. However, if the cloned file or clone is modified, the file and clone will share only blocks that are common to both of them, and the modified, unshared blocks will occupy additional space on the cluster.

Over time, the shared blocks contained in the shadow store might become useless if neither the file nor clone references the blocks. The cluster routinely deletes blocks that are no longer needed. You can force the cluster to delete unused blocks at any time by running the ShadowStoreDelete job.

Clones cannot contain alternate data streams (ADS). If you clone a file that contains alternate data streams, the clone will not contain the alternate data streams.

Shadow-store considerations

Shadow stores are hidden files that are referenced by cloned and deduplicated files. Files that reference shadow stores behave differently than other files.

- Reading shadow-store references might be slower than reading data directly. Specifically, reading non-cached shadow-store references is slower than reading non-cached data. Reading cached shadow-store references takes no more time than reading cached data.
- When files that reference shadow stores are replicated to another Isilon cluster or backed up to a Network Data Management Protocol (NDMP) backup device, the shadow stores are not transferred to the target Isilon cluster or backup device. The files are transferred as if they contained the data that they reference from shadow stores. On the target Isilon cluster or backup device, the files consume the same amount of space as if they had not referenced shadow stores.
- When OneFS creates a shadow store, OneFS assigns the shadow store to a storage pool of a file that references the shadow store. If you delete the storage pool that a shadow store resides on, the shadow store is moved to a pool occupied by another file that references the shadow store.
- OneFS does not delete a shadow-store block immediately after the last reference to the block is deleted. Instead, OneFS waits until the ShadowStoreDelete job is run to delete the unreferenced block. If a large number of unreferenced blocks exist on the cluster, OneFS might report a negative deduplication savings until the ShadowStoreDelete job is run.
- Shadow stores are protected at least as much as the most protected file that references it. For example, if one file that references a shadow store resides in a storage pool with +2 protection and another file that references the shadow store resides in a storage pool with +3 protection, the shadow store is protected at +3.
- Quotas account for files that reference shadow stores as if the files contained the data referenced from shadow stores; from the perspective of a quota, shadow-store references do not exist. However, if a quota includes data protection overhead, the quota does not account for the data protection overhead of shadow stores.

Snapshot locks

A snapshot lock prevents a snapshot from being deleted. If a snapshot has one or more locks applied to it, the snapshot cannot be deleted and is referred to as a locked snapshot. If the duration period of a locked snapshot expires, OneFS will not delete the snapshot until all locks on the snapshot have been deleted.

OneFS applies snapshot locks to ensure that snapshots generated by OneFS applications are not deleted prematurely. For this reason, it is recommended that you do not delete snapshot locks or modify the duration period of snapshot locks.

A limited number of locks can be applied to a snapshot at a time. If you create snapshot locks, the limit for a snapshot might be reached, and OneFS could be unable to apply a snapshot lock when necessary. For this reason, it is recommended that you do not create snapshot locks.

Snapshot reserve

The snapshot reserve enables you to set aside a minimum percentage of the cluster storage capacity specifically for snapshots. If specified, all other OneFS operations are unable to access the percentage of cluster capacity that is reserved for snapshots.

Note

The snapshot reserve does not limit the amount of space that snapshots can consume on the cluster. Snapshots can consume a greater percentage of storage capacity specified by the snapshot reserve. It is recommended that you do not specify a snapshot reserve.

SnapshotIQ license functionality

You can create snapshots only if you activate a SnapshotIQ license on a cluster. However, you can view snapshots and snapshot locks that are created for internal use by OneFS without activating a SnapshotIQ license.

The following table describes what snapshot functionality is available depending on whether the SnapshotIQ license is active:

	Inactive	Active
Create snapshots and snapshot schedules	No	Yes
Configure SnapshotIQ settings	No	Yes
View snapshot schedules	Yes	Yes
Delete snapshots	Yes	Yes
Access snapshot data	Yes	Yes
View snapshots	Yes	Yes

If you a SnapshotIQ license becomes inactive, you will no longer be able to create new snapshots, all snapshot schedules will be disabled, and you will not be able to modify snapshots or snapshot settings. However, you will still be able to delete snapshots and access data contained in snapshots.

Creating snapshots with SnapshotIQ

To create snapshots, you must configure the SnapshotIQ licence on the cluster. You can create snapshots either by creating a snapshot schedule or manually generating an individual snapshot.

Manual snapshots are useful if you want to create a snapshot immediately, or at a time that is not specified in a snapshot schedule. For example, if you plan to make changes to your file system, but are unsure of the consequences, you can capture the current state of the file system in a snapshot before you make the change.

Before creating snapshots, consider that reverting a snapshot requires that a SnapRevert domain exist for the directory that is being reverted. If you intend on reverting snapshots for a directory, it is recommended that you create SnapRevert domains for those

directories while the directories are empty. Creating a domain for a directory that contains less data takes less time.

Create a SnapRevert domain

Before you can revert a snapshot that contains a directory, you must create a SnapRevert domain for the directory. It is recommended that you create SnapRevert domains for a directory while the directory is empty.

The root path of the SnapRevert domain must be the same root path of the snapshot. For example, a domain with a root path of `/ifs/data/media` cannot be used to revert a snapshot with a root path of `/ifs/data/media/archive`. To revert `/ifs/data/media/archive`, you must create a SnapRevert domain with a root path of `/ifs/data/media/archive`.

Procedure

1. Run the `isi job jobs start` command.

The following command creates a SnapRevert domain for `/ifs/data/media`:

```
isi job jobs start domainmark --root /ifs/data/media \
--dm-type SnapRevert
```

Create a snapshot schedule

You can create a snapshot schedule to continuously generate snapshots of directories.

Procedure

1. Run the `isi snapshot schedules create` command.

The following command creates a snapshot schedule for `/ifs/data/media`:

```
isi snapshot schedules create hourly /ifs/data/media \
HourlyBackup_%m-%d-%Y_%H:%M "Every day every hour" \
--duration 1M
```

The following commands create multiple snapshot schedules for `/ifs/data/media` that generate and expire snapshots at different rates:

```
isi snapshot schedules create every-other-hour \
/ifs/data/media EveryOtherHourBackup_%m-%d-%Y_%H:%M \
"Every day every 2 hours" --duration 1D
isi snapshot schedules create daily /ifs/data/media \
Daily_%m-%d-%Y_%H:%M "Every day at 12:00 AM" --duration 1W
isi snapshot schedules create weekly /ifs/data/media \
Weekly_%m-%d-%Y_%H:%M "Every Saturday at 12:00 AM" --duration 1M
isi snapshot schedules create monthly /ifs/data/media \
Monthly_%m-%d-%Y_%H:%M \
"The 1 Saturday of every month at 12:00 AM" --duration 3M
```

Create a snapshot

You can create a snapshot of a directory.

Procedure

1. Run the `isi snapshot snapshots create` command.

The following command creates a snapshot for `/ifs/data/media`:

```
isi snapshot snapshots create /ifs/data/media --name media-snap
```

Snapshot naming patterns

If you schedule snapshots to be automatically generated, either according to a snapshot schedule or a replication policy, you must assign a snapshot naming pattern that determines how the snapshots are named. Snapshot naming patterns contain variables that include information about how and when the snapshot was created.

The following variables can be included in a snapshot naming pattern:

Variable	Description
%A	The day of the week.
%a	The abbreviated day of the week. For example, if the snapshot is generated on a Sunday, %a is replaced with Sun.
%B	The name of the month.
%b	The abbreviated name of the month. For example, if the snapshot is generated in September, %b is replaced with Sep.
%C	The first two digits of the year. For example, if the snapshot is created in 2014, %C is replaced with 20.
%c	The time and day. This variable is equivalent to specifying %a %b %e %T %Y.
%d	The two digit day of the month.
%e	The day of the month. A single-digit day is preceded by a blank space.
%F	The date. This variable is equivalent to specifying %Y-%m-%d
%G	The year. This variable is equivalent to specifying %Y. However, if the snapshot is created in a week that has less than four days in the current year, the year that contains the majority of the days of the week is displayed. The first day of the week is calculated as Monday. For example, if a snapshot is created on Sunday, January 1, 2017, %G is replaced with 2016, because only one day of that week is in 2017.
%g	The abbreviated year. This variable is equivalent to specifying %y. However, if the snapshot was created in a week that has less than four days in the current year, the year that contains the majority of the days of the week is displayed. The first day of the week is calculated as Monday. For example, if a snapshot is created on Sunday, January 1, 2017, %g is replaced with 16, because only one day of that week is in 2017.
%H	The hour. The hour is represented on the 24-hour clock. Single-digit hours are preceded by a zero. For example, if a snapshot is created at 1:45 AM, %H is replaced with 01.
%h	The abbreviated name of the month. This variable is equivalent to specifying %b.
%I	The hour represented on the 12-hour clock. Single-digit hours are preceded by a zero. For example, if a snapshot is created at 1:45 PM, %I is replaced with 01.

Variable	Description
%j	The numeric day of the year. For example, if a snapshot is created on February 1, %j is replaced with 32.
%k	The hour represented on the 24-hour clock. Single-digit hours are preceded by a blank space.
%l	The hour represented on the 12-hour clock. Single-digit hours are preceded by a blank space. For example, if a snapshot is created at 1:45 AM, %I is replaced with 1.
%M	The two-digit minute.
%m	The two-digit month.
%p	AM or PM.
{PolicyName}	The name of the replication policy that the snapshot was created for. This variable is valid only if you are specifying a snapshot naming pattern for a replication policy.
%R	The time. This variable is equivalent to specifying %H: %M.
%r	The time. This variable is equivalent to specifying %I: %M: %S %p.
%S	The two-digit second.
%s	The second represented in UNIX or POSIX time.
{SrcCluster}	The name of the source cluster of the replication policy that the snapshot was created for. This variable is valid only if you are specifying a snapshot naming pattern for a replication policy.
%T	The time. This variable is equivalent to specifying %H: %M: %S
%U	The two-digit numerical week of the year. Numbers range from 00 to 53. The first day of the week is calculated as Sunday.
%u	The numerical day of the week. Numbers range from 1 to 7. The first day of the week is calculated as Monday. For example, if a snapshot is created on Sunday, %u is replaced with 7.
%V	The two-digit numerical week of the year that the snapshot was created in. Numbers range from 01 to 53. The first day of the week is calculated as Monday. If the week of January 1 is four or more days in length, then that week is counted as the first week of the year.
%v	The day that the snapshot was created. This variable is equivalent to specifying %e-%b-%Y.
%W	The two-digit numerical week of the year that the snapshot was created in. Numbers range from 00 to 53. The first day of the week is calculated as Monday.
%w	The numerical day of the week that the snapshot was created on. Numbers range from 0 to 6. The first day of the week is calculated as Sunday. For example, if the snapshot was created on Sunday, %w is replaced with 0.

Variable	Description
%X	The time that the snapshot was created. This variable is equivalent to specifying %H: %M: %S.
%Y	The year that the snapshot was created in.
%y	The last two digits of the year that the snapshot was created in. For example, if the snapshot was created in 2014, %y is replaced with 14.
%Z	The time zone that the snapshot was created in.
%z	The offset from coordinated universal time (UTC) of the time zone that the snapshot was created in. If preceded by a plus sign, the time zone is east of UTC. If preceded by a minus sign, the time zone is west of UTC.
%+	The time and date that the snapshot was created. This variable is equivalent to specifying %a %b %e %X %Z %Y.
%%	Escapes a percent sign. "100%%" is replaced with 100%.

Managing snapshots

You can delete and view snapshots. You can also modify the name, duration period, and alias of an existing snapshot. However, you cannot modify the data contained in a snapshot; the data contained in a snapshot is read-only.

Reducing snapshot disk-space usage

If multiple snapshots contain the same directories, deleting one of the snapshots might not free the entire amount of space that the system reports as the size of the snapshot. The size of a snapshot is the maximum amount of data that might be freed if the snapshot is deleted.

Deleting a snapshot frees only the space that is taken up exclusively by that snapshot. If two snapshots reference the same stored data, that data is not freed until both snapshots are deleted. Remember that snapshots store data contained in all subdirectories of the root directory; if `snapshot_one` contains `/ifs/data/`, and `snapshot_two` contains `/ifs/data/dir`, the two snapshots most likely share data.

If you delete a directory, and then re-create it, a snapshot containing the directory stores the entire re-created directory, even if the files in that directory are never modified.

Deleting multiple snapshots that contain the same directories is more likely to free data than deleting multiple snapshots that contain different directories.

If multiple snapshots contain the same directories, deleting older snapshots is more likely to free disk-space than deleting newer snapshots.

Snapshots that are assigned expiration dates are automatically marked for deletion by the snapshot daemon. If the daemon is disabled, snapshots will not be automatically deleted by the system. It is recommended that you do not disable the snapshot daemon.

Delete a snapshot

You can delete a snapshot if you no longer want to access the data contained in the snapshot.

OneFS frees disk space occupied by deleted snapshots when the SnapshotDelete job is run. Also, if you delete a snapshot that contains clones or cloned files, data in a shadow store might no longer be referenced by files on the cluster; OneFS deletes unreferenced data in a shadow store when the ShadowStoreDelete job is run. OneFS routinely runs both the shadow store delete and SnapshotDelete jobs. However, you can also manually run the jobs at any time.

Procedure

1. Delete a snapshot by running the `isi snapshot snapshots delete` command.

The following command deletes a snapshot named OldSnapshot:

```
isi snapshot snapshots delete OldSnapshot
```

2. (Optional) To increase the speed at which deleted snapshot data is freed on the cluster, start the SnapshotDelete job by running the following command:

```
isi job jobs start snapshotdelete
```

3. To increase the speed at which deleted data shared between deduplicated and cloned files is freed on the cluster, start the ShadowStoreDelete job by running the following command:

```
isi job jobs start shadowstoredelete
```

Modify snapshot attributes

You can modify the name and expiration date of a snapshot.

Procedure

1. Run the `isi snapshot snapshots modify` command.

The following command causes HourlyBackup_07-15-2014_22:00 to expire on 1:30 PM on July 25th, 2014:

```
isi snapshot snapshots modify HourlyBackup_07-15-2014_22:00 \
--expires 2014-07-25T01:30
```

Modify a snapshot alias

You can modify the alias of a snapshot to assign an alternative name for the snapshot.

Procedure

1. Run the `isi snapshot snapshots modify` command.

The following command assigns an alias of LastKnownGood to HourlyBackup_07-15-2013_22:00:

```
isi snapshot snapshots modify HourlyBackup_07-15-2013_22:00 \
--alias LastKnownGood
```

View snapshots

You can view a list of snapshots or detailed information about a specific snapshot.

Procedure

1. View all snapshots by running the following command:

```
isi snapshot snapshots list
```

The system displays output similar to the following example:

ID	Name	Path
2	SIQ-c68839394a547b3fbc5c4c4b4c5673f9-latest	/ifs/data/source
6	SIQ-c68839394a547b3fbc5c4c4b4c5673f9-restore	/ifs/data/target
8	SIQ-Failover-newPol-2013-07-11_18-47-08	/ifs/data/target
12	HourlyBackup_07-15-2013_21:00	/ifs/data/media
14	HourlyBackup_07-15-2013_22:00	/ifs/data/media
16	EveryOtherHourBackup_07-15-2013_22:00	/ifs/data/media
18	HourlyBackup_07-15-2013_23:00	/ifs/data/media
20	HourlyBackup_07-16-2013_15:00	/ifs/data/media
22	EveryOtherHourBackup_07-16-2013_14:00	/ifs/data/media

2. (Optional) To view detailed information about a specific snapshot, run the `isi snapshot snapshots view` command.

The following command displays detailed information about `HourlyBackup_07-15-2013_22:00`:

```
isi snapshot snapshots view HourlyBackup_07-15-2013_22:00
```

The system displays output similar to the following example:

```
ID: 14
Name: HourlyBackup_07-15-2013_22:00
Path: /ifs/data/media
Has Locks: No
Schedule: hourly
Alias: -
Created: 2013-07-15T22:00:10
Expires: 2013-08-14T22:00:00
Size: 0b
Shadow Bytes: 0b
% Reserve: 0.00%
% Filesystem: 0.00%
State: active
```

Snapshot information

You can view information about snapshots through the output of the `isi snapshot snapshots list` command.

ID

The ID of the snapshot.

Name

The name of the snapshot.

Path

The path of the directory contained in the snapshot.

Restoring snapshot data

You can restore snapshot data through various methods. You can revert a snapshot or access snapshot data through the snapshots directory.

From the snapshots directory, you can either clone a file or copy a directory or a file. The snapshots directory can be accessed through Windows Explorer or a UNIX command line. You can disable and enable access to the snapshots directory for any of these methods through snapshots settings.

Revert a snapshot

You can revert a directory back to the state it was in when a snapshot was taken.

Before you begin

- Create a SnapRevert domain for the directory.
- Create a snapshot of a directory.

Procedure

1. (Optional) To identify the ID of the snapshot you want to revert, run the `isi snapshot snapshots view` command.

The following command displays the ID of HourlyBackup_07-15-2014_23:00:

```
isi snapshot snapshots view HourlyBackup_07-15-2014_23:00
```

The system displays output similar to the following example:

```
ID: 18
Name: HourlyBackup_07-15-2014_23:00
Path: /ifs/data/media
Has Locks: No
Schedule: hourly
Alias: -
Created: 2014-07-15T23:00:05
Expires: 2014-08-14T23:00:00
Size: 0b
Shadow Bytes: 0b
% Reserve: 0.00%
% Filesystem: 0.00%
State: active
```

2. Revert a snapshot by running the `isi job jobs start` command.

The following command reverts HourlyBackup_07-15-2014_23:00:

```
isi job jobs start snaprevert --snapid 18
```

Restore a file or directory using Windows Explorer

If the Microsoft Shadow Copy Client is installed on your computer, you can use it to restore files and directories that are stored in snapshots.

This method of restoring files and directories does not preserve the original permissions. Instead, this method assigns the file or directory the same permissions as the directory you are copying that file or directory into. To preserve permissions while restoring data from a snapshot, run the `cp` command with the `-a` option on a UNIX command line.

Note

You can access up to 64 snapshots of a directory through Windows explorer, starting with the most recent snapshot. To access more than 64 snapshots for a directory, access the cluster through a UNIX command line.

Procedure

1. In Windows Explorer, navigate to the directory that you want to restore or the directory that contains the file that you want to restore.

If the directory has been deleted, you must recreate the directory.

2. Right-click the folder, and then click **Properties**.
3. In the **Properties** window, click the **Previous Versions** tab.
4. Select the version of the folder that you want to restore or the version of the folder that contains the version of the file that you want to restore.
5. Restore the version of the file or directory.
 - To restore all files in the selected directory, click **Restore**.
 - To copy the selected directory to another location, click **Copy** and then specify a location to copy the directory to.
 - To restore a specific file, click **Open**, and then copy the file into the original directory, replacing the existing copy with the snapshot version.

Restore a file or directory through a UNIX command line

You can restore a file or directory through a UNIX command line.

Procedure

1. Open a connection to the cluster through a UNIX command line.
2. (Optional) To view the contents of the snapshot you want to restore a file or directory from, run the `ls` command for a directory contained in the snapshots root directory.

For example, the following command displays the contents of the `/archive` directory contained in Snapshot2012Jun04:

```
ls /ifs/.snapshot/Snapshot2014Jun04/archive
```

3. Copy the file or directory by using the `cp` command.

For example, the following command creates a copy of file1:

```
cp -a /ifs/.snapshot/Snapshot2014Jun04/archive/file1 \  
/ifs/archive/file1_copy
```

Clone a file from a snapshot

You can clone a file from a snapshot. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. To view the contents of the snapshot you want to restore a file or directory from, run the `ls` command for a subdirectory of the snapshots root directory.

For example, the following command displays the contents of the `/archive` directory contained in `Snapshot2014Jun04`:

```
ls /ifs/.snapshot/Snapshot2014Jun04/archive
```

3. Clone a file from the snapshot by running the `cp` command with the `-c` option.

For example, the following command clones `test.txt` from `Snapshot2014Jun04`:

```
cp -c /ifs/.snapshot/Snapshot2014Jun04/archive/test.txt \
/ifs/archive/test_clone.text
```

Managing snapshot schedules

You can modify, delete, and view snapshot schedules.

Modify a snapshot schedule

You can modify a snapshot schedule. Any changes to a snapshot schedule are applied only to snapshots generated after the modifications are made. Existing snapshots are not affected by schedule modifications.

If you modify the alias of a snapshot schedule, the alias is assigned to the next snapshot generated based on the schedule. However, if you do this, the old alias is not removed from the last snapshot that it was assigned to. Unless you manually remove the old alias, the alias will remain attached to the last snapshot that it was assigned to.

Procedure

1. Run the `isi snapshot schedules modify` command.

The following command causes snapshots created according to the snapshot schedule `hourly_media_snap` to be deleted 15 days after they are created:

```
isi snapshot schedules modify hourly_media_snap --duration 15D
```

Delete a snapshot schedule

You can delete a snapshot schedule. Deleting a snapshot schedule will not delete snapshots that were previously generated according to the schedule.

Procedure

1. Run the `isi snapshot schedules delete` command.

The following command deletes a snapshot schedule named `hourly_media_snap`:

```
isi snapshot schedules delete hourly_media_snap
```

View snapshot schedules

You can view snapshot schedules.

Procedure

1. View snapshot schedules by running the following command:

```
isi snapshot snapshots list
```

The system displays output similar to the following example:

```
ID      Name
-----
```

```

1    every-other-hour
2    daily
3    weekly
4    monthly
-----

```

2. (Optional) View detailed information about a specific snapshot schedule by running the `isi snapshot schedules view` command.

The following command displays detailed information about the snapshot schedule `every-other-hour`:

```
isi snapshot schedules view every-other-hour
```

The system displays output similar to the following example:

```

      ID: 1
      Name: every-other-hour
      Path: /ifs/data/media
      Pattern: EveryOtherHourBackup_%m-%d-%Y_%H:%M
      Schedule: Every day every 2 hours
      Duration: 1D
      Alias: -
      Next Run: 2013-07-16T18:00:00
      Next Snapshot: EveryOtherHourBackup_07-16-2013_18:00

```

Managing snapshot aliases

You can configure snapshot schedules to assign a snapshot alias to the most recent snapshot created by a snapshot schedule. You can also manually assign snapshot aliases to specific snapshots or the live version of the file system.

Configure a snapshot alias for a snapshot schedule

You can configure a snapshot schedule to assign a snapshot alias to the most recent snapshot created by the schedule.

If you configure an alias for a snapshot schedule, the alias is assigned to the next snapshot generated based on the schedule. However, if you do this, the old alias is not removed from the last snapshot that it was assigned to. Unless you manually remove the old alias, the alias will remain attached to the last snapshot that it was assigned to.

Procedure

1. Run the `isi snapshot schedules modify` command.

The following command configures the alias `LatestWeekly` for the snapshot schedule `WeeklySnapshot`:

```
isi snapshot schedules modify WeeklySnapshot --alias LatestWeekly
```

Assign a snapshot alias to a snapshot

You can assign a snapshot alias to a snapshot.

Procedure

1. Run the `isi snapshot aliases create` command.

The following command creates a snapshot alias for `Weekly-01-30-2015`:

```
isi snapshot aliases create latestWeekly Weekly-01-30-2015
```

Reassign a snapshot alias to the live file system

You can reassign a snapshot alias to redirect clients from a snapshot to the live file system. This procedure is available only through the command-line interface (CLI).

Procedure

1. Run the `isi snapshot aliases modify` command.

The following command reassigns the `latestWeekly` alias to the live file system:

```
isi snapshot aliases modify latestWeekly --target LIVE
```

View snapshot aliases

You can view a list of all snapshot aliases. This procedure is available only through the command-line interface (CLI).

Procedure

1. View a list of all snapshot aliases by running the following command:

```
isi snapshot aliases list
```

If a snapshot alias references the live version of the file system, the `Target ID` is `-1`.

2. (Optional) View information about a specific snapshot by running the `isi snapshot aliases view` command.

The following command displays information about `latestWeekly`:

```
isi snapshot aliases view latestWeekly
```

Snapshot alias information

You can view information about snapshot aliases through the output of the `isi snapshot aliases view` command.

ID

The numerical ID of the snapshot alias.

Name

The name of the snapshot alias.

Target ID

The numerical ID of the snapshot that is referenced by the alias.

Target Name

The name of the snapshot that is referenced by the alias.

Created

The date that the snapshot alias was created.

Managing with snapshot locks

You can delete, create, and modify the expiration date of snapshot locks.

CAUTION

It is recommended that you do not create, delete, or modify snapshots locks unless you are instructed to do so by Isilon Technical Support.

Deleting a snapshot lock that was created by OneFS might result in data loss. If you delete a snapshot lock that was created by OneFS, it is possible that the corresponding snapshot might be deleted while it is still in use by OneFS. If OneFS cannot access a snapshot that is necessary for an operation, the operation will malfunction and data loss might result. Modifying the expiration date of a snapshot lock created by OneFS can also result in data loss because the corresponding snapshot can be deleted prematurely.

Create a snapshot lock

You can create snapshot locks that prevent snapshots from being deleted. This procedure is available only through the command-line interface (CLI).

Although you can prevent a snapshot from being automatically deleted by creating a snapshot lock, it is recommended that you do not create snapshot locks. To prevent a snapshot from being automatically deleted, it is recommended that you extend the duration period of the snapshot.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Create a snapshot lock by running the `isi snapshot locks create` command.

For example, the following command applies a snapshot lock to "SnapshotApril2012", sets the lock to expire in one month, and adds a description of "Maintenance Lock":

```
isi snapshot locks create SnapshotApril2012 --expires 1M \
--comment "Maintenance Lock"
```

Modify a snapshot lock expiration date

You can modify the expiration date of a snapshot lock. This procedure is available only through the command-line interface (CLI).

CAUTION

It is recommended that you do not modify the expiration dates of snapshot locks.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi snapshot locks modify` command.

The following command sets an expiration date two days from the present date for a snapshot lock with an ID of 1 that is applied to a snapshot named SnapshotApril2014:

```
isi snapshot locks modify SnapshotApril2014 1 --expires 2D
```

Delete a snapshot lock

You can delete a snapshot lock. This procedure is available only through the command-line interface (CLI).

⚠ CAUTION

It is recommended that you do not delete snapshot locks.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Delete a snapshot lock by running the `isi snapshot locks delete` command.

For example, the following command deletes a snapshot lock that is applied to SnapshotApril2014 and has a lock ID of 1:

```
isi snapshot locks delete Snapshot2014Apr16 1
```

The system prompts you to confirm that you want to delete the snapshot lock.

3. Type **yes** and then press ENTER.

Snapshot lock information

You can view snapshot lock information through the `isi snapshot locks view` and `isi snapshot locks list` commands.

ID

Numerical identification number of the snapshot lock.

Comment

Description of the snapshot lock. This can be any string specified by a user.

Expires

The date that the snapshot lock will be automatically deleted by OneFS.

Count

The number of times the snapshot lock is held.

The file clone operation can hold a single snapshot lock multiple times. If multiple file clones are created simultaneously, the file clone operation holds the same lock multiple times, rather than creating multiple locks. If you delete a snapshot lock that is held more than once, you will delete only one of the instances that the lock is held. In order to delete a snapshot lock that is held multiple times, you must delete the snapshot lock the same number of times as displayed in the count field.

Configure SnapshotIQ settings

You can configure SnapshotIQ settings that determine how snapshots can be created and the methods that users can access snapshot data.

Procedure

1. (Optional) View current SnapshotIQ settings by running the following command:

```
isi snapshot settings view
```

The system displays output similar to the following example:

```

Service: Yes
Autocreate: Yes
Autodelete: Yes
Reserve: 0.00%
Global Visible Accessible: Yes
NFS Root Accessible: Yes
NFS Root Visible: Yes
NFS Subdir Accessible: Yes
SMB Root Accessible: Yes
SMB Root Visible: Yes
SMB Subdir Accessible: Yes
Local Root Accessible: Yes
Local Root Visible: Yes
Local Subdir Accessible: Yes

```

2. Configure SnapshotIQ settings by running the `isi snapshot settings modify` command:

The following command prevents snapshots from being created on the cluster:

```
isi snapshot settings modify --service disable
```

SnapshotIQ settings

SnapshotIQ settings determine how snapshots behave and can be accessed.

The following settings are displayed in the output of the `isi snapshot settings view` command:

Service

Determines whether SnapshotIQ is enabled on the cluster.

Autocreate

Determines whether snapshots are automatically generated according to snapshot schedules.

Note

Disabling snapshot generation might cause some OneFS operations to fail. It is recommended that you do not disable this setting.

Autodelete

Determines whether snapshots are automatically deleted according to their expiration dates.

Reserve

Specifies the percentage of disk space on the cluster that is reserved for snapshots.

NFS Root Accessible

Determines whether snapshot directories are accessible through NFS.

NFS Root Visible

Determines whether snapshot directories are visible through NFS.

NFS Subdir Accessible

Determines whether snapshot subdirectories are accessible through NFS.

SMB Root Accessible

Determines whether snapshot directories are accessible through SMB.

SMB Root Visible

Determines whether snapshot directories are visible through SMB.

SMB Subdir Accessible

Determines whether snapshot subdirectories are accessible through SMB.

Local Root Accessible

Determines whether snapshot directories are accessible through an SSH connection or the local console.

Local Root Visible

Determines whether snapshot directories are visible through the an SSH connection or the local console.

Local Subdir Accessible

Determines whether snapshot subdirectories are accessible through an SSH connection or the local console.

Set the snapshot reserve

You can specify a minimum percentage of cluster-storage capacity that you want to reserve for snapshots. This procedure is available only through the command-line interface (CLI).

The snapshot reserve does not limit the amount of space that snapshots are allowed to consume on the cluster. Snapshots can consume more than the percentage of capacity specified by the snapshot reserve. It is recommended that you do not specify a snapshot reserve.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Set the snapshot reserve by running the `isi snapshot settings modify` command with the `--reserve` option.

For example, the following command sets the snapshot reserve to 20%:

```
isi snapshot settings modify --reserve 20
```

Snapshot commands

You can control and access snapshots through the snapshot commands. Most snapshot commands apply specifically to the SnapshotIQ tool and are available only if a SnapshotIQ license is configured on the cluster.

Snapshot naming patterns

If you schedule snapshots to be automatically generated, either according to a snapshot schedule or a replication policy, you must assign a snapshot naming pattern that determines how the snapshots are named. Snapshot naming patterns contain variables that include information about how and when the snapshot was created.

The following variables can be included in a snapshot naming pattern:

Variable	Description
%A	The day of the week.
%a	The abbreviated day of the week. For example, if the snapshot is generated on a Sunday, %a is replaced with Sun.

Variable	Description
%B	The name of the month.
%b	The abbreviated name of the month. For example, if the snapshot is generated in September, %b is replaced with Sep.
%C	The first two digits of the year. For example, if the snapshot is created in 2014, %C is replaced with 20.
%c	The time and day. This variable is equivalent to specifying %a %b %e %T %Y.
%d	The two digit day of the month.
%e	The day of the month. A single-digit day is preceded by a blank space.
%F	The date. This variable is equivalent to specifying %Y-%m-%d
%G	The year. This variable is equivalent to specifying %Y. However, if the snapshot is created in a week that has less than four days in the current year, the year that contains the majority of the days of the week is displayed. The first day of the week is calculated as Monday. For example, if a snapshot is created on Sunday, January 1, 2017, %G is replaced with 2016, because only one day of that week is in 2017.
%g	The abbreviated year. This variable is equivalent to specifying %y. However, if the snapshot was created in a week that has less than four days in the current year, the year that contains the majority of the days of the week is displayed. The first day of the week is calculated as Monday. For example, if a snapshot is created on Sunday, January 1, 2017, %g is replaced with 16, because only one day of that week is in 2017.
%H	The hour. The hour is represented on the 24-hour clock. Single-digit hours are preceded by a zero. For example, if a snapshot is created at 1:45 AM, %H is replaced with 01.
%h	The abbreviated name of the month. This variable is equivalent to specifying %b.
%I	The hour represented on the 12-hour clock. Single-digit hours are preceded by a zero. For example, if a snapshot is created at 1:45 PM, %I is replaced with 01.
%j	The numeric day of the year. For example, if a snapshot is created on February 1, %j is replaced with 32.
%k	The hour represented on the 24-hour clock. Single-digit hours are preceded by a blank space.
%l	The hour represented on the 12-hour clock. Single-digit hours are preceded by a blank space. For example, if a snapshot is created at 1:45 AM, %I is replaced with 1.
%M	The two-digit minute.
%m	The two-digit month.

Variable	Description
%p	AM or PM.
{PolicyName}	The name of the replication policy that the snapshot was created for. This variable is valid only if you are specifying a snapshot naming pattern for a replication policy.
%R	The time. This variable is equivalent to specifying %H : %M.
%r	The time. This variable is equivalent to specifying %I : %M : %S %p.
%S	The two-digit second.
%s	The second represented in UNIX or POSIX time.
{SrcCluster}	The name of the source cluster of the replication policy that the snapshot was created for. This variable is valid only if you are specifying a snapshot naming pattern for a replication policy.
%T	The time. This variable is equivalent to specifying %H : %M : %S
%U	The two-digit numerical week of the year. Numbers range from 00 to 53. The first day of the week is calculated as Sunday.
%u	The numerical day of the week. Numbers range from 1 to 7. The first day of the week is calculated as Monday. For example, if a snapshot is created on Sunday, %u is replaced with 7.
%V	The two-digit numerical week of the year that the snapshot was created in. Numbers range from 01 to 53. The first day of the week is calculated as Monday. If the week of January 1 is four or more days in length, then that week is counted as the first week of the year.
%v	The day that the snapshot was created. This variable is equivalent to specifying %e-%b-%Y.
%W	The two-digit numerical week of the year that the snapshot was created in. Numbers range from 00 to 53. The first day of the week is calculated as Monday.
%w	The numerical day of the week that the snapshot was created on. Numbers range from 0 to 6. The first day of the week is calculated as Sunday. For example, if the snapshot was created on Sunday, %w is replaced with 0.
%X	The time that the snapshot was created. This variable is equivalent to specifying %H : %M : %S.
%Y	The year that the snapshot was created in.
%y	The last two digits of the year that the snapshot was created in. For example, if the snapshot was created in 2014, %y is replaced with 14.
%Z	The time zone that the snapshot was created in.
%z	The offset from coordinated universal time (UTC) of the time zone that the snapshot was created in. If preceded by a plus sign, the time zone is east of UTC. If preceded by a minus sign, the time zone is west of UTC.

Variable	Description
%+	The time and date that the snapshot was created. This variable is equivalent to specifying %a %b %e %X %Z %Y.
%%	Escapes a percent sign. "100%" is replaced with 100%.

isi snapshot schedules create

Creates a snapshot schedule. A snapshot schedule determines when OneFS regularly generates snapshots on a recurring basis.

Syntax

```
isi snapshot schedules create <name> <path> <pattern> <schedule>
  [--alias <alias>]
  [--duration <duration>]
  [--verbose]
```

Options

<name>

Specifies a name for the snapshot schedule.

<path>

Specifies the path of the directory to include in the snapshots.

<pattern>

Specifies a naming pattern for snapshots created according to the schedule.

<schedule>

Specifies how often snapshots are created.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`--alias <alias>`

Specifies an alias for the latest snapshot generated based on the schedule. The alias enables you to quickly locate the most recent snapshot that was generated according to the schedule.

Specify as any string.

`{--duration | -x} <duration>`

Specifies how long snapshots generated according to the schedule are stored on the cluster before OneFS automatically deletes them.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

`{--verbose | -v}`

Displays a message confirming that the snapshot schedule was created.

isi snapshot schedules modify

Modifies the attributes of an existing snapshot schedule.

If you modify a snapshot schedule, snapshots that have already been generated based on the schedule are not affected by the changes.

Syntax

```
isi snapshot schedules modify <schedule-name>
  {--name <name> | --alias <name> | --path <path>
  | --pattern <naming-pattern> | --schedule <schedule>
  | --duration <duration> | --clear-duration}...
  [--verbose]
```

Options

`<schedule-name>`

Modifies the specified snapshot schedule.

Specify as a snapshot schedule name or ID.

`--name <name>`

Specifies a new name for the schedule.
Specify as any string.

`{--alias | -a} <name>`

Specifies an alias for the latest snapshot generated based on the schedule. The alias enables you to quickly locate the most recent snapshot that was generated according to the schedule. If specified, the specified alias will be applied to the next snapshot generated by the schedule, and all subsequently generated snapshots.
Specify as any string.

`--path <path>`

Specifies a new directory path for this snapshot schedule. If specified, snapshots generated by the schedule will contain only this directory path.
Specify as a directory path.

`--pattern <naming-pattern>`

Specifies a pattern by which snapshots created according to the schedule are named.

`--schedule <schedule>`

Specifies how often snapshots are created.
Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`
- `Every [{other | <integer>}] week [on <day>]`
- `Every [{other | <integer>}] month [on the <integer>]`
- `Every [<day>[, ...] [of every [{other | <integer>}] week]]`
- `The last {day | weekday | <day>} of every [{other | <integer>}] month`
- `The <integer> {weekday | <day>} of every [{other | <integer>}] month`
- `Yearly on <month> <integer>`
- `Yearly on the {last | <integer>} [weekday | <day>] of <month>`

Specify *<frequency>* in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`{--duration | -x} <duration>`

Specifies how long snapshots generated according to the schedule are stored on the cluster before OneFS automatically deletes them.
Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

- Y**
Specifies years
- M**
Specifies months
- W**
Specifies weeks
- D**
Specifies days
- H**
Specifies hours

`--clear-duration`

Removes the duration period for snapshots created according to the schedule. If specified, generated snapshots will exist on the cluster indefinitely.

`{--verbose | -v}`

Displays a message confirming that the snapshot schedule was modified.

isi snapshot schedules delete

Deletes a snapshot schedule. Once a snapshot schedule is deleted, snapshots will no longer be generated according to the schedule. However, snapshots previously generated according to the schedule are not affected.

Syntax

```
isi snapshot schedules delete <schedule-name>
  [--force]
  [--verbose]
```

Options

`<schedule-name>`

Deletes the specified snapshot schedule.
Specify as a snapshot schedule name or ID.

`{--force | -f}`

Does not prompt you to confirm that you want to delete this snapshot schedule.

`{--verbose | -v}`

Displays a message confirming that the snapshot schedule was deleted.

isi snapshot schedules list

Displays a list of all snapshot schedules.

Syntax

```
isi snapshot schedules list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
```

```
[--no-header]
[--no-footer]
[--verbose]
```

Options

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--sort <attribute>
```

Sorts output displayed by the specified attribute.

The following values are valid:

id

Sorts output by the ID of a snapshot schedule.

name

Sorts output alphabetically by the name of a snapshot schedule.

path

Sorts output by the absolute path of the directory contained by snapshots created according to a schedule.

pattern

Sorts output alphabetically by the snapshot naming pattern assigned to snapshots generated according to a schedule.

schedule

Sorts output alphabetically by the schedule. For example, "Every week" precedes "Yearly on January 3rd"

duration

Sorts output by the length of time that snapshots created according to the schedule endure on the cluster before being automatically deleted.

alias

Sorts output alphabetically by the name of the alias assigned to the most recent snapshot generated according to the schedule.

next_run

Sorts output by the next time that a snapshot will be created according to the schedule.

next_snapshot

Sorts output alphabetically by the name of the snapshot that is scheduled to be created next.

```
{--descending | -d}
```

Displays output in reverse order.

```
--format <output-format>
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi snapshot schedules view

Displays information about a snapshot schedule.

Syntax

```
isi snapshot schedules view <schedule-name>
```

Options

<schedule-name>

Displays information about the specified snapshot schedule.
Specify as a snapshot schedule name or ID.

isi snapshot schedules pending list

Displays a list of snapshots that are scheduled to be generated by snapshot schedules.

Syntax

```
isi snapshot schedules pending list
[--begin <timestamp>]
[--end <timestamp>]
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--begin | -b} <timestamp>

Displays only snapshots that are scheduled to be generated after the specified date.
Specify <timestamp> in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

If this option is not specified, the output displays a list of snapshots that are scheduled to be generated after the current time.

{--end | -e} <time>

Displays only snapshots that are scheduled to be generated before the specified date.

Specify <time> in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

If this option is not specified, the output displays a list of snapshots that are scheduled to be generated before 30 days after the begin time.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format <output-format>

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.


```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi snapshot snapshots create

Creates a snapshot of a directory.

Syntax

```
isi snapshot snapshots create <path>
  [--name <name>]
  [--expires {<timestamp> | <duration>}]
  [--alias <name>]
  [--verbose]
```

Options

<path>

Specifies the path of the directory to include in this snapshot.

`--name <name>`

Specifies a name for the snapshot.

`{--expires | -x} {<timestamp> | <duration>}`

Specifies when OneFS will automatically delete this snapshot. If this option is not specified, the snapshot will exist indefinitely.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify *<duration>* in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

`{--alias | -a} <name>`

Specifies an alias for this snapshot. A snapshot alias is an alternate name for a snapshot.

Specify as any string.

`{--verbose | -v}`

Displays a message confirming that the snapshot was created.

isi snapshot snapshots modify

Modifies attributes of a snapshot or snapshot alias.

Syntax

```
isi snapshot snapshots modify <snapshot>
  {--name <name> | --expires {<timestamp> | <duration>}
  | --clear-expires | --alias <name>}...
  [--verbose]
```

Options

<snapshot>

Modifies the specified snapshot or snapshot alias.

Specify as the name or ID of a snapshot or snapshot alias.

--name <name>

Specifies a new name for the snapshot or snapshot alias.

Specify as any string.

{--expires | -x} {<timestamp> | <duration>}

Specifies when OneFS will automatically delete this snapshot.

Specify <timestamp> in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify <duration> in the following format:

```
<integer><time>
```

The following <time> values are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

You cannot modify the expiration date of a snapshot alias.

--clear-expires

Removes the expiration date from the snapshot, allowing the snapshot to exist on the cluster indefinitely.

You cannot modify the expiration date of a snapshot alias.

{--alias | -a} <name>

Specifies an alias for the snapshot. A snapshot alias is an alternate name for a snapshot. You cannot specify an alias for a snapshot alias.

Specify as any string.

{--verbose | -v}

Displays a message confirming that the snapshot or snapshot alias was modified.

isi snapshot snapshots delete

Deletes a snapshot. If a snapshot is deleted, it can no longer be accessed by a user or the system.

Syntax

```
isi snapshot snapshots delete [--all | --snapshot <snapshot>
| --schedule <schedule> | --type <type>}
[--force]
[--verbose]
```

Options

`--all`

Deletes all snapshots.

`--snapshot <snapshot>`

Deletes the specified snapshot.
Specify as a snapshot name or ID.

`--schedule <schedule>`

Deletes all snapshots created according to the specified schedule.
Specify as a snapshot schedule name or ID.

`--type <type>`

Deletes all snapshots of the specified type.
The following types are valid:

alias

Deletes all snapshot aliases.

real

Deletes all snapshots.

`{--force | -f}`

Does not prompt you to confirm that you want to delete the snapshot.

`{--verbose | -v}`

Displays a message confirming that the snapshot was deleted.

Examples

The following command deletes newSnap1:

```
isi snapshot snapshots delete --snapshot newSnap1
```

isi snapshot snapshots list

Displays a list of all snapshots and snapshot aliases.

Syntax

```
isi snapshot snapshots list
[--state <state>]
[--limit <integer>]
[--sort <attribute>]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
```

```
[--no-footer]
[--verbose]
```

Options

`--state <state>`

Displays only snapshots and snapshot aliases that exist in the specified state. The following states are valid:

all

Displays all snapshots and snapshot aliases that are currently occupying space on the cluster.

active

Displays only snapshots and snapshot aliases that have not been deleted.

deleting

Displays only snapshots that have been deleted but are still occupying space on the cluster. The space occupied by deleted snapshots will be freed the next time the snapshot delete job is run.

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts command output by the specified attribute. The following attributes are valid:

id

Sorts output by the ID of a snapshot.

name

Sorts output alphabetically by the name of a snapshot.

path

Sorts output by the absolute path of the directory contained in a snapshot.

has_locks

Sorts output by whether any snapshot locks have been applied to a snapshot.

schedule

If a snapshot was generated according to a schedule, sorts output alphabetically by the name of the snapshot schedule.

target_id

If a snapshot is an alias, sorts output by the snapshot ID of the target snapshot instead of the snapshot ID of the alias.

target_name

If a snapshot is an alias, sorts output by the name of the target snapshot instead of the name of the alias.

created

Sorts output by the time that a snapshot was created.

expires

Sorts output by the time at which a snapshot is scheduled to be automatically deleted.

size

Sorts output by the amount of disk space taken up by a snapshot.

shadow_bytes

Sorts output based on the amount of data that a snapshot references from shadow stores. Snapshots reference shadow store data if a file contained in a snapshot is cloned or a snapshot is taken of a cloned file.

pct_reserve

Sorts output by the percentage of the snapshot reserve that a snapshot occupies.

pct_filesystem

Sorts output by the percent of the file system that a snapshot occupies.

state

Sorts output based on the state of snapshots.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

isi snapshot snapshots view

Displays the properties of an individual snapshot.

Syntax

```
isi snapshot snapshots view <snapshot>
```

Options

<snapshot>

Displays information about the specified snapshot.
Specify as a snapshot name or ID.

isi snapshot settings modify

Modifies snapshot settings.

Syntax

```
isi snapshot settings modify
  --service {enable | disable}
  | --autocreate {enable | disable}
  | --autodelete {enable | disable}
  | --reserve <integer>
  | --global-visible-accessible {yes | no}
  | --nfs-root-accessible {yes | no}
  | --nfs-root-visible {yes | no}
```

```
| --nfs-subdir-accessible {yes | no}
| --smb-root-accessible {yes | no}
| --smb-root-visible {yes | no}
| --smb-subdir-accessible {yes | no}
| --local-root-accessible {yes | no}
| --local-root-visible {yes | no}
| --local-subdir-accessible {yes | no}}...
[--verbose]
```

Options

`--service {enable | disable}`

Determines whether snapshots can be generated.

Note

Disabling snapshot generation might cause some OneFS operations to fail. It is recommended that you do not disable this setting.

`--autocreate {enable | disable}`

Determines whether snapshots are automatically generated according to snapshot schedules.

Specifying `disable` does not prevent OneFS applications from generating snapshots.

`--autodelete {enable | disable}`

Determines whether snapshots are automatically deleted according to their expiration dates.

All snapshots that pass their expiration date while this option is disabled will immediately be deleted when the option is enabled again.

`--reserve <integer>`

Specifies the percentage of the file system to reserve for snapshot usage. Specify as a positive integer between 1 and 100.

Note

This option limits only the amount of space available to applications other than SnapshotIQ. It does not limit the amount of space that snapshots are allowed to occupy. Snapshots can occupy more than the specified percentage of system storage space.

`--global-visible-accessible {yes | no}`

Specifying `yes` causes snapshot directories and sub-directories to be visible and accessible through all protocols, overriding all other snapshot visibility and accessibility settings. Specifying `no` causes visibility and accessibility settings to be controlled through the other snapshot visibility and accessibility settings.

`--nfs-root-accessible {yes | no}`

Determines whether snapshot directories are accessible through NFS.

`--nfs-root-visible {yes | no}`

Determines whether snapshot directories are visible through NFS.

`--nfs-subdir-accessible {yes | no}`

Determines whether snapshot subdirectories are accessible through NFS.

`--smb-root-accessible {yes | no}`

Determines whether snapshot directories are accessible through SMB.

```
--smb-root-visible {yes | no}
```

Determines whether snapshot directories are visible through SMB.

```
--smb-subdir-accessible {yes | no}
```

Determines whether snapshot subdirectories are accessible through SMB.

```
--local-root-accessible {yes | no}
```

Determines whether snapshot directories are accessible through the local file system.

```
--local-root-visible {yes | no}
```

Determines whether snapshot directories are visible through the local file system.

```
--local-subdir-accessible {yes | no}
```

Determines whether snapshot subdirectories are accessible through the local file system.

```
{--verbose | -v}
```

Displays a message displaying which snapshot settings were modified.

isi snapshot settings view

Displays current SnapshotIQ settings.

Syntax

```
isi snapshot settings view
```

Options

There are no options for this command.

isi snapshot locks create

Creates a snapshot lock.

Note

It is recommended that you do not create snapshot locks and do not use this command. If the maximum number of locks on a snapshot is reached, some applications, such as SyncIQ, might not function properly.

Syntax

```
isi snapshot locks create <snapshot>
  [--comment <string>]
  [--expires {<timestamp> | <duration>}]
  [--verbose]
```

Options

<snapshot>

Specifies the name of the snapshot to apply this lock to.

```
{--comment | -c} <string>
```

Specifies a comment to describe the lock.

Specify as any string.

```
{--expires | -x} {<timestamp> | <duration>}
```

Specifies when the lock will be automatically deleted by the system.

If this option is not specified, the snapshot lock will exist indefinitely.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify *<duration>* in the following format:

```
<integer><time>
```

The following *<time>* values are valid:

- Y**
Specifies years
- M**
Specifies months
- W**
Specifies weeks
- D**
Specifies days
- H**
Specifies hours

```
{--verbose | -v}
```

Displays a message confirming that the snapshot lock was deleted.

isi snapshot locks modify

Modifies the expiration date of a snapshot lock.

CAUTION

It is recommended that you do not modify the expiration date of snapshot locks and do not run this command. Modifying the expiration date of a snapshot lock that was created by OneFS might result in data loss.

Syntax

```
isi snapshot locks modify <snapshot> <id>
  {--expires {<timestamp> | <duration>} | --clear-expires}
  [--verbose]
```

Options

<snapshot>

Modifies a snapshot lock that has been applied to the specified snapshot. Specify as a snapshot name or ID.

<id>

Modifies the snapshot lock of the specified ID.

```
{--expires | -x} {<timestamp> | <duration>}
```

Specifies when the lock will be automatically deleted by the system. If this option is not specified, the snapshot lock will exist indefinitely.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify *<duration>* in the following format:

```
<integer><time>
```

The following *<time>* values are valid:

- Y**
Specifies years
- M**
Specifies months
- W**
Specifies weeks
- D**
Specifies days
- H**
Specifies hours

```
--clear-expires
```

Removes the duration period for the snapshot lock. If specified, the snapshot lock will exist on the cluster indefinitely.

```
{--verbose | -v}
```

Displays a message confirming that the snapshot lock was modified.

Examples

The following command causes a snapshot lock applied to Wednesday_Backup to expire in three weeks:

```
isi snapshot locks modify Wednesday_Backup 1 --expires 3W
```

isi snapshot locks delete

Deletes a snapshot lock. Deleting a snapshot lock might result in data loss.

⚠ CAUTION

It is recommended that you do not delete snapshot locks and do not run this command. Deleting a snapshot lock that was created by OneFS might result in data loss.

Syntax

```
isi snapshot locks delete <snapshot> <id>
[--force]
[--verbose]
```

Options

<snapshot>

Deletes a snapshot lock that has been applied to the specified snapshot. Specify as a snapshot name or ID.

<id>

Modifies the snapshot lock of the specified ID.

```
{--force | -f}
```

Does not prompt you to confirm that you want to delete this snapshot lock.

```
{--verbose | -v}
```

Displays a message confirming that the snapshot lock was deleted.

isi snapshot locks list

Displays a list of all locks applied to a specific snapshot.

Syntax

```
isi snapshot locks list <snapshot>
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
<snapshot>
```

Displays all locks belonging to the specified snapshot.
Specify as a snapshot name.

```
{--limit | -l} <integer>
```

Displays no more than the specified number of items.

```
--sort <attribute>
```

Sorts output displayed by the specified attribute.
The following values are valid:

id

Sorts output by the ID of a snapshot lock.

comment

Sorts output alphabetically by the description of a snapshot lock.

expires

Sorts output by the length of time that a lock endures on the cluster before being automatically deleted.

count

Sorts output by the number of times that a lock is held.

```
{--descending | -d}
```

Displays output in reverse order.

```
--format <output-format>
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi snapshot locks view

Displays information about a snapshot lock.

Syntax

```
isi snapshot locks view <snapshot> <id>
```

Options

<name>

Specifies the snapshot to view locks for.
Specify as a snapshot name or ID.

<id>

Displays the specified lock.
Specify as a snapshot lock ID.

isi snapshot aliases create

Assigns a snapshot alias to a snapshot or to the live version of the file system.

Syntax

```
isi snapshot aliases create <name> <target>
[--verbose]
```

Options

<name>

Specifies the a name for the alias.

<target>

Assigns the alias to the specified snapshot or to the live version of the file system.
Specify as a snapshot ID or name. To target the live version of the file system, specify LIVE.

{--verbose | -v}

Displays more detailed information.

isi snapshot aliases modify

Modifies a snapshot alias.

Syntax

```
isi snapshot aliases modify <alias>
{--name <name> | --target <snapshot>}
[--verbose]
```

Options

<alias>

Modifies the specified snapshot alias.
Specify as a snapshot-alias name or ID.

--name <name>

Specifies a new name for the snapshot alias.

`--target <snapshot>`

Reassigns the snapshot alias to the specified snapshot or the live version of the file system.

Specify as a snapshot ID or name. To target the live version of the file system, specify `LIVE`.

`{--verbose | -v}`

Displays more detailed information.

isi snapshot aliases delete

Deletes a snapshot alias.

Syntax

```
isi snapshot aliases delete {<alias> | --all}
  [--force]
  [--verbose]
```

Options

`<alias>`

Deletes the snapshot alias of the specified name.

Specify as a snapshot-alias name or ID.

`--all`

Deletes all snapshot aliases.

`{--force | -f}`

Runs the command without prompting you to confirm that you want to delete the snapshot alias.

`{--verbose | -v}`

Displays more detailed information.

isi snapshot aliases list

Displays a list of snapshot aliases.

Syntax

```
isi snapshot aliases list
  [--limit <integer>]
  [--sort {id | name | target_id | target_name | created}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.

The following values are valid:

id

Sorts output by the ID of the snapshot alias.

name

Sorts output by the name of the snapshot alias.

target_id

Sorts output by the ID of the snapshot that the snapshot alias is assigned to.

target_name

Sorts output by the name of the snapshot that the snapshot alias is assigned to.

created

Sorts output by the date the snapshot alias was created.

{--descending | -d}

Displays output in reverse order.

--format *<output-format>*

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi snapshot aliases view

Displays detailed information about a snapshot alias.

Syntax

```
isi snapshot aliases view <alias>
```

Options

<alias>

Displays detailed information about the specified snapshot alias. Specify as a snapshot-alias name or ID.

CHAPTER 12

Deduplication with SmartDedupe

This section contains the following topics:

- [Deduplication overview](#) 520
- [Deduplication jobs](#) 520
- [Data replication and backup with deduplication](#) 521
- [Snapshots with deduplication](#) 521
- [Deduplication considerations](#) 521
- [Shadow-store considerations](#) 522
- [SmartDedupe license functionality](#) 522
- [Managing deduplication](#) 522
- [Deduplication commands](#) 526

Deduplication overview

The SmartDedupe software module enables you to save storage space on your cluster by reducing redundant data. Deduplication maximizes the efficiency of your cluster by decreasing the amount of storage required to store multiple files with similar blocks.

SmartDedupe deduplicates data by scanning an Isilon cluster for identical data blocks. Each block is 8 KB. If SmartDedupe finds duplicate blocks, SmartDedupe moves a single copy of the blocks to a hidden file called a shadow store. SmartDedupe then deletes the duplicate blocks from the original files and replaces the blocks with pointers to the shadow store.

Deduplication is applied at the directory level, targeting all files and directories underneath one or more root directories. You can first assess a directory for deduplication and determine the estimated amount of space you can expect to save. You can then decide whether to deduplicate the directory. After you begin deduplicating a directory, you can monitor how much space is saved by deduplication in real time.

SmartDedupe does not deduplicate files that are 32 KB and smaller, because doing so would consume more cluster resources than the storage savings are worth. Each shadow store can contain up to 255 blocks. Each block in a shadow store can be referenced 32000 times.

Deduplication jobs

Deduplication is performed by maintenance jobs referred to as deduplication jobs. You can monitor and control deduplication jobs as you would any other maintenance job on the cluster. Although the overall performance impact of deduplication is minimal, the deduplication job consumes 256 MB of memory per node.

When a deduplication job is first run on a cluster, SmartDedupe samples blocks from each file and creates index entries for those blocks. If the index entries of two blocks match, SmartDedupe scans the blocks adjacent to the matching pair and then deduplicates all duplicate blocks. After a deduplication job samples a file once, new deduplication jobs will not sample the file again until the file is modified.

The first deduplication job you run might take significantly longer to complete than subsequent deduplication jobs. The first deduplication job must scan all files under the specified directories to generate the initial index. If subsequent deduplication jobs take a long time to complete, this most likely indicates that a large amount of data is being deduplicated. However, it can also indicate that clients are creating a large amount of new data on the cluster. If a deduplication job is interrupted during the deduplication process, the job will automatically restart the scanning process from where the job was interrupted.

It is recommended that you run deduplication jobs when clients are not modifying data on the cluster. If clients are continually modifying files on the cluster, the amount of space saved by deduplication is minimal because the deduplicated blocks are constantly removed from the shadow store. For most clusters, it is recommended that you start a deduplication job every ten days.

The permissions required to modify deduplication settings are not the same as those needed to run a deduplication job. Although a user must have the maintenance job permission to run a deduplication job, the user must have the deduplication permission to modify deduplication settings. By default, the deduplication job is configured to run at a low priority.

Data replication and backup with deduplication

When deduplicated files are replicated to another Isilon cluster or backed up to a tape device, the deduplicated files no longer share blocks on the target Isilon cluster or backup device. However, although you can deduplicate data on a target Isilon cluster, you cannot deduplicate data on an NDMP backup device.

Shadows stores are not transferred to target clusters or backup devices. Because of this, deduplicated files do not consume less space than non-deduplicated files when they are replicated or backed up. To avoid running out of space, you must ensure that target clusters and tape devices have enough free space to store deduplicated data as if the data had not been deduplicated. To reduce the amount of storage space consumed on a target Isilon cluster, you can configure deduplication for the target directories of your replication policies. Although this will deduplicate data on the target directory, it will not allow SyncIQ to transfer shadow stores. Deduplication is still performed by deduplication jobs running on the target cluster.

The amount of cluster resources required to backup and replicate deduplicated data is the same as for non-deduplicated data. You can deduplicate data while the data is being replicated or backed up.

Snapshots with deduplication

You cannot deduplicate the data stored in a snapshot. However, you can create snapshots of deduplicated data.

If you create a snapshot for a deduplicated directory, and then modify the contents of that directory, the references to shadow stores will be transferred to the snapshot over time. Therefore, if you enable deduplication before you create snapshots, you will save more space on your cluster. If you implement deduplication on a cluster that already has a significant amount of data stored in snapshots, it will take time before the snapshot data is affected by deduplication. Newly created snapshots can contain deduplicated data, but snapshots created before deduplication was implemented cannot.

If you plan on reverting a snapshot, it is best to revert the snapshot before running a deduplication job. Restoring a snapshot can overwrite many of the files on the cluster. Any deduplicated files are reverted back to normal files if they are overwritten by a snapshot revert. However, after the snapshot revert is complete, you can deduplicate the directory and the space savings persist on the cluster.

Deduplication considerations

Deduplication can significantly increase the efficiency at which you store data. However, the effect of deduplication varies depending on the cluster.

You can reduce redundancy on a cluster by running SmartDedupe. Deduplication creates links that can impact the speed at which you can read from and write to files. In particular, sequentially reading chunks smaller than 512 KB of a deduplicated file can be significantly slower than reading the same small, sequential chunks of a non-deduplicated file. This performance degradation applies only if you are reading non-cached data. For cached data, the performance for deduplicated files is potentially better than non-deduplicated files. If you stream chunks larger than 512 KB, deduplication does not significantly impact the read performance of the file. If you intend on streaming 8 KB or less of each file at a time, and you do not plan on concurrently streaming the files, it is recommended that you do not deduplicate the files.

Deduplication is most effective when applied to static or archived files and directories. The less files are modified, the less negative effect deduplication has on the cluster. For example, virtual machines often contain several copies of identical files that are rarely modified. Deduplicating a large number of virtual machines can greatly reduce consumed storage space.

Shadow-store considerations

Shadow stores are hidden files that are referenced by cloned and deduplicated files. Files that reference shadow stores behave differently than other files.

- Reading shadow-store references might be slower than reading data directly. Specifically, reading non-cached shadow-store references is slower than reading non-cached data. Reading cached shadow-store references takes no more time than reading cached data.
- When files that reference shadow stores are replicated to another Isilon cluster or backed up to a Network Data Management Protocol (NDMP) backup device, the shadow stores are not transferred to the target Isilon cluster or backup device. The files are transferred as if they contained the data that they reference from shadow stores. On the target Isilon cluster or backup device, the files consume the same amount of space as if they had not referenced shadow stores.
- When OneFS creates a shadow store, OneFS assigns the shadow store to a storage pool of a file that references the shadow store. If you delete the storage pool that a shadow store resides on, the shadow store is moved to a pool occupied by another file that references the shadow store.
- OneFS does not delete a shadow-store block immediately after the last reference to the block is deleted. Instead, OneFS waits until the ShadowStoreDelete job is run to delete the unreferenced block. If a large number of unreferenced blocks exist on the cluster, OneFS might report a negative deduplication savings until the ShadowStoreDelete job is run.
- Shadow stores are protected at least as much as the most protected file that references it. For example, if one file that references a shadow store resides in a storage pool with +2 protection and another file that references the shadow store resides in a storage pool with +3 protection, the shadow store is protected at +3.
- Quotas account for files that reference shadow stores as if the files contained the data referenced from shadow stores; from the perspective of a quota, shadow-store references do not exist. However, if a quota includes data protection overhead, the quota does not account for the data protection overhead of shadow stores.

SmartDedupe license functionality

You can deduplicate data only if you activate a SmartDedupe license on a cluster. However, you can assess deduplication savings without activating a SmartDedupe license.

If you activate a SmartDedupe license, and then deduplicate data, the space savings are not lost if the license becomes inactive. You can also still view deduplication savings while the license is inactive. However, you will not be able to deduplicate additional data until you re-activate the SmartDedupe license.

Managing deduplication

You can manage deduplication on a cluster by first assessing how much space you can save by deduplicating individual directories. After you determine which directories are

worth deduplicating, you can configure SmartDedupe to deduplicate those directories specifically. You can then monitor the actual amount of disk space you are saving.

Assess deduplication space savings

You can assess the amount of disk space you will save by deduplicating a directory.

Procedure

1. Specify which directory to assess by running the `isi dedupe settings modify` command.

The following command configures SmartDedupe to assess deduplication savings for `/ifs/data/archive`:

```
isi dedupe settings modify --assess-paths /ifs/data/archive
```

If you assess multiple directories, disk savings will not be differentiated by directory in the deduplication report.

2. Start the assessment job by running the following command:

```
isi job jobs start dedupeassessment
```

3. Identify the ID of the assessment report by running the following command:

```
isi dedupe reports list
```

4. View prospective space savings by running the `isi dedupe reports view` command:

The following command displays the prospective savings recorded in a deduplication report with an ID of 46:

```
isi dedupe reports view 46
```

Specify deduplication settings

You can specify which directories you want to deduplicate.

Procedure

1. Specify which directories you want to deduplicate by running the `isi dedupe settings modify` command.

The following command targets `/ifs/data/archive` and `/ifs/data/media` for deduplication:

```
isi dedupe settings modify --paths /ifs/data/media,/ifs/data/archive
```

2. (Optional) To modify the settings of the deduplication job, run the `isi job types modify` command.

The following command configures the deduplication job to be run every Friday at 10:00 PM:

```
isi job types Dedupe --schedule "Every Friday at 10:00 PM"
```

View deduplication space savings

You can view the amount of disk space that you are currently saving with deduplication.

Procedure

1. Run the following command:

```
isi dedupe stats
```

View a deduplication report

After a deduplication job completes, you can view information about the job in a deduplication report.

Procedure

1. (Optional) To identify the ID of the deduplication report you want to view, run the following command:

```
isi dedupe reports list
```

2. View a deduplication report by running the `isi dedupe reports view` command.

The following command displays a deduplication report with an ID of 44:

```
isi dedupe reports view 44
```

Deduplication job report information

You can view the following deduplication specific information in deduplication job reports:

Start time

The time the deduplication job started.

End time

The time the deduplication job ended.

Iteration Count

The number of times that SmartDedupe interrupted the sampling process. If SmartDedupe is sampling a large amount of data, SmartDedupe might interrupt sampling in order to start deduplicating the data. After SmartDedupe finishes deduplicating the sampled data, SmartDedupe will continue sampling the remaining data.

Scanned blocks

The total number of blocks located underneath the specified deduplicated directories.

Sampled blocks

The number of blocks that SmartDedupe created index entries for.

Deduped blocks

The number of blocks that were deduplicated.

Dedupe percent

The percentage of scanned blocks that were deduplicated.

Created dedupe requests

The total number of deduplication requests created. A deduplication request is created for each matching pair of data blocks. For example, if you have 3 data blocks that all match, SmartDedupe creates 2 requests. One of the requests could pair file1 and file2 together and the other request could pair file2 and file3 together.

Successful dedupe requests

The number of deduplication requests that completed successfully.

Failed dedupe requests

The number of deduplication requests that failed. If a deduplication request fails, it doesn't mean that the job failed too. A deduplication request can fail for any number of reasons. For example, the file might have been modified since it was sampled.

Skipped files

The number of files that were not scanned by the deduplication job. SmartDedupe skips files for a number of reasons. For example, SmartDedupe skips files that have already been scanned and haven't been modified since. SmartDedupe also skips all files that are smaller than 4 KB.

Index entries

The number of entries that currently exist in the index.

Index lookup attempts

The total number of lookups that have been done by earlier deduplication jobs plus the number of lookups done by this deduplication job. A lookup is when the deduplication job attempts to match a block that was indexed with a block that hasn't been indexed.

Index lookup hits

The number of blocks that matched index entries.

Deduplication information

You can view information about how much disk space is being saved by deduplication.

The following information is displayed in the output of the `isi dedupe stats` command:

Cluster Physical Size

The total amount of physical disk space on the cluster.

Cluster Used Size

The total amount of disk space currently occupied by data on the cluster.

Logical Size Deduplicated

The amount of disk space that has been deduplicated in terms of reported file sizes. For example, if you have three identical files that are all 5 GB, the logical size deduplicated is 15 GB.

Logical Saving

The amount of disk space saved by deduplication in terms of reported file sizes. For example, if you have three identical files that are all 5 GB, the logical saving is 10 GB.

Estimated Size Deduplicated

The total amount of physical disk space that has been deduplicated, including protection overhead and metadata. For example, if you have three identical files that are all 5 GB, the estimated size deduplicated would be greater than 15 GB, because of the disk space consumed by file metadata and protection overhead.

Estimated Physical Saving

The total amount of physical disk space saved by deduplication, including protection overhead and metadata. For example, if you have three identical files that are all 5 GB, the estimated physical saving would be greater than 10 GB, because deduplication saved space that would have been occupied by file metadata and protection overhead.

Deduplication commands

You can control data deduplication through the deduplication commands. Deduplication commands are available only if you activate a SmartDedupe license.

isi dedupe settings modify

Modifies the settings of deduplication jobs.

Syntax

```
isi dedupe settings modify
[[-paths <path>]... | --clear-paths]
[--add-paths <path>]...
[--remove-paths <path>]...
[[-assess-paths <path>]... | --clear-assess-paths]
[--add-assess-paths <path>]...
[--remove-assess-paths <path>]...
[--verbose]
```

Options

`--paths <path>`

Deduplicates files located under the specified root directories.

`--clear-paths`

Stops deduplication for all previously specified root directories. If you run the `isi dedupe settings modify` command with this option, you must run the command again with either `--paths` or `--add-path` to resume deduplication.

`--add-paths <path>`

Deduplicates files located under the specified root directory in addition to directories that are already being deduplicated.

`--remove-paths <path>`

Stops deduplicating the specified root directory.

`--assess-paths <path>`

Assesses how much space will be saved if files located under the specified root directories are deduplicated.

`--clear-assess-paths`

Stops assessing how much space will be saved if previously specified root directories are deduplicated. If you run the `isi dedupe settings modify` command with

this option, you must run the command again with either `--paths` or `--add-path` to resume deduplication.

`--add-assess-paths <path>`

Assesses how much space will be saved if the specified root directories are deduplicated in addition to directories that are already being assessed.

`--remove-assess-paths <path>`

Stops assessing how much space will be saved if the specified root directories are deduplicated.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command starts deduplicating `/ifs/data/active` and `/ifs/data/media`:

```
isi dedupe settings modify --add-paths /ifs/data/active,/ifs/data/media
```

The following command stops deduplicating `/ifs/data/active` and `/ifs/data/media`:

```
isi dedupe settings modify --remove-paths /ifs/data/active,/ifs/data/media
```

isi dedupe settings view

Displays current deduplication settings.

Syntax

```
isi dedupe settings view
```

Options

There are no options for this command.

isi dedupe stats

Displays information about how much data is being saved by deduplication.

Syntax

```
isi dedupe stats
```

Options

There are no options for this command.

Examples

To view information about deduplication space savings, run the following command:

```
isi dedupe stats
```

The system displays output similar to the following example:

```
Cluster Physical Size: 17.019G
Cluster Used Size: 4.994G
```

```

Logical Size Deduplicated: 13.36M
      Logical Saving: 11.13M
Estimated Size Deduplicated: 30.28M
      Estimated Physical Saving: 25.23M

```

isi dedupe reports list

Displays a list of deduplication reports.

Syntax

```

isi dedupe reports list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]

```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

Examples

To view a list of deduplication reports, run the following command:

```
isi dedupe reports list
```

The system displays output similar to the following example:

```

Time                Job ID Job Type
-----
2013-05-09T11:03:37 4      Dedupe
2013-05-10T00:02:27 8      Dedupe
2013-05-15T13:03:47 12     Dedupe
2013-05-16T00:02:32 16     Dedupe
2013-05-17T00:02:32 19     Dedupe
2013-05-09T16:14:04 5      DedupeAssessment
-----
Total: 6

```


isi dedupe reports view

Displays a deduplication report.

Syntax

```
isi dedupe reports view <job-id>
```

Options

<job-id>

Displays the deduplication report for the deduplication job of the specified ID.

Examples

The following command displays a deduplication job:

```
isi dedupe reports view 12
```

The system displays output similar to the following example:

```
Time: 2013-10-14T09:39:22
Job ID: 52
Job Type: Dedupe
Reports
  Time : 2013-10-14T09:39:22
  Results :
Dedupe job report:{
  Start time = 2013-Oct-14:09:33:34
  End time = 2013-Oct-14:09:39:22
  Iteration count = 1
  Scanned blocks = 1716
  Sampled blocks = 78
  Deduped blocks = 1425
  Dedupe percent = 83.042
  Created dedupe requests = 65
  Successful dedupe requests = 65
  Failed dedupe requests = 0
  Skipped files = 0
  Index entries = 38
  Index lookup attempts = 38
  Index lookup hits = 0
}
Elapsed time:                347 seconds
Aborts:                      0
Errors:                      0
Scanned files:               6
Directories:                 2
2 paths:
/ifs/data/dir2,
/ifs/data/dir1
CPU usage:                   max 29% (dev 2), min 0% (dev 1),
avg 6%
Virtual memory size:         max 128388K (dev 1), min 106628K
(dev 1), avg 107617K
Resident memory size:        max 27396K (dev 1), min 9980K (dev
2), avg 11585K
Read:                        2160 ops, 124437504 bytes (118.7M)
Write:                        30570 ops, 222851584 bytes (212.5M)
```


CHAPTER 13

Data replication with SyncIQ

This section contains the following topics:

- [SyncIQ backup and recovery overview](#) 532
- [Replication policies and jobs](#) 532
- [Replication snapshots](#) 535
- [Data failover and failback with SyncIQ](#) 536
- [Recovery times and objectives for SyncIQ](#) 537
- [SyncIQ license functionality](#) 538
- [Creating replication policies](#) 538
- [Managing replication to remote clusters](#) 543
- [Initiating data failover and failback with SyncIQ](#) 545
- [Performing disaster recovery for SmartLock directories](#) 547
- [Managing replication policies](#) 550
- [Managing replication to the local cluster](#) 552
- [Managing replication performance rules](#) 554
- [Managing replication reports](#) 556
- [Managing failed replication jobs](#) 559
- [Managing changelists](#) 560
- [Data replication commands](#) 564

SyncIQ backup and recovery overview

OneFS enables you to replicate data from one Isilon cluster to another through the SyncIQ software module. You must activate a SyncIQ license on both Isilon clusters before you can replicate data between them.

You can replicate data at the directory level while optionally excluding specific files and sub-directories from being replicated. SyncIQ creates and references snapshots to replicate a consistent point-in-time image of a root directory. Metadata such as access control lists (ACLs) and alternate data streams (ADS) are replicated along with data.

SyncIQ enables you to maintain a consistent backup copy of your data on another Isilon cluster. SyncIQ offers automated failover and failback capabilities that enable you to continue operations on another Isilon cluster if a primary cluster becomes unavailable.

Replication policies and jobs

Data replication is coordinated according to replication policies and jobs. Replication policies specify what data is replicated, where the data is replicated to, and how often the data is replicated. Replication jobs are the operations that replicate data from one Isilon cluster to another. SyncIQ generates replication jobs according to replication policies.

A replication policy specifies two clusters: the source and the target. The cluster on which the replication policy exists is the source cluster. The cluster that data is being replicated to is the target cluster. When a replication policy starts, SyncIQ generates a replication job for the policy. When a replication job runs, files from a directory on the source cluster are replicated to a directory on the target cluster; these directories are known as source and target directories.

After the first replication job created by a replication policy finishes, the target directory and all files contained in the target directory are set to a read-only state, and can be modified only by other replication jobs belonging to the same replication policy. There is no limit to the number of replication policies that can exist on a cluster.

Note

To prevent permissions errors, make sure that ACL policy settings are the same across source and target clusters.

You can create two types of replication policies: synchronization policies and copy policies. A synchronization policy maintains an exact replica of the source directory on the target cluster. If a file or sub-directory is deleted from the source directory, the file or directory is deleted from the target cluster when the policy is run again.

You can use synchronization policies to fail over and fail back data between source and target clusters. When a source cluster becomes unavailable, you can fail over data on a target cluster and make the data available to clients. When the source cluster becomes available again, you can fail back the data to the source cluster.

A copy policy maintains recent versions of the files that are stored on the source cluster. However, files that are deleted on the source cluster are not deleted from the target cluster. Failback is not supported for copy policies. Copy policies are most commonly used for archival purposes.

Copy policies enable you to remove files from the source cluster without losing those files on the target cluster. Deleting files on the source cluster improves performance on the source cluster while maintaining the deleted files on the target cluster. This can be useful

if, for example, your source cluster is being used for production purposes and your target cluster is being used only for archiving.

After creating a job for a replication policy, SyncIQ must wait until the job completes before it can create another job for the policy. Any number of replication jobs can exist on a cluster at a given time; however, only five replication jobs can run on a source cluster at the same time. If more than five replication jobs exist on a cluster, the first five jobs run while the others are queued to run. The number of replication jobs that a single target cluster can support concurrently is dependent on the number of workers available on the target cluster.

You can replicate any number of files and directories with a single replication job. You can prevent a large replication job from overwhelming the system by limiting the amount of cluster resources and network bandwidth that data synchronization is allowed to consume. Because each node in a cluster is able to send and receive data, the speed at which data is replicated increases for larger clusters.

Automated replication policies

You can manually start a replication policy at any time, but you can also configure replication policies to start automatically based on source directory modifications or a schedule.

You can configure a replication policy to run according to a schedule, so that you can control when replication is performed. You can also configure a replication policy to start when SyncIQ detects a modification to the source directory, so that SyncIQ maintains a more current version of your data on the target cluster.

Scheduling a policy can be useful under the following conditions:

- You want to replicate data when user activity is minimal
- You can accurately predict when modifications will be made to the data

Configuring a policy to start when changes are made to the source directory can be useful under the following conditions:

- You want retain a consistent copy of your data at all times
- You are expecting a large number of changes at unpredictable intervals

For policies that are configured to start whenever changes are made to the source directory, SyncIQ checks the source directories every ten seconds. SyncIQ does not account for excluded files or directories when detecting changes, so policies that exclude files or directories from replication might be run unnecessarily. For example, assume that `newPolicy` replicates `/ifs/data/media` but excludes `/ifs/data/media/temp`. If a modification is made to `/ifs/data/media/temp/file.txt`, SyncIQ will run `newPolicy`, but will not replicate `/ifs/data/media/temp/file.txt`.

If a policy is configured to start whenever changes are made to its source directory, and a replication job fails, SyncIQ will wait one minute before attempting to run the policy again. SyncIQ will increase this delay exponentially for each failure up to a maximum delay of eight hours. You can override the delay by running the policy manually at any time. After a job for the policy completes successfully, SyncIQ will resume checking the source directory every ten seconds.

Source and target cluster association

SyncIQ associates a replication policy with a target cluster by marking the target cluster when the job runs for the first time. Even if you modify the name or IP address of the

target cluster, the mark persists on the target cluster. When a replication policy is run, SyncIQ checks the mark to ensure that data is being replicated to the correct location.

On the target cluster, you can manually break an association between a replication policy and target directory. Breaking the association between a source and target cluster causes the mark on the target cluster to be deleted. You might want to manually break a target association if an association is obsolete. If you break the association of a policy, the policy is disabled on the source cluster and you cannot run the policy. If you want to run the disabled policy again, you must reset the replication policy.

Note

Breaking a policy association causes either a full or differential replication to occur the next time you run the replication policy. During a full or differential replication, SyncIQ creates a new association between the source and target clusters. Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete.

Full and differential replication

If a replication policy encounters an issue that cannot be fixed (for example, if the association was broken on the target cluster), you might need to reset the replication policy. If you reset a replication policy, SyncIQ performs either a full or differential replication the next time the policy is run. You can specify the type of replication that SyncIQ performs.

During a full replication, SyncIQ transfers all data from the source cluster regardless of what data exists on the target cluster. A full replication consumes large amounts of network bandwidth and can take a very long time to complete. However, a full replication is less strenuous on CPU usage than a differential replication.

During a differential replication, SyncIQ first checks whether a file already exists on the target cluster and then transfers only data that does not already exist on the target cluster. A differential replication consumes less network bandwidth than a full replication; however, differential replications consume more CPU. Differential replication can be much faster than a full replication if there is an adequate amount of available CPU for the differential replication job to consume.

Controlling replication job resource consumption

You can create rules that limit the network traffic created and the rate at which files are sent by replication jobs. You can also specify the number of workers that are spawned by a replication policy to limit the amount of cluster resources that are consumed. Also, you can restrict a replication policy to connect only to a specific storage pool.

You can create network-traffic rules that control the amount of network traffic generated by replication jobs during specified time periods. These rules can be useful if, for example, you want to limit the amount of network traffic created during other resource-intensive operations.

