

Dell EMC AppSync®

Version 3.5

User and Administration Guide

302-004-015

02

Copyright © -2017 Dell Inc. or its subsidiaries. All rights reserved.

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

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

CONTENTS

	Preface	11
Chapter 1	Introduction	13
	AppSync overview	14
	Overview of service plans.....	15
	Role-based management.....	15
	AppSync reports	16
	AppSync architecture.....	16
	AppSync server	16
	AppSync agent (host plug-in) overview.....	16
	AppSync Console (user interface).....	17
	AppSync CLI.....	17
	REST interface.....	17
	AppSync and Replication Manager.....	17
	Interoperability of AppSync and Replication Manager	17
Chapter 2	AppSync Console	19
	Console overview.....	20
	Perform actions.....	20
	Times shown in the console.....	20
	Set console preferences.....	20
	Start the AppSync console.....	21
Chapter 3	AppSync CLI Utility	23
	AppSync CLI Utility.....	24
	CLI actions.....	25
	login.....	25
	logout.....	26
	refresh.....	26
	runSP.....	27
	enableSP.....	27
	disableSP.....	28
	report.....	29
	expire.....	30
	subscribe.....	31
	unsubscribe.....	32
	listCopies.....	34
	copyDetails.....	36
	mount.....	36
	unmount.....	47
Chapter 4	Service Plans	49
	Service plan overview.....	50
	Create a service plan.....	54
	Unsubscribe from a service plan.....	55
	Advanced plan settings for File systems	55
	Advanced plan settings for SQL Server.....	56

	Configuring VPLEX storage.....	56
	Configuring VMAX storage.....	57
	Exchange service plan settings.....	58
	SQL Server service plan settings.....	61
	Oracle service plan settings.....	64
	File system service plan settings.....	69
	VMware service plan settings.....	74
Chapter 5	Protect Microsoft Exchange	77
	Overview of Exchange support	78
	Deploying AppSync for Exchange protection.....	80
	Discovering Exchange databases	80
	Removing an Exchange mailbox server	81
	Protecting DAG databases in a service plan.....	81
	Convert a standalone Exchange server to a DAG member.....	81
	Protect an Exchange database	81
	Protecting an Exchange database immediately.....	82
	Subscribing an Exchange database to a service plan.....	82
	Unsubscribing a database from a service plan	82
	Expiring a copy on demand.....	83
	Creating a database copy from the Copies page.....	83
	Service plan details.....	83
	Service plan schedule.....	84
	Application discovery.....	85
	Application mapping.....	85
	Pre-copy script	85
	Create copy.....	86
	Post-copy script.....	88
	Unmount previous copy.....	89
	Mount copy.....	89
	Validate copy.....	91
	Post-mount script.....	92
	Unmount copy.....	93
	Mounting Exchange copies.....	93
	Mount and restore limitations.....	93
	Mounting an Exchange copy on-demand.....	93
	Unmounting an Exchange copy	96
	Path mapping.....	97
	Overview of Exchange copy restore.....	99
	Affected entities during restore.....	99
	Restoring from an Exchange copy.....	100
	Recovering an Exchange database manually.....	101
	Partial restore.....	102
	Restoring logs from crash-consistent (APIt) copy.....	103
	Restoring a deleted Exchange database.....	103
Chapter 6	Protect SQL Server	105
	Overview of SQL Server support.....	106
	SQL Server prerequisites.....	106
	SQL Server supported configurations.....	107
	Considerations for SQL Server in a VMware environment.....	107
	Required permissions and rights.....	108
	Update login credentials for a SQL Server instance.....	109
	Support for AlwaysOn Availability Groups.....	110

SQL Server transaction log backup.....	110
Configure SQL Server transaction log backup.....	111
Configure log backup scripts.....	113
Run log backup on demand.....	114
View log backups for a service plan.....	115
View SQL database copies.....	115
Log backup expiration.....	118
Considerations for working with SQL Server in a cluster.....	119
SQL Server User Databases folder.....	122
Discover SQL Server instances.....	122
Protect a SQL Database.....	123
Configuring protection for SQL Server database.....	123
Unsubscribing a database from a service plan.....	124
Discovering SQL Server databases	124
SQL copies page.....	125
Creating a database copy from the Copies page.....	126
Expiring an SQL database copy on demand.....	126
Service plan summary and details.....	127
Mount considerations for SQL Server.....	138
Mount SQL Server database copy on-demand.....	139
Unmounting an SQL Server copy	146
Path mapping.....	147
SQL Server database restore overview.....	148
Restore considerations for databases in an Availability Group.....	149
Affected entities during restore.....	149
Restoring a primary database or a secondary database with failover	150
Restoring a secondary database without failover.....	150
Restoring a SQL Server copy.....	151
SQL Server restore utility (assqlrestore).....	153
Chapter 7	Protect Oracle
	159
Overview of Oracle support.....	160
Oracle permissions.....	160
Red Hat Cluster Services Integration with AppSync	160
Oracle Data Guard support.....	161
Veritas Cluster Services integration.....	163
PowerHA (HACMP) cluster integration.....	165
Prerequisites and supported configurations.....	166
Protecting a database.....	174
Discovering databases.....	174
Subscribe a database to a service plan.....	174
Oracle copies page.....	175
Service plan summary and details.....	178
Service plan schedule.....	178
Overriding service plan schedules.....	178
Application discovery.....	179
Application mapping.....	179
Storage preferences.....	179
Pre-copy script.....	180
Create copy.....	181
Automatic expiration of copies.....	181
Post-copy script.....	181
Unmount previous copy.....	182
Pre-mount script.....	182

	Mount copies	183
	Overriding mount settings in a service plan.....	183
	Post mount script.....	184
	Unmount copy.....	184
	Mount an Oracle copy.....	185
	Mounting a copy using the Oracle Mount wizard.....	187
	RMAN cataloging feature	189
	Mount on standalone server and prepare scripts for manual recovery.....	190
	Mount on cluster and recover.....	191
	Path mapping.....	192
	Restoring an Oracle copy.....	194
	Affected entities during restore.....	195
	Vdisk restore with affected entities.....	196
	Restoring a RAC copy.....	196
	Restoring a RAC copy for affected entities.....	197
Chapter 8	Protect File Systems	199
	Overview of File System support.....	200
	Protect NFS file systems on VNX, VNXe, and Unity storage.....	200
	PowerHA (HACMP) cluster integration.....	202
	File system service plan settings.....	203
	Subscribing a file system to a service plan.....	208
	Overriding service plan schedules	208
	Service plan schedule.....	209
	Application discovery.....	209
	Application mapping.....	210
	Pre-copy script phase.....	210
	Create copy phase features freeze and thaw callout scripts.....	210
	Configure retry on VSS failure.....	212
	Post-copy script phase.....	212
	Unmount previous copy.....	213
	Mount copy.....	214
	Post-mount script.....	216
	Mounting a copy with the File System Mount wizard.....	216
	Changing the mount point for an affected file system.....	218
	Unmounting a file system copy	218
	Override mount settings in a service plan.....	219
	Mounting a file system after reboot	220
	Nested mount support for File systems.....	220
	Path mapping.....	221
	Restoring a file system.....	223
Chapter 9	Protect VMware Datacenters	225
	Configuration prerequisites	226
	VMware vStorage VMFS requirements	226
	Discovering datacenters	228
	List of datacenters	228
	Adding a VMware vCenter Server	228
	List of VMware datastores	229
	Protect a VMware datastore.....	229
	Considerations when mounting a VMFS copy	237
	Mounting a datastore copy on-demand.....	237
	Unmounting a VMware datastore copy	239

	Restoring a datastore from a copy.....	239
	Virtual Machine Operations during restore.....	241
	Datastore affected entities during restore.....	241
	Restoring a virtual machine from a copy.....	242
	Virtual Machine Restore options.....	244
	File or folder restore with VMFS or NFS datastores.....	245
	Restoring a file or folder from a virtual disk.....	246
Chapter 10	Repurposing	249
	Repurposing overview.....	250
	Repurpose schedule.....	251
	Modifying a repurpose plan.....	252
	Repurpose refresh.....	252
	Repurpose expire.....	253
	Data masking using scripts.....	253
	Using the Repurpose wizard.....	254
	The AppSync Service Plan In-Progress Monitor.....	255
	View or delete repurpose copy schedules.....	256
	View repurposed copies.....	256
Chapter 11	Monitor AppSync	257
	RPO concepts and best practices.....	258
	Recovery point compliance report.....	258
	Exporting an RPO compliance report to CSV.....	258
	Summary of RPO compliance.....	259
	Alerts and associated events.....	259
	Acknowledging alerts.....	260
	Email alerts.....	260
	Configure server settings for email alerts.....	261
	Specify email alert recipients.....	261
	The AppSync Service Plan In-Progress Monitor.....	261
Chapter 12	Storage considerations	263
	VNX Block	264
	Service plan considerations for applications on VNX Block storage... 265	
	Dynamic mounts	265
	Microsoft Cluster Server mounts for SQL Server.....	265
	SAN policy on Windows Server Standard Edition	266
	VNX file	266
	Service plan considerations for an application on VNX File storage... 267	
	VNX file mount.....	268
	VNXe.....	268
	Service plan considerations with VNXe.....	269
	Mount and unmount copy considerations for VNXe.....	269
	Mount/unmount VNXe NFS datastore considerations.....	270
	VMAX	270
	Service plan considerations for applications on VMAX storage....	270
	Mount and unmount VMAX copies.....	272
	Microsoft Cluster Server mounts for SQL Server.....	272
	Repurpose copies on VMAX.....	273
	VMAX restore.....	273
	VMAX 3 and VMAX All Flash	274

Service plan considerations for applications on VMAX 3 and VMAX All Flash storage.....	274
Mount/unmount VMAX 3 and VMAX All Flash copies.....	274
VMAX 3 and VMAX All Flash repurpose overview.....	275
ViPR Controller	275
Service plan considerations with ViPR Controller.....	276
Mount and unmount ViPR Controller copies.....	276
ViPR Controller copy restore.....	277
XtremIO	277
Restore options with XtremIO storage.....	279
RecoverPoint	280
Service plan considerations for applications with RecoverPoint protection.....	280
RecoverPoint prerequisites.....	280
Dynamic or static mounts.....	281
Repurpose RecoverPoint Bookmark copies of Oracle or SQL Server databases.....	282
Recovery using RecoverPoint system created Bookmark.....	283
Restore from an APIT copy.....	284
Unity	285
Service plan considerations with Unity.....	285
Mounting and unmounting Unity NFS datastore copies.....	286
Mounting and unmounting Unity copies.....	286
Mounting and unmounting Unity NFS File system copies.....	286
Repurposing copies on Unity.....	287
VPLEX.....	287
Service plan considerations for applications on VPLEX storage...	288
Mount and unmount VPLEX copies	288
VPLEX restore considerations.....	289

Chapter 13	Troubleshooting AppSync	291
	Reboot required after installing the AppSync host plug-in.....	292
	User account does not have the correct permissions.....	292
	AppSync Exchange Interface service is partially registered.....	293
	VSS timeout issue.....	293
	Host installation and deployment issue.....	294
	Oracle ASM disk groups cannot be mounted after a host reboot.....	294
	Mount of ASM disk groups fail on RHEL 6.x and 7.x MPIO configurations....	294
	AppSync fails to mount Oracle ASM disk groups (Event - ORCL_000043) ...	295
	AppSync fails to unmount Oracle ASM disk groups (Event - ORCL_000044)	295
	Oracle database discovery failure.....	296
	Oracle database discovery failure - / file system full.....	296
	Oracle database fails to start after a reboot.....	296
	Oracle restart script not removed for UNIX hosts registered using a SUDO user on agent uninstallation.....	297
	AppSync fails to create symlinks.....	297
	Checking system logs for Oracle restart.....	297
	Database fails to start with the created SPFile	298
	Oracle recovery failure.....	298
	AppSync fails to freeze the SQL Server database in a timely manner (Event - SQL_000018).....	299
	SQL Cluster - second generation copy mount failure.....	299

Timeout error during SQL database discovery.....	300
<AppSync>\jboss\standalone\tmp\vfs\ folder disk usage.....	300
XtremIO copy creation takes time.....	300
Changing an XMS IP.....	300
Error during datastore or virtual disk mount.....	301
Virtual disk mapping failure.....	301
Protection of File systems or Oracle applications on MPIO devices fail during mapping.....	301
Repurposing file systems on multiple LUNs fail.....	302
Changed file system type not updated after host discovery.....	302
Mount of a file system snapshot to RHEL7 fails.....	303
CST lockbox restore failure.....	303
Protection and repurposing failure on VMAX.....	303
AppSync services do not start after reboot.....	304
Windows server configuration issue.....	304
Flash on Windows Server 2012 R2.....	304
Browser refresh.....	305
Scheduled service plan fails.....	305
Event logging.....	306
Error handling.....	306
Glossary	309
Index	313

CONTENTS

Preface

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to <https://support.emc.com> to ensure that you are using the latest version of this document.

Purpose

This document is part of the AppSync documentation set, and includes information on using and managing AppSync.

Audience

This guide is intended for use by customers and service providers to use and configure AppSync.

Related documentation

The following publications provide additional information:

- AppSync Installation and Configuration Guide
- AppSync Release Notes

Special notice conventions used in this document

Dell EMC uses the following conventions for special notices:

DANGER

Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

WARNING

Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

CAUTION

Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

NOTICE

Addresses practices not related to personal injury.

Note

Presents information that is important, but not hazard-related.

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Table 1 Typographical conventions

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about Dell EMC products, go to <https://support.emc.com>.

Technical support

Go to <https://support.emc.com> and click Service Center. You will see several options for contacting Dell EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Online communities

Visit <https://community.emc.com/> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all Dell EMC products.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

CHAPTER 1

Introduction

This chapter includes the following topics:

- [AppSync overview](#) 14
- [AppSync architecture](#)..... 16
- [AppSync and Replication Manager](#).....17

AppSync overview

AppSync is a software that enables Integrated Copy Data Management (iCDM) with Dell EMC's primary storage systems.

AppSync simplifies and automates the process of generating and consuming copies of production data. By abstracting the underlying storage and replication technologies, and through deep application integration, AppSync empowers application owners to satisfy copy demand for operational recovery and data repurposing on their own. In turn, storage administrators need only be concerned with initial setup and policy management, resulting in an agile, frictionless environment.

AppSync automatically discovers application databases, learns the database structure, and maps it through the virtualization layer to the underlying storage LUN. It then orchestrates all the activities required from copy creation and validation through mounting at the target host and launching or recovering the application. Supported workflows also include refresh, expire, and restore production.

Key features

- Supports physical, virtual and mixed host environments across Dell EMC Block and File storage.
- Integrates with Oracle, SQL, Exchange, VMware vCenter, and more.
- Supports customer applications (EPIC, DB2, etc.) through file system copies with callout script integration to allow freeze and thaw.
- Supports application consistent, crash consistent, and virtual machine consistent (with individual virtual machine recovery) copies.
- Supports Snaps, Clone, and RecoverPoint Bookmark.
- Supports on-demand and scheduled plans.
- Repurpose wizard supports application consistent copy creation followed by manual modifications. Second generation copies of the modified copy are then distributed and optionally deleted upon configured expiration.

Supported applications and storage

AppSync supports the following applications and storage arrays:

- Applications
 - Oracle
 - Microsoft SQL Server
 - Microsoft Exchange
 - VMware VMFS datastores
 - VMware NFS datastores
 - NFS File systems
- Storage
 - VMAX
 - VMAX 3
 - VMAX All Flash
 - VNX (Block and File)
 - VNXe

- XtremIO
- ViPR Controller
- VPLEX
- Unity
- Replication Technologies
 - VNX Advanced Snapshots
 - VNXe Unified Snapshot
 - TimeFinder Clone,
 - TimeFinder VP Snap
 - SRDF
 - SnapVX
 - RecoverPoint Bookmarks
 - XtremIO Snapshot
 - ViPR Snapshot

Overview of service plans

AppSync protects an application by creating copies of application data.

You indicate to AppSync what you want to protect by subscribing an application object to a *service plan*. When the service plan runs, a copy is created. The service plan can also mount and unmount the copy, validate it, and run user-created scripts. These actions are called phases of the service plan and may differ between applications.

AppSync includes several application-specific plans that work without change. With the **Subscribe to Plan and Run** command, you apply the settings of a service plan to the data and protect it immediately.

Role-based management

AppSync supports role-based access to resources and functionality.

You can set up AppSync to have multiple users. Each user can be assigned one or more roles that correspond to their responsibilities and requirements. You can create users that are local to AppSync, and optionally add users through an LDAP server which handles the authorization.

The following table describes the user roles.

Table 2 User roles

Role	Function
Security Administrator	Manages users access to AppSync.
Resource Administrator	Manages hosts, storage systems, servers, and RecoverPoint sites.
Service Plan Administrator	Customizes and runs service plans used for data protection.
Data Administrator	Manages the protection and recovery of data.

The *AppSync Security Configuration Guide* provides more information on the specific user roles and their permissions.

AppSync reports

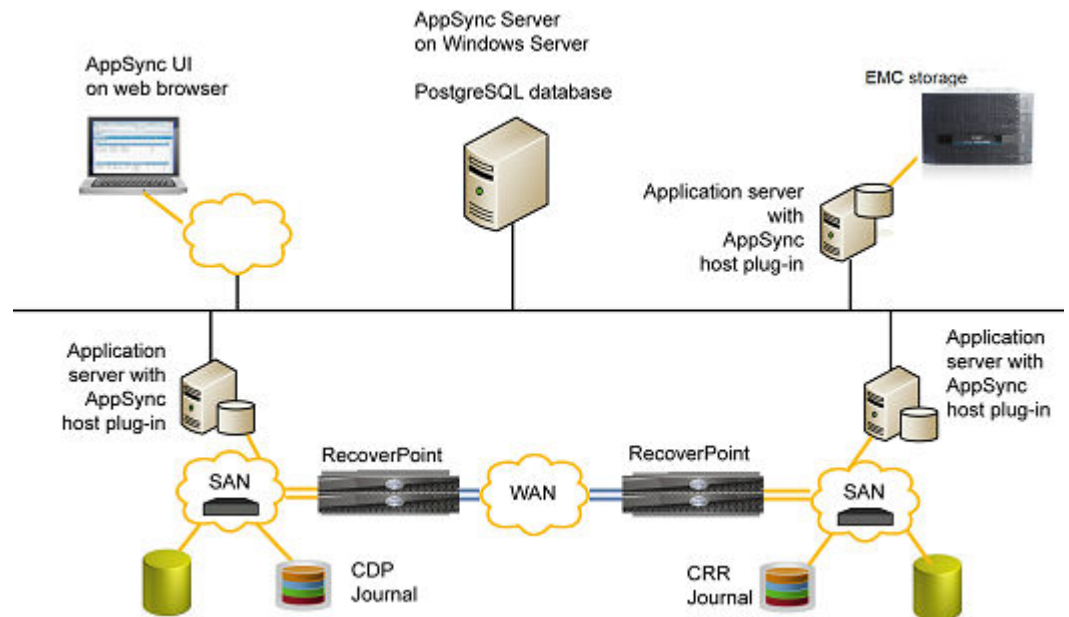
AppSync generates reports that tell you whether your data is protected, recoverable, and compliant with service level agreements.

The reports included with AppSync work without modification. Alerts and reports can be easily viewed at the top level of the AppSync dashboard. Alerts can be sent in email. AppSync can export reports to comma-separated value format.

AppSync architecture

AppSync components include the AppSync server, agent (host plug-in software), and user interfaces (UI or console).

Figure 1 AppSync architecture flow



AppSync server

The AppSync server software resides on a supported Windows system. It controls the service plans and stores data about each copy it creates.

The repository is stored in a PostgreSQL database on the AppSync server.

AppSync agent (host plug-in) overview

AppSync installs light-weight agent plug-in software on the production and mount hosts.

AppSync pushes the plug-in software from the AppSync server to the host when you add the host as a resource. In an environment that prevents the AppSync server from accessing a host, you can install the agent plug-in manually.

Note

With a push-install, the agent host plug-in remains at the same version as the AppSync server. If you want to upgrade the agent host plug-in, select **Update plug-in** with **Settings > Servers**.

For UNIX, tar bundles for AIX and Linux are pushed and extracted on the host during host registration.

Examples of hosts where the plug-in resides are Exchange mailbox servers or Exchange validation or mount hosts. The agent plug-in is not used for protection of VMware data stores.

The *AppSync Installation and Configuration Guide* provides more information on installing the AppSync agent.

AppSync Console (user interface)

The AppSync console is web-based. Supported browsers are Chrome, Internet Explorer, and Firefox. Refer to the AppSync support matrix for supported versions.

AppSync CLI

The AppSync CLI is a utility that is packaged with AppSync and is used for scripting or running tasks through a command line interface.

REST interface

AppSync has a REST interface that allows application programmers to access information controlled by AppSync.

The API is described in the *AppSync REST API Reference Guide*.

AppSync and Replication Manager

Although they are related products, there are important differences between AppSync and Replication Manager.

AppSync is primarily different from Replication Manager in that it uses the concept of service plans (protection policies) that are designed to meet specific service level requirements. The service plans are fully customizable and can be applied to the application with a single click. AppSync is easy to use for applications administrators as well as storage administrators, because it does not require prior knowledge in data replication technology.

Interoperability of AppSync and Replication Manager

Replication Manager Server and AppSync Server cannot exist on the same host.

The Replication Manager client and the AppSync agent plugin can be installed on the same host and can co-exist together.

CHAPTER 2

AppSync Console

The chapter includes the following topics:

- [Console overview](#) 20
- [Set console preferences](#) 20
- [Start the AppSync console](#) 21

Console overview

The AppSync console is arranged in sections for management, reporting, and administration.

- The **Dashboard** is a customizable view of reports and alerts. The default dashboard shows recovery point objective (RPO) status of protected applications, service plan completion status, most recent alerts, and activity in progress.
- **Copy Management** lists discovered applications such as Microsoft Exchange and Microsoft SQL Server, and provides the application-oriented entry point for protection, mount, restore, service plan subscription, and other operations.
- The **Service Plans** tab lets you view and modify service plan settings, view lists of objects that are subscribed to a service plan, the copies that were made of those objects, and the events that were generated when the service plan was run.
- **Monitoring** displays alerts, recovery point compliance reports, and service plan completion reports.
- **Settings** is for adding servers, VMware vCenter servers, and storage resources to AppSync. License management, user administration, and server settings are also found under Settings.
- **Support** provides access to how-to videos and the AppSync support page on the Dell EMC Support website.

User roles control which sections of the console are displayed and which operations are listed in menus. For example, the console does not display the **Copy Management** and **Service Plans** tabs for a user who has only the Security Administrator role.

Perform actions

To perform an action on an object, select the object and click the action buttons at the bottom of the page. Use the Shift key to select multiple objects. Use Ctrl to select noncontiguous objects. To perform an action on multiple objects, click the action button on the final selection.

Examples of objects are application hosts, storage systems, mailbox servers, databases, copies and users.

Times shown in the console

Times shown in the AppSync console reflect the local time of the AppSync server, not of the console.

Set console preferences

You can set the language displayed in the user interface and optimize the console for use over remote connections.

Procedure

1. Select the language of choice from the list of installed languages.
If you change the language, you must close the browser and launch AppSync again for the settings to take effect.
2. Check **Optimize for Remote Connection** to optimize the user interface for use over remote connections.

For example, certain visual effects may cause slow screen painting when you remotely access the console. Selecting this option enhances the remote response time but does not affect management options or functionality.

Start the AppSync console

You can run the AppSync console on a supported web browser from any system that has connectivity to the AppSync server.

Use `http://appsync_server:8085/appsync` to start the console.

If you are running the console on the AppSync server, you can start the console by clicking on the AppSync shortcut on the installer's desktop. *appsync_server* must be a host name, not an IP address.

CHAPTER 3

AppSync CLI Utility

This chapter includes the following topics:

- [AppSync CLI Utility](#) 24
- [CLI actions](#) 25

AppSync CLI Utility

The AppSync CLI is a utility that is packaged with AppSync and is used for scripting or running tasks through a command line interface

The AppSync CLI is installed in the `EMC\AppSync\appsync-cli` directory. You can run it on Windows with the file `appsync-cli.bat`. If you want to use the AppSync CLI from Unix, copy `EMC\AppSync\appsync-cli` directory to the Unix host and run `appsynccli.sh`.

Pre-requisites

- Java Runtime Environment (JRE) version 7 Update 6 (jre1.7) and above - must be installed and available in path.
- Configured AppSync installation with registered resources
- Discovered applications on registered hosts
- Configured service plans
- If you are using the CLI on a non-English host, ensure that you set the correct code page before execution. To set the code page, use `chcp` on the command prompt.

Using the CLI

You can run the AppSync CLI on the server where the AppSync installation resides. Also, you can move the `\EMC\Appsync\appsync-cli` directory to another location/host. All actions that are performed, for scripting purposes, return code zero 0 for success and local system failure code -1 for Windows or 255 for Linux). The syntax for using the AppSync CLI follows:

```
appsync-cli.bat -action options=value
```

You preface the action that you want to perform with a hyphenated `-argument`. All options specific to that action are `key=value` pairs. When using a value that contains spaces such as a file system or path, you are not required to surround the text in double quotations. Do not surround a value that ends with a trailing backslash with double quotes. Java ignores this construct.

The AppSync CLI also has two optional arguments for message handling. At any point, you can use the argument `verbose=true` for a more detailed messaging output, and `silent=true` to suppress all messages.

The Help "/"? argument

To evoke a detailed help menu for a command, add the `/?` argument. This argument displays all available CLI commands. Because of the complexity and vast number of arguments, the CLI help uses the following help menu partitions:

- `appsync-cli.bat /?` Returns information on all CLI-supported actions.
- `appsync-cli.bat -action /?` Returns non-specific application options available for the selected action.
- `appsync-cli.bat -action app=<value> /?` Returns application-specific options available for the selected action.
- `appsync-cli.bat -mount app=<value> option=<value> /?` Returns mount-specific options for the provided mount option.

Note

When using the help argument on non-English system locales you must enclose the help argument `"/?"` in double quotation marks.

CLI actions

This section describes the AppSync CLI actions.

The AppSync CLI supports the following actions:

- Login/logout
- Run a service plan
- Enable/disable a service plan
- List all copies that are created for a service plan or application object
- List all details of an application object
- Subscribe/unsubscribe an application object to/from a service plan
- Mount/unmount a copy
- Expire a copy
- Run and export AppSync reports
- Refresh

login

Authenticates the AppSync server.

Syntax

```
login
-server value
-port value
-user value
-password value
/?
```

Arguments

<code>-server <i>value</i></code>	The server you want to authenticate. The default is <code>server=localhost</code> .
<code>-port <i>value</i></code>	The AppSync server's HTTPS port. The default port is 8445.
<code>-user <i>value</i></code>	Specifies the user to be authenticated. The default user is <code>admin</code> .
<code>-password <i>value</i></code>	Specifies the password for the user. You are prompted to enter a password, if no password is set.
<code>/?</code>	Displays command line help.

Description

This command authenticates the AppSync server. It requires the server name, https communication port, AppSync user, and the corresponding password. For example:

```
appsync-cli.bat -login server=<server> port=8445 user=admin
password=<admin_pass>
```

After you log in, a file that is named `LOCAL_TOKEN` is created in the current directory containing required authentication information. If this file is deleted or the current session expires, a new session must be created by running the login command once again.

See also

logout

logout

Invalidates an open AppSync CLI connection.

Syntax

```
logout
```

```
/?
```

Arguments

/?	Displays command line help.
----	-----------------------------

Description

This command invalidates an open CLI connection. After you complete actions with the AppSync CLI, ensure that you log out. The log out command not only closes the current session, but also invalidates it. For example:

```
appsync-cli.bat -logout
```

See also

login

refresh

Refreshes the specified copy.

Syntax

```
refresh
```

```
-app value
```

```
-copy_ID value
```

```
/?
```

Arguments

-app <i>value</i>	The application that you want to refresh. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle
-copy_ID <i>value</i>	The UUID of the copy to be refreshed.
/?	Displays command line help.

Description

This command refreshes the specified application copy. For example:

```
appsync-cli.bat -refresh app=sql
```

See also

expire

runSP

Runs the specified service plan.

Syntax

```
runSP
```

```
-service_plan value
```

```
-app value
```

```
-log_backup_only value
```

```
/?
```

Arguments

<code>-service_plan <i>value</i></code>	The service plan that you want to run.
<code>-app <i>value</i></code>	Specifies the application. Values: <ul style="list-style-type: none"> • sql • oracle • filesystem • datastore • exchange
<code>-log_backup_only <i>value</i></code>	Use for on-demand SQL database log backup. Values: <ul style="list-style-type: none"> • true • false
<code>/?</code>	Displays command line help.

Description

You can run a service plan by specifying the application name and the service plan. For example:

```
appsync-cli.bat -runSP app=sql service_plan=Bronze
```

See also

enableSP

disableSP

enableSP

Enables the specified service plan.

Syntax

```
enableSP
```

```
-service_plan value
```

```
-app value
```

```
/?
```

Arguments

<code>-service_plan <i>value</i></code>	The service plan that you want to enable.
<code>-app <i>value</i></code>	Specifies the application. Values: <ul style="list-style-type: none"> • sql • oracle • filesystem • datastore • exchange
<code>/?</code>	Displays command line help.

Description

You can enable a service plan by specifying the application name and the service plan. For example:

```
appsync-cli.bat -enableSP app=sql service_plan=Bronze
```

See also

runSP

disableSP

disableSP

Disables the specified service plan.

Syntax

disableSP

`-service_plan value`

`-app value`

`/?`

Arguments

<code>-service_plan <i>value</i></code>	The service plan that you want to disable.
<code>-app <i>value</i></code>	Specifies the application. Values: <ul style="list-style-type: none"> • sql • oracle • filesystem • datastore • exchange
<code>/?</code>	Displays command line help.

Description

You can disable a service plan by specifying the application name and the service plan. For example:

```
appsync-cli.bat -disableSP app=sql service_plan=Bronze
```

See also

runSP

enableSP

report

Run and export AppSync reports.

Syntax

report

`-report_type value``-detailed value``-category value``-age value``-service_plan value``-app value`

/?

Arguments

<code>-report_type <i>value</i></code>	The type of report that you want to run. The value can be one of the following: <ul style="list-style-type: none"> • rpo • spc • alerts • activity
<code>-detailed <i>value</i></code>	The report format. You can run a detailed report or a summary report. The value can be true or false.
<code>-category <i>value</i></code>	The category of the alerts. Values: <ul style="list-style-type: none"> • all • rpo • phase_failure • other • license
<code>-age <i>value</i></code>	Specifies the duration of the events. The value can be one of the following: <ul style="list-style-type: none"> • day • week • month • all
<code>-service_plan <i>value</i></code>	Displays alerts for the specified service plan. The default value is <code>all</code> .

<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • orcale • filesystem • datastore • exchange
<code>/?</code>	Displays command line help.

Description

There are four available reports that you can run and export through the AppSync CLI. They include:

- RecoverPoint Objective (rpo)
- Service Plan Completion (spc)
- Alert
- Activity

Run reports in either summary or detailed view using the `detailed=true/false` argument. The exception to this rule occurs with an activity report which prints the activity that is currently running.

All reports are exported to a `.csv` file in the current directory with unique name from the report type and local time. For more help, use the help command (`/?`) for reports. For example:

```
appsnc-cli.bat -report report_type=rpo detailed=true
```

See also

`expire`

expire

Expires a specified copy.

Syntax

`expire`

`-app value`

`-copy_ID value`

`-force value`

`/?`

Arguments

<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore
--------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

	<ul style="list-style-type: none"> • exchange
<code>-copy_ID <i>value</i></code>	The UUID of the copy you want to expire.
<code>-force <i>value</i></code>	Removes a copy which has multiple associated copies. The value can be true or false.
<code>/?</code>	Displays command line help.

Description

To expire a copy, you must specify the application name and the copy UUID. For example: `appsync-cli.bat -expire app=datastore copy_ID=<value>`

See also

`refresh`

subscribe

Subscribes a data object to the specified service plan.

Syntax

```
subscribe
-service_plan value
-app value
/?
```

Arguments

<code>-service_plan <i>value</i></code>	The service plan that you want to subscribe to.
<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore • exchange
<code>/?</code>	Displays command line help.

Description

You can subscribe an application object to a service plan using the CLI. Options vary for each application. Run the help command `"/?"` for the application that you want to subscribe for a complete list of required arguments. For example:

```
appsync-cli.bat -subscribe app=oracle service_plan=<sp1>
oracle_server=<server> db_name=<db1>
```

Table 3 Application specific options

SQL	
<code>-sql_server <i>value</i></code>	The SQL server of the desired database. The default is <code>sql_server=localhost</code> .

Table 3 Application specific options (continued)

<code>-instance_name <i>value</i></code>	The SQL instance of the desired database. The default is <code>instance_name=MSSQLSERVER</code>
<code>-db_name <i>value</i></code>	The SQL database that you want to subscribe.
<code>-user_databases <i>value</i></code>	Allows subscription of the user database folder. The value can be true or false.
Oracle	
<code>-oracle_server <i>value</i></code>	The Oracle server of the desired database. The default is <code>oracle_server=localhost</code>
<code>-db_name <i>value</i></code>	The Oracle database that you want to subscribe.
File system	
<code>-fs_server <i>value</i></code>	The server of the desired file system. The default is <code>fs_server=localhost</code> .
<code>-fs_name <i>value</i></code>	The name of the file system. The default is <code>fs_name=C:\\.</code>
<code>-fs_type <i>value</i></code>	The format of the file system. The default is <code>fs_type=ntfs</code> .
Datastore	
<code>-datastore <i>value</i></code>	The datastore that you want to subscribe.
<code>-datacenter <i>value</i></code>	The datacenter to find the datastore.
<code>-vcenter <i>value</i></code>	The vCenter server to find the datastore. The default is <code>vcenter=localhost</code> .
Exchange	
<code>-ex_server <i>value</i></code>	The Exchange server that you want to subscribe.
<code>-db_name <i>value</i></code>	The Exchange database that you want to subscribe.

See also

unsubscribe

unsubscribe

Unsubscribes a data object from the specified service plan.

Syntax

```
unsubscribe
  -service_plan value
  -app value
  /?
```


Arguments

<code>-service_plan <i>value</i></code>	The service plan that you want to unsubscribe from.
<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore • exchange
<code>/?</code>	Displays command line help.

Description

You can unsubscribe an application object from a service plan using the CLI. Options vary for each application. Run the help command `"/?"` for the application that you want to unsubscribe for a complete list of required arguments. For example:

```
appsync-cli.bat -unsubscribe app=sql service_plan=<sp1>
sql_server=<server> instance_name=<instance> db_name=<db1>
```

Table 4 Application specific options

SQL	
<code>-sql_server <i>value</i></code>	The SQL server of the desired database. The default is <code>sql_server=localhost</code> .
<code>-instance_name <i>value</i></code>	The SQL instance of the desired database. The default is <code>instance_name=MSSQLSERVER</code>
<code>-db_name <i>value</i></code>	The SQL database that you want to unsubscribe.
<code>-user_databases <i>value</i></code>	Allows you to unsubscribe the user database folder. The value can be true or false.
Oracle	
<code>-oracle_server <i>value</i></code>	The Oracle server of the desired database. The default is <code>oracle_server=localhost</code>
<code>-db_name <i>value</i></code>	The Oracle database that you want to unsubscribe.
File system	
<code>-fs_server <i>value</i></code>	The server of the desired file system. The default is <code>fs_server=localhost</code> .
<code>-fs_name <i>value</i></code>	The name of the file system. The default is <code>fs_name=C:\.</code>
<code>-fs_type <i>value</i></code>	The format of the file system. The default is <code>fs_type=ntfs</code> .
Datastore	

Table 4 Application specific options (continued)

<code>-datastore <i>value</i></code>	The datastore that you want to unsubscribe.
<code>-datacenter <i>value</i></code>	The datacenter to find the datastore.
<code>-vcenter <i>value</i></code>	The vCenter server to find the datastore. The default is <code>vcenter=localhost</code> .
Exchange	
<code>-ex_server <i>value</i></code>	The Exchange server that you want to unsubscribe. In the case of DAG, the server name is the DAG name.
<code>-db_name <i>value</i></code>	The Exchange database that you want to unsubscribe.

See also

subscribe

listCopies

Displays all copies that meet the specified application specific properties.

Syntax

listCopies

`-service_plan value``-app value``-age value`

/?

Arguments

<code>-service_plan <i>value</i></code>	The service plan that you want to unsubscribe from.
<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore • exchange
<code>-age <i>value</i></code>	Filters viewable copies on the console by the age of a copy. The value can be one of the following: <ul style="list-style-type: none"> • day • week • month • all
/?	Displays command line help.

Description

A copy's uuid is required before you can mount the copy. To get this information, run the `-listCopies` command for either a service plan or an application object. The arguments are application-specific so ensure that you use the help command `"/?"` for details. For example:

```
appsync-cli.bat -listCopies app=sql instance_name=<value>
db_name=<value> age=all
```

Table 5 Application specific options

SQL	
<code>-instance_name <i>value</i></code>	The SQL instance of the desired database. The default is <code>instance_name=MSSQLSERVER</code> .
<code>-db_name <i>value</i></code>	Displays the specified SQL database.
<code>-log_backup_only <i>value</i></code>	Determines whether the database log backup must be displayed or not. The value can be true or false.
<code>-onlyRepurposeCopies <i>value</i></code>	Determines whether repurposed copies must be displayed or not. The value can be true or false.
Oracle	
<code>-oracle_server <i>value</i></code>	The Oracle server of the desired database. For an Oracle RAC, enter all the nodes as a comma separated string. For example, <code>oracle_server=node1,node2</code>
<code>-db_name <i>value</i></code>	Displays the specified Oracle database..
<code>-onlyRepurposeCopies <i>value</i></code>	Determines whether repurposed copies must be displayed or not. The value can be true or false.
File system	
<code>-fs_server <i>value</i></code>	The server of the desired file system. The default is <code>fs_server=localhost</code> .
<code>-fs_name <i>value</i></code>	The name of the file system. The default is <code>fs_name=C:\\</code> .
<code>-fs_type <i>value</i></code>	The format of the file system. The default is <code>fs_type=ntfs</code> .
Datastore	
<code>-datastore <i>value</i></code>	The name of the desired datastore.
<code>-datacenter <i>value</i></code>	The datacenter to find the datastore.
<code>-vcenter <i>value</i></code>	The vCenter server to find the datastore. The default is <code>vcenter=localhost</code> .
Exchange	

Table 5 Application specific options (continued)

<code>-ex_server <i>value</i></code>	Name of the Exchange server that you want to display. In the case of DAG, the server name is the DAG name.
<code>-db_name <i>value</i></code>	Displays the specified Exchange database.

See also

`copyDetails`

copyDetails

Displays information about a specified copy.

Syntax

`copyDetails`

`-app value`

`-copy_ID value`

`/?`

Arguments

<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore • exchange
<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to display.
<code>/?</code>	Displays command line help.

Description

A copy's uuid is required before you can mount the copy. To fetch additional information of an application copy, run the `-copyDetails` command for either a service plan or an application object. The arguments are application-specific so ensure that you use the help command `"/?"` for details. For example:

```
appsnc-cli.bat -copyDetails app=<app> copy_ID=<value>
```

See also

`listCopies`

mount

Mounts a specified copy.

Syntax

`mount`

`-copy_ID value`

`-app value`

/?

Arguments

<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to mount.
<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore • exchange
<code>/?</code>	Displays command line help.

Description

The AppSync CLI supports all mount options that are available through the GUI. The options vary for each application. Run the help "/?" command for the application that you want to mount to determine the mount options. For example:

- `appsync-cli.bat -mount app=filesystem copy_ID=<value> mount_host=<value>`
 - `appsync-cli.bat -mount app=sql copy_ID=<value> option=recover recovery_instance=<value> point_in_time=<value>`
 - `appsync-cli.bat -mount app=datastore copy_ID=<value> mount_host=<value> cluster_mount=yes image_access_mode=virtual_roll`
 - `appsync-cli.bat -mount app=oracle copy_ID=<value> option=rac mount_cluster=<value> mount_servers=<server1,server2>`
- `appsync-cli.bat -subscribe app=oracle service_plan=<sp1> oracle_server=<server> db_name=<db1>`

Table 6 SQL specific options

SQL	
<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to mount.
<code>-option <i>value</i></code>	Specifies the copy recovery option. The value can be mount or recover.
Mount Standalone SQL options	
<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_all_copies <i>value</i></code>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_access <i>value</i></code>	Type of access the copy must be mounted with.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy. The value can be one of the following:

Table 6 SQL specific options (continued)

	<ul style="list-style-type: none"> • Mapped path • Default path <pre>UNIX DEFAULT: /appsync-mounts WIN DEFAULT: SystemDrive \AppSyncMounts\ProdServerName</pre>
<code>-mapped_path <i>value</i></code>	<p>The map of source to target path. This option is applicable only to mapped path.</p> <pre>UNIX: /source:::/target,/abc:::/xyz WIN: R:\:::S:\,T:\ \MountPoint:::U:\</pre>
<code>-metadata_path <i>value</i></code>	The non-default path to mount copy metadata.
<code>-image_access_mode <i>value</i></code>	<p>The access mode for the image. The value can be one of the following:</p> <ul style="list-style-type: none"> • logged • virtual • virtual_roll
<code>-dedicated_sg <i>value</i></code>	Specifies the dedicated storage group.
<code>disable_rp_srm <i>value</i></code>	Disables RecoverPoint SRM (Site Recovery Manager). This option is only applicable to RecoverPoint 4.1 and later. The value can be true or false.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
Mount and Recover SQL options	
<code>-recovery_instance <i>value</i></code>	The SQL Server instance to be used for recovery.
<code>-recovery_type <i>value</i></code>	The type of recovery desired.
<code>-db_naming_suffix <i>value</i></code>	Specify a suffix that must be appended to the database after mount.
<code>-mount_path <i>value</i></code>	<p>The non-default path to mount a copy. The value can be one of the following:</p> <ul style="list-style-type: none"> • Mapped path

Table 6 SQL specific options (continued)

	<ul style="list-style-type: none"> Default path <pre>UNIX DEFAULT: /appsync-mounts WIN DEFAULT: SystemDrive \AppSyncMounts\ProdServerName</pre>
<code>-mapped_path <i>value</i></code>	<p>The map of source to target path. This option is applicable only to mapped path.</p> <pre>UNIX: /source:::/target,/abc:::/xyz WIN: R:\:::S:\,T: \MountPoint:::U:\</pre>
<code>-metadata_path <i>value</i></code>	The non-default path to mount copy metadata.
<code>-image_access_mode <i>value</i></code>	<p>The access mode for the image. The value can be one of the following:</p> <ul style="list-style-type: none"> logged virtual virtual_roll
<code>-dedicated_sg <i>value</i></code>	Specifies the dedicated storage group.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.