You can create multiple network traffic rules to enforce different limitations at different times. For example, you might allocate a small amount of network bandwidth during peak business hours, but allow unlimited network bandwidth during non-peak hours.

When a replication job runs, OneFS generates workers on the source and target cluster. Workers on the source cluster send data while workers on the target cluster write data. OneFS generates no more than 40 workers for a replication job. You can modify the maximum number of workers generated per node to control the amount of resources that a replication job is allowed to consume. For example, you can increase the maximum

number of workers per node to increase the speed at which data is replicated to the target cluster.

You can also reduce resource consumption through file-operation rules that limit the rate at which replication policies are allowed to send files. However, it is recommended that you only create file-operation rules if the files you intend to replicate are predictably similar in size and not especially large.

Replication reports

After a replication job completes, SyncIQ generates a report that contains detailed information about the job, including how long the job ran, how much data was transferred, and what errors occurred.

If a replication report is interrupted, SyncIQ might create a subreport about the progress of the job so far. If the job is then restarted, SyncIQ creates another subreport about the progress of the job until the job either completes or is interrupted again. SyncIQ creates a subreport each time the job is interrupted until the job completes successfully. If multiple subreports are created for a job, SyncIQ combines the information from the subreports into a single report.

SyncIQ routinely deletes replication reports. You can specify the maximum number of replication reports that SyncIQ retains and the length of time that SyncIQ retains replication reports. If the maximum number of replication reports is exceeded on a cluster, SyncIQ deletes the oldest report each time a new report is created.

You cannot customize the content of a replication report.

Note

If you delete a replication policy, SyncIQ automatically deletes any reports that were generated for that policy.

Replication snapshots

SyncIQ generates snapshots to facilitate replication, failover, and failback between Isilon clusters. Snapshots generated by SyncIQ can also be used for archival purposes on the target cluster.

Source cluster snapshots

SyncIQ generates snapshots on the source cluster to ensure that a consistent point-in-time image is replicated and that unaltered data is not sent to the target cluster.

Before running a replication job, SyncIQ creates a snapshot of the source directory. SyncIQ then replicates data according to the snapshot rather than the current state of the cluster, allowing users to modify source-directory files while ensuring that an exact point-in-time image of the source directory is replicated.

For example, if a replication job of `/ifs/data/dir/` starts at 1:00 PM and finishes at 1:20 PM, and `/ifs/data/dir/file` is modified at 1:10 PM, the modifications are not reflected on the target cluster, even if `/ifs/data/dir/file` is not replicated until 1:15 PM.

You can replicate data according to a snapshot generated with the SnapshotIQ tool. If you replicate data according to a SnapshotIQ snapshot, SyncIQ does not generate another snapshot of the source directory. This method can be useful if you want to replicate identical copies of data to multiple Isilon clusters.

SyncIQ generates source snapshots to ensure that replication jobs do not transfer unmodified data. When a job is created for a replication policy, SyncIQ checks whether it is the first job created for the policy. If it is not the first job created for the policy, SyncIQ compares the snapshot generated for the earlier job with the snapshot generated for the new job.

SyncIQ replicates only data that has changed since the last time a snapshot was generated for the replication policy. When a replication job is completed, SyncIQ deletes the previous source-cluster snapshot and retains the most recent snapshot until the next job is run.

Target cluster snapshots

When a replication job is run, SyncIQ generates a snapshot on the target cluster to facilitate failover operations. When the next replication job is created for the replication policy, the job creates a new snapshot and deletes the old one.

If a SnapshotIQ license has been activated on the target cluster, you can configure a replication policy to generate additional snapshots that remain on the target cluster even as subsequent replication jobs run.

SyncIQ generates target snapshots to facilitate failover on the target cluster regardless of whether a SnapshotIQ license has been configured on the target cluster. Failover snapshots are generated when a replication job completes. SyncIQ retains only one failover snapshot per replication policy, and deletes the old snapshot after the new snapshot is created.

If a SnapshotIQ license has been activated on the target cluster, you can configure SyncIQ to generate archival snapshots on the target cluster that are not automatically deleted when subsequent replication jobs run. Archival snapshots contain the same data as the snapshots that are generated for failover purposes. However, you can configure how long archival snapshots are retained on the target cluster. You can access archival snapshots the same way that you access other snapshots generated on a cluster.

Data failover and failback with SyncIQ

SyncIQ enables you to perform automated data failover and failback operations between Isilon clusters. If a cluster is rendered unusable, you can fail over to another Isilon cluster, enabling clients to access their data on the other cluster. If the unusable cluster becomes accessible again, you can fail back to the original Isilon cluster.

For the purposes of explaining failover and failback procedures, the cluster originally accessed by clients is referred to as the primary cluster, and the cluster that client data is originally replicated to is referred to as the secondary cluster. Failover is the process that allows clients to modify data on a secondary cluster. Failback is the process that allows clients to access data on the primary cluster again and begins to replicate data back to the secondary cluster.

Failover and failback can be useful in disaster recovery procedures. For example, if a primary cluster is damaged by a natural disaster, you can migrate clients to a secondary cluster until the primary cluster is repaired and then migrate the clients back to the primary cluster.

You can fail over and fail back to facilitate scheduled cluster maintenance. For example, if you are upgrading the primary cluster, you might want to migrate clients to a secondary cluster until the upgrade is complete and then migrate clients back to the primary cluster.

Note

Data failover and failback is not supported for SmartLock directories.

Data failover

Data failover is the process of preparing data on a secondary cluster to be modified by clients. After you fail over to a secondary cluster, you can redirect clients to modify their data on the secondary cluster.

Before failover is performed, you must create and run a replication policy on the primary cluster. You initiate the failover process on the secondary cluster. Failover is performed per replication policy; to migrate data that is spread across multiple replication policies, you must initiate failover for each replication policy.

You can use any replication policy to fail over. However, if the action of the replication policy is set to copy, any file that was deleted on the primary cluster will be present on the secondary cluster. When the client connects to the secondary cluster, all files that were deleted on the primary cluster will be available to the client.

If you initiate failover for a replication policy while an associated replication job is running, the failover operation completes but the replication job fails. Because data might be in an inconsistent state, SyncIQ uses the snapshot generated by the last successful replication job to revert data on the secondary cluster to the last recovery point.

If a disaster occurs on the primary cluster, any modifications to data that were made after the last successful replication job started are not reflected on the secondary cluster. When a client connects to the secondary cluster, their data appears as it was when the last successful replication job was started.

Data failback

Data failback is the process of restoring clusters to the roles they occupied before a failover operation. After data failback is complete, the primary cluster hosts clients and replicates data to the secondary cluster for backup.

The first step in the failback process is updating the primary cluster with all of the modifications that were made to the data on the secondary cluster. The next step in the failback process is preparing the primary cluster to be accessed by clients. The final step in the failback process is resuming data replication from the primary to the secondary cluster. At the end of the failback process, you can redirect users to resume accessing their data on the primary cluster.

You can fail back data with any replication policy that meets all of the following criteria:

- The source directory is not a SmartLock directory.
- The policy has been failed over.
- The policy is a synchronization policy.
- The policy does not exclude any files or directories from replication.

Recovery times and objectives for SyncIQ

The Recovery Point Objective (RPO) and the Recovery Time Objective (RTO) are measurements of the impacts that a disaster can have on business operations. You can calculate your RPO and RTO for a disaster recovery with replication policies.

RPO is the maximum amount of time for which data is lost if a cluster suddenly becomes unavailable. For an Isilon cluster, the RPO is the amount of time that has passed since

the last completed replication job started. The RPO is never greater than the time it takes for two consecutive replication jobs to run and complete.

If a disaster occurs while a replication job is running, the data on the secondary cluster is reverted to the state it was in when the last replication job completed. For example, consider an environment in which a replication policy is scheduled to run every three hours, and replication jobs take two hours to complete. If a disaster occurs an hour after a replication job begins, the RPO is four hours, because it has been four hours since a completed job began replicating data.

RTO is the maximum amount of time required to make backup data available to clients after a disaster. The RTO is always less than or approximately equal to the RPO, depending on the rate at which replication jobs are created for a given policy.

If replication jobs run continuously, meaning that another replication job is created for the policy before the previous replication job completes, the RTO is approximately equal to the RPO. When the secondary cluster is failed over, the data on the cluster is reset to the state it was in when the last job completed; resetting the data takes an amount of time proportional to the time it took users to modify the data.

If replication jobs run on an interval, meaning that there is a period of time after a replication job completes before the next replication job for the policy starts, the relationship between RTO and RPO depends on whether a replication job is running when the disaster occurs. If a job is in progress when a disaster occurs, the RTO is roughly equal to the RPO. However, if a job is not running when a disaster occurs, the RTO is negligible because the secondary cluster was not modified since the last replication job ran, and the failover process is almost instantaneous.

SyncIQ license functionality

You can replicate data to another Isilon cluster only if you activate a SyncIQ license on both the local cluster and the target cluster.

If a SyncIQ license becomes inactive, you cannot create, run, or manage replication policies. Also, all previously created replication policies are disabled. Replication policies that target the local cluster are also disabled. However, data that was previously replicated to the local cluster is still available.

Creating replication policies

You can create replication policies that determine when data is replicated with SyncIQ.

Excluding directories in replication

You can exclude directories from being replicated by replication policies even if the directories exist under the specified source directory.

Note

Failback is not supported for replication policies that exclude directories.

By default, all files and directories under the source directory of a replication policy are replicated to the target cluster. However, you can prevent directories under the source directory from being replicated.

If you specify a directory to exclude, files and directories under the excluded directory are not replicated to the target cluster. If you specify a directory to include, only the files and

directories under the included directory are replicated to the target cluster; any directories that are not contained in an included directory are excluded.

If you both include and exclude directories, any excluded directories must be contained in one of the included directories; otherwise, the excluded-directory setting has no effect. For example, consider a policy with the following settings:

- The root directory is `/ifs/data`
- The included directories are `/ifs/data/media/music` and `/ifs/data/media/movies`
- The excluded directories are `/ifs/data/archive` and `/ifs/data/media/music/working`

In this example, the setting that excludes the `/ifs/data/archive` directory has no effect because the `/ifs/data/archive` directory is not under either of the included directories. The `/ifs/data/archive` directory is not replicated regardless of whether the directory is explicitly excluded. However, the setting that excludes the `/ifs/data/media/music/working` directory does have an effect, because the directory would be replicated if the setting was not specified.

In addition, if you exclude a directory that contains the source directory, the exclude-directory setting has no effect. For example, if the root directory of a policy is `/ifs/data`, explicitly excluding the `/ifs` directory does not prevent `/ifs/data` from being replicated.

Any directories that you explicitly include or exclude must be contained in or under the specified root directory. For example, consider a policy in which the specified root directory is `/ifs/data`. In this example, you could include both the `/ifs/data/media` and the `/ifs/data/users/` directories because they are under `/ifs/data`.

Excluding directories from a synchronization policy does not cause the directories to be deleted on the target cluster. For example, consider a replication policy that synchronizes `/ifs/data` on the source cluster to `/ifs/data` on the target cluster. If the policy excludes `/ifs/data/media` from replication, and `/ifs/data/media/file` exists on the target cluster, running the policy does not cause `/ifs/data/media/file` to be deleted from the target cluster.

Excluding files in replication

If you do not want specific files to be replicated by a replication policy, you can exclude them from the replication process through file-matching criteria statements. You can configure file-matching criteria statements during the replication-policy creation process.

Note

You cannot fail back replication policies that exclude files.

A file-criteria statement can include one or more elements. Each file-criteria element contains a file attribute, a comparison operator, and a comparison value. You can combine multiple criteria elements in a criteria statement with Boolean "AND" and "OR" operators. You can configure any number of file-criteria definitions.

Configuring file-criteria statements can cause the associated jobs to run slowly. It is recommended that you specify file-criteria statements in a replication policy only if necessary.

Modifying a file-criteria statement will cause a full replication to occur the next time that a replication policy is started. Depending on the amount of data being replicated, a full replication can take a very long time to complete.

For synchronization policies, if you modify the comparison operators or comparison values of a file attribute, and a file no longer matches the specified file-matching criteria, the file is deleted from the target the next time the job is run. This rule does not apply to copy policies.

File criteria options

You can configure a replication policy to exclude files that meet or do not meet specific criteria.

You can specify file criteria based on the following file attributes:

Date created

Includes or excludes files based on when the file was created. This option is available for copy policies only.

You can specify a relative date and time, such as "two weeks ago", or specific date and time, such as "January 1, 2012." Time settings are based on a 24-hour clock.

Date accessed

Includes or excludes files based on when the file was last accessed. This option is available for copy policies only, and only if the global access-time-tracking option of the cluster is enabled.

You can specify a relative date and time, such as "two weeks ago", or specific date and time, such as "January 1, 2012." Time settings are based on a 24-hour clock.

Date modified

Includes or excludes files based on when the file was last modified. This option is available for copy policies only.

You can specify a relative date and time, such as "two weeks ago", or specific date and time, such as "January 1, 2012." Time settings are based on a 24-hour clock.

File name

Includes or excludes files based on the file name. You can specify to include or exclude full or partial names that contain specific text.

The following wildcard characters are accepted:

Note

Alternatively, you can filter file names by using POSIX regular-expression (regex) text. Isilon clusters support IEEE Std 1003.2 (POSIX.2) regular expressions. For more information about POSIX regular expressions, see the BSD man pages.

Table 17 Replication file matching wildcards

Wildcard	Description
*	Matches any string in place of the asterisk. For example, <code>m*</code> matches <code>movies</code> and <code>m123</code> .
[]	Matches any characters contained in the brackets, or a range of characters separated by a dash. For example, <code>b[aei]t</code> matches <code>bat</code> , <code>bet</code> , and <code>bit</code> . For example, <code>1[4-7]2</code> matches <code>142</code> , <code>152</code> , <code>162</code> , and <code>172</code> . You can exclude characters within brackets by following the first bracket with an exclamation mark. For example, <code>b[!ie]</code> matches <code>bat</code> but not <code>bit</code> or <code>bet</code> . You can match a bracket within a bracket if it is either the first or last character. For example, <code>[c]at</code> matches <code>cat</code> and <code>[at]</code> . You can match a dash within a bracket if it is either the first or last character. For example, <code>car[-s]</code> matches <code>cars</code> and <code>car-</code> .
?	Matches any character in place of the question mark. For example, <code>t?p</code> matches <code>tap</code> , <code>tip</code> , and <code>top</code> .

Path

Includes or excludes files based on the file path. This option is available for copy policies only.

You can specify to include or exclude full or partial paths that contain specified text.

You can also include the wildcard characters `*`, `?`, and `[]`.

Size

Includes or excludes files based on their size.

Note

File sizes are represented in multiples of 1024, not 1000.

Type

Includes or excludes files based on one of the following file-system object types:

- Soft link
- Regular file
- Directory

Configure default replication policy settings

You can configure default settings for replication policies. If you do not modify these settings when creating a replication policy, the specified default settings are applied.

Procedure

1. Run the `isi sync settings modify` command.

The following command configures SyncIQ to delete replication reports that are older than 2 years:

```
isi sync settings modify --report-max-age 2Y
```

Create a replication policy

You can create a replication policy with SyncIQ that defines how and when data is replicated to another Isilon cluster. Configuring a replication policy is a five-step process.

Configure replication policies carefully. If you modify any of the following policy settings after the policy is run, OneFS performs either a full or differential replication the next time the policy is run:

- Source directory
- Included or excluded directories
- File-criteria statement
- Target cluster name or address
This applies only if you target a different cluster. If you modify the IP or domain name of a target cluster, and then modify the replication policy on the source cluster to match the new IP or domain name, a full replication is not performed.
- Target directory

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

Procedure

1. Run the `isi sync policies create` command.

The following command creates a policy that replicates `/ifs/data/source` on the local cluster to `/ifs/data/target` on `cluster.domain.name` every week. The command also creates archival snapshots on the target cluster:

```
isi sync policies create weeklySync sync /ifs/data/source \
cluster.domain.name /ifs/data/target \
--schedule "Every Saturday at 12:00 AM" \
```

```
--target-snapshot-archive on \
--target-snapshot-pattern \
"%{PolicyName}-%{SrcCluster}-%Y-%m-%d_%H-%M"\
--target-snapshot-expiration 1Y
```

Create a SyncIQ domain

You can create a SyncIQ domain to increase the speed at which failback is performed for a replication policy. Because you can fail back only synchronization policies, it is not necessary to create SyncIQ domains for copy policies.

Failing back a replication policy requires that a SyncIQ domain be created for the source directory. OneFS automatically creates a SyncIQ domain during the failback process. However, if you intend on failing back a replication policy, it is recommended that you create a SyncIQ domain for the source directory of the replication policy while the directory is empty. Creating a domain for a directory that contains less data takes less time.

Procedure

1. Run the `isi job jobs start` command.

The following command creates a SyncIQ domain for `/ifs/data/source`:

```
isi job jobs start domainmark --root /ifs/data/media \
--dm-type SyncIQ
```

Assess a replication policy

Before running a replication policy for the first time, you can view statistics on the files that would be affected by the replication without transferring any files. This can be useful if you want to preview the size of the data set that will be transferred if you run the policy.

You can assess only replication policies that have never been run before.

Procedure

1. Run the `isi sync jobs start` command with the `--test` option.

The following command creates a report about how much data will be transferred when `weeklySync` is run:

```
isi sync jobs start weeklySync --test
```

2. To view the assessment report, run the `isi sync reports view` command.

The following command displays the assessment report for `weeklySync`:

```
isi sync reports view weeklySync 1
```

Managing replication to remote clusters

You can manually run, view, assess, pause, resume, cancel, resolve, and reset replication jobs that target other clusters.

After a policy job starts, you can pause the job to suspend replication activities. Afterwards, you can resume the job, continuing replication from the point where the job was interrupted. You can also cancel a running or paused replication job if you want to free the cluster resources allocated for the job. A paused job reserves cluster resources whether or not the resources are in use. A cancelled job releases its cluster resources and allows another replication job to consume those resources. No more than five running and paused replication jobs can exist on a cluster at a time. However, an unlimited

number of canceled replication jobs can exist on a cluster. If a replication job remains paused for more than a week, SyncIQ automatically cancels the job.

Start a replication job

You can manually start a replication job for a replication policy at any time. You can also replicate data according to a snapshot created by SnapshotIQ. You cannot replicate data according to a snapshot generated by SyncIQ.

Procedure

1. Run the `isi sync jobs start` command.

The following command starts `weeklySync`:

```
isi sync jobs start weeklySync
```

The following command replicates the source directory of `weeklySync` according to the snapshot `HourlyBackup_07-15-2013_23:00`:

```
isi sync jobs start weeklySync \
--source-snapshot HourlyBackup_07-15-2013_23:00
```

Pause a replication job

You can pause a running replication job and then resume the job later. Pausing a replication job temporarily stops data from being replicated, but does not free the cluster resources replicating the data.

Procedure

1. Run the `isi sync jobs pause` command.

The following command pauses `weeklySync`:

```
isi sync jobs pause weeklySync
```

Resume a replication job

You can resume a paused replication job.

Procedure

1. Run the `isi sync jobs resume` command.

The following command resumes `weeklySync`:

```
isi sync jobs resume weeklySync
```

Cancel a replication job

You can cancel a running or paused replication job. Cancelling a replication job stops data from being replicated and frees the cluster resources that were replicating data. You cannot resume a cancelled replication job; to restart replication, you must start the replication policy again.

Procedure

1. Run the `isi sync jobs cancel` command.

The following command cancels `weeklySync`:

```
isi sync jobs cancel weeklySync
```


View active replication jobs

You can view information about replication jobs that are currently running or paused.

Procedure

1. View all active replication jobs by running the following command:

```
isi sync jobs list
```

2. To view detailed information about a specific replication job, run the `isi sync jobs view` command.

The following command displays detailed information about a replication job created by `weeklySync`:

```
isi sync jobs view weeklySync
```

The system displays output similar to the following example:

```
Policy Name: weeklySync
           ID: 3
           State: running
           Action: run
           Duration: 5s
           Start Time: 2013-07-16T23:12:00
```

Replication job information

You can view information about replication jobs.

The following information is displayed in the output of the `isi snapshot settings view` command:

Policy Name

The name of the associated replication policy.

ID

The ID of the replication job.

State

The status of the job.

Action

The type of replication policy.

Initiating data failover and failback with SyncIQ

You can fail over from one Isilon cluster to another if, for example, a cluster becomes unavailable. You can then fail back to a primary cluster if the primary cluster becomes available again. You can revert failover if you decide that the failover was unnecessary, or if you failed over for testing purposes.

Note

Although you cannot fail over or fail back SmartLock directories, you can recover SmartLock directories on a target cluster. After you recover SmartLock directories, you can migrate them back to the source cluster.

Fail over data to a secondary cluster

You can fail over to a secondary Isilon cluster if, for example, a cluster becomes unavailable.

Complete the following procedure for each replication policy that you want to fail over.

Procedure

1. Run the `isi sync recovery allow-write` command.

The following command fails over data for `weeklySync`:

```
isi sync recovery allow-write weeklySync
```

2. On the primary cluster, modify the replication policy so that it is set to run only manually.

The following command sets `newPolicy` to run only manually:

```
isi sync policies modify newPolicy --schedule ""
```

This step will prevent the policy on the primary cluster from automatically running a replication job. If the policy on the primary cluster runs a replication job while writes are allowed to the target directory, the job will fail and the policy will be set to an unrunnable state. If this happens, modify the replication policy so that it is set to run only manually, resolve the policy, and complete the failback process. After you complete the failback process, you can modify the policy to run according to a schedule again.

Revert a failover operation

Failover reversion undoes a failover operation on a secondary cluster, enabling you to replicate data from the primary cluster to the secondary cluster again. Failover reversion is useful if the primary cluster becomes available before data is modified on the secondary cluster or if you failed over to a secondary cluster for testing purposes.

Before you begin

Fail over a replication policy.

Reverting a failover operation does not migrate modified data back to the primary cluster. To migrate data that clients have modified on the secondary cluster, you must fail back to the primary cluster.

Note

Failover reversion is not supported for SmartLock directories.

Complete the following procedure for each replication policy that you want to fail over.

Procedure

1. Run the `isi sync recovery allow-write` command with the `--revert` option.

For example, the following command reverts a failover operation for `newPolicy`:

```
isi sync recovery allow-write newPolicy --revert
```

Fail back data to a primary cluster

After you fail over to a secondary cluster, you can fail back to the primary cluster.

Before you begin

Fail over a replication policy.

Procedure

1. Create mirror policies on the secondary cluster by running the `isi sync recovery resync-prep` command on the primary cluster.

The following command creates a mirror policy for weeklySync:

```
isi sync recovery resync-prep weeklySync
```

SyncIQ names mirror policies according to the following pattern:

```
<replication-policy-name>_mirror
```

2. On the secondary cluster, replicate data to the primary cluster with the mirror policies.

You can replicate data either by manually starting the mirror policies or by modifying the mirror policies and specifying a schedule.

3. Prevent clients from accessing the secondary cluster and then run each mirror policy again.

To minimize impact to clients, it is recommended that you wait until client access is low before preventing client access to the cluster.

4. On the primary cluster, allow writes to the target directories of the mirror policies by running the `isi sync recovery allow-write` command.

The following command allows writes to the target directory of weeklySync_mirror:

```
isi sync recovery allow-write weeklySync_mirror
```

5. On the secondary cluster, complete the failback process by running the `isi sync recovery resync-prep` command for all mirror policies.

The following command completes the failback process for weeklySync:

```
isi sync recovery resync-prep weeklySync_mirror
```

Performing disaster recovery for SmartLock directories

Although you cannot fail over or fail back SmartLock directories, you can recover SmartLock directories on a target cluster. After you recover SmartLock directories, you can migrate them back to the source cluster.

Recover SmartLock directories on a target cluster

You can recover SmartLock directories that you have replicated to a target cluster.

Create and successfully run a replication policy.

Complete the following procedure for each SmartLock directory that you want to recover.

Procedure

1. On the target cluster, enable writes to the target SmartLock directories.

- If the last replication job completed successfully and a replication job is not currently running, run the `isi sync recovery allow-write` command. For example, the following command enables writes to the target directory of SmartLockSync:

```
isi sync recovery allow-write SmartLockSync
```

- If a replication job is currently running, wait until the replication job completes, and then run the `isi sync recovery allow-write` command. For example, the following command enables writes to the target directory of SmartLockSync:

```
isi sync recovery allow-write SmartLockSync
```

- If the primary cluster became unavailable while a replication job was running, select run the `isi sync target break` command. For example, the following command enables writes to the target directory of SmartLockSync:

```
isi sync target break SmartLockSync
```

2. If you ran `isi sync target break`, restore any files that are left in an inconsistent state.
 - a. Delete all files that are not committed to a WORM state from the target directory.
 - b. Copy all files from the failover snapshot to the target directory.

Failover snapshots are named according to the following naming pattern:

```
SIQ-Failover-<policy-name>-<year>-<month>-<day>_<hour>-<minute>-<second>
```

Snapshots are located under the hidden `/ifs/.snapshot` directory.

3. If any SmartLock directory configuration settings, such as an autocommit time period, were specified for the source directory of the replication policy, apply those settings to the target directory.

Because autocommit information is not transferred to the target cluster, files that were scheduled to be committed to a WORM state on the source cluster will not be scheduled to be committed at the same time on the target cluster. To ensure that all files are retained for the appropriate time period, you can commit all files in target SmartLock directories to a WORM state. For example, the following command automatically commits all files in `/ifs/data/smartlock` to a WORM state after one minute.

```
isi worm domains modify --domain /ifs/data/smartlock \
--autocommit-offset 1m
```

Migrate SmartLock directories

You might want to migrate SmartLock directories if you restored the directories on a target cluster, and want to transfer those directories either back to the source cluster or to a new cluster.

Procedure

1. On a cluster, create a replication policy for each directory that you want to migrate.

The policies must meet the following requirements:

- The source directory is the SmartLock directory that you are migrating.
- The target directory must be an empty SmartLock directory. The directory must be of the same SmartLock type as the source directory of a policy you are failing back. For example, if the target directory is a compliance directory, the source must also be a compliance directory.

2. Replicate data to the target cluster by running the policies you created.

You can replicate data either by manually starting the policies or by specifying a policy schedule.

3. (Optional) To ensure that SmartLock protection is enforced for all files, commit all files in the SmartLock directory to a WORM state.

Because autocommit information is not transferred to the target cluster, files that were scheduled to be committed to a WORM state on the source cluster will not be scheduled to be committed at the same time on the target cluster. To ensure that all files are retained for the appropriate time period, you can commit all files in target SmartLock directories to a WORM state.

For example, the following command automatically commits all files in `/ifs/data/smartlock` to a WORM state after one minute:

```
isi worm domains modify --domain /ifs/data/smartlock \
--autocommit-offset 1m
```

This step is necessarily only if you have configured an autocommit time period for the SmartLock directory.

4. Prevent clients from accessing the source cluster and run the policy that you created.

To minimize impact to clients, it is recommended that you wait until client access is low before preventing client access to the cluster.

5. On the target cluster, enable writes to the target directories of the replication policies by running the `isi sync recovery allow-writes` command.

For example, the following command enables writes to the target directory of SmartLockSync:

```
isi sync recovery allow-writes SmartLockSync
```

6. If any SmartLock directory configuration settings, such as an autocommit time period, were specified for the source directories of the replication policies, apply those settings to the target directories.
7. Delete the copy of your SmartLock data on the source cluster.

If the SmartLock directories are compliance directories or enterprise directories with the privileged delete functionality permanently disabled, you cannot recover the space consumed by the source SmartLock directories until all files are released from a WORM state. If you want to free the space before files are released from a WORM state, contact Isilon Technical Support for information about reformatting your cluster.

Managing replication policies

You can modify, view, enable and disable replication policies.

Modify a replication policy

You can modify the settings of a replication policy.

If you modify any of the following policy settings after the policy runs, OneFS performs either a full or differential replication the next time the policy runs:

- Source directory
- Included or excluded directories
- File-criteria statement
- Target cluster
This applies only if you target a different cluster. If you modify the IP or domain name of a target cluster, and then modify the replication policy on the source cluster to match the new IP or domain name, a full replication is not performed.
- Target directory

Procedure

1. Run the `isi sync policies modify` command.

Assuming that `weeklySync` has been reset and has not been run since it was reset, the following command causes a differential replication to be performed the next time `weeklySync` is run:

```
isi sync policies modify weeklySync \
--target-compare-initial-sync on
```

Delete a replication policy

You can delete a replication policy. Once a policy is deleted, SyncIQ no longer creates replication jobs for the policy. Deleting a replication policy breaks the target association on the target cluster, and allows writes to the target directory.

If you want to temporarily suspend a replication policy from creating replication jobs, you can disable the policy, and then enable the policy again later.

Procedure

1. Run the `isi sync policies delete` command.

The following command deletes `weeklySync` from the source cluster and breaks the target association on the target cluster:

```
isi sync policies delete weeklySync
```

Note

The operation will not succeed until SyncIQ can communicate with the target cluster; until then, the policy will still appear in the output of the `isi sync policies list` command. After the connection between the source cluster and target cluster is reestablished, SyncIQ will delete the policy the next time that the job is scheduled to run; if the policy is configured to run only manually, you must manually run the policy again. If SyncIQ is permanently unable to communicate with the target cluster, run the `isi sync policies delete` command with the `--local-only` option. This will delete the policy from the local cluster only and not break the target association on the target cluster.

Enable or disable a replication policy

You can temporarily suspend a replication policy from creating replication jobs, and then enable it again later.

Note

If you disable a replication policy while an associated replication job is running, the running replication job is not interrupted. However, the policy will not create another job until the policy is enabled.

Procedure

1. Run either the `isi sync policies enable` or the `isi sync policies disable` command.

The following command enables `weeklySync`:

```
isi sync policies enable weeklySync
```

The following command disables `weeklySync`:

```
isi sync policies disable weeklySync
```

View replication policies

You can view information about replication policies.

Procedure

1. View information about all replication policies by running the following command:

```
isi sync policies list
```

2. (Optional) To view detailed information about a specific replication policy, run the `isi sync policies view` command.

The following command displays detailed information about `weeklySync`:

```
isi sync policies view weeklySync
```

The system displays output similar to the following example:

```
ID: dd16d277ff995a78e9afbbba6f6e2919
Name: weeklySync
Path: /ifs/data/archive
Action: sync
```

```

Enabled: No
Target: localhost
Description:
Check Integrity: Yes
Source Include Directories: -
Source Exclude Directories: -
Source Subnet: -
Source Pool: -
Source Match Criteria:
Target Path: /ifs/data/sometarget
Target Snapshot Archive: No
Target Snapshot Pattern: SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M-%S
Target Snapshot Expiration: Never
Target Snapshot Alias: SIQ-#{SrcCluster}-#{PolicyName}-latest
Target Detect Modifications: Yes
Source Snapshot Archive: No
Source Snapshot Pattern:
Source Snapshot Expiration: Never
Schedule: Manually scheduled
Log Level: notice
Log Removed Files: No
Workers Per Node: 3
Report Max Age: 2Y
Report Max Count: 2000
Force Interface: No
Restrict Target Network: No
Target Compare Initial Sync: No
Disable Stf: No
Disable Fofb: No
Resolve: -
Last Job State: finished
Last Started: 2013-07-17T15:39:49
Last Success: 2013-07-17T15:39:49
Password Set: No
Conflicted: No
Has Sync State: Yes

```

Replication policy information

You can view information about replication policies through the output of the `isi sync policies list` command.

Name

The name of the policy.

Path

The path of the source directory on the source cluster.

Action

The type of replication policy.

Enabled

Whether the policy is enabled or disabled.

Target

The IP address or fully qualified domain name of the target cluster.

Managing replication to the local cluster

You can interrupt replication jobs that target the local cluster.

You can cancel a currently running job that targets the local cluster, or you can break the association between a policy and its specified target. Breaking a source and target

cluster association causes SyncIQ to perform a full replication the next time the policy is run.

Cancel replication to the local cluster

You can cancel a replication job that is targeting the local cluster.

Procedure

1. Run the `isi sync target cancel` command.

- To cancel a job, specify a replication policy. For example, the following command cancels a replication job created according to `weeklySync`:

```
isi sync target cancel weeklySync
```

- To cancel all jobs targeting the local cluster, run the following command:

```
isi sync target cancel --all
```

Break local target association

You can break the association between a replication policy and the local cluster. Breaking this association requires you to reset the replication policy before you can run the policy again.

Note

After a replication policy is reset, SyncIQ performs a full or differential replication the next time the policy is run. Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete.

Procedure

1. Run the `isi sync target break` command.

The following command breaks the association between `weeklySync` and the local cluster:

```
isi sync target break weeklySync
```

View replication policies targeting the local cluster

You can view information about replication policies that are currently replicating data to the local cluster.

Procedure

1. View information about all replication policies that are currently targeting the local cluster by running the following command:

```
isi sync target list
```

2. To view detailed information about a specific replication policy, run the `isi sync target view` command.

The following command displays detailed information about `weeklySync`:

```
isi sync target view weeklySync
```

The system displays output similar to the following example:

```
Name: weeklySync
Source: cluster
Target Path: /ifs/data/sometarget
Last Job State: finished
FOFB State: writes_disabled
Source Cluster GUID: 000c295159ae74fcde517c1b85adc03daff9
Last Source Coordinator IP: 127.0.0.1
Legacy Policy: No
Last Update: 2013-07-17T15:39:51
```

Remote replication policy information

You can view information about replication policies that are currently targeting the local cluster through the output of the `isi sync target list` command.

Name

The name of the replication policy.

Source

The name of the source cluster.

Target Path

The path of the target directory on the target cluster.

Last Job State

The state of the most recent replication job for the policy.

FOFB State

The failover-failback state of the target directory.

Managing replication performance rules

You can manage the impact of replication on cluster performance by creating rules that limit the network traffic created and the rate at which files are sent by replication jobs.

Create a network traffic rule

You can create a network traffic rule that limits the amount of network traffic that replication policies are allowed to generate during a specified time period.

Procedure

1. Run the `isi sync rules create` command.

The following command creates a network traffic rule that limits bandwidth consumption to 100 KB per second from 9:00 AM to 5:00 PM every weekday:

```
isi sync rules create bandwidth 9:00-17:00 M-F 100
```

Create a file operations rule

You can create a file-operations rule that limits the number of files that replication jobs can send per second.

Procedure

1. Run the `isi sync rules create` command.

The following command creates a file-operations rule that limits the file-send rate to 3 files per second from 9:00 AM to 5:00 PM every weekday: :

```
isi sync rules create file_count 9:00-17:00 M-F 3
```

Modify a performance rule

You can modify a performance rule.

Procedure

1. (Optional) To identify the ID of the performance rule you want to modify, run the following command:

```
isi sync rules list
```

2. Modify a performance rule by running the `isi sync rules modify` command.

The following command causes a performance rule with an ID of `bw-0` to be enforced only on Saturday and Sunday:

```
isi sync rules modify bw-0 --days X,S
```

Delete a performance rule

You can delete a performance rule.

Procedure

1. (Optional) To identify the ID of the performance rule you want to modify, run the following command:

```
isi sync rules list
```

2. Delete a performance rule by running the `isi sync rules delete` command.

The following command deletes a performance rule with an ID of `bw-0`:

```
isi sync rules delete bw-0
```

Enable or disable a performance rule

You can disable a performance rule to temporarily prevent the rule from being enforced. You can also enable a performance rule after it has been disabled.

Procedure

1. (Optional) To identify the ID of the performance rule you want to enable or disable, run the following command:

```
isi sync rules list
```

2. Run the `isi sync rules modify` command.

The following command enables a performance rule with an ID of `bw-0`:

```
isi sync rules modify bw-0 --enabled true
```

The following command disables a performance rule with an ID of `bw-0`:

```
isi sync rules modify bw-0 --enabled false
```

View performance rules

You can view performance rules.

Procedure

1. View information about all performance rules by running the following command:

```
isi sync rules list
```

2. (Optional) To view detailed information about a specific performance rule, run the `isi sync rules view` command.

The following command displays detailed information about a performance rule with an ID of `bw-0`:

```
isi sync rules view --id bw-0
```

The system displays output similar to the following example:

```

      ID: bw-0
  Enabled: Yes
    Type: bandwidth
    Limit: 100 kbps
    Days: Sun,Sat
  Schedule
    Begin : 09:00
    End   : 17:00
Description: Bandwidth rule for weekdays

```

Managing replication reports

In addition to viewing replication reports, you can configure how long reports are retained on the cluster. You can also delete any reports that have passed their expiration period.

Configure default replication report settings

You can configure the default amount of time that SyncIQ retains replication reports for. You can also configure the maximum number of reports that SyncIQ retains for each replication policy.

Procedure

1. Run the `isi sync settings modify` command.

The following command causes OneFS to delete replication reports that are older than 2 years:

```
isi sync settings modify --report-max-age 2Y
```

Delete replication reports

Replication reports are routinely deleted by SyncIQ after the expiration date for the reports has passed. SyncIQ also deletes reports after the number of reports exceeds the specified limit. Excess reports are periodically deleted by SyncIQ; however, you can manually delete all excess replication reports at any time. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Delete excess replication reports by running the following command:

```
isi sync reports rotate
```

View replication reports

You can view replication reports and subreports.

Procedure

1. View a list of all replication reports by running the following command:

```
isi sync reports list
```

2. View a replication report by running the `isi sync reports view` command.

The following command displays a replication report for `weeklySync`:

```
isi sync reports view weeklySync 2
```

3. (Optional) To view a list of subreports for a report, run the `isi sync reports subreports list` command.

The following command displays subreports for `weeklySync`:

```
isi sync reports subreports list weeklySync 1
```

4. (Optional) To view a subreport, run the `isi sync reports subreports view` command.

The following command displays a subreport for `weeklySync`:

```
isi sync reports subreports view weeklySync 1 2
```

The system displays output similar to the following example:

```

Policy Name: weeklySync
  Job ID: 1
Subreport ID: 2
  Start Time: 2013-07-17T21:59:10
  End Time: 2013-07-17T21:59:15
  Action: run
  State: finished
  Policy ID: a358db8b248bf432c71543e0f02df64e
  Sync Type: initial
  Duration: 5s
  Errors: -
Source Directories Visited: 0
Source Directories Deleted: 0
Target Directories Deleted: 0
Source Directories Created: 0
Target Directories Created: 0
  Source Directories Linked: 0
  Target Directories Linked: 0
Source Directories Unlinked: 0
Target Directories Unlinked: 0
  Num Retransmitted Files: 0
  Retransmitted Files: -
  Total Files: 0
  Files New: 0
Source Files Deleted: 0
  Files Changed: 0
Target Files Deleted: 0

```

```

Up To Date Files Skipped: 0
User Conflict Files Skipped: 0
Error Io Files Skipped: 0
Error Net Files Skipped: 0
Error Checksum Files Skipped: 0
Bytes Transferred: 245
Total Network Bytes: 245
Total Data Bytes: 20
File Data Bytes: 20
Sparse Data Bytes: 0
Target Snapshots: SIQ-Failover-
newPol123-2013-07-17_21-59-15, newPol123-Archive-cluster-17
Total Phases: 2
Phases
Phase : STF_PHASE_IDMAP_SEND
Start Time : 2013-07-17T21:59:11
End Time : 2013-07-17T21:59:13

```

Replication report information

You can view information about replication jobs through the **Reports** table.

Policy Name

The name of the associated policy for the job. You can view or edit settings for the policy by clicking the policy name.

Status

Displays the status of the job. The following job statuses are possible:

Running

The job is currently running without error.

Paused

The job has been temporarily paused.

Finished

The job completed successfully.

Failed

The job failed to complete.

Started

Indicates when the job started.

Ended

Indicates when the job ended.

Duration

Indicates how long the job took to complete.

Transferred

The total number of files that were transferred during the job run, and the total size of all transferred files. For assessed policies, `Assessment` appears.

Source Directory

The path of the source directory on the source cluster.

Target Host

The IP address or fully qualified domain name of the target cluster.

Action

Displays any report-related actions that you can perform.

Managing failed replication jobs

If a replication job fails due to an error, SyncIQ might disable the corresponding replication policy. For example SyncIQ might disable a replication policy if the IP or hostname of the target cluster is modified. If a replication policy is disabled, the policy cannot be run.

To resume replication for a disabled policy, you must either fix the error that caused the policy to be disabled, or reset the replication policy. It is recommended that you attempt to fix the issue rather than reset the policy. If you believe you have fixed the error, you can return the replication policy to an enabled state by resolving the policy. You can then run the policy again to test whether the issue was fixed. If you are unable to fix the issue, you can reset the replication policy. However, resetting the policy causes a full or differential replication to be performed the next time the policy is run.

Note

Depending on the amount of data being synchronized or copied, a full and differential replications can take a very long time to complete.

Resolve a replication policy

If SyncIQ disables a replication policy due to a replication error, and you fix the issue that caused the error, you can resolve the replication policy. Resolving a replication policy enables you to run the policy again. If you cannot resolve the issue that caused the error, you can reset the replication policy.

Procedure

1. Run the `isi sync policies resolve` command.

The following command resolves `weeklySync`:

```
isi sync policies resolve weeklySync
```

Reset a replication policy

If a replication job encounters an error that you cannot resolve, you can reset the corresponding replication policy. Resetting a policy causes OneFS to perform a full or differential replication the next time the policy is run.

Resetting a replication policy deletes the source-cluster snapshot.

Note

Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete. Reset a replication policy only if you cannot fix the issue that caused the replication error. If you fix the issue that caused the error, resolve the policy instead of resetting the policy.

Procedure

1. Run the `isi sync policies reset` command.

The following command resets `weeklySync`:

```
isi sync policies reset weeklySync
```

Perform a full or differential replication

After you reset a replication policy, you must perform either a full or differential replication.

Before you begin

Reset a replication policy.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in through the root or compliance administrator account.
2. Specify the type of replication you want to perform by running the `isi sync policies modify` command.

- To perform a full replication, disable the `--target-compare-initial-sync` option.

For example, the following command disables differential synchronization for `newPolicy`:

```
isi sync policies modify newPolicy \
--target-compare-initial-sync false
```

- To perform a differential replication, enable the `--target-compare-initial-sync` option.

For example, the following command enables differential synchronization for `newPolicy`:

```
isi sync policies modify newPolicy \
--target-compare-initial-sync true
```

3. Run the policy by running the `isi sync jobs start` command.

For example, the following command runs `newPolicy`:

```
isi sync jobs start newPolicy
```

Managing changelists

You can create and view changelists that describe what data was modified by a replication job. Changelists are most commonly accessed by applications through the OneFS Platform API.

To create a changelist, you must enable changelists for a replication policy. If changelists are enabled for a policy, SyncIQ does not automatically delete the repstate files generated by the policy; if changelists are not enabled for a policy, SyncIQ automatically deletes the repstate files after the corresponding replication jobs complete. SyncIQ generates one repstate file for each replication job. Because a large amount of repstate files can consume a large amount of disk space, it is recommended that you do not enable changelists for a policy unless it is necessary for your workflow.

If changelists are enabled for a policy, SyncIQ does not automatically delete source cluster snapshots for the policy. To create a changelist, you must have access to two consecutive snapshots and an associated repstate generated by a replication policy.

Create a changelist

You can create a changelist to view what data was modified by a replication job.

Before you begin

Enable changelists for a replication policy, and then run the policy at least twice. The following command enables changelists for newPolicy:

```
isi sync policies modify newPolicy --changelist true
```

Procedure

1. (Optional) Record the IDs of the snapshots generated by the replication policy.
 - a. View snapshot IDs by running the following command:

```
isi snapshot snapshots list
```

The snapshots must have been generated sequentially for a replication policy. If source-archival snapshots are not enabled for the policy, snapshots generated for the policy are named according to the following convention:

```
SIQ-Changelist-<policy-name>-<date>
```

If source-archival snapshots are enabled for the policy, snapshots are named according to the source-snapshot naming convention.

2. Create a changelist by running the `isi job jobs start` command with the `ChangelistCreate` option.

The following command creates a changelist:

```
isi job jobs start ChangelistCreate --older-snapid 2 --newer-snapid 6
```

You can also specify the `--retain-repstate` option to allow you to recreate the changelist later. If this option is not specified, the `repstate` file used to generate the changelist is deleted after the changelist is created.

View a changelist

You can view a changelist that describes what data was modified by a replication job. This procedure is available only through the command-line interface (CLI).

Procedure

1. View the IDs of changelists by running the following command:

```
isi_changelist_mod -l
```

Changelist IDs include the IDs of both snapshots used to create the changelist. If OneFS is still in the process of creating a changelist, `inprog` is appended to the changelist ID.

2. (Optional) View all contents of a changelist by running the `isi_changelist_mod` command with the `-a` option.

The following command displays the contents of a changelist named 2_6:

```
isi_changelist_mod -a 2_6
```

3. View a specific changelist entry by running the `isi_changelist_mod` command with the `-g` option.

The following command displays an entry with a LIN of 1003402c3 from a changelist named 2_6:

```
isi_changelist_mod -g 2_6 1003402c3
```

Changelist information

You can view the information contained in changelists.

Note

The information contained in changelists is meant to be consumed by applications through the OneFS Platform API. The information might be less useful when consumed through the command-line interface (CLI).

The following information is displayed in the output of the `isi_changelist_mod` command:

lin

The LIN of the changelist entry. Metadata entries are assigned a LIN of 1.

entry_type

The type of changelist entry. The field is set to either `metadata` or `file`.

size

The total size of the changelist entry, in bytes.

reserved

This field is not currently used by changelists.

root_path

The root path of the snapshots used to create the changelist.

owning_job_id

The ID of the ChangelistCreate job that created the changelist.

num_cl_entries

The number of changelist entries in the changelist.

root_path_size

The total size of the null-terminated UTF-8 string that contains the root path of the snapshots, in bytes.

root_path_offset

The number of bytes between the start of the changelist entry structure and the null-terminated UTF-8 string that contains the root path of the snapshots.

path

The path, relative to the root path, of the file or directory that was modified or removed.

type

If an item was modified, describes the type of item that was modified. The following types of items might have been modified:

regular

A regular file was modified.

directory

A directory was modified.

symlink

A symbolic link was modified.

fifo

A first-in-first-out (FIFO) queue was modified.

socket

A Unix domain socket was modified.

char device

A character device was modified.

block device

A block device was modified.

unknown

An unknown type of file was modified.

If any type of item was removed, this field is set to (REMOVED).

size

The size of the item that was modified, in bytes. If an item was removed, this field is set to 0.

path_size

The total size of the null-terminated UTF-8 string that contains the path, relative to the root path, of the file or directory that was modified or removed, in bytes.

path_offset

The number of bytes between the start of the changelist entry structure and the path, relative to the root path, of the file or directory that was modified or removed.

atime

The POSIX timestamp of when the item was last accessed.

atimensec

The number of nanoseconds past the atime that the item was last accessed.

ctime

The POSIX timestamp of when the item was last changed.

ctimensec

The number of nanoseconds past the ctime that the item was last changed.

mtime

The POSIX timestamp of when the item was last modified.

mtimensec

The number of nanoseconds past the mtime that the item was last modified.

Data replication commands

You can control data replication to other Isilon clusters through the data replication commands. Data replication commands apply specifically to the SyncIQ software module and are available only if you activate a SyncIQ license.

isi sync policies create

Creates a replication policy.

Syntax

```
isi sync policies create <name> <action>
<source-root-path> <target-host> <target-path>
[--description <string>]
[--password <password> | --set-password]
[--source-include-directories <string>]...
[--source-exclude-directories <string>]...
[--begin-filter {<predicate> <operator> <link>}]... --end-filter]
[--schedule {<schedule> | when-source-modified}]
[--enabled {true | false}]
[--check-integrity {true | false}]
[--log-level <level>]
[--log-removed-files {yes | no}]
[--workers-per-node <integer>]
[--target-snapshot-archive {on | off}]
[--target-snapshot-pattern <naming-pattern>]
[--target-snapshot-expiration <duration>]
[--target-snapshot-alias <naming-pattern>]
[--target-detect-modifications {on | off}]
[--source-snapshot-archive {on | off}]
[--source-snapshot-pattern <naming-pattern>]
[--source-snapshot-expiration <duration>]
[--report-max-age <duration>]
[--report-max-count <integer>]
[--resolve {enable | disable}]
[--restrict-target-network {on | off}]
[--source-subnet <subnet> --source-pool <pool>]
[--target-compare-initial-sync {on | off}]
[--verbose]
```

Options

<name>

Specifies a name for the replication policy.
Specify as any string.

<action>

Specifies the type of replication policy.
The following types of replication policy are valid:

copy

Creates a copy policy that adds copies of all files from the source to the target.

sync

Creates a synchronization policy that synchronizes data on the source cluster to the target cluster and deletes all files on the target cluster that are not present on the source cluster.

<source-root-path>

Specifies the directory on the local cluster that files are replicated from.
Specify as a full directory path.

<target-host>

Specifies the cluster that the policy replicates data to.
Specify as one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.
- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**
This will replicate data to another directory on the local cluster.

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

<target-path>

Specifies the directory on the target cluster that files are replicated to.
Specify as a full directory path.

--description <string>

Specifies a description of the replication policy.

--password <password>

Specifies a password to access the target cluster. If the target cluster requires a password for authentication purposes, you must specify this parameter or --set-password.

--set-password

Prompts you to specify a password for the target cluster after the command is run. This can be useful if you do not want other users on the cluster to see the password you specify. If the target cluster requires a password for authentication purposes, you must specify this parameter or --password.

{--source-include-directories | -i} <path>

Includes only the specified directories in replication.
Specify as any directory path contained in the root directory. You can specify multiple directories by specifying --source-include-directories multiple times within a command. For example, if the root directory is /ifs/data, you could specify the following:

```
--source-include-directories /ifs/data/music --source-include-directories /ifs/data/movies
```

{--source-exclude-directories | -e} <path>

Does not include the specified directories in replication. Specify as any directory path contained in the root directory. If --source-include-directories is specified, --source-exclude-directories directories must be contained in the included directories. You can specify multiple directories by specifying --source-

`exclude-directories` multiple times within a command. For example, you could specify the following:

```
--source-exclude-directories /ifs/data/music --source-exclude-directories /ifs/data/movies --exclude /ifs/data/music/working
```

```
--begin-filter {<predicate> <operator> <link>}... --end-filter
```

Specifies the file-matching criteria that determines which files are replicated. Files that do not match the file-matching criteria are not replicated. A file matching criterion consists of a predicate, an operator, and a link. The predicate specifies an attribute to filter by (for example, the size of a file). The following predicates are valid:

```
--size <integer>[{B | KB | MB | GB | TB | PB}]
```

Selects files according to the specified size.

```
--file-type <value>
```

Selects only the specified file-system object type.

The following values are valid:

f

Specifies regular files

d

Specifies directories

l

Specifies soft links

```
--name <value>
```

Selects only files whose names match the specified string.

You can include the following wildcards:

- *
- []
- ?

```
--accessed-after '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'
```

Selects files that have been accessed since the specified time. This predicate is valid only for copy policies.

```
--accessed-before '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'
```

Selects files that have not been accessed since the specified time. This predicate is valid only for copy policies.

```
--accessed-time '<integer> {days | weeks | months | years} ago'
```

Selects files that were accessed during the specified time interval. This predicate is valid only for copy policies.

```
--birth-after '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'
```

Selects files that were created after the specified time. This predicate is valid only for copy policies.

```
--birth-before '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'
```

Selects files that were created before the specified time. This predicate is valid only for copy policies.

```
--birth-time '<integer>{days | weeks | months | years} ago'
```

Selects files that were created during the specified time interval. This predicate is valid only for copy policies.

```
--changed-after '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that have been modified since the specified time. This predicate is valid only for copy policies.

```
--changed-before '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that have not been modified since the specified time. This predicate is valid only for copy policies.

```
--changed-time '<integer>{days | weeks | months | years} ago'
```

Selects files that were modified during the specified time interval. This predicate is valid only for copy policies.

```
--no-group
```

Selects files based on whether they are owned by a group.

```
--no-user
```

Selects files based on whether they are owned by a user.

```
--posix-regex-name <value>
```

Selects only files whose names match the specified POSIX regular expression. IEEE Std 1003.2 (POSIX.2) regular expressions are supported.

```
--user-id <id>
```

Selects files based on whether they are owned by the user of the specified ID.

```
--user-name <name>
```

Selects files based on whether they are owned by the user of the specified name.

```
--group-id <id>
```

Selects files based on whether they are owned by the group of the specified ID.

```
--group-name <name>
```

Selects files based on whether they are owned by the group of the specified name.

The operator specifies which files are selected in relationship to the attribute (for example, all files smaller than the given size). Specify operators in the following form:

```
--operator <value>
```

The following operator values are valid:

Value	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to

Value	Description
not	Not

The link specifies how the criterion relates to the one that follows it (for example, the file is selected only if it meets both criteria). The following links are valid:

`--and`

Selects files that meet the criteria of the options that come before and after this value.

`--or`

Selects files that meet either the criterion of the option that comes before this value or the criterion of the option that follows this value.

`{--schedule | -S} {<schedule> | when-source-modified}`

Specifies how often data will be replicated. Specifying `when-source-modified` causes OneFS to replicate data every time that the source directory of the policy is modified.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`
- `Every [{other | <integer>}] week [on <day>]`
- `Every [{other | <integer>}] month [on the <integer>]`
- `Every [<day>[, ...] [of every [{other | <integer>}] week]]`
- `The last {day | weekday | <day>} of every [{other | <integer>}] month`
- `The <integer> {weekday | <day>} of every [{other | <integer>}] month`
- `Yearly on <month> <integer>`
- `Yearly on the {last | <integer>} [weekday | <day>] of <month>`

Specify *<frequency>* in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`--enabled {true | false}`

Determines whether the policy is enabled or disabled.
The default value is `true`.

`--check-integrity {true | false}`

Specifies whether to perform a checksum on each file data packet that is affected by the SyncIQ job. If this option is set to `true`, and the checksum values do not match, SyncIQ retransmits the file data packet.

The default value is `true`.

`--log-level <level>`

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- `fatal`
- `error`
- `notice`
- `info`
- `copy`
- `debug`
- `trace`

The default value is `info`.

`--log-removed-files {yes | no}`

Determines whether SyncIQ retains a log of all files that are deleted when a synchronization policy is run. This parameter has no effect for copy policies.

The default value is `no`.

`{--workers-per-node [-w] <integer>`

Specifies the number of workers per node that are generated by SyncIQ to perform each replication job for the policy.

The default value is `3`.

`--target-snapshot-archive {on | off}`

Determines whether archival snapshots are generated on the target cluster. If this option is set to `off`, SyncIQ will still maintain exactly one snapshot at a time on the target cluster to facilitate failback. You must activate a SnapshotIQ license on the target cluster to generate archival snapshots on the target cluster.

`--target-snapshot-pattern <naming-pattern>`

Specifies the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster.

The default naming pattern is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M
```

`--target-snapshot-expiration <duration>`

Specifies an expiration period for archival snapshots on the target cluster.

If this option is not specified, archival snapshots will remain indefinitely on the target cluster.

Specify in the following format:

```
<integer><units>
```

The following `<units>` are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D
Specifies days

H
Specifies hours

`--target-snapshot-alias < naming-pattern >`

Specifies a naming pattern for the most recent archival snapshot generated on the target cluster.

The default alias is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-latest
```

`--target-detect-modifications {on | off}`

Determines whether SyncIQ checks the target directory for modifications before replicating files.



Specifying `off` could result in data loss. It is recommended that you consult Isilon Technical Support before specifying `off`.

`--source-snapshot-archive {on | off}`

Determines whether archival snapshots are retained on the source cluster. If this option is set to `off`, SyncIQ will still maintain one snapshot at a time for the policy to facilitate replication.

`--source-snapshot-pattern < naming-pattern >`

Specifies a naming pattern for the most recent archival snapshot generated on the source cluster.

For example, the following pattern is valid:

```
SIQ-source-#{PolicyName}-%Y-%m-%d_%H-%M
```

`--source-snapshot-expiration < duration >`

Specifies an expiration period for archival snapshots retained on the source cluster. If this option is not specified, archival snapshots will exist indefinitely on the source cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y
Specifies years

M
Specifies months

W
Specifies weeks

D
Specifies days

H
Specifies hours

`--report-max-age <duration>`

Specifies how long replication reports are retained before they are automatically deleted by SyncIQ.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

`--report-max-count <integer>`

Specifies the maximum number of reports to retain for the replication policy.

`--resolve {enable | disable}`

Determines whether users can manually resolve the policy if the policy encounters an error and becomes unrunnable.

`--restrict-target-network {on | off}`

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

`--source-subnet <subnet>`

Restricts replication jobs to running only on nodes in the specified subnet on the local cluster.

`--source-pool <pool>`

Restricts replication jobs to running only on nodes in the specified pool on the local cluster.

`--target-compare-initial-sync {on | off}`

Determines whether the full or differential replications are performed for this policy. Full or differential replications are performed the first time a policy is run and after a policy has been reset. If set to `on`, performs a differential replication. If set to `off`, performs a full replication.

If differential replication is enabled the first time a replication policy is run, the policy will run slower without any benefit.

The default value is `off`.

`{--verbose | -v}`

Displays a message confirming that the snapshot schedule was created.

isi sync policies modify

Modifies existing replication policies.

Syntax

```
isi sync policies modify <policy>
  {--name <new-policy-name>
  | --action <policy-type>
  | --target-host <target-cluster>
  | --target-path <target-path>
  | --source-root-path <root-path>
  | --description <string>
  | --password <password>
  | --set-password
  | --source-include-directories <string>
  | --clear-source-include-directories
  | --add-source-include-directories <string>
  | --remove-source-include-directories <string>
  | --source-exclude-directories <string>
  | --clear-source-exclude-directories
  | --add-source-exclude-directories <string>
  | --remove-source-exclude-directories <string>
  | --begin-filter <predicate> --operator <value>
  [<predicate> --operator <operator> <link>]...
  --end-filter
  | --schedule {<schedule> | when-source-modified}
  | --enabled {true | false}
  | --check-integrity {true | false}
  | --log-level <level>
  | --log-removed-files {yes | no}
  | --workers-per-node <integer>
  | --target-snapshot-archive {on | off}
  | --target-snapshot-pattern <naming-pattern>
  | --target-snapshot-expiration <duration>
  | --target-snapshot-alias <naming-pattern>
  | --target-detect-modifications {on | off}
  | --source-snapshot-archive {on | off}
  | --source-snapshot-pattern <naming-pattern>
  | --source-snapshot-expiration <duration>
  | --report-max-age <duration>
  | --report-max-count <integer>
  | --resolve {enable | disable}
  | --restrict-target-network {on | off}
  | --source-subnet <subnet> --source-pool <pool>
  | --clear-source-network
  | --target-compare-initial-sync {on | off}}...
  [--verbose]
  [--force]
```

Options

<policy>

Modifies the specified replication policy.
Specify as a replication policy name or ID.

{--name | -n} <new-policy-name>

Specifies a new name for this replication policy.

--action <policy-type>

Specifies the type of replication policy.
The following types of replication policy are valid:

copy

Creates a copy policy that adds copies of all files from the source to the target.

sync

Creates a synchronization policy that synchronizes data on the source cluster to the target cluster and deletes all files on the target cluster that are not present on the source cluster.

`{--target-host | -C} <target-cluster>`

Specifies the cluster that the policy replicates data to.
Specify as one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.
- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**
This will replicate data to another directory on the local cluster.

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

`{--target-path | -p} <target-path>`

Specifies the directory on the target cluster that files are replicated to.
Specify as a full directory path.

`--source-root-path <root-path>`

Specifies the directory on the local cluster that files are replicated from.
Specify as a full directory path.

`--description <string>`

Specifies a description of this replication policy.

`--password <password>`

Specifies a password to access the target cluster. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--set-password`.

`--set-password`

Prompts you to specify a password for the target cluster after the command is run. This can be useful if you do not want other users on the cluster to see the password you specify. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--password`.

`{--source-include-directories | -i} <path>`

Includes only the specified directories in replication.
Specify as any directory path contained in the root directory. You can specify multiple directories by specifying `--source-include-directories` multiple times within a command. For example, if the root directory is `/ifs/data`, you could specify the following:

```
--source-include-directories /ifs/data/music --source-include-directories /ifs/data/movies
```

`--clear-source-include-directories`

Clears the list of included directories.

```
--add-source-include-directories <path>
```

Adds the specified directory to the list of included directories.

```
--remove-source-include-directories <path>
```

Removes the specified directory from the list of included directories.

```
{--source-exclude-directories [-e] <path>
```

Does not include the specified directories in replication.

Specify as any directory path contained in the root directory. If `--source-include-directories` is specified, `--source-exclude-directories` directories must be contained in the included directories. You can specify multiple directories by specifying `--source-exclude-directories` multiple times within a command. For example, you could specify the following:

```
--source-exclude-directories /ifs/data/music --source-exclude-
directories /ifs/data/movies --exclude /ifs/data/music/working
```

```
--clear-source-exclude-directories
```

Clears the list of excluded directories.

```
--add-source-exclude-directories <path>
```

Adds the specified directory to the list of excluded directories.

```
--remove-source-exclude-directories <path>
```

Removes the specified directory from the list of excluded directories.

```
--begin-filter <predicate>--operator <value> [<predicate>--operator
<operator> <link>]... --end-filter
```

Specifies the file-matching criteria that determines which files are replicated. Specify `<predicate>` as one or more of the following options:

The following options are valid for both copy and synchronization policies:

```
--size <integer>[{B | KB | MB | GB | TB | PB}]
```

Selects files according to the specified size.

```
--file-type <value>
```

Selects only the specified file-system object type.

The following values are valid:

- f**
Specifies regular files
- d**
Specifies directories
- l**
Specifies soft links

```
--name <value>
```

Selects only files whose names match the specified string.

You can include the following wildcards:

- *
- []
- ?

The following options are valid only for copy policies:

```
--accessed-after '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that have been accessed since the specified time. This predicate is valid only for copy policies.

```
--accessed-before '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that have not been accessed since the specified time. This predicate is valid only for copy policies.

```
--accessed-time '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that were accessed at the specified time. This predicate is valid only for copy policies.

```
--birth-after '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that were created after the specified time. This predicate is valid only for copy policies.

```
--birth-before '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that were created before the specified time. This predicate is valid only for copy policies.

```
--birth-time '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that were created at the specified time. This predicate is valid only for copy policies.

```
--changed-after '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that have been modified since the specified time. This predicate is valid only for copy policies.

```
--changed-before '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that have not been modified since the specified time. This predicate is valid only for copy policies.

```
--changed-time '{<mm>/<dd>/<yyyy>[<HH>:<mm>] | <integer>{days | weeks | months | years} ago}'
```

Selects files that were modified at the specified time. This predicate is valid only for copy policies.

```
--no-group
```

Selects files based on whether they are owned by a group.

```
--no-user
```

Selects files based on whether they are owned by a user.

```
--posix-regex-name <value>
```

Selects only files whose names match the specified POSIX regular expression. IEEE Std 1003.2 (POSIX.2) regular expressions are supported.

```
--user-id <id>
```

Selects files based on whether they are owned by the user of the specified ID.

`--user-name <name>`

Selects files based on whether they are owned by the user of the specified name.

`--group-id <id>`

Selects files based on whether they are owned by the group of the specified ID.

`--group-name <name>`

Selects files based on whether they are owned by the group of the specified name.

The following *<operator>* values are valid:

Operator	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to
not	Not

You can use the following *<link>* values to combine and alter the options available for predicates:

`--and`

Selects files that meet the criteria of the options that come before and after this value.

`--or`

Selects files that meet either the criterion of the option that comes before this value or the criterion of the option that follows this value.

`{--schedule | -S} {<schedule> | when-source-modified}`

Specifies how often data will be replicated. Specifying `when-source-modified` causes OneFS to replicate data every time that the source directory of the policy is modified.

Specify *<schedule>* in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`
- `Every [{other | <integer>}] week [on <day>]`
- `Every [{other | <integer>}] month [on the <integer>]`
- `Every [<day>[, ...] [of every [{other | <integer>}] week]]`
- `The last {day | weekday | <day>} of every [{other | <integer>}] month`
- `The <integer> {weekday | <day>} of every [{other | <integer>}] month`
- `Yearly on <month> <integer>`
- `Yearly on the {last | <integer>} [weekday | <day>] of <month>`

Specify *<frequency>* in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

To configure a policy to be run only manually, specify the following option:

```
--schedule ""
```

```
--enabled {true | false}
```

Determines whether the policy is enabled or disabled.

```
--check-integrity {true | false}
```

Specifies whether to perform a checksum on each file data packet that is affected by the SyncIQ job. If this option is set to `true` and the checksum values do not match, SyncIQ retransmits the file data packet.

The default value is `true`.

```
--log-level <level>
```

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info
- copy
- debug
- trace

The default value is `info`.

```
--log-removed-files {yes | no}
```

Determines whether SyncIQ retains a log of all files that are deleted when a synchronization policy is run. If the policy is a copy policy, this parameter has no effect.

The default value is `no`.

```
{--workers-per-node | -w} <integer>
```

Specifies the number of workers per node that are generated by SyncIQ to perform each replication job for the policy.

The default value is 3.

```
--target-snapshot-archive {on | off}
```

Determines whether archival snapshots are generated on the target cluster. If this option is set to `off`, SyncIQ will still maintain exactly one snapshot at a time on the target cluster to facilitate failback. You must activate a SnapshotIQ license on the target cluster to generate archival snapshots on the target cluster.

```
--target-snapshot-pattern <naming-pattern>
```

Specifies the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster.

The default naming pattern is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M
```

`--target-snapshot-expiration <duration>`

Specifies an expiration period for archival snapshots on the target cluster.

If this option is not specified, archival snapshots will remain indefinitely on the target cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

`--target-snapshot-alias <naming-pattern>`

Specifies a naming pattern for the most recent archival snapshot generated on the target cluster.

The default alias is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-latest
```

`--target-detect-modifications {on | off}`

Determines whether SyncIQ checks the target directory for modifications before replicating files.



Specifying `off` could result in data loss. It is recommended that you consult Isilon Technical Support before specifying `off`.

`--source-snapshot-archive {on | off}`

Determines whether archival snapshots are retained on the source cluster. If this option is set to `off`, SyncIQ will still maintain one snapshot at a time for the policy to facilitate replication.

`--source-snapshot-pattern <naming-pattern>`

Specifies a naming pattern for the most recent archival snapshot generated on the source cluster.

For example, the following pattern is valid:

```
SIQ-source-%{PolicyName}-%Y-%m-%d_%H-%M
```

`--source-snapshot-expiration <duration>`

Specifies an expiration period for archival snapshots retained on the source cluster. If this option is not specified, archival snapshots will exist indefinitely on the source cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours

`--report-max-age <duration>`

Specifies how long replication reports are retained before they are automatically deleted by SyncIQ.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours

`--report-max-count <integer>`

Specifies the maximum number of reports to retain for the replication policy.

`--resolve {enable | disable}`

Determines whether users can manually resolve the policy if the policy encounters an error and becomes unrunnable.

`--restrict-target-network {on | off}`

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

`--source-subnet <subnet>`

Restricts replication jobs to running only on nodes in the specified subnet on the local cluster.

`--source-pool <pool>`

Restricts replication jobs to running only on nodes in the specified pool on the local cluster.

`--clear-source-network`

Runs replication jobs on any nodes in the cluster, instead of restricting the jobs to a specified subnet.

`--target-compare-initial-sync {on | off}`

Determines whether the full or differential replications are performed for this policy. Full or differential replications are performed the first time a policy is run and after a policy has been reset. If set to `on`, performs a differential replication. If set to `off`, performs a full replication.

If differential replication is enabled the first time a replication policy is run, the policy will run slower without any benefit.

The default value is `off`.

`{--verbose | -v}`

Displays a confirmation message.

`{--force | -f}`

Does not prompt you to confirm modifications.

isi sync policies delete

Deletes a replication policy.

The command will not succeed until SyncIQ can communicate with the target cluster; until then, the policy will still appear in the output of the `isi sync policies list` command. After the connection between the source cluster and target cluster is reestablished, SyncIQ will delete the policy the next time that the job is scheduled to run; if the policy is configured to run only manually, you must manually run the policy again. If SyncIQ is permanently unable to communicate with the target cluster, specify the `--local-only` option. This will delete the policy from the local cluster only and not break the target association on the target cluster.

Syntax

```
isi sync policies delete {<policy> | --all}
  [--local-only]
  [--force]
  [--verbose]
```

Options

`<policy>`

Deletes the specified replication policy.

`--all`

Deletes all replication policies.

`--local-only`

Does not break the target association on the target cluster. Not deleting a policy association on the target cluster will cause the target directory to remain in a read-only state.

Note

If SyncIQ is unable to communicate with the target cluster, you must specify this option to successfully delete the policy.

`{--force | -f}`

Deletes the policy, even if an associated job is currently running. Also, does not prompt you to confirm the deletion.

CAUTION

Forcing a policy to delete might cause errors if an associated replication job is currently running.

`{--verbose | -v}`

Displays a confirmation message.

isi sync policies list

Displays a list of replication policies.

Syntax

```
isi sync policies list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

If no options are specified, displays a table of all replication policies.

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.

The following values are valid:

name

Sorts output by the name of the replication policy.

target_path

Sorts output by the path of the target directory.

action

Sorts output by the type of replication policy.

description

Sorts output by the policy description.

enabled

Sorts output by whether the policies are enabled or disabled.

- target_host**
Sorts output by the target cluster.
- check_integrity**
Sorts output by whether the policy is configured to perform a checksum on each file data packet that is affected by a replication job.
- source_root_path**
Sorts output by the path of the source directory.
- source_include_directories**
Sorts output by directories that have been explicitly included in replication.
- source_exclude_directories**
Sorts output by directories that have been explicitly excluded in replication.
- file_matching_pattern**
Sorts output by the predicate that determines which files are replicated.
- target_snapshot_archive**
Sorts output by whether archival snapshots are generated on the target cluster.
- target_snapshot_pattern**
Sorts output by the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster.
- target_snapshot_expiration**
Sorts output by the expiration period for archival snapshots on the target cluster.
- target_detect_modifications**
Sorts output by whether full or differential replications are performed for this policy.
- source_snapshot_archive**
Sorts output by whether archival snapshots are retained on the source cluster.
- source_snapshot_pattern**
Sorts output by the naming pattern for the most recent archival snapshot generated on the source cluster.
- source_snapshot_expiration**
Sorts output by the expiration period for archival snapshots retained on the source cluster.
- schedule**
Sorts output by the schedule of the policy.
- log_level**
Sorts output by the amount of data that is recorded in logs.
- log_removed_files**
Sorts output by whether OneFS retains a log of all files that are deleted when the replication policy is run.
- workers_per_node**
Sorts output by the number of workers per node that are generated by OneFS to perform each replication job for the policy.
- report_max_age**
Sorts output by how long replication reports are retained before they are automatically deleted by OneFS
- report_max_count**
Sorts output by the maximum number of reports that are retained for the replication policy.

force_interface

Sorts output by whether data is sent over only the default interface of the subnet specified by the `--source-network` option of the `isi sync policies create` or `isi sync policies modify` commands.

restrict_target_network

Sorts output by whether replication jobs are restricted to connecting to nodes in a specified zone on the target cluster.

target_compare_initial_sync

Sorts output by whether full or differential replications are performed for the policies.

last_success

Sorts output by the last time that a replication job completed successfully.

password_set

Sorts output by whether the policy specifies a password for the target cluster.

source_network

Sorts output by the subnet on the local cluster that the replication policy is restricted to.

source_interface

Sorts output by the pool on the local cluster that the replication policy is restricted to.

`{--descending | -d}`

Displays output in reverse order.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi sync policies view

Displays information about a replication policy.

Syntax

```
isi sync policies view <policy>
```

Options

`<policy>`

Displays information about the specified replication policy. Specify as a replication policy name or ID.

isi sync policies disable

Temporarily disables a replication policy. If a replication policy is disabled, the policy will not create replication jobs. However, if a replication job is currently running for a replication policy, disabling the policy will not pause or stop the job.

Syntax

```
isi sync policies disable {<policy> | --all}
  [--verbose]
```

Options

<policy>

Disables the specified replication policy. Specify as a replication policy name or a replication policy ID.

--all

Disables all replication policies on the cluster.

--verbose

Displays more detailed information.

isi sync policies enable

Enables a disabled replication policy.

Syntax

```
isi sync policies enable {<policy> | --all}
  [--verbose]
```

Options

<policy>

Enables the specified replication policy. Specify as a replication policy name or a replication policy ID.

--all

Enables all replication policies on the cluster.

--verbose

Displays more detailed information.

isi sync jobs start

Starts a replication job for a replication policy.

Syntax

```
isi sync jobs start <policy-name>
  [--test]
  [--source-snapshot <snapshot>]
  [--verbose]
```

Options

<policy-name>

Starts a replication job for the specified replication policy.

--test

Creates a replication policy report that reflects the number of files and directories that would be replicated if the specified policy was run. You can test only policies that have not been run before.

`--source-snapshot <snapshot>`

Replicates data according to the specified SnapshotIQ snapshot. If specified, a snapshot is not generated for the replication job. Replicating data according to snapshots generated by the SyncIQ tool is not supported.

Specify as a snapshot name or ID. The source directory of the policy must be contained in the specified snapshot. This option is valid only if the last replication job completed successfully or if you are performing a full or differential replication. If the last replication job completed successfully, the specified snapshot must be more recent than the snapshot referenced by the last replication job.

`{--verbose | -v}`

Displays more detailed information.

isi sync jobs pause

Pauses a running replication job.

Syntax

```
isi sync jobs pause {<policy-name> | --all}
  [--verbose]
```

Options

`<policy-name>`

Pauses a job that was created according to the specified replication policy. Specify as a replication policy name.

`--all`

Pauses all currently running replication jobs.

`{--verbose | -v}`

Displays more detailed information.

isi sync jobs resume

Resumes paused replication jobs.

Syntax

```
isi sync jobs resume {<policy-name> | --all}
  [--verbose]
```

Options

`<policy-name>`

Resumes a paused job that was created by the specified policy. Specify as a replication policy name.

`--all`

Resumes all currently running replication jobs.

`{--verbose | -v}`

Displays more detailed information.

isi sync jobs cancel

Cancels a running or paused replication job.

Syntax

```
isi sync jobs cancel {<policy-name> | --all}
  [--verbose]
```

Options

<policy-name>

Cancels a job that was created according to the specified replication policy. Specify as a replication policy name or ID.

--all

Cancels all currently running replication jobs.

--verbose

Displays more detailed information.

isi sync jobs list

Displays information about the most recently completed and next scheduled replication jobs of replication policies.

Syntax

```
isi sync jobs list
  [--state <state>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

If no options are specified, displays information about replication jobs for all policies.

--state <state>

Displays only jobs in the specified state.

The following values are valid:

scheduled

Displays jobs that are scheduled to run.

running

Displays running jobs.

paused

Displays jobs that were paused by a user.

finished

Displays jobs that have completed successfully.

failed

Displays jobs that failed during the replication process.

canceled

Displays jobs that were cancelled by a user.

needs_attention

Displays jobs that require user intervention before they can continue.

```
--format {table | json | csv | list}
    Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
{--no-header | -a}
    Displays table and CSV output without headers.
{--no-footer | -z}
    Displays table output without footers.
{--verbose | -v}
    Displays more detailed information.
```

isi sync jobs view

Displays information about a running replication job.

Syntax

```
isi sync jobs view <policy>
```

Options

```
<policy>
    Displays information about a running replication job created according to the specified policy.
    Specify as a replication policy name or ID.
```

isi sync jobs reports list

Displays information about running replication jobs targeting the local cluster.

Syntax

```
isi sync jobs reports list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
--limit | -l} <integer>
    Displays no more than the specified number of items.
--format {table | json | csv | list}
    Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
{--no-header | -a}
    Displays table and CSV output without headers.
{--no-footer | -z}
    Displays table output without footers.
{--verbose | -v}
    Displays more detailed information.
```

isi sync jobs reports view

Displays information about a running replication job targeting the local cluster.

Syntax

```
isi sync jobs reports view <policy>
```

Options

<policy>

Displays information about a replication job created according to the specified replication policy.

Specify as a replication policy name or ID.

isi sync settings modify

Manages global replication settings.

Syntax

```
isi sync settings modify
[--service {on | off | paused}]
[--source-subnet <subnet>]
[--source-pool <pool>]
[--restrict-target-network {on | off}]
[--report-max-age <duration>]
[--report-max-count <integer>]
[--report-email <email-address>]
[--clear-report-email]
[--add-report-email <email-address>]
[--remove-report-email <email-address>]
[--verbose]
```

Options

If no options are specified, displays current default replication report settings.

`--service {on | off | paused}`

Determines the state of the SyncIQ tool.

`--source-subnet <subnet>`

Restricts replication jobs to running only on nodes in the specified subnet on the local cluster.

`--source-pool <pool>`

Restricts replication jobs to running only on nodes in the specified pool on the local cluster.

`--restrict-target-network {on | off}`

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

Note

SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

```
--report-max-age <duration>
```

Specifies the default amount of time that SyncIQ retains reports before automatically deleting them.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y

Specifies years

M

Specifies months

D

Specifies days

H

Specifies hours

```
--report-max-count <integer>
```

Specifies the default maximum number of reports to retain for a replication policy.

```
{--verbose | -v}
```

Displays more detailed information.

isi sync settings view

Displays global replication settings.

Syntax

```
isi sync settings view
```

Options

There are no options for this command.

isi sync policies resolve

Resolves a conflicted replication policy after the policy encounters an error and the cause of the error is fixed. If the cause of the error cannot be fixed, run the `isi sync policies reset` command instead.

Syntax

```
isi sync policies resolve <policy>
  [--force]
```

Options

```
<policy>
```

Resolves the specified replication policy.

Specify as a replication policy name or ID.

```
{--force | -f}
```

Suppresses command-line prompts and messages.

isi sync policies reset

Resets a replication policy after the policy encounters an error and the cause of the error cannot be identified or fixed. If you fix the cause of the error, run `isi sync policies resolve` instead.

Resetting a replication policy causes either a full replication or a differential replication to be performed the next time the policy is run.

Syntax

```
isi sync policies reset {<policy> | --all}
  [--verbose]
```

Options

<policy>

Resets the specified replication policy.
Specify as a replication policy name or ID

--all

Resets all replication policies

{--verbose | -v}

Displays more detailed information.

isi sync target cancel

Cancels running replication jobs targeting the local cluster.

Syntax

```
isi sync target cancel {<policy> | --target-path <path> | --all}
  [--verbose]
```

Options

<policy>

Cancels a replication job created according to the specified replication policy.
Specify as a replication policy name or ID.

--target-path <path>

Cancels a replication job targeting the specified directory.

--all

Cancels all running replication jobs targeting the local cluster.

--verbose

Displays more detailed information.

isi sync target list

Displays a list of replication policies targeting the local cluster.