Table 7 Oracle specific options

Oracle	
<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to mount..
<code>-option <i>value</i></code>	<p>Specifies the copy recovery option. The value can be one of the following:</p> <ul style="list-style-type: none"> mount rman recover manual rac
Mount Standalone Oracle options	

Table 7 Oracle specific options (continued)

<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_all_copies <i>value</i></code>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy.
<code>-image_access_mode <i>value</i></code>	The access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> • logged • virtual • virtual_roll
<code>disable_rp_srm <i>value</i></code>	Disables RecoverPoint SRM (Site Recovery Manager). This option is only applicable to RecoverPoint 4.1 and later. The value can be true or false.
<code>filesystem_check <i>value</i></code>	Performs a file system check during mount. This is only applicable to UNIX and LINUX hosts. The value can be true or false.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
Mount RMAN Oracle options	
<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_all_copies <i>value</i></code>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy.
<code>-image_access_mode <i>value</i></code>	The access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> • logged • virtual • virtual_roll
<code>-rman_user <i>value</i></code>	Specifies the RMAN user name.
<code>-rman_password <i>value</i></code>	Specifies the RMAN password.
<code>-rman_connect_string <i>value</i></code>	Specifies the RMAN connect string.
<code>-tns_admin <i>value</i></code>	The non-default path of TNS_ADMIN.
<code>-oracle_home <i>value</i></code>	The non-default path of ORACLE_HOME.

Table 7 Oracle specific options (continued)

<code>-asm_dg_name <i>value</i></code>	The non-default name for the ASM disk group.
<code>-skip_data_files <i>value</i></code>	Allows you to skip data files.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
Mount and Recover Oracle options	
<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy.
<code>-image_access_mode <i>value</i></code>	Access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> • logged • virtual • virtual_roll
<code>-open_mode <i>value</i></code>	Specifies the open mode for the copy after recovery.
<code>-oracle_home <i>value</i></code>	The non-default path of ORACLE_HOME.
<code>-database_name <i>value</i></code>	The non-default name for the database.
<code>-sid_name <i>value</i></code>	The non-default name for the SID.
<code>-asm_dg_name <i>value</i></code>	The non-default name for the ASM disk group.
<code>-init_params <i>value</i></code>	Specifies the custom init parameters for recovery.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
Mount Manual Recovery ORACLE	
<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy.

Table 7 Oracle specific options (continued)

<code>-image_access_mode <i>value</i></code>	Access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> logged virtual virtual_roll
<code>-open_mode <i>value</i></code>	Specifies the open mode for the copy after recovery.
<code>-oracle_home <i>value</i></code>	The non-default path of ORACLE_HOME.
<code>-database_name <i>value</i></code>	The non-default name for the database.
<code>-sid_name <i>value</i></code>	The non-default name for the SID.
<code>-asm_dg_name <i>value</i></code>	The non-default name for the ASM disk group.
<code>-init_params <i>value</i></code>	Specifies the custom init parameters for recovery.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
Mount RAC	
<code>-mount_cluster <i>value</i></code>	The cluster you wish to mount a copy to. The default is Original Cluster.
<code>-mount_server <i>value</i></code>	The server on which to mount the copy.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy.
<code>-image_access_mode <i>value</i></code>	Access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> logged virtual virtual_roll
<code>-open_mode <i>value</i></code>	Specifies the open mode for the copy after recovery.
<code>-oracle_home <i>value</i></code>	The non-default path of ORACLE_HOME.
<code>-database_name <i>value</i></code>	The non-default name for the database.
<code>-sid_name <i>value</i></code>	The non-default name for the SID.
<code>-asm_dg_name <i>value</i></code>	The non-default name for the ASM disk group.

Table 7 Oracle specific options (continued)

<code>-init_params <i>value</i></code>	Specifies the custom init parameters for recovery.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.

Table 8 File system specific options

File system	
<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to mount.
<code>-option <i>value</i></code>	Specifies the copy recovery option. The value can be mount or recover.
<code>-mount_host <i>value</i></code>	The host on which to mount the copy. The default is original host.
<code>-mount_all_copies <i>value</i></code>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_access <i>value</i></code>	Type of access the copy must be mounted with. The value can be readonly or readwrite.
<code>-mount_path <i>value</i></code>	<p>The path to mount a copy. The value can be one of the following:</p> <ul style="list-style-type: none"> • Mapped path • Default path <pre>UNIX DEFAULT: /appsync-mounts WIN DEFAULT: SystemDrive \AppSyncMounts\ProdServerName</pre>
<code>-mapped_path <i>value</i></code>	<p>The map of source to target path. This option is applicable only to mapped path.</p> <pre>UNIX: /source:::/target,/abc:::/xyz WIN: R:\:::S:\,T: \MountPoint:::U:\</pre>
<code>-image_access_mode <i>value</i></code>	<p>The access mode for the image. The value can be one of the following:</p> <ul style="list-style-type: none"> • logged • virtual

Table 8 File system specific options (continued)

	<ul style="list-style-type: none"> virtual_roll
<code>disable_rp_srm</code> <i>value</i>	Disables RecoverPoint SRM (Site Recovery Manager). This option is only applicable to RecoverPoint 4.1 and later. The value can be true or false.
<code>filesystem_check</code> <i>value</i>	Performs a file system check during mount. This is only applicable to UNIX and LINUX hosts. The value can be true or false.
<code>-point_in_time</code> <i>value</i>	Allows you to mount a point in time copy. (FORMAT: \"MM/dd/yyyy hh:mm:ss am/pm \")
<code>-desired_SLO</code> <i>value</i>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP</code> <i>value</i>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount</code> <i>value</i>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount</code> <i>value</i>	Enables VMware cluster mount. The value can be true or false.
<code>-dedicated_sg</code> <i>value</i>	Specifies the dedicated storage group. The value can be true or false.

Table 9 Datastore specific options

Datastore	
<code>-copy_ID</code> <i>value</i>	The UUID of the copy that you want to mount.
<code>-mount_host</code> <i>value</i>	The host on which to mount the copy. The default is original host.
<code>-mount_all_copies</code> <i>value</i>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_signature</code> <i>value</i>	Allows you to specify whether you want to use the original or new mount signature. The value can be new or original.
<code>-cluster_mount</code> <i>value</i>	Specifies whether you want to mount to a cluster or not. The value can be yes or no.
<code>-image_access_mode</code> <i>value</i>	The access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> logged virtual virtual_roll
<code>disable_rp_srm</code> <i>value</i>	Disables RecoverPoint SRM (Site Recovery Manager). This option is only applicable to

Table 9 Datastore specific options (continued)

	RecoverPoint 4.1 and later. The value can be true or false.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy. (FORMAT: \"MM/dd/yyyy hh:mm:ss am/pm \")
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.

Table 10 Exchange specific options

Exchange	
<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to mount.
<code>-option <i>value</i></code>	Specifies the copy recovery option. The value can be mount or validate.
Mount Standalone Exchange options	
<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_all_copies <i>value</i></code>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_access <i>value</i></code>	Type of access the copy must be mounted with.
<code>-mount_path <i>value</i></code>	The path to mount a copy. The value can be one of the following: <ul style="list-style-type: none"> • Mapped path • Default path <pre>UNIX DEFAULT: /appsync-mounts WIN DEFAULT: SystemDrive \AppSyncMounts\ProdServerName</pre>
<code>-mapped_path <i>value</i></code>	The map of source to target path. This option is applicable only to mapped path. <pre>UNIX: /source:::/target,/abc:::/xyz WIN: R:\:::S:\,T: \MountPoint:::U:\</pre>
<code>-metadata_path <i>value</i></code>	The non-default path to mount copy metadata.

Table 10 Exchange specific options (continued)

<code>-image_access_mode <i>value</i></code>	The access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> • logged • virtual • virtual_roll
<code>disable_rp_srm <i>value</i></code>	Disables RecoverPoint SRM (Site Recovery Manager). This option is only applicable to RecoverPoint 4.1 and later. The value can be true or false.
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
Mount and Validate Standalone Exchange	
<code>-mount_host <i>value</i></code>	The host on which to mount the copy.
<code>-mount_all_copies <i>value</i></code>	Determines whether to mount all copies. The value can be true or false.
<code>-mount_access <i>value</i></code>	Type of access the copy must be mounted with.
<code>-mount_path <i>value</i></code>	The non-default path to mount a copy.
<code>-metadata_path <i>value</i></code>	The non-default path to mount copy metadata.
<code>-image_access_mode <i>value</i></code>	The access mode for the image. The value can be one of the following: <ul style="list-style-type: none"> • logged • virtual • virtual_roll
<code>-point_in_time <i>value</i></code>	Allows you to mount a point in time copy.
<code>-desired_SLO <i>value</i></code>	Specifies the desired service level objectives for VMAX V3 arrays.
<code>-desired_FAST_VP <i>value</i></code>	Specifies the FAST VP policy for VMAX V2 copies.
<code>-vplex_mount <i>value</i></code>	Specifies the VPLEX mount options.
<code>-enable_cluster_mount <i>value</i></code>	Enables VMware cluster mount.
<code>-validate_copies <i>value</i></code>	Determines whether the copies will be validated as part of the mount.

Table 10 Exchange specific options (continued)

<code>-db_logs <i>value</i></code>	Validate databases and logs. The value can be Sequentially or inparallel.
<code>-log_check <i>value</i></code>	Enables you to minimize log checking. The value can be true or false.
<code>-working_dir <i>value</i></code>	Specifies the working directory.
<code>-throttle_validation <i>value</i></code>	Enables you to throttle the validation. The value can be true or false.
<code>-pause_after_I/O_count_of <i>value</i></code>	Specifies the pause after the I/O count. The default is 100.
<code>-pause_duration <i>value</i></code>	Specifies the pause duration. The default is 1000 milliseconds.
<code>-skip_db_validation <i>value</i></code>	Skips database validation. The value can be true or false.

See also

unmount

unmount

Unmounts a specified copy.

Syntax

unmount

`-copy_ID value``-app value`

/?

Arguments

<code>-copy_ID <i>value</i></code>	The UUID of the copy that you want to unmount.
<code>-app <i>value</i></code>	The application name. The value can be one of the following: <ul style="list-style-type: none"> • sql • oracle • filesystems • datastore • exchange
/?	Displays command line help.

Description

To unmount a copy you must specify the application name and the copy uuid. For example: `appsync-cli.bat -unmount app=<app> copy_ID=<value>`.

To unmount the latest or oldest mounted copy specifically for a database, filesystem, or a datastore, use the following commands:

- **For Datastores:** `appsync-cli.bat -unmount app=datastore
datastore=<value> datacenter=<value> vcenter=<value>
option=latestMountedCopy/oldestMountedCopy`
- **For SQL:** `appsync-cli.bat -unmount app=sqlinstance_name=<value>
db_name=<value> option=latestMountedCopy/oldestMountedCopy`
- **For Oracle:** `appsync-cli.bat -unmount app=oracle
oracle_server=<value> db_name=<value>
option=latestMountedCopy/oldestMountedCopy`
- **For File systems:** `appsync-cli.bat -unmount app=filesystem
fs_server=<value> fs_name=<value> fs_type=<value>
option=latestMountedCopy/oldestMountedCopy`

See also

`mount`

CHAPTER 4

Service Plans

This chapter includes the following topics:

- [Service plan overview](#) 50
- [Exchange service plan settings](#) 58
- [SQL Server service plan settings](#) 61
- [Oracle service plan settings](#) 64
- [File system service plan settings](#) 69
- [VMware service plan settings](#) 74

Service plan overview

This section describes the default service plan types, object and copy subscriptions, service plan settings, schedules, and overrides.

AppSync creates and manages copies of application data. A service plan defines the attributes of these copies. You can subscribe application data objects to a service plan, then AppSync runs the service plan and creates copies of the data from attributes that you specified in the plan. Copies that are generated by a service plan are listed in service plan **Copies** tab.

There is no limit to the number of objects you can subscribe to a service plan. AppSync automatically divides up the work for best performance. If you need fine control over which objects are grouped for mounting, scripting, and validating, consider creating multiple service plans and distributing objects among the plans. This technique works when the objects subscribed to a service plan are from the same server. It is not recommended to subscribe more than 12 objects to any one service plan when using this method.

Service plan types

AppSync provides the following application-specific tiered plans. There are three types of service plans:

- **Bronze** — You can use the Bronze service plan to create local copies of your applications.
- **Silver** — You can use the Silver service plan to create remote copies of your applications.
- **Gold** — You can use the Gold service plan to create both local and remote copies of your applications.

Note

Ensure you understand the storage capabilities when selecting a service plan type. Not all storage technologies support Remote Replication, so Silver or Gold service plans may not be successful for the application data.

Bronze, Silver and Gold service plans are provided by default, however you can customize and create your own plans.

The following table describes the service plans and applications supported.

Storage	Replication type	Bronze	Silver	Gold	RP support	Application support	Repurposing support
VNX	Advanced Snapshot	Yes	No	No	Yes	All applications AppSync supports	Yes
	File Snapshot ^a	Yes	Yes	Yes	No	VMware datastores, File systems, and Oracle	No
VMAX	VP Snap	Yes	No	No	Yes	All applications AppSync supports	Yes
	Timefinder Clone	Yes	No	No	Yes	All applications AppSync supports	Yes

Storage	Replication type	Bronze	Silver	Gold	RP support	Application support	Repurposing support
	SRDF/A Local	Yes	No	No	Yes	All applications AppSync supports	Yes
	SRDF/A Remote ^b	No	Yes	No	No	VMware datastores and Oracle	
	SRDF/S Local	Yes	No	No	Yes	All applications AppSync supports	Yes
	SRDF/S Remote	No	Yes	No	No		
VNX2e	Unified Snapshot	Yes	No	No	Yes	All applications AppSync supports	No
	File Snapshot	Yes	No	No	No	VMware datastores	No
Unity	Unified Snapshot	Yes	No	No	Yes	All applications AppSync supports	Yes
	Unified File Snapshot	Yes	No	No	No	VMware datastores, Oracle, and File systems	No
	Thin Clone ^c	No	No	No	No	Microsoft SQL, Oracle, and File systems	Yes
XtremIO	Snapshot	Yes	No	No	Yes	All applications AppSync supports	Yes
ViPR ^d	Snapshot	Yes	No	No	No	VMware datastores, Oracle, and UNIX file systems	No
VMAX 3	SnapVX Snap	Yes	No	No	No	All applications AppSync supports	Yes
	SnapVX Clone	Yes	No	No	No	All applications AppSync supports	Yes ^e
	SRDF/A Local	Yes	No	No	No	All applications AppSync supports	Yes
	SRDF/A Remote ^f	No	Yes	No	No	VMware datastores, Oracle, and UNIX file systems	
	SRDF/S Local	Yes	No	Yes	No	All applications AppSync supports	Yes
	SRDF/S Remote	No	Yes	Yes	No	All applications AppSync supports	
VPLEX ^g	VPLEX Snap ^h	Yes	No	No	No	All applications AppSync supports	XtremIO - Yes Unity - Yes VMAX 3 and VMAX All Flash - Yes
	VPLEX Clone	Yes	No	No	No		
RecoverPoint	Local bookmark	Yes	No	No	Yes	All applications AppSync supports	Yes

Storage	Replication type	Bronze	Silver	Gold	RP support	Application support	Repurposing support
	Remote bookmark	No	Yes	No	Yes	All applications AppSync supports	Yes
	Local and remote bookmark	No	No	Yes	Yes	All applications AppSync supports	No

- a. The same is also applicable to eNAS on VMAX 3.
- b. AppSync does not support remote protection of Windows applications on VMAX 3 and SRDF/A storage.
- c. The first generation copy is always a Snap. However, the second generation copy can either be a Snap or a Clone.
- d. AppSync also supports ViPR snapshots as copy technology for applications provisioned using ViPR Controller. AppSync only supports applications provisioned by ViPR with block virtual pools that are backed only by VMAX/VPLEX (with VMAX2 and XtremIO)/XtremIO storage systems.
- e. Repurposing clone from a snap is not supported.
- f. AppSync does not support remote protection of Windows applications on VMAX 3 and SRDF/A storage.
- g. VPLEX is only supported on XtremIO, VMAX 3/VMAX All Flash and Unity back-end arrays.
- h. This is the snapshot on the back-end array.

Service plan settings

When you subscribe an object to a service plan, it joins other objects that are already part of the plan. All objects in the service plan are subject to the workflow and settings that are defined in the service plan.

Service plans set a storage ordered preference which is the preferred order of storage technology the service plan uses when creating copies. If AppSync cannot satisfy a preference, it tries to use the selected preference in the storage ordered preference list. You can adjust the preferences to create service plans that use the replication technology you want on priority. If you want AppSync to skip using a particular replication technology, deselect that preference from the storage ordered preference list.

The default service plans offer tiered levels of protection. If you must change settings, modify the service plan.

Any service plan can set the automatic expiration of copies which limits the number of copies that AppSync keeps, and automatically expires older copies that exceed the number that is defined for the service plan.

Service plans also offers a few application specific copy options which can be modified. For example, Oracle service plan has the following copy options:

- Place a database in hot-backup mode (Default: enabled)
- Copy the Fast Recovery Area (Default: disabled)
- Index and copy BCT (block change tracking) file (Default: disabled)
- Create backup control file (Default: disabled)

To avoid overutilization and depletion of replication storage, when you set up a service plan, set values in the following fields:

- **RPO** in the service plan **Startup** phase
- **Always keep n Copies** in the **Create copy** phase

Note

AppSync expiry of old copies works based on the current subscription in a service plan. If applications are added or removed to a service plan, the current expire copy count and number of copies retained might not match. To avoid this, subscribe new applications to newly added service plan than altering the application subscription often.

Also, monitor the storage system with the storage system user interface on the AppSync console.

Service plan schedule overrides

You can override a service plan's run schedule settings and specify separate schedules for individual objects that are subscribed to the plan.

In the **Plan Startup** phase of the service plan, you select a recurrence type that is based on which service plan is triggered. This recurrence type is applicable for all application objects that are subscribed to a service plan. However, you can override the settings and specify separate settings for selected objects.

You can override only the settings of the **Recurrence Type** already selected for the plan. For instance, the chosen recurrence type is to run **On selected days...** and the settings are to **Run at 12:00 AM on days Fri, Sat**. When you override these settings for an object, you can change only the time and days of the week. You cannot select a different recurrence type as part of the override.

As the Service Plan Administrator, if you change the generic recurrence settings (such as the time to run, or minutes after the hour), there is no impact to the settings of the overrides. If you change the recurrence type itself, then the overrides are no longer valid. The new recurrence type now applies to all objects until you specify individual settings that are based on the new recurrence type.

Note

If an application object is subscribed to multiple plans, the plans must not be scheduled to be running simultaneously.

Service plan events

Events show the progress of an operation. They are generated when a service plan is run, and when a copy is mounted or restored.

Event information includes:

- Type (error, warning, or informational)
- Date and time of the event
- Description
- Server

You can view events at:

- Service plan **Events** tab. For example, on the AppSync console, the **Events** tab in **Service Plans > Microsoft Exchange > Bronze** shows you the events that are related to the Exchange copy under the plan.
- Application **Copies** page. For example, go to **Copy Management > VMware Datacenters** select a data center, then a datastore to view the **Copies** page. Select a copy to view its associated events in the **Events** page.
- Events are also displayed at the time they are generated in the Mount and Restore wizards and when an object is subscribed to a plan and run immediately.

By default only the top level events, which are known as milestone events, are displayed. A milestone event is generated at the completion of each phase in a service plan cycle. You can expand a milestone event to show the other events that were generated in the phase.

Create a service plan

You can create a new service plan by using an existing plan as a template.

Before you begin

This operation requires the Service Plan Administrator role in AppSync.

Procedure

1. Select **Service Plans**.
2. Click **Create**.
3. In the **Create New Plan** dialog box, select an existing plan to use as a template. Enter a name and a description for the new service plan.

Note

The new service plan contains the same schedule and other settings as the template, but there are no objects subscribed to the new service plan.

Run a service plan on demand

Service plans run on a schedule but you can also run a service plan on demand.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Select **Service Plans**.
2. Select the application to protect.
3. Select a plan from the list and click **Run**.

This service plan run is applicable to all the application objects currently subscribed to the plan.

The service plan runs immediately. The **Run Service Plan** dialog displays progress as application storage is discovered and mapped, and application protection begins according to service plan settings.

4. Click **Details** to see more events that occurred during each phase.

Enable and disable a service plan

By default all service plans are enabled. You can disable a service plan or enable a disabled service plan.

Before you begin

This operation requires the Service Plan Administrator role in AppSync.

Procedure

1. Select **Service Plans**.
2. Select the application.

3. Select the plan and click **Disable** or **Enable**.

Delete a service plan

You can delete a user-created service plan.

Before you begin

- This operation requires the Service Plan Administrator role in AppSync.
- You cannot delete a built-in service plan (for example, Bronze, Silver, Gold).
- You cannot delete a service plan if the plan has subscriptions or if there are valid copies associated with the plan.

Procedure

1. Select **Service Plans**.
2. Select the application.
3. Select a user-created plan and click **Delete**.
4. Click **Yes** to confirm.

Unsubscribe from a service plan

You can unsubscribe applications that are subscribed to bronze, silver, gold, or the custom service plans.

Procedure

1. On the AppSync console, select **Service Plans** > *application*.
The service plans available for the selected application are listed.
2. Click the desired service plan.
The **Settings** tab of the selected service plan displays.
3. Select the **Subscriptions** tab.
The **Subscriptions** tab of the selected service plan displays.
4. Select an instance or application, and click on **Unsubscribe from Plan**.
The selected application is no more subscribed to the service plan.

Note

You can also select multiple instances and unsubscribe all of them together.

Advanced plan settings for File systems

You can disable call out scripts for file system service plans.

Procedure

1. On the AppSync console, select **Service Plans** > **File systems**.
The service plans available for the selected application are listed.
2. Click the desired service plan.
The Settings tab of the selected service plan displays.
3. Select the **Advanced Plan Settings** tab.
The Advanced Settings tab of the selected service plan displays.

4. Clear **Enable callout scripts** to disable call out scripts.

Note

- By default, this option is enabled.
 - For repurposing, if you want to disable call-out scripts during refresh, edit the repurpose plan and then clear **Enable callout scripts** under **Advanced Settings**.
-

Advanced plan settings for SQL Server

Allows you to configure maximum number of databases that can be grouped together during protection.

Procedure

1. On the AppSync console, select **Service Plans > Microsoft SQL Server**.
The service plans available for the selected application are listed.
2. Click the desired service plan.
The Settings tab of the selected service plan displays.
3. Select the **Advanced Plan Settings** tab.
The Advanced Settings tab of the selected service plan displays.
4. In the **Maximum number of Databases** field, type a value to set the maximum number of databases to be protected together.

Note

By default, only a maximum of 35 databases can be protected together. It is a Microsoft recommended best practice to not exceed the 35 databases limit.

Configuring VPLEX storage

For each Service Plan, you can select where to create copies for applications running on distributed and Local RAID-1 volumes. You can also select the preferred cluster for distributed volumes and the preferred arrays to use for RAID-1 volumes.

Procedure

1. On the AppSync console, select **Service Plans > application**.
The service plans available for the selected application are listed.
2. Click the desired service plan.
The **Settings** tab of the selected service plan displays.
3. Under the plan phases on the left, select the type of copies to create (local, remote, or local and remote). Note that the type of copy depends on the type of service plan.
The appropriate copy page appears.
4. Under the **Storage Preference** section, click the **Configure storage options** link.

The **Configure storage options** dialog box appears.

5. Select the **VPLEX** tab (if not already selected).
6. For applications running on distributed volumes, click **Cluster preference** on the left to set the preferred site for protection. By default, the preferred site is cluster-1.
7. Click **Array preference** on the left. For applications running on local or distributed RAID-1 devices, you must select the array on which you want the application to be protected. The XtremIO, VMAX 3/VMAX All Flash, and Unity arrays that are configured in AppSync and connected to the VPLEX cluster configured in AppSync are displayed. By default, no arrays are selected.

Note

To set the preference order for the selected arrays, drag-and-drop the arrays to the top or bottom of the list as desired.

Configuring VMAX storage

You can select the storage pool from which to create copy devices in V2 arrays.

Before you begin

This operation requires the Resource Administrator role in AppSync.

For each service plan, you can configure the VMAX V2 storage pools. The Service Plan selects the configured storage pool to create copy devices in V2 arrays. If you do not configure the storage pools, the Service Plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays.

Procedure

1. On the AppSync console, select **Service Plans** > *application*.

The service plans available for the selected application are listed.

2. Click the desired service plan.

The **Settings** tab of the selected service plan displays.

3. Under the plan phases on the left, select the types of copies to create (local, remote, or local and remote). Note that the type of copy depends on the type of service plan.

The appropriate copy page appears.

4. Under the **Storage Preference** section, click the **Configure storage options** link.

The **Configure storage options** dialog appears.

5. Select the **VMAX** tab.

The VMAX storage arrays are listed.

6. For VMAX 3 arrays, select **Wait for Clone Sync**, if you want AppSync to wait for clone synchronization to complete during SnapVX clone copy creation.
7. Click the arrow to the left of a storage array to view the list of storage pools on the array.

The list of storage pools for selected array appears.

8. To select a storage pool, click the checkbox to the left of the storage pool. To exclude a storage pool, click the checkbox again to clear the selection (if currently selected).
9. Click **OK**.

Exchange service plan settings

The default service plan settings create an application-consistent copy every 24 hours. Only the replication technology, which is specified by the **Copy type** in the Create copy phase, is different from plan to plan.

Table 11 Exchange Service Plan - default settings

Setting	Enabled/Not enabled	Default settings	Schedule
Plan Startup	Enabled	Automatic schedule	Recurrence type: Creates a copy every 24 hours, with the first run at midnight (00:00). Recovery Point Objective (RPO): An alert is issued if the copy is not created within the specified time (default is 24 hours).
Application discovery	Enabled	None	Determined by Plan Startup phase
Application mapping	Enabled	None	Starts when Application discovery phase completes
Pre-copy script	Not enabled	None	Starts when Application mapping phase completes
Create copy	Enabled	<ul style="list-style-type: none"> • Copy type is: <ul style="list-style-type: none"> ▪ Bronze ▪ Silver ▪ Gold • Exchange backup type: Full, Copy, or Differential 	Starts when Pre-copy script phase completes

Table 11 Exchange Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<ul style="list-style-type: none"> • Storage Ordered Preference: Snapshot, Bookmark, and Clone Allows you to order, select, or clear storage preferences. By default, all the options are selected. You cannot clear all the preferences, at least one preference must be selected. • Storage Settings: <ul style="list-style-type: none"> ▪ Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations. <hr/> <p style="margin: 0;">Note</p> <p style="margin: 0;">If this option is not checked, then RecoverPoint copies accumulate, and remain until the bookmarks for them "fall off" the RecoverPoint appliance.</p> <hr/> <ul style="list-style-type: none"> ▪ Configure storage options <ul style="list-style-type: none"> ▪ VPLEX: Allows you to select the preferred cluster for distributed volumes and the preferred arrays for RAID-1 volumes. ▪ VMAX: Allows you to configure the storage pools to create VMAX V2 copy devices for the service plan. You can select or clear the desired VMAX V2 storage pools by expanding the storage array. If you do not select storage pools, the service plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays. By default, all the configured storage pools are selected. This option is only applicable for VMAX V2 arrays. <ul style="list-style-type: none"> – Wait for Clone Sync: If you select this option, the SnapVX clone copy creation waits for clone synchronization to complete. If you clear this option, the clone copy creation initiates clone synchronization and completes the copy creation without waiting for clone synchronization to complete (clone synchronization runs in the background). By default, this option is disabled. This option is only applicable for VMAX V3 arrays. • Advanced Settings - Allows you to set a retry count and retry interval to retry the failed VSS operation after an interval configured through the retry interval. This is only applicable for Windows applications. The default retry count value is 3 and the default retry interval value is 0 seconds. • Event Log Scanning: <ul style="list-style-type: none"> ▪ Fail on -1018 error ▪ Fail on -1019 error ▪ Fail on -1022 error ▪ Fail on Event ID 447 ▪ Fail on Event ID 448 ▪ Do not allow databases and logs to reside on the same volume 	

Table 11 Exchange Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
Post-copy script	Not enabled	None	Starts when Create copy phase completes
Unmount previous copy	Not enabled	None	Starts when Post-copy script phase completes
Mount copy	Not enabled	<ul style="list-style-type: none"> • Mount on server: Original Host • Mount with access: Read-only • Mount Path: Default Path or Mapped Path <hr/> <p>Note</p> <p>The drive that is specified for mount cannot be a clustered disk.</p> <hr/> <ul style="list-style-type: none"> • Copy metadata files to: Default Path • Image access mode: Logged access • Desired SLO: Select the SLO for the mount copy. This is only applicable for VMAX 3 arrays. • Desired FAST VP: Select the FAST VP policy. This is only applicable for VMAX V2 arrays. • VPLEX Mount option: Select a VPLEX mount option. • Enable VMware cluster mount: If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. This option is supported on VPLEX, XtremIO, VMAX 3/VMAX All Flash, and Unity arrays. • Disable VMware SRM: Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later. 	Starts when Unmount previous copy phase completes
Validate copy	Not enabled	<ul style="list-style-type: none"> • Check databases and logs in parallel • Do not minimize log checking • Do not perform throttle checking • Perform validation for Database (.edb) file 	Starts when Mount copy phase completes
Post-mount script	Not enabled	None	Starts when Validate copy phase completes
Unmount copy	Not enabled	None	Starts when Post-mount

Table 11 Exchange Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
			script phase completes

SQL Server service plan settings

Summary of SQL Server service plan settings.

Table 12 SQL Server Service Plan - default settings

Setting	Enabled/Not enabled	Default settings	Schedule
Plan Startup	Enabled	Automatic schedule	Recurrence type: Creates a copy every 24 hours, with the first run at midnight (00:00). Recovery Point Objective (RPO): An alert is issued if the copy is not created within the specified time (default is 24 hours).
Application discovery	Enabled	None	Determined by Plan Startup phase
Application mapping	Enabled	None	Starts when Application discovery phase completes
Pre-copy script	Not enabled	None	Starts when Application mapping phase completes
Create copy	Enabled	<ul style="list-style-type: none"> • Copy type is: <ul style="list-style-type: none"> ▪ Bronze ▪ Silver ▪ Gold • SQL Server Backup Type: Full, Copy, Non VDI, or Crash Consistent 	Starts when Pre-copy script phase completes

Table 12 SQL Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<ul style="list-style-type: none"> ▪ Full - protects the database and the active part of the transaction log. ▪ Copy - protects the database and the active part of the transaction log without affecting the sequence of backups. ▪ Non VDI - protects the database without using VDI, and depends on VSS to create crash consistent copies. ▪ Crash Consistent - protects the database without using VSS or VDI, and depends on the array to create crash consistent copies. <hr/> <p>Note</p> <p>Secondary databases are read-only and can be backed up with the Copy backup type. Auto Switch to Copy is enabled only when Full is selected as the backup type. However, it is unchecked by default. Checking Auto Switch to Copy tells AppSync to check if the database role is Secondary, and if so, to switch the backup type to Copy. If Auto Switch to Copy is not enabled, backups fail for all secondary databases. When Non VDI or Crash Consistent backup type is selected, Auto Switch to Copy and Enable log backup are disabled.</p> <hr/> <ul style="list-style-type: none"> • Storage Ordered Preference: Snapshot, Clone, and Bookmark Allows you to order, select, or clear storage preferences. By default, all the options are selected. You cannot clear all the preferences, at least one preference must be selected. • Expiration of database copies: <ul style="list-style-type: none"> ▪ Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations. <hr/> <p>Note</p> <p>If this option is not checked, then RecoverPoint copies will accumulate, and will remain until the bookmarks for them "fall off" the RecoverPoint appliance.</p> <hr/> <ul style="list-style-type: none"> • Configure storage options <ul style="list-style-type: none"> ▪ VPLEX: Allows you to select the preferred cluster for distributed volumes and the preferred arrays for RAID-1 volumes. ▪ VMAX: Allows you to configure the storage pools to create VMAX V2 copy devices for the service plan. You can select or clear the desired VMAX V2 storage pools by expanding the storage array. If you do not select storage pools, the service plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays. By default, all the configured storage pools are selected. This option is only applicable for VMAX V2 arrays. <ul style="list-style-type: none"> – Wait for Clone Sync: If you select this option, the SnapVX clone copy creation waits for clone synchronization to 	

Table 12 SQL Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<p>complete. If you clear this option, the clone copy creation initiates clone synchronization and completes the copy creation without waiting for clone synchronization to complete (clone synchronization runs in the background). By default, this option is disabled. This option is only applicable for VMAX V3 arrays.</p> <ul style="list-style-type: none"> Advanced Settings - Allows you to set a retry count and retry interval to retry the failed VSS operation after an interval configured through the retry interval. This is only applicable for Windows applications. The default retry count value is 3 and the default retry interval value is 0 seconds. Transaction Log backup options: <ul style="list-style-type: none"> Schedule immediately after database backup = default. Every 15 or 30 minutes or every 1 to 24 hours are other options. <hr/> <p>Note</p> <p>Scheduled log backups run during times between database backups.</p> <hr/> <ul style="list-style-type: none"> Backup path: Default path Free space on the volume: 5GB Backup group size: 5 Truncate logs: Selected Checksum the backup: Unselected Compression: Selected Expiration of log backups: Minimum retention hours 24 	
Post-copy script	Not enabled	None	Starts when Create copy phase completes
Unmount previous copy	Not enabled	None	Starts when Post-copy script phase completes
Mount copy	Not enabled	<ul style="list-style-type: none"> Mount Copy <ul style="list-style-type: none"> Mount on Server: Original Host Mount with access: Read only Mount Path: Default Path or Mapped Path Copy metadata files to: Default Path Image Access mode: Logged access 	Starts when Unmount previous copy phase completes

Table 12 SQL Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<ul style="list-style-type: none"> ▪ Desired SLO: Select the SLO for the mount copy. This is only applicable for VMAX 3 arrays. ▪ Desired FAST VP: Select the FAST VP policy. This is only applicable for VMAX V2 arrays. ▪ VPLEX Mount option: Select a VPLEX mount option. ▪ Enable VMware cluster mount: If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. This option is supported on VPLEX, XtremIO, VMAX 3/VMAX All Flash, and Unity arrays. ▪ Disable VMware SRM: Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later. • Mount and recover copy - Allows you to select both clustered and standalone instances to mount a SQL Server Database either as a clustered or standalone database with recovery. For mounting as a clustered or standalone database, you can mount to the original path or to the alternate mount point. 	
Post-mount script	Not enabled	None	Starts when Mount copy phase completes
Unmount copy	Not enabled	None	Starts when Post-mount script phase completes
Pre-log backup script	Not enabled	None	Starts after log backup
Post-log backup script	Not enabled	None	Starts after log backup

Oracle service plan settings

Use this list of service-plan default settings for Oracle databases in AppSync.

Table 13 Oracle Server Service Plan - default settings

Setting	Enabled/Not enabled	Default settings	Schedule
Plan Startup	Enabled	Automatic schedule	Recurrence type: AppSync creates a copy every 24

Table 13 Oracle Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
			hours, with the first run at midnight (00:00). Recovery Point Objective (RPO): An alert is issued if the copy is not created within the specified time (default is 24 hours).
Application discovery	Enabled	None	Determined by Plan Startup phase
Application mapping	Enabled	None	Starts when Application discovery phase completes.
Pre-copy script	Enabled	None	Starts when Application mapping phase completes
Create copy	Enabled	<p>The create copy options on the service plan settings provides various controls which influence how the Oracle copy is created. Copy types include:</p> <ul style="list-style-type: none"> • Bronze • Silver • Gold <p>Create copy setting options include:</p> <ul style="list-style-type: none"> • Place database in hot-backup mode (Default: enabled) When enabled, the protection puts the database in hot backup and immediately creates copies of the archive logs. If you disable this option, the database is not placed in hot backup mode and the copy is created from the live unquiesced data without any instrumentation of the database. • Copy the Fast Recovery Area. (Default: disabled) When enabled, this field tells AppSync to create a copy of the underlying storage that is used by the FRA when protecting the database's archive log files. • Index and copy the BCT (block change tracking) file. (Default: disabled) If enabled, AppSync creates an entry in the Oracle block change tracking file and re-copies the file as part of the protection. This file can then be leveraged as part of a mount and backup use-case to 	Starts when Precopy script phase completes.

Table 13 Oracle Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<p>provide accelerated incremental backup. This option requires hot backup mode.</p> <ul style="list-style-type: none"> • Create backup control file for RMAN cataloging. (Default: disabled) If enabled, AppSync creates a binary backup control file with a request to catalog the database contents in a remote RMAN catalog. This option requires hot backup mode. • Expiration: Include RecoverPoint copies in expiration rotation policy: select this option to include RecoverPoint copies when calculating rotations. If you do not select this option, RecoverPoint copies accumulate and remain until their bookmarks rotate off the RecoverPoint appliance. <hr/> <p>Note</p> <p>If RecoverPoint copies are factored in the rotation policy, bookmarks are created with the ALWAYS_CONSOLIDATE policy. Otherwise, bookmarks are created with the NEVER_CONSOLIDATE policy. Consult the RecoverPoint documentation for a definition of these consolidation policies.</p> <hr/> <ul style="list-style-type: none"> • Storage Preference: Snapshot, Clone, and Bookmark. Allows you to order, select, or clear storage preferences. By default, all the options are selected. You cannot clear all the preferences, at least one preference must be selected. • Configure storage option: <ul style="list-style-type: none"> ▪ VPLEX: Allows you to select the preferred cluster for distributed volumes and the preferred arrays for RAID-1 volumes for each Service Plan. ▪ VMAX: Allows you to configure the storage pools to create VMAX V2 copy devices for each Service Plan. If you do not select storage pools, the service plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays. By default, all the configured storage pools are selected. You can select or clear the desired VMAX V2 storage pools by expanding the storage array. This option applies only to VMAX V2 arrays. <ul style="list-style-type: none"> – Wait for Clone Sync: If you select this option, the SnapVX clone copy creation waits for clone synchronization to complete. If you clear this option, the clone copy creation initiates clone synchronization and completes the copy creation without waiting for clone synchronization to complete (clone synchronization runs in the background). By default, this option is disabled. This option is only applicable for VMAX V3 arrays. 	
Post-copy script	Not enabled	None	Starts when Create copy phase completes.

Table 13 Oracle Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
Unmount previous copy	Not enabled	None	Starts when Post-copy phase completes
Pre-mount script	Not enabled	None	Starts when Unmount previous copy phase completes appliance.
Mount and Recovery	Not enabled	<ul style="list-style-type: none"> • Mount and recovery operations: <ul style="list-style-type: none"> ▪ Mount on standalone server (RM-equivalent : No recover) ▪ Mount on standalone server and create RMAN catalog entry (RM-equivalent : Catalog with RMAN) ▪ Mount on standalone server and recover database (RM-equivalent: Recover) ▪ Mount on standalone server and prepare scripts for manual database recovery (RM-equivalent: Prepare-only/generate scripts for manual recovery) ▪ Mount on grid cluster and recover as RAC database (RM-equivalent: Mount as RAC database) • Mount Settings: <ul style="list-style-type: none"> ▪ Mount on Server: Original Host ▪ Mount Path: Default Path or Mapped Path ▪ Image Access mode: Logged access ▪ Desired SLO: Select the SLO for the mount copy. This is applicable only to VMAX 3 arrays. ▪ Desired FAST VP: Select the FAST VP policy. This is applicable only to VMAX V2 arrays. ▪ VPLEX Mount option: Select a VPLEX mount option. ▪ Enable VMware cluster mount: If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. This option is supported on VPLEX, XtremIO, VMAX 3/VMAX All Flash, and Unity arrays. ▪ Disable VMware SRM: Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later. ▪ Run Filesystem Check: During a mount operation, the AppSync agent checks file system data consistency by executing the <code>fsck</code> command. This operation can be time consuming. You can clear this option to skip file system check during a mount operation. By default, file system check is enabled. 	Starts when Pre-mount script phase completes.