Syntax

```
isi sync target list
  [--target-path <path>]
  [--limit <integer>]
```

```
[--sort <attribute>]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

If no options are specified, displays a table of all replication policies currently targeting the local cluster.

`--target-path <path>`

Displays information about the replication policy targeting the specified directory.

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.
The following values are valid:

name

Sorts output by the name of the replication policy.

source_host

Sorts output by the name of the source cluster.

target_path

Sorts output by the path of the target directory.

last_job_status

Sorts output by the status of the last replication job created according to the policy.

failover_failback_state

Sorts output by whether the target directory is read only.

`{--descending | -d}`

Displays output in reverse order.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi sync target view

Displays information about a replication policy that is targeting the local cluster.

Syntax

```
isi sync target view {<policy-name> | --target-path <path>}
```

Options`<policy-name>`

Displays information about the specified policy.

`--target-path <path>`

Displays information about the policy targeting the specified directory.

isi sync target break

Breaks the association between a local cluster and a target cluster for a replication policy.

Note

Breaking a source and target association requires you to reset the replication policy before you can run the policy again. Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete.

Syntax

```
isi sync target break {<policy> | --target-path <path>}
  [--force]
  [--verbose]
```

Options`<policy>`

Removes the association of the specified replication policy targeting this cluster. Specify as a replication policy name, a replication policy ID, or the path of a target directory.

`--target-path <path>`

Removes the association of the replication policy targeting the specified directory path.

`{--force | -f}`

Forces the replication policy association to be removed, even if an associated job is currently running.



Forcing a target break might cause errors if an associated replication job is currently running.

`{--verbose | -v}`

Displays more detailed information.

isi sync target reports list

Displays information about completed replication jobs targeting the local cluster.

Syntax

```
isi sync target reports list
  [--state <state>]
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
```



```
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

If no options are specified, displays basic information about all completed replication jobs.

`--state <state>`

Displays information about only replication jobs in the specified state. The following states are valid:

- scheduled
- running
- paused
- finished
- failed
- canceled
- needs_attention
- unknown

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.

The following values are valid:

`start_time`

Sorts output by when the replication job started.

`end_time`

Sorts output by when the replication job ended.

`action`

Sorts output by the action that the replication job performed.

`state`

Sorts output by the progress of the replication job.

`id`

Sorts output by the ID of the replication subreport.

`policy_id`

Sorts output by the ID of the replication policy

`policy_name`

Sorts output by the name of the replication policy.

`job_id`

Sorts output by the ID of the replication job.

`total_files`

Sorts output by the total number of files that were modified by the replication job.

`files_transferred`

Sorts output by the total number of files that were transferred to the target cluster.

```

bytes_transferred
Sorts output by the total number of files that were transferred to the target cluster.
duration
Sorts output by how long the replication job ran.
errors
Sorts output by errors that the replication job encountered.
warnings
Sorts output by warnings that the replication job triggered.
{--descending | -d}
    Displays output in reverse order.
--format {table | json | csv | list}
    Displays output in table (default), JavaScript Object Notation (JSON), comma-
    separated value (CSV), or list format.
{--no-header | -a}
    Displays table and CSV output without headers.
{--no-footer | -z}
    Displays table output without footers.
{--verbose | -v}
    Displays more detailed information.

```

isi sync target reports view

Displays information about a completed replication job that targeted the local cluster.

Syntax

```
isi sync target reports view <policy> <job-id>
```

Options

<policy>

Displays a replication report about the specified replication policy.

<job-id>

Displays a replication report about the job with the specified ID.

isi sync target reports subreports list

Displays subreports about completed replication jobs targeting the local cluster.

Syntax

```

isi sync target reports subreports list <policy> <job-id>
  [--limit]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]

```

Options

<policy>

Displays subreports about the specified policy.

`<job-id>`

Displays subreports about the job of the specified ID.

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.
The following values are valid:

`start_time`
Sorts output by when the replication job started.

`end_time`
Sorts output by when the replication job ended.

`action`
Sorts output by the action that the replication job performed.

`state`
Sorts output by the progress of the replication job.

`id`
Sorts output by the ID of the replication report.

`policy_id`
Sorts output by the ID of the replication policy

`policy_name`
Sorts output by the name of the replication policy.

`job_id`
Sorts output by the ID of the replication job.

`total_files`
Sorts output by the total number of files that were modified by the replication job.

`files_transferred`
Sorts output by the total number of files that were transferred to the target cluster.

`bytes_transferred`
Sorts output by the total number of files that were transferred to the target cluster.

`duration`
Sorts output by how long the replication job ran.

`errors`
Sorts output by errors that the replication job encountered.

`warnings`
Sorts output by warnings that the replication job triggered.

`{--descending | -d}`

Displays output in reverse order.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi sync target reports subreports view

Displays a subreport about a completed replication job targeting the local cluster.

Syntax

```
isi sync target reports subreports view <policy> <job-id> <subreport-id>
```

Options

<policy>

Displays a sub report about the specified replication policy. Specify as a replication policy name.

<job-id>

Displays a sub report about the specified replication job. Specify as a replication job ID.

<subreport-id>

Displays the subreport with the specified ID.

isi sync reports list

Displays information about completed replication jobs targeting a remote cluster.

Syntax

```
isi sync reports list
  [--policy-name <policy>]
  [--state <state>]
  [--reports-per-policy <integer>]
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--policy-name <policy>

Displays only replication reports that were created for the specified policy.

--state <state>

Displays only replication reports whose jobs are in the specified state.

--reports-per-policy <integer>

Displays no more than the specified number of reports per policy. The default value is 10.

{--limit | -l} <integer>

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.
The following values are valid:

`start_time`
Sorts output by when the replication job started.

`end_time`
Sorts output by when the replication job ended.

`action`
Sorts output by the action that the replication job performed.

`state`
Sorts output by the progress of the replication job.

`id`
Sorts output by the ID of the replication subreport.

`policy_id`
Sorts output by the ID of the replication policy

`policy_name`
Sorts output by the name of the replication policy.

`job_id`
Sorts output by the ID of the replication job.

`total_files`
Sorts output by the total number of files that were modified by the replication job.

`files_transferred`
Sorts output by the total number of files that were transferred to the target cluster.

`bytes_transferred`
Sorts output by the total number of bytes that were transferred to the target cluster.

`duration`
Sorts output by how long the replication job ran.

`errors`
Sorts output by errors that the replication job encountered.

`warnings`
Sorts output by warnings that the replication job triggered.

`{--descending | -d}`
Displays output in reverse order.

`--format {table | json | csv | list}`
Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`
Displays table and CSV output without headers.

`{--no-footer | -z}`
Displays table output without footers.

`{--verbose | -v}`
Displays more detailed information.

isi sync reports view

Displays information about a completed replication job that targeted a remote cluster.

Syntax

```
isi sync reports view <policy> <job-id>
```

Options

<policy>

Displays a replication report about the specified replication policy.

<job-id>

Displays a replication report about the job with the specified ID.

isi sync reports rotate

If the number of replication reports has exceeded the maximum, deletes replication reports. The system intermittently deletes excess reports automatically. However, this command causes excess reports to be deleted immediately.

Syntax

```
isi sync reports rotate
  [--verbose]
```

Options

{--verbose | -v}

Displays more detailed information.

isi sync reports subreports list

Displays subreports about completed replication jobs targeting remote clusters.

Syntax

```
isi sync reports subreports list <policy> <job-id>
  [--limit]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<policy>

Displays subreports about the specified policy.

<job-id>

Displays subreports about the job of the specified ID.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

`start_time`

Sorts output by when the replication job started.

`end_time`

Sorts output by when the replication job ended.

`action`

Sorts output by the action that the replication job performed.

`state`

Sorts output by the progress of the replication job.

`id`

Sorts output by the ID of the replication report.

`policy_id`

Sorts output by the ID of the replication policy

`policy_name`

Sorts output by the name of the replication policy.

`job_id`

Sorts output by the ID of the replication job.

`total_files`

Sorts output by the total number of files that were modified by the replication job.

`files_transferred`

Sorts output by the total number of files that were transferred to the target cluster.

`bytes_transferred`

Sorts output by the total number of files that were transferred to the target cluster.

`duration`

Sorts output by how long the replication job ran.

`errors`

Sorts output by errors that the replication job encountered.

`warnings`

Sorts output by warnings that the replication job triggered.

`{--descending | -d}`

Displays output in reverse order.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi sync reports subreports view

Displays a subreport about a completed replication job that targeted a remote cluster.

Syntax

```
isi sync reports subreports view <policy> <job-id> <subreport-id>
```

Options

<policy>

Displays a sub report about the specified replication policy. Specify as a replication policy name.

<job-id>

Displays a sub report about the specified replication job. Specify as a replication job ID.

<subreport-id>

Displays the subreport of the specified ID.

isi sync recovery allow-write

Allows modifications to data in a target directory of a replication policy without breaking the association between the local cluster and the policy. The `isi sync target allow_write` command is most commonly used in failover and failback operations.

Syntax

```
isi sync recovery allow-write <policy-name>
  [--revert]
  [--log-level <level>]
  [--workers-per-node <integer>]
  [--verbose]
```

Options

<policy-name>

Allows writes for the target directory of the specified replication policy. Specify as a replication policy name, a replication policy ID, or the path of a target directory.

--revert

Reverts an allow-writes operation on the local cluster only. This action does not affect the source cluster of the replication policy.

--log-level <level>

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info
- copy
- debug

- trace

The default value is `info`.

`{--workers-per-node | -w} <integer>`

Specifies the number of workers per node that are generated by SyncIQ to perform the allow-writes job.

The default value is 3.

`{--verbose | -v}`

Displays more detailed information.

isi sync recovery resync-prep

Disables the specified policy, reverts the source directory of the policy to the last recovery point, and creates a mirror policy on the target cluster. The `isi sync resync prep` command is most commonly used in failback operations.

Syntax

```
isi sync recovery resync-prep <policy-name>
  [--verbose]
```

Options

`<policy-name>`

Targets the following replication policy.

Specify as a replication policy name or ID. The replication policy must be a synchronization policy.

`--verbose`

Displays more detailed information.

isi sync rules create

Creates a replication performance rule.

Syntax

```
isi sync rules create <type> <interval> <days> <limit>
  [--description <string>]
  [--verbose]
```

Options

`<type>`

Specifies the type of performance rule. The following values are valid:

`file_count`

Creates a performance rule that limits the number of files that can be sent by replication jobs per second.

`bandwidth`

Creates a performance rule that limits the amount of bandwidth that replication jobs are allowed to consume.

`<interval>`

Enforces the performance rule on the specified hours of the day. Specify in the following format:

```
<hh>: <mm>-<hh>: <mm>
```

<days>

Enforces the performance rule on the specified days of the week.
The following values are valid:

X

Specifies Sunday

M

Specifies Monday

T

Specifies Tuesday

W

Specifies Wednesday

R

Specifies Thursday

F

Specifies Friday

S

Specifies Saturday

You can include multiple days by specifying multiple values separated by commas.
You can also include a range of days by specifying two values separated by a dash.

<limit>

Specifies the maximum number of files that can be sent or KBs that can be consumed per second by replication jobs.

--description <string>

Specifies a description of this performance rule.

--verbose

Displays more detailed information.

isi sync rules modify

Modifies a replication performance rule.

Syntax

```
isi sync rules modify <id>
  [--interval <interval>]
  [--days <days>]
  [--limit <integer>]
  [--enabled {true | false}]
  [--description <string>]
  [--verbose]
```

Options

<id>

Modifies the replication performance rule of the specified ID.

{--interval | -i} <interval>

Specifies which hours of the day to enforce the performance rule. Specify in the following format:

```
<hh>: <mm> - <hh>: <mm>
```

{--days | -d} <days>

Specifies which days of the week to enforce the performance rule.
The following values are valid:

- X** Specifies Sunday
- M** Specifies Monday
- T** Specifies Tuesday
- W** Specifies Wednesday
- R** Specifies Thursday
- F** Specifies Friday
- S** Specifies Saturday

You can include multiple days by specifying multiple values separated by commas.
You can also include a range of days by specifying two values separated by a dash.

`--limit <limit>`

Specifies the maximum number of files that can be sent or KBs that can be consumed per second by replication jobs.

`--enabled {true | false}`

Determines whether the policy is enabled or disabled.

`--description <string>`

Specifies a description of this performance rule.

`{--verbose | -v}`

Displays more detailed information.

isi sync rules delete

Deletes a replication performance rule.

Syntax

```
isi sync rules delete {<id> | --all | --type <type>}
  [--force]
  [--verbose]
```

Options

`<id>`

Deletes the performance rule of the specified ID.

`--all`

Deletes all performance rules.

`--type <type>`

Deletes all performance rules of the specified type. The following values are valid:

`file_count`

Deletes all performance rules that limit the number of files that can be sent by replication jobs per second.

bandwidth

Deletes all performance rules that limit the amount of bandwidth that replication jobs are allowed to consume.

`--force`

Does not prompt you to confirm that you want to delete the performance rule.

`--verbose`

Displays more detailed information.

isi sync rules list

Displays a list of replication performance rules.

Syntax

```
isi sync rules list
  [--type <type>]
  [--limit]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--type <type>`

Displays only performance rules of the specified type. The following values are valid:

`file_count`

Displays only performance rules that limit the number of files that can be sent by replication jobs per second.

bandwidth

Displays only performance rules that limit the amount of bandwidth that replication jobs are allowed to consume.

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi sync rules view

Displays information about a replication performance rule.

Syntax

```
isi sync rules view <id>
```

Options`<id>`

Displays information about the replication performance rule with the specified ID.

isi_changelist_mod

Displays, modifies, and creates changelists.

Syntax

```
isi_changelist_mod
[-l]
[-c <oldsnapshot> <newsnapshot> [{--root <string>
| --jobid <unsigned_integer>}]]
[-f <changelist>]
[-k <changelist>]
[-a <changelist> [--terse]]
[-g <changelist> <lin> [--terse]]
[-r <changelist> <low> <high> --terse]
[-s <changelist> <lin>
 [--path <string>]
 [--type <integer>]
 [--size <integer>]
 [--atime <integer>]
 [--atimensec <integer>]
 [--ctime <integer>]
 [--ctimensec <integer>]
 [--mtime <integer>]
 [--mtimensec <integer>]
 ]
[-d <changelist> <lin>]
```

Options`-l`

Displays a list of changelists.

`-c <oldsnapshot> <newsnapshot>`

Creates an empty changelist that you can manually populate for testing purposes. Specify `<oldsnapshot>` and `<newsnapshot>` as snapshot IDs.

`--root <string>`

Specifies the path of the directory the changelist is for.

`--jobid <unsigned_integer>`

Specifies the ID of the job that created the changelist.

`-f <changelist>`

Finalizes a manually-created changelist.

`-k <changelist>`

Deletes the specified changelist.

`-a <changelist>`

Displays the contents of the specified changelist.

`-g <changelist> <lin>`

Displays the specified changelist entry.

`-r <changelist> <low> <high>`

Displays the specified range of changelist entries.

`--terse`

Displays less detailed information.

`-s <changelist> <lin>`

Modifies the specified changelist entry.

`--path <string>`

Specifies the path of the file or directory that was modified or removed.

`--type <integer>`

If an item was modified, describes the type of item that was modified. The following values are valid:

1

Specifies that a regular file was modified.

2

Specifies that a directory was modified.

3

Specifies that a symbolic link was modified.

4

Specifies that a first-in-first-out (FIFO) queue was modified.

5

Specifies that a Unix domain socket was modified.

6

Specifies that a character device was modified.

7

Specifies that a block device was modified.

8

Specifies that an unknown type of file was modified.

To specify that any type of item was removed, specify 0.

`--size <integer>`

Specifies the size of the item that was modified, in bytes. If an item was removed, specify 0.

`--atime <integer>`

Specifies when the item was last accessed.
Specify as a POSIX timestamp.

`--atimensec <integer>`

Specifies the number of nanoseconds past the atime that the item was last accessed.

`--ctime <integer>`

Specifies when the item was last changed.
Specify as a POSIX timestamp.

`--ctimensec <integer>`

Specifies the number of nanoseconds past the ctime that the item was last changed.

`--mtime <integer>`

Specifies when the item was last modified.
Specify as a POSIX timestamp.

`--mtimensec <integer>`

Specifies the number of nanoseconds past the mtime that the item was last modified.

`-d <changelist> <lin>`

Deletes the specified entry from the changelist.

CHAPTER 14

Data layout with FlexProtect

This section contains the following topics:

- [FlexProtect overview](#).....610
- [File striping](#)..... 610
- [Requested data protection](#)..... 610
- [FlexProtect data recovery](#).....611
- [Requesting data protection](#)..... 612
- [Requested protection settings](#).....612
- [Requested protection disk space usage](#)..... 613

FlexProtect overview

An Isilon cluster is designed to continuously serve data, even when one or more components simultaneously fail. OneFS ensures data availability by striping or mirroring data across the cluster. If a cluster component fails, data stored on the failed component is available on another component. After a component failure, lost data is restored on healthy components by the FlexProtect proprietary system.

Data protection is specified at the file level, not the block level, enabling the system to recover data quickly. Because all data, metadata, and parity information is distributed across all nodes, the cluster does not require a dedicated parity node or drive. This ensures that no single node limits the speed of the rebuild process.

File striping

OneFS uses the internal network to automatically allocate and stripe data across nodes and disks in the cluster. OneFS protects data as the data is being written. No separate action is necessary to stripe data.

OneFS breaks files into smaller logical chunks called stripes before writing the files to disk; the size of each file chunk is referred to as the stripe unit size. Each OneFS block is 8 KB, and a stripe unit consists of 16 blocks, for a total of 128 KB per stripe unit. During a write, OneFS breaks data into stripes and then logically places the data in a stripe unit. As OneFS stripes data across the cluster, OneFS fills the stripe unit according to the number of nodes and protection level.

OneFS can continuously reallocate data and make storage space more usable and efficient. As the cluster size increases, OneFS stores large files more efficiently.

Requested data protection

The requested protection of data determines the amount of redundant data created on the cluster to ensure that data is protected against component failures. OneFS enables you to modify the requested protection in real time while clients are reading and writing data on the cluster.

OneFS provides several data protection settings. You can modify these protection settings at any time without rebooting or taking the cluster or file system offline. When planning your storage solution, keep in mind that increasing the requested protection reduces write performance and requires additional storage space for the increased number of nodes.

OneFS uses the Reed Solomon algorithm for N+M protection. In the N+M data protection model, N represents the number of data-stripe units, and M represents the number of simultaneous node or drive failures—or a combination of node and drive failures—that the cluster can withstand without incurring data loss. N must be larger than M.

In addition to N+M data protection, OneFS also supports data mirroring from 2x to 8x, allowing from two to eight mirrors of data. In terms of overall cluster performance and resource consumption, N+M protection is often more efficient than mirrored protection. However, because read and write performance is reduced for N+M protection, data mirroring might be faster for data that is updated often and is small in size. Data mirroring requires significant overhead and might not always be the best data-protection method. For example, if you enable 3x mirroring, the specified content is duplicated three times on the cluster; depending on the amount of content mirrored, this can consume a significant amount of storage space.

FlexProtect data recovery

OneFS uses the FlexProtect proprietary system to detect and repair files and directories that are in a degraded state due to node or drive failures.

OneFS protects data in the cluster based on the configured protection policy. OneFS rebuilds failed disks, uses free storage space across the entire cluster to further prevent data loss, monitors data, and migrates data off of at-risk components.

OneFS distributes all data and error-correction information across the cluster and ensures that all data remains intact and accessible even in the event of simultaneous component failures. Under normal operating conditions, all data on the cluster is protected against one or more failures of a node or drive. However, if a node or drive fails, the cluster protection status is considered to be in a degraded state until the data is protected by OneFS again. OneFS reprotects data by rebuilding data in the free space of the cluster. While the protection status is in a degraded state, data is more vulnerable to data loss.

Because data is rebuilt in the free space of the cluster, the cluster does not require a dedicated hot-spare node or drive in order to recover from a component failure. Because a certain amount of free space is required to rebuild data, it is recommended that you reserve adequate free space through the virtual hot spare feature.

As you add more nodes, the cluster gains more CPU, memory, and disks to use during recovery operations. As a cluster grows larger, data restriping operations become faster.

Smartfail

OneFS protects data stored on failing nodes or drives through a process called smartfailing.

During the smartfail process, OneFS places a device into quarantine. Data stored on quarantined devices is read only. While a device is quarantined, OneFS reprotects the data on the device by distributing the data to other devices. After all data migration is complete, OneFS logically removes the device from the cluster, the cluster logically changes its width to the new configuration, and the node or drive can be physically replaced.

OneFS smartfails devices only as a last resort. Although you can manually smartfail nodes or drives, it is recommended that you first consult Isilon Technical Support.

Occasionally a device might fail before OneFS detects a problem. If a drive fails without being smartfailed, OneFS automatically starts rebuilding the data to available free space on the cluster. However, because a node might recover from a failure, if a node fails, OneFS does not start rebuilding data unless the node is logically removed from the cluster.

Node failures

Because node loss is often a temporary issue, OneFS does not automatically start reprotecting data when a node fails or goes offline. If a node reboots, the file system does not need to be rebuilt because it remains intact during the temporary failure.

If you configure N+1 data protection on a cluster, and one node fails, all of the data is still accessible from every other node in the cluster. If the node comes back online, the node rejoins the cluster automatically without requiring a full rebuild.

To ensure that data remains protected, if you physically remove a node from the cluster, you must also logically remove the node from the cluster. After you logically remove a node, the node automatically reformats its own drives, and resets itself to the factory

default settings. The reset occurs only after OneFS has confirmed that all data has been reprotected. You can logically remove a node using the smartfail process. It is important that you smartfail nodes only when you want to permanently remove a node from the cluster.

If you remove a failed node before adding a new node, data stored on the failed node must be rebuilt in the free space in the cluster. After the new node is added, OneFS distributes the data to the new node. It is more efficient to add a replacement node to the cluster before failing the old node because OneFS can immediately use the replacement node to rebuild the data stored on the failed node.

Requesting data protection

You can request the protection of a file or directory by setting its requested protection. This flexibility enables you to protect distinct sets of data at different levels.

The default requested protection of node pools is N+2:1, which means that two drives or one node can fail without causing any data loss. For clusters or node pools containing less than two petabytes or fewer than 16 nodes, N+2:1 is the recommended requested protection. However, if the cluster or node pool is larger, you might consider higher requested protection.

OneFS allows you to request protection that the cluster is currently incapable of matching. If you request an unmatchable protection, the cluster will continue trying to match the requested protection until a match is possible. For example, in a four-node cluster, you might request a protection of 5x. In this example, OneFS would protect the data at 4x until you added a fifth node to the cluster, at which point OneFS would reprotect the data at the 5x.

Note

For 4U Isilon IQ X-Series and NL-Series nodes, and IQ 12000X/EX 12000 combination platforms, the minimum cluster size of three nodes requires a minimum of N+2:1.

Requested protection settings

Requested protection settings determine the level of hardware failure that a cluster can recover from without suffering data loss.

Requested protection setting	Minimum number of nodes required	Definition
[+1n]	3	The cluster can recover from one drive or node failure without sustaining any data loss.
[+2d:1n]	3	The cluster can recover from two simultaneous drive failures or one node failure without sustaining any data loss.
[+2n]	4	The cluster can recover from two simultaneous drive or node failures without sustaining any data loss.
[+3d:1n]	3	The cluster can recover from three simultaneous drive failures or one

Requested protection setting	Minimum number of nodes required	Definition
		node failure without sustaining any data loss.
[+3d:1n1d]	3	The cluster can recover from three simultaneous drive failures or simultaneous failures of one node and one drive without sustaining any data loss.
[+3n]	6	The cluster can recover from three simultaneous drive or node failures without sustaining any data loss.
[+4d:1n]	3	The cluster can recover from four simultaneous drive failures or one node failure without sustaining any data loss.
[+4d:2n]	4	The cluster can recover from four simultaneous drive failures or two node failures without sustaining any data loss.
[+4n]	8	The cluster can recover from four simultaneous drive or node failures without sustaining any data loss.
Nx (Data mirroring)	N For example, 5x requires a minimum of five nodes.	The cluster can recover from N - 1 drive or node failures without sustaining data loss. For example, 5x protection means that the cluster can recover from four drive or node failures.

Requested protection disk space usage

Increasing the requested protection of data also increases the amount of space consumed by the data on the cluster.

The parity overhead for N + M protection depends on the file size and the number of nodes in the cluster. The percentage of parity overhead declines as the cluster gets larger.

The following table describes the estimated amount of overhead depending on the requested protection and the size of the cluster or node pool. The table does not show recommended protection levels based on cluster size.

Number of nodes	[+1n]	[+2d:1n]	[+2n]	[+3d:1n]	[+3d:1n1d]	[+3n]	[+4d:1n]	[+4d:2n]	[+4n]
3	2 + 1 (33%)	4 + 2 (33%)	—	6 + 3 (33%)	3 + 3 (50%)	—	8 + 4 (33%)	—	—

Number of nodes	[+1n]	[+2d:1n]	[+2n]	[+3d:1n]	[+3d:1n1d]	[+3n]	[+4d:1n]	[+4d:2n]	[+4n]
4	3 + 1 (25%)	6 + 2 (25%)	—	9 + 3 (25%)	5 + 3 (38%)	—	12 + 4 (25%)	4 + 4 (50%)	—
5	4 + 1 (20%)	8 + 2 (20%)	3 + 2 (40%)	12 + 3 (20%)	7 + 3 (30%)	—	16 + 4 (20%)	6 + 4 (40%)	—
6	5 + 1 (17%)	10 + 2 (17%)	4 + 2 (33%)	15 + 3 (17%)	9 + 3 (25%)	—	16 + 4 (20%)	8 + 4 (33%)	—
7	6 + 1 (14%)	12 + 2 (14%)	5 + 2 (29%)	15 + 3 (17%)	11 + 3 (21%)	4 + 3 (43%)	16 + 4 (20%)	10 + 4 (29%)	—
8	7 + 1 (13%)	14 + 2 (12.5%)	6 + 2 (25%)	15 + 3 (17%)	13 + 3 (19%)	5 + 3 (38%)	16 + 4 (20%)	12 + 4 (25%)	—
9	8 + 1 (11%)	16 + 2 (11%)	7 + 2 (22%)	15 + 3 (17%)	15 + 3 (17%)	6 + 3 (33%)	16 + 4 (20%)	14 + 4 (22%)	5 + 4 (44%)
10	9 + 1 (10%)	16 + 2 (11%)	8 + 2 (20%)	15 + 3 (17%)	15 + 3 (17%)	7 + 3 (30%)	16 + 4 (20%)	16 + 4 (20%)	6 + 4 (40%)
12	11 + 1 (8%)	16 + 2 (11%)	10 + 2 (17%)	15 + 3 (17%)	15 + 3 (17%)	9 + 3 (25%)	16 + 4 (20%)	16 + 4 (20%)	8 + 4 (33%)
14	13 + 1 (7%)	16 + 2 (11%)	12 + 2 (14%)	15 + 3 (17%)	15 + 3 (17%)	11 + 3 (21%)	16 + 4 (20%)	16 + 4 (20%)	10 + 4 (29%)
16	15 + 1 (6%)	16 + 2 (11%)	14 + 2 (13%)	15 + 3 (17%)	15 + 3 (17%)	13 + 3 (19%)	16 + 4 (20%)	16 + 4 (20%)	12 + 4 (25%)
18	16 + 1 (6%)	16 + 2 (11%)	16 + 2 (11%)	15 + 3 (17%)	15 + 3 (17%)	15 + 3 (17%)	16 + 4 (20%)	16 + 4 (20%)	14 + 4 (22%)
20	16 + 1 (6%)	16 + 2 (11%)	16 + 2 (11%)	16 + 3 (16%)	16 + 3 (16%)	16 + 3 (16%)	16 + 4 (20%)	16 + 4 (20%)	16 + 4 (20%)
30	16 + 1 (6%)	16 + 2 (11%)	16 + 2 (11%)	16 + 3 (16%)	16 + 3 (16%)	16 + 3 (16%)	16 + 4 (20%)	16 + 4 (20%)	16 + 4 (20%)

The parity overhead for mirrored data protection is not affected by the number of nodes in the cluster. The following table describes the parity overhead for requested mirrored protection.

2x	3x	4x	5x	6x	7x	8x
50%	67%	75%	80%	83%	86%	88%

CHAPTER 15

NDMP backup

This section contains the following topics:

- [NDMP backup and recovery overview](#)..... 616
- [NDMP two-way backup](#)..... 616
- [Snapshot-based incremental backups](#)..... 617
- [NDMP protocol support](#)..... 618
- [Supported DMAs](#)..... 618
- [NDMP hardware support](#)..... 619
- [NDMP backup limitations](#)..... 619
- [NDMP performance recommendations](#)..... 620
- [Excluding files and directories from NDMP backups](#)..... 621
- [Configuring basic NDMP backup settings](#)..... 623
- [Managing NDMP user accounts](#)..... 624
- [Managing NDMP backup devices](#)..... 625
- [Managing NDMP backup ports](#)..... 626
- [Managing NDMP backup sessions](#)..... 627
- [Managing restartable backups](#)..... 629
- [Managing file list backups](#)..... 631
- [Parallel restore operation](#)..... 633
- [Sharing tape drives between clusters](#)..... 634
- [Managing default NDMP settings](#)..... 634
- [Managing snapshot based incremental backups](#)..... 638
- [View NDMP backup logs](#) 639
- [NDMP backup commands](#)..... 639

NDMP backup and recovery overview

In OneFS, you can back up and restore file-system data through the Network Data Management Protocol (NDMP). From a backup server, you can direct backup and recovery processes between an Isilon cluster and backup devices such as tape devices, media servers, and virtual tape libraries (VTLs).

OneFS supports both three-way and two-way NDMP backup models. Three-way NDMP backup is also known as the remote NDMP backup and the two-way NDMP backup is known as the local or direct NDMP backup. During a three-way NDMP backup operation, a data management application (DMA) on a backup server instructs the cluster to start backing up data to a tape media server that is either attached to the LAN or directly attached to the DMA.

During a two-way NDMP backup operation, a DMA on a backup server instructs a Backup Accelerator node on the cluster to start backing up data to a tape media server that is attached to the Backup Accelerator node.

Two-way NDMP backup is significantly faster than the three-way NDMP backup. It is also the most efficient method in terms of cluster resource consumption. However, a two-way NDMP backup requires that you attach one or more Backup Accelerator nodes to the cluster.

In both the two-way and three-way NDMP backup models, file history data is transferred from the cluster to the backup server. Before a backup begins, OneFS creates a snapshot of the targeted directory, then backs up the snapshot, which ensures that the backup image represents a specific point in time.

You do not need to activate a SnapshotIQ license on the cluster to perform NDMP backups. If you have activated a SnapshotIQ license on the cluster, you can generate a snapshot through the SnapshotIQ tool, and then back up the same snapshot to multiple tape devices. If you back up a SnapshotIQ snapshot, OneFS does not create another snapshot for the backup.

Note

If you are backing up SmartLock directories for compliance purposes, it is recommended that you do not specify autocommit time periods for the SmartLock directories. This is because, depending on the autocommit period, files in the SmartLock directories may still be subject to change.

NDMP two-way backup

The NDMP two-way backup is also known as the local or direct NDMP backup. To perform NDMP two-way backups, you must attach a Backup Accelerator node to your Isilon cluster and attach a tape device to the Backup Accelerator node. You must then use OneFS to detect the tape device before you can back up to that device.

You can connect supported tape devices directly to the Fibre Channel ports of a Backup Accelerator node. Alternatively, you can connect Fibre Channel switches to the Fibre Channel ports on the Backup Accelerator node, and connect tape and media changer devices to the Fibre Channel switches. For more information, see your Fibre Channel switch documentation about zoning the switch to allow communication between the Backup Accelerator node and the connected tape and media changer devices.

If you attach tape devices to a Backup Accelerator node, the cluster detects the devices when you start or restart the node or when you re-scan the Fibre Channel ports to

discover devices. If a cluster detects tape devices, the cluster creates an entry for the path to each detected device.

If you connect a device through a Fibre Channel switch, multiple paths can exist for a single device. For example, if you connect a tape device to a Fibre Channel switch, and then connect the Fibre Channel switch to two Fibre Channel ports, OneFS creates two entries for the device, one for each path.

Note

If you perform an NDMP two-way backup operation, you must assign static IP addresses to the Backup Accelerator node. If you connect to the cluster through a data management application (DMA), you must connect to the IP address of a Backup Accelerator node. If you perform an NDMP three-way backup, you can connect to any node in the cluster.

Snapshot-based incremental backups

You can implement snapshot-based incremental backups to increase the speed at which these backups are performed.

During a snapshot-based incremental backup, OneFS checks the snapshot taken for the previous NDMP backup operation and compares it to a new snapshot. OneFS then backs up all data that was modified since the last snapshot was made.

If the incremental backup does not involve snapshots, OneFS must scan the directory to discover which files were modified. OneFS can perform incremental backups significantly faster if snapshots are referenced.

You can perform incremental backups without activating a SnapshotIQ license on the cluster. Although SnapshotIQ offers a number of useful features, it does not enhance snapshot capabilities in NDMP backup and recovery.

Note

If you run an NDMP backup on a cluster with a SnapshotIQ license, the snapshot visibility must be turned on for SMB, NFS, and local clients for a successful completion of the operation.

Set the `BACKUP_MODE` environment variable to `SNAPSHOT` to enable snapshot-based incremental backups. If you enable snapshot-based incremental backups, OneFS retains each snapshot taken for NDMP backups until a new backup of the same or lower level is performed. However, if you do not enable snapshot-based incremental backups, OneFS automatically deletes each snapshot generated after the corresponding backup is completed or canceled.

After setting the `BACKUP_MODE` environment variable, snapshot-based incremental backup works with certain data management applications (DMAs) as listed in the next table.

Table 18 DMA support for snapshot-based incremental backups

DMA	Supported
Symantec NetBackup	Yes
EMC NetWorker	Yes
EMC Avamar	Yes

Table 18 DMA support for snapshot-based incremental backups (continued)

DMA	Supported
Commvault Simpana	No
IBM Tivoli Storage Manager	No
Symantec Backup Exec	Yes
Dell NetVault	No
ASG-Time Navigator	No

NDMP protocol support

You can back up cluster data through version 3 or 4 of the NDMP protocol.

OneFS supports the following features of NDMP versions 3 and 4:

- Full (level 0) NDMP backups
- Incremental (levels 1-10) NDMP backups

Note

In a level 10 NDMP backup, only data changed since the most recent incremental (level 1-9) backup or the last level 10 backup is copied. By repeating level 10 backups, you can be assured that the latest versions of files in your data set are backed up without having to run a full backup.

- Token-based NDMP backups
- NDMP TAR backup type
- Path-based and dir/node file history format
- Direct Access Restore (DAR)
- Directory DAR (DDAR)
- Including and excluding specific files and directories from backup
- Backup of file attributes
- Backup of Access Control Lists (ACLs)
- Backup of Alternate Data Streams (ADSs)
- Backup Restartable Extension (BRE)

OneFS supports connecting to clusters through IPv4 or IPv6.

Supported DMAs

NDMP backups are coordinated by a data management application (DMA) that runs on a backup server.

OneFS supports the following DMAs:

- Symantec NetBackup
- EMC NetWorker

- EMC Avamar
- Symantec Backup Exec
- IBM Tivoli Storage Manager
- Dell NetVault
- CommVault Simpana
- ASG-Time Navigator

Note

All supported DMAs can connect to an Isilon cluster through IPv4. CommVault Simpana is currently the only DMA that also supports connecting to an Isilon cluster through IPv6.

See the [Isilon Third-Party Software and Hardware Compatibility Guide](#) for the latest information about supported DMAs.

NDMP hardware support

OneFS can back up data to and restore data from tape devices and virtual tape libraries (VTLs).

Supported tape devices

OneFS supports the following types of emulated and physical tape devices for two-way NDMP backups:

- LTO-3
- LTO-4
- LTO-5
- LTO-6

Note

OneFS supports only the LTO-3 tape device for a FalconStor VTL. It supports only the LTO-3 and LTO-4 tape devices for a Data Domain VTL.

For three-way NDMP backups, the data management application (DMA) determines the tape devices that are supported.

Supported tape libraries

For both two-way and three-way NDMP backups, OneFS supports all the tape libraries that are supported by the DMA.

Supported virtual tape libraries

For two-way NDMP backups, OneFS supports FalconStor and Data Domain VTLs.

For three-way NDMP backups, the DMA determines the virtual tape libraries that will be supported.

NDMP backup limitations

OneFS NDMP backups have the following limitations:

- Does not support more than 4 KB path length.
- Does not back up file system configuration data, such as file protection level policies and quotas.

- Cannot back up tape blocks larger than 512 KB.
- Does not support multiplexing across multiple streams.
- Does not support multiple concurrent backups onto the same tape.
- Does not support restoring data from a file system other than OneFS. However, you can migrate data via the NDMP protocol from a NetApp or EMC VNX storage system to OneFS.
- Backup Accelerator nodes cannot interact with more than 1024 device paths, including the paths of tape and media changer devices. For example, if each device has four paths, you can connect 256 devices to a Backup Accelerator node. If each device has two paths, you can connect 512 devices.

NDMP performance recommendations

Consider the following recommendations to optimize OneFS NDMP backups.

General performance recommendations

- Install the latest patches for OneFS and your data management application (DMA).
- Run a maximum of eight NDMP concurrent sessions per A100 Backup Accelerator node and four NDMP concurrent sessions per Annapurna Backup Accelerator node to obtain optimal throughput per session.
- NDMP backups result in very high Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs). You can reduce your RPO and RTO by attaching one or more Backup Accelerator nodes to the cluster and then running two-way NDMP backups.
- The throughput for an Isilon cluster during the backup and restore operations is dependent on the dataset and is considerably reduced for small files.
- If you are backing up multiple directories that contain small files, set up a separate schedule for each directory.
- If you are performing three-way NDMP backups, run multiple NDMP sessions on multiple nodes in your Isilon cluster.
- Restore files through Direct Access Restore (DAR), especially if you restore files frequently. However, it is recommended that you do not use DAR to restore a full backup or a large number of files, as DAR is better suited to restoring smaller numbers of files.
- Restore files through Directory DAR (DDAR) if you restore large numbers of files frequently.
- Use the largest tape record size available for your version of OneFS to increase throughput.
- If possible, do not include or exclude files from backup. Including or excluding files can affect backup performance, due to filtering overhead.
- Limit the depth of nested subdirectories in your file system.
- Limit the number of files in a directory. Distribute files across multiple directories instead of including a large number of files in a single directory.
- Use path-based file history format.
- Configure multiple policies when scheduling backup operations, with each policy capturing a portion of the file system. Do not attempt to back up the entire file system through a single policy.

- If you are backing up a large number of small files, it is recommended that you distribute the files across multiple NDMP jobs. Spreading backup operations across multiple NDMP jobs can significantly improve the performance of clusters.

SmartConnect recommendations

- A two-way NDMP backup session with SmartConnect requires backup accelerators for backup and restore operations. However, a three-way NDMP session with SmartConnect does not require backup accelerators for these operations.
- For a two-way NDMP backup session with SmartConnect, connect to the NDMP session through a dedicated SmartConnect zone consisting of a pool of Network Interface Cards (NICs) on the backup accelerator nodes.
- For a two-way NDMP backup session without SmartConnect, initiate the backup session through a static IP address or fully qualified domain name of the backup accelerator node.
- For a three-way NDMP backup operation, the front-end Ethernet network or the interfaces of the nodes are used to serve the backup traffic. Therefore, it is recommended that you configure a DMA to initiate an NDMP session only using the nodes that are not already overburdened serving other workloads or connections.
- For a three-way NDMP backup operation with or without SmartConnect, initiate the backup session using the IP addresses of the nodes that are identified for running the NDMP sessions.
- For a three-way NDMP backup operation, when selecting nodes to include in the SmartConnect zone, make sure to include nodes of the same type in the zone. A SmartConnect zone with mixed node types, for example, S210 and NL400 nodes, might degrade the performance of the backup operation.

Backup Accelerator recommendations

- Assign static IP addresses to Backup Accelerator nodes.
- Attach more Backup Accelerator nodes to larger clusters. The recommended number of Backup Accelerator nodes is listed in the following table.

Table 19 Nodes per Backup Accelerator node

Node type	Recommended number of nodes per Backup Accelerator node
X-Series	3
NL-Series	3
S-Series	3
HD-Series	3

- Attach more Backup Accelerator nodes if you are backing up to more tape devices.

DMA-specific recommendations

- Enable parallelism for the DMA if the DMA supports this option. This allows OneFS to back up data to multiple tape devices at the same time.

Excluding files and directories from NDMP backups

You can exclude files and directories from NDMP backup operations by specifying NDMP environment variables through a data management application (DMA). If you include a file or directory, all other files and directories are automatically excluded from backup

operations. If you exclude a file or directory, all files and directories except the excluded one are backed up.

You can include or exclude files and directories by specifying the following character patterns:

Table 20 NDMP file and directory matching wildcards

Character	Description	Example	Includes or excludes the following directories
*	Takes the place of any character or characters	archive*	/ifs/data/archive1 /ifs/data/archive42_a/media
[]	Takes the place of a range of letters or numbers	data_store_[a-f] data_store_[0-9]	/ifs/data/data_store_a /ifs/data/data_store_c /ifs/data/data_store_8
?	Takes the place of any single character	user_?	/ifs/data/user_1 /ifs/data/user_2
\	Includes a blank space	user\ 1	/ifs/data/user 1

Note

" " are required for Symantec NetBackup when multiple patterns are specified. The patterns are not limited to directories.

Unanchored patterns such as `home` or `user1` target a string of text that might belong to many files or directories. Anchored patterns target specific file pathnames, such as `ifs/data/home`. You can include or exclude either type of pattern.

For example, suppose you want to back up the `/ifs/data/home` directory, which contains the following files and directories:

- `/ifs/data/home/user1/file.txt`
- `/ifs/data/home/user2/user1/file.txt`
- `/ifs/data/home/user3/other/file.txt`
- `/ifs/data/home/user4/emptydirectory`

If you include the `/ifs/data/home` directory, all files and directories, including `emptydirectory` would be backed up.

If you specify both the include and exclude patterns, the include pattern is first processed followed by the exclude pattern.

If you specify both the include and exclude patterns, any excluded files or directories under the included directories would not be backed up. If the excluded directories are not found in any of the included directories, the exclude specification would have no effect.

Note

Specifying unanchored patterns can degrade the performance of backups. It is recommended that you avoid unanchored patterns whenever possible.

Configuring basic NDMP backup settings

You can configure NDMP backup settings to control how these backups are performed for the cluster. You can also configure OneFS to interact with a specific data management application (DMA) for NDMP backups.

Configure and enable NDMP backup

OneFS prevents NDMP backups by default. Before you can perform NDMP backups, you must enable NDMP backups and configure NDMP settings.

Procedure

1. Enable NDMP backup by running the following command:

```
isi services ndmpd enable
```

2. Configure NDMP backup by running the `isi ndmp settings set` command.

The following command configures OneFS to interact with EMC NetWorker:

```
isi ndmp settings set dma emc
```

Disable NDMP backup

You can disable NDMP backup if you no longer want to back up data through NDMP.

Procedure

1. Run the following command:

```
isi services ndmpd disable
```

View NDMP backup settings

You can view current NDMP backup settings, which indicate whether the service is enabled, the port through which data management applications (DMAs) connect to the cluster, and the DMA vendor that OneFS is configured to interact with.

Procedure

1. Run the following command:

```
isi ndmp settings list
```

NDMP backup settings

You can configure settings that control how NDMP backups are performed on the cluster.

The following information is displayed in the output of the `isi ndmp settings list` command:

port

The number of the port through which data management applications (DMAs) can connect to the cluster.

dma

The DMA vendor that the cluster is configured to interact with.

Managing NDMP user accounts

You can create, delete, and modify the passwords of NDMP user accounts.

Create an NDMP user account

Before you can perform NDMP backups, you must create an NDMP user account through which a data management application (DMA) can access the cluster.

Procedure

1. Run the `isi ndmp user create` command.

The following command creates an NDMP user account called NDMPuser with a password of password123:

```
isi ndmp user create NDMPuser password123
```

Modify the password of an NDMP user account

You can modify the password for an NDMP user account.

Procedure

1. Run the `isi ndmp user modify` command.

The following command modifies the password of NDMPuser:

```
isi ndmp user modify NDMPuser newPassword123
```

Delete an NDMP user account

You can delete an NDMP user account.

Procedure

1. Run the `isi ndmp user delete` command.

The following command deletes NDMPuser:

```
isi ndmp user delete NDMPuser
```

View NDMP user accounts

You can view information about NDMP user accounts.

Procedure

1. Run the following command:

```
isi ndmp user list
```


Managing NDMP backup devices

After you attach a tape or media changer device to a Backup Accelerator node, you must configure OneFS to detect and establish a connection to the device. After the connection between the cluster and the backup device is established, you can modify the name that the cluster has assigned to the device, or disconnect the device from the cluster.

Detect NDMP backup devices

If you connect devices to a Backup Accelerator node, you must configure OneFS to detect the devices before OneFS can back up data to and restore data from the devices. You can scan a specific node, a specific port, or all ports on all nodes.

Procedure

1. Run the following command:

```
isi tape rescan
```

Modify an NDMP backup device entry name

You can modify the name of an NDMP device entry.

Procedure

1. Run the `isi tape rename` command.

The following command renames `tape003` to `tape005`:

```
isi tape rename tape003 tape005
```

Delete a device entry for a disconnected NDMP backup device

If you physically remove an NDMP device from a cluster, OneFS retains the entry for the device. You can delete a device entry for a removed device. You can also remove the device entry for a device that is still physically attached to the cluster; this causes OneFS to disconnect from the device.

If you remove a device entry for a device that is connected to the cluster, and you do not physically disconnect the device, OneFS will detect the device the next time it scans the ports. You cannot remove a device entry for a device that is currently being backed up to or restored from.

Procedure

1. The following command disconnects `tape001` from the cluster:

```
isi tape delete tape001
```

View NDMP backup devices

You can view information about tape and media changer devices that are currently attached to the cluster through a Backup Accelerator node.

Procedure

1. Run the following command:

```
isi tape list
```

Managing NDMP backup ports

You can manage the Fibre Channel ports that connect tape and media changer devices to a Backup Accelerator node. You can also enable, disable, or modify the settings of an NDMP backup port.

Modify NDMP backup port settings

You can modify the settings of an NDMP backup port.

Procedure

1. Run the `isi fc set` command.

The following command configures port 1 on node 5 to support a point-to-point Fibre Channel topology:

```
isi fc set 5:1 --topology ptp
```

Enable or disable an NDMP backup port

You can enable or disable an NDMP backup port.

Procedure

1. Run either the `isi fc enable` or `isi fc disable` command:

The following command disables port 1 on node 5:

```
isi fc disable 5:1
```

The following command enables port 1 on node 5:

```
isi fc enable 5:1
```

View NDMP backup ports

You can view information about Fibre Channel ports of Backup Accelerator nodes attached to a cluster.

Procedure

1. Run the following command:

```
isi fc list
```

NDMP backup port settings

OneFS assigns default settings to each port on each Backup Accelerator node attached to the cluster. These settings identify each port and specify how the port interacts with NDMP backup devices.

The following information is displayed in the output of the `isi fc list` command:

Port

The name of the Backup Accelerator node, and the number of the port.

WWNN

The world wide node name (WWNN) of the port. This name is the same for each port on a given node.

WWPN

The world wide port name (WWPN) of the port. This name is unique to the port.

State

Whether the port is enabled or disabled.

Topology

The type of Fibre Channel topology that the port is configured to support.

Rate

The rate at which data is sent through the port. The rate can be set to 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, and Auto. 8 Gb/s is available for A100 nodes only. If set to Auto, OneFS automatically negotiates with the DMA to determine the rate. Auto is the recommended setting.

Managing NDMP backup sessions

You can view the status of NDMP backup sessions or terminate a session that is in progress.

End an NDMP session

You can interrupt an NDMP backup or restore operation by ending an NDMP session.

Procedure

1. To retrieve the ID of the NDMP session that you want to end, run the `isi ndmp list` command.
2. Run the `isi ndmp kill` command.

The following command ends an NDMP session with an ID of 4.36339:

```
isi ndmp kill 4.36339
```

View NDMP sessions

You can view information about NDMP sessions that exist between the cluster and data management applications (DMAs).

Procedure

1. Run the following command:

```
isi ndmp list
```

NDMP session information

You can view information about active NDMP sessions.

The following information is displayed in the output of the `isi ndmp list` command:

Session

The unique identification number that OneFS assigned to the session.

Data

The current state of the data server. The data server is a component of OneFS that sends data during backups and receives information during restore operations.

A

Active. The data server is currently sending data.

I

Idle. The data server is not sending data.

L

Listen. The data server is waiting to connect to the data mover.

M

Mover. The current state of the data mover. The data mover is a component of the backup server that receives data during backups and sends data during restore operations.

A

Active. The data mover is currently receiving data.

I

Idle. The data mover is not receiving data.

L

Listen. The data mover is waiting to connect to the data server.

OP

The type of operation that is currently in progress. If no operation is in progress, this field is blank.

B (0-10)

Backup. Indicates that data is currently being backed up to a media server. The number indicates the level of NDMP backup.

R

Restore. Indicates that data is currently being restored from a media server.

Elapsed Time

How much time has elapsed since the session started.

Bytes Moved

The number of bytes that were transferred during the session.

Throughput

The average throughput of the session over the past five minutes.

Managing restartable backups

A restartable backup is a type of NDMP backup that you can enable in your data management application (DMA). If a restartable backup fails, for example, because of a power outage, you can restart the backup from a checkpoint close to the point of failure. In contrast, when a non-restartable backup fails, you must back up all data from the beginning, regardless of what was transferred during the initial backup process.

After you enable restartable backups from your DMA, you can manage restartable backup contexts from OneFS. These contexts are the data that OneFS stores to facilitate restartable backups. Each context represents a checkpoint that the restartable backup process can return to if a backup fails.

Restartable backups are supported only for EMC NetWorker 8.1 and later.

Configure restartable backups for EMC NetWorker

You must configure EMC NetWorker to enable restartable backups and, optionally, define the checkpoint interval.

If you do not specify a checkpoint interval, NetWorker uses the default interval of 5 GB.

Procedure

1. Configure the client and the directory path that you want to back up as you would normally.
2. In the **Client Properties** dialog box, enable restartable backups.
 - a. On the **General** page, click the **Checkpoint enabled** checkbox.
 - b. In the **Checkpoint granularity** drop-down list, select **File**.
3. In the **Application information** field, type any NDMP variables that you want to specify.

The following variable setting specifies a checkpoint interval of 1 GB:

```
CHECKPOINT_INTERVAL_IN_BYTES=1GB
```
4. Finish configuration and click **OK** in the **Client Properties** dialog box.
5. Start the backup.
6. If the backup is interrupted—for example, because of a power failure—restart it.
 - a. On the **Monitoring** page, locate the backup process in the **Groups** list.
 - b. Right-click the backup process and then, in the context menu, click **Restart**. NetWorker automatically restarts the backup from the last checkpoint.

View restartable backup contexts

You can view restartable backup contexts that have been configured.

Procedure

1. View all backup contexts by running the following command:

```
isi ndmp extensions contexts list
```

2. To view detailed information about a specific backup context, run the `isi ndmp extensions contexts view` command.

The following command displays detailed information about a backup context with an ID of 792eeb8a-8784-11e2-aa70-0025904e91a4:

```
isi ndmp extensions contexts view 792eeb8a-8784-11e2-aa70-0025904e91a4
```

Delete a restartable backup context

After a restartable backup context is no longer needed, your data management application (DMA) automatically requests that OneFS delete the context. You can manually delete a restartable backup context before the DMA requests it.

Note

It is recommended that you do not manually delete restartable backup contexts. Manually deleting a restartable backup context requires you to restart the corresponding NDMP backup from the beginning.

Procedure

1. Run the `isi ndmp extensions contexts delete` command.

The following command deletes a restartable backup context with an ID of 792eeb8a-8784-11e2-aa70-0025904e91a4:

```
isi ndmp extensions contexts delete 792eeb8a-8784-11e2-aa70-0025904e91a4
```

Configure restartable backup settings

You can specify the number of restartable backup contexts that OneFS retains at a time, up to a maximum of 1024 contexts.

Procedure

1. Run the `isi ndmp extensions settings modify` command.

The following command sets the maximum number of restartable backup contexts to 128:

```
isi ndmp extensions settings modify --bre_max_contexts 128
```

View restartable backup settings

You can view the current limit of restartable backup contexts that OneFS retains at one time.

Procedure

1. Run the following command:

```
isi ndmp extensions settings view
```

Managing file list backups

If your data management application (DMA) can pass environment variables to OneFS, you can control backups by specifying a file list.

Currently, EMC Networker and Symantec NetBackup can pass environment variables to OneFS.

With a normal NDMP level 0 (full) backup, your DMA backs up an entire source directory. With an NDMP incremental (level 1-10) backup, your DMA backs up only those files that have been created or changed since the previous incremental backup of the same level.

When you specify a file list backup, only the listed files and subdirectories in the source directory are backed up. With a level 0 file list backup, all listed files and directories in the source directory are backed up.

A backup level other than 0 triggers an incremental file list backup. In an incremental file list backup, only the listed files that were created or changed in the source directory since the last incremental backup of the same level are backed up.

To configure a file list backup, you must complete the following tasks:

- Create the file list and place it in OneFS
- Specify the path of the source directory
- Specify the file list location

The file list is an ASCII text file that lists the pathnames of files to be backed up. The pathnames must be relative to the path specified in the `FILESYSTEM` environment variable. Absolute file paths in the file list are not supported. The pathnames of all files must be included, or they are not backed up. For example, if you include the pathname of a subdirectory, only the subdirectory, not the files it contains, is backed up.

To specify the full path of the source directory to be backed up, you must specify the `FILESYSTEM` environment variable in your DMA. For example:

```
FILESYSTEM=/ifs/data/projects
```

To specify the pathname of the file list, you must specify the environment variable, `BACKUP_FILE_LIST` in your DMA. The file list must be accessible from the node performing the backup. For example:

```
BACKUP_FILE_LIST=/ifs/data/proj_list.txt
```

Format of a backup file list

You must create a file list to enable a file list backup.

A file list backup requires an ASCII text file in a particular format to identify the pathnames of files to be backed up. Following is an example of a file list with pathnames relative to `/ifs/data/projects`:

```
project-summary_rev0.xls
project-summary_rev3.xls
proj001/plan/plan001.doc
proj001/plan/plan001.pdf
proj001/plan/proj-logo.png
proj001/plan/schedule001.xls
proj001/plan/stakeholders_list.txt
proj001/plan/team_pic.png
proj002/plan/logo.png
proj002/plan/projplan002.doc
proj002/plan/projplan002.pdf
proj002/plan/sched-002.xls
proj002/plan/stakeholders.txt
proj002/plan/team.png
proj005/plan/projectlogo.png
proj005/plan/projectteam.png
proj005/plan/proj-plan005.doc
proj005/plan/proj-plan005.pdf
proj005/plan/schedule.xls
proj005/plan/stakeholders.txt
```

As shown in the example, the pathnames are relative to the full path of the source directory, which you specify in the `FILESYSTEM` environment variable. Absolute file paths are not supported in the file list.

Also as shown, the directories and files should be in sorted order.

The pathnames of all files must be included in the file list, or they are not backed up. For example, if you only include the pathname of a subdirectory, the subdirectory is backed up, but not the files the subdirectory contains. The exception is ADS (alternate data streams). All ADS associated with a file to be backed up are automatically backed up.

Placement of the file list

Before you can perform a file list backup, you must place the file list in OneFS.

For example, suppose the `FILESYSTEM` environment variable specifies the full path of the directory to be backed up as `/ifs/data/projects`. You can place the text file containing the file list anywhere within the `/ifs` path.

Start a file list backup

You can configure and start a file list backup from your data management application (DMA).

Before you begin

You should have already specified and saved the list of files to be backed up in an ASCII text file.

Configure a file list backup from your DMA as you would any backup, but with a few additional steps as described in the following procedure.

Procedure

1. Copy the file list to the OneFS file system on the cluster containing the files to be backed up.

For example, if the directory that you specify in the `FILESYSTEM` environment variable is `/ifs/data/projects`, you could place your file list at `/ifs/data`.

2. In your DMA, specify the `BACKUP_FILE_LIST` environment variable to be the full pathname of the file list.

For example, if the file list was named `proj_backup.txt`, and you placed it at `/ifs/data`, specify `/ifs/data/proj_backup.txt` as the full pathname of the file list.

3. Start your backup as you normally would.

Results

The files in your file list are backed up as specified.

Parallel restore operation

In OneFS, the default NDMP restore for any path is the parallel restore operation. Parallel (multi-threaded) restore enables faster full or partial restore operations by writing data to the cluster as fast as the data can be read from the tape. However, if you specify DAR (direct access restore), the operation reverts to serial processing.

Specify a serial restore operation

You can use the `RESTORE_OPTIONS` environment variable to specify a serial (single-threaded) restore operation.

Procedure

1. In your data management application, configure a restore operation as you normally would.
2. Make sure that the `RESTORE_OPTIONS` environment variable is set to `1` on your data management application.

If the `RESTORE_OPTIONS` environment variable is not already set to `1`, specify the `isi ndmp settings variables modify` command from the OneFS command line. The following command specifies serial restore for the `/ifs/data/projects` directory:

```
isi ndmp settings variables modify --path /ifs/data/projects --
name restore_options --value 1
```

The value of the `path` option is the `FILESYSTEM` environment variable set during the backup operation. The value that you specify for the `name` option is case sensitive.

3. Start the restore operation.

Sharing tape drives between clusters

Multiple Isilon clusters, or an Isilon cluster and a third-party NAS system, can be configured to share a single tape drive. This helps to maximize the use of the tape infrastructure in your data center.

In your data management application (DMA), you must configure NDMP to control the tape drive and ensure that it is shared properly. The following configurations are supported.

OneFS Versions	Supported DMAs	Tested configurations
<ul style="list-style-type: none"> • 7.1.1 • 7.1.0.1 (and later)* • 7.0.2.5 • 6.6.5.26 	<ul style="list-style-type: none"> • EMC NetWorker 8.0 and later • Symantec NetBackup 7.5 and later 	<ul style="list-style-type: none"> • Isilon Backup Accelerator with a second Backup Accelerator • Isilon Backup Accelerator with a NetApp storage system
* The tape drive sharing function is not supported in the OneFS 7.0.1 release.		

EMC NetWorker refers to the tape drive sharing capability as DDS (dynamic drive sharing). Symantec NetBackup uses the term SSO (shared storage option). Consult your DMA vendor documentation for configuration instructions.

Managing default NDMP settings

In OneFS, you can manage NDMP backup and restore operations by specifying default NDMP environment variables. You can also override default NDMP environment variables through your data management application (DMA). For more information about specifying NDMP environment variables through your DMA, see your DMA documentation.

Set default NDMP settings for a directory

You can set default NDMP settings for a directory.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Set default NDMP settings by running the `isi ndmp settings variables create` command.

For example, the following command enables snapshot-based incremental backups for `/ifs/data/media`:

```
isi ndmp settings variables create /ifs/data/media BACKUP_MODE
SNAPSHOT
```

Modify default NDMP settings for a directory

You can modify the default NDMP settings for a directory.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.

2. Modify default NDMP settings by running the `isi ndmp settings variables modify` command.

For example, the following command sets the default file history format to path-based format for `/ifs/data/media`:

```
isi ndmp settings variables modify /ifs/data/media HIST F
```

3. (Optional) To remove a default NDMP setting for a directory, run the `isi ndmp settings variables delete` command:

For example, the following command removes the default file history format for `/ifs/data/media`:

```
isi ndmp settings variables delete /ifs/data/media --name HIST
```

View default NDMP settings for directories

You can view the default NDMP settings for directories.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. View default NDMP settings by running the following command:

```
isi ndmp settings variables list
```

NDMP environment variables

You can specify default settings of NDMP backup and restore operations through NDMP environment variables. You can also specify NDMP environment variables through your data management application (DMA).

Table 21 NDMP environment variables

Environment variable	Valid values	Default	Description
BACKUP_MODE=	TIMESTAMP SNAPSHOT	TIMESTAMP	Enables or disables snapshot-based incremental backups. To enable snapshot-based incremental backups, specify <code>SNAPSHOT</code> . To disable snapshot-based incremental backups, specify <code>TIMESTAMP</code> .
FILESYSTEM=	<i><file-path></i>	None	Specifies the full path of the directory you want to back up. Must be specified by the DMA before starting the backup, or an error is generated.
LEVEL=	<i><integer></i>	0	Specifies the level of NDMP backup to perform.

Table 21 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			<p>The following values are valid:</p> <p>0 Performs a full NDMP backup.</p> <p>1 - 9 Performs an incremental backup at the specified level.</p> <p>10 Performs unlimited incremental backups.</p>
UPDATE=	Y N	Y	<p>Determines whether OneFS updates the dump dates file.</p> <p>Y OneFS updates the dump dates file.</p> <p>N OneFS does not update the dump dates file.</p>
HIST=	<i><file-history-format></i>	Y	<p>Specifies the file history format. The following values are valid:</p> <p>D Specifies dir/node file history.</p> <p>F Specifies path-based file history.</p> <p>Y Specifies the default file history format determined by your NDMP backup settings.</p> <p>N Disables file history.</p>
DIRECT=	Y N	N	<p>Enables or disables Direct Access Restore (DAR) and</p>

Table 21 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			Directory DAR (DDAR). The following values are valid: Y Enables DAR and DDAR. N Disables DAR and DDAR.
FILES=	<i><file-matching-pattern></i>	None	If you specify this option, OneFS backs up only files and directories that meet the specified pattern. Separate multiple patterns with a space.
EXCLUDE=	<i><file-matching-pattern></i>	None	If you specify this option, OneFS does not back up files and directories that meet the specified pattern. Separate multiple patterns with a space.
RESTORE_HARDLINK _BY_TABLE=	Y N	N	Determines whether OneFS recovers hard links by building a hard-link table during restore operations. Specify this option if hard links were incorrectly backed up, and restore operations are failing. If a restore operation fails because hard links were incorrectly backed up, the following message appears in the NDMP backup logs: Bad hardlink path for <path>
BACKUP_FILE_LIST=	<i><file-path></i>	None	Specifies the pathname in OneFS of the file list to control the backup. This variable must be passed from the DMA initiating the backup. Currently, only EMC Networker and Symantec

Table 21 NDMP environment variables (continued)

Environment variable	Valid values	Default	Description
			NetBackup can pass environment variables to OneFS.
RESTORE_OPTIONS=	0 1	0	The restore operation, by default, is multi-threaded to improve performance. To change the restore operation to single-threaded, specify RESTORE_OPTIONS=1

Managing snapshot based incremental backups

After you enable snapshot-based incremental backups, you can view and delete the snapshots created for these backups.

Enable snapshot-based incremental backups for a directory

You can configure OneFS to perform snapshot-based incremental backups for a directory by default. You can also override the default setting in your data management application (DMA).

Procedure

1. Run the `isi ndmp settings variables create` command.

The following command enables snapshot-based incremental backups for `/ifs/data/media`:

```
isi ndmp settings variables create /ifs/data/media BACKUP_MODE
SNAPSHOT
```

Delete snapshots for snapshot-based incremental backups

You can delete snapshots created for snapshot-based incremental backups.

Note

It is recommended that you do not delete snapshots created for snapshot-based incremental backups. If all snapshots are deleted for a path, the next backup performed for the path is a full backup.

Procedure

1. Run the `isi ndmp dumpdates delete` command.

The following command deletes all snapshots created for backing up `/ifs/data/media`:

```
isi ndmp dumpdates delete /ifs/data/media
```

View snapshots for snapshot-based incremental backups

You can view snapshots generated for snapshot-based incremental backups.

Procedure

1. Run the following command:

```
isi ndmp dumpdates list
```

View NDMP backup logs

You can view information about NDMP backup and restore operations through NDMP backup logs.

Procedure

1. View the contents of the `/var/log/isi_ndmp_d` directory by running the following command:

```
more /var/log/isi_ndmp_d
```

NDMP backup commands

You can control Network Data Management Protocol (NDMP) backups through the NDMP backup commands.

isi ndmp user create

Creates an NDMP user account.

Syntax

```
isi ndmp user create <name> <password>
```

Options

<name>

Specifies a name for the NDMP user account.

<password>

Specifies a password for the NDMP user account.

Examples

The following command creates an NDMP user account with a name of `ndmp_user` and a password of 1234:

```
isi ndmp user create ndmp_user 1234
```

isi ndmp user modify

Modifies the password of an NDMP user account.

Syntax

```
isi ndmp user modify <name> <password>
```

Options

<name>

Modifies the password of the specified NDMP user account.

<password>

Assigns the specified password to the given NDMP user account.

Examples

The following command sets the password of `ndmp_user` to `newpassword`:

```
isi ndmp user modify ndmp_user newpassword
```

isi ndmp user delete

Deletes an NDMP user account.

Syntax

```
isi ndmp user delete <name>
```

Options

<name>

Deletes the specified NDMP user account.

Examples

The following example deletes `ndmp_user`:

```
isi ndmp user delete ndmp_user
```

isi ndmp user list

Displays information about NDMP users.

Syntax

```
isi ndmp user list  
[--name <name>]
```


Options

If no options are specified, displays information about all NDMP users.

`--name <name>`

Displays information about only the specified NDMP user.

Examples

To view information about all NDMP user accounts, run the following command:

```
isi ndmp user list
```

isi tape rescan

Scans Fibre Channel ports for undetected NDMP backup devices that are attached to Backup Accelerator nodes. If the scan reveals new devices, the cluster creates entries for the new devices.

Syntax

```
isi tape rescan
  [--node <lnn>]
  [--port <integer>]
  [--reconcile]
```

Options

If no options are specified, scans all nodes and ports.

`--node <lnn>`

Scans only the node of the specified logical node number (LNN).

`--port <integer>`

Scans only the specified port. If you specify `--node`, scans only the specified port on the specified node. If you do not specify `--node`, scans the specified port on all nodes.

`--reconcile`

Removes entries for devices or paths that have become inaccessible.

Examples

To scan the entire cluster for NDMP devices, and remove entries for devices and paths that have become inaccessible, run the following command:

```
isi tape rescan --reconcile
```

isi tape rename

Renames an NDMP device that is currently connected to a Backup Accelerator node on the cluster.

Syntax

```
isi tape rename <devname> <rename>
```

Options

`<devname>`

Modifies the name of the specified NDMP device.

<rename>

Specifies a new name for the given NDMP device.

Examples

The following example renames tape003 to tape005:

```
isi tape rename tape003 tape005
```

isi tape delete

Disconnects the cluster from an NDMP backup device that is currently connected to a Backup Accelerator node on the cluster.

Syntax

```
isi tape delete {<dev-name> | --all}
```

Options

<devname>

Disconnects the cluster from the specified device. Specify as an NDMP device name.

--all

Disconnects the cluster from all devices.

Examples

The following command disconnects tape001 from the cluster:

```
isi tape delete tape001
```

isi tape list

Displays a list of NDMP devices that are currently connected to the cluster.

Syntax

```
isi tape list
  [--devname <name>]
  [--node <lnn>]
  [--tape]
  [--mc]
  [--verbose]
```

Options

{--devname | --n} <name>

Displays only the specified device. Specify as a device name.

--node <lnn>

Displays only devices that are attached to the node of the specified logical node number (LNN).

--tape

Displays only tape devices.

--mc

Displays only media changer devices.

```
{--verbose | -v}
```

Displays more detailed information.

Examples

To view a list of all NDMP devices, run the following command:

```
isi tape list
```

isi fc set

Configures a Fibre Channel port on a Backup Accelerator node.

This command is valid only if a Backup Accelerator node is attached to the cluster, and the specified port is disabled.

Syntax

```
isi fc set <port> {--wwnn <wwnn> | --wwpn <wwpn>
 | --topology <topology> | --rate <rate>}
```

Options

<port>

Configure the specified port.

Specify as a port ID.

--wwnn <wwnn>

Specifies the world wide node name (WWNN) of the port.

Specify as a string of 16 hexadecimal characters.

--wwpn <wwpn>

Specifies the world wide port name (WWPN) of the port.

Specify as a string of 16 hexadecimal characters.

--topology <topology>

Specifies the type of Fibre Channel topology that the port expects.

The following values are valid:

ptp

Causes the port to expect a point-to-point topology, with one backup device or Fibre Channel switch directly connected to the port.

loop

Causes the port to expect an arbitrated loop topology, with multiple backup devices connected to a single port in a circular formation.

auto

Causes the port to detect the topology automatically. This is the recommended setting. If you are using a fabric topology, specify this setting.

--rate <rate>

Specifies the rate that OneFS will attempt to send data through the port.

The following values are valid:

auto

OneFS automatically negotiates with the DMA to determine the rate. This is the recommended setting.

1

Attempts to send data through the port at a speed of 1 Gb per second.

2

Attempts to send data through the port at a speed of 2 Gb per second.

4

Attempts to send data through the port at a speed of 4 Gb per second.

8

Attempts to send data through the port at a speed of 8 Gb per second.

Examples

The following command causes port 1 on node 5 to expect a point-to-point Fibre Channel topology:

```
isi fc set 5:1 --topology ptp
```

isi fc disable

Disables a Fibre Channel port.

Syntax

```
isi fc disable <port>
```

Options

<port>

Disables the specified port.

Specify as a port ID.

Examples

The following command disables port 1 on node 5:

```
isi fc disable 5:1
```

isi fc enable

Enables a Fibre Channel port.

Syntax

```
isi fc enable <port>
```

Options

<port>

Enables the specified port.

Specify as a port ID.

Examples

The following command enables port 1 on node 5:

```
isi fc enable 5:1
```

isi fc list

Displays a list of Fibre Channel ports on Backup Accelerator nodes connected to the cluster.

Syntax

```
isi fc list
  [{<port> | --node <id>}]
```

Options

If no options are specified, displays all Fibre Channel ports on Backup Accelerator nodes connected to the cluster.

<port>

Displays the specified port. Specify in the following format:

```
<port-id>:<node-id>
```

--node *<id>*

Displays all ports on the specified node.

Specify as a node ID.

Examples

The following command displays all ports on node 5:

```
isi fc list --node 5
```

The system displays output similar to the following example:

Port	WWNN	WWPN	State	Topology	Rate
5:1	2000001b3214ccc3	2100001b3214ccc3	enabled	auto	auto
5:2	2000001b3234ccc3	2101001b3234ccc3	enabled	auto	auto
5:3	2000001b3254ccc3	2100001b3254ccc3	enabled	auto	auto
5:4	2000001b3234ccc3	2103001b3274ccc3	enabled	auto	auto

The following command displays information about port 1 on node 5:

```
isi fc list 5:1
```

The system displays output similar to the following example:

Port	WWNN	WWPN	State	Topology	Rate
5:1	2000001b3214ccc3	2100001b3214ccc3	enabled	auto	auto

isi ndmp kill

Terminates an NDMP session.

Syntax

```
isi ndmp kill <session>
```

Options

<session>

Terminates the specified session. Specify as a session ID.

Examples

The following command terminates a session with an ID of 4.36339:

```
isi ndmp kill 4.36339
```

isi ndmp list

Displays NDMP sessions.

Syntax

```
isi ndmp list
  [--session <id>]
  [--node <id>]
  [--verbose]
```

Options

If no options are specified, displays all NDMP sessions.

--session <id>

Displays only the session of the specified ID.

--node <id>

Displays only sessions running on the node of the specified ID.

{--verbose | -v}

Displays detailed information.

Examples

Run the following command to view all NDMP sessions:

```
isi ndmp list
```

The system displays output similar to the following example:

```
Session|Data|Mover|OP|Elapsed Time|Bytes Moved|Throughput
-----+-----+-----+---+-----+-----+-----
4.36339|A   | A   |R |03d 00:58:24|16.201 TB  |12.737 MB/s
```

The following list describes the values for the Data, Mover, and OP columns:

A

Active

I	Idle
P	Paused
R	Recover
B	Backup

isi ndmp probe

Displays diagnostic information about an NDMP session.

Syntax

```
isi ndmp probe <session>
```

Options

<session>

Displays diagnostic information about the specified NDMP session. Specify as a session ID.

Examples

The following command displays diagnostic information for session 4.34729:

```
isi ndmp probe 4.34729
```

isi ndmp settings set

Configures NDMP settings.

Syntax

```
isi ndmp settings set <name> <value>
```

Options

<name>

Modifies the specified setting.

The following values are valid:

dma

Configures the cluster to interact with the specified data management vendor (DMA).

port

Specifies the port through which a DMA vendor connects to the cluster.

<value>

Specifies a value for the setting. If you are modifying the `port` setting, specify a TCP/IP port.

If you are modifying the DMA setting, the following values are valid:

atempo

Atempo Time Navigator

bakbone

BakBone NetVault

commvault

CommVault Simpana

emc

EMC Networker

symantec

Symantec Netbackup

tivoli

IBM Tivoli Storage Manager

symantec-netbackup

Symantec Netbackup

symantec-backupexec

Symantec Backup Exec

generic

Unsupported DMA vendor

Examples

To set the vendor of the current data management application to EMC Networker, run the following command:

```
isi ndmp settings set --name dma --value emc
```

The following command sets the port number on which the NDMP daemon listens for incoming connections to 10001:

```
isi ndmp settings set port 10001
```

isi ndmp settings list

Displays NDMP settings and values.

Syntax

```
isi ndmp settings list
[--name <setting>]
```

Options

If no options are specified, all settings are displayed.

--name <setting>

Displays only the value of the specified setting.

The following values are valid:

port

The port through which a Data Management Application (DMA) connects.

dma

The DMA vendor that the cluster is currently configured to interact with.

Examples

To view a list of NDMP settings and values, run the following command:

```
isi ndmp settings list
```

The system displays output similar to the following example:

Setting	Value
port	10000
dma	EMC

isi ndmp settings variables create

Sets the default value for an NDMP environment variable for a given path.

Syntax

```
isi ndmp settings variables create <path> <name> <value>
```

For a list of available environment variables, see [NDMP environment variables on page 635](#).

Options

<path>

Applies the default NDMP environment variable value to the specified path.

<name>

Specifies the NDMP environment variable to define.

<value>

Specifies the value to be applied to the NDMP environment variable.

Examples

The following command causes snapshot-based incremental backups to be performed for `/ifs/data/media` by default:

```
isi ndmp settings variables create /ifs/data/media BACKUP_MODE
SNAPSHOT
```

isi ndmp settings variables modify

Modifies the default value for an NDMP environment variable for a given path.

Syntax

```
isi ndmp settings variables modify <path> <name> <value>
```

Options

For a list of available environment variables, see [NDMP environment variables on page 635](#).

<path>

Applies the default NDMP-environment-variable value to the specified path.

<name>

Specifies the NDMP environment variable to be defined.

<value>

Specifies the value to be applied to the NDMP environment variable.

Examples

The following command sets the default file history for backing up `/ifs/data/media` to be path-based:

```
isi ndmp settings variables modify /ifs/data/media HIST F
```

isi ndmp settings variables delete

Deletes the default value for an NDMP environment variable for a given path.

Syntax

```
isi ndmp settings variables delete <path>
[--name <variable>]
```

Options

<path>

Applies the default NDMP-environment-variable value to the specified path.

--name *<variable>*

Deletes the default value for the specified NDMP environment variable. The following values are valid:

- FILESYSTEM
- LEVEL
- UPDATE
- HIST
- DIRECT
- FILES
- EXCLUDE
- ENCODING
- RESTORE_HARDLINK_BY_TABLE

If this option is not specified, deletes default values for all NDMP environment variables for the given directory.

Examples

The following command removes all default NDMP settings for `/ifs/data/media`:

```
isi ndmp settings variables delete /ifs/data/media
```

The following command removes the default file-history setting for backing up `/ifs/data/media`:

```
isi ndmp settings variables delete /ifs/data/media --name HIST
```

isi ndmp settings variables list

Displays default values for NDMP environment variables for directory paths.

Syntax

```
isi ndmp settings variables list
  [--path <path>]
```

Options

`--path <path>`

Applies the default NDMP-environment-variable value to the specified path.

Examples

To view default values for NDMP environment variables for directory paths, run the following command:

```
isi ndmp settings variables list
```

The system displays output similar to the following example:

Path	Name	Value
/ifs/data/media	HIST	F
/ifs/data/media	BACKUP_MODE	SNAPSHOT

isi ndmp dumpdates delete

Deletes a snapshot created for a snapshot-based incremental backup.

Syntax

```
isi ndmp dumpdates delete <path>
  [--level <integer>]
```

Options

`<path>`

Deletes a dumpdate entry for a backup of the specified directory.

`--level <integer>`

Deletes a dumpdate entry for a backup of the specified level for the given directory. If this option is not specified, deletes all dumpdate entries for the given directory.

Examples

The following command deletes the dumpdate entry for a level 0 backup of `/ifs/data/media`:

```
isi ndmp dumpdates delete /ifs/data/media --level 0
```

isi ndmp dumpdates list

Displays snapshots created for snapshot-based incremental backups.

Syntax

```
isi ndmp dumpdates list
  [--path <path>]
```

Options`--path <path>`

Displays only dumpdate entries that relate to the specified file path.

Examples

To view NDMP dumpdate entries, run the following command:

```
isi ndmp dumpdates list
```

The system displays output similar to the following example:

Date	Lvl	SnapID	Path
Wed May 29 12:06:26 2013	0	18028	/ifs/tmp/backup
Wed May 29 12:20:56 2013	1	18030	/ifs/tmp/backup

If a snapshot was created for a non-snapshot-based incremental backup, the snapshot ID is 0.

isi ndmp extensions context delete

Deletes a restartable backup context for an NDMP backup.

Note

It is recommended that you do not manually delete restartable backup contexts. Manually deleting a backup context requires you to restart the corresponding NDMP backup from the beginning.

Syntax

```
isi ndmp extensions contexts delete <id>
```

Options`<id>`

Deletes the specified restartable backup context.

Examples

The following command deletes a restartable backup context:

```
isi ndmp extensions contexts delete 533089ed-c4c5-11e2-  
bad5-001b21a2c2dc
```

isi ndmp extensions contexts list

Displays restartable backup contexts for NDMP backups.

Syntax

```
isi ndmp extensions contexts list  
[--id <id>]
```

Options`--id <id>`

Displays only the restartable backup context of the specified ID.

Examples

To view restartable backup contexts, run the following command:

```
isi ndmp extensions contexts list
```

The system displays output similar to the following example:

```
NDMP BRE Contexts
-----
2a9a9e11-ac6c-11e2-b40d-0025904b58c6
533089ed-c4c5-11e2-bad5-001b21a2c2dc
f7224734-c7ca-11e2-bad5-001b21979310
524e8c43-c8bb-11e2-bad5-001b21979310
85074911-c8bb-11e2-bad5-001b21a2c2dc
fff87d24-c937-11e2-bad5-001b21a2c2dc
a97b2561-c938-11e2-bad5-001b21a9e862
32d06e84-c985-11e2-bad5-001b21a9e862
c4904d84-c986-11e2-bad5-001b21a2c2dc
9bc8494a-ca20-11e2-bad5-001b21a2c2dc
9f3a85da-ca21-11e2-bad5-001b21a9e862
906c60a6-ca24-11e2-bad5-001b21a2c2dc
```

isi ndmp extensions contexts view

Displays detailed information about a restartable backup context.

Syntax

```
isi ndmp extensions contexts view
[--id <id>]
```

Options

--id <id>

Displays information about the restartable backup context of the specified ID.

Examples

The following command displays information about a restartable backup context:

```
isi ndmp extensions contexts view 2a9a9e11-ac6c-11e2-b40d-0025904b58c6
```

The system displays output similar to the following example:

```
Id : 2a9a9e11-ac6c-11e2-b40d-0025904b58c6
    Snap-name : ndmp_backup_b7b9f67507dbf5cbe5e5a732
    Timestamp : 1366759134.6
    Status : 13
    Results : 0
    VERBOSE : Y
    LEVEL : 0
    UPDATE : Y
    HIST : F
    TW_NUM_STATS : 32
FT_READER_THREADS : 40
  NUM_TW_WORKERS : 20
    FILESYSTEM : /ifs/data/qadata/userdata
  NUM_CDB_WORKERS : 20
    TYPE : tar
    RECURSIVE : Y
```

isi ndmp extensions settings modify

Determines the maximum number of restartable backup contexts that are retained on a cluster.

Syntax

```
isi ndmp extensions settings modify --bre_max_contexts <integer>
```

Options

`--bre_max_contexts <integer>`

Specifies the maximum number of restartable backup contexts. Specify as an integer from 1 to 1024.

Examples

The following command sets the maximum number of restartable backup contexts to 500:

```
isi ndmp extensions settings modify --bre_max_contexts 500
```

isi ndmp extensions settings view

Displays the maximum number of restartable backup contexts that are retained on a cluster.

Syntax

```
isi ndmp extensions settings view
```

Options

There are no options for this command.

Examples

To view the maximum number of restartable backup contexts, run the following command:

```
isi ndmp extensions settings view
```

CHAPTER 16

File retention with SmartLock

This section contains the following topics:

- [SmartLock overview](#)..... 656
- [Compliance mode](#)..... 656
- [SmartLock directories](#).....656
- [Replication and backup with SmartLock](#)..... 657
- [SmartLock license functionality](#)..... 658
- [SmartLock considerations](#)..... 658
- [Set the compliance clock](#)..... 659
- [View the compliance clock](#)..... 659
- [Creating a SmartLock directory](#)..... 659
- [Managing SmartLock directories](#)..... 661
- [Managing files in SmartLock directories](#)..... 665
- [SmartLock commands](#).....668

SmartLock overview

You can prevent users from modifying and deleting files on an EMC Isilon cluster with the SmartLock software module. You must activate a SmartLock license on a cluster to protect data with SmartLock.

With the SmartLock software module, you can create SmartLock directories and commit files within those directories to a write once read many (WORM) state. You cannot erase or re-write a file committed to a WORM state. After a file is removed from a WORM state, you can delete the file. However, you can never modify a file that has been committed to a WORM state, even after it is removed from a WORM state.

Compliance mode

SmartLock compliance mode enables you to protect your data in compliance with the regulations defined by U.S. Securities and Exchange Commission rule 17a-4. You can upgrade a cluster to compliance mode during the initial cluster configuration process, before you activate the SmartLock license. To upgrade a cluster to SmartLock compliance mode after the initial cluster configuration process, contact Isilon Technical Support.

If you upgrade a cluster to compliance mode, you will not be able to log in to that cluster through the root user account. Instead, you can log in to the cluster through the compliance administrator account that is configured either during initial cluster configuration or when the cluster is upgraded to compliance mode. If you are logged in through the compliance administrator account, you can perform administrative tasks through the `sudo` command.

SmartLock directories

In a SmartLock directory, you can commit a file to a WORM state manually or you can configure SmartLock to automatically commit the file. You can create two types of SmartLock directories: enterprise and compliance. However, you can create compliance directories only if the cluster has been upgraded to SmartLock compliance mode. Before you can create SmartLock directories, you must activate a SmartLock license on the cluster.

Enterprise directories enable you to protect your data without restricting your cluster to comply with regulations defined by U.S. Securities and Exchange Commission rule 17a-4. If you commit a file to a WORM state in an enterprise directory, the file can never be modified and cannot be deleted until the retention period passes. However, if you are logged in through the root user account, you can delete the file before the retention period passes through the privileged delete feature. The privileged delete feature is not available for compliance directories. Enterprise directories reference the system clock to facilitate time-dependent operations, including file retention.

Compliance directories enable you to protect your data in compliance with the regulations defined by U.S. Securities and Exchange Commission rule 17a-4. If you commit a file to a WORM state in a compliance directory, the file cannot be modified or deleted before the specified retention period has expired. You cannot delete committed files, even if you are logged in to the compliance administrator account. Compliance directories reference the compliance clock to facilitate time-dependent operations, including file retention.

You must set the compliance clock before you can create compliance directories. You can set the compliance clock only once. After you set the compliance clock, you cannot

modify the compliance clock time. The compliance clock is controlled by the compliance clock daemon. Because root and compliance administrator users can disable the compliance clock daemon, it is possible for those users to increase the retention period of WORM committed files in compliance mode. However, it is not possible to decrease the retention period of a WORM committed file.

Replication and backup with SmartLock

You must ensure that SmartLock directories remain protected during replication and backup operations.

If you are replicating SmartLock directories with SyncIQ, it is recommended that you configure all nodes on the source and target clusters into Network Time Protocol (NTP) peer mode to ensure that the node clocks are synchronized. For compliance clusters, it is recommended that you configure all nodes on the source and target clusters into NTP peer mode before you set the compliance clock to ensure that the compliance clocks are initially set to the same time.

Note

If you replicate data to a SmartLock directory, do not configure SmartLock settings for that directory until you are no longer replicating data to the directory. Configuring an autocommit time period for a SmartLock directory that you are replicating to can cause replication jobs to fail. If the target directory commits a file to a WORM state, and the file is modified on the source cluster, the next replication job will fail because it cannot update the file.

SmartLock replication and backup limitations

Be aware of the limitations of replicating and backing up SmartLock directories with SyncIQ and NDMP.

If the source or target directory of a SyncIQ policy is a SmartLock directory, replication might not be allowed. For more information, see the following table:

Source directory type	Target directory type	Allowed
Non-SmartLock	Non-SmartLock	Yes
Non-SmartLock	SmartLock enterprise	Yes
Non-SmartLock	SmartLock compliance	No
SmartLock enterprise	Non-SmartLock	Yes; however, retention dates and commit status of files will be lost.
SmartLock enterprise	SmartLock enterprise	Yes
SmartLock enterprise	SmartLock compliance	No
SmartLock compliance	Non-SmartLock	No
SmartLock compliance	SmartLock enterprise	No
SmartLock compliance	SmartLock compliance	Yes

If you are replicating a SmartLock directory to another SmartLock directory, you must create the target SmartLock directory prior to running the replication policy. Although OneFS will create a target directory automatically if a target directory does not already

exist, OneFS will not create a target SmartLock directory automatically. If you attempt to replicate an enterprise directory before the target directory has been created, OneFS will create a non-SmartLock target directory and the replication job will succeed. If you replicate a compliance directory before the target directory has been created, the replication job will fail.

If you replicate SmartLock directories to another cluster with SyncIQ, the WORM state of files is replicated. However, SmartLock directory configuration settings are not transferred to the target directory.

For example, if you replicate a directory that contains a committed file that is set to expire on March 4th, the file is still set to expire on March 4th on the target cluster. However, if the directory on the source cluster is set to prevent files from being committed for more than a year, the target directory is not automatically set to the same restriction.

If you back up data to an NDMP device, all SmartLock metadata relating to the retention date and commit status is transferred to the NDMP device. If you restore data to a SmartLock directory on the cluster, the metadata persists on the cluster. However, if the directory that you restore to is not a SmartLock directory, the metadata is lost. You can restore to a SmartLock directory only if the directory is empty.

SmartLock license functionality

You must activate a SmartLock license on a cluster before you can create SmartLock directories and commit files to a WORM state.

If a SmartLock license becomes inactive, you will not be able to create new SmartLock directories on the cluster, modify SmartLock directory configuration settings, or delete files committed to a WORM state in enterprise directories before their expiration dates. However, you can still commit files within existing SmartLock directories to a WORM state.

If a SmartLock license becomes inactive on a cluster that is running in SmartLock compliance mode, root access to the cluster is not restored.

SmartLock considerations

- If a file is owned exclusively by the root user, and the file exists on a cluster that is in SmartLock compliance mode, the file will be inaccessible, because the root user account is disabled in compliance mode. For example, this can happen if a file is assigned root ownership on a cluster that has not been upgraded to compliance mode, and then the file is replicated to a cluster in compliance mode. This can also occur if a file is assigned root ownership before a cluster is upgraded to SmartLock compliance mode or if a root-owned file is restored on a compliance cluster after being backed up.
- It is recommended that you create files outside of SmartLock directories and then transfer them into a SmartLock directory after you are finished working with the files. If you are uploading files to a cluster, it is recommended that you upload the files to a non-SmartLock directory, and then later transfer the files to a SmartLock directory. If a file is committed to a WORM state while the file is being uploaded, the file will become trapped in an inconsistent state. Files can be committed to a WORM state while they are still open. If you specify an autocommit time period for a directory, the autocommit time period is calculated according to the length of time since the file was last modified, not when the file was closed. If you delay writing to an open file for more than the autocommit time period, the file will be automatically committed to a WORM state, and you will not be able to write to the file.

- In a Microsoft Windows environment, if you commit a file to a WORM state, you can no longer modify the hidden or archive attributes of the file. Any attempt to modify the hidden or archive attributes of a WORM committed file will generate an error. This can prevent third-party applications from modifying the hidden or archive attributes.

Set the compliance clock

Before you can create SmartLock compliance directories, you must set the compliance clock. This procedure is available only through the command-line interface (CLI).

Setting the compliance clock configures the clock to the same time as the cluster system clock. Before you set the compliance clock, ensure that the system clock is set to the correct time. If the compliance clock later becomes unsynchronized with the system clock, the compliance clock will slowly correct itself to match the system clock. The compliance clock corrects itself at a rate of approximately one week per year.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in through the compliance administrator account.
2. Set the compliance clock by running the following command.

```
isi worm cdate set
```

View the compliance clock

You can view the current time of the compliance clock. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in through the compliance administrator account.
2. View the compliance clock by running the following command:

```
isi worm cdate view
```

Creating a SmartLock directory

You can create a SmartLock directory and configure settings that control how long files are retained in a WORM state and when files are automatically committed to a WORM state. You cannot move or rename a directory that contains a SmartLock directory.

It is recommended that you set SmartLock configuration settings only once and do not modify the settings after files have been added to the SmartLock directory.

Retention periods

A retention period is the length of time that a file remains in a WORM state before being released from a WORM state. You can configure SmartLock directory settings that enforce default, maximum, and minimum retention periods for the directory.

If you manually commit a file, you can optionally specify the date that the file is released from a WORM state. You can configure a minimum and a maximum retention period for a SmartLock directory to prevent files from being retained for too long or too short a time period. It is recommended that you specify a minimum retention period for all SmartLock directories.

For example, assume that you have a SmartLock directory with a minimum retention period of two days. At 1:00 PM on Monday, you commit a file to a WORM state, and specify the file to be released from a WORM state on Tuesday at 3:00 PM. The file will be released from a WORM state two days later on Wednesday at 1:00 PM, because releasing the file earlier would violate the minimum retention period.

You can also configure a default retention period that is assigned when you commit a file without specifying a date to release the file from a WORM state.

Autocommit time periods

You can configure an autocommit time period for SmartLock directories. An autocommit time period causes files that have been in a SmartLock directory for a period of time without being modified to be automatically committed to a WORM state.

If you modify the autocommit time period of a SmartLock directory that contains uncommitted files, the new autocommit time period is immediately applied to the files that existed before the modification. For example, consider a SmartLock directory with an autocommit time period of 2 hours. If you modify a file in the SmartLock directory at 1:00 PM, and you decrease the autocommit time period to 1 hour at 2:15 PM, the file is instantly committed to a WORM state.

If a file is manually committed to a WORM state, the read-write permissions of the file are modified. However, if a file is automatically committed to a WORM state, the read-write permissions of the file are not modified.

Create a SmartLock directory

You can create a SmartLock directory and commit files in that directory to a WORM state. This procedure is available only through the command-line interface (CLI).

Before creating a SmartLock directory, be aware of the following conditions and requirements:

- You cannot create a SmartLock directory as a subdirectory of an existing SmartLock directory.
- Hard links cannot cross SmartLock directory boundaries.
- Creating a SmartLock directory causes a corresponding SmartLock domain to be created for that directory.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi worm domains create` command.

If you specify the path of an existing directory, the directory must be empty.

The following command creates a compliance directory with a default retention period of four years, a minimum retention period of three years, and a maximum retention period of five years:

```
sudo isi worm domains create /ifs/data/SmartLock/directory1 \
--compliance --default-retention 4Y --min-retention 3Y \
--max-retention 5Y --mkdir
```

The following command creates an enterprise directory with an autocommit time period of thirty minutes and a minimum retention period of three months:

```
isi worm domains create /ifs/data/SmartLock/directory2 \
--autocommit-offset 30m --min-retention 3M --mkdir
```

Managing SmartLock directories

You can modify SmartLock directory settings only 32 times per directory. SmartLock directory settings include the default, minimum, and maximum retention period and the autocommit time period.

A SmartLock directory can be renamed only if the directory is empty.

Modify a SmartLock directory

You can modify the SmartLock configuration settings for a SmartLock directory. This procedure is available only through the command-line interface (CLI).

Note

You can modify SmartLock directory settings only 32 times per directory. It is recommended that you set SmartLock configuration settings only once and do not modify the settings after files are added to the SmartLock directory.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Modify SmartLock configuration settings by running the `isi worm modify` command.

The following command sets the default retention period to one year:

```
isi worm domains modify /ifs/data/SmartLock/directory1 \
--default-retention 1Y
```

View SmartLock directory settings

You can view the SmartLock directory settings for SmartLock directories. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. View all SmartLock domains by running the following command:

```
isi worm domains list
```

The system displays output similar to the following example:

```

ID      Path                                     Type
-----
65536   /ifs/data/SmartLock/directory1 enterprise
65537   /ifs/data/SmartLock/directory2 enterprise
65538   /ifs/data/SmartLock/directory3 enterprise
-----
```

3. (Optional) To view detailed information about a specific SmartLock directory, run the `isi worm domains view` command.

The following command displays detailed information about `/ifs/data/SmartLock/directory2`:

```
isi worm domains view /ifs/data/SmartLock/directory2
```

The system displays output similar to the following example:

```

        ID: 65537
        Path: /ifs/data/SmartLock/directory2
        Type: enterprise
        LIN: 4295426060
Autocommit Offset: 30m
  Override Date: -
Privileged Delete: off
Default Retention: 1Y
  Min Retention: 3M
  Max Retention: -
  Total Modifies: 3/32 Max

```

SmartLock directory configuration settings

You can configure SmartLock directory settings that determine when files are committed to and how long files are retained in a WORM state.

ID

The numerical ID of the corresponding SmartLock domain.

Path

The path of the directory.

Type

The type of SmartLock directory.

LIN

The inode number of the directory.

Autocommit offset

The autocommit time period for the directory. After a file exists in this SmartLock directory without being modified for the specified time period, the file is automatically committed to a WORM state.

Times are expressed in the format "*<integer> <time>*", where *<time>* is one of the following values:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

m

Specifies minutes

s

Specifies seconds

Override date

The override retention date for the directory. Files committed to a WORM state are not released from a WORM state until after the specified date, regardless of the maximum retention period for the directory or whether a user specifies an earlier date to release a file from a WORM state.

Privileged delete

Indicates whether files in the directory can be deleted through the privileged delete functionality.

on

The root user can delete files committed to a WORM state by running the `isi worm files delete` command.

off

WORM committed files cannot be deleted, even through the `isi worm files delete` command.

disabled

WORM committed files cannot be deleted, even through the `isi worm files delete` command. After this setting is applied, it cannot be modified.

Default retention period

The default retention period for the directory. If a user does not specify a date to release a file from a WORM state, the default retention period is assigned. Times are expressed in the format "*<integer> <time>*", where *<time>* is one of the following values:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

m

Specifies minutes

s

Specifies seconds

`Forever` indicates that WORM committed files are retained permanently by default. `Use Min` indicates that the default retention period is equal to the minimum retention date. `Use Max` indicates that the default retention period is equal to the maximum retention date.

Minimum retention period

The minimum retention period for the directory. Files are retained in a WORM state for at least the specified amount of time, even if a user specifies an expiration date that results in a shorter retention period.

Times are expressed in the format "*<integer> <time>*", where *<time>* is one of the following values:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours
- m** Specifies minutes
- s** Specifies seconds

`Forever` indicates that all WORM committed files are retained permanently.

Maximum retention period

The maximum retention period for the directory. Files cannot be retained in a WORM state for more than the specified amount of time, even if a user specifies an expiration date that results in a longer retention period.

Times are expressed in the format "*<integer> <time>*", where *<time>* is one of the following values:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours
- m** Specifies minutes
- s** Specifies seconds

`Forever` indicates that there is no maximum retention period.

Total modifies

The total number of times that SmartLock settings have been modified for the directory. You can modify SmartLock settings only 32 times per directory.

Managing files in SmartLock directories

You can commit files in SmartLock directories to a WORM state by removing the read-write privileges of the file. You can also set a specific date at which the retention period of the file expires. Once a file is committed to a WORM state, you can increase the retention period of the file, but you cannot decrease the retention period of the file. You cannot move a file that has been committed to a WORM state, even after the retention period for the file has expired.

The retention period expiration date is set by modifying the access time of a file. In a UNIX command line, the access time can be modified through the `touch` command. Although there is no method of modifying the access time through Windows Explorer, you can modify the access time through Windows Powershell. Accessing a file does not set the retention period expiration date.

If you run the `touch` command on a file in a SmartLock directory without specifying a date on which to release the file from a SmartLock state, and you commit the file, the retention period is automatically set to the minimum retention period specified for the SmartLock directory. If you have not specified a minimum retention period for the SmartLock directory, the file is assigned a retention period of zero seconds. It is recommended that you specify a minimum retention period for all SmartLock directories.

Set a retention period through a UNIX command line

You can specify when a file will be released from a WORM state through a UNIX command line.

Procedure

1. Open a connection to any node in the cluster through a UNIX command line and log in.
2. Set the retention period by modifying the access time of the file through the `touch` command.

The following command sets an expiration date of June 1, 2015 for `/ifs/data/test.txt`:

```
touch -at 201506010000 /ifs/data/test.txt
```

Set a retention period through Windows Powershell

You can specify when a file will be released from a WORM state through Microsoft Windows Powershell.

Procedure

1. Open the Windows PowerShell command prompt.
2. (Optional) Establish a connection to the cluster by running the `net use` command.

The following command establishes a connection to the `/ifs` directory on `cluster.ip.address.com`:

```
net use "\\cluster.ip.address.com\ifs" /user:root password
```

3. Specify the name of the file you want to set a retention period for by creating an object.

The file must exist in a SmartLock directory.

The following command creates an object for `/smartlock/file.txt`:

```
$file = Get-Item "\\cluster.ip.address.com\ifs\smartlock\file.txt"
```

4. Specify the retention period by setting the last access time for the file.

The following command sets an expiration date of July 1, 2015 at 1:00 PM:

```
$file.LastAccessTime = Get-Date "2015/7/1 1:00 pm"
```

Commit a file to a WORM state through a UNIX command line

You can commit a file to a WORM state through a UNIX command line.

To commit a file to a WORM state, you must remove all write privileges from the file. If a file is already set to a read-only state, you must first add write privileges to the file, and then return the file to a read-only state.

Procedure

1. Open a connection to the cluster through a UNIX command line interface and log in.
2. Remove write privileges from a file by running the `chmod` command.

The following command removes write privileges of `/ifs/data/smartlock/file.txt`:

```
chmod ugo-w /ifs/data/smartlock/file.txt
```

Commit a file to a WORM state through Windows Explorer

You can commit a file to a WORM state through Microsoft Windows Explorer. This procedure describes how to commit a file through Windows 7.

To commit a file to a WORM state, you must apply the read-only setting. If a file is already set to a read-only state, you must first remove the file from a read-only state and then return it to a read-only state.

Procedure

1. In Windows Explorer, navigate to the file you want to commit to a WORM state.
2. Right-click the folder and then click **Properties**.
3. In the **Properties** window, click the **General** tab.
4. Select the **Read-only** check box, and then click **OK**.

Override the retention period for all files in a SmartLock directory

You can override the retention period for files in a SmartLock directory. All files committed to a WORM state within the directory will remain in a WORM state until after the specified day. This procedure is available only through the command-line interface (CLI).

If files are committed to a WORM state after the retention period is overridden, the override date functions as a minimum retention date. All files committed to a WORM state do not expire until at least the given day, regardless of user specifications.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.

2. Override the retention period expiration date for all WORM committed files in a SmartLock directory by running the `isi worm modify` command.

For example, the following command overrides the retention period expiration date of `/ifs/data/SmartLock/directory1` to June 1, 2014:

```
isi worm domains modify /ifs/data/SmartLock/directory1 \
--override-date 2014-06-01
```

Delete a file committed to a WORM state

You can delete a WORM committed file before the expiration date only if you are logged in as the root user or compliance administrator. This procedure is available only through the command-line interface (CLI).

Before you begin

Privileged delete functionality must not be permanently disabled for the SmartLock directory that contains the file.

Procedure

1. Open a connection to the cluster through a UNIX command line and log in through either the root user or compliance administrator account.
2. If privileged delete functionality was disabled for the SmartLock directory, modify the directory by running the `isi worm domains modify` command with the `--privileged-delete` option.

The following command enables privileged delete for `/ifs/data/SmartLock/directory1`:

```
isi worm domains modify /ifs/data/SmartLock/directory1 \
--privileged-delete true
```

3. Delete the WORM committed file by running the `isi worm files delete` command.

The following command deletes `/ifs/data/SmartLock/directory1/file`:

```
isi worm files delete /ifs/data/SmartLock/directory1/file
```

The system displays output similar to the following:

```
Are you sure? (yes, [no]):
```

4. Type **yes** and then press ENTER.

View WORM status of a file

You can view the WORM status of an individual file. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a connection to the cluster through a UNIX command line.
2. View the WORM status of a file by running the `isi worm files view` command.

For example, the following command displays the WORM status of a file:

```
isi worm files view /ifs/data/SmartLock/directory1/file
```

The system displays output similar to the following:

```

WORM Domains
ID      Root Path
-----
65539  /ifs/data/SmartLock/directory1

WORM State: COMMITTED
Expires: 2015-06-01T00:00:00

```

SmartLock commands

You can control file retention through the WORM commands. WORM commands apply specifically to the SmartLock tool, and are available only if you have activated a SmartLock license on the cluster.

isi worm domains create

Creates a SmartLock directory.

Syntax

```

isi worm domains create <path>
  [--compliance]
  [--autocommit-offset <duration>]
  [--override-date <timestamp>]
  [--privileged-delete {true | false}]
  [--default-retention {<duration> | forever | use_min
  | use_max}]
  [--min-retention {<duration> | forever}]
  [--max-retention <duration>]
  [--mkdir]
  [--force]
  [--verbose]

```

Options

<path>

Creates a SmartLock directory at the specified path.
Specify as a directory path.

{--compliance | -C}

Specifies the SmartLock directory as a SmartLock compliance directory. This option is valid only on clusters running in SmartLock compliance mode.

{--autocommit-offset | -a} <duration>

Specifies an autocommit time period. After a file exists in a SmartLock directory without being modified for the specified length of time, the file automatically committed to a WORM state.

Specify <duration> in the following format:

```
<integer><units>
```

Specify <units> are valid:

Y

Specifies years

M

Specifies months

- W** Specifies weeks
- D** Specifies days
- H** Specifies hours
- m** Specifies minutes
- s** Specifies seconds

To specify no autocommit time period, specify `none`. The default value is `none`.

`{--override-date | -o} <timestamp>`

Specifies an override retention date for the directory. Files committed to a WORM state are not released from a WORM state until after the specified date, regardless of the maximum retention period for the directory or whether a user specifies an earlier date to release a file from a WORM state.

Specify `<timestamp>` in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

`{--privileged-delete | -p} {true | false}`

Determines whether files in the directory can be deleted through the `isi worm files delete` command. This option is available only for SmartLock enterprise directories.

The default value is `false`.

`--disable-privileged-delete`

Permanently prevents WORM committed files from being deleted from the SmartLock directory.

Note

If you specify this option, you can never enable the privileged delete functionality for the directory. If a file is then committed to a WORM state in the directory, you will not be able to delete the file until the retention period has passed.

`{--default-retention | -d} {<duration> | forever | use_min | use_max}`

Specifies a default retention period. If a user does not explicitly assign a retention period expiration date, the default retention period is assigned to the file when it is committed to a WORM state.

Specify `<duration>` in the following format:

```
<integer><units>
```

Specify `<units>` are valid:

- Y** Specifies years

M
Specifies months

W
Specifies weeks

D
Specifies days

H
Specifies hours

m
Specifies minutes

s
Specifies seconds

To permanently retain WORM committed files by default, specify `forever`. To assign the minimum retention period as the default retention period, specify `use_min`. To assign the maximum retention period as the default retention period, specify `use_max`.

```
{--min-retention | -m} {<duration> | forever}
```

Specifies a minimum retention period. Files are retained in a WORM state for at least the specified amount of time.

Specify `<duration>` in the following format:

```
<integer><units>
```

Specify `<units>` as one of the following values:

Y
Specifies years

M
Specifies months

W
Specifies weeks

D
Specifies days

H
Specifies hours

m
Specifies minutes

s
Specifies seconds

To permanently retain all WORM committed files, specify `forever`.

```
{--max-retention | -x} {<duration> | forever}
```

Specifies a maximum retention period. Files cannot be retained in a WORM state for more than the specified amount of time, even if a user specifies an expiration date that results in a longer retention period.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* as one of the following values:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To specify no maximum retention period, specify *forever*.

```
{--mkdir | -M}
```

Creates the specified directory if it does not already exist.

```
{--force | -f}
```

Does not prompt you to confirm the creation of the SmartLock directory.

```
{--verbose | -v}
```

Displays more detailed information.

isi worm domains modify

Modifies SmartLock settings of a SmartLock directory.

Syntax

```
isi worm domains modify <domain>
  [--compliance]
  [--autocommit-offset <duration> | --clear-autocommit-offset]
  [--override-date <timestamp> | --clear-override-date]
  [--privileged-delete {true | false}]
  [--default-retention {<duration> | forever | use_min
  | use_max} | --clear-default-retention]
  [--min-retention {<duration> | forever} | --clear-min-retention]
  [--max-retention <duration> | --clear-max-retention]
  [--force]
  [--verbose]
```

Options

```
<domain>
```

Modifies the specified SmartLock directory.

Specify as a directory path, ID, or LIN of a SmartLock directory.

```
{--compliance | -C}
```

Specifies the SmartLock directory as a SmartLock compliance directory. This option is valid only on clusters running in SmartLock compliance mode.

```
{--autocommit-offset | -a} <duration>
```

Specifies an autocommit time period. After a file exists in a SmartLock directory without being modified for the specified length of time, the file automatically committed to a WORM state.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

m

Specifies minutes

s

Specifies seconds

To specify no autocommit time period, specify `none`. The default value is `none`.

```
--clear-autocommit-offset
```

Removes the autocommit time period for the given SmartLock directory.

```
{--override-date | -o} <timestamp>
```

Specifies an override retention date for the directory. Files committed to a WORM state are not released from a WORM state until after the specified date, regardless of the maximum retention period for the directory or whether a user specifies an earlier date to release a file from a WORM state.

Specify *<timestamp>* in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

```
--clear-override-date
```

Removes the override retention date for the given SmartLock directory.

```
{--privileged-delete | -p} {true | false}
```

Determines whether files in the directory can be deleted through the `isi worm files delete` command. This option is available only for SmartLock enterprise directories.

The default value is `false`.

`--disable-privileged-delete`

Permanently prevents WORM committed files from being deleted from the SmartLock directory.

Note

If you specify this option, you can never enable the privileged delete functionality for the SmartLock directory. If a file is then committed to a WORM state in the directory, you will not be able to delete the file until the retention period expiration date has passed.

`{--default-retention | -d} {<duration> | forever | use_min | use_max}`

Specifies a default retention period. If a user does not explicitly assign a retention period expiration date, the default retention period is assigned to the file when it is committed to a WORM state.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* are valid:

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

H

Specifies hours

m

Specifies minutes

s

Specifies seconds

To permanently retain WORM committed files by default, specify *forever*. To assign the minimum retention period as the default retention period, specify *use_min*. To assign the maximum retention period as the default retention period, specify *use_max*.

`--clear-default-retention`

Removes the default retention period for the given SmartLock directory.

`{--min-retention | -m} {<duration> | forever}`

Specifies a minimum retention period. Files are retained in a WORM state for at least the specified amount of time.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* as one of the following values:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours
- m** Specifies minutes
- s** Specifies seconds

To permanently retain all WORM committed files, specify *forever*.

`--clear-min-retention`

Removes the minimum retention period for the given SmartLock directory.

`{--max-retention | -x} {<duration> | forever}`

Specifies a maximum retention period. Files cannot be retained in a WORM state for more than the specified amount of time, even if a user specifies an expiration date that results in a longer retention period.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* as one of the following values:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours
- m** Specifies minutes
- s** Specifies seconds

To specify no maximum retention period, specify *forever*.

`--clear-max-retention`

Removes the maximum retention period for the given SmartLock directory.

{--force | -f}

Does not prompt you to confirm the creation of the SmartLock directory.

{--verbose | -v}

Displays more detailed information.

isi worm domains list

Displays a list of WORM directories.

Syntax

```
isi worm domains list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

id

Sorts output by the SmartLock directory ID.

path

Sorts output by the path of the SmartLock directory.

type

Sorts output based on whether the SmartLock directory is a compliance directory.

lin

Sorts output by the inode number of the SmartLock directory.

autocommit_offset

Sorts output by the autocommit time period of the SmartLock directory.

override_date

Sorts output by the override retention date of the SmartLock directory.

privileged_delete

Sorts output based on whether the privileged delete functionality is enabled for the SmartLock directory.

default_retention

Sorts output by the default retention period of the SmartLock directory.

min_retention

Sorts output by the minimum retention period of the SmartLock directory.

max_retention

Sorts output by the maximum retention period of the SmartLock directory.

total_modifies

Sorts output by the total number of times that the SmartLock directory has been modified.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

isi worm domains view

Displays WORM information about a specific directory or file.

Syntax

```
isi worm domains view <domain>
```

Options

<domain>

Displays information about the specified SmartLock directory.

Specify as a directory path, ID, or LIN of a SmartLock directory.

isi worm cdate set

Sets the SmartLock compliance clock to the current time on the system clock.

⚠ CAUTION

You can set the compliance clock only once. After the compliance clock has been set, you cannot modify the compliance clock time.

Syntax

```
isi worm cdate set
```

Options

There are no options for this command.

isi worm cdate view

Displays whether or not the SmartLock compliance clock is set. If the compliance clock is set, displays the current time on the compliance clock.

Syntax

```
isi worm cdate view
```

Options

There are no options for this command.

isi worm files delete

Deletes a file committed to a WORM state. This command can be run only by the root user or compliance administrator.

Syntax

```
isi worm files delete <path>
  [--force]
  [--verbose]
```

Options

<path>

Deletes the specified file. The file must exist in a SmartLock enterprise directory with the privileged delete functionality enabled.

Specify as a file path.

--force

Does not prompt you to confirm that you want to delete the file.

--verbose

Displays more detailed information.

isi worm files view

Displays information about a file committed to a WORM state.

Syntax

```
isi worm files view <path>
  [--no-symlinks]
```

Options

<path>

Displays information about the specified file. The file must be committed to a WORM state.

Specify as a file path.

--no-symlinks

If <path> refers to a file, and the given file is a symbolic link, displays WORM information about the symbolic link. If this option is not specified, and the file is a symbolic link, displays WORM information about the file that the symbolic link refers to.

CHAPTER 17

Protection domains

This section contains the following topics:

- [Protection domains overview](#)..... 680
- [Protection domain considerations](#)..... 680
- [Create a protection domain](#) 681
- [Delete a protection domain](#) 681

Protection domains overview

Protection domains are markers that prevent modifications to files and directories. If a domain is applied to a directory, the domain is also applied to all of the files and subdirectories under the directory. You can specify domains manually; however, OneFS usually creates domains automatically.

There are three types of domains: SyncIQ, SmartLock, and SnapRevert. SyncIQ domains can be assigned to source and target directories of replication policies. OneFS automatically creates a SyncIQ domain for the target directory of a replication policy the first time that the policy is run. OneFS also automatically creates a SyncIQ domain for the source directory of a replication policy during the failback process. You can manually create a SyncIQ domain for a source directory before you initiate the failback process, but you cannot delete a SyncIQ domain that mark the target directory of a replication policy.

SmartLock domains are assigned to SmartLock directories to prevent committed files from being modified or deleted. OneFS automatically creates a SmartLock domain when a SmartLock directory is created. You cannot delete a SmartLock domain. However, if you delete a SmartLock directory, OneFS automatically deletes the SmartLock domain associated with the directory.

SnapRevert domains are assigned to directories that are contained in snapshots to prevent files and directories from being modified while a snapshot is being reverted. OneFS does not automatically create SnapRevert domains. You cannot revert a snapshot until you create a SnapRevert domain for the directory that the snapshot contains. You can create SnapRevert domains for subdirectories of directories that already have SnapRevert domains. For example, you could create SnapRevert domains for both `/ifs/data` and `/ifs/data/archive`. You can delete a SnapRevert domain if you no longer want to revert snapshots of a directory.

Protection domain considerations

You can manually create protection domains before they are required by OneFS to perform certain actions. However, manually creating protection domains can limit your ability to interact with the data marked by the domain.

- Copying a large number of files into a protection domain might take a very long time because each file must be marked individually as belonging to the protection domain.
- You cannot move directories in or out of protection domains. However, you can move a directory contained in a protection domain to another location within the same protection domain.
- Creating a protection domain for a directory that contains a large number of files will take more time than creating a protection domain for a directory with fewer files. Because of this, it is recommended that you create protection domains for directories while the directories are empty, and then add files to the directory.
- If a domain is currently preventing the modification or deletion of a file, you cannot create a protection domain for a directory that contains that file. For example, if `/ifs/data/smartlock/file.txt` is set to a WORM state by a SmartLock domain, you cannot create a SnapRevert domain for `/ifs/data/`.

Create a protection domain

You can create replication or snapshot revert domains to facilitate snapshot revert and failover operations. You cannot create a SmartLock domain. OneFS automatically creates a SmartLock domain when you create a SmartLock directory.

Procedure

1. Run the `isi job jobs start` command.

The following command creates a SyncIQ domain for `/ifs/data/source`:

```
isi job jobs start domainmark --root /ifs/data/media \  
--dm-type SyncIQ
```

Delete a protection domain

You can delete a replication or snapshot revert domain if you want to move directories out of the domain. You cannot delete a SmartLock domain. OneFS automatically deletes a SmartLock domain when you delete a SmartLock directory.

Procedure

1. Run the `isi job jobs start` command.

The following command deletes a SyncIQ domain for `/ifs/data/source`:

```
isi job jobs start domainmark --root /ifs/data/media \  
--dm-type SyncIQ --delete
```


CHAPTER 18

Data-at-rest-encryption

This section contains the following topics:

- [Data-at-rest encryption overview](#) 684
- [Self-encrypting drives](#) 684
- [Data security on self-encrypted drives](#) 684
- [Data migration to a self-encrypted-drives cluster](#) 685
- [Chassis and drive states](#) 685
- [Smartfailed drive REPLACE state](#) 688
- [Smartfailed drive ERASE state](#) 689

Data-at-rest encryption overview

You can enhance data security with a EMC Isilon cluster that contains only self-encrypting-drive nodes, providing data-at-rest protection.

The OneFS system is available as a cluster that is composed of Isilon OneFS nodes that contain only self-encrypting drives (SEDs). The system requirements and management of data at rest on self-encrypting nodes are identical to that of nodes that do not contain self-encrypting drives. Clusters of mixed node types are not supported.

Self-encrypting drives

Self-encrypting drives store data on a EMC Isilon cluster that is specially designed for data-at-rest encryption.

Data-at-rest- encryption on self-encrypted drives occurs when data that is stored on a device is encrypted to prevent unauthorized data access. All data written to the storage device is encrypted when it is stored, and all data read from the storage device is decrypted when it is read. The stored data is encrypted with a 256-bit data AES encryption key and decrypted in the same manner. OneFS controls data access by combining the drive authentication key with on-disk data-encryption keys.

Note

All nodes in a cluster must be of the self-encrypting drive type. Mixed nodes are not supported.

Data security on self-encrypted drives

Smartfailing self-encrypted drives guarantees data security after removal.

Data on self-encrypted drives is protected from unauthorized access by authenticating encryption keys. Encryption keys never leave the drive. When a drive is locked, successful authentication unlocks the drive for data access.

The data on self-encrypted drives is rendered inaccessible in the following conditions:

- When a self-encrypting drive is smartfailed, drive authentication keys are deleted from the node. The data on the drive cannot be decrypted and is therefore unreadable, which secures the drive.
- When a drive is smartfailed and removed from a node, the encryption key on the drive is removed. Because the encryption key for reading data from the drive must be the same key that was used when the data was written, it is impossible to decrypt data that was previously written to the drive. When you smartfail and then remove a drive, it is cryptographically erased.

Note

Smartfailing a drive is the preferred method for removing a self-encrypted drive. Removing a node that has been smartfailed guarantees that data is inaccessible.

- When a self-encrypting drive loses power, the drive locks to prevent unauthorized access. When power is restored, data is again accessible when the appropriate drive authentication key is provided.

Data migration to a self-encrypted-drives cluster

You can migrate data from your existing cluster to a cluster of self-encrypted-drive nodes.

The Isilon cluster does not support the coexistence of regular and self-encrypted nodes. However, if you have data on an existing Isilon cluster that you want to migrate to a cluster of self-encrypted nodes, you can add self-encrypted nodes to your existing cluster one time only to migrate your data.

Note

Before you begin the data-migration process, both clusters must be upgraded to the same OneFS version.

During data migration, an error is generated that indicates you are running in mixed mode, which is not supported and is not secure. The data migrated to the self-encrypted drives is not secure until the smartfail process is completed for the non-encrypted drives.

⚠ CAUTION

Data migration to a cluster of self-encrypted-drive nodes must be performed by Isilon Professional Services. For more information, contact your EMC Isilon representative.

Chassis and drive states

You can view chassis and drive state details.

In a cluster, the combination of nodes in different degraded states determines whether read requests, write requests, or both work. A cluster can lose write quorum but keep read quorum. OneFS provides details about the status of chassis and drives in your cluster. The following table describes all the possible states that you may encounter in your cluster.

State	Description	Interface	Error state
HEALTHY	All drives in the node are functioning correctly.	Command-line interface, web administration interface	
SMARTFAIL or Smartfail or restripe in progress	The drive is in the process of being removed safely from the file system, either because of an I/O error or by user request. Nodes or drives in a smartfail or read-only state affect only write quorum.	Command-line interface, web administration interface	
NOT AVAILABLE	A drive is unavailable for a variety of reasons. You can click the bay to view detailed information about this condition.	Command-line interface, web administration interface	X

State	Description	Interface	Error state
	<hr/> <p>Note</p> <p>In the web administration interface, this state includes the ERASE and SED_ERROR command-line interface states.</p> <hr/>		
SUSPENDED	This state indicates that drive activity is temporarily suspended and the drive is not in use. The state is manually initiated and does not occur during normal cluster activity.	Command-line interface, web administration interface	
NOT IN USE	A node in an offline state affects both read and write quorum.	Command-line interface, web administration interface	
REPLACE	The drive was smartfailed successfully and is ready to be replaced.	Command-line interface only	
STALLED	The drive is stalled and undergoing stall evaluation. Stall evaluation is the process of checking drives that are slow or having other issues. Depending on the outcome of the evaluation, the drive may return to service or be smartfailed. This is a transient state.	Command-line interface only	
NEW	The drive is new and blank. This is the state that a drive is in when you run the <code>isi dev</code> command with the <code>-a add</code> option.	Command-line interface only	
USED	The drive was added and contained an Isilon GUID but the drive is not from this node. This drive likely will be formatted into the cluster.	Command-line interface only	
PREPARING	The drive is undergoing a format operation. The drive state changes to HEALTHY when the format is successful.	Command-line interface only	
EMPTY	No drive is in this bay.	Command-line interface only	
WRONG_TYPE	The drive type is wrong for this node. For example, a non-SED drive in a SED node, SAS instead of the expected SATA drive type.	Command-line interface only	
BOOT_DRIVE	Unique to the A100 drive, which has boot drives in its bays.	Command-line interface only	

State	Description	Interface	Error state
SED_ERROR	<p>The drive cannot be acknowledged by the OneFS system.</p> <hr/> <p>Note</p> <p>In the web administration interface, this state is included in Not available.</p> <hr/>	Command-line interface, web administration interface	X
ERASE	<p>The drive is ready for removal but needs your attention because the data has not been erased. You can erase the drive manually to guarantee that data is removed.</p> <hr/> <p>Note</p> <p>In the web administration interface, this state is included in Not available.</p> <hr/>	Command-line interface only	
INSECURE	<p>Data on the self-encrypted drive is accessible by unauthorized personnel. Self-encrypting drives should never be used for non-encrypted data purposes.</p> <hr/> <p>Note</p> <p>In the web administration interface, this state is labeled Unencrypted SED.</p> <hr/>	Command-line interface only	X
UNENCRYPTED SED	<p>Data on the self-encrypted drive is accessible by unauthorized personnel. Self-encrypting drives should never be used for non-encrypted data purposes.</p> <hr/> <p>Note</p> <p>In the command-line interface, this state is labeled INSECURE.</p> <hr/>	Web administration interface only	X

Smartfailed drive REPLACE state

You can see different drive states during the smartfail process.

If you run the `isi dev` command while the drive in bay 1 is being smartfailed, the system displays output similar to the following example:

```
Node 1, [ATTN]
  Bay 1          Lnum 11      [SMARTFAIL]   SN:Z296M8HK
000093172YE04 /dev/da1
  Bay 2          Lnum 10      [HEALTHY]     SN:Z296M8N5
00009330EYE03 /dev/da2
  Bay 3          Lnum 9       [HEALTHY]     SN:Z296LBP4
00009330EYE03 /dev/da3
  Bay 4          Lnum 8       [HEALTHY]     SN:Z296LCJW
00009327BYE03 /dev/da4
  Bay 5          Lnum 7       [HEALTHY]     SN:Z296M8XB
00009330KYE03 /dev/da5
  Bay 6          Lnum 6       [HEALTHY]     SN:Z295LXT7
000093172YE03 /dev/da6
  Bay 7          Lnum 5       [HEALTHY]     SN:Z296M8ZF
00009330KYE03 /dev/da7
  Bay 8          Lnum 4       [HEALTHY]     SN:Z296M8SD
00009330EYE03 /dev/da8
  Bay 9          Lnum 3       [HEALTHY]     SN:Z296M8QA
00009330EYE03 /dev/da9
  Bay 10         Lnum 2       [HEALTHY]     SN:Z296M8Q7
00009330EYE03 /dev/da10
  Bay 11         Lnum 1       [HEALTHY]     SN:Z296M8SP
00009330EYE04 /dev/da11
  Bay 12         Lnum 0       [HEALTHY]     SN:Z296M8QZ
00009330JYE03 /dev/da12
```

If you run the `isi dev` command after the smartfail completes successfully, the system displays output similar to the following example, showing the drive state as REPLACE:

```
Node 1, [ATTN]
  Bay 1          Lnum 11      [REPLACE]     SN:Z296M8HK
000093172YE04 /dev/da1
  Bay 2          Lnum 10      [HEALTHY]     SN:Z296M8N5
00009330EYE03 /dev/da2
  Bay 3          Lnum 9       [HEALTHY]     SN:Z296LBP4
00009330EYE03 /dev/da3
  Bay 4          Lnum 8       [HEALTHY]     SN:Z296LCJW
00009327BYE03 /dev/da4
  Bay 5          Lnum 7       [HEALTHY]     SN:Z296M8XB
00009330KYE03 /dev/da5
  Bay 6          Lnum 6       [HEALTHY]     SN:Z295LXT7
000093172YE03 /dev/da6
  Bay 7          Lnum 5       [HEALTHY]     SN:Z296M8ZF
00009330KYE03 /dev/da7
  Bay 8          Lnum 4       [HEALTHY]     SN:Z296M8SD
00009330EYE03 /dev/da8
  Bay 9          Lnum 3       [HEALTHY]     SN:Z296M8QA
00009330EYE03 /dev/da9
  Bay 10         Lnum 2       [HEALTHY]     SN:Z296M8Q7
00009330EYE03 /dev/da10
  Bay 11         Lnum 1       [HEALTHY]     SN:Z296M8SP
00009330EYE04 /dev/da11
  Bay 12         Lnum 0       [HEALTHY]     SN:Z296M8QZ
00009330JYE03 /dev/da12
```


If you run the `isi dev` command while the drive in bay 3 is being smartfailed, the system displays output similar to the following example:

```
Node 1, [ATTN]
  Bay 1      Lnum 11      [REPLACE]      SN:Z296M8HK
000093172YE04 /dev/da1
  Bay 2      Lnum 10      [HEALTHY]      SN:Z296M8N5
00009330EYE03 /dev/da2
  Bay 3      Lnum 9       [SMARTFAIL]    SN:Z296LBP4
00009330EYE03 N/A
  Bay 4      Lnum 8       [HEALTHY]      SN:Z296LCJW
00009327BYE03 /dev/da4
  Bay 5      Lnum 7       [HEALTHY]      SN:Z296M8XB
00009330KYE03 /dev/da5
  Bay 6      Lnum 6       [HEALTHY]      SN:Z295LXT7
000093172YE03 /dev/da6
  Bay 7      Lnum 5       [HEALTHY]      SN:Z296M8ZF
00009330KYE03 /dev/da7
  Bay 8      Lnum 4       [HEALTHY]      SN:Z296M8SD
00009330EYE03 /dev/da8
  Bay 9      Lnum 3       [HEALTHY]      SN:Z296M8QA
00009330EYE03 /dev/da9
  Bay 10     Lnum 2       [HEALTHY]      SN:Z296M8Q7
00009330EYE03 /dev/da10
  Bay 11     Lnum 1       [HEALTHY]      SN:Z296M8SP
00009330EYE04 /dev/da11
  Bay 12     Lnum 0       [HEALTHY]      SN:Z296M8QZ
00009330JYE03 /dev/da12
```

Smartfailed drive ERASE state

At the end of a smartfail process, OneFS attempts to delete the authentication key on a drive if it is unable to reset the key.

Note

- To securely delete the authentication key on a single drive, smartfail the individual drive.
- To securely delete the authentication key on a single node, smartfail the node.
- To securely delete the authentication keys on an entire cluster, smartfail each node and run the `isi_reformat_node` command on the last node.

Upon running the `isi dev` command, the system displays output similar to the following example, showing the drive state as ERASE:

```
Node 1, [ATTN]
  Bay 1      Lnum 11      [REPLACE]      SN:Z296M8HK
000093172YE04 /dev/da1
  Bay 2      Lnum 10      [HEALTHY]      SN:Z296M8N5
00009330EYE03 /dev/da2
  Bay 3      Lnum 9       [ERASE]        SN:Z296LBP4
00009330EYE03 /dev/da3
```

Drives showing the ERASE state can be safely retired, reused, or returned.

Any further access to a drive showing the ERASE state requires the authentication key of the drive to be set to its default manufactured security ID (MSID). This action erases the data encryption key (DEK) on the drive and renders any existing data on the drive permanently unreadable.

CHAPTER 19

SmartQuotas

This section contains the following topics:

- [SmartQuotas overview](#)..... 692
- [Quota types](#)..... 692
- [Default quota type](#).....693
- [Usage accounting and limits](#)..... 695
- [Disk-usage calculations](#)..... 696
- [Quota notifications](#)..... 697
- [Quota notification rules](#).....697
- [Quota reports](#).....698
- [Creating quotas](#).....698
- [Managing quotas](#)..... 700
- [Quota commands](#).....712

SmartQuotas overview

The SmartQuotas module is an optional quota-management tool that monitors and enforces administrator-defined storage limits. Using accounting and enforcement quota limits, reporting capabilities, and automated notifications, SmartQuotas manages storage use, monitors disk storage, and issues alerts when disk-storage limits are exceeded.

Quotas help you manage storage usage according to criteria that you define. Quotas are used as a method of tracking—and sometimes limiting—the amount of storage that a user, group, or project consumes. Quotas are a useful way of ensuring that a user or department does not infringe on the storage that is allocated to other users or departments. In some quota implementations, writes beyond the defined space are denied, and in other cases, a simple notification is sent.

The SmartQuotas module requires a separate license. For additional information about the SmartQuotas module or to activate the module, contact your EMC Isilon sales representative.

Quota types

OneFS uses the concept of quota types as the fundamental organizational unit of storage quotas. Storage quotas comprise a set of resources and an accounting of each resource type for that set. Storage quotas are also called storage domains.

Storage quotas creation requires three identifiers:

- The directory to monitor.
- Whether snapshots are to be tracked against the quota limit.
- The quota type (directory, user, or group).

You can choose a quota type from the following entities:

Directory

A specific directory and its subdirectories.

User

Either a specific user or default user (every user). Specific-user quotas that you configure take precedence over a default user quota.

Group

All members of a specific group or all members of a default group (every group). Any specific-group quotas that you configure take precedence over a default group quota. Associating a group quota with a default group quota creates a linked quota.

You can create multiple quota types on the same directory, but they must be of a different type or have a different snapshot option. You can specify quota types for any directory in OneFS and nest them within each other to create a hierarchy of complex storage-use policies.

Nested storage quotas can overlap. For example, the following quota settings ensure that the finance directory never exceeds 5 TB, while limiting the users in the finance department to 1 TB each:

- Set a 5 TB hard quota on `/ifs/data/finance`.
- Set 1 TB soft quotas on each user in the finance department.

Note

You should not create quotas of any type on the OneFS root (`/ifs`). A root-level quota may significantly degrade performance.

Default quota type

Default quotas automatically create other quotas for users or groups in a specified directory.

A default quota specifies a policy for new entities that match a trigger. The `default-user@/ifs/cs` becomes `specific-user@/ifs/cs` for each `specific-user` that is not otherwise defined.

For example, you can create a `default-user` quota on the `/ifs/dir-1` directory, where that directory is owned by the root user. The `default-user` type automatically creates a new domain on that directory for root and adds the usage there:

```
my-OneFS-1# mkdir /ifs/dir-1
my-OneFS-1# isi quota quotas create /ifs/dir-1 default-user
my-OneFS-1# isi quota quotas ls --path=/ifs/dir-1
```

Type	AppliesTo	Path	Snap	Hard	Soft	Adv	Used
default-user	DEFAULT	/ifs/dir-1	No	-	-	-	0b
user	root	/ifs/dir-1	No	-	-	-	0b

Now add a file that is owned by a different user (admin).

```
my-OneFS-1# touch /ifs/dir-1/somefile
my-OneFS-1# chown admin /ifs/dir-1/somefile
my-OneFS-1# isi quota quotas ls --path=/ifs/dir-1
```

Type	AppliesTo	Path	Snap	Hard	Soft	Adv	Used
default-user	DEFAULT	/ifs/dir-1	No	-	-	-	0b
user	root	/ifs/dir-1	No	-	-	-	26b
user	admin	/ifs/dir-1	No	-	-	-	0b

Total: 3

In this example, the `default-user` type created a new `specific-user` type automatically (`user:admin`) and added the new usage to it. `Default-user` does not have any usage because it is used only to generate new quotas automatically. `Default-user` enforcement is copied to a `specific-user` (`user:admin`), and the inherited quota is called a linked quota. In this way, each user account gets its own usage accounting.

Defaults can overlap. For example, `default-user@/ifs/dir-1` and `default-user@/ifs/cs` both may be defined. If the default enforcement changes, OneFS storage quotas propagate the changes to the linked quotas asynchronously. Because the update is asynchronous, there is some delay before updates are in effect. If a default type, such as every user or every group, is deleted, OneFS deletes all children that are marked as inherited. As an option, you can delete the default without deleting the children, but it is important to note that this action breaks inheritance on all inherited children.

Continuing with the example, add another file that is owned by the root user. Because the root type exists, the new usage is added to it.

```
my-OneFS-1# touch /ifs/dir-1/anotherfile
my-OneFS-1# isi quota ls -v --path=/ifs/dir-1 --format=list
Type: default-user
```

```

AppliesTo: DEFAULT
Path: /ifs/dir-1
Snap: No
Thresholds
    Hard : -
    Soft : -
    Adv : -
    Grace : -
Usage
    Files : 0
With Overhead : 0.00b
W/O Overhead : 0.00b
Over: -
Enforced: No
Container: No
Linked: -
-----
Type: user
AppliesTo: root
Path: /ifs/dir-1
Snap: No
Thresholds
    Hard : -
    Soft : -
    Adv : -
    Grace : -
Usage
    Files : 2
With Overhead : 3.50K
W/O Overhead : 55.00b
Over: -
Enforced: No
Container: No
Linked: Yes
-----
-
Type: user
AppliesTo: admin
Path: /ifs/dir-1
Snap: No
Thresholds
    Hard : -
    Soft : -
    Adv : -
    Grace : -
Usage
    Files : 1
With Overhead : 1.50K
W/O Overhead : 0.00b
Over: -
Enforced: No
Container: No
Linked: Yes

```

The enforcement on default-user is copied to the specific-user when the specific-user allocates within the type, and the new inherited quota type is also a linked quota.

Note

Configuration changes for linked quotas must be made on the parent quota that the linked quota is inheriting from. Changes to the parent quota are propagated to all children. To override configuration from the parent quota, you must unlink the quota first.

Usage accounting and limits

Storage quotas support two usage types that you can create to manage storage space. The usage types are accounting and enforcement limits.

You can configure OneFS quotas by usage type to track or limit storage use. The accounting option, which monitors disk-storage use, is useful for auditing, planning, and billing. Enforcement limits set storage limits for users, groups, or directories.

Accounting

The accounting option tracks but does not limit disk-storage use. Using the accounting option for a quota, you can monitor inode count and physical and logical space resources. Physical space refers to all of the space used to store files and directories, including data and metadata in the domain. Logical space refers to the sum of all files sizes, excluding file metadata and sparse regions. User data storage is tracked using logical-space calculations, which do not include protection overhead. As an example, by using the accounting option, you can do the following:

- Track the amount of disk space used by various users or groups to bill each user, group, or directory for only the disk space used.
- Review and analyze reports that help you identify storage usage patterns and define storage policies.
- Plan for capacity and other storage needs.

Enforcement limits

Enforcement limits include all of the functionality of the accounting option, plus the ability to limit disk storage and send notifications. Using enforcement limits, you can logically partition a cluster to control or restrict how much storage that a user, group, or directory can use. For example, you can set hard- or soft-capacity limits to ensure that adequate space is always available for key projects and critical applications and to ensure that users of the cluster do not exceed their allotted storage capacity. Optionally, you can deliver real-time email quota notifications to users, group managers, or administrators when they are approaching or have exceeded a quota limit.

Note

If a quota type uses the accounting-only option, enforcement limits cannot be used for that quota.

The actions of an administrator logged in as root may push a domain over a quota threshold. For example, changing the protection level or taking a snapshot has the potential to exceed quota parameters. System actions such as repairs also may push a quota domain over the limit.

The system provides three types of administrator-defined enforcement thresholds.

Threshold type	Description
Hard	Limits disk usage to a size that cannot be exceeded. If an operation, such as a file write, causes a quota target to exceed a hard quota, the following events occur: <ul style="list-style-type: none"> • the operation fails • an alert is logged to the cluster

Threshold type	Description
	<ul style="list-style-type: none"> a notification is issued to specified recipients. Writes resume when the usage falls below the threshold.
Soft	<p>Allows a limit with a grace period that can be exceeded until the grace period expires. When a soft quota is exceeded, an alert is logged to the cluster and a notification is issued to specified recipients; however, data writes are permitted during the grace period.</p> <p>If the soft threshold is still exceeded when the grace period expires, data writes fail, and a hard-limit notification is issued to the recipients you have specified.</p> <p>Writes resume when the usage falls below the threshold.</p>
Advisory	An informational limit that can be exceeded. When an advisory quota threshold is exceeded, an alert is logged to the cluster and a notification is issued to specified recipients. Advisory thresholds do not prevent data writes.

Disk-usage calculations

For each quota that you configure, you can specify whether data-protection overhead is included in future disk-usage calculations.

Most quota configurations do not need to include overhead calculations. If you do not include data-protection overhead in usage calculations for a quota, future disk-usage calculations for the quota include only the space that is required to store files and directories. Space that is required for the data-protection setting of the cluster is not included.

Consider the same example user, who is now restricted by a 40 GB quota that does not include data-protection overhead in its disk-usage calculations. If your cluster is configured with a 2x data-protection level and the user writes a 10 GB file to the cluster, that file consumes 20 GB of space but the 10GB for the data-protection overhead is not counted in the quota calculation. In this example, the user has reached 25 percent of the 40 GB quota by writing a 10 GB file to the cluster. This method of disk-usage calculation is recommended for most quota configurations.

If you include data-protection overhead in usage calculations for a quota, future disk-usage calculations for the quota include the total amount of space that is required to store files and directories, in addition to any space that is required to accommodate your data-protection settings, such as parity or mirroring. For example, consider a user who is restricted by a 40 GB quota that includes data-protection overhead in its disk-usage calculations. If your cluster is configured with a 2x data-protection level (mirrored) and the user writes a 10 GB file to the cluster, that file actually consumes 20 GB of space: 10 GB for the file and 10 GB for the data-protection overhead. In this example, the user has reached 50 percent of the 40 GB quota by writing a 10 GB file to the cluster.

Note

Cloned and deduplicated files are treated as ordinary files by quotas. If the quota includes data protection overhead, the data protection overhead for shared data is not included in the usage calculation.

You can configure quotas to include the space that is consumed by snapshots. A single path can have two quotas applied to it: one without snapshot usage, which is the default,

and one with snapshot usage. If you include snapshots in the quota, more files are included in the calculation than are in the current directory. The actual disk usage is the sum of the current directory and any snapshots of that directory. You can see which snapshots are included in the calculation by examining the `.snapshot` directory for the quota path.

Note

Only snapshots created after the QuotaScan job finishes are included in the calculation.

Quota notifications

Quota notifications are generated for enforcement quotas, providing users with information when a quota violation occurs. Reminders are sent periodically while the condition persists.

Each notification rule defines the condition that is to be enforced and the action that is to be executed when the condition is true. An enforcement quota can define multiple notification rules. When thresholds are exceeded, automatic email notifications can be sent to specified users, or you can monitor notifications as system alerts or receive emails for these events.

Notifications can be configured globally, to apply to all quota domains, or be configured for specific quota domains.

Enforcement quotas support the following notification settings. A given quota can use only one of these settings.

Limit notification settings	Description
Turn Off Notifications for this Quota	Disables all notifications for the quota.
Use Default Notification Rules	Uses the global default notification for the specified type of quota.
Use Custom Notification Rules	Enables the creation of advanced, custom notifications that apply to the specific quota. Custom notifications can be configured for any or all of the threshold types (hard, soft, or advisory) for the specified quota.

Quota notification rules

You can write quota notification rules to generate alerts that are triggered by event thresholds.

When an event occurs, a notification is triggered according to your notification rule. For example, you can create a notification rule that sends an email when a disk-space allocation threshold is exceeded by a group.

You can configure notification rules to trigger an action according to event thresholds (a notification condition). A rule can specify a schedule, such as "every day at 1:00 AM," for executing an action or immediate notification of certain state transitions. When an event occurs, a notification trigger may execute one or more actions, such as sending an email or sending a cluster alert to the interface. The following examples demonstrate the types of criteria that you can use to configure notification rules.

- Notify when a threshold is exceeded; at most, once every 5 minutes
- Notify when allocation is denied; at most, once an hour
- Notify while over threshold, daily at 2 AM
- Notify while grace period expired weekly, on Sundays at 2 AM

Notifications are triggered for events grouped by the following categories:

Instant notifications

Includes the write-denied notification, triggered when a hard threshold denies a write, and the threshold-exceeded notification, triggered at the moment a hard, soft, or advisory threshold is exceeded. These are one-time notifications because they represent a discrete event in time.

Ongoing notifications

Generated on a scheduled basis to indicate a persisting condition, such as a hard, soft, or advisory threshold being over a limit or a soft threshold's grace period being expired for a prolonged period.

Quota reports

The OneFS SmartQuotas module provides reporting options that enable administrators to manage cluster resources and analyze usage statistics.

Storage quota reports provide a summarized view of the past or present state of the quota domains. After raw reporting data is collected by OneFS, you can produce data summaries by using a set of filtering parameters and sort types. Storage-quota reports include information about violators, grouped by threshold types. You can generate reports from a historical data sample or from current data. In either case, the reports are views of usage data at a given time. OneFS does not provide reports on data aggregated over time, such as trending reports, but you can use raw data to analyze trends. There is no configuration limit on the number of reports other than the space needed to store them.

OneFS provides the following data-collection and reporting methods:

- Scheduled reports are generated and saved on a regular interval.
- Ad hoc reports are generated and saved at the request of the user.
- Live reports are generated for immediate and temporary viewing.

Scheduled reports are placed by default in the `/ifs/.isilon/smartquotas/reports` directory, but the location is configurable to any directory under `/ifs`. Each generated report includes quota domain definition, state, usage, and global configuration settings. By default, ten reports are kept at a time, and older reports are purged. You can create ad hoc reports at any time to view the current state of the storage quotas system. These live reports can be saved manually. Ad hoc reports are saved to a location that is separate from scheduled reports to avoid skewing the timed-report sets.

Creating quotas

You can create two types of storage quotas to monitor data: accounting quotas and enforcement quotas. Storage quota limits and restrictions can apply to specific users, groups, or directories.

The type of quota that you create depends on your goal.

- Enforcement quotas monitor and limit disk usage. You can create enforcement quotas that use any combination of hard limits, soft limits, and advisory limits.

Note

Enforcement quotas are not recommended for snapshot-tracking quota domains.

- Accounting quotas monitor, but do not limit, disk usage.
-

Note

After you create a new quota, it begins to report data almost immediately, but the data is not valid until the QuotaScan job completes. Before using quota data for analysis or other purposes, verify that the QuotaScan job has finished.

Create an accounting quota

You can create an accounting quota to monitor but not limit disk usage.

Optionally, you can include snapshot data, data-protection overhead, or both in the accounting quota.

For information about the parameters and options that you can use for this procedure, run the `isi quota quotas create --help` command.

Procedure

1. Create an accounting quota by running the `isi quota quotas create` command.

The following example creates a quota for the `/quota_test_1` directory. The quota sets an advisory threshold that is informative rather than enforced.

```
isi quota quotas create /ifs/data/quota_test_1 directory \
--advisory-threshold=10M --enforced=false
```

After you finish

After you create a new quota, it begins to report data almost immediately but the data is not valid until the QuotaScan job completes. Before using quota data for analysis or other purposes, verify that the QuotaScan job has finished by running the `isi job events list --job-type quotascan`.

Create an enforcement quota

You can create an enforcement quota to monitor and limit disk usage.

You can create enforcement quotas that set hard, soft, and advisory limits.

For information about the parameters and options that you can use for this procedure, run the `isi quota quotas create --help` command.

Procedure

1. Create an enforcement quota by running the `isi quota quotas create` command that sets the `--enforced` parameter to `true`.

The following command creates a quota for the `/quota_test_2` directory. The quota sets an advisory threshold that is enforced when the threshold specified is exceeded.

```
isi quota quotas create /ifs/data/quota_test_2 directory \
--advisory-threshold=100M --enforced=true
```

After you finish

After you create a new quota, it begins to report data almost immediately but the data is not valid until the QuotaScan job completes. Before using quota data for analysis or other

purposes, verify that the QuotaScan job has finished by running the `isi job events list --job-type quotascan` command.

Managing quotas

You can modify the configured values of a storage quota, and you can enable or disable a quota. You can also create quota limits and restrictions that apply to specific users, groups, or directories.

Quota management in OneFS is simplified by the quota search feature, which helps you to locate a quota or quotas by using filters. You can unlink quotas that are associated with a parent quota, and configure custom notifications for quotas. You can also disable a quota temporarily and then enable it when needed.

Note

Moving quota directories across quota domains is not supported.

Search for quotas

You can search for a quota using a variety of search parameters.

For information about the parameters and options that you can use for this procedure, run the `isi quota quotas list --help` command.

Procedure

1. Search for quotas by running the `isi quota quotas list` command.

The following command finds all quotas that monitor the `/ifs/data/quota_test_1` directory.

```
isi quota quotas list --path=/ifs/data/quota_test_1
```

Manage quotas

Quotas help you monitor and analyze the current or historic use of disk storage. You can search for quotas, and modify, delete, and unlink quotas.

An initial QuotaScan job must run for the default or scheduled quotas. Otherwise, the data displayed may be incomplete.

Before you modify a quota, consider how the changes will affect the file system and end users.

For information about the parameters and options that you can use for this procedure, run the `isi quota quotas list --help` command.

Note

- You can edit or delete a quota report only when the quota is not linked to a default quota.
 - You can unlink a quota only when the quota is linked to a default quota.
-

Procedure

1. Monitor and analyze current disk storage by running the following `isi quota quotas view` command.

The following example provides current usage information for the root user on the specified directory and includes snapshot data. For more information about the

parameters for this command, run the `isi quota quotas list --help` command.

```
isi quota quotas list -v --path=/ifs/data/quota_test_2 \
--include-snapshots="yes"
```

2. View all information in the quota report by running the `isi quota reports list` command:

To view specific information in a quota report, run the `isi quota quotas list --help` command to view the filter parameters. The following command lists all information in the quota report.

```
isi quota reports list -v
```

3. (Optional) To delete a quota, run the `isi quota quotas delete` command.

The following command deletes the specified directory-type quota. For information about parameters for this command, run the `isi quota quotas delete --help` command.

```
isi quota quotas delete /ifs/data/quota_test_2 directory
```

4. Unlink a quota by running the `isi quota quotas modify` command.

The following command example unlinks a user quota:

```
isi quota quotas modify /ifs/dir-1 user --linked=false --user=admin
```

Note

Configuration changes for linked quotas must be made on the parent (default) quota that the linked quota is inheriting from. Changes to the parent quota are propagated to all children. If you want to override configuration from the parent quota, you must first unlink the quota.

Export a quota configuration file

You can export quota settings as a configuration file, which can then be imported for reuse to another Isilon cluster. You can also store the exported quota configurations in a location outside of the cluster. This task may only be performed from the OneFS command line interface.

You can pipe the XML report to a file or directory. The file can then be imported to another cluster.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. At the command prompt, run the following command:

```
isi_classic quota list --export
```

The quota configuration file displays as raw XML.

Import a quota configuration file

You can import quota settings in the form of a configuration file that has been exported from another Isilon cluster. This task can only be performed from the OneFS command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Navigate to the location of the exported quota configuration file.
3. At the command prompt, run the following command, where *<filename>* is the name of an exported configuration file:

```
isi_classic quota import --from-file=<filename>
```

The system parses the file and imports the quota settings from the configuration file. Quota settings that you configured before importing the quota configuration file are retained, and the imported quota settings are effective immediately.

Managing quota notifications

Quota notifications can be enabled or disabled, modified, and deleted.

By default, a global quota notification is already configured and applied to all quotas. You can continue to use the global quota notification settings, modify the global notification settings, or disable or set a custom notification for a quota.

Enforcement quotas support four types of notifications and reminders:

- Threshold exceeded
- Over-quota reminder
- Grace period expired
- Write access denied

If a directory service is used to authenticate users, you can configure notification mappings that control how email addresses are resolved when the cluster sends a quota notification. If necessary, you can remap the domain that is used for quota email notifications and you can remap Active Directory domains, local UNIX domains, or both.

Configure default quota notification settings

You can configure default global quota notification settings that apply to all quotas of a specified threshold type.

The custom notification settings that you configure for a quota take precedence over the default global notification settings.

For information about the parameters and options that you can use for this procedure, run the `isi quota settings notifications modify --help` command.

Procedure

1. Run the `isi quota settings notifications modify` command to configure default quota notifications. You can run the `isi quota settings notifications modify --help` command to see the list of actions you can take when configuring default quota notification settings.

The following example configures the default quota notification settings to generate an alert when the advisory threshold is exceeded.

```
isi quota settings notifications modify advisory exceeded \
--action-alert=true
```

After you finish

After you create a new quota, it begins to report data almost immediately, but the data is not valid until the QuotaScan job completes. Before using quota data for analysis or other purposes, verify that the QuotaScan job has finished.

Configure custom quota notification rules

You can configure custom quota notification rules that apply only to a specified quota.

Before you begin

An enforcement quota must exist or be in the process of being created. To configure notifications for an existing enforcement quota, follow the procedure to modify a quota and then use these steps.

Quota-specific custom notification rules must be configured for that quota. If notification rules are not configured for a quota, the default event notification configuration is used.

For information about the parameters and options that you can use for this procedure, run the `isi quota quotas notifications create --help` command.

Procedure

1. Configure custom quota notification rules by running the `isi quota quotas notifications create` command.

The following command creates an advisory quota notification rule for the `/ifs/data/quota_test_2` directory that uses the `--holdoff` parameter to specify the length of time to wait before generating a notification:

```
isi quota quotas notifications create /ifs/data/quota_test_2 \
directory advisory exceeded --holdoff=10W
```

After you finish

After you create a quota it begins to report data almost immediately, but the data is not valid until the QuotaScan job completes. Before using quota data for analysis or other purposes, verify that the QuotaScan job has finished.

Map an email notification rule for a quota

Email notification mapping rules control how email addresses are resolved when the cluster sends a quota notification.

If necessary, you can remap the domain used for SmartQuotas email notifications. You can remap Active Directory Windows domains, local UNIX domains, or NIS domains.

Note

You must be logged in to the web administration interface to perform this task.

Procedure

1. Click **File System Management** > **SmartQuotas** > **Settings**.
2. (Optional) In the **Email Mapping** area, click **Create an email mapping rule**.

3. From the **Provider Type** list, select the authentication provider type for this notification rule. The default is `Local`. To determine which authentication providers are available on your cluster, navigate to **Access > Authentication Providers**.
4. From the **Current Domain** list, select the domain that you want to use for the mapping rule. If the list is blank, navigate to **Cluster Management > Network Configuration**, and then click **Edit** in the **DNS Settings** area to specify the domains that you want to use for mapping.
5. In the **Map-to-Domain** field, type the name of the domain that you want to map email notifications to. This can be the same domain name you selected from the **Current Domain** list. To specify multiple domains, separate the domain names with commas.
6. Click **Save Rule**.

Email quota notification messages

If email notifications for exceeded quotas are enabled, you can customize Isilon-provided templates for email notifications or create your own.

There are three email notification templates provided with OneFS. The templates are located in `/etc/ifs` and are described in the following table:

Template	Description
<code>quota_email_template.txt</code>	A notification that disk quota has been exceeded.
<code>quota_email_grace_template.txt</code>	A notification that disk quota has been exceeded (also includes a parameter to define a grace period in number of days).
<code>quota_email_test_template.txt</code>	A notification test message you can use to verify that a user is receiving email notifications.

If the default email notification templates do not meet your needs, you can configure your own custom email notification templates using a combination of text and SmartQuotas variables. Whether you choose to create your own templates or modify the existing ones, make sure that the first line of the template file is a `Subject:` line. For example:

```
Subject: Disk quota exceeded
```

If you want to include information about the message sender, include a `From:` line immediately under the subject line. If you use an email address, include the full domain name for the address. For example:

```
From: administrator@abcd.com
```

In this example of the `quota_email_template.txt` file, a `From:` line is included. Additionally, the default text "Contact your system administrator for details" at the end of the template is changed to name the administrator:

```
Subject: Disk quota exceeded
From: administrator@abcd.com
```

```
The <ISI_QUOTA_TYPE> disk quota on directory <ISI_QUOTA_PATH>
owned by <ISI_QUOTA_OWNER> on <ISI_QUOTA_NODE> was exceeded.
```

```
The quota limit is <ISI_QUOTA_THRESHOLD>, and <ISI_QUOTA_USAGE>
is currently in use. You may be able to free some disk space by
deleting unnecessary files. If your quota includes snapshot usage,
your administrator may be able to free some disk space by deleting
```


one or more snapshots. Contact Jane Anderson (janderson@abcd.com) for details.

This is an example of a what a user will see as an emailed notification (note the SmartQuotas variables are resolved):

```
Subject: Disk quota exceeded
From: administrator@abcd.com

The advisory disk quota on directory /ifs/data/sales_tools/collateral
owned by jsmith on production-Boris was exceeded.

The quota limit is 10 GB, and 11 GB is in use. You may be able
to free some disk space by deleting unnecessary files. If your
quota includes snapshot usage, your administrator may be able
to free some disk space by deleting one or more snapshots.
Contact Jane Anderson (janderson@abcd.com) for details.
```

Custom email notification template variable descriptions

An email template contains text and, optionally, variables that represent values. You can use any of the SmartQuotas variables in your templates.

Variable	Description	Example
ISI_QUOTA_PATH	Path of quota domain	/ifs/data
ISI_QUOTA_THRESHOLD	Threshold value	20 GB
ISI_QUOTA_USAGE	Disk space in use	10.5 GB
ISI_QUOTA_OWNER	Name of quota domain owner	jsmith
ISI_QUOTA_TYPE	Threshold type	Advisory
ISI_QUOTA_GRACE	Grace period, in days	5 days
ISI_QUOTA_EXPIRATION	Expiration date of grace period	Fri May 22 14:23:19 PST 2015
ISI_QUOTA_NODE	Hostname of the node on which the quota event occurred	someHost-prod-wf-1

Customize email quota notification templates

You can customize Isilon-provided templates for email notifications. This task may only be performed from the OneFS command line interface.

This procedure assumes you are using the provided templates, which are located in the `/etc/ifs` directory.

Note

We recommend that you do not edit the templates directly. Instead, copy them to another directory to edit and deploy them.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.

- Copy one of the default templates to a directory in which you will edit the file and later access it through the OneFS web administration interface. For example:

```
cp /etc/ifs/quota_email_template.txt /ifs/data/quotanotifiers/
quota_email_template_copy.txt
```

- Open the desired template file in a text editor. For example:

```
edit /ifs/data/quotanotifiers/quota_email_template_copy.txt
```

The template displays in the editor.

- Edit the template as desired. If you are using or creating a customized template, make sure the template has a `Subject:` line.
 - Save your changes. Template files must be saved as `.txt` files.
 - In the web administration interface, navigate to **File System > SmartQuotas > Quotas & Usage**.
 - Select the desired quota for which you wish to set a notification rule.
 - Click the **Settings** tab.
 - Select the notification rule you wish to use with the template you created (for example, **Advisory Limit Notification Rules**). For expanded information about setting notification rules, refer to the instructions for configuring default quota notification settings and configuring custom quota notification rules in this chapter.
 - Select the desired event for the template (for example, **Event: Advisory Limit Value Exceeded**).
 - In the **Send Email** area, select one of the owner notification type check boxes.
 - In the **Message Template** field, enter or browse to find the template you copied or customized.
 - (Optional) In the **Event** area, select **Create Cluster Event** to generate an event notification in addition to the email notification.
 - (Optional) In the **Delay** area, select the desired amount of time to wait before generating a notification. The default is zero minutes.
- Repeat steps 9 through 14 to specify an email notification template for each notification rule you wish to create for the quota.
- Click **Save**.

Managing quota reports

You can configure and schedule reports to help you monitor, track, and analyze storage use on an Isilon cluster.

You can view and schedule reports and customize report settings to track, monitor, and analyze disk storage use. Quota reports are managed by configuring settings that give you control over when reports are scheduled, how they are generated, where and how many are stored, and how they are viewed. The maximum number of scheduled reports that are available for viewing in the web-administration interface can be configured for each report type. When the maximum number of reports are stored, the system deletes the oldest reports to make space for new reports as they are generated.

Create a quota report schedule

You can configure quota report settings to generate the quota report on a specified schedule.

Quota report settings determine whether and when scheduled reports are generated, and where and how the reports are stored. If you disable a scheduled report, you can still run unscheduled reports at any time.

For information about the parameters and options that you can use for this procedure, run the `isi quota reports list --help` command.

Procedure

1. Configure a quota report schedule by running the `isi quota settings reports modify` command.

The following command creates a quota report schedule that runs every 2 days. For more information about date pattern or other schedule parameters, see `man isi-schedule`.

```
isi quota settings reports modify --schedule="Every 2 day"
```

Results

Reports are generated according to your criteria and can be viewed by running the `isi quota reports list` command.

Generate a quota report

In addition to scheduled quota reports, you can generate a report to capture usage statistics at any time.

Before you begin

Quotas must exist and the initial QuotaScan job must run and complete before you can generate a quota report.

For information about the parameters and options that you can use for this procedure, run the `isi quota reports create --help` command.

Procedure

1. Generate a quota report by running the `isi quota reports create` command.

The following command creates an ad hoc quota report.

```
isi quota reports create -v
```

Results

You can view the quota report by running the `isi quota reports list -v` command.

Locate a quota report

You can locate quota reports, which are stored as XML files, and use your own tools and transforms to view them. This task can only be performed from the OneFS command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Navigate to the directory where quota reports are stored. The following path is the default quota report location:

```
/ifs/.isilon/smartquotas/reports
```

Note

If quota reports are not in the default directory, you can run the `isi quota settings` command to find the directory where they are stored.

3. At the command prompt, run the `ls` command.

- To view a list of all quota reports in the directory, run the following command:

```
ls -a *.xml
```

- To view a specific quota report in the directory, run the following command:

```
ls <filename>.xml
```

Basic quota settings

When you create a storage quota, the following attributes must be defined, at a minimum. When you specify usage limits, additional options are available for defining your quota.

Option	Description
Directory Path	The directory that the quota is on.
User Quota	Select to automatically create a quota for every current or future user that stores data in the specified directory.
Group Quota	Select to automatically create a quota for every current or future group that stores data in the specified directory.
Include Snapshot Data	Select to count all snapshot data in usage limits; cannot be changed after the quota is created.
Include Data-Protection Overhead	Select to count protection overhead in usage limits.
No Usage Limit	Select to account for usage only.
Specify Usage Limits	Select to enforce advisory, soft, or absolute limits.

Advisory limit quota notification rules settings

You can configure custom quota notification rules for advisory limits for a quota. These settings are available when you select the option to use custom notification rules.