Table 13 Oracle Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<hr/> <p>Note</p> <ul style="list-style-type: none"> - In the case of a restore operation, the <code>Run Filesystem Check</code> option is enabled by default. You cannot disable it. - The <code>Run Filesystem Check</code> option is not applicable to ASM file systems. <hr/> <ul style="list-style-type: none"> • Recovery Settings: <ul style="list-style-type: none"> ▪ Open-mode: Read-write ▪ ORACLE_HOME: Same as production host ▪ Database name: APS is the prefix, %DB% is the variable which will be replaced with the production database name during run time. ▪ SID name: APS is the prefix, %SID% is the variable which will be replaced with the production database SID during run time. ▪ ASM diskgroup name: APS is the prefix, %DG% is the variable which will be replaced with the production ASM diskgroup name during run time. ▪ Customize Initialization Parameters: This field will be blank. You can fill in one parameter per line, for example, <code>memory_target=629145600</code> ▪ Create TEMP Tablespace: Use this option to create the Temp Tablespace on the recovery mounted database copy. This setting is enabled when you select the following mount operations with Read/Write Open-mode: Mount on standalone server and recover, Mount on standalone server and prepare scripts for manual recovery, or Mount on grid cluster and recover as RAC database. When you select the Create TEMP Tablespace option, two additional options display: ▪ Number of Tempfiles: The number of files to be added to Temp Tablespace. The size of the files are specified in the Size of each file setting. ▪ Size of each file: The size of each temp file (in kilobytes (K), megabytes (M), gigabytes (G), or terabytes (T)). ▪ Restart databases after reboot: Select this option to start the AppSync mounted Oracle databases automatically after a host reboot. By default, this option is disabled. <hr/> <p>Note</p> <p>This option is not available for RMAN and mount with generate scripts.</p> <hr/> <ul style="list-style-type: none"> ▪ Create SPFile: Select this option to create an SPFile. The SPFile is created in the default location (<code>\$ORACLE_HOME/dbs</code>), with 	

Table 13 Oracle Server Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		the name <code>spfile<SID>.ora</code> . During unmount, the SPFile is removed from the <code>\$ORACLE_HOME/dbs</code> folder.	
Post-mount script	Not enabled	None	Starts when mount phase completes.
Unmount copy	Not enabled	None	Starts when Post-mount script phase completes.

File system service plan settings

Use this table to learn default file system settings for service plan phases including startup, discovery, mapping, pre and post copy scripting, mount/unmount and copy.

Default service plan settings create an application-consistent copy every 24 hours. Only the replication technology that is specified by the Copy type in the Create copy phase varies among plans. The following table summarizes the default settings:

Table 14 Default file system Service Plan Settings

Setting	Enabled/Not enabled	Default settings	Schedule
Plan Startup	Enabled	Automatic schedule	Recurrence type: Creates a copy every 24 hours, with the first run at midnight (00:00). Recovery Point Objective (RPO): An alert is issued if the copy is not created within the specified time (default is 24 hours).
Application discovery	Enabled	None	Determined by Plan Startup phase.
Application mapping	Enabled	None	Starts when Application discovery phase completes.
Pre-copy script	Not enabled	None	Starts when Application mapping phase completes.

Table 14 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
Create copy	Enabled	<p>Copy type:</p> <ul style="list-style-type: none"> • Bronze • Silver • Gold <p>Also:</p> <ul style="list-style-type: none"> • Storage Ordered Preference: Snapshot, Clone, and Bookmark. Allows you to order, select, or clear storage preferences. By default, all the options are selected. You cannot clear all the preferences, at least one preference must be selected. • Storage Settings: Include RecoverPoint copies in expiration rotation policy —select this option to include RecoverPoint copies when calculating rotations. If you do not select this option, RecoverPoint copies accumulate and remain until the bookmarks for them "fall off" the RecoverPoint appliance. • Configure storage options <ul style="list-style-type: none"> ▪ VPLEX: Allows you to select the preferred cluster for distributed volumes and the preferred arrays for RAID-1 volumes. ▪ VMAX: Allows you to configure the storage pools to create VMAX V2 copy devices for the service plan. You can select or clear the desired VMAX V2 storage pools by expanding the storage 	Starts when Pre-copy script phase completes.

Table 14 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<p>array. If you do not select storage pools, the service plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays. By default, all the configured storage pools are selected. This option is only applicable for VMAX V2 arrays.</p> <ul style="list-style-type: none"> - Wait for Clone Sync: If you select this option, the SnapVX clone copy creation waits for clone synchronization to complete. If you clear this option, the clone copy creation initiates clone synchronization and completes the copy creation without waiting for clone synchronization to complete (clone synchronization runs in the background). By default, this option is disabled. This option is only applicable for VMAX V3 arrays. • Advanced Settings - Allows you to set a retry count and retry interval to retry the failed VSS operation after an interval configured through the retry interval. This is only applicable for Windows applications. The default 	

Table 14 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		retry count value is 3 and the default retry interval value is 0 seconds.	
Post-copy script	Not enabled	None	Starts when Create copy phase completes.
Unmount previous copy	Not enabled	None	Starts when Post-copy script phase completes.
Mount copy (A pre-mount script phase is available for file system service plans)	Not enabled	<p>Mount Copy</p> <ul style="list-style-type: none"> • Mount on Server: Original Host • Mount with access: Read/write • Mount Path: Default Path or Mapped Path • Access mode: Logged access • Copy to Mount: Local (Only for Gold Plans) • Use Dedicated Storage Group: Selected by default • Desired SLO: Select the SLO for the mount copy. This is only applicable for VMAX 3 arrays. • Desired FAST VP: Select the FAST VP policy. This is only applicable for VMAX V2 arrays. • VPLEX Mount option: Select a VPLEX mount option. • Enable VMware cluster mount: If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to 	Starts when Unmount previous copy phase completes.

Table 14 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<p>perform an ESX cluster mount, you can clear this option. This option is supported on VPLEX, XtremIO, VMAX 3/VMAX All Flash, and Unity arrays.</p> <ul style="list-style-type: none"> • Disable VMware SRM: Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later. • Run Filesystem Check: During a mount operation, the AppSync agent checks file system data consistency by executing the <code>fsck</code> command. This operation can be time consuming. You can clear this option to skip file system check during a mount operation. By default, file system check is enabled. <hr/> <p>Note</p> <p>In the case of a restore operation, the <code>Run Filesystem Check</code> option is enabled by default. You cannot disable it.</p> <hr/>	
Post-mount script	Not enabled	None	Starts when Mount copy phase completes.
Unmount copy	Not enabled	None	Starts when Post-mount script phase completes.

VMware service plan settings

The default service plan settings create an application-consistent copy every 24 hours. Only the replication technology, which is specified by the **Copy type** in the Create copy phase, is different from plan to plan.

Table 15 VMware Service Plan - default settings

Setting	Enabled/Not enabled	Default settings	Schedule
Plan Startup	Enabled	Automatic schedule	Recurrence type: Creates a copy every 24 hours, with the first run at midnight (00:00). Recovery Point Objective (RPO): An alert is issued if the copy is not created within the specified time (default is 24 hours).
Application discovery	Enabled	None	Determined by Plan Startup phase
Application mapping	Enabled	None	Starts when Application discovery phase completes
Create copy	Enabled	<ul style="list-style-type: none"> • Copy type is: <ul style="list-style-type: none"> ▪ Bronze ▪ Silver ▪ Gold • Copy Consistency: <ul style="list-style-type: none"> ▪ Virtual machine Consistent with a maximum of 4 simultaneous VM snapshots ▪ Configure VM Snapshots for VMs: Allows you to select virtual machines from the datastores added to the service plan. By default, the Exclude VMs for Snapshot option is enabled. This means that the selected VMs are ignored while taking VMware snapshots during the service plan run. If you select the Include VMs for Snapshot option, only the selected VMs are considered for VMware snapshot creation during the service plan run. 	Starts when Application mapping phase completes

Table 15 VMware Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<ul style="list-style-type: none"> ▪ Include Virtual Machine Disk: Select this checkbox to protect virtual machine disks spanning multiple data stores. By default, this option is not selected. • Storage Ordered Preference: Snapshot, Clone, and Bookmark Allows you to order, select, or clear storage preferences. By default, all the options are selected. You cannot clear all the preferences, at least one preference must be selected. • Configure storage options <ul style="list-style-type: none"> ▪ VPLEX: Allows you to select the preferred cluster for distributed volumes and the preferred arrays for RAID-1 volumes. ▪ VMAX: Allows you to configure the storage pools to create VMAX V2 copy devices for the service plan. You can select or clear the desired VMAX V2 storage pools by expanding the storage array. If you do not select storage pools, the service plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays. By default, all the configured storage pools are selected. This option is only applicable for VMAX V2 arrays. <ul style="list-style-type: none"> – Wait for Clone Sync: If you select this option, the SnapVX clone copy creation waits for clone synchronization to complete. If you clear this option, the clone copy creation initiates clone synchronization and completes the copy creation without waiting for clone synchronization to complete (clone synchronization runs in the background). By default, this option is disabled. This option is only applicable for VMAX V3 arrays. • Expiration: Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations. <hr/> <p>Note</p> <p>If this option is not checked, then RecoverPoint copies will accumulate, and remain until the bookmarks for them "fall off" the RecoverPoint appliance.</p> <hr/>	
Unmount previous copy	Not enabled	None	Starts when Create copy phase completes
Mount copy	Not enabled	<ul style="list-style-type: none"> • No default mount host • Mount using new signature • For RecoverPoint: mount with logged access, and mount local copy (in case of local and remote copy plan) • For VNX file, mount copy with read-only or read/write access for local or remote copies 	Starts when Unmount previous copy phase completes

Table 15 VMware Service Plan - default settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<ul style="list-style-type: none"> • Desired SLO: Select the SLO for the mount copy. This is only applicable for VMAX 3 arrays. • Desired FAST VP: Select the FAST VP policy. This is only applicable for VMAX V2 arrays. • VPLEX Mount option: Select a VPLEX mount option. • Enable VMware cluster mount: If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. This option is supported on VPLEX, XtremIO, VMAX 3/VMAX All Flash, and Unity arrays. • Disable VMware SRM: Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later. 	
Unmount copy	Not enabled	None	Starts when Mount copy phase completes

CHAPTER 5

Protect Microsoft Exchange

This chapter includes the following topics:

- [Overview of Exchange support](#)78
- [Deploying AppSync for Exchange protection](#)..... 80
- [Protect an Exchange database](#) 81
- [Service plan details](#)..... 83
- [Mounting Exchange copies](#)..... 93
- [Overview of Exchange copy restore](#).....99

Overview of Exchange support

Use AppSync to create application-consistent copies of Exchange data.

AppSync support for Microsoft Exchange application includes:

- Protect and manage Microsoft Exchange in standalone and DAG environments (active and passive databases).
- Mount copies to a Windows 2008 or Windows 2012 host for running consistency check or to back up to long-term storage.
- Restore from copies to production Exchange databases in the event that production databases must be brought back to a point-in-time.
- Restore individual mailboxes and mailbox items using Kroll Ontrack®.
- Support for databases on physical hosts, RDMs, and virtual disks on virtual hosts.

Note

AppSync only supports RDMs in physical compatibility mode. RDMs in virtual mode are not supported.

Exchange Server prerequisites

Verify that the Exchange configuration meets supported version requirements for AppSync, including Windows operating system requirements as well as supported service packs for Exchange. The *AppSync Support Matrix* on <https://elabnavigator.emc.com/eln/extendedSupport> is the authoritative source of information on supported software and platforms.

AppSync supports protection and operational recovery of Exchange databases in standalone and DAG configurations including:

- Exchange 2010 mailbox servers on Windows Server 2008 SP2 and Windows Server 2008 R2 or later.
- Exchange 2013 mailbox servers on Windows 2008 R2 SP1 or later.
- Microsoft Exchange 2010 and 2013 Database Availability Groups (DAGs) including active and passive copies.

Support for Exchange on virtual disks

You can protect, mount, and restore Exchange databases residing on VMware RDMs in physical compatibility mode and virtual disks. AppSync supports Full, Copy, and Differential backup types.

During protection:

- For successful mapping, the Virtual Center must be added to the AppSync server and discovery must be performed.
- For successful protection, log files and database files must reside on virtual disks. There cannot be a combination of physical and virtual storage.
- Protection of Exchange databases across virtual machines sharing the same datastore is not supported.
- Virtual Disk is supported for Exchange ESX 5.0 and above.
- AppSync versions 1.6.0.1 and above supports circular logging for Exchange Databases.

Support for Exchange on Hyper-V

In Hyper-V environments, AppSync requires the storage for Exchange to be on iSCSI direct attached devices, Virtual Fiber Channel (NPIV), or SCSI pass-through devices.

SCSI Command Descriptor Block (CDB) filtering must be turned off in the parent partition for SCSI pass-through. It is turned on by default. This is also applicable for databases in DAG configurations.

For Hyper-V SCSI pass-through, the mount host cannot be a Hyper-V host. It has to be a physical host or a virtual machine added with Virtual Fiber Channel adapter or iSCSI direct attached.

AppSync interaction with Microsoft VSS

Microsoft Volume Shadow Copy Service (VSS) is the infrastructure that enables AppSync to create application-aware copies.

When it creates a copy, AppSync coordinates with VSS and Exchange to create a shadow copy. The copy is a point-in-time copy of the volumes that contain the data, logs, and system files for Exchange databases.

AppSync coordinates with VSS and Exchange to quiesce input-output to the databases when creating the copy, and then resume the flow of data after the copy has been created. During a restore, AppSync coordinates with VSS and Exchange to recover the point-in-time shadow copy.

Permissions required by Exchange

Accounts that AppSync uses to work with Exchange require special permissions.

- On Exchange standalone servers, the account must be a domain user account with the Databases role.
- On DAG servers, the account must be a domain user account with the Database and Database Copies roles.
- On a mount host, the user account must be a domain user account that is a member of the local Administrators group.
- The account must have **Log on as a batch job** and **Log on as a service** user rights.
- The account can have the **View-only Organization** role. This role is an optional role applicable only for Microsoft Exchange 2013 if you have public folder mailboxes in the environment. AppSync uses this role to determine the database containing the public folder primary hierarchy mailbox.

AppSync Exchange Interface Service Credentials are required the first time that you access the Exchange server. You are prompted to type two sets of credentials for the AppSync Exchange Interface Service configuration.

AppSync uses the first set of credentials to install and configure the AppSync Exchange Interface service on the Exchange production or mount host. The account must have local administrator privileges. AppSync uses the second set of credentials to run the service. A user must be a domain user with the following Exchange roles:

- Database role for standalone server
- Database and Database Copies roles in DAG environment.

Changes to service plans after upgrade

After an AppSync upgrade, changes to the way service plans operate occur.

Consider the following changes:

- Storage ordered preferences — After an upgrade, all service plans will have their storage ordered preferences replaced with the new style preference. After an upgrade, check the **Create Copy** instructions for each service plan to confirm that the storage ordered preference is correct.
- RecoverPoint copy rotation — After an upgrade, all service plans do not perform replica rotation for RecoverPoint copies. If you want to enable replica rotation for

any service plan that creates RecoverPoint copies, check the **Create Copy** instructions for the service plan and check the checkbox to include RecoverPoint copies for Expiration.

Deploying AppSync for Exchange protection

A summary of steps from deployment of AppSync to setting up Exchange protection.

Procedure

1. Install the AppSync server.
2. In the AppSync console, navigate to **Settings > Storage Infrastructure** and click **Add**.

This adds the storage system where the Exchange mailbox database resides.
3. Navigate to **Settings > Servers** and click **Add**.

This adds the Exchange standalone mailbox servers or DAG member servers as hosts.
4. Navigate to **Copy Management > Microsoft Exchange** and click a server name from the list of Exchange standalone and DAG servers.
5. Enter the credentials to configure and run the AppSync Exchange Interface service.

The Exchange databases are discovered.
6. Subscribe an Exchange database for protection by choosing one of the following options:
 - Protect immediately with **Subscribe to Plan and Run**, which subscribes the database to a service plan and runs the protection immediately for the selected database only. In the case of databases in a DAG, one of the passive databases is protected by default.
 - **Subscribe to Plan**, which subscribes the database to a service plan, but does not run the plan. Protection occurs according to the service plan's schedule.

Discovering Exchange databases

To keep AppSync up-to-date, you should discover databases on the Exchange server when there is creation, deletion, or renaming of databases.

Before you begin

- This operation requires the Data Administrator role in AppSync.
- The AppSync Exchange Interface service must be running.
- If you are changing a standalone Exchange server that is part of AppSync to be part of a DAG, you should first remove the standalone server (along with its copies, if any) from AppSync prior to performing a discovery. During discovery, the erstwhile standalone server is identified to be part of the DAG.

Procedure

1. Navigate to **Copy Management > Microsoft Exchange**.
2. Click an Exchange server to display its databases.
3. Click **Discover Databases** to discover databases for this server.

Removing an Exchange mailbox server

Remove an Exchange mailbox server when there is no longer a need to manage its protection from the AppSync server.

Before you begin

This operation requires the Resource Administrator role in AppSync. There should be no copies of the mailbox server that you want to remove.

Procedure

1. Select **Settings > Servers**.
2. Select the server to remove.
3. Select **Remove > Remove servers only**.
A dialog appears asking for your confirmation.
4. Click **OK** to confirm your action.

Protecting DAG databases in a service plan

AppSync supports protection of Exchange databases that are part of a Database Availability Group (DAG).

When a DAG server is subscribed to an AppSync service plan, it is one of the passive members of the DAG that is selected for protection, by default.

Procedure

- To protect an active DAG database member, select **Active** in the **Copy to Protect** column from the plan **Subscriptions** tab.

Convert a standalone Exchange server to a DAG member

Procedure

1. Remove all the subscriptions and copies of the standalone Exchange server registered with AppSync.
2. Remove the host from the Servers page that was hosting the standalone Exchange server.
3. After the standalone Exchange server is added as a DAG member, add the host back to the AppSync server.

Protect an Exchange database

Protect an Exchange database by subscribing it to an AppSync service plan.

AppSync uses service plans as its protection mechanism for databases. You subscribe a database to a service plan and run the service plan immediately, or schedule the service plan to run at a later time.

- Choose **Subscribe to Plan and Run** when you want to protect a selected database immediately. The service plan is executed for the database alone. In the case of DAG, one of the passive databases is protected by default.
- Choose **Subscribe to Plan** when you want to schedule the protection for later. Protection for databases that are part of the service plan are executed at the scheduled time.

- Choose **Run** from the Service Plan page to run the whole plan immediately. All databases subscribed to the plan are protected.

Protecting an Exchange database immediately

Click **Subscribe to Plan and Run** to add a database to an existing service plan and run the service plan immediately for the selected database alone.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft Exchange**
2. Click an Exchange Mailbox Server or DAG from the list to display its databases.
3. From this list, select a database to protect.

When performing an operation on multiple items, be sure to keep the Shift or Ctrl key depressed.

4. From the **Protect** list, select the appropriate service plan from **Subscribe to Plan and Run**.

In DAG, a passive database is protected by default.

The **Subscribe to Plan and Run** dialog appears displaying the progress through the different phases.

Subscribing an Exchange database to a service plan

Select **Subscribe to Plan** when you want to schedule the protection for later. Protection for all databases that are part of the service plan are executed at the scheduled time.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft Exchange**.
2. Click an Exchange Mailbox Server or DAG from the list to display its databases.
3. From this list, select a database to protect.

Select multiple databases by holding down the Shift or Ctrl keys on your keyboard.

4. From the **Protect** list, select the appropriate service plan from **Subscribe to Plan**.

In DAG, a passive database is protected by default. To change the protection type with another option, specify it from the **Subscriptions** tab of the service plan.

The plan is added to the Plans column for the database.

Unsubscribing a database from a service plan

You can unsubscribe an individual database from a service plan.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft Exchange**
2. Select a mailbox server to display the list of Exchange databases.
3. Select the database to remove from a plan.
 - Select the plan to unsubscribe from **Protect > Unsubscribe from Plan**.
Only plans to which the database is subscribed to are in the popup list.
 - To unsubscribe from all service plans, select **Unsubscribe from Plan > All**.

Expiring a copy on demand

Expiring a copy removes it from the AppSync database and can free up storage, depending on the replication technology and copy state.

Before you begin

This operation requires the Data Administrator role in AppSync.

Expiring a copy that was made with RecoverPoint does not remove the corresponding bookmark from RecoverPoint itself.

Procedure

1. Select **Copy Management > Microsoft Exchange**.
2. Click an Exchange mailbox server to display its databases.
3. Click an Exchange database to display its copies.
4. Select one or more copies to delete.
You can also perform this action from the Service Plan **Copies** tab.
5. Select **Expire**.
Verify that you want to expire the copy you selected and any associated copies listed and confirm.

Creating a database copy from the Copies page

Create a copy of a database by subscribing it to an AppSync Exchange service plan from the **Copies** page.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft Exchange**.
2. Click a mailbox server instance.
3. From the list of Exchange databases, click the database to view its copies.
4. From the **Create a copy using plan** list, select the appropriate service plan.
The service plan runs immediately for the database.

Service plan details

A service plan has the following tabs: Settings, Subscriptions, Copies, and Events.

The **Settings** tab shows the name, description, and status (whether enabled or disabled) of the service plan. Apart from these, the different phases of the plan are

also part of this tab. Click on the appropriate tabs to see information regarding subscriptions, copies created by the plan, and events generated during the service plan run.

Service plan schedule

The schedule of a service plan is set in the **Plan Startup** phase.

The **Startup Type** (scheduled or on demand) determines whether the plan is run manually, or configured to run on a schedule. Options for scheduling when a service plan starts are:

- Specify a recovery point objective (RPO)
 - Set an RPO of 30 minutes or 1, 2, 3, 4, 6, 8, 12, or 24 hours
 - Minutes after the hour are set in 5 minute intervals
 - Default RPO is 24 hours
- Run every day at certain times
 - Select up to two different times during the day
 - Minutes after the hour is in 5 minute intervals
 - There is no default selected
- Run at a certain time on selected days of the week
 - One or more days of the week (up to all seven days) can be selected
 - There is no default day of the week selected. Default time of day is 12:00 AM.
- Run at a certain time on selected days of the month
 - Select one or more days of the month (up to all days)
 - Select one time of day. Available times are at 15 minute intervals.
 - Default is the first day of the month

Control replication storage utilization

When you set up a service plan, set values in the following fields so that you avoid overutilization and depletion of replication storage:

- RPO value in the Plan Startup phase
- Always keep n Copies in the Create copy phase

You should also monitor your storage system with the storage system user interface.

Overriding service plan schedules

You can set individual schedules for databases subscribed to a service plan, overriding the generic recurrence setting.

Before you begin

This operation requires the Service Plan Administrator role in AppSync.

You can override only the settings of the recurrence type already selected for the service plan.

Procedure

1. Navigate to **Service Plans > Microsoft Exchange** and select one of the plans from the list.

2. From the **Settings** tab, select the **Plan Startup** phase.
3. From the **Plan Startup Defaults** pane on the right, note the **Recurrence Type** selected for the plan.

A recurrence type can be set only if **Scheduled** is selected as the **Startup Type**.

4. Click the **Plan Startup Overrides** tab.
5. Set individual schedules for selected databases based on your requirement.

For example, if the recurrence type you selected is **On specified days of the month**, and the rule setting is to **Run at 12:00 AM on the 1st day of every month**, you can override the time and the day for individual databases.

Application discovery

Before creating the copy, AppSync examines the Exchange Mailbox Server to look for changes such as addition, deletion, renaming, or movement of databases.

There are no user settings associated with this phase and it cannot be disabled.

Application mapping

After discovering the application, AppSync maps it to array storage, and protection services such as RecoverPoint.

There are no user settings associated with this phase and it cannot be disabled.

Pre-copy script

To perform preparatory steps before creating a copy, specify a pre-copy script and parameters on a service plan's **Settings** tab.

The pre-copy script runs according to the schedule set in the **Plan Startup** phase. Valid script formats are .bat, .exe, and .ps1 (PowerShell scripts). You can optionally enter credentials to run the script as a specific user. The script runs as Local System by default.

AppSync does not support running of PowerShell scripts directly. You usually must wrap them in a .bat file. The other option is to make the default "Open" on ps1 files `C:\WINDOWS\system32\WindowsPowerShell\v1.0\powershell.exe`. When the PS script runs, you may get an error and you must set an appropriate execution policy.

To run PowerShell commands from scripts:

1. Specify the full pathname to the PowerShell command file in the .bat file:

```
powershell -command C:\PshellCommands.ps1 <nul
```
2. Set the PowerShell execution policy so you can run the script. For example, the first line in the .bat file should look like the following for an unrestricted policy:

```
powershell -command set-executionpolicy unrestricted <nul
```
3. To ensure correct termination of the PowerShell session, add <nul to the end of the line that calls the PowerShell script. The default location of the script is `%ProgramData%\EMC\AppSync\scripts\` on the application host.

Exact parameters depend on the script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Create copy

The Create Copy phase creates a copy based on the replication technology specified in the service plan.

This phase specifies the type of Exchange copy to make, whether to ignore Exchange errors in the Application event log, and if database and logs can reside on the same volume.

Review [Overview: Service Plan](#) for more service plan copy information.

Exchange backup type

AppSync uses VSS to make a consistent online copy at the volume level.

- **Full** creates a copy of the databases in the service plan using VSS, and includes the database files, transaction logs, and checkpoint files. On successful completion of the backup, the logs are truncated.
- **Copy** creates a copy of the databases in the service plan using VSS, which includes the database files, transaction logs, and checkpoint files, as it does using the **Full** option. However, it does not truncate the logs.
- **Differential** copies the entire transaction log volume. A full backup of the selected database must exist or the backup fails. The transaction logs are not truncated on completion of the backup.

Automatic expiration of copies

The automatic expiration value in a service plan's Create Copy phase specifies the maximum desired number of Snap, Clone or Bookmark copies that can exist simultaneously.

When the "Always keep x copies" value is reached, older copies are expired to free storage for the next copy in the rotation. Failed copies are not counted. AppSync does not expire the oldest copy until its replacement has been successfully created. For example, if the number of copies to keep is 7, AppSync does not expire the oldest copy until the 8th copy is created.

AppSync does not expire copies under the following circumstances:

- Mounted copies are not expired.
- A copy that contains the only replica of a database will not be expired.

This setting is independent of the VNX pool policy settings in Unisphere for automatic deletion of oldest snapshots. The service plan administrator should work with the storage administrator to ensure that the VNX pool policy settings will enable the support of the specified number of snapshot copies for the application residing in that pool.

Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations.

Note

If this option is not selected, then RecoverPoint copies will accumulate, and will remain until the bookmarks fall off the RecoverPoint appliance.

Exchange event log errors

Exchange logs certain errors in the Application event log when they occur. These errors indicate a possible corruption of the data in the .edb or log files. They can cause copy creation to fail unless you specifically instruct AppSync to ignore them.

AppSync searches the application event log for these errors every time a copy is created. The first time it runs, AppSync searches the entire log. Subsequent runs search since the last successful run. If there are no existing copies, then AppSync searches the entire log when creating the next copy.

In a service plan's Create copy phase, you can configure AppSync to ignore any or all of these errors.

Table 16 Microsoft Exchange event errors

Error	Meaning
-1018	The database tried and failed to verify information about a particular page in the database.
-1019	Similar to a -1018 error but indicates that the accessed page has returned an invalid page number (usually all zeros) rather than an invalid checksum.
-1022	Indicates major hardware problems, particularly disk subsystem problems. If the database engine requests a page from disk but instead receives an error from the I/O subsystem, a -1022 error results.
447	Indicates corruption in the logical database structure. This accompanies a message stating that the information store terminated abnormally.
448	Indicates an inconsistency or corruption in a table in the Microsoft Jet database. This accompanies a message stating that an information store data inconsistency has been detected in a table.

Database and log layout

Exchange supports environments in which the database and logs reside on the same volume when there is more than one copy of the database in a DAG environment. Service plans can be configured to ignore the restriction that prevents databases and logs from residing on the same volume.

When creating copies of Exchange databases, it is a best practice to restrict a service plan from allowing this configuration because having databases and logs on the same volume limits your restore options. However, you can choose whether service plans with this configuration should succeed or not.

This option is set in the Create Copy phase of a service plan.

When selecting this option, you are limited to restoring the database and logs together. Restore overwrites newer log files. To preserve newer log files for use during recovery, copy them to another volume before restore.

Configure retry on VSS failure

You can configure a VSS retry count in the create copy phase of a service plan. During protection, if a service plan fails because of VSS failures such as VSS timeout issue, the service plan runs the VSS freeze/thaw operation again based on the specified retry count and interval. This option is supported only on Windows applications - File system, Microsoft SQL, and Microsoft Exchange. This option is not used while creating Crash consistent copies of SQL databases.

Note

AppSync does not perform a VSS retry, if the application freeze itself fails. If the application is not in a state to create a copy, AppSync fails to quiesce it, and does not retry the VSS freeze/thaw operation. The application must be brought back to a state where it can be quiesced and then the service plan must be re-run.

Post-copy script

To perform cleanup or other post-copy steps after creating a copy, specify a post-copy script and parameters in a service plan's **Settings** tab.

The script runs on successful completion of the **Create copy** phase. Valid script formats are .bat, .exe, and .ps1 (PowerShell scripts). You can optionally enter credentials to run the script as a specific user. The script runs as Local System by default.

When AppSync creates copies of application items in a service plan, it may break up the application items and place them in separate groups for protection. This action can be for performance reasons (for example, VSS for Exchange and SQL) or because items in a service plan may be protected by different replication technologies. For example, a service plan may contain some application items that are protected by VNX Snapshots and some by RecoverPoint bookmarks. As a result, application items in these groups are protected independently.

When AppSync calls a post-copy script, it passes the copies which were created in the group by calling the script with `-appCopies <APP1> <APP2>`, where APP1 and APP2 are the names of the application items in that grouping.

AppSync does not support running of PowerShell scripts directly. You usually must wrap them in a .bat file. The other option is to make the default "Open" on ps1 files `C:\WINDOWS\system32\WindowsPowerShell\v1.0\powershell.exe`. When the PS script runs, you may get an error and you must set an appropriate execution policy.

To run PowerShell commands from scripts:

1. Specify the full pathname to the PowerShell command file in the .bat file:

```
powershell -command C:\PshellCommands.ps1 <nul
```
2. Set the PowerShell execution policy so you can run the script. For example, the first line in the .bat file should look like the following for an unrestricted policy:

```
powershell -command set-executionpolicy unrestricted <nul
```
3. To ensure correct termination of the PowerShell session, add `<nul` to the end of the line that calls the PowerShell script.

When AppSync runs the post-copy script, it is run for the application items that are part of a group. If there are multiple groups, the post-copy script runs multiple times. When AppSync runs the post-copy script, it passes the list of application items in the

replication group as arguments to the script, right after the user arguments. The syntax is:

```
-applicationCopies <ITEM1> <ITEM2> <ITEM3>
```

where <ITEMx> is the name of the application item that is being protected.

The default location of the script is %ProgramData%\EMC\AppSync\scripts\ on the application host.

Exact parameters depend on the script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Unmount previous copy

The service plan unmounts a previously mounted copy after creating the new copy. The exception is a copy that was mounted on-demand as opposed to by the service plan; in this case the on-demand mounted copy is not unmounted.

There are no user settings associated with this phase and it can be enabled or disabled.

Mount copy

The Mount copy phase mounts the copy. This phase can be enabled or disabled.

The **Mount Copy Defaults** settings for the mount host value, mount path and mount access attributes (read-only or read-write) depend on the service plan. Other mount settings determine where the Exchange metadata files are copied, the type of copy to mount and the RecoverPoint image access type.

- **Mount on Server**
Allows you to choose between Windows hosts you have access to and Original Server. If you have chosen to validate the copies, only servers that have the Exchange Management Tools installed are displayed in the drop down. These servers display on the Microsoft Exchange Protection page as "Utility Host".
- **Mount with access**
Choose the type of access the copy should be mounted with - Read/Write or Read only
- **Mount Path**
 - **Alternate mount path**
The default mount path, when the mount host is the same as the production host, is *SystemDrive:\AppSyncMounts\Production_Server_Name*.
path is represented in the console as %SystemDrive%\AppSyncMounts\%ProdServerName%.

To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). The default path also contains an AppSync variable (ProdServerName) that is delimited with two percent signs (%%).

The following characters are not valid in the path:
< > : " / | ? *
 - **Same as original path** This is another option for the mount path. You can select either of the options.

Note

When performing a DAG mount, do not select the mount path as **Same as original path** if the mount host also happens to be a DAG node having a copy of the database that you are mounting.

- **Copy metadata files to**

By default, the location to copy VSS metadata files is the default path - *SystemDrive:\AppSyncMounts\Production_Server_Name*.

The following characters are not valid in the path:

< > : " / | ? *

If you are backing up the database to another media, you must backup these metadata files as well.

- **Image access options during RecoverPoint mount**

RecoverPoint provides a target-side host application the opportunity to write data to the target-side replication volumes, while still keeping track of source changes.

- **Slow access time, fast image I/O performance (RecoverPoint access mode: Logged Access)**

Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host.

- **Fast access time, Fast after roll image I/O performance (RecoverPoint access mode: Virtual Access with Roll)**

Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing.

- **Fast access time, Slow image I/O performance (RecoverPoint access mode: Virtual Access)**

Provides nearly instant access to the image; it is not intended for heavy processing.

- **Desired Service Level Objective (SLO)**

Additionally if you are using a VMAX 3 array, a setting called Desired Service Level Objective (SLO) is available. The option appears in the Mount wizard and it specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.

- **Copy to mount**

Displayed for service plans that create both a local and remote copy. You can select the type of copy to mount.

- Additionally if you are using a VMAX 3 array, a setting called Desired Service Level Objective (SLO) is available. The option appears in the Mount wizard and it specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.

Mount host overrides in service plan

Select different mount hosts for multiple Exchange servers subscribed to a service plan.

In the Mount copy phase of a service plan, you can specify the host that the copy should be mounted on along with related mount options. If you have multiple servers as part of a service plan, you may want to host their copies on different hosts. You can

specify different mount hosts and other options from the **Mount Copy Overrides** tab of the **Mount copy** phase in a service plan.

Overriding mount hosts in a service plan

If there are multiple registered hosts and they are subscribed to the same plan, you can select a different mount host for each server, overriding the generic mount host settings.

Before you begin

This operation requires the Data Administrator role in AppSync.

Follow these steps when you have multiple hosts subscribed to a plan and you want different mounts hosts for their database copies.

Procedure

1. Navigate to **Service Plans** and select one of the plans from the list.
2. From the **Settings** tab, select **Mount copy** phase.
In the **Mount Copy Defaults** tab, the list of servers include all Exchange servers whose databases are subscribed to this plan. The mount settings display the default settings.
3. To override the default settings, click **Mount Copy Overrides**.
4. Select the server whose mount settings you wish to override and click **Set Overrides**.
5. In the **Override Default Mount Settings** dialog, select options only for those mount settings that you wish to override.
For example, if you want to mount a copy to a different path, you would select the path from the **Mount Path** list. Fields that do not have a selection retain their default settings.
6. Select **OK** to save your changes.
A pencil icon appears in the first column of the server's row whose default mount settings you changed.
7. To revert back to default settings for a server, select the server and click **Use Default Settings**.

Validate copy

Exchange management tools run a consistency check in this phase.

By default, databases and logs are checked sequentially. If the databases are not sharing the same LUN and the mount host has sufficient resources to support parallel consistency checks, use the **In parallel** option. Note that there is a limit of 16 parallel checks that Exchange can handle.

If the consistency check completes successfully, AppSync instructs Exchange to truncate the logs so only the changes that are uncommitted to the database remain.

This phase can be enabled or disabled.

Advanced options for consistency check

AppSync offers advanced options that change how Exchange consistency checks are executed. Enabling these features can impact performance.

- **Minimize log checking**

Choosing this option speeds up the log checking by instructing the consistency checking software to check only those logs that are required to recover the database. Selecting this option improves the performance of the consistency check. If you disable the option, then consistency check will be performed on all of the database's logs.

This command instructs AppSync to check only a subset of the Exchange logs that are included in the copy. The subset of the logs are actually the logs that are required to recover the database. If your backup window is small, you may find this option useful. However, the copy contains logs that have not been checked for consistency. If you attempt to restore the log volume, you may find that some log files are corrupt or the log sequence is not complete. Before restoring the log volume, you should mount the replica and run `eseutil /k Enn` against the log path.

For maximum protection, clear **Minimize log checking**. For maximum performance, select it.

You must also set a working directory, which is where the required log files will be copied for checking.

The **Minimize log checking** option is not available when the consistency method is Differential.

- **Throttle Checking**
Consistency checks can be paused to slow down the IOs during the check. You can specify the number of IOs after which to pause, and the duration of the pause.
- **Skip database validation(.edb file check only for DAG)**
If you select this option, AppSync skips database validation in the case of DAG, if it has:
 - One active and mounted database copy, and at least one passive and healthy database copy
Or
 - Two passive and healthy database copies

Post-mount script

Specify a post-mount script and parameters from the Post-mount script option in the **Settings** tab of a service plan.

The script runs on successful completion of the mount copy or mount with recovery phase. This script is typically used for backup.

From the **Server** list, select the server on which to run the script. You can optionally run it on a registered host other than the mount host, and enter credentials to run the script as a specific user.

The default location of the script is `%ProgramData%\EMC\AppSync\scripts\` on the application host.

Exact parameters depend on your script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Unmount copy

The final phase in the service plan unmounts the copy. All the mounted databases are shut down as part of this phase.

This phase is disabled if the **Unmount previous copy** phase is enabled. There are no user settings associated with this phase.

Mounting Exchange copies

AppSync can mount a copy on-demand, or as part of a plan.

Copies created on a standalone production Exchange server can be mounted to:

- An alternate host in the same location as the production host.
- An alternate host in a new location. You specify mount option by adding an alternate path to the start of the path.
- The production host in an alternate location.

Copies created in a DAG can be mounted to:

- An alternate host
- A server in another DAG
- Another server in the same DAG

Note

- Copies cannot be mounted to the same DAG server on which the copy was created.
 - A single mount host with Exchange 2013 Management Tools can be used to run consistency check for Exchange 2010 and Exchange 2013 copies.
-

Mount and restore limitations

Limitations to mount and restore or Exchange copies appear in the following list:

- When the root drive letter has mount points on it and they are all included in the same plan, mounts and restores are likely to fail. For instance, if the log and system files are on L:\ and the mailbox stores are on L:\SG1DBMP (where SG1DBMP is a mount point), mounts and restores fail.
- In Windows 2012 and later environments, when doing a restore, the data on LUNs is overwritten even if the volume is in use. This action differs from other Windows platforms in which AppSync displays a warning if the LUN is in use. Since restores overwrite everything, be sure that there is no other data on that volume and the volume is not in use.

Mounting an Exchange copy on-demand

You can initiate an on-demand mount of an Exchange copy from a copy or database.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Select **Recover > Mount a Copy** in the **Exchange Databases** page.
Alternatively, from the **Copies** page or **Copies** tab of the service plan, select a copy and click **Mount**.
2. Use the **Copies** or **Service Plan** filters to select the appropriate copy to mount.
The Copies list is refreshed based on the filters selected.
3. Select the copy to mount.
For a RecoverPoint copy, you also have the option to select a copy based on a specific time. Click **Select a point in time** to select a copy with a specific time stamp. The time displayed is the console's time. If the console is in a different time zone from the server, specify the time as per the server's time zone to mount the copy.
4. In the **Mount Additional Copies** page, select one or more additional copies to mount.
The copies listed here are of other databases that were protected at the same time and on the same host as the copy you selected in the previous step.
5. In the **Validate Copy** page, select **Yes** to validate the copies and specify validation options. See Mount validation options for details.
6. On the **Mount Options** page, select the mount options.
 - For VMAX 3 arrays, select the desired Service Level Objective (SLO) for the mount copy.

Note

The SLO values are dynamically fetched from the VMAX 3 arrays, and only the unique values are displayed.

- For VMAX 2 arrays, select the desired FAST VP policy for the mount copy.
- **Mount on Server**
Allows you to choose between Windows hosts you have access to and Original Server. If you have chosen to validate the copies, only servers that have the Exchange Management Tools installed are displayed in the drop down. These servers display on the Microsoft Exchange Protection page as "Utility Host".
- **Mount with access**
Choose the type of access the copy should be mounted with - Read/Write or Read only
- **Mount Path**
 - If you specify non-default mount path, the drive specified for mount can not be a clustered disk.
 - **Alternate mount path**
The default mount path, when the mount host is the same as the production host, is *SystemDrive:\AppSyncMounts\Production_Server_Name*.
path is represented in the console as %SystemDrive%\AppSyncMounts\%%ProdServerName%%.
To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). The default

path also contains an AppSync variable (ProdServerName) that is delimited with two percent signs (%%).

The following characters are not valid in the path:

< > : " / | ? *

- **Same as original path**

This is another option for the mount path. You can select either of the options.

Note

When performing a DAG mount, do not select the mount path as **Same as original path** if the mount host also happens to be a DAG node having a copy of the database that you are mounting.

- **Copy metadata files to**

By default, the location to copy VSS metadata files is the default path - *SystemDrive:\AppSyncMounts\Production_Server_Name*.

The following characters are not valid in the path:

< > : " / | ? *

If you are backing up the database to another media, you must backup these metadata files as well.

- **Image access options during RecoverPoint mount**

RecoverPoint provides a target-side host application the opportunity to write data to the target-side replication volumes, while still keeping track of source changes.

- **Slow access time, fast image I/O performance (RecoverPoint access mode: Logged Access)**

Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host.

- **Fast access time, Fast after roll image I/O performance (RecoverPoint access mode: Virtual Access with Roll)**

Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing.

- **Fast access time, Slow image I/O performance (RecoverPoint access mode: Virtual Access)**

Provides nearly instant access to the image; it is not intended for heavy processing.

- **VPLEX Mount options**

- **Native array:** Use this option if you want to mount the copy as native array volumes.

- **VPLEX virtual volume mount:** Use this option if you want to mount the copy as VPLEX virtual volumes.

- **Enable VMware cluster mount:** Clear this option if you do not want to perform an ESX cluster mount. By default, this option is enabled.

7. In the **Configure AppSync Exchange Interface Service** page, provide the credentials to configure the service on the mount server.

This page is displayed only if you chose to validate the copy if the service is not configured.