Option	Description	Exceeded	Remains exceeded
Send email	Specify the type of email to use.	Yes	Yes

Option	Description	Exceeded	Remains exceeded
Notify owner	Select to send an email notification to the owner of the entity.	Yes	Yes
Notify another	Select to send an email notification to another recipient and type the recipient's email address.	Yes	Yes
Message template	Select from the following template types for use in formatting email notifications: <ul style="list-style-type: none"> • Default (leave Message Template field blank to use default) • Custom 	Yes	Yes
Create cluster event	Select to generate an event notification for the quota when exceeded.	Yes	Yes
Delay	Specify the length of time (hours, days, weeks) to delay before generating a notification.	Yes	No
Frequency	Specify the notification and alert frequency: daily, weekly, monthly, yearly; depending on selection, specify intervals, day to send, time of day, multiple emails per rule.	No	Yes

Soft limit quota notification rules settings

You can configure custom soft limit notification rules for a quota. These settings are available when you select the option to use custom notification rules.

Option	Description	Exceeded	Remains exceeded	Grace period expired	Write access denied
Send email	Specify the recipient of the email notification.	Yes	Yes	Yes	Yes
Notify owner	Select to send an email notification to the owner of the entity.	Yes	Yes	Yes	Yes
Notify another	Select to send an email notification to another recipient and type the recipient's email address.	Yes	Yes	Yes	Yes
Message template	Select from the following template types for use in formatting email notifications: <ul style="list-style-type: none"> • Default (leave Message Template field blank to use default) • Custom 	Yes	Yes	Yes	Yes
Create cluster event	Select to generate an event notification for the quota.	Yes	Yes	Yes	Yes
Delay	Specify the length of time (hours, days, weeks) to delay before generating a notification.	Yes	No	No	Yes
Frequency	Specify the notification and alert frequency: daily, weekly, monthly, yearly; depending on selection,	No	Yes	Yes	No

Option	Description	Exceeded	Remains exceeded	Grace period expired	Write access denied
	specify intervals, day to send, time of day, multiple emails per rule.				

Hard limit quota notification rules settings

You can configure custom quota notification rules for hard limits for a quota. These settings are available when you select the option to use custom notification rules.

Option	Description	Write access denied	Exceeded
Send email	Specify the recipient of the email notification.	Yes	Yes
Notify owner	Select to send an email notification to the owner of the entity.	Yes	Yes
Notify another	Select to send an email notification to another recipient and type the recipient's email address.	Yes	Yes
Message template	Select from the following template types for use in formatting email notifications: <ul style="list-style-type: none"> Default (leave Message Template field blank to use default) Custom 	Yes	Yes
Create cluster event	Select to generate an event notification for the quota when exceeded.	Yes	Yes
Delay	Specify the length of time (hours, days, weeks) to delay before generating a notification.	Yes	No
Frequency	Specify the notification and alert frequency: daily, weekly, monthly, yearly; depending on selection, specify intervals, day to send, time of day, multiple emails per rule.	No	Yes

Limit notification settings

You have three notification options when you create an enforcement quota: use default notification rules, turn off notifications, or use custom notification rules. Enforcement quotas support the following notification settings for each threshold type. A quota can use only one of these settings.

Notification setting	Description
Use Default Notification Rules	Uses the default notification rules that you configured for the specified threshold type.
Turn Off Notifications for this Quota	Disables all notifications for the quota.
Use Custom Notification Rules	Provides settings to create basic custom notifications that apply to only this quota.

Quota report settings

You can configure quota report settings that track disk usage. These settings determine whether and when scheduled reports are generated, and where and how reports are stored. When the maximum number of reports are stored, the system deletes the oldest reports to make space for new reports as they are generated.

Setting	Description
Scheduled reporting	<p>Enables or disables the scheduled reporting feature.</p> <ul style="list-style-type: none"> • Off. Manually generated on-demand reports can be run at any time. • On. Reports run automatically according to the schedule that you specify.
Report frequency	<p>Specifies the interval for this report to run: daily, weekly, monthly, or yearly. You can use the following options to further refine the report schedule.</p> <p>Generate report every. Specify the numeric value for the selected report frequency; for example, every 2 months.</p> <p>Generate reports on. Select the day or multiple days to generate reports.</p> <p>Select report day by. Specify date or day of the week to generate the report.</p> <p>Generate one report per specified by. Set the time of day to generate this report.</p> <p>Generate multiple reports per specified day. Set the intervals and times of day to generate the report for that day.</p>
Scheduled report archiving	<p>Determines the maximum number of scheduled reports that are available for viewing on the SmartQuotas Reports page.</p> <p>Limit archive size for scheduled reports to a specified number of reports. Type the integer to specify the maximum number of reports to keep.</p> <p>Archive Directory. Browse to the directory where you want to store quota reports for archiving.</p>
Manual report archiving	<p>Determines the maximum number of manually generated (on-demand) reports that are available for viewing on the SmartQuotas Reports page.</p> <p>Limit archive size for live reports to a specified number of reports. Type the integer to specify the maximum number of reports to keep.</p> <p>Archive Directory. Browse to the directory where you want to store quota reports for archiving.</p>

Quota commands

You can configure quotas to track, limit, and manage disk usage by directory, user, or group. Quota commands that create and modify quotas are available only if you activate a SmartQuotas license on the cluster.

isi quota quotas create

Creates new file system quotas.

Syntax

```
isi quota quotas create
  --path <path>
  --type {directory | user | group | default-user | default-group}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--hard-threshold <size>]
  [--advisory-threshold <size>]
  [--soft-threshold <size>] [--soft-grace <duration>]
  [--container {yes | no}]
  [--include-snapshots {yes | no}]
  [--thresholds-include-overhead {yes | no}]
  [--enforced {yes | no}] [--zone <zone>]
  [--verbose]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.



You should not create quotas of any type on the `/ifs` directory. A root-level quota may result in significant performance degradation.

`--type`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--user <name>`
 Specifies a user name.

`--group <name>`
 Specifies a group name.

`--gid <id>`
 Specifies the numeric group identifier (GID).

`--uid <id>`
 Specifies a numeric user identifier (UID).

`--sid <sid>`
 Sets a security identifier (SID). For example, S-1-5-21-13.

`--wellknown <name>`
 Specifies a well-known user, group, machine, or account name.

`--hard-threshold <size>`
 Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit. Size is a capacity value formatted as `<integer>[b | K | M | G | T | P]`.

`--advisory-threshold <size>`
 Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests. Size is a capacity value formatted as `<integer>[b | K | M | G | T | P]`.

`--soft-threshold <size>`
 Specifies the soft threshold, which allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter. Size is a capacity value formatted as `<integer>[b | K | M | G | T | P]`.

`--soft-grace <duration>`
 Specifies the soft threshold grace period, which is the amount of time to wait before disk write requests are denied.

Specify `<duration>` in the following format:

```
<integer><units>
```

The following `<units>` are valid:

Y
 Specifies years

M
 Specifies months

W
 Specifies weeks

D
 Specifies days

H
 Specifies hours

```
--container {yes | no}
    Specifies that threshold be shown as the available space on the SMB share, instead
    of the whole cluster. The setting applies only to hard thresholds. When setting this
    value, you must specify --enforced.
```

```
--include-snapshots {yes | no}
    Includes snapshots in the quota size.
```

```
--thresholds-include-overhead {yes | no}
    Includes OneFS storage overhead in the quota threshold when set to yes.
```

```
--enforced {yes | no}
    Enforces this quota when set to yes. Specifying any threshold automatically sets this
    value to yes on create.
```

```
--zone <zone>
    Specifies an access zone.
```

```
{--verbose | -v}
    Displays more detailed information.
```

isi quota quotas delete

Deletes a file system quota or multiple quotas.

Syntax

```
isi quota quotas delete
  --path <path>
  --type {directory | user | group | default-user | default-group
        | --all}
  [--user <name>] | [--group <name>] | [--gid <id>] | [--uid <id>]
  | [--sid <sid>]
  [--recurse-path-parents]
  [--recurse-path-children]
  [--include-snapshots {yes | no}]
  [--zone <zone>]
  [--verbose]
```

Options

--path <path>

Specifies an absolute path within the /ifs file system.

--type

Deletes quotas of the specified type. Argument must be specified with the *--path* option. The following values are valid:

directory

Specifies a quota for all data in the directory, regardless of owner.

user

Specifies a quota for one specific user. Requires specification of *--user*, *--uid*, or *--sid*.

group

Specifies a quota for one specific group. Requires specification of the *--group*, *--gid*, or *--sid* option.

default-user

Specifies a master quota that creates a linked quota for every user who has data in the directory.

default-group

Specifies a master quota that creates a linked quota for every group that owns data in the directory.

--all

Deletes all quotas. Flag may not be specified with `--type` or `--path`.

`--user <name>`

Deletes a quota associated with the user identified by name.

`--gid <id>`

Deletes a quota by the specified numeric group identifier (GID).

`--uid <id>`

Deletes a quota by the specified numeric user identifier (UID).

`--sid <sid>`

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

`--recurse-path-parents`

Searches parent paths for quotas.

`--recurse-path-children`

Searches child paths for quotas.

`--include-snapshots {yes | no}`

Deletes quotas that include snapshot data usage.

`--zone <zone>`

Specifies an access zone.

`{--verbose | -v}`

Displays more detailed information.

isi quota quotas modify

Modifies a file system quota.

Syntax

```
isi quota quotas modify
--path <path>
--type {directory | user | group | default-user | default-group}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--hard-threshold <size>]
[--advisory-threshold <size>]
[--soft-threshold <size>]
[--soft-grace <duration>]
[--container {yes | no}]
[--include-snapshots {yes | no}]
[--thresholds-include-overhead {yes | no}]
[--enforced {yes | no}]
[--linked {yes | no}]
[--clear-hard-threshold]
[--clear-advisory-threshold]
```

```
[--clear-soft-threshold]
[--zone <string>]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.

`--type`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, or `--sid` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, or `--sid` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--user <name>`

Specifies a user name.

`--group <name>`

Specifies a group name.

`--gid <id>`

Specifies the numeric group identifier (GID).

`--uid <id>`

Specifies a numeric user identifier (UID).

`--sid <sid>`

Specifies a security identifier (SID) for selecting the quota that you want to modify. For example, `S-1-5-21-13`.

`--wellknown <name>`

Specifies a well-known user, group, machine, or account name.

`--hard-threshold <size>`

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit. Size is a capacity value formatted as `<integer>{[b | K | M | G | T | P]}`.

`--advisory-threshold <size>`

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests. Size is a capacity value formatted as `<integer>{b | K | M | G | T | P}`.

`--soft-threshold <size>`

Specifies the soft threshold, which allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter. Size is a capacity value formatted as `<integer>{b | K | M | G | T | P}`.

`--soft-grace <duration>`

Specifies the soft threshold grace period, which is the amount of time to wait before disk write requests are denied.

Specify `<duration>` in the following format:

```
<integer><units>
```

The following `<units>` are valid:

Y
Specifies years

M
Specifies months

W
Specifies weeks

D
Specifies days

H
Specifies hours

`--container {yes | no}`

Specifies that threshold be shown as the available space on the SMB share, instead of the whole cluster. The setting applies only to hard thresholds. When setting this value, you must specify `--enforced`.

`--include-snapshots {yes | no}`

Includes snapshots in the quota size.

`--thresholds-include-overhead {yes | no}`

Includes OneFS storage overhead in the quota threshold when set to `yes`.

`--enforced {yes | no}`

Enforces this quota when set to `yes`. Specifying any threshold automatically sets this value to `yes` on create.

`--linked {yes | no}`

Unlinks a linked quota created automatically by a default-user or default-group quota. Unlinking allows the quota to be modified separately. To modify a linked quota, you must modify the original default-user or default-group quota it originated from, instead of the linked quota itself.

`--clear-hard-threshold`

Clears an absolute limit for disk usage.

`--clear-advisory-threshold`

Clears the advisory threshold.

`--clear-soft-threshold`

Clears the soft threshold.

`--zone <string>`

The zone used by the quota. Use this parameter only to resolve personas used by the quota.

isi quota quotas list

Displays a list of quotas.

Syntax

```
isi quota quotas list
  [--user <name> | --group <name> | --gid <id> | --uid <id>
  | --sid <sid> | --wellknown <name>]
  [--type {directory | user | group | default-user
  | default-group}]
  [--path]
  [--recurse-path-parents]
  [--recurse-path-children]
  [--include-snapshots {yes | no}]
  [--exceeded]
  [--enforced {yes | no}]
  [--zone <zone>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--user <name>`

Specifies a user name.

`--group <name>`

Specifies a group name.

`--gid <id>`

Specifies the numeric group identifier (GID).

`--uid <id>`

Specifies a numeric user identifier (UID).

`--sid <sid>`

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

`--wellknown <name>`

Specifies a well-known user, group, machine, or account name.

`<type>`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--path`

Specifies quotas on the specified path.

`--recurse-path-parents`

Specifies parent paths for quotas.

`--recurse-path-children`

Specifies child paths for quotas.

`--include-snapshots {yes | no}`

Specifies quotas that include snapshot data usage.

`--exceeded`

Specifies only quotas that have an exceeded threshold.

`--enforced {yes | no}`

Specifies quotas that have an enforced threshold.

`--zone <zone>`

Specifies quotas in the specified zone.

`--limit <integer>`

Specifies the number of quotas to display.

`--format`

Displays quotas in the specified format. The following values are valid:

`table`

`json`

`csv`

`list`

`{--no-header | -a}`

Suppresses headers in CSV or table formats.

`{--no-footer | -z}`

Suppresses table summary footer information.

`{--verbose | -v}`

Displays more detailed information.

isi quota quotas view

Displays detailed properties of a single file system quota.

Syntax

```
isi quota quotas view
  --path <path>
  --type {directory | user | group | default-user | default-group}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
  [--zone <string>]
```

Options

`--path <path>`

Specifies an absolute path within the /ifs file system.

`--type`

Specifies quotas of the specified type. Argument must be specified with the `--path` option. The following values are valid:

directory

Specifies a quota for all data in the directory, regardless of owner.

user

Specifies a quota for one specific user. Requires specification of `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Specifies a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Specifies a master quota that creates a linked quota for every user who has data in the directory.

default-group

Specifies a master quota that creates a linked quota for every group that owns data in the directory.

`--user <name>`

Specifies a quota associated with the user identified by name.

`--group <name>`

Specifies a quota associated with the group identified by name.

`--gid <id>`

Specifies a quota by the numeric group identifier (GID).

`--uid <id>`

Specifies a quota by the specified numeric user identifier (UID).

`--sid <sid>`

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

`--wellknown <name>`

Specifies a well-known user, group, machine, or account name.

`--include-snapshots {yes | no}`

Specifies quotas that include snapshot data usage.

`--zone <zone>`

Specifies an access zone.

isi quota quotas notifications clear

Clears rules for a quota and uses system notification settings.

Note

Use the `isi quota quotas notifications disable` command to disable all notifications for a quota.

Syntax

```
isi quota quotas notifications clear
  --path <path>
  --type {directory | user | group | default-user | default-group}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.

`<type> --type`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--user <name>`

Specifies a user name.

`--group <name>`

Specifies a group name.

`--gid <id>`

Specifies the numeric group identifier (GID).

`--uid <id>`

Specifies a numeric user identifier (UID).

`--sid <sid>`

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

`--wellknown <name>`

Specifies a well-known user, group, machine, or account name.

`--include-snapshots {yes | no}`

Includes snapshots in the quota size.

isi quota quotas notifications create

Creates a notification rule for a quota.

Syntax

```
isi quota quotas notifications create
--path <path>
--type {directory | user | group | default-user | default-group}
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
[--schedule <name>]
[--holdoff <duration>]
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address <address>]
[--verbose]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.

`--type`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--threshold`

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Sets a security identifier (SID). For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Specifies quotas that include snapshot data usage.

--schedule <name>

Specifies the date pattern at which recurring notifications are made. Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--holdoff *<duration>*

Specifies the length of time to wait before generating a notification. Specify *<duration>* in the following format:

<integer><units>

The following *<units>* are valid:

- Y**
Specifies years
- M**
Specifies months
- W**
Specifies weeks
- D**
Specifies days
- H**
Specifies hours
- S**
Specifies seconds

--action-alert {yes | no}

Generates an alert when the notification condition is met.

--action-email-owner {yes | no}

Specifies that an email be sent to a user when the threshold is crossed. Requires --action-email-address.

--action-email-address *<address>*

Specifies the email address of user to be notified.

{--verbose | -v}

Displays more detailed information.

isi quota quotas notifications delete

Deletes a quota notification rule.

Syntax

```
isi quota quotas notifications delete
--path <path>
--type {directory | user | group | default-user | default-group}
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
[--verbose]
```

Options

--path *<path>*

Deletes quota notifications set on an absolute path within the `/ifs` file system.

--type

Deletes a quota notification by specified type. The following values are valid:

directory

Specifies a quota for all data in the directory, regardless of owner.

user

Specifies a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Specifies a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Specifies a master quota that creates a linked quota for every user who has data in the directory.

default-group

Specifies a master quota that creates a linked quota for every group that owns data in the directory.

--threshold

Deletes a quota notification by specified threshold. The following values are valid:

hard

Specifies an absolute limit for disk usage.

soft

Specifies the soft threshold.

advisory

Specifies the advisory threshold..

--condition

Deletes a quota notification by the specified condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

`--user <name>`

Deletes a quota notification by the specified user name.

`--group <name>`

Deletes a quota notification by the specified group name.

`--gid <id>`

Deletes a quota notification by the specified numeric group identifier (GID).

`--uid <id>`

Deletes a quota notification by the specified numeric user identifier (UID).

`--sid <sid>`

Deletes a quota notification by the specified security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

`--wellknown <name>`

Deletes a quota notification by the specified well-known user, group, machine, or account name.

`--include-snapshots {yes | no}`

Deletes a quota notification by the specified settings for Included snapshots in the quota size.

`{--verbose | -v}`

Displays more detailed information.

isi quota quotas notifications disable

Disables all quota notifications.



When you disable all quota notifications, system notification behavior is disabled also. Use the `--clear` options to remove specific quota notification rules and fall back to the system default.

Syntax

```
isi quota quotas notifications disable
  --path <path>
  --type {directory | user | group | default-user | default-group}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.

`--type`

Disables quotas of the specified type. Argument must be specified with the `--path` option. The following values are valid:

directory

Specifies a quota for all data in the directory, regardless of owner.

user

Specifies a quota for one specific user. Requires specification of `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Specifies a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Specifies a master quota that creates a linked quota for every user who has data in the directory.

default-group

Specifies a master quota that creates a linked quota for every group that owns data in the directory.

`--user <name>`

Disables a quota associated with the user identified by name.

`--gid <id>`

Disables a quota by the specified numeric group identifier (GID).

`--uid <id>`

Disables a quota by the specified numeric user identifier (UID).

`--sid <sid>`

Specifies a security identifier (SID) for selecting a quota. For example, S-1-5-21-13.

`--wellknown <name>`

Specifies a well-known user, group, machine, or account name.

`--include-snapshots {yes | no}`

Disables quotas that include snapshot data usage.

isi quota quotas notifications list

Displays a list of quota notification rules.

Syntax

```
isi quota quotas notifications list
  --path <path>
  --type {directory | user | group | default-user | default-group}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.

`--type`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--threshold`

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

{--limit |-1} <integer>

Specifies the number of quota notification rules to display.

--format

Displays quota notification rules in the specified format. The following values are valid:

table

json

csv

list

- `{--no-header | -a}`
Suppresses headers in CSV or table formats.
- `{--no-footer | -z}`
Suppresses table summary footer information.
- `{--verbose | -v}`
Displays more detailed information.

isi quota quotas notifications modify

Modifies a notification rule for a quota.

Syntax

```
isi quota quotas notifications modify
--path <path>
--type {directory | user | group | default-user | default-group}
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
[--schedule <string>]
[--holdoff <duration>]
[--clear-holdoff]
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address <address>]
[--email-template <path>]
[--clear-email-template]
[--verbose]
```

Options

- `--path <path>`
Specifies an absolute path within the `/ifs` file system.
- `--type`
Specifies a quota type. The following values are valid:
 - directory**
Creates a quota for all data in the directory, regardless of owner.
 - user**
Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.
 - group**
Creates a quota for one specific group. Requires specification of `--group`, `--gid`, `--sid`, or `--wellknown` option.
 - default-user**
Creates a master quota that creates a linked quota for every user who has data in the directory.
 - default-group**
Creates a master quota that creates a linked quota for every group that owns data in the directory.
- `--threshold`
Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Sets a security identifier (SID). For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

--schedule <name>

Specifies the date pattern at which recurring notifications are made. Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--holdoff *<duration>*

Specifies the length of time to wait before generating a notification. Specify *<duration>* in the following format:

<integer><units>

The following *<units>* are valid:

- Y**
Specifies years
- M**
Specifies months
- W**
Specifies weeks
- D**
Specifies days
- H**
Specifies hours
- S**
Specifies seconds

--clear-holdoff

Clears the value for the `--holdoff` duration.

--action-alert {yes | no}

Generates an alert when the notification condition is met.

--action-email-owner {yes | no}

Specifies that an email be sent to a user when the threshold is crossed. Requires `--action-email-address`.

`--action-email-address <address>`

Specifies the email address of user to be notified.

`{--verbose | -v}`

Displays more detailed information.

isi quota quotas notifications view

Displays the properties of a quota notification rule.

Syntax

```
isi quota quotas notifications view
  --path <path>
  --type {directory | user | group | default-user | default-group}
  --threshold {hard | soft | advisory}
  --condition {exceeded | denied | violated | expired}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
```

Options

`--path <path>`

Specifies an absolute path within the `/ifs` file system.

`--type`

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user

Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.

group

Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.

default-user

Creates a master quota that creates a linked quota for every user who has data in the directory.

default-group

Creates a master quota that creates a linked quota for every group that owns data in the directory.

`--threshold`

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

isi quota reports create

Generates a quota report.

Syntax

```
isi quota reports create
  [--verbose]
```

Options

```
{--verbose | -v}
```

Displays more detailed information.

isi quota reports delete

Deletes a specified report.

Syntax

```
isi quota reports delete
--time <string>
--generated {live | scheduled | manual}
--type {summary | detail}
[--verbose]
```

Options

`--time <string>`

Specifies the timestamp of the report.

Specify *<time-and-date>* in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

Specify *<time>* as one of the following values.

Y

Specifies years

M

Specifies months

W

Specifies weeks

D

Specifies days

h

Specifies hours

s

Specifies seconds

`--generated`

Specifies the method used to generate the report. The following values are valid:

live

scheduled

manual

`--type`

Specifies a report type. The following values are valid:

summary

detail

`{--verbose | -v}`

Displays more detailed information.

isi quota reports list

Displays a list of quota reports.

Syntax

```
isi quota reports list
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

`--limit <integer>`

Specifies the number of quotas to display.

`--format`

Displays quotas in the specified format. The following values are valid:

table

json

csv

list

`{--no-header | -a}`

Suppresses headers in CSV or table formats.

`{--no-footer | -z}`

Suppresses table summary footer information.

`{--verbose | -v}`

Displays more detailed information.

isi quota settings notifications clear

Clears all default quota notification rules.

When you clear all default notification rules, the system reverts to system notification behavior. Use the `--disable` option to disable notification settings for a specific quota notification rule.

Syntax

```
isi quota settings notifications clear
```

isi quota settings notifications create

Creates a default notification rule.

Syntax

```
isi quota settings notifications create
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
--schedule <string>
--holdoff <duration>
```



```
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address {yes | no}]
[--email-template <path>]
[--verbose]
```

Options

--threshold

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--schedule <string>

Specifies the date pattern at which recurring notifications are made. Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month

- Yearly on `<month> <integer>`
- Yearly on the {last | `<integer>`} [weekday | `<day>`] of `<month>`

Specify `<frequency>` in one of the following formats:

- at `<hh>[:<mm>] [{AM | PM}]`
- every [`<integer>`] {hours | minutes} [between `<hh>[:<mm>] [{AM | PM}]` and `<hh>[:<mm>] [{AM | PM}]`]
- every [`<integer>`] {hours | minutes} [from `<hh>[:<mm>] [{AM | PM}]` to `<hh>[:<mm>] [{AM | PM}]`]

You can optionally append "st", "th", or "rd" to `<integer>`. For example, you can specify "Every 1st month"

Specify `<day>` as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`--holdoff <duration>`

Specifies the length of time to wait before generating a notification. Specify `<duration>` in the following format:

`<integer> <units>`

The following `<units>` are valid:

- Y** Specifies years
- M** Specifies months
- W** Specifies weeks
- D** Specifies days
- H** Specifies hours
- S** Specifies seconds

`--action-alert {yes | no}`

Generates an alert when the notification condition is met.

`--action-email-owner {yes | no}`

Specifies that an email be sent to a user when the threshold is crossed. Requires `--action-email-address`.

`--action-email-address <address>`

Specifies the email address of user to be notified.

`--email-template <path>`

Specifies the path in `/ifs` to the email template.

`{--verbose | -v}`

Displays more detailed information.

isi quota settings notifications delete

Delete a default quota notification rule.

Syntax

```
isi quota settings notifications delete
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--verbose]
```

Options

`--threshold`

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

`--condition`

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

`{--verbose | -v}`

Displays more detailed information.

isi quota settings notifications list

Displays a list of global quota notification rules.

Syntax

```
isi quota settings notifications list
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

`{--limit |-l} <integer>`

Specifies the number of quota notification rules to display.

`--format`

Displays quotas in the specified format. The following values are valid:

table

json

csv

list

`{--no-header |-a}`

Suppresses headers in CSV or table formats.

`{--no-footer |-z}`

Suppresses table summary footer information.

`{--verbose |-v}`

Displays more detailed information.

isi quota settings notifications modify

Modifies a quota notification rule.

Syntax

```
isi quota settings notifications modify
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--schedule <string>]
[--holdoff <duration>]
[--clear-holdoff]
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address <address>]
[--email-template <path>]
[--clear-email-template]
[--verbose]
```

Options

`--threshold`

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--schedule <string>

Specifies the date pattern at which recurring notifications are made.

--holdoff <duration>

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]

- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

Specifies the length of time to wait before generating a notification. Specify *<duration>* in the following format:

```
<integer><units>
```

The following *<units>* are valid:

- Y**
Specifies years
- M**
Specifies months
- W**
Specifies weeks
- D**
Specifies days
- H**
Specifies hours
- S**
Specifies seconds

```
--clear-holdoff
```

Clears the value for the `--holdoff` duration.

```
--action-alert {yes | no}
```

Generates an alert when the notification condition is met.

```
--action-email-owner {yes | no}
```

Specifies that an email be sent to a user when the threshold is crossed. Requires `--action-email-address`.

```
--action-email-address <address>
```

Specifies the email address of user to be notified.

```
{--verbose | -v}
```

Displays more detailed information.

```
--clear-email-template
```

Clears the setting for the path to the email template.

isi quota settings notifications view

Displays properties of a system default notification rule.

Syntax

```
isi quota settings notifications view
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
```

Options

`--threshold`

Specifies the threshold type. The following values are valid:

hard

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.

soft

Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.

advisory

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

`--condition`

Specifies the quota condition on which to send a notification. The following values are valid:

denied

Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.

exceeded

Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.

violated

Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired

Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

isi quota settings reports modify

Modifies cluster-wide quota report settings.

Syntax

```
isi quota settings reports modify
  [--schedule <schedule>]
  [--revert-schedule]
  [--scheduled-dir <dir>]
  [--revert-scheduled-dir]
  [--scheduled-retain <integer>]
  [--revert-scheduled-retain]
  [--live-dir <dir> | --revert-live-dir]
  [--live-retain <integer> | --revert-live-retain]
  [--verbose]
```

Options

`--schedule <schedule>`

Specifies the date pattern at which recurring notifications are made. Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`--revert-schedule`

Sets the `--schedule` value to system default.

`--scheduled-dir <dir>`

Specifies the location where scheduled quota reports are stored.

`--revert-scheduled-dir`

Sets the `--scheduled-dir` value to system default.

`--scheduled-retain <integer>`

Specifies the maximum number of scheduled reports to keep.

`--revert-scheduled-retain`

Sets the `--scheduled-retain` value to system default.

`--live-dir <dir>`

Specifies the location where live quota reports are stored.

`--revert-live-dir`

Sets the `--live-dir` value to system default.

`--live-retain <integer>`

Specifies the maximum number of live quota reports to keep.

`--revert-live-retain`

Sets the `--live-retain` value to system default.

`{--verbose | -v}`

Displays more detailed information.

isi quota settings reports view

Displays cluster-wide quota report settings.

Syntax

```
isi quota settings reports view
```

Options

There are no options for this command.

CHAPTER 20

Storage Pools

This section contains the following topics:

- [Storage pools overview](#)..... 748
- [Storage pool functions](#)..... 748
- [Autoprovisioning](#)..... 750
- [Node pools](#)..... 750
- [Virtual hot spare](#)..... 751
- [Spillover](#)..... 752
- [Suggested protection](#)..... 752
- [Protection policies](#)..... 753
- [SSD strategies](#)..... 753
- [Global namespace acceleration](#)..... 754
- [L3 cache overview](#)..... 755
- [Tiers](#)..... 756
- [File pool policies](#)..... 756
- [Managing node pools through the command-line interface](#)..... 757
- [Managing L3 cache from the command-line interface](#)..... 761
- [Managing tiers](#)..... 762
- [Creating file pool policies](#)..... 763
- [Managing file pool policies](#)..... 771
- [Monitoring storage pools](#)..... 774
- [Storage pool commands](#)..... 775

Storage pools overview

OneFS organizes different node types into separate node pools, and you can further organize these node pools into logical tiers of storage. By activating a SmartPools license, you can create file pool policies that store files in these tiers automatically, based on file matching criteria that you specify.

Without an active SmartPools license, OneFS manages all node pools as a single pool of storage and stores file data and metadata across the entire cluster, ensuring that data is protected, secure, and readily accessible. All files belong to the default file pool and are governed by the default file pool policy. In this mode, OneFS provides a number of functions that you can deploy, such as autoprovisioning, compatibilities, virtual hot spare, SSD strategies, global namespace acceleration (GNA), L3 cache, and storage tiers.

When you activate a SmartPools license, you can take further control of your data set to optimize the performance of your cluster. Additional functions become available, including custom file pool policies that take precedence over the default policy, and spillover management.

The following table summarizes storage pool functions based on whether a SmartPools license is inactive or active.

Function	Without active SmartPools license	With active SmartPools license
Automatic storage pool provisioning	Yes	Yes
Compatibilities (node equivalence)	Yes	Yes
Virtual hot spare	Yes	Yes
SSD strategies	Yes	Yes
L3 cache	Yes	Yes
Tiers	Yes	Yes
GNA	Yes	Yes
File pool policies	No	Yes
Spillover management	No	Yes

Storage pool functions

OneFS automatically groups equivalence-class nodes into node pools when the system is installed and whenever nodes are added to the cluster. This autoprovisioning of nodes into node pools enables OneFS to optimize reliability and data protection on the cluster.

Without an active SmartPools license, OneFS uses a default file pool policy to organize all data into a single file pool. In this mode, OneFS distributes data across the entire cluster so that data is protected and readily accessible.

OneFS includes the following functions with or without an active SmartPools license:

Autoprovisioning of node pools

OneFS automatically groups equivalence-class nodes into node pools for optimal storage efficiency and protection.

Compatibilities (node equivalence)

Enable certain nodes that are not equivalence-class to join existing node pools. OneFS supports compatibilities between Isilon S200 and S210 nodes, and between Isilon X400 and X410 nodes.

Tiers

Group node pools into logical tiers of storage. It is recommended that you activate a SmartPools license for this feature. An active SmartPools license enables you to create custom file pool policies and direct different file pools to appropriate storage tiers.

Default file pool policy

Governs all file types and can store files anywhere on the cluster. Custom file pool policies, which require an SmartPools license, take precedence over the default file pool policy.

Requested protection

Specify a requested protection setting for the default file pool, per node pool, or even on individual files. You can leave the default setting in place, or choose the suggested protection calculated by OneFS to ensure optimal data protection.

Virtual hot spare

Reserve a percentage of available storage for data repair in the event of a disk failure.

SSD strategies

Define the type of data that is stored on SSDs in the cluster, for example, storing metadata for read/write acceleration.

L3 cache

Specify that SSDs in nodes exclusively cache data and metadata, thus enabling faster access to the most frequently requested files.

Global namespace acceleration

Activate global namespace acceleration (GNA), which allows data stored on node pools without SSDs to access SSDs elsewhere in the cluster to store extra metadata mirrors. Extra metadata mirrors accelerate metadata read operations.

When you activate a SmartPools license, OneFS provides access to the following additional functions:

Custom file pool policies

Create custom file pool policies to identify different classes of files, and store these file pools in logical storage tiers. For example, you could define a high-performance tier of Isilon S-series node pools, and an archival tier of high-capacity Isilon NL400 and HD400 node pools. Then, with custom file pool policies, you could identify file pools based on matching criteria, and define actions to perform on these pools. For example, one file pool policy could identify all JPEG files older than a year and store them in an archival tier. Another policy could move all files that have been created or modified within the last three months to a performance tier.

Storage pool spillover

Enable automated capacity overflow management for storage pools. Spillover defines how to handle write operations when a storage pool is not writable for some reason. If spillover is enabled, data is redirected to a specified storage pool. If spillover is disabled, new data writes fail and an error message is sent to the client attempting the write operation.

Autoprovisioning

When you add a node to your cluster, OneFS automatically assigns the node to a node pool. With node pools, OneFS can ensure optimal performance, load balancing, and reliability of the file system. Autoprovisioning reduces the time required for the manual management tasks associated with resource planning and configuration.

Nodes are not provisioned, meaning they are not associated with each other and are not writable, until at least three nodes of an equivalence class are added to the cluster. If you have added only two nodes of an equivalence class to your cluster, no data is stored on those nodes until you add a third node of the same equivalence class.

Similarly, if a node goes down or is removed from the cluster so that fewer than three equivalence-class nodes remain, the node pool becomes under-provisioned. However, the two remaining nodes are still writable. If only one node remains, that node is not writable, but remains readable.

OneFS offers a compatibility function, also referred to as node equivalency, that enables certain node types to be provisioned to existing node pools, even when there are fewer than three equivalence-class nodes. For example, an S210 node could be provisioned and added to a node pool of S200 nodes. Similarly, an X410 node could be added to a node pool of X400 nodes. The compatibility function ensures that you can add nodes one at a time to your cluster and still have them be fully functional peers within a node pool.

Node pools

A node pool is a collection of three or more nodes. As you add nodes to an Isilon cluster, OneFS automatically provisions them into node pools based on characteristics such as series, drive size, RAM, and SSD-per-node ratio. Nodes with identical characteristics are called equivalence-class nodes.

If you add fewer than three nodes of a node type, OneFS cannot autoprovision the nodes to your cluster. In these cases, you can create compatibilities. Compatibilities enable OneFS to provision nodes that are not equivalence-class to a compatible node pool.

After provisioning, each node in the OneFS cluster is a peer, and any node can handle a data request. Each provisioned node increases the aggregate disk, cache, CPU, and network capacity on the cluster.

You can move nodes from an automatically managed node pool into one that you define manually. This capability is available only through the OneFS command-line interface. If you attempt to remove nodes from a node pool such that the removal would leave fewer than three nodes in the pool, the removal fails. When you remove a node from a manually defined node pool, OneFS attempts to move the node into a node pool of the same equivalence class, or into a compatible node pool.

Node compatibilities

OneFS requires that a node pool contain at least three nodes so that the operating system can write data and perform the necessary load balancing and data protection operations. You can enable certain nodes to be provisioned to an existing node pool by defining a compatibility.

Note

The compatibility function is also referred to as node equivalency.

If you add fewer than three Isilon S210 or X410 nodes to your cluster, and you have existing S200 or X400 node pools, you can create compatibilities to provision the new nodes and make them functional within the cluster. Only S210 and X410 nodes are eligible for compatibility.

To be provisioned, the S210 or X410 nodes must have the same drive configurations as their S200 and X400 counterparts, and must have compatible RAM amounts, as shown in the following table:

S200/S210 Compatibility		X400/X410 Compatibility	
S200 RAM	S210 RAM	X400 RAM	X410 RAM
24 GB	32 GB	24 GB	32 GB
48 GB	64 GB	48 GB	64 GB
96 GB	128 GB	96 GB	128 GB
	256 GB	192 GB	256 GB

Note

After you have added three or more S210 or X410 nodes to your cluster, you should consider removing the compatibilities that you have created. This step enables OneFS to autoprovision new S210 or X410 node pools and take advantage of the performance specifications of the newer node types.

Manual node pools

If the node pools automatically provisioned by OneFS do not meet your needs, you can configure node pools manually. You do this by moving nodes from an existing node pool into the manual node pool.

This capability enables you to store data on specific nodes according to your purposes, and is available only through the OneFS command-line interface.

CAUTION

It is recommended that you enable OneFS to provision nodes automatically. Manually created node pools might not provide the same performance and efficiency as automatically managed node pools, particularly if your changes result in fewer than 20 nodes in the manual node pool.

Virtual hot spare

Virtual hot spare (VHS) settings enable you to reserve disk space to rebuild the data in the event that a drive fails.

You can specify both a number of virtual drives to reserve and a percentage of total storage space. For example, if you specify two virtual drives and 15 percent, each node pool reserves virtual drive space equivalent to two drives or 15 percent of their total capacity (whichever is larger).

You can reserve space in node pools across the cluster for this purpose by specifying the following options:

- At least 1–4 virtual drives.

- At least 0–20% of total storage.

OneFS calculates the larger number of the two factors to determine the space that is allocated. When configuring VHS settings, be sure to consider the following information:

- If you deselect the option to **Ignore reserved space when calculating available free space** (the default), free-space calculations include the space reserved for VHS.
- If you deselect the option to **Deny data writes to reserved disk space** (the default), OneFS can use VHS for normal data writes. We recommend that you leave this option selected, or data repair can be compromised.
- If **Ignore reserved space when calculating available free space** is enabled while **Deny data writes to reserved disk space** is disabled, it is possible for the file system to report utilization as more than 100 percent.

Note

VHS settings affect spillover. If the VHS option **Deny data writes to reserved disk space** is enabled while **Ignore reserved space when calculating available free space** is disabled, spillover occurs before the file system reports 100% utilization.

Spillover

When you activate a SmartPools license, you can designate a node pool or tier to receive spillover data when the hardware specified by a file pool policy is full or otherwise not writable.

If you do not want data to spill over to a different location because the specified node pool or tier is full or not writable, you can disable this feature.

Note

Virtual hot spare reservations affect spillover. If the setting **Deny data writes to reserved disk space** is enabled, while **Ignore reserved space when calculating available free space** is disabled, spillover occurs before the file system reports 100% utilization.

Suggested protection

Based on the configuration of your Isilon cluster, OneFS automatically calculates the amount of protection that is recommended to maintain EMC Isilon's stringent data protection requirements.

OneFS includes a function to calculate the suggested protection for data to maintain a theoretical mean-time to data loss (MTTDL) of 5000 years. Suggested protection provides the optimal balance between data protection and storage efficiency on your cluster.

By configuring file pool policies, you can specify one of multiple requested protection settings for a single file, for subsets of files called file pools, or for all files on the cluster.

It is recommended that you do not specify a setting below suggested protection. OneFS periodically checks the protection level on the cluster, and alerts you if data falls below the recommended protection.

Protection policies

OneFS provides a number of protection policies to choose from when protecting a file or specifying a file pool policy.

The more nodes you have in your cluster, up to 20 nodes, the more efficiently OneFS can store and protect data, and the higher levels of requested protection the operating system can achieve. Depending on the configuration of your cluster and how much data is stored, OneFS might not be able to achieve the level of protection that you request. For example, if you have a three-node cluster that is approaching capacity, and you request +2n protection, OneFS might not be able to deliver the requested protection.

The following table describes the available protection policies in OneFS.

Protection policy	Summary
+1n	Tolerate the failure of 1 drive or the failure of 1 node
+2d:1n	Tolerate the failure of 2 drives or the failure of 1 node
+2n	Tolerate the failure of 2 drives or the failure of 2 nodes
+3d:1n	Tolerate the failure of 3 drives or the failure of 1 node
+3d:1n1d	Tolerate the failure of 3 drives or the failure of 1 node and 1 drive
+3n	Tolerate the failure of 3 drives or the failure of 3 nodes
+4d:1n	Tolerate the failure of 4 drives or the failure of 1 node
+4d:2n	Tolerate the failure of 4 drives or the failure of 2 nodes
+4n	Tolerate the failure of 4 drives or the failure of 4 nodes
Mirrors: 2x 3x 4x 5x 6x 7x 8x	<p>Duplicates, or mirrors, data over the specified number of nodes. For example, 2x results in two copies of each data block.</p> <hr/> <p>Note</p> <p>Mirrors can use more data than the other protection policies, but might be an effective way to protect files that are written non-sequentially or to provide faster access to important files.</p> <hr/>

SSD strategies

OneFS clusters can contain nodes that include solid-state drives (SSD). OneFS autoprovisions equivalence-class nodes with SSDs into one or more node pools. The SSD strategy defined in the default file pool policy determines how SSDs are used within the cluster, and can be set to increase performance across a wide range of workflows.

You can configure file pool policies to apply specific SSD strategies as needed. When you select SSD options during the creation of a file pool policy, you can identify the files in the OneFS cluster that require faster or slower performance. When the SmartPools job runs, OneFS uses file pool policies to move this data to the appropriate storage pool and drive type.

The following SSD strategy options that you can set in a file pool policy are listed in order of slowest to fastest choices:

Avoid SSDs

Writes all associated file data and metadata to HDDs only.

CAUTION

Use this option to free SSD space only after consulting with Isilon Technical Support personnel. Using this strategy can negatively affect performance.

Metadata read acceleration

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if available. The SSD mirror is in addition to the number of mirrors, if any, required to satisfy the requested protection.

Metadata read/write acceleration

Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata read acceleration** setting. Enabling GNA does not affect read/write acceleration.

Data on SSDs

Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly increased storage requirements compared with the other SSD strategy options.

Global namespace acceleration

Global namespace acceleration (GNA) allows data stored on node pools without SSDs to use SSDs elsewhere in the cluster to store extra metadata mirrors. Extra metadata mirrors can improve file system performance by accelerating metadata read operations.

You can only enable GNA if 20% or more of the nodes in the cluster contain at least one SSD and 1.5% or more of the total cluster storage is SSD-based. For best results, ensure that at least 2.0% of the total cluster storage is SSD-based before enabling global namespace acceleration.

If the ratio of SSDs to non-SSDs on the cluster falls below the 1.5% threshold, GNA becomes inactive even if enabled. GNA is reactivated when the ratio is corrected. When GNA is inactive, existing SSD mirrors are readable but newly written metadata does not include the extra SSD mirror.

Note

If GNA is enabled for the cluster, file pool policies that direct data to node pools with L3 cache enabled should also set the SSD strategy to `Avoid SSDs`. Otherwise, additional SSD mirrors would be created for data that is already accelerated by L3 cache. This is an inefficient use of SSD storage space and is not recommended.

L3 cache overview

You can configure nodes with solid-state drives (SSDs) to increase cache memory and speed up file system performance across larger working file sets.

OneFS caches file data and metadata at multiple levels. The following table describes the types of file system cache available on an Isilon cluster.

Name	Type	Profile	Scope	Description
L1 cache	RAM	Volatile	Local node	Also known as front-end cache, holds copies of file system metadata and data requested by the front-end network through NFS, SMB, HTTP, and so on.
L2 cache	RAM	Volatile	Global	Also known as back-end cache, holds copies of file system metadata and data on the node that owns the data.
SmartCache	Variable	Non-volatile	Local node	Holds any pending changes to front-end files waiting to be written to storage. This type of cache protects write-back data through a combination of RAM and stable storage.
L3 cache	SSD	Non-volatile	Global	Holds file data and metadata released from L2 cache, effectively increasing L2 cache capacity.

OneFS caches frequently accessed file and metadata in available random access memory (RAM). Caching enables OneFS to optimize data protection and file system performance. When RAM cache reaches capacity, OneFS normally discards the oldest cached data and processes new data requests by accessing the storage drives. This cycle is repeated each time RAM cache fills up.

You can deploy SSDs as L3 cache to reduce the cache cycling issue and further improve file system performance. L3 cache adds significantly to the available cache memory and provides faster access to data than hard disk drives (HDD).

As L2 cache reaches capacity, OneFS evaluates data to be released and, depending on your workflow, moves the data to L3 cache. In this way, much more of the most frequently accessed data is held in cache, and overall file system performance is improved.

For example, consider a cluster with 128GB of RAM. Typically the amount of RAM available for cache fluctuates, depending on other active processes. If 50 percent of RAM is available for cache, the cache size would be approximately 64GB. If this same cluster had three nodes, each with two 200GB SSDs, the amount of L3 cache would be 1.2TB, approximately 18 times the amount of available L2 cache.

L3 cache is enabled by default for new node pools. A node pool is a collection of nodes that are all of the same equivalence class, or for which compatibilities have been created. L3 cache applies only to the nodes where the SSDs reside. For the HD400 node, which is primarily for archival purposes, L3 cache is on by default and cannot be turned off. On the HD400, L3 cache is used only for metadata.

If you enable L3 cache on a node pool, OneFS manages all cache levels to provide optimal data protection, availability, and performance. In addition, in case of a power failure, the data on L3 cache is retained and still available after power is restored.

Note

Although some benefit from L3 cache is found in workflows with streaming and concurrent file access, L3 cache provides the most benefit in workflows that involve random file access.

Migration to L3 cache

L3 cache is enabled by default on new nodes. If you are upgrading your cluster from an older release (pre-OneFS 7.1.1), you must enable L3 cache manually on node pools that have SSDs. When you enable L3 cache, OneFS activates a process that migrates SSDs from storage disks to cache. File data currently on SSDs is moved elsewhere in the cluster.

You can enable L3 cache as the default for all new node pools or manually for a specific node pool, either through the command line or from the web administration interface. You can enable L3 cache only on node pools whose nodes have SSDs.

Depending on the amount of data stored in your SSDs, the migration process can take some time. OneFS displays a message informing you that the migration is about to begin:

```
WARNING: Changes to L3 cache configuration can have a long completion time. If this is a concern, please contact EMC Isilon Support for more information.
```

You must confirm whether OneFS should proceed with the migration. After you do, OneFS handles the migration intelligently as a background process. You can continue to administer your cluster during the migration.

If you choose to disable L3 cache on a node pool, the migration process is very fast.

L3 cache on HD400 node pools

The HD400 is a high-capacity node designed primarily for archival workflows. L3 cache is turned on by default on HD400 node pools and cannot be turned off.

Archival workflows feature a higher percentage of data writes compared to data reads. Consequently, L3 cache on HD400 node pools holds only metadata, which improves the speed of file system traversal activities such as directory lookup. L3 cache on HD400 nodes does not contain file data.

Tiers

A tier is a user-defined collection of node pools that you can specify as a storage pool for files. A node pool can belong to only one tier.

You can create tiers to assign your data to any of the node pools in the tier. For example, you can assign a collection of node pools to a tier specifically created to store data that requires high availability and fast access. In a three-tier system, this classification may be Tier 1. You can classify data that is used less frequently or that is accessed by fewer users as Tier-2 data. Tier 3 usually comprises data that is seldom used and can be archived for historical or regulatory purposes.

File pool policies

File pool policies define sets of files—file pools—and where and how they are stored on your cluster. You can configure multiple file pool policies with filtering rules that identify

specific file pools and the requested protection and I/O optimization settings for these file pools. Creating custom file pool policies requires an active SmartPools license.

The initial installation of OneFS places all files into a single file pool, which is subject to the default file pool policy. Without an active SmartPools license, you can configure only the default file pool policy, which controls all files and stores them anywhere on the cluster.

With an active SmartPools license, OneFS augments basic storage functions by enabling you to create custom file pool policies that identify, protect, and control multiple file pools. With a custom file pool policy, for example, you can define and store a file pool on a specific node pool or tier for fast access or archival purposes.

When you create a file pool policy, flexible filtering criteria enable you to specify time-based attributes for the dates that files were last accessed, modified, or created. You can also define relative time attributes, such as 30 days before the current date. Other filtering criteria include file type, name, size, and custom attributes. The following examples demonstrate a few ways you can configure file pool policies:

- A file pool policy to set stronger protection on a specific set of important files.
- A file pool policy to store frequently accessed files in a node pool that provides the fastest reads or read/writes.
- A file pool policy to evaluate the last time files were accessed, so that older files are stored in a node pool best suited for regulatory archival purposes.

When the SmartPools job runs, typically once a day, it processes file pool policies in priority order. You can edit, reorder, or remove custom file pool policies at any time. The default file pool policy, however, is always last in priority order. Although you can edit the default file pool policy, you cannot reorder or remove it. When custom file pool policies are in place, the settings in the default file pool policy apply only to files that are not covered by another file pool policy.

When new files are created, OneFS temporarily chooses a storage pool based on file pool policies in place when the last SmartPools job ran. OneFS might apply new storage settings and move these new files when the next SmartPools job runs, based on a new matching file pool policy.

Managing node pools through the command-line interface

Node pools can be managed through the command-line interface. You can work with node pools that are automatically provisioned, create compatibilities (equivalencies) for new nodes, and create and manage node pools manually.

A node pool, whether automatically provisioned or manually created, must contain a minimum of three equivalent or compatible nodes. Nodes are provisioned when at least three equivalent or compatible nodes have been added to the cluster. If you have added only two equivalent or compatible nodes to a cluster, you cannot store data on the nodes until adding a third node.

OneFS allows you to create compatibilities (also known as node equivalencies) for S200 and S210 nodes, and for X400 and X410 nodes, so that compatible nodes can be members of the same node pool. After you create a compatibility, any time a new compatible node is added to the cluster, OneFS provisions the new node to the appropriate node pool.

You can create a node pool manually only by selecting a subset of equivalence-class or compatible nodes from an existing autoprovisioned node pool. You cannot create a manual node pool that takes some nodes from one node pool and some nodes from another.

Also, when you establish a manual node pool, the original node pool should have at least three remaining nodes, or zero nodes.

You must have the `ISI_PRIV_SMARTPOOLS` or a higher administrative privilege to manage node pools.

Add a compatible node to a node pool

OneFS automatically adds a new equivalence-class node to an existing node pool. For a new node that is not equivalence-class, you can create a compatibility to add the node to an existing node pool.

Compatibilities work only between S200 and S210 nodes, and X400 and X410 nodes. For example, if you have a node pool made up of three or more S200 nodes, you can create a compatibility so that new S210 nodes are automatically added to the S200 node pool. Similarly, if you have a node pool made up of three or more X400 nodes, you can create a compatibility so that new X410 nodes are automatically added to the X400 node pool.

Procedure

1. Run the `isi storagepool compatibilities active create` command.

The following command creates a compatibility between Isilon S200 and S210 nodes:

```
isi storagepool compatibilities active create S200 S210
```

OneFS provides a summary of the results of executing the command, and requires you to confirm the operation.

2. To proceed, type **yes**, and then press ENTER. To cancel, type **no**, and then press ENTER.

Results

If you proceeded to create the compatibility, OneFS adds any unprovisioned S210 nodes to the S200 node pool.

Merge compatible node pools

You can merge multiple compatible node pools to optimize storage efficiency on your cluster.

For example, if you had six S200 nodes in one node pool and three S210 nodes in a second node pool, you could create a compatibility to merge the two node pools into one pool of nine nodes. Larger node pools, up to approximately 20 nodes, enable OneFS to protect data more efficiently, thus preserving more storage space for new data.

Procedure

1. Run the `isi storagepool compatibilities active create` command.

The following command creates a compatibility between Isilon X400 and X410 node pools:

```
isi storagepool compatibilities active create X400 X410
```

OneFS provides a summary of the results of executing the command, including the node pools that will be merged, and requires you to confirm the operation.

2. To proceed, type **yes**, and then press ENTER. To cancel, type **no**, and then press ENTER.

Results

If you allowed the operation to proceed, the compatible node pools are merged into one node pool.

Delete a compatibility

You can delete a compatibility, and any nodes that are part of a node pool because of this compatibility are removed from the node pool.

⚠ CAUTION

Deleting a compatibility could result in unintended consequences. For example, if you delete a compatibility, and fewer than three compatible nodes are removed from the node pool, those nodes would be removed from your cluster's available pool of storage. The next time the SmartPools job runs, data on those nodes would be restriped elsewhere on the cluster, which could be a lengthy process. If three or more compatible nodes are removed from the node pool, these nodes would form their own node pool, but data could still be restriped. Any file pool policy pointing to the original node pool would now point to the node pool's tier if one existed, or otherwise to a new tier created by OneFS.

Procedure

1. Run the `isi storagepool compatibilities active delete` command.

The following command deletes a compatibility with an ID number of 1:

```
isi storagepool compatibilities active delete 1
```

OneFS provides a summary of the results of executing the command, including the node pools that will be affected by the compatibility removal, and requires you to confirm the operation.

2. To proceed, type **yes**, and then press ENTER. To cancel, type **no**, and then press ENTER.

Results

If you allowed the operation to proceed, OneFS splits any merged node pools, or unprovisions any previously compatible nodes fewer than three in number.

Create a node pool manually

You can create node pools manually if autoprovisioning does not meet your requirements.

When you add new nodes to your cluster, OneFS places these nodes into node pools. This process is called autoprovisioning. For some workflows, you might prefer to create node pools manually. A manually created node pool must have at least three nodes, identified by the logical node numbers (LNNs).

⚠ CAUTION

It is recommended that you enable OneFS to provision nodes automatically. Manually created node pools might not provide the same performance and efficiency as automatically managed node pools, particularly if your changes result in fewer than 20 nodes in the manual node pool.

Procedure

1. Run the `isi storagepool nodepools create` command.

You can specify the nodes to be added to a nodepool by a comma-delimited list of LNNs (for example, `--lnns 1,2,5`) or by using ranges (for example, `--lnns 5-8`).

The following command creates a node pool by specifying the LNNs of three nodes to be included.

```
isi storagepool nodepools create PROJECT-1 --lnns 1,2,5
```

Add a node to a manually managed node pool

You can add a node to a manually managed node pool.

If you specify a node that is already part of another node pool, OneFS removes the node from the original node pool and adds it to the manually managed node pool.

Procedure

1. Run the `isi storagepool nodepools modify` command.

The following command adds nodes with the LNNs (logical node numbers) of 3, 4, and 10 to an existing node pool:

```
isi storagepool nodepools modify PROJECT-1 --lnns 3-4, 10
```

Change the name or protection policy of a node pool

You can change the name or protection policy of a node pool.

Procedure

1. Run the `isi storagepool nodepools modify` command.

The following command changes the name and protection policy of a node pool:

```
isi storagepool nodepools modify PROJECT-1 --set-name PROJECT-A \
--protection-policy +2:1
```

Remove a node from a manually managed node pool

You can remove a node from a manually managed node pool.

If you attempt to remove nodes from either a manually managed or automatically managed node pool so that the removal leaves only one or two nodes in the pool, the removal fails. You can, however, move all nodes from an autoprovisioned node pool into one that is manually managed.

When you remove a node from the manually managed node pool, OneFS autoprovisions the node into another node pool of the same equivalence class.

Procedure

1. Run the `isi storagepool nodepools modify` command.

The following command removes two nodes, identified by its LNNs (logical node numbers) from a node pool.

```
isi storagepool nodepools modify ARCHIVE_1 --remove-lnns 3,6
```


LNN values can be specified as a range, for example, `--lanns=1-3`, or in a comma-separated list, for example, `--lanns=1,2,5,9`.

Managing L3 cache from the command-line interface

L3 cache can be administered globally or on specific node pools. If you choose to, you can also revert SSDs back to storage drives. In Isilon HD400 node pools, SSDs are exclusively for L3 cache purposes. On these nodes, L3 cache is turned on by default and cannot be turned off.

Set L3 cache as the default for new node pools

You can set L3 cache as the default, so that when new node pools are created, L3 cache is enabled automatically.

Before you begin

L3 cache is effective only on nodes that include SSDs. If none of your nodes has SSD storage, there is no need to enable L3 cache as the default.

Procedure

1. Run the `isi storagepool settings modify` command.

The following command sets L3 cache enabled as the default for new node pools that are added.

```
isi storagepool settings modify --ssd-l3-cache-default-enabled yes
```

2. Run the `isi storagepool settings view` command to confirm that the **SSD L3 Cache Default Enabled** attribute is set to **Yes**.

Enable L3 cache on a specific node pool

You can enable L3 cache for a specific node pool. This is useful when only some of your node pools are equipped with SSDs.

Procedure

1. Run the `isi storagepool nodepools modify` command on a specific node pool.

The following command enables L3 cache on a node pool named `hq_datastore`:

```
isi storagepool nodepools modify hq_datastore --l3 true
```

If the SSDs on the specified node pool previously were used as storage drives, a message appears asking you to confirm the change.

2. If prompted, type **yes**, and then press ENTER.

Restore SSDs to storage drives for a node pool

You can disable L3 cache for SSDs on a specific node pool and restore those SSDs to storage drives.

Note

On HD400 node pools, SSDs are used only for L3 cache, which is turned on by default and cannot be turned off. If you attempt to turn off L3 cache on an HD400 node pool through the command-line interface, OneFS generates this error message: `Disabling L3 not supported for the given node type.`

Procedure

1. Run the `isi storagepool nodepools modify` command on a specific node pool.

The following command disables L3 cache on a node pool named `hq_datastore`:

```
isi storagepool nodepools modify hq_datastore --l3 false
```

2. At the confirmation prompt, type **yes**, and then press ENTER.

Managing tiers

You can move node pools into tiers to optimize file and storage management.

Managing tiers requires `ISI_PRIV_SMARTPOOLS` or higher administrative privileges.

Create a tier

You can create a tier to group together one or more node pools for specific storage purposes.

Depending on the types of nodes in your cluster, you can create tiers for different categories of storage, for example, an archive tier, performance tier, or general-use tier. After creating a tier, you need to add the appropriate node pools to the tier.

Procedure

1. Run the `isi storagepool tiers create` command.

The following command creates a tier named `ARCHIVE_1`, and adds node pools named `hq_datastore1` and `hq_datastore2` to the tier.

```
isi storagepool tiers create ARCHIVE_1 --children hq_datastore1
--children hq_datastore2
```

Add or move node pools in a tier

You can group node pools into tiers and move node pools from one tier to another.

Procedure

1. Run the `isi storagepool nodepools modify` command.

The following example adds a node pool named `PROJECT-A` to a tier named `ARCHIVE_1`.

```
isi storagepool nodepools modify PROJECT-A --tier ARCHIVE_1
```

If the node pool, PROJECT-A, happened to be in another tier, the node pool would be moved to the ARCHIVE_1 tier.

Rename a tier

A tier name can contain alphanumeric characters and underscores but cannot begin with a number.

Procedure

1. Run the `isi storagepool tiers modify` command.

The following command renames a tier from ARCHIVE_1 to ARCHIVE_A:

```
isi storagepool tiers modify ARCHIVE_1 --set-name ARCHIVE_A
```

Delete a tier

When you delete a tier, its node pools remain available and can be added to other tiers.

Procedure

1. Run the `isi storagepool tiers delete` command.

The following command deletes a tier named **ARCHIVE_A**:

```
isi storagepool tiers delete ARCHIVE_A
```

Creating file pool policies

You can configure file pool policies to identify logical groups of files called file pools, and you can specify storage operations for these files.

You must activate a SmartPools license before you can create file pool policies, and the administrator must have `ISI_PRIV_SMARTPOOLS` or higher administrative privileges to create file pool policies.

File pool policies have two parts: file-matching criteria that define a file pool, and the actions to be applied to the file pool. You can define file pools based on characteristics, such as file type, size, path, birth, change, and access timestamps, and combine these criteria with Boolean operators (AND, OR).

In addition to file-matching criteria, you can identify a variety of actions to apply to the file pool. These actions include:

- Setting requested protection and data-access optimization parameters
- Identifying data and snapshot storage targets
- Defining data and snapshot SSD strategies
- Enabling or disabling SmartCache

For example, to free up disk space on your performance tier (S-series node pools), you could create a file pool policy to match all files greater than 25 MB in size, which have not been accessed or modified for more than a month, and move them to your archive tier (NL-series node pools).

You can configure and prioritize multiple file pool policies to optimize file storage for your particular work flows and cluster configuration. When the SmartPools job runs, by default once a day, it applies file pool policies in priority order. When a file pool matches the criteria defined in a policy, the actions in that policy are applied, and lower-priority custom policies are ignored for the file pool.

After the list of custom file pool policies is traversed, if any of the actions are not applied to a file, the actions in the default file pool policy are applied. In this way, the default file pool policy ensures that all actions apply to every file.

Note

You can reorder the file pool policy list at any time, but the default file pool policy is always last in the list of file pool policies.

OneFS also provides customizable template policies that you can copy to make your own policies. These templates, however, are only available from the OneFS web administration interface.

Create a file pool policy

You can create a file pool policy to match specific files and apply SmartPools actions to the matched file pool. SmartPools actions include moving files to certain storage tiers, changing the requested protection levels, and optimizing write performance and data access.

CAUTION

If existing file pool policies direct data to a specific storage pool, do not configure other file pool policies that match this data with `anywhere` for the `--data-storage-target` setting. Because the specified storage pool is included when you use `anywhere`, you should target specific storage pools to avoid unintentional file storage locations.

Procedure

1. Run the `isi filepool policies create` command.

The following command creates a file pool policy that archives older files to a specific storage tier:

```
isi filepool policies create ARCHIVE_OLD
--description "Move older files to archive storage"
--data-storage-target ARCHIVE_TIER --data-ssd-strategy metadata
--begin-filter --file-type=file --and --birth-time=2013-09-01
--operator=lt --and --accessed-time=2013-12-01 --operator=lt
--end-filter
```

Results

The file pool policy is applied when the next scheduled SmartPools job runs. By default, the SmartPools job runs once a day; however, you can also start the SmartPools job manually.

File-matching options for file pool policies

You can configure a file pool policy for files that match specific criteria.

The following file-matching options can be specified when you create or edit a file pool policy.

Note

OneFS supports UNIX shell-style (glob) pattern matching for file name attributes and paths.

The following table lists the file attributes that you can use to define a file pool policy.

File attribute	Specifies
Name	<p>Includes or excludes files based on the file name.</p> <p>You can specify whether to include or exclude full or partial names that contain specific text. Wildcard characters are allowed.</p>
Path	<p>Includes or excludes files based on the file path.</p> <p>You can specify whether to include or exclude full or partial paths that contain specified text. You can also include the wildcard characters *, ?, and [].</p>
File type	<p>Includes or excludes files based on one of the following file-system object types:</p> <ul style="list-style-type: none"> • File • Directory • Other
Size	<p>Includes or excludes files based on their size.</p> <hr/> <p>Note</p> <p>File sizes are represented in multiples of 1024, not 1000.</p> <hr/>
Modified	<p>Includes or excludes files based on when the file was last modified.</p> <p>In the web administration interface, you can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.</p>
Created	<p>Includes or excludes files based on when the file was created.</p> <p>In the web administration interface, you can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.</p>
Metadata changed	<p>Includes or excludes files based on when the file metadata was last modified. This option is available only if the global access-time-tracking option of the cluster is enabled.</p> <p>In the web administration interface, you can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.</p>
Accessed	<p>Includes or excludes files based on when the file was last accessed based on the following units of time:</p> <p>In the web administration interface, you can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.</p>

File attribute	Specifies
	<p>Note</p> <p>Because it affects performance, access time tracking as a file pool policy criterion is disabled by default.</p>
File attribute	Includes or excludes files based on a custom user-defined attribute.

Valid wildcard characters

You can combine wildcard characters with file-matching options to define a file pool policy.

OneFS supports UNIX shell-style (glob) pattern matching for file name attributes and paths.

The following table lists the valid wildcard characters that you can combine with file-matching options to define a file pool policy.

Wildcard	Description
*	Matches any string in place of the asterisk. For example, <code>m*</code> matches <code>movies</code> and <code>m123</code> .
[a-z]	<p>Matches any characters contained in the brackets, or a range of characters separated by a hyphen. For example, <code>b[aei]t</code> matches <code>bat</code>, <code>bet</code>, and <code>bit</code>, and <code>1[4-7]2</code> matches <code>142</code>, <code>152</code>, <code>162</code>, and <code>172</code>.</p> <p>You can exclude characters within brackets by following the first bracket with an exclamation mark. For example, <code>b[!ie]</code> matches <code>bat</code> but not <code>bit</code> or <code>bet</code>.</p> <p>You can match a bracket within a bracket if it is either the first or last character. For example, <code>[c]at</code> matches <code>cat</code> and <code>[at]</code>.</p> <p>You can match a hyphen within a bracket if it is either the first or last character. For example, <code>car[-s]</code> matches <code>cars</code> and <code>car-</code>.</p>
?	Matches any character in place of the question mark. For example, <code>t?p</code> matches <code>tap</code> , <code>tip</code> , and <code>top</code> .

SmartPools settings

SmartPools settings include directory protection, global namespace acceleration, L3 cache, virtual hot spare, spillover, requested protection management, and I/O optimization management.

Settings in Web Admin	Settings in CLI	Description	Notes
Increase directory protection to a higher level than its contents	<code>--protect-directories-one-level-higher</code>	Increases the amount of protection for directories at a higher level than the directories and files that they contain, so that	This setting should be enabled (the default). When this setting is disabled, the directory that contains a file pool

Settings in Web Admin	Settings in CLI	Description	Notes
		<p>data that is not lost can still be accessed.</p> <p>When device failures result in data loss (for example, three drives or two nodes in a +2:1 policy), enabling this setting ensures that intact data is still accessible.</p>	<p>is protected according to your protection-level settings, but the devices used to store the directory and the file may not be the same. There is potential to lose nodes with file data intact but not be able to access the data because those nodes contained the directory.</p> <p>As an example, consider a cluster that has a +2 default file pool protection setting and no additional file pool policies. OneFS directories are always mirrored, so they are stored at 3x, which is the mirrored equivalent of the +2 default.</p> <p>This configuration can sustain a failure of two nodes before data loss or inaccessibility. If this setting is enabled, all directories are protected at 4x. If the cluster experiences three node failures, although individual files may be inaccessible, the directory tree is available and provides access to files that are still accessible.</p> <p>In addition, if another file pool policy protects some files at a higher level, these too are accessible in the event of a three-node failure.</p>
<p>Enable global namespace acceleration</p>	<p>--global-namespace-acceleration-enabled</p>	<p>Specifies whether to allow per-file metadata to use SSDs in the node pool.</p> <ul style="list-style-type: none"> • When disabled, restricts per-file metadata to the storage pool policy of the file, except in the case of spillover. This is the default setting. • When enabled, allows per-file metadata to use the SSDs in any node pool. 	<p>This setting is available only if 20 percent or more of the nodes in the cluster contain SSDs and at least 1.5 percent of the total cluster storage is SSD-based.</p> <p>If nodes are added to or removed from a cluster, and the SSD thresholds are no longer satisfied, GNA becomes inactive. GNA remains enabled, so that when the SSD thresholds are met again, GNA is reactivated.</p> <hr/> <p>Note</p> <p>Node pools with L3 cache enabled are effectively invisible for GNA purposes. All ratio calculations for GNA are done exclusively for node pools without L3 cache enabled.</p>


Settings in Web Admin	Settings in CLI	Description	Notes
Use SSDs as L3 Cache by default for new node pools	--ssd-l3-cache-default-enabled	For node pools that include solid-state drives, deploy the SSDs as L3 cache. L3 cache extends L2 cache and speeds up file system performance across larger working file sets.	L3 cache is enabled by default on new node pools. When you enable L3 cache on an existing node pool, OneFS performs a migration, moving any existing data on the SSDs to other locations on the cluster. OneFS manages all cache levels to provide optimal data protection, availability, and performance. In case of a power failure, the data on L3 cache is retained and still available after power is restored.
Virtual Hot Spare	--virtual-hot-spare-deny-writes --virtual-hot-spare-hide-spare --virtual-hot-spare-limit-drives --virtual-hot-spare-limit-percent	Reserves a minimum amount of space in the node pool that can be used for data repair in the event of a drive failure. To reserve disk space for use as a virtual hot spare, select from the following options: <ul style="list-style-type: none"> • Ignore reserved disk space when calculating available free space. Subtracts the space reserved for the virtual hot spare when calculating available free space. • Deny data writes to reserved disk space. Prevents write operations from using reserved disk space. • VHS Space Reserved. You can reserve a minimum number of virtual drives (1-4), as well as a minimum percentage of total disk space (0-20%). 	If you configure both the minimum number of virtual drives and a minimum percentage of total disk space when you configure reserved VHS space, the enforced minimum value satisfies both requirements. If this setting is enabled and Deny new data writes is disabled, it is possible for the file system utilization to be reported at more than 100%.
Enable global spillover	--no-spillover	Specifies how to handle write operations to a node pool that is not writable.	<ul style="list-style-type: none"> • When enabled, redirects write operations from a node pool that is not writable either to another node pool or anywhere on the cluster (the default). • When disabled, returns a disk space error for write operations to a node pool that is not writable.

Settings in Web Admin	Settings in CLI	Description	Notes
Spillover Data Target	--spillover-target --spillover-anywhere	Specifies another storage pool to target when a storage pool is not writable.	When spillover is enabled, but it is important that data writes do not fail, select anywhere for the Spillover Data Target setting, even if file pool policies send data to specific pools.
Manage protection settings	--automatically-manage-protection	When this setting is enabled, SmartPools manages requested protection levels automatically.	When Apply to files with manually-managed protection is enabled, overwrites any protection settings that were configured through File System Explorer or the command-line interface.
Manage I/O optimization settings	--automatically-manage-io-optimization	When enabled, uses SmartPools technology to manage I/O optimization.	When Apply to files with manually-managed I/O optimization settings is enabled, overwrites any I/O optimization settings that were configured through File System Explorer or the command-line interface.

Default file pool requested protection settings

Default protection settings include specifying the data storage target, snapshot storage target, requested protection, and SSD strategy for files that are filtered by the default file pool policy.

Settings (Web Admin)	Settings (CLI)	Description	Notes
Storage Target	--data-storage-target --data-ssd-strategy	<p>Specifies the storage pool (node pool or tier) that you want to target with this file pool policy.</p> <p>CAUTION</p> <p>If existing file pool policies direct data to a specific storage pool, do not configure other file pool policies with <code>anywhere</code> for the Data storage target option. Because the specified storage pool is included when you use <code>anywhere</code>, target specific storage pools to avoid unintentional file storage locations.</p> <p>Select one of the following options to define your SSD strategy:</p>	<p>Note</p> <p>If GNA is not enabled and the storage pool that you choose to target does not contain SSDs, you cannot define an SSD strategy.</p> <p>Use SSDs for metadata read acceleration writes both file data and metadata to HDD storage pools but adds an additional SSD mirror if possible to accelerate read performance. Uses HDDs to provide reliability and an extra metadata mirror to SSDs, if available, to improve read</p>

Settings (Web Admin)	Settings (CLI)	Description	Notes
		<p>Use SSDs for metadata read acceleration Default. Write both file data and metadata to HDDs and metadata to SSDs. Accelerates metadata reads only. Uses less SSD space than the Metadata read/write acceleration setting.</p> <p>Use SSDs for metadata read/write acceleration Write metadata to SSD pools. Uses significantly more SSD space than Metadata read acceleration, but accelerates metadata reads and writes.</p> <p>Use SSDs for data & metadata Use SSDs for both data and metadata. Regardless of whether global namespace acceleration is enabled, any SSD blocks reside on the storage target if there is room.</p> <p>Avoid SSDs Write all associated file data and metadata to HDDs only.</p> <p> CAUTION</p> <p>Use this to free SSD space only after consulting with Isilon Technical Support personnel; the setting can negatively affect performance.</p>	<p>performance. Recommended for most uses.</p> <p>When you select Use SSDs for metadata read/write acceleration, the strategy uses SSDs, if available in the storage target, for performance and reliability. The extra mirror can be from a different storage pool using GNA enabled or from the same node pool.</p> <p>Neither the Use SSDs for data & metadata strategy nor the Use SSDs for data & metadata strategy result in the creation of additional mirrors beyond the normal requested protection. Both file data and metadata are stored on SSDs if available within the file pool policy. This option requires a significant amount of SSD storage.</p>
Snapshot storage target	--snapshot-storage-target --snapshot-ssd-strategy	Specifies the storage pool that you want to target for snapshot storage with this file pool policy. The settings are the same as those for data storage target, but apply to snapshot data.	Notes for data storage target apply to snapshot storage target
Requested protection	--set-requested-protection	<p>Default of storage pool. Assign the default requested protection of the storage pool to the filtered files.</p> <p>Specific level. Assign a specified requested protection to the filtered files.</p>	To change the requested protection, select a new value from the list.

Default file pool I/O optimization settings

You can manage the I/O optimization settings that are used in the default file pool policy, which can include files with manually managed attributes.

To allow SmartPools to overwrite optimization settings that were configured using File System Explorer or the `isi set` command, select the **Including files with manually-managed I/O optimization settings** option in the **Default Protection Settings** group. In the CLI, use the `--automatically-manage-io-optimization` option with the `isi storagepool settings modify` command.

Setting (Web Admin)	Setting (CLI)	Description	Notes
Write Performance	--enable-coalescer	Enables or disables SmartCache (also referred to as the coalescer).	Enable SmartCache is the recommended setting for optimal write performance. With asynchronous writes, the Isilon server buffers writes in memory. However, if you want to disable this buffering, we recommend that you configure your applications to use synchronous writes. If that is not possible, disable SmartCache.
Data Access Pattern	--data-access-pattern	Defines the optimization settings for accessing concurrent, streaming, or random data types.	Files and directories use a concurrent access pattern by default. To optimize performance, select the pattern dictated by your workflow. For example, a workflow heavy in video editing should be set to Optimize for streaming access . That workflow would suffer if the data access pattern was set to Optimize for random access .

Managing file pool policies

You can perform a number of file pool policy management tasks.

File pool policy management tasks include:

- Modifying file pool policies
- Modifying the default file pool policy
- Creating a file pool policy from a template
- Reordering file pool policies
- Deleting file pool policies

Note

You can create a file pool policy from a template only in the OneFS web administration interface.

Modify a file pool policy

You can modify the name, description, filter criteria, and the protection and I/O optimization settings applied by a file pool policy.

CAUTION

If existing file pool policies direct data to a specific storage pool, do not configure other file pool policies with *anywhere* for the Data storage target option. Because the specified storage pool is included when you use *anywhere*, target specific storage pools to avoid unintentional file storage locations.

Procedure

1. Run the `isi filepool policies list` command to view a list of available file pool policies.
A tabular list of policies and their descriptions appears.
2. Run the `isi filepool policies view` command to view the current settings of a file pool policy.

The following example displays the settings of a file pool policy named **ARCHIVE_OLD**.

```
isi filepool policies view ARCHIVE_OLD
```

3. Run the `isi filepool policies modify` command to change a file pool policy.

The following example modifies the settings of a file pool policy named **ARCHIVE_OLD**.

```
isi filepool policies modify ARCHIVE_OLD --description
"Move older files to archive storage" --data-storage-target TIER_A
--data-ssd-strategy metadata-write --begin-filter --file-type=file
--and --birth-time=2013-01-01 --operator=lt --and --accessed-time=
2013-09-01 --operator=lt --end-filter
```

Results

Changes to the file pool policy are applied when the next SmartPools job runs. However, you can also manually run the SmartPools job immediately.

Modify default storage pool settings

You can modify default storage pool settings for requested protection, I/O optimization, global namespace acceleration, virtual hot spare, and spillover.

Procedure

1. Run the `isi storagepool settings modify` command.

The following command specifies automatic file protection and I/O optimization, disables global namespace acceleration, specifies a percentage of storage for a virtual hot spare, and enables L3 cache for node pools with SSDs:

```
isi storagepool settings modify
--automatically-manage-protection files_at_default
--automatically-manage-io-optimization files_at_default
--global-namespace-acceleration-enabled no
--virtual-hot-spare-limit-percent 5
--ssd-l3-cache-default-enabled yes
```

Results

OneFS applies your changes to any files managed by the default file pool policy the next time the SmartPools job runs.

Configure default file pool policy settings

Files that are not managed by custom file pool policies are managed by the default file pool policy. You can configure the default file pool policy settings.

Procedure

1. Run the `isi filepool default-policy view` command to display the current default file pool policy settings.

Output similar to the following example appears:

```
Set Requested Protection: default
    Data Access Pattern: random
        Enable Coalescer: True
    Data Storage Target: anywhere
        Data SSD Strategy: metadata
    Snapshot Storage Target: anywhere
        Snapshot SSD Strategy: metadata
```

2. Run the `isi filepool default-policy modify` command to change default settings.

The following command modifies all default settings:

```
isi filepool default-policy modify --set-requested-protection +2 \
--data-access-pattern concurrency --enable-coalescer false \
--data-storage-target ARCHIVE_A --data-ssd-strategy avoid \
--snapshot-storage-target ARCHIVE_A --snapshot-ssd-strategy avoid
```

3. Run the `isi filepool default-policy view` command again to ensure that default file pool policy settings reflect your intentions.

Results

OneFS implements the new default file pool policy settings when the next scheduled SmartPools job runs and applies these settings to any files that are not managed by a custom file pool policy.

Prioritize a file pool policy

You can change the priority order of a file pool policy.

File pool policies are evaluated in descending order according to their position in the file pool policies list. By default, when you create a new policy, it is inserted immediately above the default file pool policy. You can assign a policy a different priority by moving it up or down in the list. The default policy is always the last in priority, and applies to all files that are not matched by any other file pool policy.

Procedure

1. Run the `isi filepool policies list` command to view the list of available file pool policies and their priority order.

Output similar to the following appears:

```
Name          Description
-----
ARCHIVE_1     Move older files to archive tier
MOVE-LARGE    Move large files to archive tier
PERFORM_1     Move recent files to perf. tier
-----
Total: 3
```

2. Run the `isi filepool policies modify` command to change the priority of a file pool policy.

The following example changes the priority of a file pool policy named `PERFORM_1`.

```
isi filepool policies modify PERFORM_1 --apply-order 1
```

3. Run the `isi filepool policies list` command again to ensure that the policy list displays the correct priority order.

Delete a file pool policy

You can delete a file pool policy.

Delete a file pool policy only if you are aware of, or unconcerned with, the consequences.

Procedure

1. Run the `isi filepool policies delete` command.

The following example deletes a file pool policy named `ARCHIVE_1`.