8. Review the **Summary** and click **Finish** to mount the copy.
9. In the **Results** page, select **View Details** to see progress of the different phases that are part of mounting a copy.

The last phase completed is displayed at the bottom of the list.

Validation options for a mount copy

Validation for differential backup copies is not supported.

Validate database and logs

When you create a replica of one or more Microsoft Exchange databases, you should mount the replica and test it for consistency. If you choose to automatically mount the replica to an alternate host once it has been created, you should run a consistency check on the replica. The options to validate are:

- **Sequentially** — Run tests on one database at a time in order (serial mode). Select this option if you have several Exchange databases on one LUN.
- **In Parallel** — Run tests on several databases simultaneously (parallel mode).

Minimize log checking

By selecting Minimize log checking, AppSync checks a subset of the Exchange logs that are included in the replica. If your backup window is small, you may find this option useful. However, the replica may contain logs that have not been checked for consistency.

For maximum protection, clear Minimize log checking. For maximum performance, select it.

Working directory — This field allows you to specify the directory to which the relevant log files will be moved in order to run the check, since a consistency check can only be run on all logs in a single directory.

Throttle validation

Select this to throttle the I/Os during a consistency check. This option is for advanced users and typically should not be selected unless you are working with Dell EMC Support to resolve an issue related to I/O throughput. Typically, the throttling option is not required.

If you choose to throttle I/Os, you have the following two options.

- **Pause after I/O count of: 100** — This option allows you to choose how many I/Os can occur between pauses. You can choose any value between 100 and 10,000 I/Os.
- **Duration of pause (in milliseconds): 1000** — You can specify the duration of the pause in milliseconds. 1000 milliseconds = 1 second. If this option is not available, the pause will be one second long.)

Skip database validation(.edb file check only for DAG)

If you select this option, AppSync skips database validation in the case of DAG.

Unmounting an Exchange copy

When you select a copy to unmount, other copies that were mounted along with the selected copy will also be unmounted.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can unmount a copy only from a list of copies made for a database.

Procedure

1. Navigate to the **Copies** page from the **Data Protection** or **Service Plan** pages:
 - **Copy Management > Microsoft Exchange** > select the Exchange Mailbox Server that hosts the database, then select the database with the copy to unmount.
 - **Service Plans > Microsoft Exchange** > select a service plan, then select the **Copies** tab.

2. From the list of copies, select the copy and click **Unmount**.

The **Unmount Confirmation** dialog displays copies of other databases that were mounted along with the selected copy to be unmounted.

3. Click **Yes** to confirm the unmount of all the copies shown in the dialog.

The **Unmount** window displays the progress of the unmount operation. All copies associated with the selected copy will be unmounted.

Path mapping

The path mapping option mounts the copy to a host using a path mapping table set to user-defined locations. When you use a path mapping table, you have more control over where data is located.

You must specify the path where you want to mount a specific file system. You must provide a path map where the source file system and the target mount point is specified.

The following is a sample path mapping table for UNIX.

Source file system	Target mount path
/opt/emc/db1	/opt/emc/copies/db1
/opt/emc/db2	/opt/emc/test/db2
/opt/emc/db3	/opt/emc/reserved/db3
/prd/db1	/tst/db1

The following is a sample path mapping table for Windows.

The first two target paths, G:\ and H:\ drives must already be available on the mount host. That is, the root drive for the mount path must pre-exist before attempting a mount.

Source file system	Target mount path
D:\Test1	G:\Test1
E:\	H:\Test2
F:\Test3	I:\
L:\	N:\

Note

- If a target path is not provided for a source path, then it is mounted to a path same as the source path on the mount host.
 - Ensure that you type in the absolute mount path on the target host. If the path is invalid, mount fails.
 - Mount copy overrides is unavailable, if you select the mount path as Mapped path.
 - For Windows, if one of the entered path is invalid, VSS import fails. Therefore, the entire mount fails. Partial failed scenarios are not supported for Windows mount.
 - For Windows and NFS file systems on Unix, nested target mount points are not supported.
 - Path mapping is not supported for Oracle ASM disk groups.
 - Path mapping is not supported for Oracle NFS RAC database.
 - Path Mapping is not applicable to metadata paths for Microsoft Exchange and Microsoft SQL Server.
-

Specify path mapping settings

You can specify the path where you want to mount a specific copy.

Procedure

1. On the AppSync console, go to the **Copies** page to display available copies.
2. Select a copy to mount, and then click **Mount**.
The Select Copy to Mount page of the Mount wizard launches.
3. Select the copy to mount, and click **Next**.
The Specify Mount Settings page appears.
4. On the **Specify Mount Settings** page:
 - a. Select the mount host.
 - b. From the **Mount Path** list, select **Mapped Path**.
The Path Mapping Settings link appears.
5. Click on the link to open the Path Mapping Settings window.
6. From the **Select Source Host** list, select a host.
All the file systems on the selected host are displayed in the source path column.
7. Specify the target path.
8. Click **Save** to save your settings.
If you want to set the target path for a file system on another source host, repeat steps 6 to 8.
9. Click **Reset**, to clear all the entered target paths for the selected source host.
10. Click **OK** to exit the Path Mapping window.

Note

If you change the path mapping settings, the earlier saved path mapping settings is not valid and the new path mapping settings takes precedence. Therefore, ensure that you save the path mapping settings for all the hosts before changing it.

Overview of Exchange copy restore

Learn about Exchange restore features along with associated storage copy levels.

With AppSync you can restore the following objects:

- A database with its logs.
- A database .edb file.
- Only the logs for a database.
- An active or passive database (in conjunction with any one of the three points already mentioned), if the server is a member of a DAG (Database Availability Group).

AppSync restores VNX/VMAX copies at the LUN level, VNXe copies at the LUN group level, and Unity copies at the consistency group level. In a RecoverPoint environment, restore is at the consistency group level.

Note

Ensure that no virtual machine snapshots are present before protecting a datastore. If virtual machine snapshots are present, protection succeeds, but AppSync fails to perform a file or virtual machine restore.

Affected entities during restore

When restoring from a copy, you may be prompted to restore items in addition to the ones you selected.

An affected entity is data that resides on your production host that unintentionally becomes part of a replica because of its proximity to the data you intend to protect. You can prevent affected entity situations by properly planning your data layout based on replica granularity. The granularity of a replica depends upon the environment.

If there are *affected entities* in your underlying storage configuration, the Restore Wizard notifies you of these items. The following scenarios produce *affected entities* that require you to acknowledge that additional items will be restored:

- For RecoverPoint, if the databases are in the same consistency group they become *affected entities* when the other database is protected.
- For VNXe, if the databases are in the same LUN group, they become affected entities when another database in the group is protected.
- For Unity, if the databases are in the same consistency group, they become affected entities when another database in the group is protected.
- For VMAX, VNX, VNXe, Unity, or XtremIO, if the databases are on the same LUN they become *affected entities* when the other database is protected.
- For VMware virtual disks, since restore involves a datastore, restore of all applications residing on the same datastore (virtual disks on the same datastore) are also affected entities.

If the affected entity was protected along with the database selected for restore, AppSync restores it. Any other database that was not protected but is an affected entity is overwritten. AppSync calculates affected entities for the consistency groups, LUN groups or LUNs of the database that is selected for restore. If the affected databases partially reside on other consistency groups, LUN groups or LUNs, AppSync does not calculate affected entities on those consistency groups, LUN groups or LUNs.

Affected entities are calculated on the basis of restore granularity. If both data and logs are selected for restore, then affected entities are calculated for all the consistency groups, LUN groups, or LUNs on which the database resides. If only data or only log restore is selected, then the affected entities are only calculated for the selected component's consistency group, LUN Group, or LUN only.

If the database data and log components reside on the same consistency group, LUN group, or LUN, the option to restore only logs or restore only data is not available. You have the option only to restore data and logs. The only exception to this scenario is when you perform a differential copy restore.

Since restore involves a datastore with VMware virtual disks, restore of all applications residing on the same datastore (virtual disks on the same datastore) are also affected entities.

Restoring from an Exchange copy

You can perform a restore of an Exchange copy from a copy or a database.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Select **Recover > Restore > Databases and logs** from the action buttons at the bottom of the page.
2. Select the copy to restore, and click **Next**.

The **Configure Storage Options** screen appears.

- **Wait for mirror rebuild to complete** - This option is applicable for VPLEX Snap copies whose production data resides on local or distributed RAID-1 volumes. By default, this option is selected.
- **Disable VMware SRM** - Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later.
- **Perform device restore in background** - Allows you to optimize restore of VMAX devices. If you select this option, AppSync restore operation does not wait for VMAX track synchronization to complete. The production application is available instantly.

3. Click **Next**.
4. From the **Select Options** page, select the appropriate restore options and click **Next**.

Option	Description
Both data and logs	Available when restoring from full and copy backup.
Data	Available when restoring from full and copy backup.

Option	Description
Logs	Available when restoring from full, copy, and differential backup.
Recover and mount databases after restore	Leave this selected. If not selected, you must recover the database manually.
Allow AppSync to activate databases	Available for DAG only. If selected, the database copy is activated on the server prior to restoring it. <hr/> Note Exchange restores only to the active copy and AppSync restores to the server that created the copy. <hr/>

The **Restore Warnings** page is displayed. This page may be displayed if the selected copy has affected entities.

5. Read the warning messages for the affected databases.

In case of Exchange 2013, AppSync attempts to determine if the selected database or any of the affected entities contain the public folder primary hierarchy mailbox. If a match is found, you see an error message that the restore operation cannot continue.

To determine where the public folder primary hierarchy mailbox resides, AppSync requires your user profile to have an additional Exchange role - the View-only Organization role. If your account does not have this role, you see a warning message.

6. Select the checkbox to indicate your agreement to restore other entities along with the selected copy.

You must manually unmount the databases that will be overwritten.

7. Review the **Summary** page and click **Finish** to restore the copy.
8. In the **Results** page, click **View Details** to see progress of the different phases that are part of restoring a copy.

The last phase completed is displayed at the bottom of the list.

Recovering an Exchange database manually

Perform a manual recovery when you have not selected the **Recover and mount the databases after restore** option in the Restore wizard.

Before you begin

When you are recovering just a database file, verify that the transaction log files needed for recovery are present. An unbroken sequence is required. To determine the minimum required range of logs, run the following command against each database after the restore and before running recovery: `ESEUTIL /mh <database name>`. Look for the Log Required information in the ESEUTIL output.

If the database is the active copy, it must first be unmounted in order to run the ESEUTIL command successfully.

Procedure

1. Delete the checkpoint file (Enn.chk).

This is optional.

2. Delete the restore.env file (EnnRESTORE.env).
3. Recover the databases manually in soft recovery mode using the ESEUTIL command.

```
eseutil /r E<nn> /l <logpath> /s <chkpt file path> /d
<database path>
```

4. Use Exchange Management Console to mount all the restored databases.

Partial restore

In a partial restore, you restore data alone or restore data and then restore the logs separately.

Before you perform a partial restore, ensure that the database layout fulfills some conditions.

Partial restore considerations

In a RecoverPoint environment, the granularity of restore is at the consistency group level. When you restore a database from a bookmark, any bookmarks that are newer than the bookmark being restored are deleted. The corresponding application copies are also deleted. The following best practices are recommended:

- The database and logs must reside in different consistency groups. If you have data and logs for an Exchange database in the same consistency group, partial restore is not supported.
- The logs should be restored from a newer Differential backup copy. AppSync does not support restoring just the logs from a Full or Copy backup in a RecoverPoint environment.

In a VMAX/VNX environment, the database and logs must reside on different LUNs.

Restore data

Restore data from a Full or Copy backup. You can restore data only to preserve the logs that are on the production host.

In the Restore wizard, restore data from the most recent copy and select the **Recover and mount the databases after restore** option.

Restore logs

Restore data from a Full or Copy backup and then restore the logs from a later copy to make the copy current.

Restoring a copy from the logs is a two-step process. Run the Restore wizard and select a full backup copy to restore only data. Do not opt to **Recover and mount the databases after restore** in this run.

Run the Restore wizard again and select a backup copy (a differential backup in case of RecoverPoint) to restore only the logs. This time, select the **Recover and mount the databases after restore** option. This copy must be later than the backup copy that you selected during the first run.

Note

If the restore operation includes restoring logs, the restore overwrites any logs that are created since the copy was created. Therefore, after the restore, the database reflects the point in time when the copy was created. If you want to preserve logs that are created since the copy, restore only the databases, preventing AppSync from restoring older logs over the newer logs. You can also make a copy of the current log files on another volume.

Restoring logs from crash-consistent (APiT) copy

Restore an any point in time (APiT) copy using logs.

Before you begin

This is applicable only in a RecoverPoint environment.

Note

Restoring logs from a crash-consistent copy is not a recommended practice as the backup is not taken with the Exchange writer. However, the option can be used to minimize data loss when application consistent copies for that time window are not available to restore from.

Procedure

1. Restore a database from an application-consistent copy without recovering it.
2. Mount a copy from a newer point in time.
3. Copy the newer log files to the production log volume.
4. Use `ESEUTIL /k Enn` (Enn is the log prefix for the database) to check the logs, then recover and mount the database.

Restoring a deleted Exchange database

AppSync can restore a database even if it is deleted from Exchange in standalone and DAG environments.

Before you begin

- If you deleted the database files and created an empty database, dismount the database and delete its files. The database that you are restoring should not have data and log files at the original location where they were when the empty database was created. The log file signatures will not match those in the AppSync copy and the restore will fail.
- If you completely remove the database and recreate it, the database name and its file path and names should be exactly the same as those in your AppSync copy. If you do not recreate the deleted database, AppSync recreates it.
- In a DAG environment:
 - There should be no active or passive copies of the deleted DAG database.
 - AppSync recreates and restores only the active database copy to the server that created the AppSync copy. After the database has been restored and recovered, you can recreate the DAG passive copies.

If you have not selected the **Recover and mount the databases after restore** option in the Restore wizard, perform the following manual steps to recover the database.

Procedure

1. Copy the required logs from `_restoredLogs` directory to the directory where the current logs reside.
2. If the log file prefix changed, rename the required log files to use the new prefix.
3. Delete the `E<nn>restore.env` file.
4. Recover the databases manually in soft recovery mode using the `ESEUTIL` command.

```
eseutil /r E<nn> /l <logpath> /s <chkpt file path> /d  
<database path>
```

5. Delete the `_restoredLogs` directory that should be empty after the database is recovered.

CHAPTER 6

Protect SQL Server

This chapter includes the following topics:

- [Overview of SQL Server support](#)..... 106
- [Support for AlwaysOn Availability Groups](#)..... 110
- [SQL Server transaction log backup](#)..... 110
- [Considerations for working with SQL Server in a cluster](#)..... 119
- [SQL Server User Databases folder](#)..... 122
- [Protect a SQL Database](#)..... 123
- [Mount considerations for SQL Server](#)..... 138
- [SQL Server database restore overview](#)..... 148

Overview of SQL Server support

Use AppSync to create and manage application-consistent or crash-consistent copies of Microsoft SQL Server databases.

AppSync support for Microsoft SQL applications includes:

- AlwaysOn Availability Group support.
- Dynamic discovery of user databases during service plan run.
- Support for databases on physical hosts, RDMs, and virtual disks on virtual hosts.

Note

AppSync only supports RDMs in physical compatibility mode. There is no support for RDMs in virtual mode.

-
- Protection for standalone and clustered production SQL Server instances.
 - Mount on a standalone server or cluster nodes of alternate cluster or production cluster as non-clustered resource. Mount with recovery on an alternate clustered instance belonging to an alternate or production failover cluster.
 - Support for Repurposing SQL server database copies.

SQL Server prerequisites

Verify that the SQL Server configuration meets the prerequisites that are listed here. The AppSync Support Matrix on <https://elabnavigator.emc.com/eln/modernHomeDataProtection> is the authoritative source of information on supported software and platforms.

- SQL Server database and its transaction logs must be on disks in the same storage array.
- The SQL Server database must be online during replication.
- Full-text catalogs that are associated with a file group are included as part of a replica of that file group. If the full-text catalogs are not located on supported storage, protection fails. When using full-text catalogs, ensure that the storage device where the catalog is located does not include data that is not related to the database.
- If you want to recover databases from the mounted copy, the mount host must have an installed SQL Server. It is recommended to use the same version of SQL Server on the production and mount hosts.
- When mounting with recovery to an alternate clustered instance, you must add all the owner nodes of the SQL Server clustered instance to AppSync.
- When restoring SQL Server clustered databases, you must add all the owner nodes of the SQL Server clustered instance to AppSync.
- In Hyper-V environments, AppSync requires the storage for SQL database and log files to be on iSCSI direct attached devices, Virtual Fiber Channel (NPIV), or SCSI pass-through devices. SCSI Command Descriptor Block (CDB) filtering must be turned off in the parent partition for SCSI pass-through. It is turned on by default. This is also applicable for SQL cluster servers.

- For Hyper-V SCSI pass-through, the mount host cannot be a Hyper-V host. It has to be a physical host or a virtual machine added with Virtual Fiber Channel adapter or iSCSI direct attached.
- System databases are not supported.
- SQL Server database snapshots are not discovered.
- Creating a copy of a database mirror is not supported. Trying to do so results in an error that the database is not in a valid state.

SQL Server supported configurations

AppSync provides support for the SQL configurations listed here.

- Multiple SQL Server databases can exist on the same volume, or across multiple volumes. However, it is best practice to not mix databases from more than one SQL Server instance on a volume.
- Multiple SQL Server instances can coexist on the same host.

Considerations for SQL Server in a VMware environment

You can protect, mount and restore SQL Server standalone and clustered databases residing on VMware virtual disks (VDisks) or raw device mappings (RDMs).

During protection of SQL on RDM or VDisks:

- For successful mapping, the Virtual Center must be added to the AppSync server and discovery must be performed.
- For successful protection, log files and database files must reside on virtual disks. There cannot be a combination of physical and virtual storage.
- Protection of SQL Server databases across virtual machines sharing the same datastore is not supported.
 - If multiple SQL databases from different virtual machines on the same datastore are subscribed to the same service plan, protection works, but multiple copies of the same datastore are created. This causes issues during mount and restore operations.
 - If multiple SQL databases from different virtual machines on the same datastore are subscribed to a different service plan, protection, mount, and restore works. However, restore causes problems for databases in other virtual machines that are hosted on the same datastore. It is recommended that multiple SQL databases on different virtual machines be hosted on different datastores.
- Multiple SQL databases from the same host can reside on the same or different virtual disk that are on the same VMFS datastore, and protected by the same service plan. However, restore of one database on this VMFS datastore restores other databases on the same datastore, which are listed as affected entities during restore.
- When restoring SQL Server clustered databases, you must add all the owner nodes of the SQL Server clustered instance to AppSync.
- If the mount host is a virtual machine, the Virtual Center must be registered with AppSync. This is required to mount RDMs.
- For virtual disks:
 - Production mount is not supported, if the ESX host version is earlier than 5.0.

- Non-persistent virtual disks are not supported.
- For datastore and virtual disk mounts on ESXi 5.x and RecoverPoint 4.1.7.7 environments, disable hardware acceleration to ensure successful virtual access type mounts. For more details, refer the relevant VMware Knowledge Base article. For Hyper-V SCSI pass-through, the mount host cannot be a Hyper-V host. It must be a physical host or virtual machine with NPIV or iSCSI direct attached.
- If databases reside on raw device mappings in VMWare environments, the SQL Server cluster nodes must reside across different ESXi. This is a requirement from VMWare. For database on virtual disks, SQL Server cluster nodes can reside on the same ESX server.
- AppSync only supports RDMs in physical compatibility mode. There is no support for RDMs in virtual compatibility mode.
- For VMware virtual disks, restore of all applications residing on the same datastore (virtual disks on the same datastore) are also affected entities because restore involves a datastore.

Required permissions and rights

Users require certain permissions and rights to protect databases in a SQL Server environment. The user account must be configured to use either SQL Server authentication or Windows authentication.

The Windows user account can either be a member of the local Administrators group or a non-Administrator account with the restrictions outlined next.

In SQL Server 2012, the default virtual account used in the service startup account of the database engine does not have the requisite file system permissions for accessing the mounted or restored database files. Therefore, recovery of SQL databases may fail. To overcome this, you must change the service startup account for the SQL Server database engine to use a domain user account with appropriate privileges and permissions.

Setting up permissions for a domain account that does not have local administrator privileges

Additional setup is required if you need to use a domain account that does not have local administrator privileges.

Procedure

1. Create a Windows domain user (for example, sqluser) and make it part of the Domain Users group.
2. In SQL Server Management Studio, create a new login, using the newly created domain account and select Windows authentication.
3. In the **General** page, select **master** as the default database.
4. In the **Server Roles** page, select **sysadmin** and **public**.
5. In the **User Mapping** page, set the database role membership to **public**.
6. Add the user to each SQL Server instance on which this user needs access:
 - a. On the domain controller: On the hosts added to the domain: **Start > Programs > Administrative Tools > Domain Controller Security Policy**
On the hosts added to the domain: **Start > Programs > Administrative Tools > Local Security Policy**
 - b. Access security settings and allow login locally (**Security Settings > Local Policies > User Rights Assignment > Allow log on locally**)

- c. Add the user (the example is sqluser) you created earlier.
7. Log in to the domain controller machine for each host added to that domain that uses AppSync and set the Security policy.
8. Grant this user read and write permissions on the directory where the AppSync plug-in is installed (typically C:\Program Files\EMC\AppSync Host Plug-in).
9. Use this user from AppSync when you configure protection or perform other actions that require access to SQL Server.
10. At the time of restore, if you select the option to back up the transaction logs to a file, the user must have rights to the target directory.

Setting permissions for a local, non-administrator user

A user account that does not have local administrator privileges needs certain permissions before it can be used to access SQL Server from AppSync.

Procedure

1. Create a Windows user and make it part of the Users group.
2. In SQL Server Management Studio, create a new login, using the newly created account. For the authentication type, select Windows authentication.
3. In the **Server Roles** page, select **sysadmin** and **public**.
4. In the **User Mapping** page, set the database role membership to **public**.
5. Add the user to each SQL Server instance on which this user needs access:
 - a. On the host running the plug-in, set the security policy. On the domain controller, run **Start > Programs > Administrative Tools > Local Security Policy**.

On the hosts added to the domain: **Start > Programs > Tools > Local Security Policy**.
 - b. Access security settings and allow login locally (**Security Settings > Local Policies > User Rights Assignment > Allow log on locally**).
 - c. Add the user (the example is sqluser) you created earlier.
6. Grant this user read and write permissions on the folder where the AppSync plug-in is installed.
7. If you select the restore option to back up the transaction logs to a file, the user must have rights to the target directory.

Update login credentials for a SQL Server instance

If the credentials for a SQL Server instance have changed, you need to update them in AppSync.

Before you begin

This operation requires the Data Administrator role in AppSync. In addition, you should know the new credentials for the SQL Server instance.

Procedure

1. Select **Copy Management**.
2. Select **Microsoft SQL Server**.
3. Select an instance.

4. Click **Connection Settings** from the row of buttons below.
5. Enter the SQL Server credentials.

The credentials can be a Windows user or a SQL user with required privileges.

Support for AlwaysOn Availability Groups

The Availability Groups can be part of clustered and non-clustered SQL Server instances installed on AlwaysOn Failover clusters.

AppSync supports Full or Copy backups of primary databases and Copy backups of secondary databases. The **Auto Switch to Copy** option in the SQL Server service plan's **Create copy** phase allows you to switch from **Full** to **Copy** for secondary databases.

Special considerations when you are using AlwaysOn Availability Groups:

- To protect secondary databases, they must be read-only. The `ReadableSecondary` option in the SQL Server Management Studio must be set to `Yes`; `Read-intent only` is not supported.
- Do not use the original path when mounting an AppSync copy to a node in the same cluster if that node hosts a copy of the database.
- It is recommended to protect replicas in the Synchronous-commit mode.
- The considerations for working with SQL Server in a cluster also apply to Availability Groups. See [Considerations for SQL in a cluster](#).
- Multi-subnets are supported for AlwaysOn Availability Groups as long as none of the database copies belong to a clustered SQL Server instance.

SQL Server transaction log backup

AppSync 2.1 and above supports SQL Server transaction log backup. Get key considerations as well as restrictions before implementing your backups.

Every SQL Server database has a transaction log. Write the log backups to Dell EMC storage systems that are supported by AppSync so you can create copies of the log backup volume. If you back up logs for databases in a failover cluster environment, use shared storage or a network share so the log backups are written to the same location.

You can use transaction log backups during recovery of a production database or when making a copy of a production database. Depending on the database recovery model, the transaction log can become full. To prevent the accumulation of logs, regularly run transaction log backups with truncation enabled.

AppSync can backup transaction logs in AlwaysOn Availability Group (AAG) environments. It can back up primary or secondary database copies. If truncation is enabled, to initiate truncation, back up either the primary or secondary database transaction log.

Transaction log backups are supported using only streaming back up; they are not supported using VSS hardware snapshot technology. You can use AppSync to back up transaction logs to a file. The file can be written to a local volume or network share using a UNC path.

Note

AppSync supports UNC path of a network share only if both the machines are in the same domain.

Restrictions

- To back up a transaction log, the database recovery model must be either “Full” or “Bulk-logged.” AppSync skips backing up the log for any database with the simple recovery model.
- To create any log backups with log truncation, first create at least one full database backup.
- To truncate transaction logs, AppSync must have a Full database backup copy.
- Subscribe a database to only one service plan with log backup enabled.
- To truncate logs in an AAG environment, subscribe only one copy of a database to a service plan that is configured for Full database backups and transaction log backups with log truncation.
- To back up transaction logs for databases that belong to an availability group, alter the schedule so that different copies of the database are not backed up at the same time.

Related topics

- [Configure SQL Server transaction log backup](#)
- [Run log backup on demand](#)
- [View log backups for a service plan](#)
- [View log backup list for a single database](#)

Configure SQL Server transaction log backup

Learn how to enable transaction log backups for an SQL Server service plan, by selecting the **Enable log backup** checkbox on the Create Copy options page of the AppSync console.

Before you begin

Verify that the user account you select for backups has full control of the directory. This account is the user account that you entered when discovering databases. Also verify that the account configured for the SQL Server Database Engine Service of the SQL Server instance being protected has full control of the backup directory.

After you select this checkbox, the **Transaction Log Backup Options** dialog box is enabled where you can customize when and how to run log backups and where to write the log backup files. Transaction log backups run sequentially.

Figure 2 Transaction Log Backup Options dialog box

Transaction Log Backup Options

Schedule Immediately after database backup

Every

Log Backup schedule is disabled when service plan is On demand.

Backup path

(DefaultPath: SQL Server default backup directory)

Free space on the volume

Backup group size

Truncate the logs

Checksum the backup

Compression

Expiration of Log Backups

Minimum Retention Hours

Procedure

1. Use the **Schedule** field to set log backup runs.

You can select to run the transaction log backup once, immediately after a database backup is run, or you can select to schedule log backups. You can set log backup schedules to run every 15 or 30 minutes or every 1 to 24 hours. If you set a service plan to run on demand, you disable the log backup schedule.

When you schedule log backups to run at a specified interval, the service plan will have two schedules associated with it: one for database backups and one for log backups. The log backup is referred to as the alternate schedule. Log backups run between database backups using the alternate schedule.

2. Edit the **Backup path** field to set the location where AppSync writes log backup files.

Default path uses the SQL Server instance default backup directory. You can also enter a path on any volume on the server or the UNC path of a network share.

AppSync creates the directory if it does not exist. It creates a subdirectory using the name of the SQL Server instance. The log backup file names have the following format: `EMC_AppSync_databasename_timestamp.trn`, for example, `EMC_AppSync_AdventureWorks_2014_10_18_15_38_32.trn`

3. Use the **Free space on volume** field to set a value to verify the amount of free space on the volume before AppSync begins a transaction log backup.

If not enough free space is available, an alert is generated and the log backup fails.

4. Use the **Backup group size** field to control the number of parallel log backups for an SQL Server instance. The default value is 5, (AppSync runs log backups in groups of five).

For example, if you subscribe 15 databases from the same SQL Server instance to a service plan, three log backups will run in parallel. Transaction log backups run sequentially.

5. Select or clear the **Truncate the logs** field when you create Full database backups.

This field is checked by default when you select Full backup type, and it is disabled when you select Copy . To protect secondary databases, truncate logs, select **Auto switch to Copy** and **Truncate the logs**.

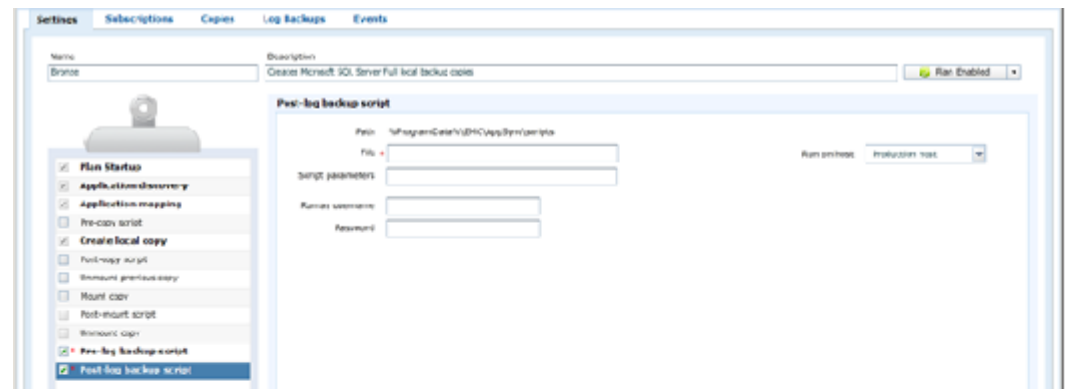
6. To perform a checksum on the log backup, select the **Checksum the backup** field.
7. Set **Minimum Retention Hours** option to control when transaction log backup files are deleted.

Transaction log backup expiration is done when no older database backups exist. AppSync deletes the log backup files and the log backup information contained in the AppSync database. The default setting is 24 hours which means that AppSync will not expire any log backup before it is a minimum of 24 hours old. The valid range is 0 to 10,000 hours.

Configure log backup scripts

You can run scripts before and after log backups by enabling the pre- and post- log backup scripts.

The pre-log backup script runs on the production host. The post-log backup script can run on the production host or the mount host (if mount is enabled), or you can specify a server. The server must have the AppSync host plug-in installed.

Figure 3 Configure log backup scripts

Run log backup on demand

You can schedule and run SQL Server log backups, or you can run log backups on demand.

Before you begin

To run log backups, make sure you check the service plan's **Enable Log Backup** option.

You can run a log backup on demand for an entire service plan, or run a log backup on demand for a single database instance.

Procedure

1. To run a log backups on demand for an entire service plan, go to **Service Plans > Microsoft SQL Server**, select the desired service plan enabling the **Run Log Backups** button, and then click **Run Log Backups** to run the backup for the entire service plan.
2. To run a log backup for a single database go to **Copy Management > Microsoft SQL Server**, select an SQL Server instance, select the desired database, and then select the **Log Backups** tab. Finally, select the appropriate service plan from the **Create log backup using plan** list to run the log backup.

View log backups for a service plan

The list of SQL Server log backups can be viewed from the Service Plan Log Backups tab or from the Database Log Backups tab.

Before you begin

This operation requires the Data Administrator role in AppSync.

The list of copies can be filtered by time of creation, and by service plan. In the Service Plan Copies tab, you can also filter by instance.

Procedure

1. To view the list of all log backups for a service plan, navigate to **Service Plans > Microsoft SQL Server**.
2. Select a service plan.
3. Click the **Log Backups** tab.

Results

You can now view the log backup list for the service plan. The following table describes details about the log backup:

Table 17 Service Plan log backup details

Column	Description
Status	<ul style="list-style-type: none"> • Green: successful • Yellow: some log backups completed with errors when the service plan ran. • Red: failed
Instance	SQL Server instance name
Database	SQL Server database name
Name	Name of the log backup copy. The copy is named with the time at which it was made.
Service Plan	Name of the service plan associated with the log backup.
Truncated	Indicates if the transaction log was truncated by the log backup. Yes, if the log was truncated, otherwise No.
Backup File	The name of the log backup file and its location.

View SQL database copies

View the list of database copies by browsing to **Copy Management > Microsoft SQL Server** and selecting a SQL Server, then a database.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can also see details of a copy from the Copies tab of the service plan.

You can filter the list of copies by time of creation, and by service plan. In the Service Plan Copies tab, you can also filter by instance.

Table 18 Service Plan Copy details

Column	Description
Status	<ul style="list-style-type: none"> Green: successful Yellow: completed with errors Red: failed
Name	Name of the copy. The copy is named with the time at which it was made.
Service Plan	<p>Name of the service plan that is associated with the copy. For repurposed copies, a Repurpose link displays in this column. Click this link to edit the Service Plan for 1st or 2nd generation copies.</p> <hr/> <p>Note</p> <p>In the service plan for repurposed copies, the options to schedule and mount overrides will be disabled.</p> <hr/>
SQL Server Backup Type	<p>Type of SQL backup: Full, Copy, Non VDI, or Crash Consistent</p> <ul style="list-style-type: none"> Full protects the database, and the active part of the transaction log. Copy protects the database and the active part of the transaction log without affecting the sequence of backups. Non VDI protects the database without using VDI and depends on VSS to create crash consistent copies. Crash Consistent protects the database without using VSS or VDI, and depends on the array functionality to create crash consistent copies. Therefore, there are no VSS timeout issues during copy creation. Secondary databases are read-only and can only be backed up with the Copy backup type. Auto Switch to Copy is enabled only when Full is selected as the backup type. However it is unchecked by default. Checking Auto Switch to Copy tells AppSync to check if the database role is Secondary, and if so, to switch the backup type to Copy. If Auto Switch to Copy is not enabled, backups fail for all secondary databases. When Non VDI or Crash Consistent backup type is selected, Auto Switch to Copy and Enable log backup are disabled.
Mount Status	Shows if the copy is mounted. If mounted, the name of the mount host displays.
Recovery Status	<p>Available values:</p> <ul style="list-style-type: none"> Not Recovered - when copy is not mounted or it is a file system mount Successful - when Recovery is successful Failed - when Recovery failed

Table 18 Service Plan Copy details (continued)

Column	Description
Availability Group	The Availability Group column lists the availability group the database belongs to.
Generation	Used for repurposed copies, this column describes how many generations removed the copy is from the production database.
Source	This column displays the source database or copy from which a copy was created.
Copy Type	<p>Type of copy can be one of the following:</p> <ul style="list-style-type: none"> • RecoverPoint Continuous Data Protection Bookmark • RecoverPoint Continuous Remote Replication Bookmark • VNX Snap • VNXeSnap • Unity Snap • VMAX Snap, VMAX Clone • XtremIO snapshot • VMAX 3 SnapVXClone, SnapVXSnap • VPLEX Snap <p>The following additional details are displayed in the Service Plan Copies tab:</p> <ul style="list-style-type: none"> • Instance: The SQL Server instance that hosts the database. • Database name: The name of the copy's database. • Time: The time at which the database copy was made. • Server/cluster: Name of the server or the cluster that hosts the SQL Server instance. • Site: RecoverPoint

Note

A **Repurpose** button on this page is enabled. When you select a **1st Generation copy**, the Repurpose wizard launches where you can create 2nd Generation copies.

View log backup list for a single database

You can also view log backups for a single database.

Follow these steps:

Procedure

1. Navigate to **Copy Management > Microsoft SQL Server**, and then select an SQL Server instance.
2. Click the **User Databases** folder.
3. Click on a database in the list and select the **Log Backups** tab.

Results

You can now view log the log backup list for the database. The following table describes details about the log backup:

Table 19 Database log backup details: SQL Server instance

Column	Description
Status	<ul style="list-style-type: none"> Green: successful Yellow: some log backups completed with errors when the service plan ran. Red: failed
Name	Name of the log backup copy. The copy is named with the time at which it was made.
Service Plan	Name of the service plan associated with the log backup.
Truncated	Indicates if the transaction log was truncated by the log backup. Yes, if the log was truncated, otherwise No.
Backup File	The name of the log backup file and its location.

Log backup expiration

AppSync expires log backups when the service plan runs to create a new log backup. During expiration, AppSync deletes the log backup file and removes information about the backup from the AppSync database.

Log backups are always based off the previous Full database backup. However, you do not have to use AppSync to create the Full database backup. You can use AppSync to create a Copy database and log backup.

Additionally, AppSync can create Full database backups and log backups with, or without log truncation. Log backup expiration behavior depends on the type of database backup you create.

Log backups are eligible for expiration when the following conditions occur:

- The log backup is older than the service plan Minimum Retention Hours setting.
- All older database backups are expired. The database backups included in this check depends on the SQL Server Backup Type.
 - If the log backup service plan has SQL Server Backup Type set to Copy, only database backups created by that service plan are considered when looking for older database backups.
 - If the log backup service plan has SQL Server Backup Type set to Full, then Full database backups created by any service plan are considered.

Example 1: consider the following scenario:

- Service plan has log backup enabled.
- Database backup type set to Copy.
- Rotation set to one.
- Log backup minimum retention is set to 24 hours.

The service plan has run several times, creating a database backup and several log backups. The service plan runs again, creating a database backup and expiring the first database backup. This leaves several log backups with no older database backup. The service plan runs again, creating a log backup and expiring all of the previous log backups that are at least 24 hours old.

Example 2: consider the following scenario:

- You have two service plans.
- Both have database backup type set to Full.
- Service plan 1 is scheduled to run a database backup once a week with rotation set to four.
- Service plan 2 is scheduled to run daily at 8 PM with a rotation of seven.
- Service plan 2 has log backup enabled to run every hour and the log backup minimum retention is set to 24 hours.
- Both service plans have been running.
- Service plan 1 has four database copies and service plan 2 has seven database copies. Service plan 2 also has many log backups that were run between each of the seven database copies.
- Service plan 2 runs again and creates a database copy and then expires its oldest copy. It runs an hour later to create a log backup and looks for log backups that are eligible for expiration.

No log backups are eligible because service plan 1 has Full database backups that are older than all of the log backups. The next time service plan 1 runs, the oldest database backup will be expired. Log backups will then be eligible for expiration.

Manual expiration of log backups

You can also expire log backups manually.

To expire log backups for several databases:

1. Navigate to **Service Plans > Microsoft SQL Server**, and click on a service plan.
2. Click the **Log Backups** tab.
3. Select the log backups that you would like to expire and then click **Expire**.
4. Click **OK** on the confirmation dialog. AppSync will delete the log backup file and remove information about the backup from the AppSync database.

To expire log backups for a single database:

1. Navigate to **Copy Management > Microsoft SQL Server** and select an SQL Server instance.
2. Click the **User Databases** folder.
3. Click on a database in the list and select the **Log Backups** tab.
4. Select the log backups that you would like to expire and then click **Expire**.
5. Click **OK** on the confirmation dialog.

Considerations for working with SQL Server in a cluster

There are special considerations when working with SQL Server in a cluster.

When protecting SQL Server databases in a clustered environment, you must install the AppSync host plug-in on all of the nodes that are possible owners of a clustered SQL Server instance. You can use the AppSync console to install the plug-in or manually install the plug-in on each server.

Protecting clustered SQL Server instances:

- It is mandatory to add all the possible nodes of a clustered SQL Server instance and then add SQL Server virtual server (network name or IP address) to AppSync after installing the AppSync host plug-in software on each node.

Note

The SQL Server virtual server is different from the failover cluster name.

- Only single subnets are supported.
- Production VMWare virtual disk with multi writer option enabled is not supported. Protection might succeed, but mount fails.
- Mounting AppSync copies as a standalone resource:
 - You can mount AppSync copies of clustered databases as a standalone database to a standalone server or any cluster node.
 - You can mount AppSync copies of standalone databases as a standalone database to a standalone server or any cluster node.

Mounting a SQL Server copy as a clustered resource:

- Supports mount to either an alternate cluster or a production cluster as a clustered resource. On the production cluster, you must select an alternate clustered instance for mount with recovery.
- Mount is supported in the environments of VMAX, VNX, VNXe, Unity, XtremIO or RecoverPoint. The *AppSync Installation and Configuration Guide* describes the required storage configuration steps.
- Select the appropriate mount option that applies for cluster mount based on your cluster and storage configuration.
- Manually disable `automount`. Run `diskpart` at a command prompt then enter `automount disable` at the `DISKPART>` prompt.

Special considerations for mount to production cluster:

- Mounting to a production cluster node using the original path is not supported.
- Virtual servers are filtered out while using the "mount copy" to server option. Only cluster nodes or standalone servers are visible.
- Mount with Recovery as a clustered resource to a clustered production server is supported.
- For Mount with Recovery as a clustered resource to a production virtual server, consider the following:
 - Mounting to a different clustered SQL server instance is supported.
 - Mounting to a production clustered SQL server instance is not supported.
 - Mounting to an alternate mount path is supported. The root disk for the alternate mount point must be a clustered disk, and must be added as a dependency to SQL server.
 - Mounting to the original path is not supported.
- Performing a RecoverPoint mounted restore while the copy is mounted to a production cluster is not supported.

Example 1 Protect databases owned by clustered instances of SQL Server

Example 1 Protect databases owned by clustered instances of SQL Server (continued)

In this example, you want to protect databases owned by clustered instances of SQL Server. In addition, some of those databases belong to AlwaysOn availability groups. Refer to the following figure for this example:

Figure 4 Cluster information

Name	Virtual Machine	Virtual Server	Cluster	Platform	Plug-in Ver...	Last Discovery
Irmq093	Yes	Yes	AlwaysOnCluster	Microsoft(R) Windows(R) Server 2012 Standard Edition 64-bit (build 9200)	2.2.0.0	Wed 05/29/2014 12:50:19 PM
Irmq094	Yes	Yes	AlwaysOnCluster	Microsoft(R) Windows(R) Server 2012 Standard Edition 64-bit (build 9200)	2.2.0.0	Wed 06/26/2014 13:00:19 PM
Irmq120	Yes	No	AlwaysOnCluster	Microsoft(R) Windows(R) Server 2012 Standard Edition 64-bit (build 9200)	2.2.0.0	Wed 05/29/2014 12:27:08 PM
Irmq121	Yes	No	AlwaysOnCluster	Microsoft(R) Windows(R) Server 2012 Standard Edition 64-bit (build 9200)	2.2.0.0	Wed 05/29/2014 12:27:06 PM
Irmq122	Yes	No	None	Microsoft(R) Windows(R) Server 2012 Standard Edition 64-bit (build 9200)	2.2.0.0	Wed 05/29/2014 12:33:08 PM
Irmq126	Yes	No	None	Microsoft(R) Windows(R) Server 2012 Standard Edition 64-bit (build 9200)	2.2.0.0	Wed 05/29/2014 12:27:04 PM

This cluster has the following configuration:

- Irmq093 is a SQL Server virtual server and it hosts databases belonging to AlwaysOn Availability Groups.
- Irmq094 is a SQL Server virtual server.
- Irmq120, Irmq121, and Irmq122 belong to a SQL Server AlwaysOn failover cluster.
- Irmq120 and Irmq121 are possible owners of the clustered SQL Server instances owned by Irmq093 and Irmq094.
- Irmq122 has a standalone instance of SQL Server installed that hosts databases belonging to AlwaysOn availability groups.
- Irmq126 is the mount host with a standalone instance of SQL Server installed

To protect the databases belonging to the clustered and standalone instance, follow these steps:

1. Use the AppSync console to add Irmq120, Irmq121, and Irmq122. AppSync will install the plug-in on these servers and discover any non-clustered instances. If you need a mount host, you can add Irmq126 now.
2. With the AppSync console, add the virtual servers for the clustered instances, add Irmq093 and Irmq094.

SQL Server User Databases folder

The SQL Server User Database folder contains all the user databases for this SQL Server instance that have been discovered and stored in the AppSync database.

From the **Protect** button, you can subscribe the folder to a plan. By doing so, all the databases part of this folder are also protected. Once protected, the **Service Plan** column displays the name of the plan.

Clicking on the **User Databases** folder lists the individual databases part of this SQL Server instance.

In the Databases page, an entry in the **Service Plans** column tells you that all the databases that are part of the folder are protected. Any user databases added to the instance will also be protected. AppSync will automatically stop protecting any databases removed from the instance.

Note

If one or more user databases for an SQL Server instance are subscribed to a service plan, you cannot subscribe the User Databases folder to the same service plan. Conversely, if the User Databases folder is subscribed to a service plan, you cannot subscribe individual user database instances to the same service plan.

Discover SQL Server instances

To keep AppSync up-to-date, you should discover SQL Server instances when there is creation or deletion of instances.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft SQL Server**.
2. From the **Discover Instances** popup button below:
 - Select **On Server** and select one of the servers to discover instances from.
 - Select **Add Servers** to add a new server to AppSync.

Note

Clustered instances are discovered and displayed in the AppSync GUI only if SQL Virtual Server (network name or IP address) is registered to the AppSync server.

Discovering SQL Server databases

AppSync discovers new user databases on demand or automatically on a service plan run.

When you click the User Databases folder the first time, AppSync discovers databases and lists them. To manually discover databases again, click **Discover Databases** in the **Databases** page.

On the other hand, when you subscribe the User Databases folder to a plan, databases are automatically discovered on each run of the plan. All databases that are currently

ONLINE, including those that were added to the SQL instance after the last service plan run, are automatically protected.

If individual databases are subscribed to a plan instead of the User Databases folder, AppSync does not automatically discover any new databases that were created after the last run of the plan. In this case, AppSync rediscovers the database information of all the databases originally subscribed to the plan and protects the ones that are ONLINE.

Protect a SQL Database

Protect a SQL database by subscribing it to an AppSync service plan.

To optimize performance, AppSync creates copies of a maximum of 35 databases per instance. If more than 35 databases are subscribed per instance, AppSync breaks them into groups of 35 and creates copies of the groups sequentially. If more than 35 databases are subscribed to a service plan, and the databases reside on same storage unit (CG, LUN, DS, and so on), the split into groups with 35 databases does not occur. A single copy is desirable for a configuration when storage is on the same storage unit.

This number (35) is a server setting and can be modified, if required. Contact Dell EMC Support to do so.

You can protect objects in different ways from different places in AppSync:

- Choose **Subscribe to Plan and Run** when you want to protect a selected database immediately. The service plan is run for the database alone.
- Choose **Subscribe to Plan** when you want to schedule the protection for later. Protection for databases that are part of the service plan are run at the scheduled time.
- Choose an appropriate service plan from **Create copy using plan** in the database Copies page.
- Choose **Run** from the SQL Server Service Plans page to run the whole plan immediately.

Note

Ensure that the database you are protecting is not configured for backup at the same time using a non-AppSync backup tool. This might interfere with AppSync copy operation and result in unexpected errors.

Configuring protection for SQL Server database

You subscribe a database or the User Databases folder to a service plan and run the service plan immediately, or schedule the service plan to run at a later time.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft SQL Server**.
2. Click a server instance to display its databases.
 - To protect all databases within the User Database folder, select the **User Database** folder.
 - To protect an individual database, click **User Databases** and select a database from the list.

3. From the **Protect** popup button below, select the appropriate service plan from:

Option	Description
Subscribe to Plan and Run	To subscribe the database for protection and run the plan immediately for the selected database(s).
Subscribe to Plan	To subscribe the database for protection. Protection for all databases that are part of the service plan are executed at the scheduled time.

Unsubscribing a database from a service plan

When you unsubscribe an individual database from a service plan, all existing database copies will be retained; only further protection will be removed.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft SQL Server**.
2. Click the SQL server instance.
3. Click **User Databases**.
4. Select the database to unsubscribe from a service plan.
 - Select the plan to unsubscribe from **Protect > Unsubscribe from Plan**. Only plans to which the database is subscribed to are in the popup list.
 - To unsubscribe from all service plans, select **Unsubscribe from Plan > All**. Alternatively, you can now easily unsubscribe SQL database from the **Subscription** tab of a service plan.

Discovering SQL Server databases

Use the **Discover New Databases** command to update the SQL Server databases known to AppSync.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft SQL Server**.
2. Click a server instance, then the **User Databases** folder.
3. In the User Databases page, click **Discover New Databases** from the row of buttons below.

Discovery can take several minutes to complete depending on the size of the instance.

SQL copies page

You can see details of a copy from the Copies tab of a database by clicking on that particular database. The list of copies can be filtered by time of creation and by type (protection or repurposed).

Table 20 Copy page fields

Column	Description
Status	<ul style="list-style-type: none"> Green: successful Yellow: completed with errors Red: failed
Name	Name of the copy. The copy name is the time AppSync created it.
Service Plan	<p>Name of the service plan associated with the copy. In the case of a repurpose plan, select a copy and click on the Repurpose link to edit it.</p> <hr/> <p>Note</p> <p>Each copy is associated with a unique repurposing service plan.</p> <hr/>
Label	Label assigned to the copy in case of repurposing.
SQL Server Backup Type	<ul style="list-style-type: none"> Full Copy Non VDI Crash Consistent
Mount Status	Status of the copy: mounted or not mounted. If mounted, the name of the mount host displays.
Recovery Status	<p>Was copy recovered post mount or not. Values are:</p> <ul style="list-style-type: none"> Not Recovered - copy was not mounted or copy was a file system mount. Successful - recovery was successful. Failed - recovery failed.
Availability Group	Lists the AlwaysOn Availability Groups to which the database belongs.
Generation	First or second generation copy - for repurposing

Table 20 Copy page fields (continued)

Column	Description
Source	Production database (for first generation copy) or a copy of a copy (second generation) copies.
Copy Type	<ul style="list-style-type: none"> • Local Bookmark • Remote Bookmark • VNX Snap, VNXeSnap • Unity Snap • VMAX Clone • VMAX Snap • VMAX V3/VMAX AFA: SnapVXSnap, SnapVXClone • XtremIO Snap • ViPRSnap
Site	RecoverPoint site information.

Creating a database copy from the Copies page

Create a copy of a database by subscribing it to an AppSync SQL Server service plan from the Copies page.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Microsoft SQL Server**.
2. Click a server instance and then click **User Databases** to display its databases.
3. Click on a folder to display **User Databases**.
4. From this list, click the database to view its copies.
5. From the **Create a copy using plan** list, select the appropriate service plan.

The service plan runs immediately for the database.

Expiring an SQL database copy on demand

Expiring a database copy removes it from the AppSync database and can free up storage, depending on the replication technology and copy state.

Before you begin

This operation requires the Data Administrator role in AppSync.

Expiring a copy that was made with RecoverPoint does not remove the corresponding bookmark from RecoverPoint itself.

Procedure

1. Select **Copy Management > Microsoft SQL Server**.

2. Click a SQL Server instance to display its database folders.
3. Click the **User Database** folder.
4. Click the database whose copies you want to expire.
5. From the **Copies** page, select one or more copies to expire.

You can also perform this action from the Service Plan's **Copies** tab.

6. Select **Expire** from the row of buttons below.

Verify that you want to expire the copy you selected and any associated copies listed and confirm.

Service plan summary and details

The service plan **Settings** tab shows the name, description, schedule, and status of the service plan. Click the phases for detailed service plan settings and other tabs for information about subscriptions, lists of copies and events generated by the plan.

Service plan schedule

The schedule of a service plan is set in the **Plan Startup** phase.

The **Startup Type** (scheduled or on demand) determines whether the plan is run manually, or configured to run on a schedule. Options for scheduling when a service plan starts are:

- Specify a recovery point objective (RPO)
 - Set an RPO of 30 minutes or 1, 2, 3, 4, 6, 8, 12, or 24 hours
 - Minutes after the hour are set in 5 minute intervals
 - Default RPO is 24 hours
- Run every day at certain times
 - Select up to two different times during the day
 - Minutes after the hour is in 5 minute intervals
 - There is no default selected
- Run at a certain time on selected days of the week
 - One or more days of the week (up to all seven days) can be selected
 - There is no default day of the week selected. Default time of day is 12:00 AM.
- Run at a certain time on selected days of the month
 - Select one or more days of the month (up to all days)
 - Select one time of day. Available times are at 15 minute intervals.
 - Default is the first day of the month

Overriding service plan schedules

You can set individual schedules for databases subscribed to a service plan, overriding the generic recurrence setting.

Before you begin

This operation requires the Service Plan Administrator role in AppSync.

You can override only the settings of the recurrence type already selected for the service plan.

Procedure

1. Navigate to **Service Plans** and select one of the plans from the list.
2. From the **Settings** tab, select the **Plan Startup** phase.
You will see the **Plan Startup Defaults** pane on the right.
3. Note the **Recurrence Type** selected for the plan.
A recurrence type can be set only if **Scheduled** is selected as the **Startup Type**.
4. Select the **Start service plan** phase.
You will see the **Start service plan** pane on the right.
5. Note the **Recurrence Type** selected for the plan.
A recurrence type can be set only if **Automatic** is selected in the **Startup**.
6. Click the **Plan Startup Overrides** tab.
You can see the list of all databases subscribed to the plan.
7. Select one or more databases and click **Override Schedule**.
The **Override Schedule** dialog is displayed.
8. Set the schedule based on your requirement and click **OK**.
For example, if the default recurrence type is **On specified days of the month**, and the rule setting is to **Run at 12:00 AM on the 1st day of every month**, you can override the time and the day for individual datastores.
A Pencil icon indicates that default settings have been overridden.

Application discovery

Before creating the User Database folder's copy, AppSync examines the SQL Server instance to look for changes such as addition, deletion, renaming, or movement of databases. If individual databases are being protected, AppSync rediscovers information about the selected database. A database is protected only if it is in the ONLINE state.

There are no user settings associated with this phase and it cannot be disabled.

Application mapping

After discovering the application, AppSync maps it to array storage, and protection services such as RecoverPoint.

There are no user settings associated with this phase and it cannot be disabled.

Pre-copy script

To perform preparatory steps before creating a copy, specify a pre-copy script and parameters on a service plan's **Settings** tab.

The pre-copy script runs according to the schedule set in the **Plan Startup** phase. Valid script formats are .bat, .exe, and .ps1 (PowerShell scripts). You can optionally enter credentials to run the script as a specific user. The script runs as Local System by default.

AppSync does not support running of PowerShell scripts directly. You usually must wrap them in a .bat file. The other option is to make the default "Open" on ps1 files `C:\WINDOWS\system32\WindowsPowerShell\v1.0\powershell.exe`. When the PS script runs, you may get an error and you must set an appropriate execution policy.

To run PowerShell commands from scripts:

1. Specify the full pathname to the PowerShell command file in the .bat file:

```
powershell -command C:\PshellCommands.ps1 <nul
```
2. Set the PowerShell execution policy so you can run the script. For example, the first line in the .bat file should look like the following for an unrestricted policy:

```
powershell -command set-executionpolicy unrestricted <nul
```
3. To ensure correct termination of the PowerShell session, add <nul to the end of the line that calls the PowerShell script. The default location of the script is %ProgramData%\EMC\AppSync\scripts\ on the application host.

Exact parameters depend on the script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Create copy

The **Create Copy** phase creates a copy based on the replication technology specified in the service plan.

This phase specifies the backup type of SQL Server copy to make. For VNX Snapshot copies, this phase also sets the period for automatic expiration of the copies.

Review [Overview: Service Plan](#) for more service plan copy information.

SQL Server backup type

The four main backup types supported are: Full, Copy, Non VDI, and Crash Consistent.

- **Full** protects the database, and the active part of the transaction log. This copy type is typically used when the copy will be considered a backup of the database or when the copy will be mounted in order to use a third-party product to create a backup of the database. This type of copy allows you to restore transaction logs to bring the database forward to a point in time that is newer than the copy, assuming you have backed up those transaction logs. AppSync uses Microsoft SQL Server's VDI snapshot feature to create this type of copy.
 - **Auto Switch to Copy** is enabled only when **Full** is selected as the backup type. However it is unchecked by default. Checking **Auto Switch to Copy** tells AppSync to check if the database role is Secondary, and if so, to switch the backup type to **Copy**.

Note

If **Auto Switch to Copy** is not enabled, backups fail for all secondary databases.

- **Copy** protects the database and the active part of the transaction log without affecting the sequence of backups. This provides DBAs with a way to create a copy without interfering with third-party backup applications that may be creating full and/or differential backups of the SQL Server databases. AppSync uses Microsoft SQL Server's VDI snapshot feature to create this type of copy.

Note

Secondary databases are read-only and can only be backed up with the **Copy** backup type.

- **Non VDI** protects the database with the non VDI approach. This creates crash consistent copies of SQL using the VSS freeze/thaw framework. No VDI meta

data is generated for non VDI copies. You can mount Non VDI SQL copies using the Attach Database and Mount Copy options. You can restore a Non VDI copy using the No Recovery mode.

- **Crash Consistent** protects the database without any agent involvement, that is, VSS or VDI are not used. This backup type creates crash consistent copies of SQL databases using array level features. You can use this backup type to re-snap, if the copy recovery fails. This eliminates intervention on the production side. There is no VSS and VDI metadata generated for Crash Consistent copies. Crash Consistent SQL copies are mounted using VDS mount and are recovered with the `Attach Database` recovery option. You can restore a Crash Consistent copy using the `No Recovery` mode. Crash Consistent backup type is supported on all supported arrays and RecoverPoint except ViPR with the following restrictions:
 - For VNX, VNXe, and Unity, the SQL database must reside on the LUNs that are in a consistency group.
 - For VPLEX virtual volumes on Unity, the SQL database must reside on the Unity LUNS that are in a consistency group.
 - There is no restriction for VMAX, RecoverPoint, XtremIO, and VPLEX virtual volumes on XtremIO.
 - RecoverPoint bookmarks shows Snap consistency as **Crash Consistent** in the RecoverPoint GUI.

Automatic expiration of copies

The automatic expiration value in a service plan's Create Copy phase specifies the maximum desired number of Snap, Clone or Bookmark copies that can exist simultaneously.

When the "Always keep *x* copies" value is reached, older copies are expired to free storage for the next copy in the rotation. Failed copies are not counted. AppSync does not expire the oldest copy until its replacement has been successfully created. For example, if the number of copies to keep is 7, AppSync does not expire the oldest copy until the 8th copy is created.

AppSync does not expire copies under the following circumstances:

- Mounted copies are not expired.
- A copy that contains the only replica of a database will not be expired.

This setting is independent of the VNX pool policy settings in Unisphere for automatic deletion of oldest snapshots. The service plan administrator should work with the storage administrator to ensure that the VNX pool policy settings will enable the support of the specified number of snapshot copies for the application residing in that pool.

Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations.

Note

If this option is not selected, then RecoverPoint copies will accumulate, and will remain until the bookmarks fall off the RecoverPoint appliance.

Configure retry on VSS failure

You can configure a VSS retry count in the create copy phase of a service plan. During protection, if a service plan fails because of VSS failures such as VSS timeout issue, the service plan runs the VSS freeze/thaw operation again based on the specified retry count and interval. This option is supported only on Windows applications - File

system, Microsoft SQL, and Microsoft Exchange. This option is not used while creating Crash consistent copies of SQL databases.

Note

AppSync does not perform a VSS retry, if the application freeze itself fails. If the application is not in a state to create a copy, AppSync fails to quiesce it, and does not retry the VSS freeze/thaw operation. The application must be brought back to a state where it can be quiesced and then the service plan must be re-run.

Post-copy script

To perform cleanup or other post-copy steps after creating a copy, specify a post-copy script and parameters in a service plan's **Settings** tab.

The script runs on successful completion of the **Create copy** phase. Valid script formats are .bat, .exe, and .ps1 (PowerShell scripts). You can optionally enter credentials to run the script as a specific user. The script runs as Local System by default.

When AppSync creates copies of application items in a service plan, it may break up the application items and place them in separate groups for protection. This action can be for performance reasons (for example, VSS for Exchange and SQL) or because items in a service plan may be protected by different replication technologies. For example, a service plan may contain some application items that are protected by VNX Snapshots and some by RecoverPoint bookmarks. As a result, application items in these groups are protected independently.

When AppSync calls a post-copy script, it passes the copies which were created in the group by calling the script with `-appCopies <APP1> <APP2>`, where APP1 and APP2 are the names of the application items in that grouping.

AppSync does not support running of PowerShell scripts directly. You usually must wrap them in a .bat file. The other option is to make the default "Open" on ps1 files `C:\WINDOWS\system32\WindowsPowerShell\v1.0\powershell.exe`. When the PS script runs, you may get an error and you must set an appropriate execution policy.

To run PowerShell commands from scripts:

1. Specify the full pathname to the PowerShell command file in the .bat file:

```
powershell -command C:\PshellCommands.ps1 <nul
```
2. Set the PowerShell execution policy so you can run the script. For example, the first line in the .bat file should look like the following for an unrestricted policy:

```
powershell -command set-executionpolicy unrestricted <nul
```
3. To ensure correct termination of the PowerShell session, add `<nul` to the end of the line that calls the PowerShell script.

When AppSync runs the post-copy script, it is run for the application items that are part of a group. If there are multiple groups, the post-copy script runs multiple times. When AppSync runs the post-copy script, it passes the list of application items in the replication group as arguments to the script, right after the user arguments. The syntax is:

```
-applicationCopies <ITEM1> <ITEM2> <ITEM3>
```

where `<ITEMx>` is the name of the application item that is being protected.

The default location of the script is `%ProgramData%\EMC\AppSync\scripts\` on the application host.

Exact parameters depend on the script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Unmount previous copy

The service plan unmounts a previously mounted copy after creating the new copy. The exception is a copy that was mounted on-demand as opposed to by the service plan; in this case the on-demand mounted copy is not unmounted.

All the recovered databases are shut down as part of this phase. There are no user settings associated with this phase and it can be enabled or disabled.

Mount copy

The Mount copy phase either mounts the copy or mounts and recovers the copy. This phase can be enabled or disabled.

In the **Mount Copy Defaults** settings, you can set values to Mount copy or Mount and recover copy.

In the **Mount copy** settings, you set the mount host value, mount path and mount permissions (read-only or read-write). Other mount settings determine where the SQL metadata files are copied and the RecoverPoint image access type.

Field	Description
Mount on Server	The server on which to mount the copy. Only the nodes of the cluster or standalone hosts are available for selection. SQL virtual servers are filtered out.
Mount with access	Type of access the copy should be mounted with.
Mount path	<ul style="list-style-type: none"> The Default Mount Path is %SystemDrive%\AppSyncMounts\%%ProdServerName%%. To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). The default path also contains an AppSync variable (ProdServerName) which is delimited with 2 percent signs (% %). The following characters are not valid in the path: < > : " / ? * The mount path could also be Same as Original Path. However, this option is not available when the mount host is the same as production host. If you specify a non-default mount path, the drive that is specified for mount cannot be a clustered disk. Select Mapped Path to specify the path where you want to mount the database.
Copy metadata files to	<ul style="list-style-type: none"> The Default Path is the location to copy VDI and VSS metadata files: %SystemDrive%\AppSyncMounts\%%ProdServerName%% The following characters are not valid in the path: < > : " / ? *

Field	Description
	<ul style="list-style-type: none"> • If you back up the database to another media, back up the metadata files as well. • AppSync can integrate with third-party backup software to create tape backups of SQL Server copies. The target directory that is specified here must be part of the backup. <hr/> <p>Note</p> <ul style="list-style-type: none"> • Metadata is not created for Non VDI copies. • VSS or VDI metadata is not generated for Crash Consistent copies.
Image access mode (during RecoverPoint mount)	<ul style="list-style-type: none"> • Logged access: Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. Logged access is the only option available when you mount to the production host. • Virtual access with roll: Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. With RP VMAX, and RP XtremIO, virtual access with roll is not supported. • Virtual access: Provides nearly instant access to the image. Virtual access is not intended for heavy processing. Virtual access with RP VMAX and RP XtremIO is not supported.
Service Level Objective (SLO)	For VMAX 3 arrays only, a setting called Desired Service Level Objective (SLO) appears in the Mount wizard and specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.
VPLEX Mount option	<ul style="list-style-type: none"> • Native array: Use this option if you want to mount the copy as native array volumes. • VPLEX virtual volume mount: Use this option if you want to mount the copy as VPLEX virtual volumes. • Enable VMware cluster mount: Clear this option if you do not want to perform an ESX cluster mount. By default, this option is enabled.
Use Dedicated Storage Group	<ul style="list-style-type: none"> • Applicable only for physical hosts or virtual machines with direct iSCSI as part of cluster. • Checked by default, enabling this option allows AppSync to enforce a dedicated VMAX , VNX storage group, or XtremIO initiator group for a mount. (A dedicated VMAX or VNX storage group contains the selected mount host only.) For XtremIO, this option applies to an XtremIO initiator group that only contains an initiator for the mount host. The mount fails if you are mounting to a node of a cluster that is in a storage group that is shared with the other nodes.

Field	Description
	<p>Note</p> <p>Use this option to mount the copy to a node for copy validation or backup to tape. In this scenario, you need two storage groups. One storage group is dedicated to the passive node being used as a mount host and the other storage group is for the remainder of the nodes in the cluster. Both storage groups contain the shared storage for the cluster.</p> <hr/> <ul style="list-style-type: none"> If unchecked, AppSync does not enforce the use of a dedicated storage group for a mount. <hr/> <p>Note</p> <p>Uncheck this option for manually adding the target devices as clustered storage and presenting them to clustered SQL Server instances for data repurposing and data mining.</p>

In the **Mount and recover copy** settings, you specify the recovery instance, the type of recovery, and the database naming details. Other settings are similar to the Mount copy settings such as mount path and image access type.

Field	Description
Recovery Instance	<p>The SQL Server instance to be used for recovery. If the connection settings are not set or are invalid for the instance, the SQL Server Connection Settings dialog appears. Click Connection Settings to reset the credentials.</p> <p>If you are using a VMAX 3 array, a setting called Desired Service Level Objective (SLO) is available. The option appears in the Mount wizard and it specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group</p>
Recovery Type	Available options are: Recovery (default), No Recovery, Standby, and Attach Database
Database renaming	<p>This drop down includes:</p> <ul style="list-style-type: none"> Use original database names (default if alternate instance): This is not available for selection if the Recovery Instance is the production instance. Use original database names with suffix: This is the default if Recovery Instance is the production instance.
Naming Suffix	Only displayed when Original database names with Suffix is selected in the Database renaming dropdown. The default value is AppSync .
Mount path	<ul style="list-style-type: none"> The default mount path, when the mount SQL instance is a standalone instance (%SystemDrive%\AppSyncMounts\%ProdServerName%).

Field	Description
	<ul style="list-style-type: none"> • To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). • The default path also contains an AppSync variable (ProdServerName) which is delimited with two percent signs (%%). • The following characters are not valid in the path: < > : " / ? * • The mount path could also be Same as Original Path. You can select either of the options. • If you specify a non-default mount path for mounting to a standalone instance, the drive specified for mount cannot be a clustered disk. • For mounting to a clustered SQL instance, the “Same as Original Path” and alternate mount paths are supported. “Default mount path” is not supported. Instead, you can type this option to specify an alternate mount path. The root disk for the alternate mount path must be clustered and a dependency must exist for SQL server on the clustered disk.
Copy metadata files to	<ul style="list-style-type: none"> • By default, the location to copy VSS metadata files is the same as the mount path. • If the mount path is Same as Original Path, then this defaults to %SystemDrive%\AppSyncMounts\%%ProdServerName%%. • The following characters are not valid in the path: < > : " / ? * • If you are backing up the database to another media, you must backup these metadata files as well. • AppSync can integrate with third-party backup software to create tape backups of SQL Server copies. The target directory specified here must be part of the backup. <hr/> <p data-bbox="667 1331 722 1356">Note</p> <ul style="list-style-type: none"> • Metadata is not created for Non VDI copies. • VSS or VDI metadata is not generated for Crash Consistent copies.
Image access mode (during RecoverPoint mount)	<ul style="list-style-type: none"> • Logged Access: Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host. • Virtual Access with Roll: Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. • Virtual Access: Provides nearly instant access to the image; it is not intended for heavy processing.

Field	Description
Use Dedicated Storage Group	<ul style="list-style-type: none"> • Applicable only for physical hosts or virtual machines with direct iSCSI part of cluster. • Checked by default, enabling this option allows AppSync to enforce a dedicated VMAX, VNX, or XtremIO storage group. For XtremIO, this option applies to an XtremIO initiator group that only contains an initiator for the mount host. The storage group contains the selected mount host only for a mount and the mount will fail if you are mounting to a node of a cluster that is in a storage group shared with the other nodes. <hr/> <p>Note</p> <p>Use this option to mount the copy to a node for copy validation or backup to tape. In this scenario, you will need two storage groups. One storage group is dedicated to the passive node being used as a mount host and the other storage group is for the remainder of the nodes in the cluster. Both storage groups contain the shared storage for the cluster.</p> <hr/> <ul style="list-style-type: none"> • If unchecked, AppSync does not enforce the use of a dedicated storage group for a mount and the mount will proceed. Host initiators can only belong in one initiator group in XtremIO, so use this option to ensure that you mount to a mount host that is the only host in the initiator group. <hr/> <p>Note</p> <p>Uncheck this option for manually adding the target devices as clustered storage and presenting them to clustered SQL Server instances for data repurposing and data mining.</p>

Overriding mount settings in a service plan

If multiple registered SQL Servers are subscribed to the same plan, you can select different mount and recover settings for each SQL Server, overriding the generic settings.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Service Plans > Microsoft SQL Server** and click one of the plans from the list.
2. From the **Settings** tab, select the **Mount copy** phase.
3. On the right pane, select the **Mount Copy Overrides** tab.

The list of servers include all SQL servers whose databases are subscribed to this plan.

Based on whether **Mount copy** or **Mount and recover copy** is selected, the default settings display for all the Servers.

4. Select the Server whose settings you want to override and click **Set Overrides**.

The **Override Default Mount Settings** dialog is displayed.

5. Select options only for those mount settings that you wish to override.

Fields that do not have a selection retain their default settings.

6. Click **OK**.

A pencil icon appears in the first column of the Server's row whose default mount settings you changed.

7. To revert back to default settings for a server, click **Use Default Settings**.

Post-mount script

Specify a post-mount script and parameters from the Post-mount script option in the **Settings** tab of a service plan.

The script runs on successful completion of the mount copy or mount with recovery phase. This script is typically used for backup.

From the **Server** list, select the server on which to run the script. You can optionally run it on a registered host other than the mount host, and enter credentials to run the script as a specific user.

The default location of the script is `%ProgramData%\EMC\AppSync\scripts\` on the application host.

Exact parameters depend on your script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Unmount copy

The final phase in the service plan unmounts the copy. This phase is disabled if the **Unmount previous copy** phase is enabled. There are no user settings associated with this phase.

If you have chosen to **Mount and recover copy** in the **Mount copy** phase, all the mounted databases are shut down as part of this phase.

Custom shutdown script prior to unmount

Prior to unmount, if you wish to perform a customized shut down of the databases, you can place a script at the following location: `%ProgramData%\EMC\AppSync\script`.

The script name must be in this format:

```
<ServicePlanName>_<host_ProductionInstanceName OR
ProductionInstanceName>_ ShutdownSQL.bat where:
```

- `ServicePlanName` is the name of the service plan that the database is subscribed to
- `host_ProductionInstanceName OR ProductionInstanceName`:
 - In `host_ProductionInstanceName`, you can replace `host` by another name, the `ProductionInstanceName` is needed irrespective of whether there are different SQL instances or not.
 - Use `ProductionInstanceName` in case of default production instance which is equal to the host name.

Note

- It is recommended that you run the script as a Windows user. To run the script as a SQL Server user in SQL Server 2012 environment, the Local System user must have the sysadmin role.
 - Using the _ as a separator in the script file name is mandatory.
-

In the absence of a customized script, AppSync will perform a shut down of the databases prior to unmount.

Mount considerations for SQL Server

This section describes the mount host requirements, including rules for mount and production host versions and virtual machine mount host support.

The mount host requires the same versions of the AppSync agent plug-in, SQL Server, and HBA drivers as the production host. Mount hosts must have an SQL Server installed if you want to recover databases from the mounted copy. If database recovery is not performed, then SQL Server is not required on the mount host.

Note

- When you mount a replica of a SQL Server database to the production server, do not mount it using the same instance of SQL Server that the production database is using. You must use a different instance of SQL Server.
 - The mount path must not exceed 32,767 characters in length.
-

Mount and production host versions

- If you are mounting to the node of Windows failover cluster, please see the section [Microsoft Cluster Server mounts for SQL Server](#).
- If the major version of the SQL Server instance on the production mount host is later than that of the mount host, recovery will fail for all databases belonging to that instance.
- If the major version of the SQL Server instance on the production mount host is earlier than that of the mount host, recovery will succeed only if the recovery type is either RECOVERY or NORECOVERY. Recovery will fail if recovery type is STANDBY.
- If the major version of the SQL Server instance on the production mount host is same as that of the mount host, but the minor version is earlier, recovery will fail for all databases belonging to that instance.
- If the major version of the SQL Server instance on the production mount host is same as that of the mount host, but the minor version is later, recovery will succeed only if the recovery type is either RECOVERY or NORECOVERY. Recovery will fail if recovery type is STANDBY.
- If an AppSync created SQL database copy is recovered on a higher version of SQL Server instance on the mount host, recovery of the same copy on a lower version of SQL Server instance fails.

Virtual disk support

If the mount host is a virtual machine, the Virtual Center must be registered with AppSync. This is needed to mount RDMs.

For virtual disks:

- Production mount is not supported if the ESX host version is prior to 5.0.
- Non-persistent virtual disks are not supported.
- For datastore and virtual disk mounts on ESXi 5.x and RecoverPoint 4.1.7.7 environments, disable hardware acceleration to ensure successful virtual access type mounts. For more details, refer VMware Knowledge Base article 2006858.

For Hyper-V SCSI pass-through, the mount host cannot be a Hyper-V host it has to be a physical host or VM with NPIV or iSCSI direct attached.

Mount SQL Server Cluster as a clustered resource

[Considerations for working with SQL Server in a cluster](#) provides information on adding and discovering clustered resources.

- To mount a copy from a production cluster to an alternate cluster as a clustered resource, you must select a clustered SQL server instance of the alternate cluster on the **Mount with recovery** page. Mount as a clustered resource to any other clustered instance on the production cluster is supported. Mount as a clustered resource to the production cluster instance is not supported.
- Mount as a clustered resource is supported for SQL Server databases that reside on paths starting with drive letters such as P:\mysqlpdb\ or Q:\mysqlpdb. Mount as a clustered resource is not supported if production databases reside on clustered mount points such as I:\mount_point\, where I: is a clustered drive and another drive is mounted at I:\mount_point\.
- Mount to Same as Original Path is supported.
- Mount to an alternate path on the mount host is supported. You must specify the alternate mount path in the Mount path options. The root disk for the alternate mount point must be a clustered disk and SQL Server must have a dependency on the clustered disk.
- Multiple copies of the same database can be mounted to an alternate cluster at the same time.
- All recovery types are supported.
- Repurposing is supported.
- Databases can reside on any storage supported by AppSync except ViPR.
- If databases reside on raw device mappings in VMWare environments, the SQL Server cluster nodes must reside across different ESXi. This is a requirement from VMWare. For database on virtual disks, SQL Server cluster nodes can reside on the same ESX server.
- Raw device mapping in virtual compatibility mode is not supported.
- Static mounts are supported for RecoverPoint.

Mount SQL Server database copy on-demand

You can initiate an on-demand mount of a database copy from a copy or a database.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. In the Databases page, select **Recover > Mount a Copy**.
From the Copies page, select a copy and click **Mount**.
The SQL Server Mount wizard launches.

2. Use the **Database**, **Copies** or **Service Plan** filters to select the appropriate copy to mount.

The copies list is refreshed based on the filters selected.

3. Select the copy to mount.

For a RecoverPoint copy, you also have the option to select a bookmark based on a specific time. However, there should be a copy available in AppSync prior to the time you select **Mount**.

Click **Select a point in time** to select a copy with a specific time stamp. The time shown here is the console's time. If the console is in a different time zone from the RecoverPoint Appliance (RPA), specify the time as per the server's time zone to mount the copy.

For VMAX 3 arrays, select the desired Service Level Objective (SLO) for the mount copy.

Note

The SLO values are dynamically fetched from the VMAX 3 arrays, and only the unique values are displayed.

For VMAX V2 arrays, select the desired FAST VP policy for the mount copy.

For VPLEX, select a mount option:

- **VPLEX virtual volume** - The snapshot on the back-end array is made visible to VPLEX. AppSync creates VPLEX virtual volumes on the underlying native array snapshots and provisions these virtual volumes to the mount host. The provisioned virtual volumes are added to the storage view of the mount host. During unmount, the virtual volumes are removed from the storage view of the mount host, and it tears down the created VPLEX virtual volume on the underlying native array snapshots. The snapshot on the back-end array is de-provisioned from VPLEX.
 - **Native array volume** - AppSync provisions the native array snapshots to the mount host. The mount host must be zoned to the native array where the snapshot is created. All other mount considerations of the native array are applicable.
 - **Enable VMware cluster mount** checkbox - If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. Then the target LUN is made visible only to the ESX cluster on which the mount host resides. This is applicable for both RDM and vDisk device types.
4. In the **Mount Additional Copies** page, select one or more additional copies to mount. The copies listed here are of other databases that were protected at the same time and on the same SQL Server as the copy you selected in the previous step.
 5. On the **Select Mount Option** page, select one of the following:
 - **Mount copy** - You can choose to mount the copy to the mount host as VPLEX virtual volumes or as native array volumes. If the mount host is an ESX Cluster, you can perform a standalone or cluster mount. If the mount host is a node of a Windows Cluster, you can perform the mount by selecting a dedicated storage group where the copy target LUNs are masked to a different storage group than the storage group which has Cluster volumes already added.

- **Mount and recover copy** - Provides the option to select clustered instances to mount a SQL Server Database as a clustered resource. When you select this option, the **Recovery Instance** field becomes enabled. Select the clustered instance from the drop down box. Note that until you make a selection, all other options will be grayed out. Once you select an instance, AppSync prompts you to connect to that instance, and then the rest of the options become available.
6. In the **Summary** page, review the choices you made in the previous pages and click **Finish** to mount the copy.
 7. In the **Results** page, select **View Details** to see progress of the different phases that are part of mounting a copy.

The last phase completed is displayed at the bottom of the list.

SQL Server Mount Copy options

Review SQL server mount copy fields and descriptions.

Field	Description
Mount on Server	The server on which to mount the copy. Only the nodes of the cluster or standalone hosts are available for selection. SQL virtual servers are filtered out.
Mount with access	Type of access the copy should be mounted with.
Mount path	<ul style="list-style-type: none"> • The Default Mount Path is %SystemDrive%\AppSyncMounts\%%ProdServerName%%. • To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). • The default path also contains an AppSync variable (ProdServerName) which is delimited with 2 percent signs (% %). • The following characters are not valid in the path: < > : " / ? * • The mount path could also be Same as Original Path. However, this option is not available when the mount host is the same as production host. • If you specify a non-default mount path, the drive that is specified for mount cannot be a clustered disk. • Select Mapped Path to specify the path where you want to mount the database.
Copy metadata files to	<ul style="list-style-type: none"> • The Default Path is the location to copy VDI and VSS metadata files: %SystemDrive%\AppSyncMounts\%%ProdServerName%% • The following characters are not valid in the path: < > : " / ? * • If you back up the database to another media, back up the metadata files as well.

Field	Description
	<ul style="list-style-type: none"> AppSync can integrate with third-party backup software to create tape backups of SQL Server copies. The target directory that is specified here must be part of the backup. <hr/> <p>Note</p> <ul style="list-style-type: none"> Metadata is not created for Non VDI copies. VSS or VDI metadata is not generated for Crash Consistent copies. <hr/>
Image access mode (during RecoverPoint mount)	<ul style="list-style-type: none"> Logged access: Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. Logged access is the only option available when you mount to the production host. Virtual access with roll: Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. With RP VMAX, and RP XtremIO, virtual access with roll is not supported. Virtual access: Provides nearly instant access to the image. Virtual access is not intended for heavy processing. Virtual access with RP VMAX and RP XtremIO is not supported.
Service Level Objective (SLO)	For VMAX 3 arrays only, a setting called Desired Service Level Objective (SLO) appears in the Mount wizard and specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.
VPLEX Mount option	<ul style="list-style-type: none"> Native array: Use this option if you want to mount the copy as native array volumes. VPLEX virtual volume mount: Use this option if you want to mount the copy as VPLEX virtual volumes. Enable VMware cluster mount: Clear this option if you do not want to perform an ESX cluster mount. By default, this option is enabled.
Use Dedicated Storage Group	<ul style="list-style-type: none"> Applicable only for physical hosts or virtual machines with direct iSCSI as part of cluster. Checked by default, enabling this option allows AppSync to enforce a dedicated VMAX , VNX storage group, or XtremIO initiator group for a mount. (A dedicated VMAX or VNX storage group contains the selected mount host only.) For XtremIO, this option applies to an XtremIO initiator group that only contains an initiator for the mount host. The mount fails if you are mounting to a node of a cluster that is in a storage group that is shared with the other nodes.

Field	Description
	<p>Note</p> <p>Use this option to mount the copy to a node for copy validation or backup to tape. In this scenario, you need two storage groups. One storage group is dedicated to the passive node being used as a mount host and the other storage group is for the remainder of the nodes in the cluster. Both storage groups contain the shared storage for the cluster.</p> <hr/> <ul style="list-style-type: none"> If unchecked, AppSync does not enforce the use of a dedicated storage group for a mount. <hr/> <p>Note</p> <p>Uncheck this option for manually adding the target devices as clustered storage and presenting them to clustered SQL Server instances for data repurposing and data mining.</p>

SQL Server Mount and Recover copy options

SQL Server mount and recover copy options are explained in the following table:

Field	Description
Recovery Instance	<p>The SQL Server instance to be used for recovery. If the connection settings are not set or are invalid for the instance, the SQL Server Connection Settings dialog appears. Click Connection Settings to reset the credentials.</p> <p>If you are using a VMAX 3 array, a setting called Desired Service Level Objective (SLO) is available. The option appears in the Mount wizard and it specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group</p>
Recovery Type	Available options are: Recovery (default), No Recovery, Standby, and Attach Database
Database renaming	<p>This drop down includes:</p> <ul style="list-style-type: none"> Use original database names (default if alternate instance): This is not available for selection if the Recovery Instance is the production instance. Use original database names with suffix: This is the default if Recovery Instance is the production instance.
Naming Suffix	Only displayed when Original database names with Suffix is selected in the Database renaming dropdown. The default value is AppSync.
Mount path	<ul style="list-style-type: none"> The default mount path, when the mount SQL instance is a standalone instance (%SystemDrive%\AppSyncMounts\%ProdServerName%).

Field	Description
	<ul style="list-style-type: none"> • To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). • The default path also contains an AppSync variable (ProdServerName) which is delimited with two percent signs (%%). • The following characters are not valid in the path: < > : " / ? * • The mount path could also be Same as Original Path. You can select either of the options. • If you specify a non-default mount path for mounting to a standalone instance, the drive specified for mount cannot be a clustered disk. • For mounting to a clustered SQL instance, the “Same as Original Path” and alternate mount paths are supported. “Default mount path” is not supported. Instead, you can type this option to specify an alternate mount path. The root disk for the alternate mount path must be clustered and a dependency must exist for SQL server on the clustered disk.
Copy metadata files to	<ul style="list-style-type: none"> • By default, the location to copy VSS metadata files is the same as the mount path. • If the mount path is Same as Original Path, then this defaults to %SystemDrive%\AppSyncMounts\%%ProdServerName%%. • The following characters are not valid in the path: < > : " / ? * • If you are backing up the database to another media, you must backup these metadata files as well. • AppSync can integrate with third-party backup software to create tape backups of SQL Server copies. The target directory specified here must be part of the backup. <hr/> <p data-bbox="667 1331 724 1356">Note</p> <ul style="list-style-type: none"> • Metadata is not created for Non VDI copies. • VSS or VDI metadata is not generated for Crash Consistent copies.
Image access mode (during RecoverPoint mount)	<ul style="list-style-type: none"> • Logged Access: Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host. • Virtual Access with Roll: Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. • Virtual Access: Provides nearly instant access to the image; it is not intended for heavy processing.