```
isi filepool policies delete ARCHIVE_1
```

The system asks you to confirm the deletion.

2. Type **yes**, then press ENTER.

Results

The file pool policy is removed. When you delete a policy, its file pool will be controlled either by another policy or by the default file pool policy the next time the SmartPools job runs.

Monitoring storage pools

You can access information on storage pool health and usage.

The following information is available:

- File pool policy health
- SmartPools health, including tiers, node pools, and subpools
- For each storage pool, percentage of HDD and SSD disk space usage
- SmartPools job status

Monitor storage pools

You can view storage pool status and details.

Details include the names of tiers and associated node pools, requested protection, HDD and SSD capacities and usage.

Procedure

1. Run the `isi storagepool list` command.

Output similar to the following example appears:

Name	Nodes	Protect	HDD	Total	%	SSD	Total	%
PERF_TIER	1-3	-	12.94T	17.019T	26.99%	0.4T	1.2T	33.00%
- s-series	1-3	+2:1	12.94T	17.019T	26.99%	0.4T	1.2T	33.00%
HOME_TIER	4-6	-	16.59T	19.940T	77.73%	0b	0b	0.00%
- x-series	4-6	+2:1	16.59T	19.940T	77.73%	0b	0b	0.00%
ARCHIVE_1	7-9	-	100.8T	200.60T	49.88%	0b	0b	0.00%
- nl-serie	7-9	+2:1	100.8T	200.60T	49.88%	0b	0b	0.00%
Total: 6			200.5G	17.019G	26.99%	0b	0b	0.00%

View the health of storage pools

You can view the health of storage pools.

Procedure

1. Run the `isi storagepool health` command.

The following command, using the verbose option, displays a tabular description of storage pool health:

```
isi storagepool health --verbose
```

View results of a SmartPools job

You can review detailed results from the last time the SmartPools job ran.

The SmartPools job, by default, runs once a day. It processes the file pool policies that you have created to manage storage on your cluster.

Procedure

1. Run the `isi job events list` command.

A tabular listing of the most recent system jobs appears. The listing for the SmartPools job is similar to the following example:

```
2014-04-28T02:00:29 SmartPools [105] Succeeded
```

2. Locate the SmartPools job in the listing, and make note of the number in square brackets.

This is the job ID number.

3. Run the `isi job reports view` command, using the job ID number.

The following example displays the report for a SmartPools job with the job ID of 105.

```
isi job reports view 105
```

Results

The SmartPools report shows the outcome of all of the file pool policies that were run, including summaries for each policy, and overall job information such as elapsed time, LINs traversed, files and directories processed, and memory and I/O statistics.

Storage pool commands

You can monitor and manage node pools, tiers, and file pool policies and settings through the OneFS command-line interface. Some storage pool functions are available only if you activate a SmartPools license on the cluster.

isi filepool apply

Applies all file pool policies to the specified file or directory path. If no policy matches the file or directory path, OneFS applies the default file pool policy.

Syntax

```
isi filepool apply <path>
  [--path] <path>
  [--dont-restripe]
  [--nop]
  [--stats]
  [--quiet]
  [--recurse]
  [--verbose]
```

Options

{--path | -p} <path>

Specifies the path to the file to be processed. This parameter is required.

{--dont-restripe | -d}

Changes the per-file policies without restriping the file.

{--nop | -n}

Calculates the specified settings without actually applying them. This option is best used with `--verbose` or `--stats`.

{--stats | -s}

Displays statistics on the files processed.

```
{--quiet | -q}
```

Suppresses warning messages.

```
{--recurse | -r}
```

Specifies recursion through directories.

```
{--verbose | -v}
```

Displays the configuration settings to be applied. We recommend using verbose mode. Otherwise the command would not display any screen output except for error messages.

Examples

These examples show the results of running `isi filepool apply` in verbose mode. In the examples, the output shows the results of comparing the path specified with each file pool policy. The `recurse` option is set so that all files in the `/ifs/data/projects` path are matched against all file pool policies. The first policy listed is always the system default policy. In this example, the second match is to the file pool policy

Technical Data.

```
isi filepool apply --path=/ifs/data/projects --verbose --recurse
```

```
Processing file /ifs/data/projects
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0
New File Attributes
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0

{'default' :
  {'Policy Number': -2,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':1, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 1,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
  },
 'system':

'Technical Data':
  {'Policy Number': 0,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':0, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 0,
```



```
'File data placed on HDDs': {'head':0, 'snapshot': 0},
'File data placed on SSDs': {'head':0, 'snapshot': 0},
```

This example shows the result of using the `--nop` option to calculate the results that would be produced by applying the file pool policy.

```
isi filepool apply --path=/ifs/data/projects --nop --verbose
```

```
Processing file /ifs/data/projects
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0
New File Attributes
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0

{'default' :
  {'Policy Number': -2,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':1, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 1,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
  },
'system':
  {'Policy Number': -1,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':0, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 0,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
  },
```

isi filepool default-policy modify

Modifies default file pool policy settings. The default file pool policy specifies storage settings for all files to which a higher-priority file pool policy does not apply.

Syntax

```
isi filepool default-policy modify
  [--data-access-pattern {random | concurrency | streaming}]
  [--set-requested-protection {default | +1 | +2:1 | +2 | +3:1 | +3 |
+4 | 2x | 3x | 4x | 5x | 6x | 7x | 8x}]
  [--data-storage-target <string>]
```

```
[--data-ssd-strategy {metadata | metadata-write | data | avoid}]
[--snapshot-storage-target <string>]
[--snapshot-ssd-strategy {metadata | metadata-write | data | avoid}]
[--enable-coalescer {yes | no}]
[{-verbose | -v}]
```

Options

--data-access-pattern *<string>*

Specifies the preferred data access pattern, one of `random`, `streaming`, or `concurrent`.

--set-requested-protection *<string>*

Specifies the requested protection for files that match this filepool policy (for example, `+2:1`).

--data-storage-target *<string>*

Specifies the node pool or tier to which the policy moves files on the local cluster.

--data-ssd-strategy *<string>*

Specifies how to use SSDs to store local data.

avoid

Writes all associated file data and metadata to HDDs only.

metadata

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if SSDs are available. The SSD mirror is in addition to the number required to satisfy the requested protection. Enabling global namespace acceleration (GNA) makes read acceleration available to files in node pools that do not contain SSDs.

metadata-write

Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata** setting. Enabling GNA does not affect read/write acceleration.

data

Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly more storage space compared with the other SSD strategy options.

--snapshot-storage-target *<integer>*

The ID of the node pool or tier chosen for storage of snapshots.

--snapshot-ssd-strategy *<string>*

Specifies how to use SSDs to store snapshots. Valid options are `metadata`, `metadata-write`, `data`, `avoid`. The default is `metadata`.

--enable-coalescer {yes | no}

Enable or disable the coalescer, also referred to as SmartCache. The coalescer protects write-back data through a combination of RAM and stable storage. It is enabled by default, and should be disabled only in cooperation with EMC Isilon Customer Support.

--verbose

Displays more detailed information.

Example

The command shown in the following example modifies the default file pool policy in several ways. The command sets the `requested-protection-level` to `+2:1`, sets the `data-storage-target` to `anywhere` (the system default), and changes the `data--ssd-strategy` to `metadata-write`.

```
isi filepool default-policy modify --set-requested-protection=+2:1 --
data-storage-target=anywhere --data-ssd-strategy=metadata-write
```

isi filepool default-policy view

View default file pool policy settings. The default file pool policy specifies storage settings for all files to which a higher-priority file pool policy does not apply.

Syntax

```
isi filepool default-policy view
```

The following display shows sample command output:

```
Set Requested Protection: default
  Data Access Pattern: random
    Enable Coalescer: True
  Data Storage Target: anywhere
    Data SSD Strategy: metadata
  Snapshot Storage Target: anywhere
    Snapshot SSD Strategy: metadata
```

isi filepool policies create

Create a custom file pool policy to identify a specific storage target and perform other actions on matched files and directories.

Syntax

```
isi filepool policies create <name>
  [--description <string>]
  [--begin-filter{<predicate> <operator> <link>}...--end-filter]
  [--apply-order <integer>]
  [--data-access-pattern {random | concurrency | streaming}]
  [--set-requested-protection {default | +1 | +2:1 | +2 | +3:1 | +3 |
+4 | 2x | 3x | 4x | 5x | 6x | 7x | 8x}]
  [--data-storage-target <string>]
  [--data-ssd-strategy {metadata | metadata-write | data | avoid}]
  [--snapshot-storage-target <string>]
  [--snapshot-ssd-strategy {metadata | metadata-write | data | avoid}]
  [--enable-coalescer {Yes | No}]
  [{--verbose | -v}]
```

Options

<name>

Specifies the name of the file pool policy to create.

--begin-filter {<predicate> <operator> <link>}... --end-filter

Specifies the file-matching criteria that determine the files to be managed by the filepool policy.

Each file matching criterion consists of three parts:

- Predicate. Specifies what attribute(s) to filter on. You can filter by path, name, file type, timestamp, or custom attribute, or use a combination of these attributes.
- Operator. Qualifies an attribute (for example, birth time) to describe a relationship to that attribute (for example, before).
- Link - Combines attributes using AND and OR statements.

The following predicates are valid:

`--size=<nn>{{B | KB | MB | GB | TB | PB}}`

Selects files according to the specified size.

`--path=<pathname>`

Selects files relative to the specified pathname.

`--file-type= <value>`

Selects only the specified file-system object type.

The following values are valid:

file

Specifies regular files.

directory

Specifies directories.

other

Specifies links.

`--name= <value> [--case-sensitive= {true | false}]`

Selects only files whose names match the specified string. Use `--case-sensitive=true` to enable case-sensitivity.

When forming the name, you can include the following wildcards:

- *
- []
- ?

`--birth-time=<timestamp>`

Selects files that were created relative to the specified date and time. Timestamp arguments are formed as **YYYY-MM-DDTHH:MM:SS**. For example, **2013-09-01T08:00:00** specifies a timestamp of September 1, 2013 at 8:00 A.M. You can use `--operator=` with an argument of `gt` to mean after the timestamp or `lt` to mean before the timestamp.

`--changed-time=<timestamp>`

Selects files that were modified relative to the specified date and time.

`--metadata-changed-time=<timestamp>`

Selects files whose metadata was modified relative to the specified date and time.

`--accessed-time=<timestamp>`

Selects files that were accessed at the specified time interval.

`--custom-attribute=<value>`

Selects files based on a custom attribute.

You can use the `operator=` option to specify a qualifier for the file-matching criterion. Specify operators in the following form:

```
--operator=<value>
```

The following operator values are valid:

Value	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to
not	Not

Link arguments can be used to specify multiple file-matching criteria. The following links are valid:

```
--and
```

Connects two file-matching criteria where files must match both criteria.

```
--or
```

Connects two file-matching criteria where files must match one or the other criteria.

```
--description <string>
```

Specifies a description of the filepool policy

```
--apply-order <integer>
```

Specifies the order index for execution of this policy.

```
--data-access-pattern <string>
```

Data access pattern random, streaming or concurrent.

```
--set-requested-protection <string>
```

Specifies a protection level for files that match this filepool policy (e.g., +3, +2:3, 8x).

```
--data-storage-target <string>
```

The name of the node pool or tier to which the policy moves files on the local cluster. If you do not specify a data storage target, the default is **anywhere**.

```
--data-ssd-strategy <string>
```

Specifies how to use SSDs to store local data.

avoid

Writes all associated file data and metadata to HDDs only.

metadata

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if SSDs are available. The SSD mirror is in addition to the number required to satisfy the requested protection. Enabling GNA makes read acceleration available to files in node pools that do not contain SSDs.

metadata-write

Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata** setting. Enabling GNA does not affect read/write acceleration.

data

Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly increases storage requirements compared with the other SSD strategy options.

--snapshot-storage-target *<string>*

The name of the node pool or tier chosen for storage of snapshots. If you do not specify a snapshot storage target, the default is **anywhere**.

--snapshot-ssd-strategy *<string>*

Specifies how to use SSDs to store snapshots. Valid options are `metadata`, `metadata-write`, `data`, `avoid`. The default is `metadata`.

--enable-coalescer {Yes | No}

Enable the coalescer.

--verbose

Displays more detailed information.

Examples

The following example creates a file pool policy that moves all files in directory `/ifs/data/chemical/arco/finance` to the local storage target named `Archive_2`.

```
isi filepool policies create Save_Fin_Data --begin-filter
--path=/ifs/data/chemical/arco/finance --end-filter
--data-storage-target Archive_2 --data-ssd-strategy=metadata
```

The following example matches older files that have not been accessed or modified later than specified dates, and moves the files to an archival tier of storage.

```
isi filepool policies create archive_old
--data-storage-target ARCHIVE_1 --data-ssd-strategy avoid
--begin-filter --file-type=file --and --birth-time=2013-09-01
--operator=lt --and --accessed-time=2013-12-01 --operator=lt
--and --changed-time=2013-12-01 --operator=lt --end-filter
```

isi filepool policies delete

Delete a custom file pool policy. The default file pool policy cannot be deleted.

To list all file pool policies, run the `isi filepool policies list` command.

Syntax

```
isi filepool policies delete <name>
[--force | -f]
[--verbose | -v]
```

Options

<name>

Specifies the name of the file pool policy to be deleted.

`--force`

Deletes the file pool policy without asking for confirmation.

`--verbose`

Displays more detailed information.

Example

The following command deletes a file pool policy named `ARCHIVE_OLD`. The `--force` option circumvents the requirement to confirm the deletion:

```
isi filepool policies delete ARCHIVE_OLD --force
```

isi filepool policies list

List all custom file pool policies configured on the system.

Syntax

```
isi filepool policies list
[--format {table | json | csv | list}]
[{-#no-header | -a}]
[{-#no-footer | -z}]
[{-#verbose | -v}]
```

Options

`--format`

Output the list of file pool policies in a variety of formats. The following values are valid:

- table
- json
- csv
- list

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

Example

The following example lists custom file pool policies in `.csv` format and outputs the list to a file in the OneFS file system.

```
isi filepool policies list --format csv > /ifs/data/policy.csv
```

isi filepool policies modify

Modify a custom file pool policy to identify a specific storage target and perform other actions on matched files and directories.

Syntax

```
isi filepool policies modify <name>
```

```

[--description <string>]
[--begin-filter{<predicate> <operator> <link>}...--end-filter]
[--apply-order <integer>]
[--data-access-pattern {random | concurrency | streaming}]
[--set-requested-protection {default | +1 | +2:1 | +2 | +3:1 | +3 |
+4 | 2x | 3x | 4x | 5x | 6x | 7x | 8x}]
[--data-storage-target <string>]
[--data-ssd-strategy {metadata | metadata-write | data | avoid}]
[--snapshot-storage-target <string>]
[--snapshot-ssd-strategy {metadata | metadata-write | data | avoid}]
[--enable-coalescer {Yes | No}]
[--verbose]

```

Options

<name>

Specifies the name of the file pool policy to create.

--begin-filter {<predicate> <operator> <link>}... --end-filter

Specifies the file-matching criteria that determine the files to be managed by the filepool policy.

Each file matching criterion consists of three parts:

- Predicate. Specifies what attribute(s) to filter on. You can filter by path, name, file type, timestamp, or custom attribute, or use a combination of these attributes.
- Operator. Qualifies an attribute (for example, birth time) to describe a relationship to that attribute (for example, before).
- Link - Combines attributes using AND and OR statements.

The following predicates are valid:

--size=<nn>{B | KB | MB | GB | TB | PB}}

Selects files according to the specified size.

--path=<pathname>

Selects files relative to the specified pathname.

--file-type= <value>

Selects only the specified file-system object type.

The following values are valid:

file

Specifies regular files.

directory

Specifies directories.

other

Specifies links.

--name= <value> [--case-sensitive= {true | false}]

Selects only files whose names match the specified string. Use --case-sensitive=true to enable case-sensitivity.

When forming the name, you can include the following wildcards:

- *
- []

- ?

`--birth-time=<timestamp>`

Selects files that were created relative to the specified date and time. Timestamp arguments are formed as **YYYY-MM-DDTHH:MM:SS**. For example, **2013-09-01T08:00:00** specifies a timestamp of September 1, 2013 at 8:00 A.M. You can use `--operator=` with an argument of `gt` to mean after the timestamp or `lt` to mean before the timestamp.

`--changed-time=<timestamp>`

Selects files that were modified relative to the specified date and time.

`--metadata-changed-time=<timestamp>`

Selects files whose metadata was modified relative to the specified date and time.

`--accessed-time=<timestamp>`

Selects files that were accessed at the specified time interval.

`--custom-attribute=<value>`

Selects files based on a custom attribute.

You can use the `operator=` option to specify a qualifier for the file-matching criterion. Specify operators in the following form:

```
--operator=<value>
```

The following operator values are valid:

Value	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to
not	Not

Link arguments can be used to specify multiple file-matching criteria. The following links are valid:

`--and`

Connects two file-matching criteria where files must match both criteria.

`--or`

Connects two file-matching criteria where files must match one or the other criteria.

`--description <string>`

Specifies a description of the filepool policy

`--apply-order <integer>`

Specifies the order index for execution of this policy.

`--data-access-pattern <string>`

Data access pattern random, streaming or concurrent.

`--set-requested-protection <string>`

Specifies a protection level for files that match this filepool policy (for example, +3, +2:3, 8x).

`--data-storage-target <string>`

The name of the node pool or tier to which the policy moves files on the local cluster.

`--data-ssd-strategy <string>`

Specifies how to use SSDs to store local data.

avoid

Writes all associated file data and metadata to HDDs only.

metadata

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if SSDs are available. The SSD mirror is in addition to the number required to satisfy the requested protection. Enabling GNA makes read acceleration available to files in node pools that do not contain SSDs.

metadata-write

Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata** setting. Enabling GNA does not affect read/write acceleration.

data

Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly increases storage requirements compared with the other SSD strategy options.

`--snapshot-storage-target <string>`

The name of the node pool or tier chosen for storage of snapshots.

`--snapshot-ssd-strategy <string>`

Specifies how to use SSDs to store snapshots. Valid options are `metadata`, `metadata-write`, `data`, `avoid`. The default is `metadata`.

`--enable-coalescer {Yes | No}`

Enable the coalescer.

--verbose

Display more detailed information.

Examples

The following example modifies a file pool policy to move matched files to a different local storage target named `Archive_4`. The next time the SmartPools job runs, matched files would be moved to the new storage target.

```
isi filepool policies modify Save_Fin_Data --begin-filter
--path=/ifs/data/chemical/arco/finance --end-filter
--data-storage-target Archive_4 --data-ssd-strategy=metadata
```

The following example matches older files that have not been accessed or modified later than specified dates, and moves the files to an archival tier of storage.

```
isi filepool policies modify archive_old
--data-storage-target ARCHIVE_1 --data-ssd-strategy avoid
--begin-filter --file-type=file --and --birth-time=2013-06-01
--operator=lt --and --accessed-time=2013-09-01 --operator=lt
--and --changed-time=2013-09-01 --operator=lt --end-filter
```

isi filepool policies view

Displays detailed information on a custom file pool policy.

Run the `isi filepool policies list` command to list the names of all custom file pool policies.

Syntax

```
isi filepool policies view <name>
```

Options

<name>

Specifies the name of the file pool policy. Names must begin with a letter or an underscore and contain only letters, numbers, hyphens, underscores or periods.

isi filepool templates list

Lists available file pool policy templates.

Syntax

```
isi filepool templates list
[--limit <integer>]
[--sort <string>]
[--descending <string>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

`--limit <integer>`

Specifies the number of templates to display.

`--sort <string>`

Sorts data by the field specified.

`--descending <integer>`

Sorts data in descending order.

`--format`

Displays file pool templates in the specified format. The following values are valid:

table

json

csv

list

```
--no-header
    Displays table and CSV output without headers.
--no-footer
    Displays table output without footers.
--verbose
    Displays more detailed information.
```

isi filepool templates view

View the detailed settings in a file pool policy template.

Syntax

```
isi filepool templates view <name>
```

Options

<name>
The name of the template to view.

isi storagepool compatibilities active create

Creates a compatibility to enable an unprovisioned node to join a node pool.

Syntax

```
isi storagepool compatibilities active create <class-1> <class-2>
  [--assess {yes|no}]
  [--verbose]
  [--force]
```

Options

<class-1>

An existing node pool class, one of S200 or X400.

<class-2>

The node class that is compatible with the existing node pool, one of S210 or X410. Note that S210 nodes are only compatible with S200 node pools, and X410 nodes are only compatible with X400 node pools.

```
{--assess | -a} {yes | no}
```

Checks whether the compatibility is valid without actually creating the compatibility.

```
{--verbose | -v}
```

Displays more detailed information.

```
{--force | -f}
```

Performs the action without asking for confirmation.

Examples

The following command creates a compatibility between S200 and S210 nodes without asking for confirmation:

```
isi storagepool compatibilities active create S200 S210 --force
```

isi storagepool compatibilities active delete

Deletes a node compatibility. If fewer than three compatible nodes had been added to an existing node pool, they are removed and become unprovisioned.

Syntax

```
isi storagepool compatibilities active delete <ID>
  [--assess {yes | no}]
  [--verbose]
  [--force]
```

Options

<ID>

The ID number of the compatibility. You can use the `isi storagepool compatibilities active list` command to view the ID numbers of active compatibilities.

{--assess | -a} {yes | no}

Checks the results without actually deleting the compatibility.

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Performs the action without asking for confirmation.

Example

The following command provides information about the results of deleting a compatibility without actually performing the action:

```
isi storagepool compatibilities active delete 1 --assess yes
```

Provided that a compatibility with the ID of 1 exists, OneFS displays information similar to the following example:

```
Deleting compatibility with id 1 is possible.
This delete will cause these nodepools to split:
1: Nodepool s200_0b_0b will be split. A tier will be created and all
resultant nodepools from this split will be incorporated into it. All
filepool policies targeted at the splitting pool will be redirected
towards this new tier. That tier's name is s200_0b_0b-tier
```

isi storagepool compatibilities active list

Lists node compatibilities that have been created.

Syntax

```
isi storagepool compatibilities active list
  [--limit <integer>]
  [--format {table | json |
  csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of compatibilities that are listed.

`{--format | -f}`

Lists active compatibilities in the specified format. The following values are valid:

table

json

csv

list

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

Example

The following command lists active compatibilities:

```
isi storagepool compatibilities active list
```

Command output appears similar to the following example:

```

ID      Class 1   Class 2
-----
1       S200     S210
2       X400     X410
-----
Total: 2

```

isi storagepool compatibilities active view

Displays the details of an active node compatibility.

Syntax

```
isi storagepool compatibilities active view <ID>
```

Options

<ID>

The ID number of the compatibility to view. You can use the `isi storagepool compatibilities active list` command to display the ID numbers of active compatibilities.

Example

The following command displays information about an active compatibility with ID number 1:

```
isi storagepool compatibilities active view 1
```

Output from the command would be similar to the following:

```
ID: 1
Class 1: S200
Class 2: S210
```

isi storagepool compatibilities available list

Lists compatibilities that are available, but not yet created.

Syntax

```
isi storagepool compatibilities available list <name>
  [--limit <integer>]
  [--format {table | json |
  csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of available compatibilities that are listed.

{--format | -f}

Lists available compatibilities in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Example

The following command lists available compatibilities:

```
isi storagepool compatibilities available list
```

If available compatibilities exist, command output appears similar to the following example:

```
Class 1   Class 2
-----
S200     S210
X400     X410
-----
Total: 2
```

isi storagepool health

Displays the health information of storage pools.

Syntax

```
isi storagepool health
```

Options

```
{--verbose | -v}
```

Displays more detailed information.

isi storagepool list

Displays node pools and tiers in the cluster.

Syntax

```
isi storagepool list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
--format
```

Displays node pools and tiers in the specified format. The following values are valid:

```
table
```

```
json
```

```
csv
```

```
list
```

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi storagepool nodepools create

Creates a manually managed node pool. This command should only be used by experienced OneFS administrators or the with assistance of technical support personnel.

Syntax

```
isi storagepool nodepools create <name>
  [--lnns <lnns>]
  [--verbose]
```

Options

```
<name>
```


Specifies the name for the node pool. Names must begin with a letter or an underscore and may contain only letters, numbers, hyphens, underscores, or periods.

```
{--lanns <lanns> | -n <lanns>}
```

Specifies the nodes in this pool. Nodes can be a comma-separated list or range of LNNs—for example, `1,4,10,12,14,15` or `1-6`.

```
{--verbose | -v}
```

Displays more detailed information.

isi storagepool nodepools delete

Deletes a node pool and autoprovisions the affected nodes into the appropriate node pool. This command is used only for manually managed node pools and should be executed by experienced OneFS administrators or with direction from technical support personnel.

Syntax

```
isi storagepool nodepools delete <name>
  [--force]
  [--verbose]
```

Options

<name>

Specifies the name of the node pool to be deleted.

```
{--force | -f}
```

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

```
{--verbose | -v}
```

Displays more detailed information.

isi storagepool nodepools list

Displays a list of node pools.

Syntax

```
isi storagepool nodepools list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

```
{--limit | -l} <integer>
```

Specifies the number of node pools to display.

```
--format
```

Displays tiers in the specified format. The following values are valid:

```
table
```

```
json
```

```
csv
```

```
list
{--no-header | -a}
    Displays table and CSV output without headers.
{--no-footer | -z}
    Displays table output without footers.
{--verbose | -v}
    Displays more detailed information.
```

isi storagepool nodepools modify

Modifies a node pool.

Syntax

```
isi storagepool nodepools modify <name>
[--protection-policy <string>]
[--lnns <integer>]
[--clear-lnns]
[--add-lnns <integer>]
[--remove-lnns <integer>]
[--tier <string>]
[--clear-tier]
[--l3 {yes | no}]
[--set-name <string>]
```

Options

<string>

Name of the node pool to be modified.

`--protection-policy <string>`

Requested protection for the node pool. Possible protection policy values are:

- +1n
- +2d:1n
- +2n
- +3d:1n
- +3d:1n1d
- +3n
- +4d:1n
- +4d:2n
- +4n
- Mirror values: 2x, 3x, 4x, 5x, 6x, 7x, 8x

OneFS calculates the optimal protection policy (referred to as suggested protection). If the value you set is lower than the suggested protection, OneFS displays an alert.

`--lnns <integer>`

Nodes for the manually managed node pool. Specify `--lnns` for each additional node for the manually managed node pool.

`--clear-lnns`

Clear value for nodes for the manually managed node pool.

`--add-lnns <integer>`

Add nodes for the manually managed node pool. Specify `--add-lnns` for each additional node to add.

`--remove-lnns` *<integer>*

Remove nodes for the manually managed node pool. Specify `--remove-lnns` for each additional node to remove.

`--tier` *<string>*

Set parent for the node pool. Node pools can be grouped into a tier to service particular file pools.

`--clear-tier`

Remove the specified node pool from its parent tier.

`--l3` {yes | no}

Use SSDs in the specified node pool as L3 cache. Note that, on Isilon HD400 node pools, L3 cache is on by default and you cannot disable it. If you try to disable L3 cache on an HD400 node pool, OneFS generates the following error message: Disabling L3 not supported for the given node type.

`--set-name` *<string>*

New name for the manually managed node pool.

Examples

The following command specifies that SSDs in a node pool named `hq_datastore` are to be used as L3 cache:

```
isi storagepool nodepools modify hq_datastore --l3 yes
```

The following command adds the node pool `hq_datastore` to an existing tier named `archive-1`:

```
isi storagepool nodepools modify hq_datastore --tier archive-1
```

isi storagepool nodepools view

Displays details for a node pool.

Syntax

```
isi storagepool nodepools view <name>
  [--verbose]
```

Options

<name>

Specifies the name of the storage pool.

{--verbose | -v}

Displays more detailed information.

isi storagepool settings modify

Modifies global SmartPools settings.

Syntax

```
isi storagepool settings modify
[--automatically-manage-protection {all | files_at_default | none}]
[--automatically-manage-io-optimization {all | files_at_default |
none}]
[--protect-directories-one-level-higher {yes | no}]
[--global-namespace-acceleration-enabled {yes | no}]
[--virtual-hot-spare-deny-writes {yes | no}]
[--virtual-hot-spare-hide-spare {yes | no}]
[--virtual-hot-spare-limit-drives <integer>]
[--virtual-hot-spare-limit-percent <integer>]
[--snapshot-disk-pool-policy-id <integer>]
[--spillover-target <string>| --no-spillover | --spillover-anywhere]
[--ssd-l3-cache-default-enabled {yes | no}]
[--verbose]
```

Required Privileges

ISI_PRIV_SMARTPOOLS

Options

--automatically-manage-protection {all | files_at_default | none}

Specifies whether SmartPools manages files' protection settings.

--automatically-manage-io-optimization {all | files_at_default | none}

Specifies whether SmartPools manages I/O optimization settings for files.

--protect-directories-one-level-higher {yes | no}

Protects directories at one level higher.

--global-namespace-acceleration-enabled {yes | no}

Enables or disables global namespace acceleration.

--virtual-hot-spare-deny-writes {yes | no}

Denies new data writes to the virtual hot spare.

--virtual-hot-spare-hide-spare {yes | no}

Reduces the amount of available space for the virtual hot spare.

--virtual-hot-spare-limit-drives <integer>

Specifies the maximum number of virtual drives.

--virtual-hot-spare-limit-percent <integer>

Limits the percentage of node resources that is allocated to virtual hot spare.

--spillover-target <string>

Specifies the target for spillover.

--no-spillover

Globally disables spillover.

--spillover-anywhere

Globally sets spillover to anywhere.

```
--ssd-l3-cache-default-enabled {yes | no}
```

Enables or disables SSDs on new node pools to serve as L3 cache.

```
--verbose
```

Enables verbose messaging.

Examples

The following command specifies that SSDs on newly created node pools are to be used as L3 cache:

```
isi storagepool settings modify --ssd-l3-cache-default-enabled yes
```

The following command specifies that 20 percent of node resources can be used for the virtual hot spare:

```
isi storagepool settings modify --virtual-hot-spare-limit-percent 20
```

isi storagepool settings view

Displays global SmartPools settings.

Syntax

```
isi storagepool settings view
```

Options

There are no options for this command.

Example

The following command displays the global SmartPools settings on your cluster:

```
isi storagepool settings view
```

The system displays output similar to the following example:

```
Automatically Manage Protection: files_at_default
Automatically Manage Io Optimization: files_at_default
Protect Directories One Level Higher: Yes
  Global Namespace Acceleration: disabled
  Virtual Hot Spare Deny Writes: Yes
  Virtual Hot Spare Hide Spare: Yes
  Virtual Hot Spare Limit Drives: 1
  Virtual Hot Spare Limit Percent: 0
  Global Spillover: anywhere
  SSD L3 Cache Default Enabled: No
```

isi storagepool tiers create

Creates a tier.

Syntax

```
isi storagepool tiers create <name>
  [--children <string>]
  [--verbose]
```

Options

<name>

Specifies the name for the storage pool tier. Specify as any string.

`--children <string>`

Specifies a node pool to be added to the tier. For each node pool that you intend to add, include a separate `--children` argument.

`--verbose`

Displays more detailed information.

Example

The following command creates a tier and adds two node pools to the tier:

```
isi storagepool tiers create ARCHIVE_1 --children hq_datastore1
--children hq_datastore2
```

isi storagepool tiers delete

Deletes a tier.

Syntax

```
isi storagepool tiers delete {<name> | --all}
[--verbose]
```

Options

{<name> | --all}

Specifies the tier to delete. The acceptable values are the name of the tier or all.

{--verbose | -v}

Displays more detailed information.

isi storagepool tiers list

Displays a list of tiers.

Syntax

```
isi storagepool tiers list
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

`--format`

Displays tiers in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information.

isi storagepool tiers modify

Renames a tier.

Syntax

```
isi storagepool tiers modify <name>
  [--set-name <string>]
  [--verbose]
```

Options

<name>

Specifies the tier to be renamed.

```
{--set-name | -s} <string>
```

Sets the new name for the tier.

```
{--verbose | -v}
```

Displays more detailed information.

isi storagepool tiers view

Displays details for a tier.

Syntax

```
isi storagepool tiers view <name>
```

Options

<name>

Specifies the name of the tier.

```
{--verbose | -v}
```

Displays more detailed information.

CHAPTER 21

System jobs

This section contains the following topics:

- [System jobs overview](#).....802
- [System jobs library](#).....802
- [Job operation](#)..... 805
- [Job performance impact](#)..... 806
- [Job priorities](#)..... 807
- [Managing system jobs](#)..... 807
- [Managing impact policies](#)..... 811
- [Viewing job reports and statistics](#).....813
- [Job management commands](#)..... 815

System jobs overview

The most critical function of OneFS is maintaining the integrity of data on your Isilon cluster. Other important system maintenance functions include monitoring and optimizing performance, detecting and mitigating drive and node failures, and freeing up available space.

Because maintenance functions use system resources and can take hours to run, OneFS performs them as jobs that run in the background through a service called Job Engine. The time it takes for a job to run can vary significantly depending on a number of factors. These include other system jobs that are running at the same time; other processes that are taking up CPU and I/O cycles while the job is running; the configuration of your cluster; the size of your data set; and how long since the last iteration of the job was run.

Up to three jobs can run simultaneously. To ensure that maintenance jobs do not hinder your productivity or conflict with each other, Job Engine categorizes them, runs them at different priority and impact levels, and can temporarily suspend them (with no loss of progress) to enable higher priority jobs and administrator tasks to proceed.

In the case of a power failure, Job Engine uses a checkpoint system to resume jobs as close as possible to the point at which they were interrupted. The checkpoint system helps Job Engine keep track of job phases and tasks that have already been completed. When the cluster is back up and running, Job Engine restarts the job at the beginning of the phase or task that was in process when the power failure occurred.

As system administrator, through the Job Engine service, you can monitor, schedule, run, terminate, and apply other controls to system maintenance jobs. The Job Engine provides statistics and reporting tools that you can use to determine how long different system jobs take to run in your OneFS environment.

Note

To initiate any Job Engine tasks, you must have the role of SystemAdmin in the OneFS system.

System jobs library

OneFS contains a library of jobs that runs in the background to maintain your Isilon cluster. Some jobs are automatically started by OneFS when particular conditions arise, and some jobs have a default schedule. However, you can run all jobs manually or schedule them according to your workflow.

Job name	Description	Exclusion Set	Impact Policy	Priority	Operation
AutoBalance	Balances free space in a cluster, and is most efficient in clusters that contain only hard disk drives (HDDs). Run as part of MultiScan, or automatically by the system if MultiScan is disabled.	Restripe	Low	4	Auto

Job name	Description	Exclusion Set	Impact Policy	Priority	Operation
AutoBalanceLin	Balances free space in a cluster, and is most efficient in clusters when file system metadata is stored on solid state drives (SSDs). Run as part of MultiScan, or automatically by the system if MultiScan is disabled.	Restripe	Low	4	Auto
AVScan	Performs an antivirus scan on all files.	None	Low	6	Manual
Collect	Reclaims free space that previously could not be freed because the node or drive was unavailable. Run as part of MultiScan, or automatically by the system if MultiScan is disabled.	Mark	Low	4	Auto
Dedupe*	Scans a directory for redundant data blocks and deduplicates all redundant data stored in the directory. Available only if you activate a SmartDedupe license.	None	Low	4	Manual
DedupeAssessment	Scans a directory for redundant data blocks and reports an estimate of the amount of space that could be saved by deduplicating the directory.	None	Low	6	Manual
DomainMark	Associates a path, and the contents of that path, with a domain.	None	Low	5	Manual
FlexProtect	Scans the file system after a device failure to ensure that all files remain protected. FlexProtect is most efficient in clusters that contain only HDDs.	Restripe	Medium	1	Auto

Job name	Description	Exclusion Set	Impact Policy	Priority	Operation
	<p>Note</p> <p>Unlike HDDs and SSDs that are used for storage, when an SSD used for L3 cache fails, the drive state should immediately change to REPLACE without a FlexProtect job running. An SSD drive used for L3 cache contains only cache data that does not have to be protected by FlexProtect. After the drive state changes to REPLACE, you can pull and replace the failed SSD.</p>				
FlexProtectLin	Scans the file system after a node failure to ensure that all files remain protected. Most efficient when file system metadata is stored on SSDs.	Restripe	Medium	1	Auto
FSAnalyze	Gathers information about the file system.	None	Low	1	Scheduled
IntegrityScan	Verifies file system integrity.	Mark	Medium	1	Manual
MediaScan	Locates and clears media-level errors from disks.	Restripe	Low	8	Scheduled
MultiScan	Performs the work of the AutoBalance and Collect jobs simultaneously.	Restripe Mark	Low	4	Auto
PermissionRepair	Corrects file and directory permissions in the <code>/ifs</code> directory.	None	Low	5	Manual
QuotaScan*	Updates quota accounting for domains created on an existing file tree. Available only if you activate a SmartQuotas license.	None	Low	6	Auto

Job name	Description	Exclusion Set	Impact Policy	Priority	Operation
SetProtectPlus	Applies a default file policy across the cluster. Runs only if a SmartPools license is not active.	Restripe	Low	6	Manual
ShadowStoreDelete	Frees space that is associated with a shadow store.	None	Low	2	Scheduled
SmartPools*	Enforces SmartPools file policies. Available only if you activate a SmartPools license.	Restripe	Low	6	Scheduled
SnapRevert	Reverts an entire snapshot back to head.	None	Low	5	Manual
SnapshotDelete	Creates free space associated with deleted snapshots.	None	Medium	2	Auto
TreeDelete	Deletes a specified file path in the /ifs directory.	None	Medium	4	Manual
Upgrade	<p>Upgrades the file system after a software version upgrade.</p> <hr/> <p>Note</p> <p>The Upgrade job should be run only when you are updating your cluster with a major software version. For complete information, see the <i>Isilon OneFS Upgrade Planning and Process Guide</i>.</p> <hr/>	Restripe	Medium	3	Manual
* Available only if you activate an additional license					

Job operation

OneFS includes system maintenance jobs that run to ensure that your Isilon cluster performs at peak health. Through the Job Engine, OneFS runs a subset of these jobs automatically, as needed, to ensure file and data integrity, check for and mitigate drive and node failures, and optimize free space. For other jobs, for example, Dedupe, you can use Job Engine to start them manually or schedule them to run automatically at regular intervals.

The Job Engine runs system maintenance jobs in the background and prevents jobs within the same classification (exclusion set) from running simultaneously. Two exclusion sets are enforced: restripe and mark.

Restripe job types are:

- AutoBalance
- AutoBalanceLin
- FlexProtect
- FlexProtectLin
- MediaScan
- MultiScan
- SetProtectPlus
- SmartPools

Mark job types are:

- Collect
- IntegrityScan
- MultiScan

Note that MultiScan is a member of both the restripe and mark exclusion sets. You cannot change the exclusion set parameter for a job type.

The Job Engine is also sensitive to job priority, and can run up to three jobs, of any priority, simultaneously. Job priority is denoted as 1–10, with 1 being the highest and 10 being the lowest. The system uses job priority when a conflict among running or queued jobs arises. For example, if you manually start a job that has a higher priority than three other jobs that are already running, Job Engine pauses the lowest-priority active job, runs the new job, then restarts the older job at the point at which it was paused. Similarly, if you start a job within the restripe exclusion set, and another restripe job is already running, the system uses priority to determine which job should run (or remain running) and which job should be paused (or remain paused).

Other job parameters determine whether jobs are enabled, their performance impact, and schedule. As system administrator, you can accept the job defaults or adjust these parameters (except for exclusion set) based on your requirements.

When a job starts, the Job Engine distributes job segments—phases and tasks—across the nodes of your cluster. One node acts as job coordinator and continually works with the other nodes to load-balance the work. In this way, no one node is overburdened, and system resources remain available for other administrator and system I/O activities not originated from the Job Engine.

After completing a task, each node reports task status to the job coordinator. The node acting as job coordinator saves this task status information to a checkpoint file. Consequently, in the case of a power outage, or when paused, a job can always be restarted from the point at which it was interrupted. This is important because some jobs can take hours to run and can use considerable system resources.

Job performance impact

The Job Engine service monitors system performance to ensure that maintenance jobs do not significantly interfere with regular cluster I/O activity and other system administration

tasks. Job Engine uses impact policies that you can manage to control when a job can run and the system resources that it consumes.

Job Engine has four default impact policies that you can use but not modify. The default impact policies are:

Impact policy	Allowed to run	Resource consumption
LOW	Any time of day.	Low
MEDIUM	Any time of day.	Medium
HIGH	Any time of day.	High
OFF_HOURS	Outside of business hours. Business hours are defined as 9AM to 5pm, Monday through Friday. OFF_HOURS is paused during business hours.	Low

If you want to specify other than a default impact policy for a job, you can create a custom policy with new settings.

Jobs with a low impact policy have the least impact on available CPU and disk I/O resources. Jobs with a high impact policy have a significantly higher impact. In all cases, however, the Job Engine uses CPU and disk throttling algorithms to ensure that tasks that you initiate manually, and other I/O tasks not related to the Job Engine, receive a higher priority.

Job priorities

Job priorities determine which job takes precedence when more than three jobs of different exclusion sets attempt to run simultaneously. The Job Engine assigns a priority value between 1 and 10 to every job, with 1 being the most important and 10 being the least important.

The maximum number of jobs that can run simultaneously is three. If a fourth job with a higher priority is started, either manually or through a system event, the Job Engine pauses one of the lower-priority jobs that is currently running. The Job Engine places the paused job into a priority queue, and automatically resumes the paused job when one of the other jobs is completed.

If two jobs of the same priority level are scheduled to run simultaneously, and two other higher priority jobs are already running, the job that is placed into the queue first is run first.

Managing system jobs

The Job Engine enables you to control periodic system maintenance tasks that ensure OneFS file system stability and integrity. As maintenance jobs run, the Job Engine constantly monitors and mitigates their impact on the overall performance of the cluster.

As system administrator, you can tailor these jobs to the specific workflow of your Isilon cluster. You can view active jobs and job history, modify job settings, and start, pause, resume, cancel, and update job instances.

Start a job

Although OneFS runs several critical system maintenance jobs automatically when necessary, you can also manually start any job at any time.

The Collect job, used here as an example, reclaims free space that previously could not be freed because the node or drive was unavailable.

Procedure

1. Run the `isi job jobs start` command.

The following command runs the Collect job with a stronger impact policy and a higher priority.

```
isi job jobs start Collect --policy MEDIUM --priority 2
```

Results

When the job starts, a message such as `Started job [7]` appears. In this example, 7 is the job ID number, which you can use to run other commands on the job.

Pause a job

To free up system resources, you can pause a job temporarily.

Before you begin

To pause a job, you need to know the job ID number. If you are unsure of the job ID number, you can use the `isi job jobs list` command to see a list of running jobs.

Procedure

1. Run the `isi job jobs pause` command.

The following command pauses a job with an ID of 7.

```
isi job jobs pause 7
```

If there is only one instance of a job type currently active, you can specify the job type instead of the job ID.

```
isi job jobs pause Collect
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job pause Collect
```

Modify a job

You can change the priority and impact policy of an active, waiting, or paused job.

Before you begin

To modify a job, you need to know the job ID number. If you are unsure of the job ID number, you can use the `isi job jobs list` command to see a list of running jobs.

When you modify a job, only the current instance of the job runs with the updated settings. The next instance of the job returns to the default settings for that job type.

Procedure

1. Run the `isi job jobs modify` command.

The following command updates the priority and impact policy of an active job (job ID number 7).

```
isi job jobs modify 7 --priority 3 --policy medium
```

If there is only one instance of a job type currently active, you can specify the job type instead of the job ID.

```
isi job jobs modify Collect --priority 3 --policy medium
```

Resume a job

You can resume a paused job.

Before you begin

To resume a job, you need to know the job ID number. If you are unsure of the job ID number, you can use the `isi job jobs list` command.

Procedure

1. Run the `isi job jobs resume` command.

The following command resumes a job with the ID number 7.

```
isi job jobs resume 7
```

If there is only one instance of a job type currently active, you can specify the job type instead of the job ID.

```
isi job jobs resume Collect
```

Cancel a job

If you want to free up system resources, or for any reason, you can cancel a running, paused, or waiting job.

Before you begin

To cancel a job, you need to know the job ID number. If you are unsure of the job ID number, you can use the `isi job jobs list` command.

Procedure

1. Run the `isi job jobs cancel` command.

The following command cancels a job with the ID number 7.

```
isi job jobs cancel 7
```

If there is only one instance of a job type currently active, you can specify the job type instead of the job ID.

```
isi job jobs cancel Collect
```

Modify job type settings

You can customize system maintenance jobs for your administrative workflow by modifying the default priority level, impact level, and schedule for a job type.

The job type ID is the job name, for example, `MediaScan`.

Procedure

1. Run the `isi job types modify` command.

The following command modifies the default priority level and impact level for the `MediaScan` job type.

```
isi job types modify mediascan --priority 2 --policy medium
```

When you run this command, the system prompts you to confirm the change. Type **yes** or **no**, and then press ENTER.

2. Establish a regular schedule for a job type.

The following command schedules the `MediaScan` job to run every Saturday morning at 9 AM. The `--force` option overrides the confirmation step.

```
isi job types modify mediascan --schedule 'every Saturday at 09:00' --force
```

3. Remove a regular schedule for a job type.

The following command removes the schedule for a job type that is scheduled.

```
isi job types modify mediascan --clear-schedule --force
```

Results

All subsequent iterations of the `MediaScan` job type run with the new settings. If a `MediaScan` job is in progress, it continues to use the old settings.

View active jobs

You can view information about jobs that are currently running on your Isilon cluster.

You might want to check active jobs if you are noticing slower system response or to see what jobs are active before starting a new job.

Procedure

1. Run the `isi job jobs list` command.

View job history

You can view recent activity for system maintenance jobs.

You might want to check the last time a critical job ran, view all job history within a recent time period, or output job history for a certain time period into a comma-delimited format file.

Procedure

1. Run the `isi job events list` command for a specific job type.

The following command displays the activity of the MultiScan job type.

```
isi job events list --job-type multiscan
```

2. View all jobs within a specific time frame.

The following command displays all jobs that ran since September 16, 2013.

```
isi job events list --begin 2013-09-16
```

3. For reporting purposes, redirect output to a comma-delimited file.

The following command outputs job history for a specific two-week period to a specified path name.

```
isi job events list --begin 2013-09-15 --end 2013-09-16 > /ifs/  
data/report1.txt
```

Time	Message
2013-09-15T12:55:55	MultiScan[4] Phase 1: end lin scan and mark
2013-09-15T12:55:57	MultiScan[4] Phase 2: begin lin repair scan
2013-09-15T12:56:10	MultiScan[4] Phase 2: end lin repair scan
2013-09-16T01:47:12	SetProtectPlus[3] System Cancelled
2013-09-16T07:00:00	SmartPools[5] Waiting

Managing impact policies

For system maintenance jobs that run through the Job Engine service, you can create and assign policies that help control how jobs affect system performance.

As system administrator, you can create, copy, modify, and delete impact policies, and view their settings.

Create an impact policy

The Job Engine includes four default impact policies, which you cannot modify or delete. However, you can create new impact policies.

You can create custom impact policies to define the best times for system maintenance jobs to run and mitigate their impact on system resources.

Procedure

1. Run the `isi job policies create` command.

The following command creates a custom policy defining a specific time frame and impact level. You can apply the custom policy to any job instance to enable the job to run at a higher impact over the weekend.

```
isi job policies create MY_POLICY --impact medium  
--begin 'Saturday 00:00' --end 'Sunday 23:59'
```

2. View a list of available impact policies to see if your custom policy was created successfully.

The following command displays a list of impact policies.

```
isi job policies list
```

The displayed list appears as follows.

```

ID           Description
-----
HIGH        Isilon template: high impact at all times
LOW         Isilon template: high impact at all times
MEDIUM     Isilon template: high impact at all times
OFF-HOURS  Isilon template: Paused M-F 9-5, low impact otherwise
MY_POLICY
-----

```

3. Add a description to the custom policy.

The following command adds a description to the custom policy.

```
isi job policies modify MY_POLICY --description 'Custom
policy: medium impact when run on weekends'
```

View impact policy settings

You can view the settings of any impact policy.

If you intend to modify an impact policy, you can view the current policy settings. In addition, after you have modified an impact policy, you can view the policy settings to ensure that they are correct.

Procedure

1. Run the `isi job policies view` command.

The following command displays the impact policy settings of the custom impact policy `MY_POLICY`.

```
isi job policies view MY_POLICY
```

Modify an impact policy

You can change the description and policy intervals of a custom impact policy.

Before you begin

You cannot modify the default impact policies, `HIGH`, `MEDIUM`, `LOW`, and `OFF_HOURS`. You can only modify policies that you create.

Procedure

1. Run the `isi job policies modify` command to reset current settings to base defaults.

Policy settings are cumulative, so defining a new impact level and time interval adds to any existing impact level and interval already set on the custom policy. The following command resets the policy interval settings to the base defaults: low impact and anytime operation.

```
isi job policies modify MY_POLICY --reset-intervals
```

2. Run the `isi job policies modify` command to establish new impact level and interval settings for the custom policy.

The following command defines the new impact level and interval of a custom policy named `MY_POLICY`.

```
isi job policies modify MY_POLICY --impact high --begin
'Saturday 09:00' --end 'Sunday 11:59'
```

3. Verify that the custom policy has the settings that you intended.

The following command displays the current settings for the custom policy.

```
isi job policies view MY_POLICY
```

Delete an impact policy

You can delete impact policies that you have created.

Before you begin

You cannot delete default impact policies, `HIGH`, `MEDIUM`, `LOW`, and `OFF_HOURS`.

Procedure

1. Run the `isi job policies delete` command.

The following command deletes a custom impact policy named `MY_POLICY`.

```
isi job policies delete MY_POLICY
```

OneFS displays a message asking you to confirm the deletion of your custom policy.

2. Type **yes** and press **ENTER**.

Viewing job reports and statistics

You can generate reports for system jobs and view statistics to better determine the amounts of system resources being used.

Most system jobs controlled by the Job Engine run at a low priority and with a low impact policy, and generally do not have a noticeable impact on cluster performance.

A few jobs, because of the critical functions they perform, run at a higher priority and with a medium impact policy. These jobs include FlexProtect and FlexProtect Lin, FSAnalyze, SnapshotDelete, and TreeDelete.

As a system administrator, if you are concerned about the impact a system job might have on cluster performance, you can view job statistics and reports. These tools enable you to view detailed information about job load, including CPU and memory usage and I/O operations.

View statistics for a job in progress

You can view statistics for a job in progress.

Before you begin

You need to specify the job ID to view statistics for a job in progress. The `isi job jobs list` command displays a list of active jobs, including job IDs.

Procedure

1. Run the `isi job statistics view` command with a specific job ID.

The following command displays statistics for a Collect job with the ID of 857:

```
isi job statistics view --job-id 857
```

The system displays output similar to the following example:

```
Job ID: 857
Phase: 1
Nodes
Node: 1
  PID: 26224
  CPU: 7.96% (0.00% min, 28.96% max, 4.60% avg)

Virtual: 187.23M (187.23M min, 187.23M max, 187.23M avg)
Physical: 19.01M (18.52M min, 19.33M max, 18.96M avg)

Read: 931043 ops, 7.099G
Write: 1610213 ops, 12.269G
Workers: 1 (0.00 STW avg.)
```

View a report for a completed job

After a job finishes, you can view a report about the job.

Before you begin

You need to specify the job ID to view the report for a completed job. The `isi job reports list` command displays a list of all recent jobs, including job IDs.

Procedure

1. Run the `isi job reports view` command with a specific job ID.

The following command displays the report of a Collect job with an ID of 857:

```
isi job reports view 857
```

The system displays output similar to the following example:

```
Collect[857] phase 1 (2014-03-11T11:39:57)
-----
LIN scan
Elapsed time:                6506 seconds
LINs traversed:              433423
Files seen:                  396980
Directories seen:           36439
Errors:                       0
Total blocks:                27357443452 (13678721726 KB)
CPU usage:                   max 28% (dev 1), min 0% (dev
1), avg 4%
Virtual memory size:         max 193300K (dev 1), min
191728K (dev 1), avg 1925
Resident memory size:        max 21304K (dev 1), min 18884K
(dev 2), avg 20294K
Read:                         11637860 ops, 95272875008 bytes
(90859.3M)
Write:                        20717079 ops, 169663891968
bytes (161804.1M)
```

Job management commands

OneFS provides a library of system maintenance jobs to ensure the integrity of your data and the health of your Isilon cluster. You can control these system maintenance jobs through the job management commands.

isi job events list

Lists recent job events.

Syntax

```
isi job events list
  [--job-type <string>]
  [--job-id <integer>]
  [--begin <timestamp>]
  [--end <timestamp>]
  [--state {failed | running | cancelled_user | succeeded |
  paused_user | unknown | paused_priority | cancelled_system |
  paused_policy | paused_system}]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--job-type <string>`

Displays all events of all instances of a specific job type (for example, SmartPools).

`--job-id <integer>`

Displays all events of a specific job instance.

`--begin <timestamp>`

Specifies the beginning of the time period for which job events should be listed. For example: `--begin "2013-09-17T00:00"`. This means that job events beginning at the first moment of September 17, 2013 should be listed.

`--end <timestamp>`

Specifies the end of the time period for job events to be listed. For example, `--end "2013-09-17T23:59"` means that job events right up to the last minute of September 17, 2013 should be listed.

`--state {failed | running | cancelled_user | succeeded | paused_user | unknown | paused_priority | cancelled_system | paused_policy | paused_system}`

Specifies that events of the given state or states should be listed.

`{--limit | -l} <integer>`

Displays no more than the specified number of job events. If no `timestamp` parameters are specified, the most recent job events of the specified number are listed.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information about job events.

Examples

The following command lists all FSAnalyze events that happened in the month of September.

```
isi job events list --job-type fsanalyze --begin "2013-09-01" --end "2013-09-30"
```

The system displays output similar to the following example.

```
Time                Message
-----
2013-09-16T22:00:21 FSAnalyze[7]  Waiting
2013-09-16T22:00:23 FSAnalyze[7]  Running
2013-09-16T22:00:25 FSAnalyze[7]  Phase 1: begin scan
2013-09-16T22:01:45 FSAnalyze[7]  Phase 1: end scan
2013-09-16T22:01:46 FSAnalyze[7]  Phase 2: begin merge
2013-09-16T22:02:30 FSAnalyze[7]  Phase 2: end merge
2013-09-16T22:02:31 FSAnalyze[7]  Succeeded
2013-09-17T22:00:05 FSAnalyze[9]  Waiting
2013-09-17T22:00:08 FSAnalyze[9]  Running
2013-09-17T22:00:11 FSAnalyze[9]  Phase 1: begin scan
2013-09-17T22:01:37 FSAnalyze[9]  Phase 1: end scan
2013-09-17T22:01:38 FSAnalyze[9]  Phase 2: begin merge
2013-09-17T22:02:24 FSAnalyze[9]  Phase 2: end merge
2013-09-17T22:02:26 FSAnalyze[9]  Succeeded
-----
Total: 14
```

The following command lists all the job events that happened on a specific day.

```
isi job events list --begin "2013-09-17T00:00" --end "2013-09-17T23:59"
```

The system displays output similar to the following example:

```
Time                Message
-----
2013-09-17T22:00:04 SmartPools[8]  Waiting
2013-09-17T22:00:05 FSAnalyze[9]  Waiting
2013-09-17T22:00:06 SmartPools[8]  Running
2013-09-17T22:00:07 SmartPools[8]  Phase 1: begin lin policy update
2013-09-17T22:00:08 FSAnalyze[9]  Running
2013-09-17T22:00:11 FSAnalyze[9]  Phase 1: begin scan
2013-09-17T22:01:01 SmartPools[8]  Phase 1: end lin policy update
2013-09-17T22:01:03 SmartPools[8]  Phase 2: begin sin policy update
2013-09-17T22:01:06 SmartPools[8]  Phase 2: end sin policy update
2013-09-17T22:01:09 SmartPools[8]  Succeeded
2013-09-17T22:01:37 FSAnalyze[9]  Phase 1: end scan
2013-09-17T22:01:38 FSAnalyze[9]  Phase 2: begin merge
2013-09-17T22:02:24 FSAnalyze[9]  Phase 2: end merge
2013-09-17T22:02:26 FSAnalyze[9]  Succeeded
-----
Total: 14
```


isi job jobs cancel

Cancels an active job.

Syntax

```
isi job jobs cancel <job>
```

Options

<job>

Specifies the job to cancel. You can specify the job by job ID or job type. Specify a job type only if one instance of that job type is active.

Examples

The following command cancels an active MultiScan job.

```
isi job jobs cancel multiscan
```

The following command cancels an active job with an instance ID of 14.

```
isi job jobs cancel 14
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job cancel 14
```

isi job jobs list

Displays information about active jobs.

Syntax

```
isi job jobs list
  [--state {running | paused_user | paused_priority | paused_policy |
  paused_system}]
  [--limit <integer>]
  [--sort {id | type | state | impact | policy | priority |
  start_time | running_time}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--state {running | paused_user | paused_priority | paused_policy | paused_system}`

Controls which jobs are listed according to status.

`{--limit | -l} <integer>`

Displays no more than the specified number of items. If no other parameters are specified, displays the most recently activated jobs up to the specified number.

`--sort {id | type | state | impact | policy | priority | start_time | running_time}`

Sorts the output by the specified attribute.

```
--descending
```

Sorts the output in descending order of activation time.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

```
{--no-header}
```

Displays table and CSV output without headers.

```
{--no-footer}
```

Displays table output without footers.

```
{--verbose}
```

Displays more detailed information about active jobs.

Examples

The following example lists jobs that have been manually paused.

```
isi job jobs list --state paused_user
```

The system displays output similar to the following example.

ID	Type	State	Impact	Pri	Phase	Running Time
12	Collect	Paused by user	Low	4	1/2	11s
23	SmartPools	Paused by user	Low	6	1/8	40s

Total: 2

The following example outputs a CSV-formatted list of jobs to a file in the `/ifs/data` path.

```
isi job jobs list --format csv > /ifs/data/joblist.csv
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job list --format csv > /ifs/data/joblist.csv
```

isi job jobs modify

Changes the priority level or impact policy of a queued, running, or paused job.

Syntax

```
isi job jobs modify <job>
  [--priority <integer>]
  [--policy <string>]
```

Options

<job>

Specifies the job ID or job type to modify. If you specify job type (for example, FlexProtect), only one instance of that type can be active.

```
{--priority | -p} <integer>
```

Sets the priority level for the specified job.

```
{--policy|-o} <string>
```

Sets the impact policy for the specified job.

Examples

The following command changes the impact policy of an active MultiScan job. This command example, which specifies the job type, works only when a single instance of MultiScan is active.

```
isi job jobs modify multiscan --policy high
```

If more than one instance of a job type is active, you can specify the job ID number instead of job type. The following command changes the priority of an active job with an ID of 7.

```
isi job jobs modify 7 --priority 2
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job modify 7 --priority 2
```

isi job jobs pause

Pauses an active job.

Syntax

```
isi job jobs pause <job>
```

Options

<job>

Specifies the job to pause. You can specify the job by job type or job ID. If you use job type, only one instance of the job type can be active.

Examples

The following command pauses an active AutoBalance job.

```
isi job jobs pause autobalance
```

The following command pauses an active job with an ID of 18.

```
isi job jobs pause 18
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job pause 18
```

To resume a paused job, use the `isi job resume` command.

isi job jobs resume

Resumes a paused job.

You can confirm that a job has resumed by using the `isi job jobs list` command. Actual resumption of the job can take a while, depending on other activity in the Job Engine queue.

Syntax

```
isi job jobs resume <job>
```

Options

<job>

Specifies the job to resume. You can specify the job by job type or job ID. If you use the job type parameter, only one instance of this job type can be in the Job Engine queue.

Examples

The following command resumes a paused AutoBalance job.

```
isi job jobs resume autobalance
```

The following command resumes a paused job with an ID of 16.

```
isi job jobs resume 16
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job resume 16
```

isi job jobs start

Starts a new job.

The `isi job jobs start` command does not control jobs that are already in progress. If an active job is paused, you can use the `isi job jobs resume` command to start it from the point it was paused.

Syntax

```
isi job jobs start <type>
  [--policy <string>]
  [--priority <integer>]
  [--no-dup]
  [--paths <path>]
  [--delete]
  [--root <path>]
  [--dm-type {snaprevert | synciq}]
  [--mapping-type {clone | sid | unix | native}]
  [--mode {clone | inherit | convert}]
  [--template <path>]
  [--zone <string>]
  [--snapid <integer>]
  [--verbose]
```

Options

`<type>`

Specifies the type of job to add to the job queue (for example, MediaScan).

`{--priority} <integer>`

Sets the priority level for the specified job, with 1 being the highest priority and 10 being the lowest.

`{--policy} <string>`

Sets the impact policy for the specified job.

`{--no-dup}`

Disallows duplicate jobs. If an instance of the specified job is already in the queue, the new job does not start.

`--paths <path>`

Specifies the path of the job, which must be within `/ifs`. This option is valid only for the TreeDelete and PermissionRepair jobs.

`--delete`

Valid for the DomainMark job only. Deletes the domain mark.

`--root <path>`

Valid for the DomainMark job only. Specifies the root path location for the DomainMark job.

`--dm-type {snaprevert | synciq}`

Valid for the DomainMark job only. Specifies the domain type for the DomainMark job.

`--mapping-type {global | sid | unix | native}`

Valid for the PermissionRepair job only, and is only used with the `--mode convert` option. Specifies the type for PermissionRepair.

`--mode {clone | inherit | convert}`

Valid for the PermissionRepair job only. Specifies the mode for PermissionRepair.

`--template <path>`

Valid for the PermissionRepair job only. Specifies the pathname of a template file to use as a model for the PermissionRepair job. Must be within the `/ifs` path.

`--zone <string>`

Valid for the PermissionRepair job only. Specifies the access zone for PermissionRepair.

`--snapid <integer>`

Valid for the SnapRevert job only. Specifies a snapshot ID for the SnapRevert job.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command starts an AutoBalance job.

```
isi job jobs start autobalance
```

The following command starts a MultiScan job with a priority of 8 and a high impact policy.

```
isi job jobs start multiscan --priority 8 --policy high
```

The following command starts a TreeDelete job with a priority of 10 and a low impact policy that deletes the /ifs/data/old directory.

```
isi job jobs start treedelete --path /ifs/data/old --priority 10 --policy low
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job start autobalance
```

isi job jobs view

Displays information about a running or queued job, including the state, impact policy, priority, and schedule.

Syntax

```
isi job jobs view <job>
```

Options

<job>

Specifies the job to view. You can specify the job by job type or job ID. If you specify a job type, only one instance of this job can be active.

Examples

The following command displays information about an AutoBalance job with a job ID of 15.

```
isi job jobs view 15
```

The system displays information similar to the following example.

```

      ID: 15
      Type: AutoBalance
      State: Paused by user
      Impact: Low
      Policy: LOW
      Pri: 4
      Phase: 1/5
      Start Time: 2013-09-19T09:08:28
      Running Time: 24s
      Participants: 1, 2, 3
      Progress: Drives: 6 done, 0 in progress; last updated 3:0;
      Processed 4624 LINS and 918669 KB; 0 ECCs and 0 errors
      Waiting on job ID: -
      Description:
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job view 15
```

isi job policies create

Creates a custom job impact policy.

By default, the new impact policy is assigned a low impact level. You can specify multiple time periods (intervals) during which the job can run at higher impact levels or be paused.

Syntax

```
isi job policies create <name>
  [--description <string>]
  [--impact {Low | Medium | High | Paused }]
  [--begin <interval_time>]
  [--end <interval_time>]
```

Options

<name>

Specifies a name for the new impact policy. The following names are reserved and cannot be used: LOW, MEDIUM, HIGH, and OFF_HOURS.

--description <string>

Describes the job policy.

--impact {Low | Medium | High | Paused}

Specifies an impact level for the policy: Low, Medium, High, or Paused. You can specify an --impact parameter for each impact interval that you define.

--begin <interval_time>

Specifies the beginning time, on a 24-hour clock, of the period during which a job can run. For example: --begin "Friday 20:00".

--end <interval_time>

Specifies the ending time, on a 24-hour clock, of the period during which a job can run. For example: --end "Sunday 11:59".

Examples

The following command creates a new impact policy named HIGH-WKEND.

```
isi job policies create HIGH-WKEND --impact high --begin "Saturday
00:01" --end "Sunday 23:59"
```

The following command creates a more complex impact policy named HI-MED-WKEND. This policy includes multiple impact levels and time intervals. At the end of the specified intervals, a job running with this policy would automatically return to LOW impact.

```
isi job policies create HI-MED-WKEND --description "High to medium
impact when run on the weekend" --impact high --begin "Friday 20:00"
--end "Monday 03:00" --impact medium --begin "Monday 03:01" --end
"Monday 08:00"
```

isi job policies delete

Deletes a job impact policy.

The following policies are reserved and cannot be deleted: LOW, MEDIUM, HIGH, and OFF_HOURS.

Syntax

```
isi job policies delete <id>
[--force]
```

Options

<id>

Specifies the name of the impact policy to delete. If you are unsure of the name, you can use the `isi job policies list` command.

--force

Forces deletion of the impact policy without the system asking for confirmation.

Examples

The following command deletes a custom impact policy named HIGH-MED.

```
isi job policies delete HIGH-MED
```

When you press ENTER, OneFS displays a confirmation message: Are you sure you want to delete the policy HIGH-MED? (yes/[no]):

Type **yes**, and then press ENTER.

The following command deletes a custom impact policy named HIGH-WKEND without the confirmation message being displayed.

```
isi job policies delete HIGH-WKEND --force
```

isi job policies list

Displays the names and descriptions of job impact policies.

Syntax

```
isi job policies list
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

The following command displays a list of available impact policies.

```
isi job policies list
```

The system displays output similar to the following example:

```

ID           Description
-----
HIGH        Isilon template: high impact at all times
LOW         Isilon template: low impact at all times
MEDIUM     Isilon template: medium impact at all times
OFF_HOURS   Isilon template: paused M-F 9-5, low-impact
            at other times
HI-MED      High to medium to low
HI-WKEND    High impact when run on weekends
MED-WKEND   Medium impact when run on weekends
-----
Total: 7

```

The following command displays more information about available policies.

```
isi job policies list --verbose
```

The system displays verbose output in a list format as shown in the following partial example:

```

                ID: HIGH
Description: Isilon template: high impact at all times
System: True
Impact Intervals
                Impact : High
                  Begin : Sunday 00:00
                  End   : Sunday 00:00
-----
                ID: LOW
                Description: Isilon template: low impact at all times
                System: True
Impact Intervals
                Impact : Low
                  Begin : Sunday 00:00
                  End   : Sunday 00:00
-----

```

isi job policies modify

Change the description, impact levels and time intervals of a custom impact policy.

To confirm that the custom policy reflects your changes, you can use the `isi job policy view` command.

Syntax

```

isi job policy modify <ID>
  [--description<string>]
  [--impact {Low | Medium | High | Paused}]
  [--begin <interval_time>]
  [--end <interval_time>]
  [--reset_intervals]

```

Options`<ID>`

Specifies the name of the policy to modify.

`--description <string>`

Specifies a description for the policy. Replaces an older description if one was in place.

`--impact {Low | Medium | High | Paused}`

Specifies an impact level for the policy: Low, Medium, High, or Paused. Specify an `--impact` parameter for each additional impact interval that you define.

`--begin <interval_time>`

Specifies the beginning time, on a 24-hour clock, of the period during which a job can run. For example: `--begin "Friday 20:00"`.

`--end <interval_time>`

Specifies the ending time, on a 24-hour clock, of the period during which a job can run. For example: `--end "Sunday 11:59"`.

`--reset-intervals`

Clears all job policy intervals and restores the defaults.

Examples

The following command clears the custom intervals from a custom policy named MY_POLICY as the first step to adding new intervals.

```
isi job policies modify MY_POLICY --reset-intervals
```

The following command adds new intervals to a custom policy.

```
isi job policies modify MY_POLICY --impact high --begin "Friday 20:00" --end "Sunday 11:59"
```

isi job policies view

Displays the details for a specific Job Engine job policy.

Syntax

```
isi job policies view
  [<id> <string>]
```

Options`<id> <string>`

Specifies the job policy to display by policy ID.

Examples

The following command displays the details for the default job policy, HIGH.

```
isi job policies view HIGH
```

The system displays the following policy details:

```

                                     ID: HIGH
Description: Isilon template: high impact at all times
                                     System: True
```

```
Impact Intervals
      Impact : High
      Begin  : Sunday 00:00
      End    : Sunday 00:00
```

isi job reports list

Displays information about successful job operations, including date and time, job ID, job type, and job phases that fully completed.

Syntax

```
isi job reports list
  [--job-type <string>]
  [--job-id <integer>]
  [--begin <timestamp>]
  [--end <timestamp>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

`--job-type <string>`

Displays reports for all instances of the specified job type.

`--job-id <integer>`

Displays the report for a job with the specified job ID. If a job has multiple phases, Job Engines displays a report for each phase of the specified job ID.

`{--begin | -b} <interval_time>`

Specifies the beginning of the time period for the job reports list. For example: `--begin "2013-09-19"`.

`{--end | -e} <interval_time>`

Specifies the end of the time period for the job reports list. For example: `--end "2013-09-20"`.

`{--limit | -l} <integer>`

Displays no more than the specified number of reports.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

Examples

The following command displays reports for all MultiScan jobs within a specified time period.

```
isi job reports list --job-type multiscan --begin "2013-9-19" --end "2013-9-20"
```

The system displays output similar to the following example.

```

Time                Job ID  Job Type  Phase
-----
2013-09-19T10:00:08 1      MultiScan 1
2013-09-19T10:00:20 1      MultiScan 2
2013-09-19T10:00:21 1      MultiScan 3
2013-09-19T10:00:34 1      MultiScan 4
2013-09-19T10:02:50 2      MultiScan 1
2013-09-19T10:03:06 2      MultiScan 2
2013-09-19T10:03:09 2      MultiScan 3
2013-09-19T10:03:12 2      MultiScan 4
2013-09-20T10:04:53 4      MultiScan 1
2013-09-20T10:05:11 4      MultiScan 2
2013-09-20T10:05:15 4      MultiScan 3
2013-09-20T10:05:20 4      MultiScan 4
-----
Total: 12

```

isi job reports view

Displays a detailed report for a specific job. Reports can be displayed only for successful jobs or for successful phases of a job.

Syntax

```
isi job reports view <id>
```

Options

<id>

Specifies the job ID for the reports you want to view.

Examples

The following command requests reports for an FSAnalyze job with an ID of 7.

```
isi job reports view 7
```

The system displays output similar to the following example. Note that when a job has more than one phase, a report for each phase is provided.

```

FSAnalyze[7] phase 1 (2013-09-19T22:01:58)
-----
FSA JOB QUERY PHASE
Elapsed time:                83 seconds
LINS traversed:              433
Errors:                      0
CPU usage:                   max 30% (dev 2), min 0% (dev 1), avg 10%
Virtual memory size:         max 111772K (dev 1), min 104444K (dev 2),
avg 109423K
Resident memory size:        max 14348K (dev 1), min 9804K (dev 3),
avg 12706K
Read:                        9 ops, 73728 bytes (0.1M)
Write:                       3035 ops, 24517120 bytes (23.4M)

FSAnalyze[7] phase 2 (2013-09-19T22:02:47)
-----
FSA JOB MERGE PHASE
Elapsed time:                47 seconds
Errors:                      0
CPU usage:                   max 33% (dev 1), min 0% (dev 1), avg 8%
Virtual memory size:         max 113052K (dev 1), min 110748K (dev 2),
avg 111558K

```

```
Resident memory size:      max 16412K (dev 1), min 13424K (dev 3),
                           avg 14268K
Read:                     2 ops, 16384 bytes (0.0M)
Write:                    2157 ops, 16871424 bytes (16.1M)
```

isi job statistics list

Displays a statistical summary of active jobs in the Job Engine queue.

Syntax

```
isi job statistics list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information about active jobs, including node activity, CPU and memory usage, and number of workers (processes) involved.

Examples

The following command requests a statistical summary for active jobs.

```
isi job statistics list
```

The system displays output similar to the following example.

```
JobID Phase CPU Avg. Virt.Mem.Avg. Phys.Mem.Avg. I/O Ops I/O Bytes
-----
16     1     5.91%    102.25M      9.72M      9133    67.22M
-----
Total: 1
```

The following command requests more detailed statistics about active jobs.

```
isi job statistics list --verbose
```

The system displays output similar to the following example. In the example, PID is the process ID and CPU indicates CPU utilization by the job. Also indicated are how many worker threads exist for the job on each node and what the sleep-to-work (STW) ratio is for each thread. The statistics represent how the system throttles the job based on impact policies.

```
Job ID: 16
Phase: 1
```

```

Nodes
Node : 1
  PID : 30977
  CPU : 0.00% (0.00% min, 5.91% max, 2.84% avg)
  Memory
    Virtual : 102.25M (102.12M min, 102.25M max, 102.23M avg)
    Physical : 9.99M (9.93M min, 9.99M max, 9.98M avg)
  I/O
    Read : 5637 ops, 62.23M
    Write : 3601 ops, 23.11M
  Workers : 2 (0.60 STW avg.)
Node : 2
  PID : 27704
  CPU : 0.00% (0.00% min, 5.91% max, 2.18% avg)
  Memory
    Virtual : 102.25M (102.00M min, 102.25M max, 102.22M avg)
    Physical : 9.57M (9.46M min, 9.57M max, 9.56M avg)
  I/O
    Read : 4814 ops, 53.30M
    Write : 1658 ops, 7.94M
  Workers : 2 (0.60 STW avg.)
Node : 3
  PID : 27533
  CPU : 7.96% (1.95% min, 7.96% max, 5.62% avg)
  Memory
    Virtual : 102.25M (102.25M min, 102.25M max, 102.25M avg)
    Physical : 9.57M (9.57M min, 9.61M max, 9.59M avg)
  I/O
    Read : 5967 ops, 65.31M
    Write : 5721 ops, 39.69M
  Workers : 2 (0.60 STW avg.)

```

isi job statistics view

Displays statistics for an active job or jobs on an entire cluster or a specific node.

Syntax

```

isi job statistics view
  [--job-id <integer>]
  [--devid <integer>]
  [--verbose]
  [--format {table | json | csv | list}]

```

Options

`--job-id <integer>`

Displays statistics for a specific job ID.

`--devid <integer>`

Displays statistics for a specific node (device) in the cluster.

`{--verbose | -v}`

Displays more detailed statistics for an active job or jobs.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

Examples

The following command requests statistics for an AutoBalance job with an ID of 6.

```

isi job statistics view --job-id 6

```

The system displays output similar to the following example. In the example, PID is the process ID, and CPU indicates CPU utilization by the job. Also indicated are how many worker threads exist for the job on each node and what the sleep-to-work (STW) ratio is for each thread. The statistics represent how the system throttles the job based on impact policies.

```

Job ID: 6
Phase: 2
Nodes
  Node : 1
    PID : 17006
    CPU : 0.00% (0.00% min, 7.91% max, 4.50% avg)
  Memory
    Virtual : 104.62M (104.37M min, 104.62M max, 104.59M avg)
    Physical : 10.08M (10.01M min, 10.11M max, 10.09M avg)
  I/O
    Read : 4141 ops, 45.33M
    Write : 5035 ops, 35.28M
  Workers : 2 (0.60 STW avg.)
  Node : 2
    PID : 16352
    CPU : 13.96% (1.95% min, 13.96% max, 9.61% avg)
  Memory
    Virtual : 104.62M (104.37M min, 104.62M max, 104.59M avg)
    Physical : 10.01M (9.90M min, 10.01M max, 10.00M avg)
  I/O
    Read : 3925 ops, 43.39M
    Write : 4890 ops, 34.13M
  Workers : 2 (0.60 STW avg.)
  Node : 3
    PID : 15929
    CPU : 0.98% (0.98% min, 12.89% max, 6.82% avg)
  Memory
    Virtual : 104.62M (104.37M min, 104.62M max, 104.57M avg)
    Physical : 9.86M (9.84M min, 9.94M max, 9.92M avg)
  I/O
    Read : 3354 ops, 36.77M
    Write : 772 ops, 2.12M
  Workers : 2 (0.60 STW avg.)

```

isi job types list

Displays a list of job types and default settings.

Syntax

```

isi job types list
  [--all]
  [--sort {id | policy | exclusion_set | priority}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]

```

Options

--all

Displays all job types available in the Job Engine.

--sort {id|policy|exclusion_set|priority}

Sorts the output by the specified parameter.

--descending

In conjunction with `--sort` option, specifies that output be sorted descending order. By default, output is sorted in ascending order.

```
--format {table | json | csv | list}
```

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated values (CSV), or list format.

```
{--no-header | -a}
```

Displays table and CSV output without headers.

```
{--no-footer | -z}
```

Displays table output without footers.

```
{--verbose | -v}
```

Displays more detailed information about a specific job type or all job types.

Examples

The following command provides detailed information about job types.

```
isi job types list --sort id --verbose
```

The system displays output similar to the following example.

```

          ID: AVScan
Description: Perform an antivirus scan on all files.
  Enabled: Yes
   Policy: LOW
  Schedule:
Exclusion Set: None
  Priority: 6
-----
          ID: AutoBalance
Description: Balance free space in a cluster. AutoBalance is most
efficient in clusters that contain only HDDs.
  Enabled: Yes
   Policy: LOW
  Schedule:
Exclusion Set: Restripe
  Priority: 4
-----
          ID: AutoBalanceLin
Description: Balance free space in a cluster. AutoBalanceLin is
most efficient if file system metadata is stored on
SSDs.
  Enabled: Yes
   Policy: LOW
  Schedule:
Exclusion Set: Restripe
  Priority: 4

```

isi job status

Displays a summary of active, completed, and failed jobs.

Syntax

```
isi job status
  [--verbose]
```

Options

```
{--verbose | -v}
```


Displays more detailed job status information, including information about the cluster and nodes.

Examples

The following command provides basic job status.

```
isi job status
```

The system displays output that is similar to the following example.

```
The job engine is running.
No running or queued jobs.
Recent finished jobs:
ID   Type           State           Time
-----
1    MultiScan      System Cancelled 2013-09-24T08:23:44
3    MultiScan      Succeeded       2013-09-24T08:26:37
2    SetProtectPlus Succeeded       2013-09-24T08:27:16
4    FlexProtect    Succeeded       2013-09-24T09:14:27
-----
Total: 4
```

The following command provides more detailed job status information.

```
isi job status --verbose
```

The system displays additional output that includes cluster and node information.

```
The job engine is running.
    Coordinator: 1
    Connected: True
    Disconnected Nodes: -
Down or Read-Only Nodes: False
    Statistics Ready: True
    Cluster Is Degraded: False
    Run Jobs When Degraded: False
No running or queued jobs.
Recent finished jobs:
ID   Type           State           Time
-----
1    MultiScan      System Cancelled 2013-09-24T08:23:44
3    MultiScan      Succeeded       2013-09-24T08:26:37
2    SetProtectPlus Succeeded       2013-09-24T08:27:16
4    FlexProtect    Succeeded       2013-09-24T09:14:27
-----
Total: 4
```

isi job types modify

Modifies the parameters of a specified job type.

You can view the current parameters of any job type by using the `isi job types view` command.

Syntax

```
isi job types modify <id>
    [--enabled <boolean>]
```

```
[--policy <string>]
[--schedule<string>]
[--priority <integer>]
[--clear-schedule]
```

Options

<id>

Specifies the job type to modify.

--enabled <boolean>

Specifies whether the job type is enabled or disabled.

--policy<string>

Sets the policy for the specified job type.

--schedule <string>

Sets a recurring date pattern to run the specified job type.

--priority<integer>

Sets the priority level for the specified job type. Job types have a priority value between 1 and 10, with 1 being the highest priority and 10 being the lowest.

--clear-schedule

Clears any schedule associated with the specified job type.

--force

Forces the modification without a confirmation message.

Examples

The following command adds a recurring schedule to the MultiScan command.