Field	Description
Use Dedicated Storage Group	<ul style="list-style-type: none"> Applicable only for physical hosts or virtual machines with direct iSCSI part of cluster. Checked by default, enabling this option allows AppSync to enforce a dedicated VMAX, VNX, or XtremIO storage group. For XtremIO, this option applies to an XtremIO initiator group that only contains an initiator for the mount host. The storage group contains the selected mount host only for a mount and the mount will fail if you are mounting to a node of a cluster that is in a storage group shared with the other nodes. <hr/> <p>Note</p> <p>Use this option to mount the copy to a node for copy validation or backup to tape. In this scenario, you will need two storage groups. One storage group is dedicated to the passive node being used as a mount host and the other storage group is for the remainder of the nodes in the cluster. Both storage groups contain the shared storage for the cluster.</p> <hr/> <ul style="list-style-type: none"> If unchecked, AppSync does not enforce the use of a dedicated storage group for a mount and the mount will proceed. Host initiators can only belong in one initiator group in XtremIO, so use this option to ensure that you mount to a mount host that is the only host in the initiator group. <hr/> <p>Note</p> <p>Uncheck this option for manually adding the target devices as clustered storage and presenting them to clustered SQL Server instances for data repurposing and data mining.</p>

Supported mount recovery modes

The following mount recovery types are available when you are recovering a SQL database copy.

Recovery Type	Description
Recovery	Instructs the restore operation to roll back any uncommitted transactions. After the recovery process, the database is ready for use.
No Recovery	Instructs the restore operation not to roll back any uncommitted transactions. When in No Recovery mode, the database is unusable. This option is useful when the Database Administrator needs to restore one or more transaction log backups. Database is attached to the instance selected for recovery and is left in the "Restoring" state.
Standby	Restores files and opens the database in read-only mode. Subsequently, the Database Administrator can manually apply additional transaction log backups.

Recovery Type	Description
	<p>Note</p> <p>If you are restoring a database from an older version of SQL Server onto a newer SQL Server version, do not use standby mode. If you use standby, the upgrade to the newer version cannot happen and that will result in a failure of the operation.</p>
Attach Database	<p>Mounts the file system on which the database files are located, and then attaches the database to the SQL Server. The Attach Database option is only available for Non VDI and Crash Consistent copies because all the data necessary to attach the database is part of the copy. You might have to perform additional steps for full recovery of the database.</p>

Note

- Recovery, No recovery, and Standby modes are not supported for Non VDI and Crash Consistent copies.
 - Attach Database is not supported for Full or Copy SQL copies.
-

Unmounting an SQL Server copy

When you select an SQL Server copy to unmount, other copies that were mounted along with the selected copy will also be unmounted.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can unmount a copy only from a list of copies made for a database.

Procedure

1. Navigate to the Copies page from the Copy Management or Service Plan pages:
 - **Copy Management > Microsoft SQL Server** > select the server which hosts the filesystem you want to unmount, then select the database instance with the copy to unmount.
 - **Service Plans > Microsoft SQL Server**> select a service plan, then select the **Copies** tab.
2. From the list of copies, select the copy and click **Unmount** from the button in the lower part of the page.

The **Unmount Confirmation** dialog displays all the copies of other databases that were mounted along with the selected copy to be unmounted.

3. Click **Yes** to confirm the unmount of all the copies shown in the dialog.

The **Unmount** page displays the progress of the unmount operation. All copies associated with the selected copy will be unmounted.

Path mapping

The path mapping option mounts the copy to a host using a path mapping table set to user-defined locations. When you use a path mapping table, you have more control over where data is located.

You must specify the path where you want to mount a specific file system. You must provide a path map where the source file system and the target mount point is specified.

The following is a sample path mapping table for UNIX.

Source file system	Target mount path
/opt/emc/db1	/opt/emc/copies/db1
/opt/emc/db2	/opt/emc/test/db2
/opt/emc/db3	/opt/emc/reserved/db3
/prd/db1	/tst/db1

The following is a sample path mapping table for Windows.

The first two target paths, G:\ and H:\ drives must already be available on the mount host. That is, the root drive for the mount path must pre-exist before attempting a mount.

Source file system	Target mount path
D:\Test1	G:\Test1
E:\	H:\Test2
F:\Test3	I:\
L:\	N:\

Note

- If a target path is not provided for a source path, then it is mounted to a path same as the source path on the mount host.
- Ensure that you type in the absolute mount path on the target host. If the path is invalid, mount fails.
- Mount copy overrides is unavailable, if you select the mount path as Mapped path.
- For Windows, if one of the entered path is invalid, VSS import fails. Therefore, the entire mount fails. Partial failed scenarios are not supported for Windows mount.
- For Windows and NFS file systems on Unix, nested target mount points are not supported.
- Path mapping is not supported for Oracle ASM disk groups.
- Path mapping is not supported for Oracle NFS RAC database.
- Path Mapping is not applicable to metadata paths for Microsoft Exchange and Microsoft SQL Server.

Specify path mapping settings

You can specify the path where you want to mount a specific copy.

Procedure

1. On the AppSync console, go to the **Copies** page to display available copies.
2. Select a copy to mount, and then click **Mount**.
The Select Copy to Mount page of the Mount wizard launches.
3. Select the copy to mount, and click **Next**.
The Specify Mount Settings page appears.
4. On the **Specify Mount Settings** page:
 - a. Select the mount host.
 - b. From the **Mount Path** list, select **Mapped Path**.
The Path Mapping Settings link appears.
5. Click on the link to open the Path Mapping Settings window.
6. From the **Select Source Host** list, select a host.
All the file systems on the selected host are displayed in the source path column.
7. Specify the target path.
8. Click **Save** to save your settings.
If you want to set the target path for a file system on another source host, repeat steps 6 to 8.
9. Click **Reset**, to clear all the entered target paths for the selected source host.
10. Click **OK** to exit the Path Mapping window.

Note

If you change the path mapping settings, the earlier saved path mapping settings is not valid and the new path mapping settings takes precedence. Therefore, ensure that you save the path mapping settings for all the hosts before changing it.

SQL Server database restore overview

Review and consider the following sections regarding SQL Server database restore options.

These include:

- Restore considerations for databases in an Availability Group
- Affected entities during restore
- Restoring a primary database or a secondary database with failover
- Restoring a secondary database without failover
- How AppSync manages damaged SQL databases
- Restoring an SQL Server copy

- Restoring an SQL Server copy on XtremIO
- SQL restore utility (`assqlrestore`)

Note

Ensure that no virtual machine snapshots are present before protecting a datastore. If virtual machine snapshots are present, protection succeeds, but AppSync fails to perform a file or virtual machine restore.

Restore considerations for databases in an Availability Group

AppSync restores copies of primary and secondary databases. Consider the following when restoring a database in an Availability Group.

- Restore is at the LUN level and must be restored back to the source LUN that was used to create the AppSync copy.
- AppSync suspends data movement as part of the restore process.
- A database cannot be restored if it is part of an Availability Group. AppSync removes the database from the Availability Group as part of the restore process.
- AppSync does not put the database back in the Availability Group. For more information on restoring databases in an Availability Group, see "Restoring a primary database or a secondary database with failover" and "Restoring a secondary database without failover".

Affected entities during restore

When restoring from a copy, you may be prompted to restore items in addition to the ones you selected.

An affected entity is data that resides on your production host that unintentionally becomes part of a replica because of its proximity to the data you intend to protect. You can prevent affected entity situations by properly planning your data layout based on replica granularity. The granularity of a replica depends upon the environment.

If there are *affected entities* in your underlying storage configuration, the Restore Wizard notifies you of these items. The following scenarios produce *affected entities* that require you to acknowledge that additional items will be restored:

- For RecoverPoint, if the databases are in the same consistency group they become *affected entities* when the other database is protected.
- For VNXe, if the databases are in the same LUN group they become affected entities when the other database is protected.
- For Unity, if the databases are in the same consistency group they become affected entities when the other database is protected.
- For VNX/VMAX, VNX, VNXe, Unity, or XtremIO, if the databases are on the same LUN they become *affected entities* when the other database is protected.
- For VMware virtual disks, since restore involves a datastore, restore of all applications residing on the same datastore (virtual disks on the same datastore) are also *affected entities*.

If the affected entity was protected along with the database that is selected for restore, it will be restored by AppSync. Any other database that was not protected but is an affected entity will be overwritten.

AppSync calculates affected entities for the consistency groups or LUN groups of the database that is selected for restore. If the affected databases in turn partially reside

on other consistency groups or LUNs groups, AppSync does not calculate affected entities on those consistency groups or LUN groups.

Depending upon the type of affected entity, the affected databases are detached by AppSync or you must manually detach them from the SQL Server instance.

Affected entities are calculated only for the SQL Server instances where the credentials are configured. AppSync does a fresh database discovery for all these instances before calculating the affected entities.

Restoring a primary database or a secondary database with failover

Once you click the **Finish** button in the **SQL Server Restore** wizard, AppSync performs the following actions:

1. If you had selected the **Failover the Availability Group if the current role is Secondary** checkbox, AppSync verifies the health of the databases in the Availability Group that are not being restored. If they are not healthy, AppSync cannot perform the failover and the restore operation fails. You must retry the restore operation without selecting the checkbox.
2. If you had chosen to backup the transaction log, AppSync backs up the transaction log.
3. AppSync suspends data movement for all replicas of the selected database before removing all replicas of the selected database from the Availability Group.
4. If the database being restored is secondary, AppSync initiates the failover.
5. AppSync restores the LUNs of the selected database.
6. Finally, AppSync recovers the database and leaves it in the Recovery state that you selected in the **SQL Server Restore** wizard.

After AppSync completes the restore, you must perform the following steps.

Procedure

1. Restore any log backups and recover the primary database.
2. Add the database back into the Availability Group.
3. If the primary database was rolled forward so it is at the same time as the secondary database, re-join the secondary copies to the Availability Group.
4. If the primary database was not rolled forward:
 - a. Delete any secondary copies of the restored database.
 - b. Reseed and re-join the secondary database replicas to the availability group.

Note

After AppSync removes the primary database copy, the copy is in the recovered state if it is healthy. If you restored a secondary copy with failover, the primary role will have moved to another SQL Server instance. You must delete the original primary database and reseed it.

Restoring a secondary database without failover

Once you click the **Finish** button in the **SQL Server Restore** wizard, AppSync performs the following actions:

1. If you had chosen to backup the transaction log, AppSync backs up the transaction log.
2. AppSync suspends data movement for the selected secondary database replica. Replication continues to work for other replicas of the database.
3. AppSync removes the selected secondary database replica from the Availability Group.
4. AppSync restores the LUNs of the selected database.
5. Finally, AppSync recovers the database and leaves it in the Recovery state that you selected in the **SQL Server Restore** wizard.

After AppSync completes the restore, you must perform the following steps.

Procedure

1. Restore any log backups and leave the secondary database in a "NO RECOVERY" state.
2. Join the secondary database back into the Availability Group.

How AppSync manages damaged SQL databases

Damaged databases may have data files missing or damaged with their log files intact. AppSync can take tail log backups for damaged databases. A damaged database must not contain bulk-logged changes and it must not be in OFFLINE state.

If the production database is damaged and you select the **Database is damaged** checkbox during restore, AppSync backs up the tail log of the damaged database before proceeding with restore. If the damaged database is in RECOVERY_PENDING or SUSPECT state, AppSync first tries to detach the database by setting the EMERGENCY mode on it. If AppSync fails to set EMERGENCY mode on the database, it drops the database and then proceeds with the restore. Once the restore is successful, you can recover the database manually using the tail log backup.

Restoring a SQL Server copy

You can perform a restore of an SQL Server copy from the Server's Copies page, service plan's Copies page or from the Databases page.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. In the Databases page, select **Recover > Restore** .
In the Copies page, select a copy and click **Restore**.
The **SQL Server Restore** wizard launches.
2. Select the copy to restore.
Use the **Time** or **Service Plan** filters to select the appropriate copy to restore. The copies list is refreshed based on the filters selected.
3. Click **Next**.
If the selected copy has affected entities, the **Restore Warnings** page is displayed.
4. Read the warning messages for the affected databases. Select the checkbox to indicate your agreement to restore other entities along with the selected copy.

5. In the **Backup Transaction Logs** step, select **Yes** to backup logs prior to restore.
 - a. In the **Back up to** box, enter the location where the logs will be backed up. The files will bear the name of the database.
 - b. Select the **Add a file extension to the backup file name** checkbox and specify an extension for the backed up files. The default extension is .trn.
 - c. Select the **Add a prefix to the backup file names** checkbox and specify a prefix for the backed up files. The default prefix is AppSync.
 - d. Select the **Database is damaged** checkbox to backup tail log files.
 - e. Select the **Truncate the transaction logs** checkbox as required.
This checkbox is not available for selection if you selected the **Database is damaged** checkbox.
 - f. Select the **Overwrite existing backup files** checkbox as required.
This checkbox is not available for selection if you selected the **Database is damaged** checkbox.

Note

You cannot backup transaction logs for Non VDI and Crash Consistent copies.

6. Click **Next** for the **Restore Options**.
The **Restore Options** page is displayed.
7. Select the appropriate recovery or restore options.
 - a. To recover the database, choose from one of these options:
 - **Recovery:** This option is not available if you have chosen to backup transaction logs.
 - **No Recovery:** Leaves the database in a non-operational mode and requires manual intervention. For Non VDI and Crash Consistent copies, you must manually attach and recover the database.

Note

Select this option when restoring a secondary database without failover. This leaves the secondary database in the restoring state so transaction logs can be restored allowing the database to rejoin the Availability Group.

- **Standby:** If you select this option, specify the location in the **Standby file location** box where the standby files must be stored. The default path is C:\temp.
-

Note

Non VDI and Crash Consistent copies only support the **No recovery** restore option.

- b. To force a restore by overwriting the existing database, select the **Overwrite the existing databases** checkbox.

This option is not required in the normal circumstances.

This option is not available if you have chosen to backup transaction logs.

Note

You cannot overwrite the existing database for Non VDI and Crash Consistent copies.

- c. In case of Availability Group, select the **Failover the Availability Group if the current role is Secondary** checkbox to initiate a failover before restore if you are restoring to a secondary database.

A warning message is displayed that the database will be removed from the Availability Group and that you must rejoin it after AppSync restores the volume.

- d. Select the **I have read and understand the warning above, and want to continue with the restore** checkbox to acknowledge the message.

8. Click **Next**.

The **Configure Storage Options** screen appears.

- **Wait for mirror rebuild to complete** - This option is applicable for VPLEX Snap copies whose production data resides on local or distributed RAID-1 volumes. By default, this option is selected.
- **Disable VMware SRM** - Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later.
- **Perform device restore in background** - Allows you to optimize restore of VMAX devices. If you select this option, AppSync restore operation does not wait for VMAX track synchronization to complete. The production application is available instantly.

9. Click **Next**.

The **Summary** page is displayed.

10. Review the **Summary** page and click **Finish** to perform the restore.

11. In the **Results** page, click **View Details** to see progress of the different phases that are part of restoring a copy.

The last phase completed is displayed at the bottom of the list.

12. After AppSync restores a database in an Availability Group, perform additional steps as needed:
- [Restore a primary database or a secondary database with failover.](#)
 - [Restore a secondary database without failover.](#)

SQL Server restore utility (assqlrestore)

AppSync includes a SQL Server restore utility called `assqlrestore`. This section describes its function and uses.

The `assqlrestore` utility lets you restore individual SQL Server databases from a tape backup or mounted copy without reverse-syncing the target device over the source device. It can restore a database, filegroup, or file. The utility can restore to the original database or to a new database. SQL Server VDI metadata that was

created as part of the replication activity is required to restore a database using `assqlrestore`.

Note

For Non VDI and Crash Consistent copies, you cannot restore a database using `assqlrestore` because no metadata is created.

`assqlrestore` is a command line interface that you run from a command prompt window on the AppSync client. It is installed on the client as part of the AppSync installation.

Restoring an individual database from a mounted copy is especially useful when you need to recover only one database and do not want to overwrite an entire device which occurs with a normal AppSync restore. This utility supports item level restore from a mounted copy.

Assqlrestore command syntax with examples

This topics lists the command syntax for the `assqlrestore` command followed by examples of the commands.

Command syntax

The following table lists the command syntax for the `assqlrestore` command.

Table 21 `assqlrestore` Command Syntax

Option	Description
Required	
-s	SQL Server name including instance name (host\instance).
-f	Metadata filename and location (the path selected in the GUI under Copy metadata files to).
-d	Database name.
Connection Types (-E or -U)	
-E	User used for Windows Authentication (specify username)
-U	SQL Server login ID.
-P	Clear text password (used with -E and -U options).
-p	Encrypted password (used with -E and -U options).
Optional	
-r	Recover option - RECOVERY, NORECOVERY (default), or STANDBY.
-u	Undo filename, required for STANDBY
-m	Move file. Option has two parameters: logical_file_name and

Table 21 `assqlrestore` Command Syntax (continued)

Option	Description
	operating_system_file_name. Pathnames must exist. Repeat option for each file, including log files or full text catalog files. If you are restoring to a new database name, use the -m option so you do not overwrite the original files. For example: -m logicalfilename S:\existingdir \newfilename.mdf
-fg	Filegroup to restore. Repeat option for each filegroup.
-lf	Logical file to restore. Repeat option for each logical file.
-e	Displays encrypted password when unencrypted password is specified as an argument. Not used with other parameters.
-v	Verbose mode.
-q	Quiet mode. Will not ask questions.
-l <log_dir>	Creates log files in the specified directory.
-h	Help.

Example 2 Command syntax examples

Command options are case-sensitive. Refer to the "SQL Server books online" for a description of the T-SQL

Note

Command parameters have changed from the Replication Manager utility (`rmsqlrestore`)

- Using Windows authentication, restore without applying logs.

```
assqlrestore.exe -E Administrator -P password -s
sql1\instance1 -d custinfo
-f "C:\AppSyncMounts\sql1\APPSYNC_VDI_INSTANCE1_
custinfo.bin" -
r RECOVERY
```

- Restore to a new database name and move files using a SQL login and encrypted password:

```
assqlrestore -s sql1\instance1
-d custinfoTest
-f "C:\AppSyncMounts\sql1\APPSYNC_VDI_INSTANCE1_
custinfoTest.bin"
-r RECOVERY
```

Example 2 Command syntax examples (continued)

```
-m custinfo_Data S:\custinfoTest.mdf
-m custinfo_Log T:\custinfoTest.ldf
-U sa -p 1EMC_4roJdyU5;x
```

- To get the encrypted password:

```
assqlrestore -e <unencrypted_password>
```

Restoring an SQL Server database with `assqlrestore`

The basic steps to restore a database are provided here. You may need additional steps but use these as a framework.

Before you begin

Log in to the SQL Server system as a user with Administrator rights, then back up the SQL Server transaction log.

Procedure

1. Take the target SQL Server database offline.
2. Restore the database files (.ldf, .ndf, and .mdf) from tape, or copy them from a mounted replica. You can copy them over the original files or to a new location.
3. Open a command prompt window and cd to: C:\Program Files\EMC\AppSync Host Plug-in
4. Run the `assqlrestore` command.

Refer to the Assqlrestore command syntax with examples section for sample commands. The basic command syntax is:

```
assqlrestore -s <SQLservername> -d <databasename> -f
<metadata file> -r <recovery_type>
```

5. If required, apply transaction logs and recover the database.

Restoring a file or filegroup with the SQL Server restore utility

Learn how to restore a file or filegroup with the SQL Server `assqlrestore` utility.

Before you begin

Be sure you understand how restore of files and filegroups behave in SQL Server before proceeding.

Note

You cannot use the `assqlrestore` utility to restore a SQL Server filegroup if the filegroup name contains non-ASCII characters.

Log in to the SQL Server system as a user with Administrator rights, then back up the SQL Server transaction log. For file or filegroup restore, the database must be online.

Procedure

1. Open a command prompt window and cd to: C:\Program Files\EMC\AppSync Host Plug-in
2. Run the `assqlrestore` command.
 - a. When `assqlrestore` displays the restore command that it is about to run, verify with **Y** if it is correct.
 - b. When `assqlrestore` prompts, restore the files you are recovering, enter **Y** to continue.

To restore two files, for example, run:

```
assqlrestore -s <SQLservername> -d <databasename>  
-f <metadatafile> -lf <logical_filename1>  
-lf <logical_filename2> -r norecovery
```

To restore two filegroups, run:

```
assqlrestore  
-s <SQLservername>  
-d <databasename>  
-f <metadatafile>  
-lf <logical_filename1>  
-fg <logical_filegroupname1>  
-fg <logical_filegroupname2>  
-r norecovery
```

Do not use the quiet mode for a file or filegroup restore. You can use `-lf` and `-fg` in the same restore command.

CHAPTER 7

Protect Oracle

This chapter includes the following topics:

- [Overview of Oracle support](#)..... 160
- [Protecting a database](#)..... 174
- [Service plan summary and details](#)..... 178
- [Mount an Oracle copy](#)..... 185
- [Restoring an Oracle copy](#)..... 194

Overview of Oracle support

Use AppSync to create and manage application consistent (using hot backup mode) and crash consistent (without hot backup mode) copies of Oracle® databases. The copies can be used for mount (with/without recovery) and restore.

The *AppSync Support Matrix* on <https://elabnavigator.emc.com/eln/extendedSupport> is the authoritative source of information on supported software and platforms..

AppSync supports:

- Oracle - (Standalone and Oracle Real Application Cluster) and on Linux and AIX.
- Oracle installations on physical hosts as well as virtual machines (with pRDMs and Vdisks) - There is no support for RDMs in virtual mode.
- Oracle databases residing on NFS file systems with VNX File, VNXe File, and Unity File storage.
- Oracle databases residing on ASM disks.
- Oracle databases residing on file systems.
- RMAN cataloging of databases to a remote catalog.
- Repurposing of Oracle database copies.

Note

AppSync does not support file systems or ASM diskgroups on Linux operating system devices which are not full block devices (such as /dev/sdc) or primary first partition (such as /dev/sdc1).

Oracle permissions

These permissions are required for AppSync to work with Oracle.

- Root or sudo access to Oracle production server and mount server.
- When connecting to Oracle databases, AppSync uses a bequeath connection and always connects as SYSDBA.
- When connecting to Oracle ASM, AppSync uses a bequeath connection and always connects as SYSASM.

Red Hat Cluster Services Integration with AppSync

AppSync can work with standalone Oracle databases that are configured to failover from node to node in an RHCS (Red Hat Cluster Services) environment.

Overview

During a replication process, if the node you used to create a service plan is not accessible, AppSync runs the replication on another node in the cluster. If the node you used to create the original copy is not accessible, AppSync does not rely on the Virtual IP of the Oracle service group. Therefore, ensure that you register all nodes in the RHCS cluster in the AppSync server for database replication.

From a restore perspective, AppSync can only restore to the node where the copy was originally created, therefore the original node must be active, otherwise the restore process fails.

Requirements

Review the following requirements to use a standalone database that fails over as part of an RHCS cluster:

- The AppSync host plug-in must be installed on all nodes of the cluster.
- The IP resource must be configured in the Oracle service group for the clustered database.
- If a failover occurs while running a replication or restore process, the operation fails. Node failover should occur before running the service plan, before the start of a replication, or start of a restore.
- The Oratab file should have an entry for all possible SIDs that can run on the specified node (passive and active instances).
- The package `sg3_utils`, which contains utilities for accessing devices that use SCSI command sets, must be installed on all nodes.

Mount considerations

- The mount host must not be part of the RHCS cluster.
- The mount host run the same Oracle version as the copy host.
- The AppSync host plug-in must be installed on the mount host.
- The package `sg3_utils`, which contains utilities for accessing devices that use SCSI command sets, must be installed on the mount host.

Restore considerations

- AppSync can only restore to the node where the copy was originally created, therefore the original node must be active. Otherwise, the restore process fails, and corrupts the database. The console provides a detailed warning message before the restart of the restore.
- To perform a restore in an RHCS environment, follow these steps:
 1. Perform the restore.
 2. Start the mount instance.
 3. Perform a manual recovery.
 4. Shut down abort database.
 5. Enable the Service Group.

Oracle Data Guard support

AppSync supports an Oracle Data Guard configuration for a primary (source database) and a physical standby (target database) which is open in active or passive/non-active mode.

There are three types of standby databases:

- Physical standby
- Logical standby
- Snapshot standby

All three configurations can be opened in one of the following modes:

- Active standby mode—Standby database in read-only or read/write mode
- Passive/non-active standby mode—Standby database in mounted mode

AppSync currently only supports Data Guard physical standby configuration in active or non-active mode.

When a physical standby database is open in active mode, the standby database can be opened in read-only mode while logs are applied. This action allows you to query the database for information while Data Guard applies logs.

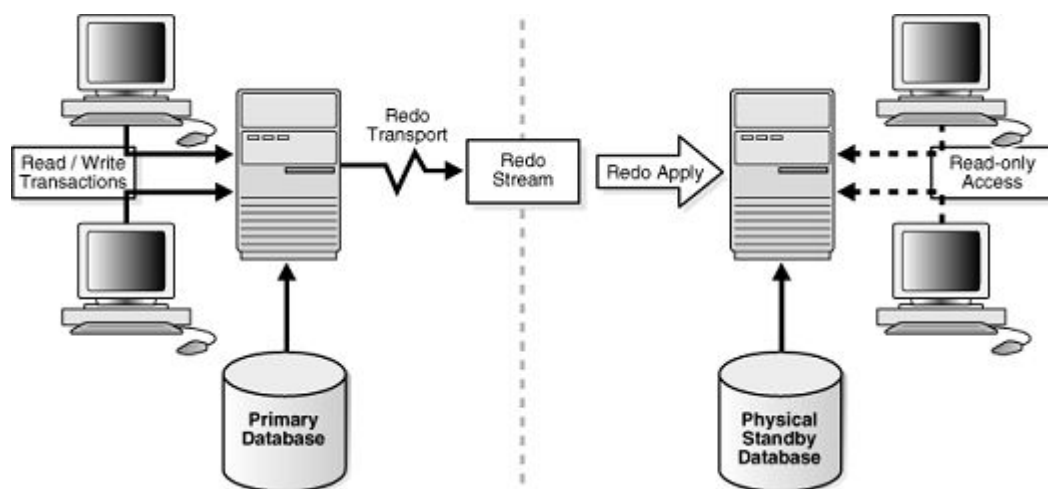
Snapshot and logical standby configurations also allow the database to be open in read/write mode. A passive/non-active setting means that the database can start in mounted mode and logs can be applied in the background.

Physical standby

In a physical standby environment, archive logs are applied when they are received. A physical standby has a 1:1 mapping of the file and storage layout from primary to standby. A physical standby database can be open in both read-only or mounted mode which means it can be either an active or passive/non-active configuration.

The following diagram displays a typical primary/standby (source/target) Data Guard configuration:

Figure 5 Physical standby environment



Copy Management

On the AppSync console, go to the **Copy Management > Oracle** page. A Data Guard relationship column now displays. If you have an existing Data Guard relationship, you can view two databases that are part of a Data Guard configuration. One database is the primary database and one is the physical standby (non-active) database.

Review the following copy management considerations for Data Guard:

- To protect a primary Data Guard database (source database), create a copy like any other standalone database. You can take a hot backup copy.
- For protection of an active standby Data Guard database (Target Database): Protection in hot backup mode of an active standby database is not allowed because the standby database is in read-only mode. Also, the standby database contains up-to-date archive logs and is an exact copy of the primary and does not require archive logs to be copied for recovery. You can however take a non-hot backup copy of a Standby database.
- For protection of a passive/non-active standby Data Guard database (target database): A passive/non-active standby database operates the same way as an active standby database. Hot backup copy of the database is not allowed. The difference here is that the copy is created from the mounted database without opening the database in read-only or read/write mode.
- Creating a copy of a mounted database only succeeds for a passive/non-active Data Guard standby database in mounted state. Standalone Oracle databases that

are mounted cannot be protected. They appear as offline on the database protection page of the console.

Mount and restore (recover)

Review the following mount and restore considerations for Data Guard:

- For a primary database (Source database): Mount and restore operate the same way with a Primary Data Guard database as any Oracle Standalone database. If you use RAC to configure the Primary database then the RAC mount/restore rules for AppSync apply.
- For an active standby database (target database): Mount and restore operate the same for an active standby Data Guard database as any other Oracle standalone database. If the standby database is configured using RAC then the RAC mount/restore rules for AppSync apply.
- For a passive/non-active standby database (target database): Mount and restore operate the same for a passive/non-active standby Data Guard database as any other Oracle standalone database. If the standby database is configured using RAC, then the RAC mount/restore rules for AppSync apply.

Note

If you mount and restore either a primary or standby database, the database appears on the console as a standalone Oracle database. No Data Guard configuration persists.

Repurposing (copy or a copy) Data Guard databases

For general repurposing information, refer to the AppSync user documentation.

Review the following repurpose considerations for Data Guard:

- Repurposing a primary database (source database): Repurposing operates the same for a primary Data Guard database like any Oracle standalone database.
- Repurposing an active standby database (target database): Repurposing operates the same for an active standby Data Guard database as any Oracle standalone database. You cannot hot backup a standby database for a repurposed copy.
- Repurposing a passive/non-active standby database (target database): repurposing operates the same for a passive/non-active standby Data Guard database as any Oracle standalone database. You cannot hot backup a standby database for a repurposed copy.

Restore Data Guard databases

Restore for a primary database (source database): Restore for a primary Data Guard database operates the same way for any Oracle standalone database. Manually recover the database and then resynchronize the primary and standby databases after the AppSync restore process completes.

Veritas Cluster Services integration

AppSync can work with standalone Oracle databases that are configured to failover from node to node in a VCS (Veritas Cluster Services) environment.

Introduction

During a replication process, if the node that was used to create the service plan is not accessible, AppSync runs the replication on another node in the cluster. AppSync does not rely on the Virtual IP of the Oracle service group. Therefore, register all nodes in the VCS cluster to the AppSync server before you replicate the database.

From a restore perspective, AppSync can only restore to the node where the copy was originally created. The original node must be active, otherwise the restore process fails.

Requirements

The following are the requirements for using a standalone database that fails over as part of a VCS cluster:

- Install the AppSync host plug-in on all nodes of the cluster.
- Configure the IP resource in the Oracle service group for a clustered database.
- If a failover occurs while running a replication or restore process, then the operation fails. Node failover occurs before running a service plan, before the start of a replication, or a restore.
- The Oratab file should have an entry for all possible SIDs that can run on the specified node (passive and active instances).
- Ensure `tnsnames.ora` files on all nodes contain entries of all standalone instances, including the virtual IP address of the Oracle service group (per Symantec documentation).
- The following files should be accessible to all nodes on the cluster where the database runs:
 - Database `init/spfile`
 - Password file
- Install package `sg3_utils`, which contain utilities to access devices that use SCSI command sets, on all nodes.

Mount considerations

- The mount host must not be part of the VCS cluster.
- The mount host requires installation of VxVM Storage Foundations minimum 6.1.
- The package `sg3_utils`, which contain utilities for accessing devices that use SCSI command sets, must be installed on the mount host.

Restore considerations

AppSync can only restore to the node where the copy was originally created, therefore the original node must be active. To perform a restore in a VCS environment, follow these steps:

1. Freeze the Oracle service group: `>hagrp -freeze <service_group_name>`
2. Perform the restore.
3. Start the instance.
4. Perform a manual recovery.
5. Open the database.
6. Unfreeze the Oracle service group: `>hagrp -unfreeze <service_group_name>`.

Note

AppSync can only restore to the node where the copy was originally created, therefore the original node must be active. Otherwise, the restore process fails, and leaves the database in a corrupt state. The console provides a detailed message warning you of this scenario before the restart of the restore.

PowerHA (HACMP) cluster integration

AppSync can work with standalone Oracle databases that are configured to failover from node to node in an IBM® PowerHA (HACMP) cluster environment.

Introduction

AppSync protects the database on the node where the current state is active before the Service Plan run. AppSync relies on the service label IP of the Oracle database resource group.

Note

Starting in AppSync 3.1, AppSync mandates the use of service label IP for Oracle database protection. The copies created using AppSync versions earlier than 3.1 can be mounted or restored even after an upgrade to AppSync 3.1. After an upgrade, scheduled service plan run fails, if the service label IP is not registered with the AppSync server.

Consider the following when protecting an application that fails over as part of a PowerHA (HACMP) cluster:

- The nodes of the PowerHA cluster must be registered with AppSync before registering the service label IP.

Note

You must register the objects with AppSync in the following order:

- Active node
 - Passive node
 - Service label IP of the Oracle database
-
- The service label IP/name must be configured for the Oracle database resource group.
 - If the resource group has multiple service label IPs configured, register only one service label IP with the AppSync server.
 - If a failover occurs while running a replication or restore process, the operation fails. Node failover must occur before running the service plan, or at the start of a restore.
 - The Oratab file entry must be the same on all the cluster nodes.
 - The following files should be accessible to all nodes on the cluster where the database runs:
 - Database `init/spfile`
 - Password file

Mount considerations

- The mount host must not be part of the PowerHA cluster.
- The AppSync host plug-in must be installed on the mount host.
- Mount to service label IP is not supported.

Restore in a PowerHA environment

Perform a restore. After a restore, the volume group is not concurrent. You must manually make them concurrent before performing a host or file system rediscovery.

Note

AppSync can only restore to the node where the copy was originally created, therefore the original node must be active, otherwise the restore process fails, and leaves the database in a corrupted state. The console provides a detailed message warning you of this scenario before the restart of a restore.

Post restore procedure in a PowerHA environment

Learn how to perform manual steps with a restore in a PowerHA environment after a restore.

After restore, a file system mounts to the production host in non-concurrent mode. Remove the file system from the resource group, make it a concurrent volume group, and then add it back to the resource group.

Perform these steps on an active node:

Procedure

1. Unmount the file system.
2. Type the `varyoffvg` command.
3. Type the `varyonvg` command with `-c` option (to make it concurrent).

Verification:

The `lspv` command must show `vg` as concurrent on both nodes as follows:

```
node 2
hdiskpower8      00c2bfb0f1ee76ca  oradata concurrent
hdiskpower9      00c2bfb0f1f434e3  oralogs concurrent

node 1
hdiskpower18     00c2bfb0f1ee76ca  oradata concurrent
hdiskpower19     00c2bfb0f1f434e3  oralogs concurrent
```

4. Add the file system back to the resource group.
5. Verify and synchronize the configuration.

Prerequisites and supported configurations

Learn about prerequisites and supported configurations for Oracle with AppSync. Included is information about supported device configurations, Oracle on file systems, logical volume managers and ASM-based storage, RecoverPoint consistency group-based storage, Linux and AIX-based configurations including sudo user, and support for virtualization setups.

AppSync can create application-consistent (using Oracle hot backup) and crash-consistent (without hot backup) copies. For AppSync to create application-consistent copies of Oracle databases, the data files, fast recovery area, and archive logs must not share the file system, volume group, ASM disk group, RP consistency group, or data store. If the Oracle configuration is such that the data files and archive logs share any of these groupings, then AppSync can create crash-consistent copies for such databases.

During copy creation, if hot backup mode is not selected, AppSync creates crash consistent copies, and does not quiesce the database. You must use this method to create copies, if you have archive logs or fast recovery area sharing the same file system, volume group, ASM disk group, RP consistency group, or data store as the data files and/or control files and/or redo logs.

If all Oracle files, including archive logs are on one disk group, AppSync can protect that database without hot backup mode. The copy is crash consistent.

Note

Read-only Oracle databases can only be protected in no hot backup mode.

AppSync does not provide an option to protect the archive log location separately (as part of the Oracle service plan). However, if the location is shared with other database components, use init overrides (see the Custom initialization parameters field under Mount options for details) and point to that location during mount with recovery. Ensure that you specify the correct path in init overrides, especially if ASM disk group rename or alternate path mount is used.

If the database is running in NOARCHIVELOG mode, do not select the hot backup mode option when creating copies.

When using VNX, ensure that all consistency groups are VNX consistency groups. When using ViPR controller, ensure that all consistency groups are ViPR consistency group. Additionally, the archive log files must be on a different CG from the rest of the database files.

Note

Database files refer to data files, and/or control files, and/or redo logs. Archive log files refer to archive log destination and/or Fast Recovery Area .

ASM rebalancing

AppSync switches off ASM rebalancing power factor before taking a snapshot or clone of the underlying disks and turns it back on, after the copy is created. This is to ensure that no automatic rebalancing occurs during protection. AppSync disables ASM rebalancing using the `alter system set asm_power_limit=0` command. AppSync checks for active rebalancing operations only on ASM disk groups that are involved in protection.

```
select count(group_number) from gv$asm_operation where
operation='REBAL' and group_number in <group_number of DGs in
protection set>
```

The value returned must be 0 (which means that no disk group involved in protection is undergoing a rebalance operation) for AppSync to continue. If there are disk groups undergoing rebalancing, AppSync polls for the operation to terminate for a maximum of one hour, after which protection fails.

AppSync does not issue a manual rebalance operation per disk group, and it cannot disable ongoing manual rebalancing operations issued outside AppSync. Manual rebalancing operations might take a long time to complete. In order to avoid timeout failures during protection, you must ensure that no manual rebalancing operations are active prior to a Service Plan run.

AppSync also enables and disables ASM rebalancing for no hot backup copies.

Removal of Oracle deleted database

When you remove or delete a database that has copies associated with it, the database is marked for pending delete.

Note

- For Oracle RAC, all instances of the database must be removed.
 - A database can be removed only if the entry for the database is absent in the `/etc/oratab` file.
 - Subscribed or protected databases are marked for pending delete only if all the instances are removed.
-

If the database is recreated with the same name after deletion, and rediscovery is performed after the database has been recreated, AppSync displays duplicate entries with the same database name. This is not a concern if you intend to retain the copies of the deleted database. However, if you do not intend to retain the copies of the deleted database, host or database rediscovery must be performed immediately after the removal of an Oracle database from a cluster or host.

Do the following to remove the duplicate entries:

1. Comment out the `/etc/oratab` entry for the database using a `#` at the beginning of the entry.
2. Remove all copies or subscriptions of the deleted database, and rediscover host or database in the AppSync GUI. If the created database has copies or subscriptions, it is marked pending delete. Otherwise, it does not show up after this step.
3. Uncomment the entry that was commented in step 1.
4. Rediscover host or database again in the AppSync GUI. If the database was marked pending delete in step 2, the pending delete flag is cleared. Otherwise, a single database entry for the new database is displayed in the GUI.

Oracle on file system-based storage configurations

Some examples of Oracle configurations for which AppSync can offer both app-consistent as well as crash-consistent copies follow:

- Single database: database files on, for example, `/data`; archive log files on, for example, `/archive`.
 - Multiple databases sharing single archive log location: for example, Database 1 on `/db1`, Database 2 on `/db2`, archive logs on `/arch`.
 - Multiple databases sharing data location and archive log locations: for example, Database 1, 2, 3 files on `/data`, database 1, 2, 3 archive log locations on `/archive`.
 - Affected databases scenario: Two file systems on one volume group with two more file systems on another volume group, such that one Oracle database has data on `fs1` in `vg1` and logs on `fs1` on `vg2` and second Oracle database has data on `fs2` on `vg1` and logs on `fs2` on `vg2`.
-

Note

AppSync does not support the following configuration: one oracle database has data files on `fs1` in `vg1` and logs on `fs1` on `vg2`, and a second Oracle database has data files on `fs2` on `vg2` and logs on `fs2` on `vg1`.

Oracle on logical volume managers-based storage configurations (LVM/VxVM)

- Single database: Database files on a volume in, for example, `datavg`, and then archive log files in a volume on, for example, `archvg`.

- Multiple databases sharing single archive log location: Database 1 files on a volume in, for example, `data1vg`, and Database 2 files on a volume in, for example, `data2vg`, and then archive logs in a volume on, for example, `archvg`.
- Multiple databases sharing data location and archive log locations: Databases 1, 2, 3 files in a volume on, for example, `datavg`, and then Database 1, 2, 3 archive log locations in a volume on, for example `archvg`.

Oracle on ASM-based storage configurations

- Single database: Database files on, for example, `diskgroup +data`, then archive log files on, for example, `diskgroup +arch`.
- Multiple databases sharing a single archive log location: Database 1 files on, for example, `diskgroup +data1`, and database 2 files on, for example, `diskgroup +data2`, then archive logs on, for example, `diskgroup +fra`.
- Multiple databases sharing data location and archive log location: Database 1,2,3 files on, for example, `diskgroup +data1`, and database 1,2,3 archive logs on, for example, `diskgroup +data2`.

Oracle on RecoverPoint consistency group-based storage

- Single database: Database files on LUNs in RP consistency group, for example, `DATA2CG` and archive log files in RP consistency group, for example, `ARCHCG`.
- Multiple databases sharing single archive log location: Database 1 files on LUNs in RP consistency group, for example, `DATA1CG`, then database 2 files on LUNs in RP consistency group `DATA2CG` and then archive log files in RP consistency group, for example, `ARCHCG`.
- Multiple databases sharing data location and archive log locations: Database 1, 2, 3 files on LUNs in RP consistency group, for example, `DATA2CG`, then database 1, 2, 3 archive logs on LUNs in RP consistency group, for example, `ARCHCG`.