```
isi job types modify multiscan --schedule "Every Friday at 22:00"
```

When you run this command, the system prompts you to confirm the change. Type **yes** or **no**, and then press ENTER.

isi job types view

Displays the parameters of a specific job type, including the description, schedule, policy, priority, and whether the job type is a member of an exclusion set.

Syntax

```
isi job types view <id>
```

Options

<id>

Specifies the job type to view.

Examples

The following command displays the parameters of the job type MultiScan.

```
isi job types view multiscan
```

The system displays output similar to the following example.

```
      ID: MultiScan
Description: Perform the work of the AutoBalance and Collect jobs
simultaneously.
      Enabled: Yes
      Policy: LOW
      Schedule:
Exclusion Set: Restripe, Mark
      Priority: 4
```


CHAPTER 22

Networking

This section contains the following topics:

- [Networking overview](#) 838
- [Internal network overview](#) 838
- [External client network overview](#) 839
- [Managing internal network settings](#) 847
- [Managing external network settings](#) 849
- [Managing external network subnets](#) 850
- [Managing IP address pools](#) 853
- [Managing SmartConnect Settings](#) 857
- [Managing network interface members](#) 860
- [Managing network interface provisioning rules](#) 863
- [Managing routing options](#) 865
- [Networking commands](#) 866

Networking overview

After you determine the topology of your network, you can set up and manage your internal and external networks.

There are two types of networks associated with an EMC Isilon cluster:

Internal

Nodes communicate with each other using a high speed low latency InfiniBand network. You can optionally configure a second InfiniBand network as a failover for redundancy.

External

Clients connect to the cluster through the external network with Ethernet. The Isilon cluster supports standard network communication protocols, including NFS, SMB, HTTP, and FTP. The cluster includes various external Ethernet connections, providing flexibility for a wide variety of network configurations. External network speeds vary by product.

Internal network overview

The EMC Isilon cluster must connect to at least one high-speed, low-latency InfiniBand switch for internal communications and data transfer. The connection to the InfiniBand switch is also referred to as an internal network. The internal network is separate from the external network (Ethernet) by which users access the cluster.

Upon initial configuration of your cluster, OneFS creates a default internal network for the InfiniBand switch. The interface to the default internal network is int-a. An internal network for a second InfiniBand switch can be added for redundancy and failover. Failover allows continuous connectivity during path failures. The interface to the default internal network is int-b, which is referred to as int-b/failover in the web administration interface.

Internal IP address ranges

The number of IP addresses assigned to the internal network determines how many nodes can be joined to the EMC Isilon cluster.

When you initially configure the cluster, you specify one or more IP address ranges for the internal InfiniBand network. This range of addresses is used by the nodes to communicate with each other. It is recommended that you create a range of addresses large enough to accommodate adding additional nodes to your cluster. If the IP address range defined during the initial configuration is too restrictive for the size of the internal network, you can add ranges to the int-a network and int-b network. For certain configuration changes, such as deleting an IP address assigned to a node, the cluster must be restarted.

While all clusters will have, at minimum, one internal InfiniBand network (int-a), to enable a second internal network (int-b) you must assign another IP address range to it. To enable internal network failover, assign an IP address range to the failover network. This range is used to refer to the actual IP addresses in use to provide seamless internal IP address failover.

Internal network failover

You can configure an internal switch as a failover network to provide redundancy for intra-cluster communications.

Enable an internal failover network by connecting the int-a interfaces of each node in the cluster to one switch, connecting the int-b ports on each node to another switch, and then restarting the EMC Isilon cluster.

In addition to the IP address range assigned to the int-a internal network, if you enable failover on a second InfiniBand switch, you must assign an IP address range that points to actual IP addresses used by the cluster. These addresses enable seamless failover in the event that either the int-a or int-b switches fail.

External client network overview

You connect a client computer to the EMC Isilon cluster through the external network. OneFS supports network subnets, IP address pools, and features network provisioning rules to simplify configuration.

Subnets simplify external (front-end) network management and provide flexibility in implementing and maintaining the cluster network. You can create IP address pools within subnets to partition your network interfaces according to workflow or node type. You can configure external network settings through provisioning rules and then those rules are applied to nodes that are added to the cluster.

You must initially configure the default external IP subnet in IPv4 format. After configuration is complete, you can configure additional subnets using IPv4 or IPv6.

IP address pools can be associated with a node or a group of nodes as well as with the NIC ports on the nodes. For example, based on the network traffic that you expect, you might decide to establish one subnet for storage nodes and another subnet for accelerator nodes.

How you set up your external network subnets depends on your network topology. In a basic network topology where all client-node communication occurs through a single gateway, only a single external subnet is required. If clients connect through multiple subnets or internal connections, you must configure multiple external network subnets.

External network settings

A default external network subnet is created during the initial set up of your EMC Isilon cluster. You can make modifications to this subnet, create new subnets, and make additional configuration changes to the external network.

During initial cluster setup, OneFS performs the following actions:

- Creates a default external network subnet called subnet0, with the specified netmask, gateway, and SmartConnect service address.
- Creates a default IP address pool called pool0 with the specified IP address range, the SmartConnect zone name, and the external interface of the first node in the cluster as the only member.
- Creates a default network provisioning rule called rule0, which automatically assigns the first external interface for all newly added nodes to pool0.
- Adds pool0 to subnet0 and configures pool0 to use the virtual IP of subnet0 as its SmartConnect service address.
- Sets the global, outbound DNS settings to the domain name server list and DNS search list, if provided.

Once the initial external network has been established, you can configure the following information about your external network:

- Netmask
- IP address range
- Gateway
- Domain name server list (optional)
- DNS search list (optional)
- SmartConnect zone name (optional)
- SmartConnect service address (optional)

You can make modifications to the external network through the web administration interface and the command-line interface.

IP address pools

You can partition EMC Isilon cluster nodes and external network interfaces into logical IP address pools. IP address pools are also utilized when configuring SmartConnect zones and IP failover support for protocols such as NFS. Multiple pools for a single subnet are available only if you activate a SmartConnect Advanced license.

IP address pools:

- Map available addresses to configured interfaces.
- Belong to external network subnets.
- Partition network interfaces on your cluster into pools.
- Can be assigned to groups in your organization.

The IP address pool of a subnet consists of one or more IP address ranges and a set of cluster interfaces. All IP address ranges in a pool must be unique.

A default IP address pool is configured during the initial cluster setup through the command-line configuration wizard. You can modify the default IP address pool at any time. You can also add, remove, or modify additional IP address pools.

If you add external network subnets to your cluster through the subnet wizard, you must specify the IP address pools that belong to the subnet.

IP address pools are allocated to external network interfaces either dynamically or statically. The static allocation method assigns one IP address per pool interface. The IP addresses remain assigned, regardless of that interface's status, but the method does not guarantee that all IP addresses are assigned. The dynamic allocation method distributes all pool IP addresses, and the IP address can be moved depending on the interface's status and connection policy settings.

IPv6 support

You can configure dual stack support for IPv6.

With dual-stack support in OneFS, you can configure both IPv4 and IPv6 addresses. However, configuring an EMC Isilon cluster to use IPv6 exclusively is not supported. When you set up the cluster, the initial subnet must consist of IPv4 addresses.

The following table describes important distinctions between IPv4 and IPv6.

IPv4	IPv6
32-bit addresses	128-bit addresses

IPv4	IPv6
Subnet mask	Prefix length
Address Resolution Protocol (ARP)	Neighbor Discovery Protocol (NDP)

SmartConnect module

SmartConnect is a module that specifies how the DNS server on the EMC Isilon cluster handles connection requests from clients and the methods used to assign IP addresses to network interfaces.

Settings and policies configured for SmartConnect are applied per IP address pool. You can configure basic and advanced SmartConnect settings.

SmartConnect Basic

SmartConnect Basic is included with OneFS as a standard feature and does not require a license.

SmartConnect Basic supports the following settings:

- Specification of the DNS zone
- Round robin connection balancing method
- Service subnet to answer DNS requests

SmartConnect Basic has the following limitations to IP address pool configuration:

- You may only specify a static IP address allocation policy.
- You cannot specify an IP address failover policy.
- You cannot specify an IP address rebalance policy.
- You may only assign one IP address pool per external network subnet.

SmartConnect Advanced

SmartConnect Advanced extends the settings available from SmartConnect Basic. It requires an active license.

SmartConnect Advanced supports the following settings:

- Round robin, CPU utilization, connection counting, and throughput balancing methods.
- Static and dynamic IP address allocation.

SmartConnect Advanced allows you to specify the following IP address pool configuration options:

- You can define an IP address failover policy for the IP address pool.
- You can define an IP address rebalance policy for the IP address pool.
- SmartConnect Advanced supports multiple IP address pools per external subnet to allow multiple DNS zones within a single subnet.

Connection balancing

The connection balancing policy determines how the DNS server handles client connections to the EMC Isilon cluster.

You can specify one of the following balancing methods:

Round robin

Selects the next available node on a rotating basis. This is the default method. Without a SmartConnect license for advanced settings, this is the only method available for load balancing.

Connection count

Determines the number of open TCP connections on each available node and selects the node with the fewest client connections.

Network throughput

Determines the average throughput on each available node and selects the node with the lowest network interface load.

CPU usage

Determines the average CPU utilization on each available node and selects the node with lightest processor usage.

IP address allocation

The IP address allocation policy ensures that all of the IP addresses in the pool are assigned to an available network interface.

You can specify whether to use static or dynamic allocation.

Static

Assigns one IP address to each network interface added to the IP address pool, but does not guarantee that all IP addresses are assigned.

Once assigned, the network interface keeps the IP address indefinitely, even if the network interface becomes unavailable. To release the IP address, remove the network interface from the pool or remove it from the cluster.

Without a license for SmartConnect Advanced, static is the only method available for IP address allocation.

Dynamic

Assigns IP addresses to each network interface added to the IP address pool until all IP addresses are assigned. This guarantees a response when clients connect to any IP address in the pool.

If a network interface becomes unavailable, its IP addresses are automatically moved to other available network interfaces in the pool as determined by the IP address failover policy.

This method is only available with a license for SmartConnect Advanced.

Allocation recommendations based on file sharing protocols

It is recommended that you select a static allocation method if your clients connect through stateful protocols and a dynamic allocation method with stateless protocols.

The following table displays several common protocols and the recommended allocation method:

File sharing protocol	Recommended allocation method
<ul style="list-style-type: none"> • SMB • NFSv4 • HTTP • FTP 	Static

File sharing protocol	Recommended allocation method
<ul style="list-style-type: none"> • sFTP • FTPS • HDFS • SyncIQ 	
<ul style="list-style-type: none"> • NFSv2 • NFSv3 	Dynamic

IP address failover

The IP address failover policy specifies how to handle the IP addresses of network interfaces that become unavailable

To define an IP address failover policy, you must have a license for SmartConnect Advanced, and the IP address allocation policy must be set to dynamic. Dynamic IP allocation ensures that all of the IP addresses in the pool are assigned to available network interfaces.

When a network interface becomes unavailable, the IP addresses that were assigned to it are redistributed to available network interfaces according to the IP address failover policy. Subsequent client connections are directed to the new network interfaces.

You can select one of the following the connection balancing methods to determine how the IP address failover policy selects which network interface receives a redistributed IP address:

- Round robin
- Connection count
- Network throughput
- CPU usage

IP address rebalancing

The IP address rebalance policy specifies when to redistribute IP addresses if one or more previously unavailable network interfaces becomes available again.

To define an IP address rebalance policy, you must have a license for SmartConnect Advanced, and the IP address allocation policy must be set to dynamic. Dynamic IP addresses allocation ensures that all of the IP addresses in the pool are assigned to available network interfaces.

You can set rebalancing to occur manually or automatically:

Manual

Does not redistribute IP addresses until you manually issue a rebalance command through the command-line interface.

Upon rebalancing, IP addresses will be redistributed according to the connection balancing method specified by the IP address failover policy defined for the IP address pool.

Automatic

Automatically redistributes IP addresses according to the connection balancing method specified by the IP address failover policy defined for the IP address pool.

Automatic rebalance may also be triggered by changes to cluster nodes, network interfaces, or the configuration of the external network.

Note

Rebalancing can disrupt client connections. Ensure the client workflow on the IP address pool is appropriate for automatic rebalancing.

SmartConnect DNS service

The SmartConnect service IP address handles client DNS requests and is configured as a subnet setting.

You must have at least one subnet configured with a SmartConnect service IP address in order to handle client DNS requests.

Do not designate an IP address from a pool as the SmartConnect service IP. The SmartConnect service IP should only answer DNS requests; client connections through the SmartConnect service IP result in unexpected behavior or disconnection.

When configuring IP address pool settings, you can designate any subnet with a service IP address to act as the SmartConnect DNS service for the pool. You can assign a SmartConnect DNS service to multiple pools.

The SmartConnect DNS service handles all incoming DNS requests on behalf of each associated pool's SmartConnect zone and it distributes the requests according to each pool's connection balancing policy.

Any pool that does not specify a SmartConnect DNS service is excluded when answering incoming DNS requests.

Note

SmartConnect requires that you add a new name server (NS) record to the existing authoritative DNS zone that contains the cluster, and you must provide the fully qualified domain name (FQDN) of the SmartConnect zone.

DNS name resolution

You can designate up to three DNS servers and up to six search domains for your external network.

You can configure the DNS server settings during initial cluster configuration with the command-line Configuration wizard. After the initial configuration, you can modify the DNS server settings through the web administration interface or through the `isi networks` command.

NIC aggregation

Network interface card (NIC) aggregation, also known as link aggregation, is optional, and enables you to combine the bandwidth of a node's physical network interface cards into a single logical connection. NIC aggregation provides improved network throughput.

Note

Configuring link aggregation is an advanced function of network switches. Consult your network switch documentation before configuring your EMC Isilon cluster for link aggregation.

NIC aggregation can be configured during the creation of a new external network subnet. Alternatively, you can configure NIC aggregation on the existing IP address pool of a subnet.

- OneFS provides support for the following link aggregation methods:

Link Aggregation Control Protocol (LACP)

This method supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP). Balances outgoing traffic across the interfaces based on hashed protocol header information that includes the source and destination address and the VLAN tag, if available. Also assembles interfaces of the same speed into groups called Link Aggregated Groups (LAGs) and balances traffic across the fastest LAGs. This option is the default mode for new pools.

Legacy Fast EtherChannel (FEC) mode

This method aggregates network interfaces through an older FEC driver recommended for OneFS 6.5 and earlier.

Etherchannel (FEC)

This method provides static balancing on aggregated interfaces through the Cisco Fast EtherChannel (FEC) driver, which is found on older Cisco switches. Capable of load balancing traffic across Fast Ethernet links. Allows multiple physical Fast Ethernet links to combine into one logical channel.

Active / Passive Failover

This method switches to the next active interface when the primary interface becomes unavailable. Manages traffic only through a primary interface. The second interface takes over the work of the first as soon as it detects an interruption in communication.

Round-Robin

This method rotates connections through the nodes in a first-in, first-out sequence, handling all processes without priority. Balances outbound traffic across all active ports in the aggregated link and accepts inbound traffic on any port.

- Some NICs may allow aggregation of ports only on the same network card.
- For LACP and FEC aggregation modes, the switch must support IEEE 802.3ad link aggregation. Since the trunks on the network switch must also be configured, the node must be connected with the correct ports on the switch.

Routing options

By default, outgoing client traffic on the EMC Isilon cluster is destination-based; traffic is routed to a particular gateway based on where the traffic is going. OneFS supports source-

based routing and static routes; these options allow for more granular control of the direction of outgoing client traffic.

Source-based routing

Source-based routing selects which gateway to direct outgoing client traffic through based on the source IP address in each packet header.

In the following example, you enable source-based routing on an Isilon cluster that is connected to SubnetA and SubnetB. Each subnet is configured with a SmartConnect zone and a gateway, also labeled A and B. When a client on SubnetA makes a request to SmartConnect ZoneB, the response originates from ZoneB. This results in a ZoneB address as the source IP in the packet header, and the response is routed through GatewayB. Without source-based routing, the default route is destination-based, so the response is routed through GatewayA.

In another example, a client on SubnetC, which is not connected to the Isilon cluster, makes a request to SmartConnect ZoneA and ZoneB. The response from ZoneA is routed through GatewayA, and the response from ZoneB is routed through GatewayB. In other words, the traffic is split between gateways. Without source-based routing, both responses are routed through the same gateway.

When enabled, source-based routing automatically scans your network configuration to create client traffic rules. If you make modifications to your network configuration, such as changing the IP address of a gateway server, source-based routing adjusts the rules. Source-based routing is applied across the entire cluster and only supports the IPv4 protocol.

Enabling or disabling source-based routing goes into effect immediately. Packets in transit continue on their original courses, and subsequent traffic is routed based on the status change. Transactions composed of multiple packets might be disrupted or delayed if the status of source-based routing changes during transmission.

Source-based routing can conflict with static routes. If a routing conflict occurs, source-based routing rules are prioritized over the static route.

You might enable source-based routing if you have a large network with a complex topology. For example, if your network is a multi-tenant environment with several gateways, traffic is more efficiently distributed with source-based routing.

Static routing

A static route directs outgoing client traffic to a specified gateway based on the IP address the client is connected through.

You configure static routes by IP address pool, and each route applies to all network interfaces that are members of the IP address pool. Static routes only support the IPv4 protocol.

You might configure static routing in order to connect to networks that are unavailable through the default routes or if you have a small network that only requires one or two routes.

Note

If you have upgraded from a version earlier than OneFS 7.0, existing static routes that were added through rc scripts will no longer work and must be re-created by running the `isi networks modify pool` command with the `--add-static-routes` option.

VLANs

Virtual LAN (VLAN) tagging is an optional setting that enables an EMC Isilon cluster to participate in multiple virtual networks.

You can partition a physical network into multiple broadcast domains, or virtual local area networks (VLANs). You can enable a cluster to participate in a VLAN which allows multiple cluster subnet support without multiple network switches; one physical switch enables multiple virtual subnets.

VLAN tagging inserts an ID into packet headers. The switch refers to the ID to identify from which VLAN the packet originated and to which network interface a packet should be sent.

Managing internal network settings

You can modify internal IP address ranges and configure an Infiniband switch for failover.

Add or remove an internal IP address range

You can configure IP address ranges for the int-a, int-b, and failover networks.

Each internal Infiniband switch requires an IP address range. The ranges should have a sufficient number of IP addresses for present operating conditions as well as future expansion and addition of nodes.

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem.

2. Modify the internal IP address ranges by running the `iprange` command.

The following command adds an IP range to the int-a internal network:

```
iprange int-a 192.168.206.10-192.168.206.20
```

The following command deletes an existing IP address range from the int-a internal network:

```
deliprange int-a 192.168.206.15-192.168.206.20
```

3. Run the `commit` command to complete the configuration changes and exit `isi config`.

Modify an internal network netmask

You can modify the subnet mask, or netmask, value for the int-a and int-b internal network interfaces.

If the netmask is too restrictive for the size of the internal network, you must modify the netmask settings. It is recommended that you specify a class C netmask, such as 255.255.255.0, for the internal netmask, that is large enough to accommodate future growth of your Isilon clusters.

It is recommended that the netmask values you specify for int-a and int-b/failover are the same. If you modify the netmask value of one, modify the other.

Note

You must reboot the cluster to apply modifications to the netmask.

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem.

2. Modify the internal network netmask by running the `netmask` command.

The following command changes the `int-a` internal network netmask:

```
netmask int-a 255.255.255.0
```

The system displays output similar to the following example:

```
!! WARNING: The new netmask will not take effect until the nodes
are rebooted.
```

3. Run the `commit` command to complete the configuration changes and exit `isi config`.

Configure and enable internal network failover

You can configure the `int-b` internal interfaces to provide backup in the event of an `int-a` network failure.

Failover configuration involves enabling the `int-b` interface, specifying a valid netmask, and adding IP address ranges for the `int-b` interface and the failover network. By default, the `int-b` interface and failover network are disabled.

Note

You must reboot the EMC Isilon cluster to apply modifications to internal network failover.

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem.

2. Set a netmask for the second interface by running the `netmask` command.

The following command changes the `int-b` internal network netmask:

```
netmask int-b 255.255.255.0
```

The system displays output similar to the following example:

```
!! WARNING: The new netmask will not take effect until the nodes
are rebooted.
```

3. Set an IP address range for the second interface by running the `iprange` command.

The following command adds an IP range to the `int-b` internal network:

```
iprange int-b 192.168.206.21-192.168.206.30
```

4. Set an IP address range for the failover interface by running the `iprange` command.

The following command adds an IP range to the internal failover network:

```
iprange failover 192.168.206.31-192.168.206.40
```

5. Enable a second interface by running the `interface` command.

The following command specifies the interface name as `int-b` and enables it:

```
interface int-b enable
```

6. Run the `commit` command to complete the configuration changes and exit `isi config`.
7. Restart the cluster to apply netmask modifications.

Disable internal network failover

You can disable internal network failover by disabling the `int-b` interface.

Note

You must reboot the cluster to apply modifications to internal network failover.

Procedure

1. Run the `isi config` command.

The command-line prompt changes to indicate that you are in the `isi config` subsystem.

2. Disable the `int-b` interface by running the `interface` command.

The following command specifies the `int-b` interface and disables it:

```
interface int-b disable
```

3. Run the `commit` command to complete the configuration changes and exit `isi config`.
4. Restart the cluster to apply failover modifications.

Managing external network settings

You can configure settings that are applicable to the entire external network.

Configure DNS settings

You can configure the domain name servers (DNS) and DNS search list to resolve host names for the EMC Isilon cluster.

Procedure

1. Specify a list of DNS servers by running the `isi networks --dns-servers` command.

The following command specifies IP addresses for two DNS servers:

```
isi networks --dns-servers 192.168.205.26,192.168.205.27
```

2. Specify a list of DNS search suffixes by running the `isi networks --dns-search` command.

The following command specifies three DNS search suffixes:

```
isi networks --dns-search
storage.company.com,data.company.com,support.company.com
```

Managing external network subnets

You can configure subnets on an external network to manage connections between the EMC Isilon cluster and client computers.

Create a subnet

You can add a subnet to the external network of a cluster.

When you create a subnet, a netmask is required. The first subnet you create must use IPv4 protocol; however, subsequent subnets can use IPv6. Your Isilon cluster must always have at least one IPv4 subnet.

Procedure

1. Run the `isi networks create subnet` command.

The following command creates a subnet named `subnet10` with an associated IPv4 netmask:

```
isi networks create subnet --name subnet10 --netmask 255.255.255.0
```

Modify a subnet

You can modify a subnet on the external network.

Note

Modifying an external network subnet that is in use can disable access to the cluster.

Procedure

1. (Optional) To identify the name of the external subnet you want to modify, run the following command:

```
isi networks list subnets
```

2. Modify the external subnet by running the `isi networks modify subnet` command

The following command changes the name of the subnet:

```
isi networks modify subnet --name subnet10 --new-name subnet20
```

The following command sets MTU to 1500, specifies the gateway address as 198.162.205.10, and sets the gateway priority to 1 on `subnet10`:

```
isi networks modify subnet --name subnet10 --mtu 1500 --gateway
198.162.205.10 --gateway-prio 1
```

Delete a subnet

You can delete an external network subnet that you no longer need.

Note

Deleting an external network subnet also deletes any associated IP address pools. Deleting a subnet that is in use can prevent access to the cluster.

Procedure

1. (Optional) To identify the name of the subnet you want to delete, run the following command:

```
isi networks list subnets
```

2. Delete an external subnet by running the `isi networks delete subnet` command.

The following command deletes subnet10:

```
isi networks delete subnet --name subnet10
```

3. At the prompt, type **yes**.

View subnets

You can view all subnets on the external network.

You can also view subnets by the following criteria:

- Subnet name
- IP address in the subnet range

Procedure

1. Run the `isi networks list subnets` command.

The system displays output similar to the following example:

Name	Subnet	Gateway:Prio	SC Service	Pools
subnet0	192.168.205.0/24	192.168.205.2:1	0.0.0.0	1
subnet10	192.158.0.0/16	192.158.0.2:2	192.158.205.40	1
subnet5	192.178.128.0/22	192.178.128.2:3	0.0.0.0	1

The following command displays subnets with a pool containing 192.168.205.5 in its range:

```
isi networks list subnets --has-addr 192.168.205.5
```

The system displays output similar to the following example:

Name	Subnet	Gateway:Prio	SC Service	Pools
subnet0	192.168.205.0/24	192.168.205.2:1	0.0.0.0	1

Enable or disable VLAN tagging

You can partition the external network into Virtual Local Area Networks or VLANs.

In this network configuration, you can enable VLAN tagging.

VLAN tagging requires a VLAN ID that corresponds to the ID number for the VLAN set on the switch. Valid VLAN IDs are 2 to 4094.

Procedure

1. (Optional) To identify the name of the external subnet you want to modify for VLAN tagging, run the following command:

```
isi networks list subnets
```

2. Enable or disable VLAN tagging on the external subnet by running the `isi networks modify subnet` command.

The following command enables VLAN tagging on subnet10 and sets the required VLAN ID to 256:

```
isi networks modify subnet --name subnet10 --enable-vlan --vlan-id 256
```

The following command disables VLAN tagging on subnet10:

```
isi networks modify subnet --name subnet10 --disable-vlan
```

3. At the prompt, type **yes**.

Add or remove a DSR address

You can specify a Direct Server Return (DSR) address for a subnet if your Isilon cluster contains an external hardware load balancing switch that uses DSR.

A DSR implementation routes all client traffic to the cluster through the switch. The switch determines which node handles traffic for the client and passes traffic to that node.

Procedure

1. (Optional) To identify the name of the external subnet you want to modify for DRS addresses, run the following command:

```
isi networks list subnets
```

2. Add or remove DSR addresses on the subnet by running the `isi networks modify subnet` command.

The following command adds a DSR address to subnet10:

```
isi networks modify subnet --name subnet10 --add-dsr-addr 198.162.205.30
```

The following command removes a DSR address from subnet10:

```
isi networks modify subnet --name subnet10 --remove-dsr-addr 198.162.205.30
```

Managing IP address pools

IP address pools allow you to manage the IP addresses clients use to connect to an EMC Isilon cluster. You can also add network interfaces to IP address pools. Each IP address pool is associated with a subnet.

Create an IP address pool

You can partition the external network interface into groups, or pools, of unique IP address ranges.

Note

If you have not activated a SmartConnect Advanced license, the cluster is allowed one IP address pool per subnet. If you activate a SmartConnect Advanced license, the cluster is allowed unlimited IP address pools per subnet.

When you create an address pool, you must assign it to a subnet.

Procedure

1. Run the `isi networks create pool` command.

The following command creates a pool named `pool5` and assigns it to `subnet10`:

```
isi networks create pool --name subnet10:pool5
```

Modify an IP address pool

You can modify IP address pools to update pool settings.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify, run the following command:

```
isi networks list pools
```

2. Modify the IP address pool by running the `isi networks modify pool` command.

The following command changes the name of the pool:

```
isi networks modify pool --name subnet0:pool5 --new-name pool15
```

Delete an IP address pool

You can delete an IP address pool that you no longer need.

When a pool is deleted, the pool and pool settings are removed from the assigned subnet.

Procedure

1. (Optional) To identify the name of the IP address pool you want to delete, run the following command:

```
isi networks list pools
```

2. Delete an IP address pool by running the `isi networks delete pool` command.

The following command deletes the specified IP address pool:

```
isi networks delete pool --name subnet10:pool5
```

3. At the prompt, type **yes**.

View IP address pools

You can view all IP address pools on the external network.

You can also view pools by the following criteria:

- Pool name
- Assigned subnet
- Network interface name
- Rule name
- IP address

Procedure

1. Run the `isi networks list pools` command.

The system displays output similar to the following example:

Subnet	Pool	SmartConnect Zone	Ranges	Alloc
subnet0	pool2	data.company.com	192.168.205.5-192.1...	Static
subnet10	pool5		192.158.0.5-192.158...	Static
subnet5	pool8	support.company.com	192.178.128.0-192.1...	Static

The following command displays pools that have been assigned the first external network interface of a node:

```
isi networks list pools --iface ext-1
```

The system displays output similar to the following example:

Subnet	Pool	SmartConnect Zone	Ranges	Alloc
subnet0	pool2	data.company.com	192.168.205.5-192.1...	Static
subnet10	pool5		192.158.0.5-192.158...	Static

Add or remove an IP address range

You can configure a range of IP addresses for a pool.

All IP address ranges in a pool must be unique.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for IP address ranges, run the following command:

```
isi networks list pools
```

2. Configure a pool's IP address range by running the `isi networks modify pool` command.

The following command adds an address range to pool5:

```
isi networks modify pool --name subnet10:pool5 --add-ranges
192.168.205.128-192.168.205.256
```

The following command deletes an address range from pool5:

```
isi networks modify pool --name subnet10:pool5 --remove-ranges
192.168.205.128-192.168.205.256
```

Configure IP address allocation

You can specify the method by which IP addresses are allocated to a pool's network interfaces.

Static

Select this IP allocation method to assign IP addresses when network interfaces are added to a IP pool. As network interfaces are added to the pool, this method allocates the next unused IP address from the pool to each new interface. After an IP address is allocated, the network interface keeps the address indefinitely unless the interface is removed from the IP address pool or the IP address itself is removed.

Dynamic

Select this IP allocation method to ensure that all IP addresses in the IP address pool are assigned to network interfaces, which allows clients to connect to any IP addresses in the pool and be guaranteed a response. If a node or a network interface becomes unavailable, the IP addresses are automatically moved to other available network interfaces in the pool.

Note

To configure dynamic IP address allocation, you must activate a SmartConnect Advanced license.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for IP address distribution, run the following command:

```
isi networks list pools
```

2. Specify a distribution method for IP addresses in a pool by running the `isi networks modify pool` command.

The following command specifies static distribution of IP addresses in pool5:

```
isi networks modify pool --name subnet10:pool5 --static
```

Configure an IP rebalance policy

You can configure a rebalance policy for an IP address pool.

Before you begin

You can configure an IP rebalance policy only if you activate a SmartConnect Advanced license and the pool uses dynamic IP address allocation.

The rebalance policy dictates when to distribute IP addresses if a network interfaces that was previously unavailable becomes available. You can set rebalancing to occur manually or automatically.

Manual

Does not redistribute IP addresses until you manually run the `--sc-rebalance` option from `isi networks modify pool`. IP addresses will be redistributed according to the pool's IP failover policy.

Automatic

Automatically redistributes IP addresses according to the pool's IP failover policy. Automatic rebalance is triggered by changes to cluster nodes, network interfaces, or the configuration of the external network.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for a rebalance policy, run the following command:

```
isi networks list pools
```

2. Specify a manual or automatic rebalance policy for a pool by running the `isi networks modify pool` command.

- a. Run the `--manual-failback` option to set a manual policy.

The following command specifies manual rebalancing of IP addresses for pool5:

```
isi networks modify pool --name subnet10:pool5 --manual-failback
```

- b. Run the `--auto-failback` option to set an automatic policy.

The following command specifies automatic rebalancing of IP addresses for pool5:

```
isi networks modify pool --name subnet10:pool5 --auto-failback
```

Configure an IP failover policy

You can specify how to redistribute IP addresses among the remaining network interfaces of an IP address pool when one or more interfaces are unavailable.

Before you begin

You can configure an IP failover policy only if you activate a SmartConnect Advanced license and the pool uses dynamic IP address allocation.

The failover policy ensures that all of the IP addresses in a pool are assigned to an available network interface. When one or more network interfaces become unavailable, the failover policy specifies how to redistribute those IP addresses among the remaining available interfaces. Subsequent client connections are directed to the interfaces that were assigned to those IP addresses.

Note

A failover policy supports the following distribution methods:

- Round robin
- Connection count
- Network throughput

- CPU usage

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for an IP failover policy, run the following command:

```
isi networks list pools
```

2. Specify a failover policy for a pool by running the `isi networks modify pool` command.

The following command specifies IP redistribution by CPU usage upon failback for pool5:

```
isi networks modify pool --name subnet10:pool5 --failover-policy
cpu-usage
```

Managing SmartConnect Settings

You can balance external client connections to your EMC Isilon cluster with SmartConnect.

SmartConnect is an optional module that is available in two modes:

Basic

If you have not activated a SmartConnect Advanced license, SmartConnect operates in Basic mode. In Basic mode, client connection balancing is limited to round robin. Basic mode is limited to static IP address allocation and to one IP address pool per external network subnet. This mode is included with OneFS as a standard feature.

Advanced

If you activate a SmartConnect Advanced license, SmartConnect operates in Advanced mode. Advanced mode enables client connection balancing based on round robin, CPU utilization, connection counting, or network throughput. Advanced mode supports IP failover and allows IP address pools to support multiple DNS zones within a single subnet.

Note

SmartConnect requires that a new name server (NS) record is added to the existing authoritative DNS zone containing the cluster.

Configure a SmartConnect zone

You can specify one SmartConnect zone to handle DNS requests for an IP address pool.

Before you begin

A subnet configured with a service IP address must be assigned to the pool.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify, run the following command:

```
isi networks list pools
```

2. Configure a SmartConnect zone on a pool by running the `isi networks modify pool` command:

The following command specifies a SmartConnect zone name in pool5:

```
isi networks modify pool --name subnet10:pool5 --zone
data.company.com
```

It is recommended that the SmartConnect zone be a fully-qualified domain name (FQDN).

After you finish

SmartConnect requires that you add a new name server (NS) record to the existing authoritative DNS zone that contains the cluster and that you delegate the FQDN of the SmartConnect zone.

Add or remove a SmartConnect zone alias

You can specify one or more zone aliases for an IP address pool.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify, run the following command:

```
isi networks list pools
```

2. Add or remove one or more zone aliases in a pool by running the `isi networks modify pool` command.

The following command specifies two zone aliases in pool5:

```
isi networks modify pool --name subnet10:pool5 --add-zone-aliases
storage.company.com,data.company
```

The following command removes a zone aliases from pool5:

```
isi networks modify pool --name subnet10:pool5 --remove-zone-
aliases data.company
```

Configure a SmartConnect connection balancing policy

You can specify a connection balancing policy for an IP address pool.

The policy dictates how SmartConnect handles client connections to the cluster. You can specify one of the following balancing methods:

- Round robin

Note

Round robin is the only connection policy available without activating a SmartConnect Advanced license.

- Connection count
- Network throughput
- CPU usage

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for a connection balancing policy, run the following command:

```
isi networks list pools
```

2. Specify the connection balancing policy for a pool by running the `isi networks modify pool` command.

The following command specifies a round robin balancing policy for pool5:

```
isi networks modify pool --name subnet10:pool5 --connect-policy round-robin
```

Configure a SmartConnect service IP address

You can specify a SmartConnect service IP address on a subnet.

The SmartConnect service IP address receives all incoming DNS requests for a pool and distributes the requests according to the pool's connection-balancing policy.

Note

You must have at least one subnet configured with a SmartConnect service IP in order to balance DNS requests.

A SmartConnect service IP address can handle DNS requests on behalf of any pool on the Isilon cluster if it is specified as the pool's SmartConnect service subnet.

Procedure

1. (Optional) To identify the name of the external subnet you want to modify for a SmartConnect service IP address, run the following command:

```
isi networks list subnets
```

2. Specify a SmartConnect service IP address on an external subnet by running the `isi networks modify subnet` command.

The following command specifies a SmartConnect service IP address on subnet20:

```
isi networks modify subnet --name subnet20 --sc-service-addr 192.158.205.40
```

After you finish

Assign this subnet to one or more pools in order to handle DNS requests for those pools.

Configure a SmartConnect service subnet

You can specify a subnet with a SmartConnect service IP address for a pool.

Before you begin

At least one subnet must first be configured with a SmartConnect service IP address.

A subnet with a SmartConnect service IP address (or service subnet) can be assigned to any pool. The service subnet answers all DNS requests on behalf of the pool's SmartConnect zone.

Any pool that does not specify a SmartConnect service subnet is excluded when answering incoming DNS requests.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for a SmartConnect service subnet, run the following command:

```
isi networks list pools
```

2. Specify the name of the SmartConnect service subnet that will answer a pool's DNS requests by running the `isi networks modify pool` command.

The following command specifies that subnet20 is responsible for the SmartConnect zone in pool5:

```
isi networks modify pool --name subnet10:pool5 --sc-subnet subnet20
```

Managing network interface members

You can assign nodes and network interfaces to specific IP address pools.

You can also aggregate network interfaces and specify the aggregation method.

Add or remove a network interface

You can configure which network interfaces are assigned to an IP address pool.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for network interfaces, run the following command:

```
isi networks list pools
```

2. Configure a pool's network interfaces by running the `isi networks modify pool` command.

The following command adds the first external network interfaces on nodes 1 through 3 to pool5:

```
isi networks modify pool --name subnet10:pool5 --add-ifaces 1-3:ext-1
```

The following command removes the first network interface on node 3 from pool5:

```
isi networks modify pool --name subnet10:pool5 --remove-ifaces 3:ext-1
```

Configure NIC Aggregation

You can configure IP address pools to use network interface card (NIC) aggregation.

NIC aggregation combines multiple, physical external network interfaces on a node into a single logical interface, but only with interfaces of the same type. NIC aggregation cannot be used with mixed-interface types. You can add an aggregated interface to a pool and specify one of the following aggregation modes:

- LACP
- Round robin
- Failover

- FEC
- Legacy

Note

You must enable NIC aggregation on the cluster before you can enable NIC aggregation on the network switch. If the cluster is configured but the node is not, the cluster can continue to communicate. If the node is configured but the cluster is not, the cluster cannot communicate.

Procedure

1. (Optional) To identify the name of the IP address pool you want to modify for NIC aggregation, run the following command:

```
isi networks list pools
```

2. Configure an aggregated interface and specify a NIC aggregation method for a pool by running the `isi networks modify pool` command.

The following command adds interfaces `ext-1` and `ext-2` on node 1 to `pool5` and aggregates them underneath `ext-agg` using the FEC driver:

```
isi networks modify pool --name subnet10:pool5 --iface 1:ext-agg
--aggregation-mode fec
```

NIC and LNI aggregation options

Network interface card (NIC) and logical network interface (LNI) mapping options can be configured for aggregation.

The following list provides guidelines for interpreting the aggregation options.

- Nodes support multiple network card configurations.
- LNI numbering corresponds to the physical positioning of the NIC ports as found on the back of the node. LNI mappings are numbered from left to right.
- Aggregated LNIs are listed in the order in which they are aggregated at the time they are created.
- NIC names correspond to the network interface name as shown in command-line interface tools such as `ifconfig` and `netstat`.

LNI	NIC	Aggregated LNI	Aggregated NIC	Aggregated NIC (Legacy FEC mode)
ext-1 ext-2	em0 em1	ext-agg = ext-1 + ext-2	lagg0	fec0
ext-1 ext-2 ext-3 ext-4	em2 em3 em0 em1	ext-agg = ext-1 + ext-2 ext-agg-2 = ext-3 + ext-4 ext-agg-3 = ext-3 + ext-4 + ext-1 + ext-2	lagg0 lagg1 lagg2	fec0 fec1 fec2
ext-1 ext-2 10gige-1	em0 em1 cxgb0	ext-agg = ext-1 + ext-2 10gige-agg-1 = 10gige-1 + 10gige-2	lagg0 lagg1	fec0 fec1

LNI	NIC	Aggregated LNI	Aggregated NIC	Aggregated NIC (Legacy FEC mode)
10gige-1	cxgb1			

Add or remove a static route

You can configure static routes to direct outgoing client traffic through a specific gateway.

Before you begin

Source-based routing rules are prioritized over static routes. You can check if source-based routing is enabled on the cluster by running the following command:

```
isi networks
```

Configure a static route on a per-pool basis. Static routes support only the IPv4 protocol. You can only add or remove a static route through the command-line interface.

Procedure

1. (Optional) Identify the name of the IP address pool you want to modify for static routes by running the following command:

```
isi networks list pools
```

2. Configure static routes on a pool by running the `isi networks modify pool` command.

Specify the route in classless inter-domain routing (CIDR) notation format.

The following command adds a static route to `pool5` and assigns the route to all network interfaces that are members of the pool:

```
isi networks modify pool subnet10:pool5 --add-static-routes=192.168.205.128/24-192.168.205.2
```

The following command removes a static route from `pool5`:

```
isi networks modify pool subnet10:pool5 --remove-static-routes=192.168.205.128/24-192.168.205.2
```

View network interfaces

You can view all network interfaces on the external network.

You can also view network interfaces by the following criteria:

- Node number
- Inactive status

Procedure

1. Run the `isi networks list interfaces` command.

The system displays output similar to the following example:

```
-----
Iface      Stat      Membership      Addresses
-----
```

```

1:ext-1 up subnet0:pool2 192.168.205.5
1:int-a up int-a-subnet:int-a-pool 192.168.206.10
2:ext-1 up
2:int-a up int-a-subnet:int-a-pool 192.168.206.11
3:ext-1 up subnet5:pool8 192.178.128.0
3:int-a up int-a-subnet:int-a-pool 192.168.206.12

```

The following command displays interfaces on nodes 1 and 3:

```
isi networks list interfaces --nodes 1,3
```

The system displays output similar to the following example:

Iface	Stat	Membership	Addresses
1:ext-1	up	subnet0:pool2	192.168.205.5
1:int-a	up	int-a-subnet:int-a-pool	192.168.206.10
3:ext-1	up	subnet5:pool8	192.178.128.0
3:int-a	up	int-a-subnet:int-a-pool	192.168.206.12

Managing network interface provisioning rules

You can configure provisioning rules to automate the configuration of external network interfaces.

Provisioning rules specify how new nodes are configured when they are added to an EMC Isilon cluster.

If the new node type matches the type defined in a rule, the new node's interface name is added to the subnet and the IP address pool specified in the rule.

For example, you can create a provisioning rule that configures new Isilon storage nodes, and another rule that configures new accelerator nodes.

OneFS automatically checks for multiple provisioning rules when new rules are added to ensure there are no conflicts.

Create a network interface provisioning rule

You can create a network interface provisioning rule to specify how new nodes are configured when they are added to a cluster.

Procedure

1. Run the `isi networks create rule` command.

The following command creates a rule named `rule3` that assigns the first external network interface on each new Accelerator-I node to `pool5` on `subnet10`:

```
isi networks create rule --name subnet10:pool5:rule3 --iface ext-1
```

Modify a network interface provisioning rule

You can modify network interface provisioning rules to update rule settings.

Procedure

1. (Optional) To identify the name of the network interface provisioning rule you want to modify, run the following command:

```
isi networks list rules
```

2. Modify the network interface provisioning rule by running the `isi networks modify rule` command.

The following command changes the name of rule3 to rule3accelerator:

```
isi networks modify rule --name subnet10:pool5:rule3 --new-name rule3accelerator
```

The following command changes rule3 so that it applies only to new storage node types added to the cluster.

```
isi networks modify rule --name subnet10:pool5:rule3 --storage
```

Delete a network interface provisioning rule

You can delete an external interface provisioning rule that you no longer need.

Procedure

1. (Optional) To identify the name of the external interface provisioning rule you want to delete, run the following command:

```
isi networks list rules
```

2. Delete an external interface provisioning rule by running the `isi networks delete rule` command.

The following command deletes rule3 from pool5:

```
isi networks delete rule --name subnet10:pool5:rule3
```

3. At the prompt, type **yes**.

View network interface provisioning rules

You can view all network interface provisioning rules on the external network.

You can also view provisioning rules by the following criteria:

- Rule name
- Assigned subnet
- Assigned pool
- Network interface name
- Node type

Procedure

1. Run the `isi networks list rules` command:

The system displays output similar to the following example:

Name	Pool	Filter	Iface
rule1accel	subnet0:pool2	Accelerator	ext-1
rule3any	subnet10:pool5	Any	ext-1
rule4accel	subnet10:pool5	Accelerator	ext-1
rule2storage	subnet5:pool8	Storage	ext-1

The following command displays rules in pool5 that apply to new accelerator nodes:

```
isi networks list rules --pool subnet10:pool5 --accelerator
```

The system displays output similar to the following example:

Name	Pool	Filter	Iface
rule4accel	subnet10:pool5	Accelerator	ext-1

Managing routing options

You can control the direction of outgoing client traffic through source-based routing or static route configuration.

Enable or disable source-based routing

You can enable source-based routing to ensure that outgoing client traffic is routed to the gateway of the source IP address in the packet header. If you disable source-based routing, outgoing traffic is destination-based or it follows static routes.

Before you begin

Source-based routing rules are prioritized over static routes. You can check if there are static routes configured in any IP address pools by running the following command:

```
isi networks list pools -v
```

Source-based routing is enabled or disabled on the entire EMC Isilon cluster and supports only the IPv4 protocol.

Enabling and disabling source-based routing is only supported through the command-line interface.

Procedure

1. Enable source-based routing by running the following command:

```
isi networks sbr enable
```

2. Disable source-based routing by running the following command:

```
isi networks sbr disable
```

Add or remove a static route

You can configure static routes to direct outgoing client traffic through a specific gateway.

Before you begin

Source-based routing rules are prioritized over static routes. You can check if source-based routing is enabled on the cluster by running the following command:

```
isi networks
```

Configure a static route on a per-pool basis. Static routes support only the IPv4 protocol. You can only add or remove a static route through the command-line interface.

Procedure

1. (Optional) Identify the name of the IP address pool you want to modify for static routes by running the following command:

```
isi networks list pools
```

2. Configure static routes on a pool by running the `isi networks modify pool` command.

Specify the route in classless inter-domain routing (CIDR) notation format.

The following command adds a static route to pool5 and assigns the route to all network interfaces that are members of the pool:

```
isi networks modify pool subnet10:pool5 --add-static-  
routes=192.168.205.128/24-192.168.205.2
```

The following command removes a static route from pool5:

```
isi networks modify pool subnet10:pool5 --remove-static-  
routes=192.168.205.128/24-192.168.205.2
```

Networking commands

You can view and configure settings for the internal and external networks on an EMC Isilon cluster through the networking commands.

isi networks

Manages external network configuration settings.

Syntax

```
isi networks  
[--dns-servers<ip-address-list>]  
[--add-dns-servers<ip-address-list>]  
[--remove-dns-servers<ip-address-list>]  
[--dns-search<dns-domain-list>]  
[--add-dns-search<dns-domain-list>]  
[--remove-dns-search<dns-domain-list>]  
[--dns-options<dns-option-list>]  
[--add-dns-option<dns-option-list>]  
[--remove-dns-option<dns-option-list>]  
[--tcp-ports<number>]  
[--add-tcp-port<number>]  
[--remove-tcp-port<number>]  
[--server-side-dns-search<boolean>]  
[--sc-rebalance-all]  
[--sc-rebalance-delay<number>]
```

Options

If no options are specified, displays the current settings for each option and lists all subnets configured on the EMC Isilon cluster.

`--dns-servers <ip-address-list>`

Sets a list of DNS IP addresses. Nodes issue DNS requests to these IP addresses. The list cannot contain more than three IP addresses. This option overwrites the current list of DNS IP addresses.

`--add-dns-servers <ip-address-list>`

Adds one or more DNS IP addresses, separated by commas. The list cannot contain more than three IP addresses.

`--remove-dns-servers <ip-address-list>`

Removes one or more DNS IP addresses, separated by commas.

`--dns-search <dns-domain-list>`

Sets the list of DNS search suffixes. Suffixes are appended to domain names that are not fully qualified. The list cannot contain more than six suffixes. This option overwrites the current list of DNS search suffixes.

Note

Do not begin suffixes with a leading dot; leading dots are automatically added.

`--add-dns-search <dns-domain-list>`

Adds one or more DNS search suffixes, separated by commas. The list cannot contain more than six search suffixes.

`--remove-dns-search <dns-domain-list>`

Removes one or more DNS search suffixes, separated by commas.

`--dns-options <dns-option-list>`

Sets the DNS resolver options list. DNS resolver options are described in the `/etc/resolv.conf` file.

Note

Setting DNS resolver options is not recommended. Most clusters do not need DNS resolver options and setting them may change how OneFS performs DNS lookups.

`--add-dns-option <dns-option-list>`

Adds DNS resolver options.

`--remove-dns-option <dns-option-list>`

Removes DNS resolver options.

`--tcp-ports <number>`

Sets one or more recognized client TCP ports. This option overwrites the current list of TCP ports.

`--add-tcp-port <number>`

Adds one or more recognized client TCP ports, separated by commas.

`--remove-tcp-port <number>`

Removes one or more recognized client TCP ports, separated by commas.

`--server-side-dns-search <boolean>`

Specifies whether to append DNS search suffixes to a SmartConnect zone that does not use a fully qualified domain name when attempting to map incoming DNS requests. This option is enabled by default.

`--sc-rebalance-all`

Rebalances all dynamic IP address pools on the EMC Isilon cluster. This option requires an active license for SmartConnect Advanced.

`--sc-rebalance-delay <number>`

Specifies a period of time (in seconds) that should pass after a qualifying event before an automatic rebalance is performed. The default value is 0 seconds.

Examples

Run the `isi networks` command without any specified options to display the current settings for each option and list all subnets configured on the cluster.

The system displays output similar to the following example:

```
Domain Name Server: 10.52.0.1,10.52.0.2
DNS Search List:   company.com,storage.company.com
DNS Resolver Opti... N/A
Server-side DNS S... Enabled
DNS Caching:      Enabled
Client TCP ports: 2049, 445, 20, 21, 80
Rebalance delay:  0

Subnet: subnet0 - Default ext-1 subnet (192.168.205.0/24)
```

The following command sets 10.52.0.1 and 10.52.0.2 as the current DNS IP addresses::

```
isi networks --dns-servers=10.52.0.1,10.52.0.2
```

The following command sets "company.com" and "storage.company.com" as DNS search suffixes:

```
isi networks --dns-search=company.com,storage.company.com
```

The following command sets 2049, 445, 20, 21, and 80 as recognized client TCP ports:

```
isi networks --tcp-port=2049,445,20,21,80
```

isi networks create pool

Creates IP address pools that enable you to partition the network interfaces on your cluster into groups.

Syntax

```
isi networks create pool <name>
  [--ranges <ip-address-range-list>]...
  [--ifaces <node-interface>]...
  [--sc-subnet <string>]
  [--desc <description>]
  [--dynamic]
  [--static]
  [--aggregation-mode <mode>]
  [--add-static-routes <route>]...
  [--remove-static-routes <route>]...
  [--ttl <number>]
  [--auto-unsuspend-delay <integer>]
  [--zone <zone>]
  [--add-zone-aliases <aliases>]...
  [--remove-zone-aliases <aliases>]...
  [--access-zone <zone>]
  [--connect-policy <policy>]
  [--failover-policy <policy>]
  [--manual-failback]
  [--auto-failback]
  [--sc-suspend-node <node>]
  [--sc-resume-node <node>]
  [--verbose]
  [--force]
```

Options

<name>

Specifies the name of the new pool that you want to create. The name includes the name of the subnet and the name of the pool separated by a colon—for example, subnet1:pool0. The pool name must be unique in the subnet. Specify the name in the following format:

```
<subnet>:<pool>
```

`--ranges` *<ip-address-range-list>...*

Specifies one or more IP address ranges for the pool. IP addresses within these ranges are assigned to the network interfaces that are members of the IP address pool.

Specify the IP address range in the following format:

```
<low-ip-address>-<high-ip-address>
```

`--ifaces` *<node-interface>...*

Specifies which network interfaces should be members of the IP address pool. The interface values specified through this option override any previously set values. Specify network interfaces in the following format:

```
<node-number>:<interface>
```

To specify multiple nodes, separate each node ID with a comma. To specify a range of nodes, separate the lower and upper node IDs with a dash. To specify multiple network interfaces, separate each interface name with a comma.

The following example adds interfaces ext-1 and ext-2 on nodes 1, 2, 3 and 5 to the IP address pool:

```
--ifaces 1-3,5:ext-1,ext-2
```

`--sc-subnet` *<subnet>*

Specifies the name of the service subnet that is responsible handling DNS requests for the SmartConnect zone.

`--desc` *<string>*

Specifies a description of the pool.

`--dynamic`

Specifies that all pool IP addresses must be assigned to a network interface at all times. Allows multiple IP addresses to be assigned to an interface. If a network interface becomes unavailable, this option ensures that the assigned IP address are redistributed to another interface.

Note

This option is only available if a SmartConnect Advanced license is active on the cluster.

`--static`

Assigns each network interface in the IP address pool a single, permanent IP address from the pool. Depending on the number of IP addresses available, some IP addresses might go unused. The static option is the default setting.

`--aggregation-mode` *<mode>*

Specifies how outgoing traffic is distributed across aggregated network interfaces. The aggregation mode is applied only if at least one aggregated network interface is a member of the IP address pool.

The following values are valid:

lacp

Supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP). Balances outgoing traffic across the interfaces based on hashed protocol header information that includes the source and destination address and the VLAN tag, if available. Also assembles interfaces of the same speed into groups called Link Aggregated Groups (LAGs) and balances traffic across the fastest LAGs. This option is the default mode for new pools.

roundrobin

Rotates connections through the nodes in a first-in, first-out sequence, handling all processes without priority. Balances outbound traffic across all active ports in the aggregated link and accepts inbound traffic on any port.

failover

Switches to the next active interface when the primary interface becomes unavailable. Manages traffic only through a primary interface. The second interface takes over the work of the first as soon as it detects an interruption in communication.

fec

Provides static balancing on aggregated interfaces through the Cisco Fast EtherChannel (FEC) driver, which is found on older Cisco switches. Capable of load balancing traffic across Fast Ethernet links. Allows multiple physical Fast Ethernet links to combine into one logical channel.

legacy

Aggregates network interfaces through an older FEC driver recommended for OneFS 6.5 and earlier.

`--add-static-routes <route>..`

Designates the specified IP addresses as a static route and specifies the destination gateway. If a client connects through a static route IP address, outgoing client traffic is routed through the specified gateway. Multiple routes can be specified in a comma-separated list.

This command is limited to adding IPv4 routes and is available from the command line only.

Specify the route in the following classless inter-domain routing (CIDR) notation format:

```
<network-address>/<subnet-mask>-<gateway-ip-address>
```

`--remove-static-routes <route>..`

Removes a static route from the pool. Multiple routes can be specified in a comma-separated list.

Specify the route in the following CIDR notation format:

```
<network-address>/<subnet-mask>-<gateway-ip-address>
```

`--ttl <integer>`

Specifies the time to live value for SmartConnect DNS query responses (in seconds). DNS responses are only valid for the time specified. The default value is 0 seconds.

`--auto-unsuspend-delay <integer>`

Specifies the time delay (in seconds) before a node that is automatically unsuspended resumes SmartConnect DNS query responses for the node. During certain cluster operations such as rolling upgrades, general node splits, or node reboots, a node is automatically suspended and then unsuspended by the system. This setting is only available through the command line interface; you can view the current setting by listing the current pools in verbose mode.

`--zone <zone>`

Specifies the SmartConnect zone name for this pool. Pool IP addresses are returned in response to DNS queries to this zone. The `--connect-policy` option determines which pool IP addresses are returned.

`--add-zone-aliases <aliases>...`

Adds specified DNS names to the pool as SmartConnect zone aliases. Multiple aliases can be specified in a comma-separated list.

`--remove-zone-aliases <aliases>...`

Removes SmartConnect zone aliases from the pool as comma separated string of DNS names.

`--access-zone <zone>`

Sets access zone for connections to the pool.

`--connect-policy <connection-policy>`

Specifies how incoming client connections are balanced across IP addresses. The following values are valid:

round-robin

Rotates connections through nodes equally. This value is the default policy.

conn-count

Assigns connections to the node that has least connections.

throughput

Assigns connections to the node with the least throughput.

cpu-usage

Assigns connections to the node with the least CPU usage.

`--failover-policy <failover-policy>`

Specifies how IP addresses that belong to an unavailable interface are rebalanced across the remaining network interfaces.

The following values are valid:

round-robin

Assigns IP addresses across nodes equally. This is the default policy.

conn-count

Assigns IP addresses to the node that has least connections.

throughput

Assigns IP addresses to the node with least throughput.

cpu-usage

Assigns IP addresses to the node with least CPU usage.

`--manual-failback`

Requires that connection rebalancing be performed manually after failback.

To manually rebalance a pool, run the following command:

```
isi networks modify pool --name <subnet>:<pool> --sc-rebalance
```

`--auto-failback`

Causes connections to be rebalanced automatically after failback. This is the default setting.

`sc-suspend-node <node>`

Suspends SmartConnect DNS query responses for a node.

`sc-resume-node <node>`

Resumes SmartConnect DNS query responses for a node.

`{--verbose | -v}`

Displays more detailed information.

`{--force | -f}`

Forces commands without warnings.

Examples

The following command creates a new IP address pool called pool1 under subnet0 that assigns IP addresses 192.168.8.10-192.168.8.15 to ext-1 network on nodes 1, 2, and 3. The SmartConnect zone name of this pool is storage.company.com, but it accepts the alias of storage.company:

```
isi networks create pool subnet0:pool1 --ifaces=1-3:ext-1 --
ranges=192.168.8.10-192.168.8.15 --zone=storage.company.com --add-
zone-aliases=storage.company
```

The following command creates a new IP address pool named pool2 under subnet0 that includes interfaces ext-1 and ext-2 on node 1, and aggregates them underneath ext-agg, alternating equally between them.

```
isi networks create pool subnet0:pool2 --iface=1:ext-agg --
ranges=192.168.8.10-192.168.8.15 --aggregation-mode=roundrobin
```

The following command creates a new IP address pool named pool3 under a subnet named subnet0 and specifies that connection rebalancing must be performed manually:

```
isi networks create pool subnet0:pool3 --ifaces=1,2:ext-1,ext-2 --
ranges=192.168.8.10-192.168.8.15 --manual-failback
```

isi networks create rule

Creates a provisioning rule for automatically configuring network interfaces. With provisioning rules, field-replaced interfaces and interfaces on newly added nodes are automatically added to subnets and pools.

Syntax

```
isi networks create rule <name> <iface>
  [--desc <description>]
  [--any]
  [--storage]
  [--accelerator]
  [--backup-accelerator]
  [--verbose]
```


Options`<name>`

Specifies the name and location of the new provisioning rule. Valid names include the subnet, pool, and a unique rule name, separated by colons. The rule name must be unique throughout the given pool.

Specify in the following format:

```
<subnet>:<pool>:<rule>
```

`<iface>`

Specifies the interface name the rule applies to. To view a list of interfaces on your system, run the `isi networks list interfaces` command.

`--desc <description>`

Specifies an optional description of the rule.

`--any`

Sets the provisioning rule to apply to all nodes. This is the default setting.

`--storage`

Sets the provisioning rule to apply to storage nodes.

`--accelerator`

Sets the provisioning rule to apply to Accelerator nodes.

`--backup-accelerator`

Sets the provisioning rule to apply to Backup Accelerator nodes.

`{--verbose | -v}`

Displays more detailed information.

isi networks create subnet

Creates network subnets. Subnets simplify external network management, and provide flexibility in implementing and maintaining cluster network operations.

Syntax

```
isi networks create subnet --name {--netmask <ip-address> | --
prefixlen <integer>}
  [--dsr-address <ip-address-list>]
  [--desc <description>]
  [--gateway <ip-address>]
  [--gateway-prio <integer>]
  [--mtu <mtu>]
  [--sc-service-addr <ip-address>]
  [--vlan-id <vlan-identifier>]
  [--verbose]
```

Options`--name`

Specifies the name of the subnet. Must be unique throughout the cluster.

`--netmask <ip-address>`

Specifies the netmask for an IPv4 subnet. You must specify either a netmask or an IPv6 subnet prefix length.

`--prefixlen <number>`

Sets the prefix length of an IPv6 subnet. You must specify either a prefix length or an IPv4 netmask.

`--dsr-addr <ip-address-list>`

Sets the Direct Server Return address(es) for the subnet. If an external hardware load balancer is used, this parameter is required.

`--desc <description>`

Sets a description for the subnet.

`{--gateway | -g} <ip-address>`

Specifies the gateway IP address used by the subnet. If unspecified, the default gateway is used.

Note

The IP address must belong to the appropriate gateway. If an incorrect IP address is specified, communication with the cluster might be disabled.

`{--gateway-prio | -p} <number>`

Specifies the gateway priority for the subnet. Valid values are numbers between 1 and the total number of existing subnets. The default priority is the lowest possible number.

`--mtu <mtu>`

Sets the maximum transmission unit (MTU) of the subnet. Common values are 1500 and 9000.

Note

Although OneFS supports both 1500 MTU and 9000 MTU, using a larger frame size for network traffic permits more efficient communication on the external network between clients and cluster nodes. For example, if a subnet is connected through a 10 GbE interface and NIC aggregation is configured for IP address pools in the subnet, it is recommended you set the MTU to 9000. To benefit from using jumbo frames, all devices in the network path must be configured to use jumbo frames.

`--sc-service-addr <ip-address>`

Specifies the IP address on which the SmartConnect module listens for domain name server (DNS) requests on this subnet.

`--vlan-id <vlan-identifier>`

Specifies the VLAN ID for all interfaces in the subnet.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command creates a subnet named `example1` with a netmask of `255.255.255.0`:

```
isi networks create subnet --name example1 --netmask 255.255.255.0
```

The system displays output similar to the following example:

```
Creating subnet 'example1:' OK
Saving: OK
```

isi networks delete pool

Deletes IP address pools.

Note

Deleting all IP address pools may result in connection issues.

Syntax

```
isi networks delete pool --name <subnet>:<pool>
 [--force]
```

Options

`--name <subnet>:<pool>`

Required. Specifies the name of the IP address pool to be delete.

`{--force | -f}`

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

Examples

The following command deletes an IP address pool named pool0 from subnet1:

```
isi networks delete pool subnet1:pool0
```

isi networks delete rule

Deletes provisioning rules.

Syntax

```
isi networks delete rule --name:<pool>:<rule>
 [--force]
```

Options

`{name | -n} <subnet>:<pool>:<rule>`

Required. Specifies the provisioning rule to delete.

`{--force | -f}`

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

Examples

The following command deletes a provisioning rule named rule0 from subnet1:pool2:

```
isi networks delete rule subnet1:pool2:rule0
```

isi networks delete subnet

Deletes a subnet.

Note

Deleting all subnets may result in connection issues.

Syntax

```
isi networks delete subnet --name  
[--force]
```

Options

--name

Required. Specifies the subnet for deletion.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

Examples

The following command deletes a subnet named subnet1:

```
isi networks delete subnet subnet1
```

isi networks dnscache disable

Disables the DNS cache.

Syntax

```
isi networks dnscache disable
```

Options

This command has no options.

isi networks dnscache enable

Enables the DNS cache.

Syntax

```
isi networks dnscache enable
```

Options

This command has no options.

isi networks dnscache flush

Flushes the DNS cache.

Syntax

```
isi networks dnscache flush
```

Options

This command has no options.

isi networks dnscache modify

Modifies the DNS cache.

Syntax

```
isi networks dnscache modify
 [--ttl-max-noerror <integer>]
 [--ttl-min-noerror <integer>]
 [--ttl-max-nxdomain <integer>]
 [--ttl-min-nxdomain <integer>]
 [--ttl-max-other <integer>]
 [--ttl-min-other <integer>]
 [--eager-refresh <integer>]
 [--cache-entry-limit <integer>]
 [--testping-delta <integer>]
```

Options

`--ttl-max-noerror <integer>`

Specifies the upper boundary on ttl for cache hits.

`--ttl-min-noerror <integer>`

Specifies the lower boundary on ttl for cache hits.

`--ttl-max-nxdomain <integer>`

Specifies the upper boundary on ttl for nxdomain.

`--ttl-min-nxdomain <integer>`

Specifies the lower boundary on ttl for nxdomain.

`--ttl-max-other <integer>`

Specifies the upper boundary on ttl for non-nxdomain failures.

`--ttl-min-other <integer>`

Specifies the lower boundary on ttl for non-nxdomain failures.

`--eager-refresh <integer>`

Specifies the lead time to refresh cache entries that are nearing expiration.

`--cache-entry-limit <integer>`

Specifies the entry limit for the DNS cache.

`--testping-delta <integer>`

Specifies the delta for checking the cbind cluster health.

isi networks dnscache statistics

Shows the DNS cache statistics.

Syntax

```
isi networks dnscache statistics
```

Options

There are no options for this command.

isi networks list interfaces

Displays a list of network interfaces within a subnet's IP address pool.

Syntax

```
isi networks list interfaces
  [--verbose]
  [--wide]
  [--show-inactive]
  [--nodes <number>]
```

Options

{--verbose | -v}

Displays more detailed information.

{--wide | -w}

Displays entries without enforcing truncation.

--show-inactive

Includes inactive interfaces in the list.

{--nodes | -n} <integer>

Displays a list of nodes to retrieve interfaces from.

Examples

The following command lists network interfaces, including those that are inactive:

```
isi networks list interfaces --show-inactive
```

The system displays output similar to the following example:

Interface	Status	Membership	Addresses
1:ext-1	up	subnet0:pool0	11.22.3.45
1:ext-2	no carrier		
1:ext-agg	inactive		
2:ext-1	up	subnet0:pool0	11.22.34.56
2:ext-2	no carrier		
2:ext-agg	inactive		

isi networks list pools

Displays available IP address pools. IP address pools enable you to partition your cluster's network interfaces into groups, and then assign ranges of IP addresses to logical or functional groups within your organization.

Syntax

```
isi networks list pools --name <subnet>:<pool>
  [--subnet <string>]
  [--iface <node-interface>]
  [--rule <string>]
  [--has-addr <ip-address>]
  [--verbose]
```

Options

If you run this command without options or with only the `--verbose` option, the system displays a list of all available IP address pools.

`--name <subnet>:<pool>`

Displays only pool names that match the specified string, or specifies a full pool name.

`--subnet <string>`

Displays only pools within a subnet whose name matches the specified string.

`--iface <node-interface>`

Displays only pools containing the specified member interface.

`--rule <string>`

Displays only pools containing a rule name that matches the specified string.

`--has-addr <ip-address>`

Displays only the pool that contains the specified IP address.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command displays a list all available IP address pools:

```
isi networks list pools
```

The system displays output similar to the following example:

Subnet	Pool	SmartConnect Zone	Ranges	Alloc
subnet0	pool0		10.22.136.1-6	Static
subnet1	pool0		10.22.136.1-6	Static
subnet1	pool01		10.22.136.1-6	Static
subnet1	pool10		10.22.136.1-6	Static
subnet0	example1	example.site.com	10.33.150.20-30	Dynamic

The following command displays a list of all pools whose names contain the string 'pool0.'

```
isi networks list pools --name pool0
```

The system displays output similar to the following example:

Subnet	Pool	SmartConnect Zone	Ranges	Alloc
subnet0	pool0		10.22.136.1-6	Static
subnet1	pool0		10.22.136.1-6	Static
subnet1	pool01		10.22.136.1-6	Static

isi networks list rules

Displays provisioning rules.

Syntax

```
isi networks list rules [--name <subnet>:<pool>:<rule>]
  [--subnet <string>]
  [--pool <string>]
  [--iface <node-interface>]
  [--any]
  [--storage]
  [--accelerator]
  [--backup-accelerator]
  [--verbose]
```

Options

If no options are specified, the command displays a list of all provisioning rules.

`--name <subnet>:<pool>:<rule>`

Specifies the name of the rule.

`--pool <subnet>:<pool>`

Name of the pool the provisioning rule applies to.

`--iface <node-interface>`

Names the interface that the provisioning rule applies to.

`--any`

Sets the provisioning rule to apply to any type of node.

`--storage`

Sets the provisioning rule to apply to storage nodes.

`--accelerator`

Sets the provisioning rule to apply to accelerator nodes.

`--backup-accelerator`

Sets the provisioning rule to apply to backup accelerator nodes.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following example displays a list of provisioning rules on a node:

```
isi networks list rules
```


The system displays the list of rules in output similar to the following example:

Name	Pool	Node Type	Interface
rule0	subnet0:pool0	All	ext-1

isi networks list subnets

Displays available subnets. Subnets simplify external network management, and provide flexibility when implementing and maintaining efficient cluster network operations.

Syntax

```
isi networks list subnets --name
  [--has-addr <ip-address>]
  [--verbose]
```

Options

If you run this command without options or with only the `--verbose` option, the system displays a list of all available subnets.

`--name`

Displays only subnets that contain the specified string.

`--has-addr <ip-address>`

Displays only pools containing the specified member interface.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command displays a list of all subnets:

```
isi networks list subnets
```

The system displays output similar to the following example:

Name	Subnet	Gateway:Prio	SC Service	Pools
subnet0	11.22.3.0/24	11.22.0.1:1	1.22.100.10	2
10gbe	11.22.33.0/24	N/A	N/A	1

isi networks modify pool

Modifies IP address pool settings.

Syntax

```
isi networks modify pool <name>
  [--new-name <pool-name>]
  [--sc-rebalance]
  [--ranges <ip-address-range-list>]...
  [--add-ranges <ip-address-range-list>]...
  [--remove-ranges <ip-address-range-list>]...
  [--ifaces <node-interface>]...
  [--add-ifaces <node-interface>]...
  [--remove-ifaces <node-interface>]...
  [--sc-subnet <string>]
```

```

[--desc <string>]
[--dynamic]
[--static]
[--aggregation-mode <mode>]
[--add-static-routes <route>]...
[--remove-static-routes <route>]...
[--ttl <number>]
[--auto-unsuspend-delay <number>]
[--zone <string>]
[--add-zone-aliases <aliases>]...
[--remove-zone-aliases <aliases>]...
[--access-zone <zone>]
[--connect-policy <policy>]
[--failover-policy <policy>]
[--manual-failback]
[--auto-failback]
[--sc-suspend-node <node>]
[--sc-resume-node <node>]
[--verbose]
[--force]

```

Options

You must specify at least one IP address pool setting to modify.

<name>

Specifies the name of the pool to modify. Must be unique throughout the subnet.
Specify the pool name in the following format:

```
<subnet>:<pool>
```

`--new-name <string>`

Required. Specifies a new name for the IP address pool.
Specify the new pool name in the following format:

```
<subnet>:<pool>
```

`--sc-rebalance`

Rebalances IP addresses for the pool.

`--ranges <ip-address-range-list>...`

Specifies one or more IP address ranges for the pool. IP addresses within these ranges are assigned to the network interfaces that are members of the IP address pool.

Note

Specifying new ranges with this option will remove any previously entered ranges from the pool.

Specify the IP address range in the following format:

```
<low-ip-address>-<high-ip-address>
```

`--add-ranges <ip-address-range-list>...`

Adds specified IP address ranges to the pool.
Specify the IP address range in the following format:

```
<low-ip-address>-<high-ip-address>
```

`--remove-ranges <ip-address-range-list>...`

Removes specified IP address ranges from the pool.
Specify the IP address range in the following format:

```
<low-ip-address>-<high-ip-address>
```

`--ifaces <node-interface>`

Specifies which network interfaces should be members of the IP address pool. The interface values specified through this option override any previously set values. Specify network interfaces in the following format:

```
<node-number>:<interface>
```

To specify multiple nodes, separate each node ID with a comma. To specify a range of nodes, separate the lower and upper node IDs with a dash. To specify multiple network interfaces, separate each interface name with a comma.

The following example adds interfaces ext-1 and ext-2 on nodes 1, 2, 3 and 5 to the IP address pool:

```
--ifaces 1-3,5:ext-1,ext-2
```

`--add-ifaces <node-interface>`

Adds a network interface to the pool.
Specify network interfaces in the following format:

```
<node-number>:<interface>
```

`--remove-ifaces <node-interface>`

Removes a network interface from the pool.
Specify network interfaces in the following format:

```
<node-number>:<interface>
```

`--sc-subnet <subnet>`

Specifies the name of the service subnet that is responsible handling DNS requests for the SmartConnect zone.

`--desc <string>`

Specifies a description of the pool.

`--dynamic`

Specifies that all pool IP addresses must be assigned to a network interface at all times. Allows multiple IP addresses to be assigned to an interface. If a network interface becomes unavailable, this option ensures that the assigned IP address are redistributed to another interface.

Note

This option is only available if a SmartConnect Advanced license is active on the cluster.

`--static`

Assigns each network interface in the IP address pool a single, permanent IP address from the pool. Depending on the number of IP addresses available, some IP addresses might go unused. The static option is the default setting.

`--aggregation-mode <mode>`

Specifies how outgoing traffic is distributed across aggregated network interfaces. The aggregation mode is applied only if at least one aggregated network interface is a member of the IP address pool.

The following values are valid:

lacp

Supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP). Balances outgoing traffic across the interfaces based on hashed protocol header information that includes the source and destination address and the VLAN tag, if available. Also assembles interfaces of the same speed into groups called Link Aggregated Groups (LAGs) and balances traffic across the fastest LAGs. This option is the default mode for new pools.

roundrobin

Rotates connections through the nodes in a first-in, first-out sequence, handling all processes without priority. Balances outbound traffic across all active ports in the aggregated link and accepts inbound traffic on any port.

failover

Switches to the next active interface when the primary interface becomes unavailable. Manages traffic only through a primary interface. The second interface takes over the work of the first as soon as it detects an interruption in communication.

fec

Provides static balancing on aggregated interfaces through the Cisco Fast EtherChannel (FEC) driver, which is found on older Cisco switches. Capable of load balancing traffic across Fast Ethernet links. Allows multiple physical Fast Ethernet links to combine into one logical channel.

legacy

Aggregates network interfaces through an older FEC driver recommended for OneFS 6.5 and earlier.

`--add-static-routes <route>..`

Designates the specified IP addresses as a static route and specifies the destination gateway. If a client connects through a static route IP address, outgoing client traffic is routed through the specified gateway. Multiple routes can be specified in a comma-separated list.

This command is limited to adding IPv4 routes and is available from the command line only.

Specify the route in the following classless inter-domain routing (CIDR) notation format:

```
<network-address>/<subnet-mask>-<gateway-ip-address>
```

`--remove-static-routes <route>..`

Removes a static route from the pool. Multiple routes can be specified in a comma-separated list.

Specify the route in the following CIDR notation format:

```
<network-address>/<subnet-mask>-<gateway-ip-address>
```

`--ttl <integer>`

Specifies the time to live value for SmartConnect DNS query responses, in seconds. DNS responses are only valid for the time specified. The default value is 0.

`--auto-unsuspend-delay <integer>`

Specifies the time delay (in seconds) before a node that is automatically unsuspended resumes SmartConnect DNS query responses for the node. During certain cluster operations such as rolling upgrades, general node splits, or node reboots, a node is automatically suspended and then unsuspended by the system. This setting is only available through the command line interface; you can view the current setting by listing the current pools in verbose mode.

`--zone <zone>`

Specifies the SmartConnect zone name for the pool. DNS queries to this zone return pool IP addresses. The `--connect-policy` setting determines which pool IP addresses are returned.

`--add-zone-aliases <aliases>...`

Adds specified DNS names to the pool as SmartConnect zone aliases. Multiple aliases can be specified in a comma-separated list.

`--remove-zone-aliases <aliases>...`

Removes specified DNS names from the pool as SmartConnect zone aliases. Multiple aliases can be specified in a comma-separated list.

`--access-zone <zone>`

Sets the access zone for connections to the pool.

`--connect-policy <connection-policy>`

Specifies how incoming client connections are balanced across IP addresses. The following values are valid:

round-robin

Rotates connections through nodes equally. This value is the default policy.

conn-count

Assigns connections to the node that has least connections.

throughput

Assigns connections to the node with the least throughput.

cpu-usage

Assigns connections to the node with the least CPU usage.

`--failover-policy <failover-policy>`

Specifies how IP addresses that belong to an unavailable interface are rebalanced across the remaining network interfaces.

The following values are valid:

round-robin

Assigns IP addresses across nodes equally. This is the default policy.

conn-count

Assigns IP addresses to the node that has least connections.

throughput

Assigns IP addresses to the node with least throughput.

cpu-usage

Assigns IP addresses to the node with least CPU usage.

`--manual-failback`

Requires that connection rebalancing be performed manually after failback.

You can manually rebalance a pool by running the following command:

```
isi networks modify pool --name <subnet>:<pool> --sc-rebalance
```

`--auto-failback`

Rebalances connections automatically after failback. This option is the default setting.

`--sc-suspend-node <node>`

Suspends SmartConnect DNS query responses for the specified node. While suspended, SmartConnect does not return IP addresses for this node, but allows active clients to remain connected.

`--sc-resume-node <node>`

Resumes SmartConnect DNS query responses for a node.

`{--verbose | -v}`

Displays the results of running this command.

`{--force | -f}`

Suppresses warning messages about pool modification.

Examples

The following command removes node 6 from participating in the SmartConnect profile for subnet0:pool0:

```
isi networks modify pool subnet0:pool0 --sc-suspend-node=6
```

You can confirm that a node has been suspended by running the following command:

```
isi networks list pools --verbose
```

The following command removes the IP address 192.168.9.84 from pool subnet0:pool01:

```
isi networks modify pool subnet0:pool01 --remove-ranges=192.168.9.84
```

The following command causes subnet0:pool1 to rotate equally through aggregated interfaces:

```
isi networks modify pool subnet0:pool1 --aggregation-mode=roundrobin
```

isi networks modify rule

Modifies network provisioning rule settings.

Syntax

```
isi networks modify rule --name <subnet>:<pool>:<rule>
  [--new-name <rule>]
  [--pool <subnet>:<pool>]
  [--iface <node-interface>]
  [--desc <description>]
  [--any]
  [--storage]
  [--accelerator]
  [--backup-accelerator]
  [--verbose]
```

Options

You must specify at least one network provisioning rule setting to modify.

`--name <subnet>:<pool>`

Required. Specifies the name and location of the rule being modified. Must be unique throughout the cluster.

`--new-name`

Specifies a new name for the rule. This name must be unique throughout the pool.

Note

This option does not include the name of the subnet or the pool.

`--pool <subnet>:<pool>`

Changes the pool to which the rule belongs. You must specify both the name of the subnet and the name of the pool.

`--iface <node-interface>`

Specifies the node interface to which the rule applies.

`--desc <description>`

Specifies an optional description of the rule.

`--any`

Applies this rule to all nodes. This is the default setting.

`--storage`

Sets the provisioning rule to apply to storage nodes.

`--accelerator`

Sets the provisioning rule to apply to Accelerator nodes.

`--backup-accelerator`

Sets the provisioning rule to apply to Backup Accelerator nodes.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following example applies rule3 on subnet0:pool0 only to storage nodes:

```
isi networks modify rule subnet0:pool0:rule3 --storage
```

The system displays detailed output similar to the following example:

```
Modifying rule 'subnet0:pool0:rule3':
Saving: OK
```

isi networks modify subnet

Modifies network subnet settings.

Syntax

```
isi networks modify subnet --name <string>
  [--new-name <subnet>]
  [--netmask <ip-address>]
```

```

[--prefixlen <number>]
[--disable-vlan]
[--enable-vlan]
[--dsr-address <ip-address-list>]
[--add-dsr-addr <ip-address-list>]
[--remove-dsr-addr <ip-address-list>]
[--desc <description>]
[--gateway <ip-address>]
[--gateway-prio <number>]
[--mtu <mtu>]
[--sc-service-addr <ip-address>]
[--vlan-id <vlan-identifier>]
[--verbose]
[--force]

```

Options

You must specify at least one network subnet setting to modify.

`--name <string>`

Required. Specifies the name of the subnet to modify.

`--new-name <subnet>`

Specifies a new name for the subnet. Must be unique throughout the cluster.

`--netmask <ip-address>`

Sets the netmask of the subnet.

`--prefixlen <number>`

Sets the prefix length of an IPv6 subnet.

`--enable-vlan`

Enables all VLAN tagging on the subnet.

`--disable-vlan`

Disables all VLAN tagging on the subnet.

`--dsr-addr <ip-address-list>`

Specifies the Direct Server Return addresses for the subnet.

`--add-dsr-addr <ip-address-list>`

Adds one or more Direct Server Return addresses to the subnet.

`--remove-dsr-addr <ip-address-list>`

Removes one or more Direct Server Return addresses from the subnet.

`--desc <description>`

Specifies an optional description for this subnet.

`{--gateway | -g} <ip-address>`

Specifies the gateway IP address used by the subnet. If not specified, the default gateway is used.

Note

IP address must belong to the appropriate gateway. If an incorrect IP address is specified, communication with the cluster might be disabled.

`{--gateway-prio | -p} <number>`

Specifies the gateway priority for the subnet. Valid values are numbers between 1 and the total number of existing subnets. The default priority is the lowest possible.

`--mtu <mtu>`

Specifies the maximum transmission unit (MTU) of the subnet (in bytes). Valid values are 1500 or 9000.

`--sc-service-addr <ip address>`

Specifies the address on which SmartConnect listens for DNS requests on this subnet.

`--vlan-id <vlan-identifier>`

Specifies the VLAN ID or tag for all interfaces on this subnet.

`{--verbose | -v}`

Displays more detailed information.

`{--force | -f }`

Suppresses any prompts or warnings messages that would otherwise appear before or during the subnet modification operation.

Examples

The following command changes the name of subnetOld to subnetNew:

```
isi networks modify subnet subnetOld --new-name=subnetNew
```

The system displays output similar to the following example:

```
Creating subnet 'example1': OK
Saving: OK
```

isi networks sbr enable

Enables source-based routing on the EMC Isilon cluster.

Syntax

```
isi networks sbr enable
```

Options

There are no options for this command.

isi networks sbr disable

Disables source-based routing on the EMC Isilon cluster.

Syntax

```
isi networks sbr disable
```

Options

There are no options for this command.

CHAPTER 23

Hadoop

This section contains the following topics:

- [Hadoop overview](#)..... 892
- [Hadoop architecture](#).....892
- [How Hadoop is implemented on OneFS](#).....893
- [Hadoop distributions supported by OneFS](#)..... 893
- [WebHDFS](#)..... 894
- [Secure impersonation](#)..... 894
- [Ambari agent](#).....895
- [Virtual HDFS racks](#).....895
- [HDFS implementation considerations](#).....896
- [Managing the HDFS service](#)..... 897
- [Managing HDFS access zone settings](#).....900
- [Configuring secure impersonation](#).....903
- [Managing virtual HDFS racks](#).....905
- [HDFS commands](#).....908

Hadoop overview

Hadoop is an open-source platform that runs analytics on large sets of data across a distributed file system.

In a Hadoop implementation on an EMC Isilon cluster, OneFS acts as the distributed file system and HDFS is supported as a native protocol. Clients from a Hadoop cluster connect to the Isilon cluster through the HDFS protocol to manage and process data.

Hadoop support on the cluster requires you to activate an HDFS license. To obtain a license, contact your EMC Isilon sales representative.

Hadoop architecture

Hadoop consists of a compute layer and a storage layer.

In a typical Hadoop implementation, both layers exist on the same cluster.

Hadoop compute layer

MapReduce 2.0 (YARN) is the task processing engine of the Hadoop compute layer.

MapReduce runs a variety of jobs (also known as applications), or queries, on large sets of data and pulls information out. MapReduce relies on the following key components:

ResourceManager

Global authority that allocates resources (such as CPU, memory, disk, network) to NodeManagers, and schedules jobs based on their resource requirements.

ApplicationMaster

Per-job component that negotiates job resources from the ResourceManager and tracks job status.

NodeManager

Per-node component that launches jobs and monitors job resource consumption.

HDFS storage layer

The storage layer is known as the Hadoop distributed file system (HDFS).

The storage layer contains the data accessed and processed by the compute layer. HDFS relies on two key components:

DataNode

Node that stores data and serves read and write requests as directed by the NameNode component.

NameNode

Node that stores in-memory maps of every file, including information about which DataNode the file resides on and the location of the file on the DataNode.

A typical Hadoop implementation contains one NameNode that acts as a master and routes requests for data access to the proper DataNode.

How Hadoop is implemented on OneFS

In a Hadoop implementation on the EMC Isilon cluster, data is stored on OneFS. HDFS is supported as a protocol, which is used by Hadoop compute clients to access data.

A Hadoop implementation with OneFS differs from a typical Hadoop implementation in the following ways:

- The compute and storage layers are on separate clusters instead of the same cluster.
- Instead of storing data within a Hadoop distributed file system, the storage layer functionality is fulfilled by OneFS on an EMC Isilon cluster. Nodes on the Isilon cluster function as both a NameNode and a DataNode.
- The compute layer is established on a Hadoop compute cluster that is separate from the Isilon cluster. MapReduce and its components are installed on the Hadoop compute cluster only.
- Rather than a storage layer, HDFS is implemented on OneFS as a native, lightweight protocol layer between the Isilon cluster and the Hadoop compute cluster. Clients from the Hadoop compute cluster authenticate over HDFS to access data on the Isilon cluster.
- In addition to HDFS, clients from the Hadoop compute cluster can connect to the Isilon cluster over any protocol that OneFS supports such as NFS, SMB, FTP, and HTTP.
- Hadoop compute clients can connect to any node on the Isilon cluster that functions as a NameNode instead of being routed by a single NameNode.

Hadoop distributions supported by OneFS

You can run most of the common Hadoop distributions with the EMC Isilon cluster.

OneFS supports the following Hadoop distributions:

Hadoop distribution	Versions supported
Cloudera CDH	<ul style="list-style-type: none"> • 3 (Updates 2–5) • 4.2 • 5.0 • 5.1 • 5.2
Cloudera Manager	<ul style="list-style-type: none"> • 4.0 • 5.2
Greenplum GPHD	<ul style="list-style-type: none"> • 1.1 • 1.2
HAWQ	<ul style="list-style-type: none"> • 1.1.0.1
Hortonworks Data Platform	<ul style="list-style-type: none"> • 1.1.1–1.3.3 (non-GUI) • 2.1

Hadoop distribution	Versions supported
Pivotal HD	<ul style="list-style-type: none"> • 1.0.1 • 2.0
Apache Hadoop	<ul style="list-style-type: none"> • 0.20.203 • 0.20.205 • 1.0.0–1.0.3 • 1.2.1 • 2.0.x • 2.2–2.4

WebHDFS

OneFS supports access to HDFS data through WebHDFS client applications.

WebHDFS is a RESTful programming interface based on HTTP operations such as GET, PUT, POST, and DELETE that is available for creating client applications. WebHDFS client applications allow you to access HDFS data and perform HDFS operations through HTTP and HTTPS.

WebHDFS is supported by OneFS on a per-access zone basis and is enabled by default.

WebHDFS supports simple authentication and Kerberos authentication. If the HDFS authentication method for an access zone is set to `ALL`, OneFS uses simple authentication by default.

Note

To prevent unauthorized client access through simple authentication, disable WebHDFS in each access zone that should not support it.

Secure impersonation

Secure impersonation enables you to create proxy users that can impersonate other users to run Hadoop jobs.

You might configure secure impersonation if you use applications, such as Apache Oozie, to automatically schedule, manage, and run Hadoop jobs. For example, you can create an Oozie proxy user that securely impersonates a user called `HadoopAdmin`, which allows the Oozie user to request that Hadoop jobs be performed by the `HadoopAdmin` user.

You configure proxy users for secure impersonation on a per-zone basis, and users or groups of users that you assign as members to the proxy user must be from the same access zone. A member can be one or more of the following identity types:

- User specified by user name or UID
- Group of users specified by group name or GID
- User, group, machine, or account specified by SID
- Well-known user specified by name

If the proxy user does not present valid credentials or if a proxy user member does not exist on the cluster, access is denied. The proxy user can only access files and sub-

directories located in the HDFS root directory of the access zone. It is recommended that you limit the members that the proxy user can impersonate to users that have access only to the data the proxy user needs.

Ambari agent

The Ambari client/server framework is a third-party tool that enables you to configure, manage, and monitor a Hadoop cluster through a browser-based interface. The OneFS Ambari agent allows you to monitor the status of HDFS services on the EMC Isilon cluster through the Ambari interface.

The Ambari agent is configured per access zone; you can configure the OneFS Ambari agent in any access zone that contains HDFS data. To start an Ambari agent in an access zone, you must specify the address of the external Ambari server and the address of a NameNode that acts as the point of contact for the access zone.

The external Ambari server receives communications from the OneFS Ambari agent. Once the Ambari agent assigned to the access zone registers with the Ambari server, the agent provides a heartbeat status at regular intervals. The OneFS Ambari agent does not provide metrics or alerts to the Ambari server. The external Ambari server must be specified by a resolvable hostname, FQDN, or IP address and must be assigned to an access zone.

The NameNode is the designated point of contact in an access zone that Hadoop services managed through the Ambari interface should connect through. For example, if you manage services such as YARN or Oozie through the Ambari interface, the services will connect to the access zone through the specified NameNode. The Ambari agent communicates the location of the designated NameNode to the Ambari server, and to the Ambari interface, the NameNode represents the access zone. If you change the designated NameNode address, the Ambari agent will inform the Ambari server. The NameNode must be a resolvable SmartConnect zone name or an IP address from the IP address pool associated with the access zone.

Note

The specified NameNode value maps to the NameNode, secondary NameNode, and DataNode components on the Ambari interface.

The OneFS Ambari agent is based on the Apache Ambari framework and is compatible with Ambari server versions 1.5.1.110 and 1.6.0.

Virtual HDFS racks

You can create a virtual HDFS rack of nodes on the EMC Isilon cluster to optimize performance and reduce latency when accessing HDFS data.

A virtual HDFS rack enables you to specify a pool of preferred HDFS nodes on the EMC Isilon cluster and specify an associated pool of Hadoop compute clients.

When a Hadoop compute client from the defined pool connects to the cluster, OneFS returns at least two IP addresses from the pool of preferred HDFS nodes.

Virtual HDFS racks allow you to fine-tune client connectivity by directing Hadoop compute clients to go through quicker, less-busy switches or to faster nodes, depending on your network topology.

HDFS implementation considerations

Implementing HDFS requires you to take other areas of OneFS into consideration to ensure successful connections and access to HDFS data.

HDFS directories and Hadoop user accounts

Before implementing Hadoop, ensure that the directories and user accounts that you will need for Hadoop are configured on the EMC Isilon cluster.

When you set up directories, files, accounts, and permissions, ensure that they have the correct permissions so that Hadoop clients and applications can access the directories and files. Directories and permissions will vary by Hadoop distribution, environment, requirements, and security policies.

You must also ensure that the user accounts that your Hadoop distribution requires are configured on the Isilon cluster on a per-zone basis. The user accounts that you need and the associated owner and group settings vary by distribution, requirements, and security policies. The profiles of the accounts, including UIDs and GIDS, on the Isilon cluster should match those of the accounts on your Hadoop compute clients.

OneFS must be able to look up a local Hadoop user by name. If there are no directory services, such as Active Directory or LDAP, that can perform a user lookup, you must create a local Hadoop user. If directory services are available, a local user account is not required.

HDFS settings in access zones

Some HDFS attributes must be configured for each access zone on the EMC Isilon cluster.

You configure one HDFS root directory for each access zone. When a Hadoop compute client connects to the cluster, the user can access all files and sub-directories in the specified root directory. Unlike NFS mounts or SMB shares, clients connecting to the cluster through HDFS cannot be given access to individual folders within the root directory. The default HDFS directory is `/ifs`.

If you have multiple Hadoop workflows that require separate sets of data, you can create multiple access zones and configure a unique HDFS root directory for each zone.

You configure the authentication method for each access zone. HDFS supports simple authentication, Kerberos authentication, or both. When a Hadoop compute client connects to an access zone on the Isilon cluster, the client must authenticate with the method specified for that access zone. By default, HDFS accepts both simple and Kerberos authentication.

Proxy users are configured on a per-zone basis. Members assigned to the proxy user must belong to the same access zone.

When HDFS is licensed, the HDFS service is enabled on the entire cluster. You cannot disable HDFS on a per-access zone basis. If you create multiple access zones, you must configure HDFS settings for each zone that you want Hadoop compute clients to access through HDFS.

HDFS and SmartConnect

You can set up a SmartConnect zone for connections from Hadoop compute clients.

SmartConnect is a module that specifies how the DNS server on the EMC Isilon cluster handles connection requests from clients.

Each SmartConnect zone represents a specific pool of IP addresses. When you associate a SmartConnect zone with an access zone, OneFS only allows Hadoop clients connecting through the IP addresses in the SmartConnect zone to reach the HDFS data in the access zone. A root HDFS directory is specified for each access zone. This configuration isolates data within access zones and allows you to restrict client access to the data.

A SmartConnect zone evenly distributes NameNode requests from Hadoop compute clients across the access zone. When a Hadoop compute client makes an initial DNS request to connect to the SmartConnect zone, the Hadoop client is routed to an Isilon node that serves as a NameNode. Subsequent requests from the Hadoop compute client go to the same node. When a second Hadoop client makes a DNS request for the SmartConnect zone, SmartConnect balances the traffic and routes the client connection to a different node than that used by the previous Hadoop compute client.

If you create a SmartConnect zone, you must add a new name server (NS) record as a delegated domain to the authoritative DNS zone that contains the Isilon cluster. On the Hadoop compute cluster, you must add the name of the DNS entry of the SmartConnect zone to the `core-site.xml` file so that your Hadoop compute clients connect to a NameNode with the DNS name of the zone.

SmartConnect is discussed in further detail in the *Networking* section of this guide.

Implementing Hadoop with OneFS

To support Hadoop on the EMC Isilon cluster, you must configure HDFS on the Isilon cluster to communicate with a Hadoop cluster.

The process for configuring HDFS on the Isilon cluster is summarized in the following list:

- Activate a license for HDFS. When a license is activated, the HDFS service is enabled by default.
- Create directories on the cluster that will be set as HDFS root directories.
- Create a SmartConnect zone for balancing connections from Hadoop compute clients.
- Create local Hadoop users in access zones that do not have directory services such as Active Directory or LDAP.
- Set the HDFS root directory in each access zone that supports HDFS connections.
- Enable or disable WebHDFS in each access zone.
- Set an authentication method in each access zone that supports HDFS connections.
- Configure HDFS service settings on the cluster.
- Configure proxy users for secure impersonation.
- Configure virtual HDFS racks.

Managing the HDFS service

You can configure HDFS service settings on the EMC Isilon cluster to improve performance for HDFS workflows.

Configure HDFS service settings

You can configure HDFS service settings to improve performance for HDFS workflows.

HDFS service settings can be configured only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and then log in.
2. Run the `isi hdfs settings modify` command.

The following example command sets the block size to 1 GB:

```
isi hdfs settings modify --default-block-size=1G
```

You must specify the block size in bytes. Suffixes K, M, and G are allowed.

The following example command sets the checksum type to `crc32`:

```
isi hdfs settings modify --default-checksum-type=crc32
```

The following example command sets the log level to `CRIT`:

```
isi hdfs settings modify --server-log-level=CRIT
```

The following example command sets the number of server threads to 32:

```
isi hdfs settings modify --server-threads=32
```

HDFS service settings

HDFS service settings affect the performance of HDFS workflows.

You can configure the following HDFS service settings:

Setting	Description
Block size	<p>The HDFS block size setting on the EMC cluster determines how the HDFS service returns data upon read requests from Hadoop compute client. You can modify the HDFS block size on the cluster to increase the block size from the default of 64 MB up to 128 MB. Increasing the block size enables the Isilon cluster nodes to read and write HDFS data in larger blocks and optimize performance for most use cases.</p> <p>The Hadoop cluster maintains a different block size that determines how a Hadoop compute client writes a block of file data to the Isilon cluster. The optimal block size depends on your data, how you process your data, and other factors. You can configure the block size on the Hadoop cluster in the <code>hdfs-site.xml</code> configuration file in the <code>dfs.block.size</code> property.</p>
Checksum type	<p>The HDFS service sends the checksum type to Hadoop compute clients, but it does not send any checksum data, regardless of the checksum type. The default checksum type is set to <code>None</code>. If your Hadoop distribution requires a checksum type other than <code>None</code> to the client, you can set the checksum type to <code>CRC32</code> or <code>CRC32C</code>.</p>
Service threads	<p>The HDFS service generates multiple threads to handle HDFS traffic from EMC Isilon nodes.</p> <p>By default, the service thread value is set to <code>auto</code>, which calculates the thread count by multiplying the number of cores on a node by eight and adding a minimum threshold of thirteen. It generates a maximum of 96 threads on a node.</p> <p>To support a large system of Hadoop compute clients, you might need to increase the number of threads. If you are distributing HDFS traffic across all of</p>

Setting	Description
	the nodes in an Isilon cluster through a SmartConnect zone, the total number of HDFS service threads should equal at least half of the total number of maps and reduces on the Hadoop compute cluster. The maximum thread count is 256 per node.
Logging level	<p>The HDFS service supports the following logging levels:</p> <ul style="list-style-type: none"> • Emergency—panic conditions broadcast to all users • Alert—conditions that must be corrected immediately, such as a corrupt system database • Critical—critical conditions, such as a hard device error • Error—general errors • Notice—conditions that are not errors, but might require special handling • Information—information messages that do not require action • Debug—information typically useful only when debugging a program

View HDFS service settings

You can view configuration details for the HDFS service.

HDFS service settings can be viewed only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and then log in.
2. Run the `isi hdfs settings view` command.

The system displays output similar to the following example:

```
Default Block Size:      64M
Default Checksum Type:  none
Server Log Level:       crit
Server Threads:         auto
```

Enable or disable the HDFS service

The HDFS service, which is enabled by default after you activate an HDFS license, can be enabled or disabled by running the `isi services` command.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in by using the root user account.
2. At the command prompt, run the `isi service` command to enable or disable the HDFS service, `isi_hdfs_d`.
 - To enable the HDFS service, run the following command:

```
isi services isi_hdfs_d enable
```
 - To disable the HDFS service, run the following command:

```
isi services isi_hdfs_d disable
```

Managing HDFS access zone settings

Some HDFS attributes must be configured for each access zone on the EMC Isilon cluster. You can specify values for the following HDFS attributes within each access zone:

- HDFS root directory
- Authentication method
- WebHDFS support
- Ambari agent settings

Supported HDFS authentication methods

The authentication method determines what credentials are required by OneFS to establish a Hadoop compute client connection.

An HDFS authentication method is specified for each access zone. OneFS supports the following authentication methods for HDFS:

Authentication method	Description
Simple only	Requires only a user name to establish client connections.
Kerberos only	Requires Kerberos credentials to establish client connections. Note You must configure Kerberos as an authentication provider on the EMC Isilon cluster, and you must modify the <code>core-site.xml</code> file on clients running Hadoop 2.2 and later.
All (default value)	Accepts both simple authentication and Kerberos credentials. If Kerberos settings and file modifications are not completed, client connections default to simple authentication. CAUTION To prevent unintended access through simple authentication, set the authentication method to <code>Kerberos only</code> to enforce client access through Kerberos.

Set the HDFS authentication method in an access zone

You can configure the HDFS authentication method within each access zone on the EMC Isilon cluster.

Before you begin

If you want to Hadoop clients to connect to an access zone through Kerberos, a Kerberos authentication provider must be configured on the cluster.

HDFS access zone settings can only be configured through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.

2. (Optional) To identify the name of the access zone that you want to modify for HDFS, run the following command:

```
isi zone zones list
```

3. Set the HDFS authentication method for the access zone by running the `isi zone zones modify <zone>` command, where `<zone>` is the name of the zone.

The following command specifies that Hadoop compute clients connecting to `zone3` must be identified through the simple authentication method:

```
isi zone zones modify zone3 --hdfs-authentication=simple_only
```

The following example command specifies that Hadoop compute clients connecting to `zone3` must be identified through the Kerberos authentication method:

```
isi zone zones modify zone3 --hdfs-authentication=kerberos_only
```

After you finish

To ensure that users can authenticate through Kerberos, you must modify the `core-site.xml` file on clients running Hadoop 2.2 and later.

Configure HDFS authentication properties on the Hadoop client

If you want clients running Hadoop 2.2 and later to connect to an access zone through Kerberos, you must make some modifications to the `core-site.xml` and `hdfs-site.xml` files on the Hadoop clients.

Before you begin

Kerberos must be set as the HDFS authentication method and a Kerberos authentication provider must be configured on the cluster.

Procedure

1. Go to the `$HADOOP_CONF` directory on your Hadoop client.
2. Open the `core-site.xml` file in a text editor.
3. Set the value of the `hadoop.security.token.service.use_ip` property to `false` as shown in the following example:

```
<property>
  <name>hadoop.security.token.service.use_ip</name>
  <value>>false</value>
</property>
```

4. Save and close the `core-site.xml` file.
5. Open the `hdfs-site.xml` file in a text editor.
6. Set the value of the `dfs.namenode.kerberos.principal.pattern` property to the Kerberos realm as shown in the following example:

```
<property>
  <name>dfs.namenode.kerberos.principal.pattern</name>
  <value>hdfs/*@storage.company.com</value>
</property>
```

7. Save and close the `hdfs-site.xml` file.

Create a local Hadoop user

OneFS must be able to look up a local Hadoop user by name. If there are no directory services in an access zone that can perform a user lookup, you must create a local Hadoop user that maps to a user on a Hadoop compute client for that access zone. If directory services are available, a local user account is not required.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster.
2. Run the `isi auth users create` command.

The following example command creates a user named `hadoop-user1` assigned to a local authentication provider within the `zone3` access zone:

```
isi auth users create --name=hadoop-user1 --provider=local --
zone=zone3
```

Set the HDFS root directory in an access zone

You configure one HDFS root directory for each access zone.

Before you begin

The directory structure you want to set as the root path should already exist on the OneFS file system.

HDFS access zone settings can only be configured through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the EMC Isilon cluster and log in.
2. (Optional) To identify the name of the zone you want to modify for HDFS, run the following command:

```
isi zone zones list
```

3. Configure the HDFS root directory for this access zone by running the `isi zone zones modify <zone>` command, where `<zone>` is the name of the zone.

The following command specifies that Hadoop compute clients connecting to `zone3` are given access to the `/ifs/hadoop/` directory:

```
isi zone zones modify zone3 --hdfs-root-directory=/ifs/hadoop
```

Enable or disable WebHDFS within an access zone

You can specify whether WebHDFS is supported per access zone.

HDFS access-zone settings can only be configured through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the EMC Isilon cluster and log in.
2. (Optional) To identify the name of the zone you want to modify for HDFS, run the following command:

```
isi zone zones list
```

3. Enable or disable WebHDFS in the access zone by running the `isi zone zones modify` command.

The following command enables WebHDFS in zone3:

```
isi zone zones modify zone3 --webhdfs-enabled=yes
```

The following command disables WebHDFS in zone3:

```
isi zone zones modify zone3 --webhdfs-enabled=no
```

Configure Ambari agent settings

You can configure Ambari agent support in each access zone that contains HDFS data. Ambari agent settings can only be configured through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. (Optional) To identify the name of the access zone that you want to configure for Ambari agent settings, run the following command:

```
isi zone zones list
```

3. Run the `isi zone zones modify` command.

The following example specifies `company.ambari.server.com` as the external Ambari server that receives communication from the Ambari agent running in the zone3 access zone:

```
isi zone zones modify zone3 \
--hdfs-ambari-server=company.ambari.server.com
```

The following example designates the IP address `192.168.205.5` as the point of contact in the zone3 access zone for Hadoop services managed through the Ambari interface:

```
isi zone zones modify zone3 --hdfs-ambari-namenode=192.168.205.5
```

Configuring secure impersonation

Configure and manage proxy users that can securely impersonate other users and groups.

Create a proxy user

You can create a proxy user that securely impersonates another user.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi hdfs proxyusers create` command.

The following command designates `hadoop-user23` in zone1 as a new proxy user:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1
```

The following command designates `hadoop-user23` in `zone1` as a new proxy user and adds the group `hadoop-users` to the list of members that the proxy user can impersonate:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1 --add-group=hadoop-users
```

The following command designates `hadoop-user23` in `zone1` as a new proxy user and adds UID 2155 to the list of members that the proxy user can impersonate:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1 --add-UID=2155
```

Modify a proxy user

You can modify a proxy user that securely impersonates another user.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi hdfs proxyusers modify` command.

The following command adds the well-known user `local` to, and removes the user whose UID is 2155 from, the list of members for proxy user `hadoop-user23` in `zone1`:

```
isi hdfs proxyusers modify hadoop-user23 --zone=zone1 --add-wellknown=local --remove-uid=2155
```

Delete a proxy user

You can delete a proxy user that securely impersonates another user.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi hdfs proxyusers delete` command.

The following command deletes `hadoop-user23` in `zone1` from the list of proxy users:

```
isi hdfs proxyusers delete hadoop-user23 --zone=zone1
```

List the members of a proxy user

You can display all groups of users and individual users, known as members, that can be impersonated by a specific proxy user.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi hdfs proxyusers members list` command.

The following command displays a detailed list of the users and groups of users that are members of proxy user `hadoop-user23` in `zone1`:

```
isi hdfs proxyusers members list hadoop-user23 --zone=zone1 -v
```


The system displays output similar to the following example:

```
Type: user
Name: krb_user_005
   ID: UID:1004
-----
Type: group
Name: krb_users
   ID: SID:S-1-22-2-1003
-----
Type: wellknown
Name: LOCAL
   ID: SID:S-1-2-0
```

View proxy users

You can view all proxy users in a specific zone or you can view information for a specific proxy user.

This procedure is available only through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. To view a list of all proxy users configure in a specific access zone, run the `isi hdfs proxyusers list` command.

The following command displays a list of all proxy users configured in zone1:

```
isi hdfs proxyusers list --zone=zone1
```

The system displays output similar to the following example:

```
Name
-----
hadoop-user23
hadoop-user25
hadoop-user28
-----
Total: 3
```

3. To view the configuration details for a specific proxy user, run the `isi hdfs proxyusers view` command.

The following command displays the configuration details for the `hadoop-user23` proxy user in zone1:

```
isi hdfs proxyusers view hadoop-user23 --zone=zone1
```

The system displays output similar to the following example:

```
Name: hadoop-user23
Members: krb_users
         LOCAL
         krb_user_004
```

Managing virtual HDFS racks

You can manage virtual HDFS racks of nodes on the EMC Isilon cluster.

A virtual HDFS rack is a pool of nodes on the Isilon cluster associated with a pool of Hadoop compute clients. You can create, modify, and delete virtual racks.

Create a virtual HDFS rack

You can create a virtual HDFS rack on the EMC Isilon cluster.

Virtual HDFS racks can only be created through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Create a virtual HDFS rack by running the `isi hdfs racks create` command.

A rack name begins with a forward slash—for example, `/hdfs-rack2`.

The following command creates a rack named `/hdfs-rack2`:

```
isi hdfs racks create /hdfs-rack2
```

The following command creates a rack named `hdfs-rack2`, specifies `120.135.26.10-120.135.26.20` as the IP address range of Hadoop compute client associated with the rack, and specifies `subnet0:pool0` as the pool of Isilon nodes assigned to the rack:

```
isi hdfs racks create /hdfs-rack2 --client-ip-ranges=120.135.26.10-120.135.26.20 --ip-pools=subnet0:pool0
```

Modify a virtual HDFS rack

You can modify the settings of a virtual HDFS rack.

Virtual HDFS racks can only be modified through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. (Optional) To identify the name of the virtual HDFS rack you want to modify, run the following command:

```
isi hdfs racks list
```

3. Modify the virtual HDFS rack by running the `isi hdfs racks modify` command.

A rack name begins with a forward slash—for example, `/hdfs-rack2`.

The following example command renames a rack named `/hdfs-rack2` to `/hdfs-rack5`:

```
isi hdfs racks modify /hdfs-rack2 --new-name=/hdfs-rack5
```

The following example command adds `120.135.26.30-120.135.26.40` to the list of existing Hadoop compute client IP addresses on the rack named `/hdfs-rack2`:

```
isi hdfs racks modify /hdfs-rack2 --add-client-ip-ranges=120.135.26.30-120.135.26.40
```

In addition to adding a new range to the list of existing ranges, you can modify the client IP address ranges by replacing the current ranges, deleting a specific range or deleting all ranges.

The following example command replaces any existing IP pools with subnet1:pool1 and subnet2:pool2 on the rack named /hdfs-rack2:

```
isi hdfs racks modify /hdfs-rack2 --ip-
pools=subnet1:pool1,subnet2:pool2
```

In addition to replacing the list of existing pools with new pools, you can modify the IP pools by adding pools to list of current pools, deleting a specific pool or deleting all pools.

Delete a virtual HDFS rack

You can delete a virtual HDFS rack from an EMC Isilon cluster.

Virtual HDFS racks can only be deleted through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. (Optional) To identify the name of the virtual HDFS rack you want to delete, run the following command:

```
isi hdfs racks list
```

3. Delete a virtual HDFS rack by running the `isi hdfs racks delete` command.

A rack name begins with a forward slash—for example, /hdfs-rack2.

The following command deletes the virtual HDFS rack named /hdfs-rack2:

```
isi hdfs racks delete /hdfs-rack2
```

4. At the prompt, type **yes**.

View virtual HDFS racks

You can view information for all virtual HDFS racks or for a specific rack.

Virtual HDFS rack settings can only be viewed through the command-line interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. To view a list of all virtual HDFS racks configured on the cluster, run the `isi hdfs racks list` command.

The system displays output similar to the following example:

```

Name          Client IP Ranges      IP Pools
-----
/hdfs-rack1  10.20.30.40-10.20.30.50  subnet0:pool0
/hdfs-rack2  20.10.30.10-20.10.30.20  subnet1:pool1
-----
Total: 2
```

The following example command displays setting details for all virtual HDFS racks configured on the cluster:

```
isi hdfs racks list -v
```

The system displays output similar to the following example:

```
Name: /hdfs-rack1
Client IP Ranges: 10.20.30.40-10.20.30.50
IP Pools: subnet0:pool0
-----
Name: /hdfs-rack2
Client IP Ranges: 20.10.30.10-20.10.30.20
IP Pools: subnet1:pool1
```

3. To view the setting details for a specific virtual HDFS rack, run the `isi hdfs racks view` command:

Each rack name begins with a forward slash—for example `/hdfs-rack2`.

The following example command displays setting details for the virtual HDFS rack named `/hdfs-rack2`:

```
isi hdfs racks view /hdfs-rack2
```

The system displays output similar to the following example:

```
Name: /hdfs-rack2
Client IP Ranges: 20.10.30.10-20.10.30.20
IP Pools: subnet1:pool1
```

HDFS commands

You can access and configure the HDFS service through the HDFS commands.

isi hdfs settings modify

Modifies the settings of the HDFS service on the EMC Isilon cluster.

Syntax

```
isi hdfs settings modify
  [--default-block-size <size>]
  [--default-checksum-type {none | crc32 | crc32c}]
  [--server-log-level {emerg | alert | crit | err | notice | info |
  debug} ]
  [--server-threads <integer>]
```

Options

`--default-block-size <size>`

Specifies the block size (in bytes) reported by the HDFS service. K, M, and G; for example, 64M, 512K, 1G, are valid suffixes. The default value is 64M.

`--default-checksum-type {none | crc32 | crc32c}`

Specifies the checksum type reported by the HDFS service. The default value is `none`.

`--server-log-level {emerg | alert | crit | err | notice | info | debug}`

Sets the default logging level for the HDFS service on the cluster. The following values are valid:

EMERG

A panic condition. This is normally broadcast to all users.

ALERT

A condition that should be corrected immediately, such as a corrupted system database.

CRIT

Critical conditions, such as hard device errors.

ERR

Errors.

NOTICE

Conditions that are not error conditions, but may need special handling.

INFO

Information messages.

DEBUG

Messages that contain information typically of use only when debugging a program.

The default value is NOTICE.

```
--server-threads {<integer>|auto}
```

Specifies the number of worker threads generated by the HDFS service. The default value is `auto`, which enables the HDFS service to determine the number of necessary worker threads.

isi hdfs settings view

Displays the current settings of the HDFS service on the EMC Isilon cluster.

Syntax

```
isi hdfs settings view
```

Options

There are no options for this command.

isi hdfs proxyusers create

Creates a proxy user that can securely impersonate another user or group.

Syntax

```
isi hdfs proxyusers create <proxyuser-name>
  [--zone <zone-name>]
  [--add-group <group-name>...]
  [--add-gid <group-identifier>...]
  [--add-user <user-name>...]
  [--add-uid <user-identifier>...]
  [--add-sid <security-identifier>...]
  [--add-wellknown <well-known-name>...]
  [--verbose]
```

Options

<proxyuser-name>

Specifies the user name of a user currently configured on the cluster to be designated as a proxy user.

`--zone <zone-name>`

Specifies the access zone the user authenticates through.

`--add-group <group-name>...`

Adds the group specified by name to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple group names in a comma-separated list.

`--add-gid <group-identifier>...`

Adds the group by specified by UNIX GID to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple UNIX GIDs in a comma-separated list.

`--add-user <user-name>...`

Adds the user specified by name to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple user names in a comma-separated list.

`--add-uid <user-identifier>...`

Adds the user specified by UNIX UID to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple UNIX UIDs in a comma-separated list.

`--add-sid <security-identifier>...`

Adds the user, group of users, machine or account specified by Windows SID to the list of proxy user members. The object must authenticate to the same access zone as the proxy user. You can specify multiple Windows SIDs in a comma-separated list.

`--add-wellknown <well-known-name>...`

Adds the well-known user specified by name to the list of members the proxy user can impersonate. The well-known user must authenticate to the same access zone as the proxy user. You can specify multiple well-known user names in a comma-separated list.

`{ --verbose | -v }`

Displays more detailed information.

Examples

The following command designates `hadoop-user23` in `zone1` as a new proxy user:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1
```

The following command designates `hadoop-user23` in `zone1` as a new proxy user and adds the group of users named `hadoop-users` to the list of members that the proxy user can impersonate:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1 \
--add-group=hadoop-users
```

The following command designates `hadoop-user23` in `zone1` as a new proxy user and adds UID 2155 to the list of members that the proxy user can impersonate:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1 --add-UID=2155
```

isi hdfs proxyusers modify

Modifies a proxy user that can securely impersonate another user or group.

Syntax

```
isi hdfs proxyusers modify <proxyuser-name>
  [--zone <zone-name>]
  [--add-group <group-name>...]
  [--add-gid <group-identifier>...]
  [--add-user <user-name>...]
  [--add-uid <user-identifier>...]
  [--add-sid <security-identifier>...]
  [--add-wellknown <well-known-name>...]
  [--remove-group <group-name>...]
  [--remove-gid <group-identifier>...]
  [--remove-user <user-name>...]
  [--remove-uid <user-identifier>...]
  [--remove-sid <security-identifier>...]
  [--remove-wellknown <well-known-name>...]
  [--verbose]
```

Options

<proxyuser-name>

Specifies the user name of the proxy user to be modified.

--zone <zone-name>

Specifies the access zone that the proxy user authenticates through.

--add-group <group-name>...

Adds the group specified by name to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple group names in a comma-separated list.

--add-gid <group-identifier>...

Adds the group specified by UNIX GID to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple UNIX GIDs in a comma-separated list.

--add-user <user-name>...

Adds the user specified by name to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple user names in a comma-separated list.

--add-uid <user-identifier>...

Adds the user specified by UNIX UID to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple UNIX UIDs in a comma-separated list.

--add-sid <security-identifier>...

Adds the user, group of users, machine or account specified by Windows SID to the list of proxy user members. The object must authenticate to the same access zone as the proxy user. You can specify multiple Windows SIDs in a comma-separated list.

--add-wellknown <well-known-name>...

Adds the well-known user specified by name to the list of members the proxy user can impersonate. The well-known user must authenticate to the same access zone as

the proxy user. You can specify multiple well-known user names in a comma-separated list.

`--remove-group <group-name>...`

Removes the group specified by name from the list of proxy user members so that the proxy user can no longer impersonate any user in the group. You can specify multiple group names in a comma-separated list.

`--remove-gid <group-identifier>...`

Removes the group specified by UNIX GID from the list of proxy user members so that the proxy user can no longer impersonate any user in the group. You can specify multiple UNIX GIDs in a comma-separated list.

`--remove-user <user-name>...`

Removes the user specified by name from the list of members the proxy user can impersonate. You can specify multiple user names in a comma-separated list.

`--remove-uid <user-identifier>...`

Removes the user specified by UNIX UID from the list of members the proxy user can impersonate. You can specify multiple UNIX UIDs in a comma-separated list.

`--remove-sid <security-identifier>...`

Removes the user, group of users, machine or account specified by Windows SID from the list of proxy user members. You can specify multiple Windows SIDs in a comma-separated list.

`--remove-wellknown <well-known-name>...`

Removes the well-known user specified by name from the list of members the proxy user can impersonate. You can specify multiple well-known user names in a comma-separated list.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command adds the well-known local user to, and removes the user whose UID is 2155 from, the list of members for proxy user `hadoop-user23` in zone1:

```
isi hdfs proxyusers modify hadoop-user23 --zone=zone1 \
--add-wellknown=local --remove-uid=2155
```

isi hdfs proxyusers delete

Deletes a proxy user.

Syntax

```
isi hdfs proxyusers delete <proxyuser-name>
[--zone <zone-name>]
[--force]
[--verbose]
```

Options

`<proxyuser-name>`

Specifies the user name of the proxy user to be deleted.

`--zone <zone-name>`

Specifies the access zone that the proxy user authenticates through.


```
{ --force | -f }
```

Deletes the specified proxy user without requesting confirmation.

```
{ --verbose | -v }
```

Displays more detailed information.

Examples

The following command deletes `hadoop-user23` in `zone1` from the list of proxy users:

```
isi hdfs proxyusers delete hadoop-user23 --zone=zone1
```

isi hdfs proxyusers members list

Displays the users and groups of users, known as members, that can be impersonated by a proxy user.

Syntax

```
isi hdfs proxyusers members list <proxyuser-name>
  [--zone <zone-name>]
  [--format {table | json | csv | list}]
  [--no-header ]
  [--no-footer ]
  [--verbose]
```

Options

<proxyuser-name>

Specifies the name of the proxy user.

`--zone <zone-name>`

Specifies the access zone the proxy user authenticates through.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`--no-header`

Displays table and CSV output without headers.

`--no-footer`

Displays table output without footers.

```
{ --verbose | -v }
```

Displays more detailed information.

Examples

The following command displays a detailed list of the users and groups that are members of proxy user `hadoop-user23` in `zone1`:

```
isi hdfs proxyusers members list hadoop-user23 --zone=zone1 -v
```

The system displays output similar to the following example:

```
Type: user
Name: krb_user_005
  ID: UID:1004
-----
Type: group
Name: krb_users
```

```
ID: SID:S-1-22-2-1003
```

```
-----
Type: wellknown
Name: LOCAL
ID: SID:S-1-2-0
```

isi hdfs proxyusers list

Displays all proxy users that are configured in an access zone.

Syntax

```
isi hdfs proxyusers list
  [--zone <zone-name>]
  [--format {table | json | csv | list}]
  [--no-header ]
  [--no-footer ]
  [--verbose]
```

Options

`--zone <zone-name>`

Specifies the name of the access zone.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`--no-header`

Displays table and CSV output without headers.

`--no-footer`

Displays table output without footers.

`{ --verbose | -v }`

Displays more detailed information.

Examples

The following command displays a list of all proxy users that are configured in zone1:

```
isi hdfs proxyusers list --zone=zone1
```

The system displays output similar to the following example:

```
Name
-----
hadoop-user23
hadoop-user25
hadoop-user28
-----
Total: 3
```

isi hdfs proxyusers view

Displays the configuration details of a specific proxy user.

Syntax

```
isi hdfs proxyusers view <proxyuser-name>
  [--zone <zone-name>]
```

Options

<proxyuser-name>

Specifies the user name of the proxy user.

`--zone <zone-name>`

Specifies the access zone the proxy user authenticates through.

Examples

The following command displays the configuration details for the `hadoop-user23` proxy user in `zone1`:

```
isi hdfs proxyusers view hadoop-user23 --zone=zone1
```

The system displays output similar to the following example:

```
Name: hadoop-user23
Members: krb_users
         LOCAL
         krb_user_004
```

isi hdfs racks create

Creates a new virtual HDFS rack.

Syntax

```
isi hdfs racks create <rack-name>
  [--client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--ip-pools <subnet>:<pool>]...
```

Options

<rack-name>

Specifies the name of the virtual HDFS rack. The rack name must begin with a forward slash—for example, `/example-name`.

`--client-ip-ranges <low-ip-address>-<high-ip-address>...`

Specifies IP address ranges of external Hadoop compute clients assigned to the virtual rack.

`--ip-pools <subnet>:<pool>...`

Assigns a pool of Isilon cluster IP addresses to the virtual rack.

isi hdfs racks modify

Modifies a virtual HDFS rack.

Syntax

```
isi hdfs racks modify <rack-name>
  [--new-name <rack-name>]
  [--client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--add-client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--remove-client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--clear-client-ip-ranges]
  [--ip-pools <subnet>:<pool>]...
  [--add-ip-pools <subnet>:<pool>]...
  [--remove-ip-pools <subnet>:<pool>]...
  [--clear-ip-pools]
```

Options*<rack-name>*

Specifies the virtual HDFS rack to be modified. Each rack name begins with a forward slash—for example **/example-name**.

`--new-name <rack-name>`

Assigns a new name to the specified virtual rack. The rack name must begin with a forward slash—for example **/example-name**.

`--client-ip-ranges <low-ip-address><high-ip-address>...`

Specifies IP address ranges of external Hadoop compute clients assigned to the virtual rack. The value assigned through this option overwrites any existing IP address ranges. You can add a new range through the `--add-client-ip-ranges` option.

`--add-client-ip-ranges <low-ip-address><high-ip-address>...`

Adds a specified IP address range of external Hadoop compute clients to the virtual rack.

`--remove-client-ip-ranges <low-ip-address><high-ip-address>...`

Removes a specified IP address range of external Hadoop compute clients from the virtual rack. You can only remove an entire range; you cannot delete a subset of a range.

`--clear-client-ip-ranges`

Removes all IP address ranges of external Hadoop compute clients from the virtual rack.

`--ip-pools <subnet>:<pool>...`

Assigns pools of Isilon node IP addresses to the virtual rack. The value assigned through this option overwrites any existing IP address pools. You can add a new pool through the `--add-ip-pools` option.

`--add-ip-pools <subnet>:<pool>...`

Adds a specified pool of Isilon cluster IP addresses to the virtual rack.

`--remove-ip-pools <subnet>:<pool>...`

Removes a specified pool of Isilon cluster IP addresses from the virtual rack.

`--clear-ip-pools`

Removes all pools of Isilon cluster IP addresses from the virtual rack.

isi hdfs racks delete

Deletes a virtual HDFS rack.

Syntax

```
isi hdfs racks delete <rack-name>
```

Options*<rack-name>*

Deletes the specified virtual HDFS rack. Each rack name begins with a forward slash—for example, **/example-name**.

isi hdfs racks list

Lists the existing HDFS racks.

Syntax

```
isi hdfs racks list  
  [--verbose]
```

Options

{--verbose | -v}

Displays more detailed information.

isi hdfs racks view

Displays information for a specific virtual HDFS rack.

Syntax

```
isi hdfs racks view <rack-name>
```

Options

<rack-name>

Specifies the name of the virtual HDFS rack to view. Each rack name begins with a forward slash—for example, **/example-name**.

CHAPTER 24

Antivirus

This section contains the following topics:

- [Antivirus overview](#) 920
- [On-access scanning](#) 920
- [Antivirus policy scanning](#) 921
- [Individual file scanning](#) 921
- [Antivirus scan reports](#) 921
- [ICAP servers](#) 921
- [Supported ICAP servers](#) 922
- [Antivirus threat responses](#) 922
- [Configuring global antivirus settings](#) 923
- [Managing ICAP servers](#) 924
- [Create an antivirus policy](#) 926
- [Managing antivirus policies](#) 926
- [Managing antivirus scans](#) 927
- [Managing antivirus threats](#) 928
- [Managing antivirus reports](#) 929
- [Antivirus commands](#) 930

Antivirus overview

You can scan the files you store on an Isilon cluster for computer viruses and other security threats by integrating with third-party scanning services through the Internet Content Adaptation Protocol (ICAP). OneFS sends files through ICAP to a server running third-party antivirus scanning software. These servers are referred to as ICAP servers. ICAP servers scan files for viruses.

After an ICAP server scans a file, it informs OneFS of whether the file is a threat. If a threat is detected, OneFS informs system administrators by creating an event, displaying near real-time summary information, and documenting the threat in an antivirus scan report. You can configure OneFS to request that ICAP servers attempt to repair infected files. You can also configure OneFS to protect users against potentially dangerous files by truncating or quarantining infected files.

Before OneFS sends a file to be scanned, it ensures that the scan is not redundant. If a file has already been scanned and has not been modified, OneFS will not send the file to be scanned unless the virus database on the ICAP server has been updated since the last scan.

Note

Antivirus scanning is available only if all nodes in the cluster are connected to the external network.

On-access scanning

You can configure OneFS to send files to be scanned before they are opened, after they are closed, or both. Sending files to be scanned after they are closed is faster but less secure. Sending files to be scanned before they are opened is slower but more secure.

If OneFS is configured to ensure that files are scanned after they are closed, when a user creates or modifies a file on the cluster, OneFS queues the file to be scanned. OneFS then sends the file to an ICAP server to be scanned when convenient. In this configuration, users can always access files without any delay. However, it is possible that after a user modifies or creates a file, a second user might access the file before the file is scanned. If a virus was introduced to the file from the first user, the second user will be able to access the infected file. Also, if an ICAP server is unable to scan a file, the file will still be accessible to users.

If OneFS ensures that files are scanned before they are opened, when a user attempts to download a file from the cluster, OneFS first sends the file to an ICAP server to be scanned. The file is not sent to the user until the scan is complete. Scanning files before they are opened is more secure than scanning files after they are closed, because users can access only scanned files. However, scanning files before they are opened requires users to wait for files to be scanned. You can also configure OneFS to deny access to files that cannot be scanned by an ICAP server, which can increase the delay. For example, if no ICAP servers are available, users will not be able to access any files until the ICAP servers become available again.

If you configure OneFS to ensure that files are scanned before they are opened, it is recommended that you also configure OneFS to ensure that files are scanned after they are closed. Scanning files as they are both opened and closed will not necessarily improve security, but it will usually improve data availability when compared to scanning files only when they are opened. If a user wants to access a file, the file may have already

been scanned after the file was last modified, and will not need to be scanned again if the ICAP server database has not been updated since the last scan.

Antivirus policy scanning

You can create antivirus scanning policies that send files from a specified directory to be scanned. Antivirus policies can be run manually at any time, or configured to run according to a schedule.

Antivirus policies target a specific directory on the cluster. You can prevent an antivirus policy from sending certain files within the specified root directory based on the size, name, or extension of the file. Antivirus policies do not target snapshots. Only on-access scans include snapshots. Antivirus scans are handled by the OneFS job engine, and function the same as any system job.

Individual file scanning

You can send a specific file to an ICAP server to be scanned at any time.

If a virus is detected in a file but the ICAP server is unable to repair it, you can send the file to the ICAP server after the virus database had been updated, and the ICAP server might be able to repair the file. You can also scan individual files to test the connection between the cluster and ICAP servers.

Antivirus scan reports

OneFS generates reports about antivirus scans. Each time that an antivirus policy is run, OneFS generates a report for that policy. OneFS also generates a report every 24 hours that includes all on-access scans that occurred during the day.

Antivirus scan reports contain the following information:

- The time that the scan started.
- The time that the scan ended.
- The total number of files scanned.
- The total size of the files scanned.
- The total network traffic sent.
- The network throughput that was consumed by virus scanning.
- Whether the scan succeeded.
- The total number of infected files detected.
- The names of infected files.
- The threats associated with infected files.
- How OneFS responded to detected threats.

ICAP servers

The number of ICAP servers that are required to support an Isilon cluster depends on how virus scanning is configured, the amount of data a cluster processes, and the processing power of the ICAP servers.

If you intend to scan files exclusively through antivirus scan policies, it is recommended that you have a minimum of two ICAP servers per cluster. If you intend to scan files on

access, it is recommended that you have at least one ICAP server for each node in the cluster.

If you configure more than one ICAP server for a cluster, it is important to ensure that the processing power of each ICAP server is relatively equal. OneFS distributes files to the ICAP servers on a rotating basis, regardless of the processing power of the ICAP servers. If one server is significantly more powerful than another, OneFS does not send more files to the more powerful server.

Supported ICAP servers

OneFS supports ICAP servers running the following antivirus scanning software:

- Symantec Scan Engine 5.2 and later.
- Trend Micro Interscan Web Security Suite 3.1 and later.
- Kaspersky Anti-Virus for Proxy Server 5.5 and later.
- McAfee VirusScan Enterprise 8.7 and later with VirusScan Enterprise for Storage 1.0 and later.

Antivirus threat responses

You can configure the system to repair, quarantine, or truncate any files that the ICAP server detects viruses in.

OneFS and ICAP servers react in one or more of the following ways when threats are detected:

Alert

All threats that are detected cause an event to be generated in OneFS at the warning level, regardless of the threat response configuration.

Repair

The ICAP server attempts to repair the infected file before returning the file to OneFS.

Quarantine

OneFS quarantines the infected file. A quarantined file cannot be accessed by any user. However, a quarantined file can be removed from quarantine by the root user if the root user is connected to the cluster through secure shell (SSH).

If you backup your cluster through NDMP backup, quarantined files will remain quarantined when the files are restored. If you replicate quarantined files to another Isilon cluster, the quarantined files will continue to be quarantined on the target cluster. Quarantines operate independently of access control lists (ACLs).

Truncate

OneFS truncates the infected file. When a file is truncated, OneFS reduces the size of the file to zero bytes to render the file harmless.

You can configure OneFS and ICAP servers to react in one of the following ways when threats are detected:

Repair or quarantine

Attempts to repair infected files. If an ICAP server fails to repair a file, OneFS quarantines the file. If the ICAP server repairs the file successfully, OneFS sends the file to the user. Repair or quarantine can be useful if you want to protect users from accessing infected files while retaining all data on a cluster.

Repair or truncate

Attempts to repair infected files. If an ICAP server fails to repair a file, OneFS truncates the file. If the ICAP server repairs the file successfully, OneFS sends the file to the user. Repair or truncate can be useful if you do not care about retaining all data on your cluster, and you want to free storage space. However, data in infected files will be lost.

Alert only

Only generates an event for each infected file. It is recommended that you do not apply this setting.

Repair only

Attempts to repair infected files. Afterwards, OneFS sends the files to the user, whether or not the ICAP server repaired the files successfully. It is recommended that you do not apply this setting. If you only attempt to repair files, users will still be able to access infected files that cannot be repaired.

Quarantine

Quarantines all infected files. It is recommended that you do not apply this setting. If you quarantine files without attempting to repair them, you might deny access to infected files that could have been repaired.

Truncate

Truncates all infected files. It is recommended that you do not apply this setting. If you truncate files without attempting to repair them, you might delete data unnecessarily.

Configuring global antivirus settings

You can configure global antivirus settings that are applied to all antivirus scans by default.

Exclude files from antivirus scans

You can prevent files from being scanned by antivirus policies.

Procedure

1. Run the `isi sync avscan settings` command.

The following command configures OneFS to scan only files with the `.txt` extension:

```
isi avscan settings --glob-enable true --glob-include true \
--glob-filter .txt
```

Configure on-access scanning settings

You can configure OneFS to automatically scan files as they are accessed by users. On-access scans operate independently of antivirus policies.

Procedure

1. Run the `isi sync avscan settings` command.

The following command configures OneFS to scan files and directories under `/ifs/data/media` when they are closed:

```
isi avscan settings --scan-on-close true \
--path-prefix /ifs/data/media
```

Configure antivirus threat response settings

You can configure how OneFS responds to detected threats.

Procedure

1. Run the `isi avscan settings` command.

The following command configures OneFS and ICAP servers to attempt to repair infected files and quarantine files that cannot be repaired:

```
isi avscan settings --repair true --quarantine true
```

Configure antivirus report retention settings

You can configure how long OneFS retains antivirus reports before automatically deleting them.

Procedure

1. Run the `isi avscan settings` command.

The following command configures OneFS to delete antivirus reports older than 12 weeks.

```
isi avscan settings --report-expiry 12w
```

Enable or disable antivirus scanning

You can enable or disable all antivirus scanning. This procedure is available only through the web administration interface.

Procedure

1. Click **Data Protection** > **Antivirus** > **Summary**.
2. In the **Service** area, click **Enable** or **Disable**.

Managing ICAP servers

Before you can send files to be scanned on an ICAP server, you must configure OneFS to connect to the server. You can test, modify, and remove an ICAP server connection. You can also temporarily disconnect and reconnect to an ICAP server.

Add and connect to an ICAP server

You can add and connect to an ICAP server. After a server is added, OneFS can send files to the server to be scanned for viruses.

Procedure

1. Run the `isi avscan settings` command.

The following command adds and connects to an ICAP server at 10.7.180.108:

```
isi avscan settings --add-server 10.7.180.108
```

Test an ICAP server connection

You can test the connection between the cluster and an ICAP server. This procedure is available only through the web administration interface.

Procedure

1. Click **Data Protection** > **Antivirus** > **Summary**.
2. In the **ICAP Servers** table, in the row for the ICAP server, click **Test connection**.

If the connection test succeeds, the **Status** column displays a green icon. If the connection test fails, the **Status** column displays a red icon.

Temporarily disconnect from an ICAP server

If you want to prevent OneFS from sending files to an ICAP server, but want to retain the ICAP server connection settings, you can temporarily disconnect from the ICAP server.

Procedure

1. To identify the ID of the ICAP server you want to disconnect from, run the following command:

```
isi avscan settings
```

The system displays the ID of an ICAP server as an integer in the `ICAP server` field.

2. To disconnect from an ICAP server, run the `isi avscan settings` command.

The following command temporarily disconnects from an ICAP server with an ID of 1:

```
isi avscan settings --disable-server 1
```

Reconnect to an ICAP server

You can reconnect to an ICAP server that you have temporarily disconnected from.

Procedure

1. To identify the ID of the ICAP server you want to reconnect to, run the following command:

```
isi avscan settings
```

The system displays the ID of an ICAP server as an integer in the `ICAP server` field.

2. To reconnect to an ICAP server, run the `isi avscan settings` command.

The following command reconnects to an ICAP server with an ID of 1:

```
isi avscan settings --enable-server 1
```

Remove an ICAP server

You can permanently disconnect from the ICAP server.

Procedure

1. To identify the ID of the ICAP server you want to remove, run the following command:

```
isi avscan settings
```

The system displays the ID of an ICAP server as an integer in the `ICAP server` field.

2. To remove an ICAP server, run the `isi avscan settings` command.

The following command removes an ICAP server with an ID of 1:

```
isi avscan settings --del-server 1
```

Create an antivirus policy

You can create an antivirus policy that causes specific files to be scanned for viruses each time the policy is run.

Procedure

1. Run the `isi avscan policy add` command.

The following command creates an antivirus policy that scans `/ifs/data` every Friday at 12:00 PM:

```
isi avscan policy add --name WeekendVirusScan --path /ifs/data \
--schedule "Every Friday at 12:00 PM"
```

Managing antivirus policies

You can modify and delete antivirus policies. You can also temporarily disable antivirus policies if you want to retain the policy but do not want to scan files.

Modify an antivirus policy

You can modify an antivirus policy.

Procedure

1. To identify the ID of the antivirus policy you want to modify, run the following command:

```
isi avscan policy
```

2. Run the `isi avscan policy edit` command.

The following command modifies a policy with an ID of `51e96a9a207c1` to be run on Saturday at 12:00 PM:

```
isi avscan policy edit --id 51e96a9a207c1 \
--schedule "Every Friday at 12:00 PM"
```

Delete an antivirus policy

You can delete an antivirus policy.

Procedure

1. To identify the ID of the antivirus policy you want to delete, run the following command:

```
isi avscan policy
```

2. Run the `isi avscan policy delete` command.

The following command deletes a policy with an ID of `51e96a9a207c1`:

```
isi avscan policy delete --id 51e96a9a207c1
```

Enable or disable an antivirus policy

You can temporarily disable antivirus policies if you want to retain the policy but do not want to scan files.

Procedure

1. To identify the ID of the antivirus policy you want to enable or disable, run the following command:

```
isi avscan policy
```

2. Run the `isi avscan policy edit` command.

The following command enables a policy with an ID of 51e96a9a207c1:

```
isi avscan policy edit --id 51e96a9a207c1 --enable true
```

The following command disables a policy with an ID of 51e96a9a207c1:

```
isi avscan policy edit --id 51e96a9a207c1 --enable false
```

View antivirus policies

You can view antivirus policies.

Procedure

1. Run the following command:

```
isi avscan policy
```

Managing antivirus scans

You can scan multiple files for viruses by manually running an antivirus policy, or scan an individual file without an antivirus policy. You can also stop antivirus scans.

Scan a file

You can manually scan an individual file for viruses. This procedure is available only through the command-line interface (CLI).

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the `isi avscan manual` command.

For example, the following command scans `/ifs/data/virus_file`:

```
isi avscan manual /ifs/data/virus_file
```

Manually run an antivirus policy

You can manually run an antivirus policy at any time. This procedure is available only through the web administration interface.

Procedure

1. Click **Data Protection** > **Antivirus** > **Policies**.
2. In the **Policies** table, in the row for a policy, click **Start**.

Stop a running antivirus scan

You can stop a running antivirus scan. This procedure is available only through the web administration interface.

Procedure

1. Click **Data Protection** > **Antivirus** > **Summary**.
2. In the **Currently Running** table, in the row for an antivirus scan, click **Cancel**.

Managing antivirus threats

You can repair, quarantine, or truncate files in which threats are detected. If you think that a quarantined file is no longer a threat, you can rescan the file or remove the file from quarantine.

Manually quarantine a file

You can quarantine a file to prevent the file from being accessed by users.

Procedure

1. Run the `isi avscan quarantine` command.

The following command quarantines `/ifs/data/badFile.txt`:

```
isi avscan quarantine /ifs/data/badFile.txt
```

Remove a file from quarantine

You can remove a file from quarantine if, for example, you believe that the file is no longer a threat.

Procedure

1. Run the `isi avscan unquarantine` command.

The following command removes `/ifs/data/badFile.txt` from quarantine:

```
isi avscan unquarantine /ifs/data/badFile.txt
```

View threats

You can view files that have been identified as threats by an ICAP server.

Procedure

1. Run the `isi avscan report threat` command.

The following command displays all recently detected threats:

```
isi avscan report threat
```

The system displays output similar to the following example:

```
Report ID: R:4db6cdbc:29bc
File: /ifs/eicar.com
Time: 04-26-2011 13:51:29
Remediation: Quarantined
Threat: EICAR Test String
```



```
Infected: eicar.com
Policy ID: MANUAL
1 records displayed.
```

Antivirus threat information

You can view information about the antivirus threats that are reported by an ICAP server.

The following information is displayed in the output of the `isi avscan report threat` command.

Report ID

The ID of the antivirus report.

File

The path of the potentially infected file.

Time

The time that the threat was detected.

Remediation

How OneFS responded to the file when the threat was detected. If OneFS did not quarantine or truncate the file, `Infected` is displayed.

Threat

The name of the detected threat as it is recognized by the ICAP server.

Infected

The name of the potentially infected file.

Policy ID:

The ID of the antivirus policy that detected the threat. If the threat was detected as a result of a manual antivirus scan of an individual file, `MANUAL` is displayed.

Managing antivirus reports

In addition to viewing antivirus reports through the web administration interface, you can export reports to a comma-separated values (CSV) file. You can also view events that are related to antivirus activity.

Export an antivirus report

You can export an antivirus report to a comma separated values (CSV) file.

Procedure

1. Run the `isi avscan report scan` command.

The following command exports a report with an ID of `R:4dadd90d:16e10` to `/ifs/data/antivirusReport.csv`:

```
isi avscan report scan --report-id R:4dadd90d:16e10 \
--export /ifs/data/antivirusReport.csv
```

View antivirus reports

You can view antivirus reports.

Procedure

1. Run the `isi avscan report scan` command.

The following command displays information about recent antivirus scans:

```
isi avscan report scan
```

View antivirus events

You can view events that relate to antivirus activity.

Procedure

1. Run the following command:

```
isi events
```

All events related to antivirus scans are classified as warnings. The following events are related to antivirus activities:

Anti-Virus scan found threats

A threat was detected by an antivirus scan. These events refer to specific reports on the Antivirus Reports page but do not provide threat details.

No ICAP Servers available

OneFS is unable to communicate with any ICAP servers.

ICAP Server Unresponsive or Invalid

OneFS is unable to communicate with an ICAP server.

Antivirus commands

You can control antivirus scanning activity on an Isilon cluster through the antivirus commands.

isi avscan policy add

Creates an antivirus scan policy.

Syntax

```
isi avscan policy add --name <name>
  [--id <id>]
  [--enable {true | false}]
  [--description <string>]
  [--path <path>...]
  [--recurse <integer>]
  [--force {true | false}]
  [--schedule <schedule>]
```

Options

`--name <name>`

Specifies a name for the policy.

`--id <id>`

Specifies an ID for the policy. If no ID is specified, one is dynamically generated.

`--enable {true | false}`

Determines whether the policy is enabled or disabled. If set to **true**, the policy is enabled. The default value is **false**.

`--description <string>`

Specifies a description for the policy.

`--path <path>`

Specifies a directory to scan when the policy is run.

`--recurse <integer>`

Note

This option has been deprecated and will not impact antivirus scans if specified.

Specifies the depth of subdirectories to include in the scan.

`--force {true | false}`

Determines whether to force policy scans. If a scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

`--schedule <schedule>`

Specifies when the policy is run.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "Saturday" and "sat" are valid.

isi avscan policy edit

Modifies an antivirus scan policy.

Syntax

```
isi avscan policy edit --id <id>
  [--enable {true | false}]
```

```
[--name <new-name>]
[--description <string>]
[--path <path>...]
[--recurse <integer>]
[--force {true | false}]
[--schedule <schedule>]
```

Options

`--option <id>`

Modifies the policy of the specified ID.

`--name <new-name>`

Specifies a new name for this policy.

`--enable {true | false}`

Determines whether this policy is enabled or disabled. If set to `true`, the policy is enabled. The default value is `false`.

`--description <string>`

Specifies a new description for this policy.

`--path <path>`

Specifies a directory to scan when this policy is run.

`--recurse <integer>`

Note

This option has been deprecated and will not impact antivirus scans if specified.

Specifies the depth of subdirectories to include in the scan.

`--force {true | false}`

Determines whether to force policy scans. If a scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

`--schedule <schedule>`

Specifies when the policy is run.
Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`
- `Every [{other | <integer>}] week [on <day>]`
- `Every [{other | <integer>}] month [on the <integer>]`
- `Every [<day>[, ...] [of every [{other | <integer>}] week]]`
- `The last {day | weekday | <day>} of every [{other | <integer>}] month`
- `The <integer> {weekday | <day>} of every [{other | <integer>}] month`
- `Yearly on <month> <integer>`
- `Yearly on the {last | <integer>} [weekday | <day>] of <month>`

Specify *<frequency>* in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "Saturday" and "sat" are valid.

isi avscan policy delete

Deletes an antivirus scan policy.

Syntax

```
isi avscan policy delete --id <id>
```

Options

`--id <id>`

Deletes the policy of the specified ID.

isi avscan policy

Displays information about antivirus scan policies.

Syntax

```
isi avscan policy --id <id>
```

Options

`--id <id>`

Displays information on only the policy of the specified ID.

isi avscan policy run

Runs an antivirus policy.

Syntax

```
isi avscan policy run --id <policy-id>
  [--report <id>]
  [--force {true | false}]
  [--update {yes | no}]
```

Options

`--id <policy-id>`

Runs the policy of the specified ID.

`--report <id>`

Assigns the specified ID to the report generated for this run of the avscan policy.

`--force {true | false}`

Determines whether to force the scan. If the scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

```
--update {yes | no}
```

Specifies whether to update the last run time in the policy file. The default value is **yes**.

isi avscan manual

Manually scans a file for viruses.

Syntax

```
isi avscan manual <name>
  [--policy <id>]
  [--report <id>]
  [--force {yes | no}]
```

Options

<name>

Scans the specified file. Specify as a file path.

```
{--policy | -p} <id>
```

Assigns a policy ID for this scan. The default ID is `MANUAL`.

```
{--report | -r} <id>
```

Assigns the specified report ID to the report that will include information about this scan. If this option is not specified, the report ID is generated dynamically.

```
--force {yes | no}
```

Determines whether to force the scan. If the scan is forced, the scan will complete regardless of whether OneFS has marked the file as having been scanned, or if global settings specify that the file should not be scanned.

isi avscan quarantine

Quarantines a file manually. Quarantined files cannot be read or written to.

Syntax

```
isi avscan quarantine <name>
```

Options

<name>

Quarantines the specified file. Specify as a file path.

isi avscan unquarantine

Removes a file from quarantine. Quarantined files cannot be read or written to.

Syntax

```
isi avscan unquarantine <name>
```

Options

<name>

Removes the specified file from quarantine. Specify as a file path.

isi avscan report threat

Displays information about recently detected threats.

Syntax

```
isi avscan report threat
  [--detail]
  [--wide]
  [--export <path>]
  [--max-results <integer>]
  [--all]
  [--report-id <id>]
  [--file <path>]
  [--remediation <action>]
```

Options

{--detail | -d}

Displays detailed information.

{--wide | -w}

Displays output in a wide table without truncations.

{--export | -e} <path>

If specified, exports the output to the specified file path.

{--max-result | -m} <integer>

If specified, displays no more than the specified number of results.

{--all | -a}

If specified, displays all threats, regardless of when the threats were detected.

{--report-id | -r} <id>

Displays only threats included in the report of the specified ID.

{--file | -f} <path>

Displays information about only the specified file.

{--remediation | -R} <action>

Displays information about threats that caused the specified action.

The following values are valid:

- infected
- truncated
- repaired
- quarantined

isi avscan report scan

Displays information about recent antivirus scans.

Syntax

```
isi avscan report scan
  [--detail]
  [--wide]
  [--export <path>]
  [--max-results <integer>]
```

```
[{--all | --report-id <id> | --policy-id <id>}]
[--running]
```

Options

If no options are specified, displays a summary of recently completed scans.

```
{--detail | -d}
```

Displays detailed output.

```
{--wide | -w}
```

Displays output in a wide table.

```
{--export | -e} <path>
```

If specified, exports the output to the specified file path.

```
{--max-result | -m} <integer>
```

If specified, displays no more than the specified number of results.

```
{--all | -a}
```

If specified, displays all scans, regardless of when the scans were run.

```
{--report-id | -r} <id>
```

Displays only the report of the specified ID.

```
{--policy-id | -p} <id>
```

Displays only reports about the policy of the specified ID.

```
{--running | -R}
```

Displays only scans that are still in progress.

isi avscan report purge

Deletes antivirus reports.

Syntax

```
isi avscan report purge
  [--expire <integer>]
  <time>
```

Options

If no options are specified, deletes reports that are older than the value specified by the `isi avscan config --report-expiry` option.

```
{--expire | -e} <integer> <time>
```

Sets the minimum age of reports to be deleted.

The following *<time>* values are valid:

- s** Specifies seconds
- d** Specifies days
- m** Specifies minutes
- w** Specifies weeks

isi avscan settings

Sets and displays global configuration settings for anti-virus scanning.

Syntax

```
isi avscan settings
  [--scan-on-open {true | false}]
  [--fail-open {true | false}]
  [--scan-on-close {true | false}]
  [--max-scan-size <float> [{B | KB | MB | GB}]]
  [--repair {true | false}]
  [{--quarantine {true | false} | --truncate {true | false}}]
  [--report-expiry <integer><time>]
  [--glob-enable {true | false}]
  [--glob-include {true | false}]
  [--glob-filter <string>]...
  [--path-prefix <path>]...
  [--add-server <url>]
  [--del-server <id>]
  [--enable-server <id>]
  [--disable-server <id>]
```

Options

`--scan-on-open {true | false}`

Determines whether files are scanned before the files are sent to users.

`--fail-open {true | false}`

If `--scan-on-open` is set to `true`, determines whether users can access files that cannot be scanned. If this option is set to `false`, users cannot access a file until the file is scanned by an ICAP server.

If `--scan-on-open` is set to `true`, this option has no effect.

`--scan-on-close {true | false}`

Determines whether files are scanned after the files are closed.

`--max-scan-size <float> [{B | KB | MB | GB}]]`

If specified, OneFS will not send files larger than the specified size to an ICAP server to be scanned.

Note

Although the parameter accepts values larger than 2GB, OneFS does not scan files larger than 2GB.

`--repair {true | false}`

Determines whether OneFS attempts to repair files that threats are detected in.

`--quarantine {true | false}`

Determines whether OneFS quarantines files that threats are detected in. If `--repair` is set to `true`, OneFS will attempt to repair the file before quarantining it.

`--truncate {true | false}`

Determines whether OneFS truncates files that threats are detected in. If `--repair` is set to `true`, OneFS will attempt to repair the file before truncating it.

`--report-expiry <integer> <time>`

Determines how long OneFS will retain antivirus scan reports before deleting them.

The following *<time>* values are valid:

- s** Specifies seconds
- d** Specifies days
- m** Specifies minutes
- w** Specifies weeks

`--glob-enable {true | false}`

Determines whether glob filters are enabled. If no glob-filters are specified, glob-filters will remain disabled even if this option is set to `true`.

`--glob-include {true | false}`

Determines how glob-filters are interpreted by OneFS. If set to `true`, OneFS will scan only files that match a glob-filter. If set to `false`, OneFS will scan only files that do not match any glob-filters.

`--glob-filter <string>`

Specifies a file name or extension. To specify multiple filters, you must include multiple `--glob-filter` options within the same command. Specifying this option will remove any existing glob filters.

You can include the following wildcards:

Wildcard	Description
*	Matches any string in place of the asterisk. For example, specifying "m*" would match "movies" and "m123"
[]	Matches any characters contained in the brackets, or a range of characters separated by a dash. For example, specifying "b[aei]t" would match "bat", "bet", and "bit" For example, specifying "1[4-7]2" would match "142", "152", "162", and "172" You can exclude characters within brackets by following the first bracket with an exclamation mark. For example, specifying "b[!ie]" would match "bat" but not "bit" or "bet" You can match a bracket within a bracket if it is either the first or last character. For example, specifying "[[c]at" would match "cat", and "[at" You can match a dash within a bracket if it is either the first or last character. For example, specifying "car[-s]" would match "cars", and "car-"
?	Matches any character in place of the question mark. For example, specifying "t?p" would match "tap", "tip", and "top"

`--path-prefix <path>`

If specified, only files contained in the specified directory path will be scanned. This option affects only on-access scans. To specify multiple directories, you must include multiple `--path-prefix` options within the same command. Specifying this option will remove any existing path prefixes.

`--add-server <url>`

Adds an ICAP server of the specified URL.

`--del-server <id>`

Removes an ICAP server of the specified ID.

`--enable-server <id>`

Enables an ICAP server of the specified ID.

`--disable-server <id>`

Disables an ICAP server of the specified ID.

isi avscan get

Displays information about the scan status of files.

Syntax

```
isi avscan get <name>
```

Options

<name>

Displays information about the file of the specified name. Specify as a file path.

CHAPTER 25

VMware integration

This section contains the following topics:

- [VMware integration overview](#).....942
- [VAAI](#)..... 942
- [VASA](#).....942
- [Configuring VASA support](#)..... 943
- [Disable or re-enable VASA](#).....945
- [Troubleshooting VASA storage display failures](#).....945

VMware integration overview

OneFS integrates with VMware infrastructures, including vSphere, vCenter, and ESXi. VMware integration enables you to view information about and interact with Isilon clusters through VMware applications.

OneFS interacts with VMware infrastructures through VMware vSphere API for Storage Awareness (VASA) and VMware vSphere API for Array Integration (VAAI).

OneFS integrates with VMware vCenter through the Isilon for vCenter plug-in. The Isilon for vCenter plug-in enables you to locally backup and restore virtual machines on an Isilon cluster. For more information about Isilon for vCenter, see the following documents:

- *Isilon for vCenter Release Notes*
- *Isilon for vCenter Installation Guide*
- *Isilon for vCenter User Guide*

VAAI

OneFS uses VMware vSphere API for Array Integration (VAAI) to support offloading specific virtual machine storage and management operations from VMware ESXi hypervisors to an Isilon cluster.

VAAI support enables you to accelerate the process of creating virtual machines and virtual disks. For OneFS to interact with your vSphere environment through VAAI, your VMware environment must include ESXi 5.0 or later hypervisors.

If you enable VAAI capabilities for an Isilon cluster, when you clone a virtual machine residing on the cluster through VMware, OneFS clones the files related to that virtual machine.

To enable OneFS to use VMware vSphere API for Array Integration (VAAI), you must install the VAAI NAS plug-in for Isilon on the ESXi server. For more information on the VAAI NAS plug-in for Isilon, see the *VAAI NAS plug-in for Isilon Release Notes*.

VASA

OneFS communicates with VMware vSphere through VMware vSphere API for Storage Awareness (VASA).

VASA support enables you to view information about Isilon clusters through vSphere, including Isilon-specific alarms in vCenter. VASA support also enables you to integrate with VMware profile driven storage by providing storage capabilities for Isilon clusters in vCenter. For OneFS to communicate with vSphere through VASA, your VMware environment must include ESXi 5.0 or later hypervisors.

Isilon VASA alarms

If the VASA service is enabled on an Isilon cluster and the cluster is added as a VMware vSphere API for Storage Awareness (VASA) vendor provider in vCenter, OneFS generates alarms in vSphere.

The following table describes the alarm that OneFS generates:

Alarm name	Description
Thin-provisioned LUN capacity exceeded	There is not enough available space on the cluster to allocate space for writing data to thinly provisioned LUNs. If this condition persists, you will not be able to write to the virtual machine on this cluster. To resolve this issue, you must free storage space on the cluster.

VASA storage capabilities

OneFS integrates with VMware vCenter through VMware vSphere API for Storage Awareness (VASA) to display storage capabilities of Isilon clusters in vCenter.

The following storage capabilities are displayed through vCenter:

Archive

The Isilon cluster is composed of Isilon NL-Series nodes. The cluster is configured for maximum capacity.

Performance

The Isilon cluster is composed of Isilon i-Series, Isilon X-Series, or Isilon S-Series nodes. The cluster is configured for maximum performance.

Note

If a node type supports SSDs but none are installed, the cluster is recognized as a capacity cluster.

Capacity

The Isilon cluster is composed of Isilon X-Series nodes that do not contain SSDs. The cluster is configured for a balance between performance and capacity.

Hybrid

The Isilon cluster is composed of nodes associated with two or more storage capabilities. For example, if the cluster contained both Isilon S-Series and NL-Series nodes, the storage capability of the cluster is displayed as *Hybrid*.

Configuring VASA support

To enable VMware vSphere API for Storage Awareness (VASA) support for a cluster, you must enable the VASA daemon on the cluster, download the Isilon vendor provider certificate and add the Isilon vendor provider in vCenter.

Enable VASA

You must enable an Isilon cluster to communicate with VMware vSphere API for Storage Awareness (VASA) by enabling the VASA daemon.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Enable VASA by running the following command:

```
isi services isi_vasa_d enable
```

Download the Isilon vendor provider certificate

To add an Isilon cluster VASA vendor provider in VMware vCenter, you must use a vendor provider certificate.

Procedure

1. In a supported web browser, connect to an Isilon cluster at `https://<IPAddress>`, where `<IPAddress>` is the IP address of the Isilon cluster.
2. Add a security exception and view the security certificate to make sure that it is current.
3. Download the security certificate and save it to a location on your machine.

For more information about exporting a security certificate, see the documentation of your browser.

Note

Record the location of where you saved the certificate. You will need this file path when adding the vendor provider in vCenter.

Add the Isilon vendor provider

You must add an Isilon cluster as a vendor provider in VMware vCenter before you can view information about the storage capabilities of the cluster through vCenter.

Before you begin

Download a vendor provider certificate.

Procedure

1. In vCenter, navigate to the **Add Vendor Provider** window.
2. Fill out the following fields in the **Add Vendor Provider** window:

Name

Type a name for this VASA provider. Specify as any string. For example, type **EMC Isilon Systems**.

URL

Type `https://<IPAddress>:8081/vasaprovider`, where `<IPAddress>` is the IP address of a node in the Isilon cluster.

Login

Type `root`.

Password

Type the password of the root user.

Certificate location

Type the file path of the vendor provider certificate for this cluster.

3. Select the **Use Vendor Provider Certificate** box.
4. Click **OK**.

Disable or re-enable VASA

You can disable or re-enable an Isilon cluster to communicate with VMware vSphere through VMware vSphere API for Storage Awareness (VASA).

To disable support for VASA, you must disable both the VASA daemon and the Isilon web administration interface. You will not be able to administer the cluster through an internet browser while the web interface is disabled. To re-enable support for VASA, you must enable both the VASA daemon and the web interface.

Procedure

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Disable or enable the web interface by running one of the following commands:
 - `isi services apache2 disable`
 - `isi services apache2 enable`
3. Disable or enable the VASA daemon by running one of the following commands:
 - `isi services isi_vasa_d disable`
 - `isi services isi_vasa_d enable`

Troubleshooting VASA storage display failures

If you are unable to view information about Isilon clusters through vSphere, follow the troubleshooting tips given below to fix the issue.

- Verify that the vendor provider certificate is current and has not expired.
- Verify that the Isilon cluster is able to communicate with VASA through the VASA daemon. If the VASA daemon is disabled, run the following command to enable it:

```
isi services isi_vasa_d enable
```

- Verify that the date and time on the cluster is set correctly.
- Verify that data has been synchronized properly from vCenter.