Oracle on datastore-based storage layouts

- Single database: Database files on vDISKs from data store, for example, `DATADS` and archive log files on vDISKs from data store, for example, `ARCHDS`.
- Multiple databases sharing single archive log location: Database 1 files on vDISKs from data store, for example, `DATA1DS`, then database 2 files on vDISKs from data store `DATA2DS` and then archive log files on vDISKs from data store, for example, `ARCHDS`.
- Multiple databases sharing data location and archive log locations: Database 1, 2, 3 files on vDISKs from data store, for example, `DATADS`, then database 1, 2, 3 archive logs on vDISKs from data store, for example, `ARCHDS`.

Oracle on VNXe LUN Group-based storage

The following configurations are supported:

- Single database: Database files on LUNs in VNXe LUN group, for example, data files in LUN Group `DATALUNGRP` and archive log files in LUN group `ARCHCG`.
- Multiple databases sharing single archive log location: Database 1 files on LUNs in VNXe LUN group, for example, `DATA1LUNGRP`, then database 2 files on LUNs in VNXe LUN group `DATA2LUNGRP`, and then archive log files in VNXe LUN group, for example, `ARCHLUNGRP`.
- Multiple databases sharing data location and archive log locations: Database 1, 2, 3 files on LUNs in VNXe LUN group, for example, `DATALUNGRP`, then database 1, 2, 3 archive logs on LUNs in VNXe LUN group, for example, `ARCHLUNGRP`.

Oracle on Unity Consistency Group-based storage

The following configurations are supported:

- Single database: Database files on LUNs in Unity consistency group, for example, data files in consistency group DATALUNGRP and archive log files in consistency group ARCHCG.
- Multiple databases sharing single archive log location: Database 1 files on LUNs in Unity consistency group, for example, DATA1LUNGRP, then database 2 files on LUNs in Unity consistency group DATA2LUNGRP, and then archive log files in Unity consistency group, for example, ARCHLUNGRP.
- Multiple databases sharing data location and archive log locations: Database 1, 2, 3 files on LUNs in Unity consistency group, for example, DATALUNGRP, then database 1, 2, 3 archive logs on LUNs in Unity consistency group, for example, ARCHLUNGRP.

Supported virtualization configurations

AppSync supports protection, mount, and restore of Oracle databases on vDisks in standalone and RAC.

AppSync does not support configuration where data and archive logs are on mix of RDM and VDisks.

Considerations:

- For Oracle databases on vDisks on VMs on ESX:
 - If your vDisk reside on ESX 5.0, disable the `VMFS3hardwareaccelerated` locking flag on the ESX that is hosting the VMs hosting the Oracle databases on vDisks.
 - If ATS locking is enabled for VMFS3/5 datastore, AppSync datastore mount fails. This is not applicable for ESX versions later than 5.0.
- To run SCSI commands from AppSync, set `disk.EnableUUID` on the VM.
- Ensure your VM datastore does not share the same VMFS as your Oracle databases

Note

AppSync does not support the following configurations:

- ASM database on VXVM volume groups
 - ASM database on VXDMP devices
 - ASM database on raw devices under the control of VXVM
 - Non-ASM database on Native LVM volume group residing on VXDMP devices
-

Support for Oracle on VMware virtual disks

You can protect, mount and restore Oracle standalone and clustered databases residing on VMware virtual disks.

Consider the following information when working with Oracle and VMware virtual disks.

- For successful mapping, add the vCenter to the AppSync server and then perform discovery before adding the Oracle host. Otherwise you must rediscover the Oracle host after adding the vCenter.
- For successful protection, log files and database files must reside on virtual disks. There cannot be a combination of physical and virtual storage.

- AppSync **does not** support:
 - NFS datastores
 - Protection of Oracle databases across virtual machines sharing the same datastore
- Production VMWare virtual disk with multi writer option enabled is not supported. Protection might succeed, but mount fails.
- To perform Oracle mount and recovery to a virtualized host, you need VMware permissions to modify the VMware configuration of the mount VM (create RDM / SCSI adapter), as well as rescan datastores/VMFS.

Refer also to [Oracle vDisk restore with affected entities](#).

Support for VIO vSCSI

Oracle with AIX LPARs can now also use "virtual" connections to the storage.

Overview: Support for VIO (Virtual I/O disk) vSCSI

Previously, AppSync supported Oracle on AIX physical machines and on AIX virtual machines (LPARs) that use physical or NPIV connections to the array storage. AppSync now supports Oracle AIX LPARs with virtual connections.

All supported applications and use cases for AIX Hosts using physical or NPIV storage connections are now also supported on VIO VSCSI devices. Two restrictions apply to this support:

- Mounting of replicas must be done to mount hosts using physical or NPIV storage connections. Mounts cannot be created as virtual disks .
- The VIO Server must map whole raw disks to the VIO Clients. Do not map logical volumes from the VIO Server.

In addition AppSync can coexist with AIX Live Partition Mobility. AppSync will continue to protect and repurpose applications after the migration of a client partition to a new managed server.

Supported versions

When referring to an AppSync support matrix, AIX Virtual I/O disks are supported as a valid virtual disk type known as Virtualization Server Solutions.

Oracle supported configurations

The following table describes the Oracle supported configurations.

Table 22 Oracle supported configurations

Oracle Features/ Environments	XtremIO	VMAX	VNX	VNX file	VNXe	Unity	Unity File	VPLEX	RP	ViPR
Oracle Standalone	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Oracle on file systems	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Oracle on ASM	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Oracle RAC with NFS (Non ASM)	No	No	No	Yes	No	No	Yes	No	N.A	No
Oracle RAC with ASM	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Oracle Dataguard (Primary and Secondary)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 22 Oracle supported configurations (continued)

Oracle Features/ Environments	XtremIO	VMAX	VNX	VNX file	VNXe	Unity	Unity File	VPLEX	RP	ViPR
Host cluster support for Oracle standalone (PowerHA - AIX and VCS/ RHCS -Linux) ^a	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Hot backup mode ^b	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No hot backup mode/crash consistent	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Oracle on physical and virtual machines (with pRDMs and Vdisks) - no support for RDMs in virtual mode ^c	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Oracle databases residing on NFS file systems with VNX, Unity, or eNAS	N.A	N.A	N.A	Yes	N.A	N.A	Yes	N.A	N.A	N.A
Oracle with AIX LPARs - virtual connections and physical or NPIV connections	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Repurposing of Oracle databases	Yes	Yes	Yes	No	No	Yes	No	Yes	Yes	No
mknode with ASM ^d	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
UDEV with ASM ^e	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
ASMLib with ASM ^f	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Mounting Oracle standalone to standalone	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mounting standalone ASM to standalone ASM	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Mounting RAC NFS (non ASM) to alternate RAC NFS (non ASM)	N.A	N.A	N.A	Yes	N.A	N.A	Yes	N.A	N.A	N.A
Mounting RAC ASM to alternate RAC ASM	Yes	Yes	Yes	N.A	Yes	Yes	N.A	Yes	Yes	Yes
Mounting to production RAC as a cluster	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mounting back to production RAC as a single instance/non-clustered	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RMAN	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RMAN with BCT	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 22 Oracle supported configurations (continued)

Oracle Features/ Environments	XtremIO	VMAX	VNX	VNX file	VNXe	Unity	Unity File	VPLEX	RP	ViPR
ASM RAC or standalone databases created using ASM Filter Driver (ASMFD)	Yes	Yes	Yes	No	Yes	Yes	No	Yes	Yes	Yes
Mixed layout of an Oracle database which has data files, control files, and redo logs on ASM disk groups and archive logs and/or FRA on file systems (such as ext4 on Linux)	Yes	Yes	Yes	No	Yes	Yes	No	Yes	Yes	Yes

- a. RHCS is applicable for RHEL and VCS for both SuSE and RHEL.
- b. See the Prerequisites and supported configurations section for information on database layout.
- c. This is only applicable for LINUX.
- d. This is only applicable for AIX.
- e. This is only applicable for LINUX.
- f. This is only applicable for LINUX.

AppSync does not support the following Oracle environments:

- Oracle on any cluster file systems (such as ACFS, OCFS, GFS, GFS2, QFS, and so on)
- Cold backup
- Nested file systems on NFS
- Mounting to different OS versions and different Oracle versions
- Oracle RAC one node database
- Mix of data and archive logs on RDM and Vdisks
- Oracle on VMware virtual disks on NFS datastores
- Multiple databases residing on different virtual machines sharing the same datastore (for example, VM1 with DB1 and VM2 with DB2)
- ASM mounting to non-ASM
- ASM database on VXVM volume groups
- ASM database on VXDMP devices
- ASM database on raw devices under the control of VXVM
- Non-ASM database on Native LVM volume group residing on VXDMP devices
- Pluggable databases (PDB)
- Oracle Flex cluster
- Oracle Flex ASM
- Oracle GoldenGate
- ASM Dynamic Volume Manager (ADVM)
- Oracle Multitenant
- Oracle RAC database on NFS 4 file system is not supported by Oracle

Protecting a database

To protect a database, subscribe it to an AppSync service plan.

You can protect objects in different ways from different places in AppSync:

- Select **Subscribe to Plan and Run** when you want to protect a selected database immediately. The service plan is executed for that database alone.
- Select **Subscribe to Plan** when you want to schedule protection for later. Protection for databases that are part of a service plan is executed at a scheduled time.
- Select an appropriate service plan from **Create copy** using a plan in the database Copies page.
- Select **Run** from the Oracle Service Plans page to run the entire plan immediately.

Discovering databases

To keep AppSync up-to-date, you should discover databases on the Oracle server when there is creation, deletion, or renaming of databases.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Oracle** to display the Databases page.

Only databases that are started and are in an open state show up as online on the databases page. Databases that do not have an entry in the `/etc/oratab` file as well as shutdown databases do not appear.

2. From the **Discover Databases** drop-down menu on the bottom left of the screen, click **On Server**, and then click on the desired server where the database you want to discover resides.

Subscribe a database to a service plan

You can subscribe a database to a service plan and run the service plan immediately, or schedule the service plan to run at a later time.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Oracle**.
2. Select one or more Oracle databases.
3. From the **Protect** popup button, select the appropriate service plan, for example:

Table 23 Service plan protection options

Option	Description
Subscribe to Plan and Run	To subscribe the database for protection and run the plan immediately for any selected database(s).
Subscribe to Plan	To subscribe the database for protection. Protection for all databases that are part of the service plan is executed at the scheduled time.

Unsubscribe database from a service plan

When you unsubscribe an individual database from a service plan, AppSync retains all existing database copies; only further protection will be removed.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Oracle**.
2. Select the database to unsubscribe from a service plan.
 - Select the plan to unsubscribe from: **Protect > Unsubscribe from Plan**. Only plans to which the database was subscribed appear in the popup list.
 - To unsubscribe from all service plans, select **Unsubscribe from Plan > All**.

Oracle copies page

You can see details of a copy from the Copies tab of the Service Plan. The list of copies can be filtered by time of creation, and by service plan.

Table 24 Copy page fields

Column	Description
Status	<ul style="list-style-type: none"> • Green: successful • Yellow: completed with errors • Red: failed
Name	Name of the copy. The copy name is the time AppSync created it.
Service Plan	<p>Name of the service plan associated with the copy. In the case of a repurpose plan, select a copy and click on the Repurpose link to edit it.</p> <hr/> <p>Note</p> <p>Each copy is associated with a unique repurposing service plan.</p> <hr/>

Table 24 Copy page fields (continued)

Column	Description
Label	Label assigned to the copy in case of repurposing.
Application Consistent	<ul style="list-style-type: none"> • Yes, if database was successfully put in hot backup mode while creating the copy. • No, if hot backup mode was not selected in the Create Copy phase. • No, if hot backup mode was selected and database failed to go into hot backup mode.
Mount Status	Status of the copy: mounted or not mounted. If mounted, the name of the mount host displays.
Recovery Status	Was copy recovered post mount or not. Values are: <ul style="list-style-type: none"> • Not Recovered - copy was not mounted or copy was a file system mount. • Successful - recovery was successful. • Failed - recovery failed.
Copy Type	<ul style="list-style-type: none"> • Local Bookmark • Remote Bookmark • VNX Snap, VNXeSnap • VNX File Snap • VNXe File Snap • Unity Snap • Unity File Snap • VMAX Clone • VMAX Snap • VMAX V3/VMAX AFA: SnapVXSnap, SnapVXClone • XtremIO Snap • ViPRSnap
Generation	First or second generation copy - for repurposing
Source	Production database (for first generation copy) or a copy of a copy (second generation) copies.
Site	RecoverPoint site information.
Storage System	Array serial number/name

Viewing database copies

Follow these steps to view an Oracle database copy on the AppSync console.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Oracle**
2. Click a database to view existing copies of the database.

You can see details of a copy from the Copies tab of the Service Plan. The list of copies can be filtered by time of creation, and by service plan.

Creating a database copy from the Copies page

Create a copy of a database by subscribing it to an AppSync Oracle service plan from the **Copies** page.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > Oracle**
2. Click a database to view existing copies.
3. From the **Create a copy using a service plan** list, select the appropriate service plan.

The service plan runs immediately for the selected database.

Expiring a database copy on demand

Expiring a database copy removes it from the AppSync database and can free up storage, depending on the replication technology and copy state.

Before you begin

This operation requires the Data Administrator role in AppSync.

Expiring a copy that was made with RecoverPoint does not remove the corresponding bookmark from RecoverPoint itself.

Procedure

1. Navigate to **Copy Management > Oracle**.
2. Click the desired database for copy expiration.
3. From the **Copies** page, select one or more copies to expire.
You can also perform this action from the service plan's **Copies** tab.
4. Select **Expire** from the row of buttons on the lower part of the screen.
5. Verify that you selected the appropriate copy, and any associated copies that are also listed and confirm.

Service plan summary and details

The service plan **Settings** tab shows the name, description, schedule, and status of the service plan.

Click the phases for detailed service plan settings and other tabs for information about subscriptions, lists of copies and events generated by the plan.

Review [Overview: Service Plan](#) for more service plan copy information.

Service plan schedule

The schedule of a service plan is set in the **Plan Startup** phase. The **Startup Type** (scheduled or on demand) determines whether the plan is run manually, or configured to run on a schedule.

Options for scheduling when a service plan starts include:

- Specify a recovery point objective (RPO).
 - Set an RPO of 30 minutes or 1, 2, 3, 4, 6, 8, 12, or 24 hours
 - Set minutes after the hour in 5 minute intervals.
 - Default RPO is 24 hours.
- Runs every day at specific times.
 - Select up to two different times during the day.
 - Select minutes after the hour in 5 minute intervals.
 - There is no default selected.
- Run at a certain time on selected days of the week.
 - You can select one or more days of the week (up to seven days).
 - There is no default for day of the week. Default time of day is 12:00 AM.
- Runs at a certain time on selected days of the month.
 - Select one or more days of the month (up to all days).
 - Select one time of day. Available times are at 15 minute intervals.
 - Default is the first day of the month.

Overriding service plan schedules

You can set individual schedules for databases subscribed to a service plan by overriding the generic recurrence setting.

Before you begin

This operation requires the Service Plan Administrator role in AppSync.

You can only override the settings of the recurrence type previously selected for the service plan.

Procedure

1. Navigate to **Service Plans** and select one of the plans from the list.
2. From the **Settings** tab, select the **Plan Startup** phase.

The **Plan Startup Defaults** pane appears on the right.

3. Note the **Recurrence Type** selected for the plan.
A recurrence type can be set only if **Scheduled** was set as the **Startup Type**.
4. Select the **Start service plan phase**.
You will see the Start service plan pane on the right.
5. Note the Recurrence Type selected for the plan.
A recurrence type can be set only if **Automatic** is selected in the **Startup** phase.
6. Click the Plan Startup Overrides tab.
You can view the list of all databases subscribed to the plan.
7. Select one or more databases and click **Override Schedule**.
The Override Schedule dialog appears.
8. Set the schedule based on your requirement and then click **OK**.
For example, if the default recurrence type is for specified days of the month, and the rule setting is to Run at 12:00 AM on the 1st day of every month, you can override the time and the day for individual datastores.

Results

A Pencil icon indicates that default settings have been overridden.

Application discovery

Before creating the database copy, AppSync examines the Oracle database on the host to look for changes such as addition, removal, and shutdown or for changes in open status.

A database is protected only if it is in the ONLINE state. This means the database(s) must be in **open** mode (databases started in **nomount** mode or in **mount** mode and offline databases are not protected).

There are no user settings associated with this phase and it cannot be disabled.

Application mapping

After discovering the application, AppSync maps it to array storage, and protection services such as RecoverPoint.

There are no user settings associated with this phase and it cannot be disabled.

Storage preferences

Sets the preferred order of storage technology to use while creating copies, for example, VNX Snapshot or VMAX-Clone/Snap or RecoverPoint Bookmark.

Use the **Move Up** and **Move Down** buttons. Copies are made using the first technology preference when possible. If conditions are such that the first technology can no longer be used, then any remaining copies will be handled by the next preference instead. For example, if your first preference was a bookmark but not all the application data in the service plan could be mapped to RecoverPoint, then AppSync uses Snap instead.

Note

A single service plan can contain a mix of datasets configured on VNX/VMAX block/file and RecoverPoint. For example, with VNX, if you have a Bronze service plan for Oracle, the databases subscribed can on a mix of RecoverPoint and VNX/VMAX block objects.

A database mix of VNX and VMAX is not supported. Also to get an RP bookmark copy for a database, all LUNs in that database should be configured with RecoverPoint protection; if not Snap copies are created for that database.

Pre-copy script

To perform preparatory steps before creating a copy, specify a pre-copy script and parameters on a service plan's **Settings** tab.

The pre-copy script runs according to the schedule set in the **Plan Startup** phase. AppSync executes this script once per host per service plan run on the production host.

All script phases are non-blocking, which means that even if they fail, service plan execution does not terminate and the next phase continues.

This operation requires the Data Administrator role in AppSync.

For a successful script run ensure:

- The script phase is enabled.
- The script exists in the specified path. You provide absolute path to script; there is no default location.
- You use valid script formats: all executables on UNIX are supported. The script requires execute permissions for the specified user.
- The pre-copy script runs per the schedule set in the Plan Startup phase.
- The script runs as Local System by default for Windows only.
- The script does not put the database/tablespaces in backup mode.
- The script does not shut down the database.

Table 25 Pre-copy script console fields

Field in UI	Description
Full path to script	The complete path to the script location.
Script parameters	Parameters that will be passed to the script during the run.
Run as username	User that has execute permissions on the script.
Password	Password of the user.

Create copy

The Create Copy phase creates a copy that is based on the replication technology that is specified in the service plan.

The Create copy phase specifies the backup type for the Oracle database copy that AppSyncs creates. This phase also sets the period for automatic expiration of the copies.

Automatic expiration of copies

The automatic expiration value in a service plan **Create Copy** phase specifies the maximum desired number of Snap, Clone or Bookmark that can exist simultaneously.

When the "Always keep x copies" value is reached, older copies are expired to free storage for the next copy in the rotation. Failed copies are not counted. AppSync does not expire the oldest copy until its replacement has been successfully created. For example, if the number of copies to keep before expiration is 7; AppSync does not expire the oldest copy until the 8th copy is created. AppSync does not expire copies under the following circumstances:

- Mounted copies are not expired.
- A copy that contains the only replica of a database will not be expired.

This setting is independent of any storage policy setting (for example the VNX pool policy settings in Unisphere for automatic deletion of oldest snapshots.) The service plan administrator should work with the storage administrator to ensure that the Storage policy settings will enable the support of the specified number of snap copies for that application.

Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations.

Note

If this option is not selected, then RecoverPoint copies accumulate, and remain until the bookmarks fall off the RecoverPoint appliance.

Post-copy script

To perform cleanup or other post-copy steps after creating a copy, specify a post-copy script and parameters in the service plan **Settings** tab.

The pre-copy script runs as per the schedule set in the **Plan Startup** phase. You can execute this phase once per host per service plan run. If this script phase is enabled but the permissions to run it are improper, or if the script does not exist in the specified path, the Service Plan run fails with appropriate error.

This process requires the role of AppSync Data Administrator. AppSync executes this script once per host per service plan run on the production host.

All script phases are non-blocking, which means that even if they fail, service plan execution does not terminate and the next phase continues.

For a successful script run ensure:

- The script exists in the specified path. You provide absolute path to script; there is no default location.
- You use valid script formats: all executables on UNIX are supported. The script requires execute permissions for the specified user.

Table 26 Post copy script console fields

Field in UI	Description
Full path to script	The complete path to the script location.
Script parameters	Parameters passed to the script during the run.
Run as username	User that has execute permissions on the script.
Password	Password of the user.

Unmount previous copy

The service plan unmounts a previously mounted copy after creating the new copy.

The exception is a copy that was mounted on-demand instead of mounted by the service plan; in this case the on-demand mounted copy is not unmounted.

All the recovered databases are shut down as part of this phase. There are no user settings associated with this phase and it can be enabled or disabled.

Pre-mount script

You can enable this phase if you want to run a script prior to AppSync performing a mount operation.

This script will be executed once per host per service plan run. If you enable the script phase but the permissions to run it are improper, or if the script does not exist in the specified path, the service plan run fails with appropriate error.

Show caution when using several mount hosts in a Service Plan run. (Refer to [Overriding mount settings on a service plan](#). You must select **Same as mount host** in the **Run on host** option so that the script runs on all mount hosts.

Table 27 Pre-mount script field descriptions

Field in UI	Description
Full path to script	The complete path to the script location
Script parameters	Parameters passed to the script during the run
Run as username	User with execute permissions on the script
Password	Password of the user
Run on host	Host where the script needs to run. Select Same as mount host if several mount hosts are involved.

Mount copies

The Mount copy phase either mounts the copy or mounts and recovers the copy. This phase can be enabled or disabled.

In Mount Copy Defaults settings, you can set values to Mount copy or Mount and recover copy.

For **Mount copy settings**, you can set the mount host value and mount path and the RecoverPoint image access type.

For **Mount and recover copy settings**, you specify the recovery instance, the type of recovery, and the database naming details. Other settings are similar to the Mount copy settings such as mount path and image access type.

For **Mount on standalone server and prepare scripts for Manual Recovery** Oracle mount option, if you enable script phase after the mount operations completes AppSync creates scripts on the mount host that you run to recover the database. The scripts are two types, RMAN and SQL. The scripts are created under `/tmp/<MOUNTED_SID_NAME>/RecoveryScripts`.

Console field descriptions:

- **Host name:** This field is used to specify the host where you want to mount the Oracle copy.
- **Mount to path:** The path on which to mount database files and file systems. For ASM RAC, this setting is unused/ignored.
- **Service Level Objective (SLO):** If you are using a VMAX 3 array, a setting called Desired Service Level Objective (SLO) is available. The option appears in the Mount wizard and it specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.
- **Database name:** This field represents the format of the mounted database name. To specify the original database name use the token `%DB%`. For example: To use the original name that is prefixed by TEST, use `TEST%DB%`.
- **SID name:** This field represents the format of the mounted instance name. To specify the original instance name use the token `%SID%`. For example: To use the original name that is prefixed by TEST, use `TEST%SID%`.
- **ASM Diskgroup:** This field represents the format of the ASM disk group. To specify the original disk group name use the token `%DG%`. For example: To use the original name that is prefixed by TEST, use `TEST%DG%`.
- **Custom initialization parameters:** This field is a multi-line field which allows you to specify settings which override any original database setting on the mounted database copy. This field is useful for editing options such as memory settings.
- **Restart databases after reboot:** This option is used to start the AppSync mounted Oracle databases automatically after a host reboot. By default, this option is disabled.

Overriding mount settings in a service plan

If multiple registered databases are subscribed to the same plan, you can select different mount settings for each database, overriding the generic settings. Recovery settings cannot be overridden.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Service Plans > Oracle** and click one of the service plans from the list.
2. From the **Settings** tab, select **Mount copy** phase.
3. On the right pane, select the **Mount Copy Overrides** tab.
The list of servers includes all Oracle hosts whose databases are subscribed to this plan.
4. Select the server for settings override, and then click **Set Overrides**.
The Override Default Mount Settings dialog displays.
5. Select options for the mount settings that you want to override.
Fields that do not have a selection; they retain their default settings.
6. Click **OK**.
A pencil icon appears in the first column of the row of the server with your changed settings.
7. To revert back to default settings for a server, click **Use Default Settings**.

Post mount script

You can enable this phase if you want to run a script after AppSync performs a mount operation.

This script will be executed once per host per service plan run. If you enable the script phase but the permissions to run it are improper, or if the script does not exist in the specified path, the service plan run fails with appropriate error.

Show caution when using several mount hosts in a Service Plan run. (Refer to the Overriding mount settings on a service plan section. You must select **Same as mount host** in the **Run on host** option so that the script runs on all mount hosts.

Table 28 Post-mount script field descriptions

Field in UI	Description
Full path to script	The complete path to the script location
Script parameters	Parameters passed to the script during the run
Run as username	User with execute permissions on the script
Run on host	Host where the script needs to run. Select Same as mount host if several mount hosts are involved.

Unmount copy

The final phase in the service plan unmounts the copy.

This phase is disabled if the **Unmount previous copy** phase is enabled. There are no user settings associated with this phase.

If you have chosen to mount with recovery options (standalone, RMAN, or cluster mount) in the **Mount copy** phase, all the mounted databases are shut down as part of this phase.

Mount an Oracle copy

This section describes the AppSync console mount fields and their descriptions.

Mount operations

Table 29 Console field descriptions

Field	Description
Mount on Server	The server on which to mount the copy.
Mount path	The Default Mount Path is <code>/appsync</code> . The mount path could also be Same as Original Path . However, this option is not available when the mount host = production host. You can also change Default Mount Path, for example, <code>/EMC</code> instead of <code>/AppSync</code> . Select Mapped Path to specify the path where you want to mount the database.
Image access mode (during RecoverPoint mount)	<ul style="list-style-type: none"> • Logged access: Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host. Virtual access with RP-VMAX, is not supported. • Virtual access with roll: Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. With RP-VMAX, and RP-XtremIO, virtual access with roll is not supported. • Virtual access: provides nearly instant access to the image; it is not intended for heavy processing. With RP-VMAX, and RP-XtremIO, virtual access is not supported.
Restart databases after reboot	Use this option to start the AppSync mounted Oracle databases automatically after a host reboot. By default, this option is disabled.
Desired SLO	For VMAX 3 arrays only, a setting called Desired SLO appears in the Mount wizard and specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.

Table 29 Console field descriptions (continued)

Field	Description
VPLEX Mount option	<ul style="list-style-type: none"> • Native array: Use this option if you want to mount the copy as native array volumes. • VPLEX virtual volume mount: Use this option if you want to mount the copy as VPLEX virtual volumes. • Enable VMware cluster mount: Clear this option if you do not want to perform an ESX cluster mount. By default, this option is enabled.

You can mount a copy created on any multipathing device production host, and mount it on any multipathing device mount host. This means you can create a copy on Block/PowerPath/MPIO devices and mount it on a mount host with any of these combinations.

For DMP, make sure you install DMP on both production and mount hosts.

Note

It is recommended to install Oracle 12c Patch 19404068 for reliable recovery of an Oracle database.

Additional information

- With AppSync 2.2.1 and above, you can configure a temporary location per UNIX host from the AppSync console in the Servers page.
 - AppSync uses the set temporary location during Oracle mount operations for storing information that previously resided in `/tmp/<SID>/`.
 - `/tmp/` is the default temporary location unless you specify otherwise.
- For UNIX hosts, you can configure a command execution timeout value from the Servers page of the AppSync console. AppSync uses this value to wait for each operating system command that is executed by AppSync on a UNIX platform. The default value is 60 minutes. For example, if fsck during file system copy mount takes more than 60 minutes on a host, you can increase the command execution timeout value.
- AIX multiple mounts
 - Multiple copies can be mounted to the same AIX host only if the copies are created using AppSync 3.0.1 and later.
 - If copies were created using AppSync 3.0 or earlier, you cannot mount multiple copies to the same AIX host, even after you upgrade both the sever and agent to AppSync 3.0.1 and later.
 - If you have copies created using both AppSync 3.0 and 3.0.1 and later, it is recommended that you mount the copy created using AppSync 3.0.1 and later for successful concurrent mounts. If you intend to mount the AppSync 3.0 copy, only one copy can be mounted.
 - If you mount the copy created from AppSync 3.0.1 and later, the mount of AppSync 3.0 copy might fail.

- After you upgrade the AppSync server to 3.0.1 and later, ensure that you upgrade the agent to AppSync 3.0.1 and later.

Mounting a copy using the Oracle Mount wizard

From the AppSync console, you can perform a mount of a copy using the Oracle Mount wizard .

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. On the Databases page, select **Recover > Mount a Copy**.
A list of the Oracle database instances appears.
2. From the Copies page, select a copy and click **Mount**.
The Oracle Mount wizard launches.
3. Use the **Copies** or **Service Plan** or **Type** filters to select the copy to mount.
The copies list is refreshed based on the filters selected.
4. Select the copy that you want to mount. (For a RecoverPoint copy, you also have the option to select a bookmark that is based on a specific time, however, ensure that there is a copy available in AppSync.) Select the copy, and click **Mount** to launch the **Mount Copy of Oracle** wizard.
5. Click **Select a point in time** to select a copy with a specific timestamp. The time that is shown here is the console's time. If the console is in a different time zone from the RecoverPoint Appliance (RPA), specify the time in the server's time zone to mount the copy.
6. On the Mount Options page, from the Mount operation list, select one of the following options:
 - Mount on standalone server
 - Mount on standalone server and create RMAN catalog entry
 - Mount on standalone server and recover
 - Mount on standalone server and prepare scripts for manual recovery
 - Mount on grid cluster and recover as RAC database

If you select **Mount on standalone server and recover**, **Mount on standalone server and prepare scripts for manual recovery**, or **Mount on grid cluster and recover as RAC database**, with read/write open mode for recovery, the **Create TempTable Space** option is enabled. This option is used to create the TEMP TableSpace on the recovery-mounted database copy. After you select Create TEMP TableSpace, AppSync shows two other options:

- a. Number of TEMPFILES': Number of files to be added to TEMP TableSpace, each of size specified in 'Size of each file' option
- b. The `size_clause` specifies a number of bytes, kilobytes (K), megabytes (M), gigabytes (G), terabytes (T), petabytes (P), or exabytes (E) . The `size_clause` allows you to establish amounts of disk or memory space, for example 10M. The size of the TempTable Space equals the Temp table file that is multiplied by the size of each file . For example, if the Temp table file count = 2 and the size of each file = 10M, the TempTable Space Size = 20M.

AppSync generates the name of the TempTable Space in the form of `<DBNAME>_TEMP`. This newly created TableSpace is set as the default TEMP

TableSpace of the mounted database instance. During unmount, AppSync drops the created TEMP TableSpace.

Note

- With manual recovery mount, scripts are prepared to both create ('Step-5_createTempTableSpace.sql') and drop ('Step-6_dropTempTableSpace.txt') TEMP TableSpace. You should drop the created TEMP TableSpace manually before unmounting a copy with AppSync.
 - If AppSync fails to drop the TEMP Tablespace during unmount, and if a restore operation is performed using this copy, the tablespace is restored.
 - If you attempt to restore a RecoverPoint copy, the TEMP TableSpace, if created during mount with recovery, is also restored to production. You should drop the TEMP TableSpace manually from the mounted database copy, and then attempt a restore.
-

7. Under **Mount Settings**, do the following:
 - a. From the **Mount on Server** list, select the server on which to mount the copy.
 - b. From the **Mount to Path** list, select a mount path location either to **Default path, Same as original path**, or **Mapped Path**. The mount path is the location where the copy is mounted on the mount host.
 8. Under **Recovery Settings**:
 - a. Review the default settings and make the required changes.
 - b. Select **Restart databases after reboot** to start the AppSync mounted Oracle databases automatically after a host reboot. By default, this option is disabled.
 - c. Select **Create SPFile** to create an SPFile. The SPFile is created in the default location (`$ORACLE_HOME/dbs`), with the name `spfile<SID>.ora`. During unmount, the SPFile is removed from the `$ORACLE_HOME/dbs` folder.
-

Note

The **Restart databases after reboot** and **Create SPFile** options are not available for RMAN and mount with generate scripts.

9. Under **RMAN Settings**:
 - a. Specify the RMAN user name, password, and the connect string.
 - b. Review other default settings and make the required changes.
 - c. Select **Skip Data Files** to skip cataloging of database data files.
-

Note

This is applicable only to the RMAN catalog entry option.

10. Review the default Mount and Recovery settings and make changes, if required.

-- For VPLEX, select a mount option:

- **VPLEX virtual volume** - The snapshot on the back-end array is made visible to VPLEX. AppSync creates VPLEX virtual volumes on the underlying native array snapshots and provisions these virtual volumes to the mount host. The provisioned virtual volumes are added to the storage view of the mount host. During unmount, the virtual volumes are removed from the storage view of the mount host, and it tears down the created VPLEX virtual volume on the underlying native array snapshots. The snapshot on the back-end array is de-provisioned from VPLEX.
- **Native array volume** - AppSync provisions the native array snapshots to the mount host. The mount host must be zoned to the native array where the snapshot is created. All other mount considerations of the native array are applicable.
- **Enable VMware cluster mount** checkbox - If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. Then the target LUN is made visible only to the ESX cluster on which the mount host resides. This is applicable for both RDM and vDisk device types.

-- For VMAX 3 arrays, select the Service Level Objective (SLO) for the mount copy.

-- For VMAX V2 arrays, select the desired FAST VP Policy. Each FAST VP Policy is associated with a storage group on the array. Select the storage group to use for the mount operation by selecting the FAST VP policy associated with that storage group.

11. Click **Next** to display the **Summary** page.
12. Review the mount settings and click **Finish** to complete the mount.
13. In the **Results** page, you can view the progress of the different phases that are part of mounting a copy.

The last phase that is completed displays at the bottom of the list.

RMAN cataloging feature

This section includes prerequisites and restrictions for creating RMAN catalog entry, and copying BCT file.

Mount Operation: Mount on standalone server and create RMAN catalog entry

Table 30 Console field descriptions

Field in UI	Description
RMAN user	Catalog owner
RMAN password	Catalog owner's password
RMAN connect string	The TNS alias used to connect to remote RMAN catalog
TNS_ADMIN	Path of the tnsnames.ora file where the TNS alias is specified. (Default Path : \$ORACLE_HOME\network\admin\)

Table 30 Console field descriptions (continued)

Field in UI	Description
ORACLE_HOME	ORACLE_HOME path for the Oracle binaries. Default: Same as production host
ASM Diskgroup Name	Specify prefix or suffix to rename diskgroups on mount host or %DG% (if production ASM diskgroup name is to be used during mount). Default: APS%DG%
Skip Data Files	Skip cataloging of database data files. Default: Not selected.

Notes on prerequisites

- RMAN catalog database must exist and be accessible on the same network as the mount host.
- The `tnsnames.ora` file on the mount host must contain a TNS alias that points to the RMAN catalog database where AppSync should catalog the copy.
- The catalog and catalog owner must be created prior to mounting a copy to be cataloged.
- Production database must be registered in the RMAN catalog before mounting the copy.
- The Oracle version running the RMAN catalog database must be equal to or greater than the highest Oracle version of all production databases registered to that catalog.
- Copies mounted with RMAN integration cannot be renamed using the database rename option. This also implies that only one copy per database can be mounted on a mount host for RMAN cataloging, and **Mount to Original Host** is not possible.
- Copies mounted with Read-only access cannot be cataloged using RMAN.
- Database must be put in hot backup mode.
- **Create backup controlfile** must be selected in **Create Copy** phase.

Mount on standalone server and prepare scripts for manual recovery

This action overrides mount settings on a service plan. This section includes prerequisites and details for performing a standalone mount of an Oracle copy for use with script-assisted manual recovery steps.

Console field description:

- **Mount to server:** This field is used to specify the host where you want to mount the Oracle copy.
- **Mount to path:** The path on which to mount database files and filesystems. For ASM RAC, this setting is unused/ignored.
- **Database name:** This field represents the format of the mounted database name. To specify the original database name use the token `%DB%`. For example: to use the original name prefixed by TEST, use `TEST%DB%`. The length of the database name is eight characters.
- **SID name:** This field represents the format of the mounted instance name. To specify the original instance name use the token `%SID%`. For example: to use the

original name prefixed by TEST, use `TEST%SID%`. The length of the SID name is 16 characters and the length of the RAC SID is 15 characters.

- **ASM Diskgroup:** This field represents the format of the ASM diskgroup. To specify the original diskgroup name use the token `%DG%`. For example: to use the original name prefixed by TEST, use `TEST%DG%`.
- **Custom initialization parameters:** This field is a multi-line field which allows the you to specify settings which will override any original database setting on the mounted database copy. This is useful for editing options such as memory settings.

After the mount operations complete AppSync will create scripts on the mount host that you must execute to recover the database. The scripts are RMAN scripts and SQL scripts. The scripts are created in `/tmp/<MOUNTED_SID_NAME>/RecoveryScripts`. The script files are named as `Step-
<number>_<operation>.<extension>`. The `<number>` represents the file that must be run first and so on. The `<operation>` signifies what the script does. The `<extension>` specifies the type of script, either RMAN or SQL. Depending on the type of script, either execute it in RMAN or execute through SQLPlus. The generated filenames follow:

```
Step-1_DatabaseRename.sql
Step-1_DatabaseFileRename.sql
Step-2_RecoverDatabase.rman
Step-3_RecoverDatabase.sql
Step-4_OpenDatabase.sql
```

There is only one Step-1 file created depending on whether the recovery operation was performed using the production SID name or an altered SID name. In order to execute the scripts, follow these steps as an Oracle user:

1. Export the Oracle SID as the SID used during recovery.
2. When executing an SQL script, login to SQLPlus using `sqlplus / as sysdba`. You can then run the script: `@/tmp/<MOUNTED_SID_NAME>/RecoveryScripts/Step-
<number>_<operation>.sql`
3. When executing an RMAN script, login to RMAN using `rman target=/. You can then run the script as, @/tmp/<MOUNTED_SID_NAME>/RecoveryScripts/
Step-
<number>_<operation>.rman.`

Note

Make sure you follow the order of these steps during recovery.

Mount on cluster and recover

This section includes prerequisites and details for performing a mount of a copy containing an Oracle RAC database as a RAC database on another cluster or, if renamed, back to the same cluster. The settings for this are as follows:

Console field descriptions:

- **Mount to cluster:** This field is used to specify the cluster where you want to mount the copy. Alternatively, it can be Original cluster to mount back to the production cluster.
- **Mount to servers:** You can select a subset of nodes from the selected cluster, or alternatively, all nodes in the cluster that have been added to AppSync.

Note

AppSync will only mount to cluster nodes which have been registered; unregistered nodes will not be used.

- **Mount to path:** For ASM RAC, ignore this setting
- **Database name:** This field represents the format of the mounted database name. To specify the original database name use the token `%DB%`. For example: to use the original name prefixed by TEST, use `TEST%DB%`. The length of the database name is eight characters.
- **SID name:** This field represents the format of the mounted instance name. To specify the original instance name use the token `%SID%`. For example: to use the original name prefixed by TEST, use `TEST%SID%`. The length of the SID name is 16 characters and the length of the RAC SID is 15 characters.

Note

For RAC mounts, each node in the cluster receives a unique instance name, postfixed by an numeral.

- **ASM Diskgroup:** This field represents the format of the ASM diskgroup. To specify the original diskgroup name use the token `%DG%`. For example: to use the original name prefixed by TEST, use `TEST%DG%`.
- **Custom initialization parameters:** This field is a multi-line field which allows you to specify settings which will override any original database setting on the mounted database copy. This is useful for editing options such as memory settings.

Path mapping

The path mapping option mounts the copy to a host using a path mapping table set to user-defined locations. When you use a path mapping table, you have more control over where data is located.

You must specify the path where you want to mount a specific file system. You must provide a path map where the source file system and the target mount point is specified.

The following is a sample path mapping table for UNIX.

Source file system	Target mount path
/opt/emc/db1	/opt/emc/copies/db1
/opt/emc/db2	/opt/emc/test/db2
/opt/emc/db3	/opt/emc/reserved/db3
/prd/db1	/tst/db1

The following is a sample path mapping table for Windows.

The first two target paths, `G:\` and `H:\` drives must already be available on the mount host. That is, the root drive for the mount path must pre-exist before attempting a mount.

Source file system	Target mount path
D:\Test1	G:\Test1

Source file system	Target mount path
E:\	H:\Test2
F:\Test3	I:\
L:\	N:\

Note

- If a target path is not provided for a source path, then it is mounted to a path same as the source path on the mount host.
- Ensure that you type in the absolute mount path on the target host. If the path is invalid, mount fails.
- Mount copy overrides is unavailable, if you select the mount path as Mapped path.
- For Windows, if one of the entered path is invalid, VSS import fails. Therefore, the entire mount fails. Partial failed scenarios are not supported for Windows mount.
- For Windows and NFS file systems on Unix, nested target mount points are not supported.
- Path mapping is not supported for Oracle ASM disk groups.
- Path mapping is not supported for Oracle NFS RAC database.
- Path Mapping is not applicable to metadata paths for Microsoft Exchange and Microsoft SQL Server.

Specify path mapping settings

You can specify the path where you want to mount a specific copy.

Procedure

1. On the AppSync console, go to the **Copies** page to display available copies.
2. Select a copy to mount, and then click **Mount**.
The Select Copy to Mount page of the Mount wizard launches.
3. Select the copy to mount, and click **Next**.
The Specify Mount Settings page appears.
4. On the **Specify Mount Settings** page:
 - a. Select the mount host.
 - b. From the **Mount Path** list, select **Mapped Path**.
The Path Mapping Settings link appears.
5. Click on the link to open the Path Mapping Settings window.
6. From the **Select Source Host** list, select a host.
All the file systems on the selected host are displayed in the source path column.
7. Specify the target path.
8. Click **Save** to save your settings.

If you want to set the target path for a file system on another source host, repeat steps 6 to 8.

9. Click **Reset**, to clear all the entered target paths for the selected source host.
10. Click **OK** to exit the Path Mapping window.

Note

If you change the path mapping settings, the earlier saved path mapping settings is not valid and the new path mapping settings takes precedence. Therefore, ensure that you save the path mapping settings for all the hosts before changing it.

Restoring an Oracle copy

You can perform a restore of an Oracle copy using the Oracle Restore wizard from the AppSync console.

Before you begin

This operation requires the Data Administrator role in AppSync.

Note

- If a copy is mounted or recovered with database rename, it is not recommended to use this copy for restore.
- Ensure that no virtual machine snapshots are present before protecting a datastore. If virtual machine snapshots are present, protection succeeds, but AppSync fails to perform a file or virtual machine restore.

Procedure

1. On the AppSync console, click the **Copy Management** tab, then select **Oracle** from the drop-down list.

A list of the Oracle database instances appears.

2. Select a database to open the **Copies Page** page for the selected Oracle database which lists available copies with dates of copy.
3. Select the desired copy, and then click **Restore** to launch the **Oracle Restore** wizard.

You may receive the following warning message: You are attempting to perform a restore on a cluster. Please follow the instructions in the AppSync documentation for specific cluster restore procedures.

4. Select the copy to restore, and then click **Next**.

The Restore Options page appears.

5. Click the Restore drop-down list and select one of the following options to restore: **Data**, **Archive logs**, or **Both Data and Archive logs**.

If the database being restored affects any other database. you may receive an affected entity warning message.

6. Click **Next**.

The **Configure Storage Options** screen appears.

- **Wait for mirror rebuild to complete** - This option is applicable for VPLEX Snap copies whose production data resides on local or distributed RAID-1 volumes. By default, this option is selected.
 - **Disable VMWare SRM** - Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later.
 - **Perform device restore in background** - Allows you to optimize restore of VMAX devices. If you select this option, AppSync restore operation does not wait for VMAX track synchronization to complete. The production application is available instantly.
7. Click **Next** to display the **Summary** page.
 8. Review your restore settings and click **Finish** to complete the restore.

On the **Results** page you can view the progress of the different phases that are part of restoring a copy.

Results

Appsync only displays restore warnings for databases discoverable by AppSync that are common to that host. No warnings display for any databases which either are not common to the host or not discoverable.

Refer also to [Restoring a RAC copy](#).

Affected entities during restore

When restoring from a copy, you may be prompted to restore items in addition to the ones you selected.

An affected entity is data that resides on your production host that unintentionally becomes part of a replica because of its proximity to the data you intend to protect. You can prevent affected entity situations by properly planning your data layout based on replica granularity. The granularity of a replica depends upon the environment.

For Oracle, an affected entity can only be another Oracle Database data file(s) or archive logs. You can choose to restore using one of three options. This will determine the level to which affected entities are determined.

- Data only
- Archive Logs only
- Data and Archive Logs

Affected entities only display according to the restore option. If you select, **Data**, Appsync looks for affected entities with respect to the Oracle database data filesystems and storage. AppSync does not use Oracle database(s) archive log storage for checking for affected entities.

If you select, **Archive logs**, the reverse is true. Only the Oracle database archive logs filesystems and storage are used for checking affected entities and not the Oracle database(s) data filesystems.

If you select both **Data** and **Archive logs**, then filesystems and storage from both the Oracle database(s) data files and archive logs will be used for checking for affected entities.

If there are *affected entities* in your underlying storage configuration, the Restore Wizard notifies you of these items. The following scenarios produce *affected entities* that require you to acknowledge that additional items will be restored:

- For RecoverPoint and ViPR Controller, if the databases are in the same consistency group they become *affected entities* when the other database is protected.
- For VNX, VNXe, Unity, VMAX, XtremIO, and ViPR Controller, if the databases are on the same LUN they become *affected entities* when the other database is protected. For VNXe, if the databases are in the same LUN group they become affected entities when the other database is protected.
- For Unity, if the databases are in the same consistency group they become affected entities when the other database is protected.
- For vDISK/datastore - If data files of two data bases: DB1 and DB2 reside on datastore [DS1] and or similarly archive logs of same two databases resides on datastore [DS2], then both become affected entities.

If the affected entity was protected along with the Oracle database selected for restore, AppSync restores it. Any other Oracle database that was not protected but is an affected entity is overwritten.

AppSync calculates affected entities for the consistency groups or LUNs of the Oracle database that is selected for restore. If the affected databases partially reside on other consistency groups, LUN groups, or LUNs, AppSync does not calculate affected entities on those consistency groups, LUN groups, or LUNs.

Affected entities are calculated on the basis of restore granularity. If both data and log are selected for restore, then affected entities are calculated for all the consistency groups, LUN groups, LUNs, or datastores on which the database resides. If only data or only log restore is selected, then the affected entities are only calculated for the selected component's consistency group, LUN group, LUN, or datastore.

If the database's data and log components reside on the same consistency group or LUN, the option to restore only logs or restore only data is not available. You have the option only to restore data and logs. The only exception to this scenario is when you choose to do a differential copy restore.

Vdisk restore with affected entities

Review this information for a Vdisk restore with affected entities.

- During restore, if there are affected databases on virtual disks that are not protected by AppSync, shutdown the these databases including all unmounted filesystems. Additionally, remove Vdisks from VM before proceeding with LUN level restore.
- If affected databases reside on any volume or disk groups, then deport or dismount VGs and DGs before restore and then manually import and mount them post-restore. (Since Appsync does not control these entities, a post storage LUN restore can fail when attempting import/mount of affected VGs and DGs on the production host.)
- Affected entity databases on Vdisks with VG or ASM are not supported.

Restoring a RAC copy

Follow this procedure to restore a RAC copy.

Before you begin

On remote nodes follow these steps:

1. Shutdown all impacted databases as oracle user: `oracle> srvctl stop instance -d <RACDB> -i <DbInstanceOnRemoteNodes>`

2. Dismount all impacted ASM disk groups as grid user: `grid> asmcmd umount <DG>`

On the restore node, perform the restore. Follow these steps:

Procedure

1. On the AppSync console, go to **Copy Management > Oracle**.
2. Select the desired database, and then from the drop-down menu in the lower center of the screen click **Recover** and select **Restore**.
The Oracle Restore wizard launches.
3. On the Select Copy page, select the copy you want to restore, and then click **Next**.
4. Under **Restore Options**, select the option you want, and then click **Next**.
The Summary page opens.
5. Review your restore actions and click **Finish**.
6. Verify that the selected database is being shut down.
7. Verify the message that disk groups and devices are being unmounted.
8. Verify that the restore was successful.

Results

After the restore:

1. Remount the diskgroups on the remote nodes as grid user: `grid> asmcmd mount <DG>`.
2. On any node, perform recovery of the restored database using redo or archive with resetlogs:
`Oracle > startup mount`
`Oracle > recover database`
3. Open the database on the recovery node:
`Oracle > alter database open`
4. Bring up the instances on the additional nodes:
`srvctl start instance -d <RACDB> -i <DbInstanceOnRemoteNodes>`

Restoring a RAC copy for affected entities

Follow these steps to create your restore.

Before you begin

On remote nodes follow these steps:

1. Shutdown all impacted databases as oracle user: `oracle> srvctl stop instance -d <RACDB> -i <DbInstanceOnRemoteNodes>`
2. Shutdown other affected databases: `oracle> srvctl stop database -d <RACDB2>`
3. Dismount all impacted ASM disk groups as grid user: `grid> asmcmd umount <DG>`

On the restore node, perform the restore. Follow these steps:

Procedure

1. On the AppSync console, go to **Copy Management > Oracle**.

2. Select the desired database, and then from the drop-down menu in the lower center of the screen click **Recover** and select **Restore**.

The Oracle Restore wizard launches.

Verify the Warning:

You are attempting to perform a restore on a cluster. Please follow the instructions in the AppSync documentation for specific cluster restore procedures.

3. On the Select Copy page, copy should already be selected. Click **Next**.
4. Under **Restore Options**, select **Data**, and then click **Next**.

Verify that the Affected Entities warning has been displayed that there is a database that is impacted and that you need to shutdown the database manually.

Note

A database may have associated entities on the same storage but in a different node. If this is the case this warning will not display.

5. Click **Next** at the Affected Entities screen.
6. Review the Summary and then click **Finish**.
7. Verify that the selected database is being shut down.
8. Verify the message that disk groups and devices are being unmounted.
9. Verify that the restore was successful.

Results

After the restore:

1. Remount the diskgroups on the remote nodes as grid user:grid> `asmcmd mount <DG>`.
2. On any node, perform recovery of the restored database using redo or archive with resetlogs:


```
Oracle > startup mount
Oracle > recover database
```
3. Open the database on the recovery node:


```
Oracle > alter database open
```
4. Bring up the instances on the additional nodes:


```
srvctl start instance -d <RACDB> -i <DbInstanceOnRemoteNodes>
```
5. Repeat steps 2 and 3 to recover affected database <RACDB2>
6. Bring up affected database <RACDB2>

CHAPTER 8

Protect File Systems

This chapter includes the following topics:

- [Overview of File System support](#)..... 200
- [File system service plan settings](#)..... 203
- [Mounting a copy with the File System Mount wizard](#).....216
- [Restoring a file system](#)..... 223

Overview of File System support

Use AppSync to create and manage application-consistent copies of file systems.

File system features include:

- Dynamic discovery of file systems during service plan run.
- Protection of file systems with service plan or with copy now option. You can select one or more file systems to protect at one time or click **SELECT ALL** to protect all the file systems on the list of file systems page.
- List copies that you can filter by time of creation, copy status, and service plan.
- Mount on a standalone server

Note

AppSync does not support file systems on Linux operating system devices which are not full block devices (such as /dev/sdc) or primary first partition (such as /dev/sdc1).

Hyper-V support

In Hyper-V environments, AppSync requires the storage for File systems to be on iSCSI direct attached devices, Virtual Fiber Channel (NPIV), or SCSI pass-through devices. SCSI Command Descriptor Block (CDB) filtering must be turned off in the parent partition for SCSI pass-through. It is turned on by default.

For Hyper-V SCSI pass-through, the mount host cannot be a Hyper-V host it has to be a physical host or a virtual machine added with Virtual Fiber Channel adapter or iSCSI direct attached.

Resilient File System (ReFS) support

AppSync can discover, protect, mount, and restore ReFS file systems, and AppSync supported standalone and clustered applications residing on ReFS file systems.

Repurposing is supported for ReFS file systems and for SQL server databases on ReFS. Storage spaces are not supported.

Protect NFS file systems on VNX, VNXe, and Unity storage

Learn how AppSync supports protection of NFS file systems on VNX , VNXe File, and Unity File storage.

AppSync supports protecting NFS file systems on Linux (RHEL, SUSE, and OEL) and AIX. You can use these copies for operational recovery.

In the case of service plans configured for VNX file remote protection, the NFS copy is created as a SnapSure Snapshot on the local and/or remote file system. Copies of NFS data stores can be created from service plans configured for local, remote, and local and remote protection. AppSync can also create copies for file system on an Oracle database for Bronze, Silver, and Gold service plans.

During restore from a VNX NFS copy, AppSync creates a roll back snapshot for every file system that has been restored. The name of each roll back snapshot can be found in the restore details. You can manually delete the roll back snapshot after verifying the contents of the restore. Retaining these snapshots beyond their useful life can fill the VNX snap cache and cause resource issues.

VNXe and Unity file snap only support local (Bronze) copies. AppSync can create copies for file system on an Oracle database for Bronze service plan.

Review the following pre-requisites for Silver and Gold copies:

- Register remote VNX arrays with AppSync.
- Create Remote Replication sessions with corresponding remote arrays for each NFS file system where you want creation of Silver and Gold copies. Ensure array status is OK.

Host file systems page

The Host file systems page shows all the host instances currently registered with AppSync.

You can add a host (Windows and UNIX) from this page and can discover file systems available on the hosts.

For further information on file systems, refer to:

- [File systems page](#)
- [File system copies page](#)
- [Mounting a copy with the File System Mount wizard](#)

Filesystem page

The Filesystem page lists all the available filesystems that are discovered for the selected server instance.

Click on a filesystem name to display copies of the filesystem.

File system information includes:

- Status of service plan run, for example checkmark in a green circle = successful
- Name
- Type, for example, NTFS
- Format, for example MBR
- Service plan, for example Bronze
Some file systems can be subscribed to multiple service plans.
- Storage size in GB
- Send alerts to (if requested)

You can select one or more file systems to protect at one time. Click **SELECT ALL** to protect all the file systems on this page (except a file system C:\ which contains host system information).

You can also select a file system, and click on **Repurpose** to repurpose the file system. Select a file system and click **Repurpose > View Repurpose Copy Schedule** to view the repurpose schedule.

For further information on filesystems, refer to:

- [Filesystem copies page](#)
- [Filesystem hosts page](#)
- [Mounting a copy with the Filesystem Mount wizard](#)
- [Restoring a file system](#)
- [Repurposing overview](#)

Filesystem Copies Page

In this page you can view the list of filesystem copies.

The list of copies can be filtered by time of creation, the status of the copies that are created and service plan.

Select a copy to display events for that copy in the Details panel located on the bottom of the Copies page.

From the Copies page, you can select to mount, restore, expire, refresh or repurpose copies.

For further information on filesystems, refer to:

- [Filesystems page](#)
- [Filesystem hosts page](#)
- [Mounting a copy with the Filesystem Mount wizard](#)

PowerHA (HACMP) cluster integration

AppSync can protect file systems residing in a PowerHA cluster environment.

AppSync protects applications on the node where the resource group is available at the time of a service plan run. AppSync relies on the service label IP (virtual IP) of the resource group to create copies or restore across failovers. All the nodes to which clustered file systems can fail over must be registered with AppSync before registering the resource group's service label IP.

Consider the following when protecting an application that fails over as part of a PowerHA cluster:

- The nodes of the PowerHA cluster must be registered with AppSync before registering the service label IP.
- The service label IP/name must be configured in the resource group for the clustered application.
- If a failover occurs while running a replication or restore process, the operation fails. Node failover must occur before running the service plan, or at the start of a restore.
- Each resource group must have a unique service label IP and that must be registered with AppSync for failover awareness.
- File systems that belong to a particular resource group can be discovered and protected only if the corresponding service label IP is registered with AppSync.
- Callout scripts must be present on both the nodes of a PowerHA cluster.
- Clustered file systems must be discovered and subscribed by using the service label IP/name and not through the physical nodes.

Mount considerations

- Mount to the original path is not supported on any of the nodes of the production PowerHA cluster.
- The AppSync host plug-in must be installed on the mount host.
- Mount is only supported in a standalone mode. Cluster mount is not supported.

Restore in a PowerHA environment

Perform a restore. After a restore, the volume group is not concurrent. You must manually make them concurrent before performing a host or file system rediscovery.

Post restore procedure in a PowerHA environment

Learn how to perform manual steps with a restore in a PowerHA environment after a restore.

After restore, a file system mounts to the production host in non-concurrent mode. Remove the file system from the resource group, make it a concurrent volume group, and then add it back to the resource group.

Perform these steps on an active node:

Procedure

1. Unmount the file system.
2. Type the `varyoffvg` command.
3. Type the `varyonvg` command with `-c` option (to make it concurrent).

Verification:

The `lspv` command must show `vg` as concurrent on both nodes as follows:

```
node 2
hdiskpower8      00c2bf0f1ee76ca  oradata concurrent
hdiskpower9      00c2bf0f1f434e3  oralogs concurrent

node 1
hdiskpower18     00c2bf0f1ee76ca  oradata concurrent
hdiskpower19     00c2bf0f1f434e3  oralogs concurrent
```

4. Add the file system back to the resource group.
5. Verify and synchronize the configuration.

File system service plan settings

Use this table to learn default file system settings for service plan phases including startup, discovery, mapping, pre and post copy scripting, mount/unmount and copy.

Default service plan settings create an application-consistent copy every 24 hours. Only the replication technology that is specified by the Copy type in the Create copy phase varies among plans. The following table summarizes the default settings:

Table 31 Default file system Service Plan Settings

Setting	Enabled/Not enabled	Default settings	Schedule
Plan Startup	Enabled	Automatic schedule	Recurrence type: Creates a copy every 24 hours, with the first run at midnight (00:00). Recovery Point Objective (RPO): An alert is issued if the copy is not created within the specified

Table 31 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
			time (default is 24 hours).
Application discovery	Enabled	None	Determined by Plan Startup phase.
Application mapping	Enabled	None	Starts when Application discovery phase completes.
Pre-copy script	Not enabled	None	Starts when Application mapping phase completes.
Create copy	Enabled	<p>Copy type:</p> <ul style="list-style-type: none"> • Bronze • Silver • Gold <p>Also:</p> <ul style="list-style-type: none"> • Storage Ordered Preference: Snapshot, Clone, and Bookmark. Allows you to order, select, or clear storage preferences. By default, all the options are selected. You cannot clear all the preferences, at least one preference must be selected. • Storage Settings: Include RecoverPoint copies in expiration rotation policy —select this option to include RecoverPoint copies when calculating rotations. If you do not select this option, RecoverPoint copies accumulate and remain until the bookmarks for them "fall off" the RecoverPoint appliance. • Configure storage options <ul style="list-style-type: none"> ▪ VPLEX: Allows you to select the preferred cluster for distributed 	Starts when Pre-copy script phase completes.

Table 31 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<p>volumes and the preferred arrays for RAID-1 volumes.</p> <ul style="list-style-type: none"> ▪ VMAX: Allows you to configure the storage pools to create VMAX V2 copy devices for the service plan. You can select or clear the desired VMAX V2 storage pools by expanding the storage array. If you do not select storage pools, the service plan creates copy devices from storage pools enabled in AppSync for VMAX V2 arrays. By default, all the configured storage pools are selected. This option is only applicable for VMAX V2 arrays. – Wait for Clone Sync: If you select this option, the SnapVX clone copy creation waits for clone synchronization to complete. If you clear this option, the clone copy creation initiates clone synchronization and completes the copy creation without waiting for clone synchronization to complete (clone synchronization runs in the background). By 	

Table 31 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<p>default, this option is disabled. This option is only applicable for VMAX V3 arrays.</p> <ul style="list-style-type: none"> Advanced Settings - Allows you to set a retry count and retry interval to retry the failed VSS operation after an interval configured through the retry interval. This is only applicable for Windows applications. The default retry count value is 3 and the default retry interval value is 0 seconds. 	
Post-copy script	Not enabled	None	Starts when Create copy phase completes.
Unmount previous copy	Not enabled	None	Starts when Post-copy script phase completes.
Mount copy (A pre-mount script phase is available for file system service plans)	Not enabled	<p>Mount Copy</p> <ul style="list-style-type: none"> Mount on Server: Original Host Mount with access: Read/write Mount Path: Default Path or Mapped Path Access mode: Logged access Copy to Mount: Local (Only for Gold Plans) Use Dedicated Storage Group: Selected by default Desired SLO: Select the SLO for the mount copy. This is only applicable for VMAX 3 arrays. Desired FAST VP: Select the FAST VP policy. This is only applicable for VMAX V2 arrays. 	Starts when Unmount previous copy phase completes.

Table 31 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<ul style="list-style-type: none"> <li data-bbox="887 369 1211 457">• VPLEX Mount option: Select a VPLEX mount option. <li data-bbox="887 478 1211 1010">• Enable VMware cluster mount: If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. This option is supported on VPLEX, XtremIO, VMAX 3/VMAX All Flash, and Unity arrays. <li data-bbox="887 1031 1211 1310">• Disable VMware SRM: Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later. <li data-bbox="887 1331 1211 1738">• Run Filesystem Check: During a mount operation, the AppSync agent checks file system data consistency by executing the <code>fsck</code> command. This operation can be time consuming. You can clear this option to skip file system check during a mount operation. By default, file system check is enabled. 	

Table 31 Default file system Service Plan Settings (continued)

Setting	Enabled/Not enabled	Default settings	Schedule
		<hr/> <p>Note</p> <p>In the case of a restore operation, the Run Filesystem Check option is enabled by default. You cannot disable it.</p> <hr/>	
Post-mount script	Not enabled	None	Starts when Mount copy phase completes.
Unmount copy	Not enabled	None	Starts when Post-mount script phase completes.

Subscribing a file system to a service plan

This sections shows you how to subscribe a file system to a service plan. Protection for all file systems that are part of a service plan runs at the scheduled time.

Procedure

1. Browse to **Copy Management > Filesystem**.
2. Click the desired server.
The File Systems Page loads for the selected server.
3. Select the file systems that you want to protect, then click **Subscribe to Service Plan** on the **Protect** drop-down list.
4. Select **Gold**, **Silver**, or **Bronze** service plan.

You can also select **Subscribe to Service Plan and run** for immediate subscription and protection.

Overriding service plan schedules

You can set different schedules for individual applications that are subscribed to a service plan, overriding the generic recurrence setting.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can override only the settings of the recurrence type that is already selected for the service plan.

Procedure

1. Browse to **Service Plans** and select one of the plans from the list.
2. From the **Settings** tab, select the **Plan Startup** phase.
3. In the **Plan Startup Defaults** pane on the right, note the **Recurrence Type** selected for the plan.

A recurrence type can be set only if **Scheduled** is selected as the **Startup Type**.

4. Click the **Plan Startup Overrides** tab.

You can see the list of all applications that are subscribed to the plan.

5. Select one or more applications and click **Override Schedule**.
6. In the **Override Schedule** dialog box, set the schedule that is based on your requirement and click **OK**.

For example, if the default recurrence type is **On specified days of the month**, and the rule setting is to **Run at 12:00 AM on the 1st day of every month**, you can override the time and the day for individual applications.

A Pencil icon indicates that default settings have been overridden.

Service plan schedule

The schedule of a service plan is set in the **Plan Startup** phase.

The **Startup Type** (scheduled or on demand) determines whether the plan is run manually, or configured to run on a schedule. Options for scheduling when a service plan starts are:

- Specify a recovery point objective (RPO)
 - Set an RPO of 30 minutes or 1, 2, 3, 4, 6, 8, 12, or 24 hours
 - Minutes after the hour are set in 5 minute intervals
 - Default RPO is 24 hours
- Run every day at certain times
 - Select up to two different times during the day
 - Minutes after the hour is in 5 minute intervals
 - There is no default selected
- Run at a certain time on selected days of the week
 - One or more days of the week (up to all seven days) can be selected
 - There is no default day of the week selected. Default time of day is 12:00 AM.
- Run at a certain time on selected days of the month
 - Select one or more days of the month (up to all days)
 - Select one time of day. Available times are at 15 minute intervals.
 - Default is the first day of the month

Application discovery

Before creating a file system copy, AppSync examines the file system to look for changes such as addition, deletion, renaming, or movement of file systems. If individual file systems are being protected, AppSync rediscovers information about the selected file system.

There are no user settings associated with this phase and it cannot be disabled.

Application mapping

After discovering the application, AppSync maps it to array storage, and protection services such as RecoverPoint.

There are no user settings associated with this phase and it cannot be disabled.

Pre-copy script phase

To perform preparatory steps before creating a copy, specify a pre-copy script and parameters on a service plan's Settings tab.

This operation requires the Service Plan Administrator role in.

The pre-copy script runs according to the schedule set in the Plan Startup phase. Valid script formats are .bat, .exe, and .ps1 (PowerShell scripts) for Windows and .sh for UNIX.

AppSync does not support running of PowerShell scripts directly. You usually have to wrap them in a .bat file. The other option is to make the default "Open" on ps1 files C:\WINDOWS\system32\WindowsPowerShell\v1.0\powershell.exe. When the PS script runs, you may get an error and you will need to set an appropriate execution policy.

To run PowerShell commands from scripts:

1. Specify the full path name to your PowerShell command file in the .bat file:

```
powershell -command C:\PshellCommands.ps1 <nul
```
2. Set the PowerShell execution policy so you can run your script. For example, the first line in the .bat file should look like the following for an unrestricted policy:

```
powershell -command set-executionpolicy unrestricted <nul
```
3. To ensure correct termination of your PowerShell session, add <nul to the end of the line that calls your PowerShell script.

For Windows you can optionally enter credentials to run the script as a specific user. The script runs as Local System by default. For UNIX the credentials are mandatory.

For Windows the default location of the script is %ProgramData%\EMC\AppSync\scripts\ on the application host.

For UNIX the File field should have the full path to the scripts that are to be executed.

Exact parameters depend on your script. Parameters with spaces must be enclosed in double quotes. This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Create copy phase features freeze and thaw callout scripts

The Create copy phase creates a copy that is based on the replication technology that is specified in the service plan. You can configure and run freeze and thaw callout scripts in this phase

Configuring and running freeze and thaw callout scripts

AppSync provides two scripting opportunities during the execution of Create copy phase, called the freeze and thaw callout scripts. Unlike pre-copy and post-copy script phases which are run before and after the Create copy phase, freeze and thaw scripts are unconfigurable by the GUI. The scripts are placed in a pre-defined location with a pre-defined name. You can use these scripts to quiesce (suspend I/O) and thaw on any AppSync unsupported databases residing on the subscribed file systems for a short period (usually few seconds). During this time the copy is activated. The scripts are run with the user credentials used to register host-plugin with AppSync.

During the Create copy phase, when AppSync executes these scripts, a temporary XML file is provided as the only argument to these callout scripts. This XML file has the list of file systems being protected by the Create copy phase.

AppSync continues with normal copy creation if no callout script is found or if the callout script is not executable. If either of the callout scripts fail (non-zero exit value), copy creation fails and the Create copy phase ends with an error. If the freeze callout script runs successfully, and then copy creation fails due to any storage issue, the thaw callout script is run before ending the create copy phase with an error.

For a Windows host-plugin, place the callout executable scripts in the %ProgramData%\EMC\AppSync\scripts folder and name it

appsync_freeze_filesystem_<service plan name in lower case>.bat for the freeze callout. Name the thaw callout script:

appsync_thaw_filesystem_<service plan name in lower case>.bat

For example, AppSync runs this script as follows:

```
C:\ProgramData\EMC\AppSync\scripts
\appsync_freeze_filesystem_bronze.bat, C:\Windows\TEMP
\d575f2e6-7dc4-4389-87c9-491effc57318.xml
```

Where C:\Windows\TEMP\d575f2e6-7dc4-4389-87c9-491effc57318.xml file content is in the following form:

```
<Application type='Filesystem'><sourceVolumePath>F:\</
sourceVolumePath><sourceVolumePath>G:\</sourceVolumePath></
Application>
```

For a UNIX host-plugin, place the callout executable scripts in the /var/opt/emc/appsnc/scripts folder name as appsync_freeze_filesystem_<service plan name in lower case> for the freeze callout and appsync_thaw_filesystem_<service plan name in lower case> for the thaw callout. Do not use a file name extension such as .pl or .sh. The scripts should be executable.

Unix scripts must have #!/bin/sh, #!/usr/bin/python, or any other type as the first line based on the script type being used.

AppSync runs this script as follows:

```
/var/opt/emc/appsnc/scripts/
appsync_freeze_filesystem_bronze /tmp/904f510f-47ce-402f-a27a-
b3a48840a279ybo61k.xml
```

Where /tmp/904f510f-47ce-402f-a27a-b3a48840a279ybo61k.xml file content is in the following form:

```
<Application
type='Filesystem'><filesystems><filesystem><name>/FS1</name></
filesystem><filesystem><name>/FS2</name></filesystem></
filesystems></Application>
```

Repurposing example

- **Windows - Place the callout executable scripts in the %ProgramData%\EMC\AppSync\scripts folder and name the scripts:**
 - **Freeze callout script** - appsync_freeze_filesystem_<first generation copy label name in lower case>.bat
 - **Thaw callout script** - appsync_thaw_filesystem_<first generation copy label name in lower case>.bat

- **Unmount callout script** - `appsync_unmount_filesystem_<first generation copy label name in lower case>.bat`
- **Unix - Place the callout executable scripts in the `/var/opt/emc/appsync/scripts` folder and name the scripts:**
 - **Freeze callout script** - `appsync_freeze_filesystem_<first generation copy label name in lower case>`
 - **Thaw callout script** - `appsync_thaw_filesystem_<first generation copy label name in lower case>`
 - **Unmount callout script** - `appsync_unmount_filesystem_<first generation copy label name in lower case>`

Do not use a file name extension such as `.pl` or `.sh`. The scripts should be executable.

Unix scripts must have `#!/bin/sh`, `#!/usr/bin/python`, or any other type as the first line based on the script type being used.

Configure retry on VSS failure

You can configure a VSS retry count in the create copy phase of a service plan. During protection, if a service plan fails because of VSS failures such as VSS timeout issue, the service plan runs the VSS freeze/thaw operation again based on the specified retry count and interval. This option is supported only on Windows applications - File system, Microsoft SQL, and Microsoft Exchange. This option is not used while creating Crash consistent copies of SQL databases.

Note

AppSync does not perform a VSS retry, if the application freeze itself fails. If the application is not in a state to create a copy, AppSync fails to quiesce it, and does not retry the VSS freeze/thaw operation. The application must be brought back to a state where it can be quiesced and then the service plan must be re-run.

Post-copy script phase

To perform cleanup or other post-copy steps after creating a copy, specify a post-copy script and parameters in a service plan's **Settings** tab.

This operation requires the Service Plan Administrator role in AppSync.

The script runs on successful completion of the **Create copy** phase. Valid script formats are `.bat`, `.exe`, and `.ps1` (PowerShell scripts) for Windows and `.sh` for UNIX.

AppSync does not support running of PowerShell scripts directly. You usually must wrap them in a `.bat` file. The other option is to make the default "Open" on `ps1` files `C:\WINDOWS\system32\WindowsPowerShell\v1.0\powershell.exe`. When the PS script runs, you may get an error and you must set an appropriate execution policy.

To run PowerShell commands from scripts:

1. Specify the full pathname to the PowerShell command file in the `.bat` file:
`powershell -command C:\PshellCommands.ps1 <nul`
2. Set the PowerShell execution policy so you can run the script. For example, the first line in the `.bat` file should look like the following for an unrestricted policy:
`powershell -command set-executionpolicy unrestricted <nul`
3. To ensure correct termination of the PowerShell session, add `<nul` to the end of the line that calls your PowerShell script.

For Windows, you can optionally enter credentials to run the script as a specific user. The script runs as Local System by default. For UNIX, the credentials are mandatory.

For Windows, the default location of the script is %ProgramData%\EMC\AppSync\scripts\ on the application host.

For UNIX, the File field should have the full path to the scripts to be run.

Exact parameters depend on the script. Parameters with spaces must be enclosed in double quotes. This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Unmount previous copy

The service plan unmounts a previously mounted copy after creating the new copy. The exception is a copy that was mounted on-demand as opposed to by the service plan; in this case the on-demand mounted copy is not unmounted.

There are no user settings associated with this phase and it can be enabled or disabled.

Unmount callout script

Appsync provides scripting opportunity during the Unmount previous copy and Unmount copy phases. The unmount callout scripts are placed in a pre-defined location with a pre-defined name. You can use these scripts to shutdown the unsupported database before the actual unmounting of the file system starts. The scripts are run with the user credentials used to register the host-plugin with AppSync. During the unmount phase, when AppSync executes these scripts, a temporary XML file is provided as the only argument to these callout scripts. This XML file has the list of mount points that are being unmounted by the Unmount copy phase.

For a UNIX host-plugin, place the callout executable scripts in the /var/opt/emc/appsync/scripts folder and name it

appsync_unmount_filesystem_<serviceplan name in lower case>.

Do not use file name extensions such as .pl or .sh. The scripts must be executable.

Unix scripts must have #!/bin/sh, #!/usr/bin/python, or any other type as the first line based on the script type being used.

For example, AppSync runs the following script:

```
/var/opt/emc/appsync/scripts/appsync_unmount_filesystem_bronze /tmp/904f510f-47ce-402f-a27ab3a48840a279ybo61k.xml
```

The content of the /tmp/904f510f-47ce-402f-a27a-b3a48840a279ybo61k.xml file is in the following format:

```
<Application type='Filesystem'>
<filesystems>
<filesystem><mountPoint>/appsync-mounts/ppvgl</mountPoint></
filesystem>
</filesystems>
</Application>
```

For a Windows host-plugin, place the callout executable scripts in the %ProgramData%\EMC\AppSync\scripts folder and name it appsync_unmount_filesystem_<service plan name in lower case>.bat for the unmount callout.

For example, AppSync runs the following script:

```
C:\ProgramData\EMC\AppSync\scripts
\appsync_unmount_filesystem_bronze.bat, C:\Windows\TEMP
\d575f2e6-7dc4-4389-87c9-491effc57318.xml
```

The content of the C:\Windows\TEMP\d575f2e6-7dc4-4389-87c9-491effc57318.xml file is in the following format:

```
<Application type="Filesystem">
<mountpath>C:\AppSyncMounts\HDrive\</mountpath>
<mountpath>C:\AppSyncMounts\GDrive\</mountpath>
</Application>
```

Mount copy

The Mount copy phase mounts the copy. This phase can be enabled or disabled.

Field	Description
Mount on Server	The server on which to mount the copy. Only the nodes of the cluster or standalone hosts are available for selection.
Mount with access	Type of access the copy should be mounted with.
Mount path	The Default Mount Path is %SystemDrive%\AppSyncMounts\%%ProdServerName%. To specify the value of a Windows environment variable in the mount path, delimit the variable name with single percent signs (%). The default path also contains an AppSync variable (ProdServerName) which is delimited with two percent signs (%%). The following characters are not valid in the path: < > " / ? * . The mount path could also be Same as Original Path . However, this option is not available when the mount host is the same as production host. Select Mapped Path to specify the path where you want to mount the file system.
Desired SLO (VMAX 3 only)	Select the desired SLO for the target LUN. If there is a storage group for the mount host with the desired SLO, the LUN will be added to the storage group. If it does not exist, AppSync will add it to any storage group that is masked to the host.
Image access mode (during RecoverPoint mount)	<ul style="list-style-type: none"> • Logged Access: Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host. • Virtual Access with Roll: Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. With RP-VMAX and RP-XtremIO, virtual access with roll is not supported. • Virtual Access: Provides nearly instant access to the image; it is not intended for heavy processing. Virtual access with RPVMAX and RP XtremIO is not supported.

Field	Description
Use Dedicated Storage Group	<ul style="list-style-type: none"> Applicable only for physical hosts or virtual machines with direct iSCSI part of cluster. Checked by default, enabling this option allows AppSync to enforce a dedicated VMAX or VNX storage group, or an XtremIO initiator group for a mount. (A dedicated VMAX or VNX storage group, or an XtremIO initiator group contains the selected mount host only). The mount will fail if you are mounting to a node of a cluster that is in a storage group shared with the other nodes. <hr/> <p>Note</p> <p>Use this option to mount the copy to a node for copy validation or backup to tape. In this scenario, you will need two storage groups. One storage group is dedicated to the passive node being used as a mount host and the other storage group is for the remainder of the nodes in the cluster. Both storage groups contain the shared storage for the cluster.</p> <hr/> <ul style="list-style-type: none"> If unchecked, AppSync does not enforce the use of a dedicated storage group for a mount.

Override mount settings in a service plan

If there are multiple file systems subscribed from different hosts to the same plan, you can select different mount settings for each file system, overriding the generic mount settings.

Before you begin

This operation requires the Data Administrator role in AppSync.

Note

Mount overrides are not supported for multiple file systems on the same host.

Procedure

1. Browse to **Service Plans > FileSystems** and click one of the plans from the list.
2. From the **Settings** tab, select the **Mount copy** phase.
3. On the right pane, select the **Mount Copy Overrides** tab.

The list of file systems includes all file systems subscribed to this plan. The mount settings display the default settings.

4. Select the file system whose settings you want to override and click **Set Overrides**.

Press and hold the Shift or Ctrl keys to select multiple file systems.

5. On the **Mount Copy Overrides** dialog box, select options only for those mount settings that you want to override.

Fields that do not have a selection retain their default settings.

6. Click **OK**.

A pencil icon appears in the first column of the file system's row whose default mount settings you changed.

7. To revert to default settings, select the file systems and click **Use Default Settings**.

Post-mount script

Specify a post-mount script and parameters from the Post-mount script option in the **Settings** tab of a service plan.

The script runs on successful completion of the mount copy or mount with recovery phase. This script is typically used for backup.

From the **Server** list, select the server on which to run the script. You can optionally run it on a registered host other than the mount host, and enter credentials to run the script as a specific user.

The default location of the script is `%ProgramData%\EMC\AppSync\scripts\` on the application host.

Exact parameters depend on your script. Parameters with spaces must be enclosed in double quotes.

This phase can be enabled or disabled. This operation requires the Service Plan Administrator role in AppSync.

Mounting a copy with the File System Mount wizard

Use the File System Mount wizard to create any point-in-time mount for RecoverPoint copies.

Before you begin

This task requires the Data Administrator role in AppSync.

For UNIX, the File field should have a full path to the scripts to be run. You need user credentials to run the script.

Follow these steps:

Procedure

1. On the AppSync console, go to **File systems > Copies** to display available copies.
2. Select a copy to mount, and then click **Mount**.
The Select Copy to Mount page of the File System Mount wizard launches. You can select any copy to be mounted. By default the copy that you selected before selecting Mount is highlighted, however other copy instances also appear.
3. Select the copy to mount, and click **Next**.
The Specify Mount Settings page appears.
4. On the Mount Settings page:
 - a. From the **Mount on Server** list, select the server on which to mount the copy.
 - b. From the **Mount with access** list, select the type of access the copy must be mounted with: **Read-only** or **Read-write**.
 - c. From the **Mount Location** list, select a mount path location either **To original path**, **Mount to alternate path**, or **Mapped Path**. The mount path is

the location where the copy is mounted on the mount host. By default AppSync displays the path of the mount host you selected. You can also edit and mount the copy to a user-defined location.

- d. In case the selected copy is a RecoverPoint bookmark, from the **Image access mode** list, select one of the following options:
- **Logged access:** Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. Logged access is the only option available when you mount to the production host.
 - **Virtual access with roll:** Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. With RP VMAX, and RP XtremIO, virtual access with roll is not supported.
 - **Virtual access:** Provides nearly instant access to the image. Virtual access is not intended for heavy processing. Virtual access with RP VMAX and RP XtremIO is not supported.
- e. For VMAX 3 arrays, from the **Desired SLO** list, select the desired Service Level Objective (SLO) for the mount copy.

Note

The SLO values are dynamically fetched from the VMAX 3 arrays, and only the unique values are displayed.

- f. For VMAX V2 arrays, select the desired FAST VP policy for the mount copy.
- g. Clear the **Use Dedicated Storage Group** option, if you do not want AppSync to enforce the use of a dedicated storage group for a mount. By default, this option is enabled.
- h. From the **VPLEX Mount option**, select one of the following:
- **Native array:** Use this option if you want to mount the copy as native array volumes.
 - **VPLEX virtual volume mount:** Use this option if you want to mount the copy as VPLEX virtual volumes.
 - **Enable VMware cluster mount:** Clear this option if you do not want to perform an ESX cluster mount. By default, this option is enabled.
5. From the **Summary** page, review the copy and mount settings that you chose in the previous pages and click **Finish** to mount the copy.

The Results page launches with mount result information.

Note

- For UNIX hosts, you can configure a command execution timeout value from the Servers page of the AppSync console. AppSync uses this value to wait for each operating system command that is executed by AppSync on a UNIX platform. The default value is 60 minutes. For example, if fsck during file system copy mount takes more than 60 minutes on a host, you can increase the command execution timeout value.
 - AIX multiple mounts
 - Multiple copies can be mounted to the same AIX host only if the copies are created using AppSync 3.0.1 and later.
 - If copies were created using AppSync 3.0 or earlier, you cannot mount multiple copies to the same AIX host, even after you upgrade both the sever and agent to AppSync 3.0.1 and later.
 - If you have copies created using both AppSync 3.0 and 3.0.1 and later, it is recommended that you mount the copy created using AppSync 3.0.1 and later for successful concurrent mounts. If you intend to mount the AppSync 3.0 copy, only one copy can be mounted.
 - If you mount the copy created from AppSync 3.0.1 and later, the mount of AppSync 3.0 copy might fail.
 - After you upgrade the AppSync server to 3.0.1 and later, ensure that you upgrade the agent to AppSync 3.0.1 and later.
-

Changing the mount point for an affected file system

Follow this procedure to manually change the mount point for an affected file system.

Assume VG1 is the source volume group.

Procedure

1. Get the list of LVs using the `lsvg -l VG1` command, and check which file systems show mount point on `/tmp/EMCAppsync **` directory.
2. Run `chfs -m <Original MT Pt> /tmp/EMCAppsync6922/vg1_logs` command where `<Original Mt Pt>` is the mount point where the file system was originally mounted.
3. Run `fsck` on the source Logical Volume `fsck -y /dev/fslv01`.
4. Run `mount` command using the log logical volume and make sure that the source has been mounted successfully `mount -v jfs2 -o rw,log=/dev/loglv00 /dev/fslv01 <Orig Mt Pt>`

Unmounting a file system copy

When you select a copy to unmount, other copies that were mounted along with the selected copy will also be unmounted.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can unmount a copy only from a list of copies made for a file system.

Procedure

1. Navigate to the Copies page from the Copy Management or Service Plan pages:
 - **Copy Management > FileSystems** > select the server which hosts the file system you want to unmount, then select the file system with the copy to unmount.
 - **Service Plans > File systems**> select a service plan, then select the **Copies** tab.
2. From the list of copies, select the copy and click **Unmount** from the button in the lower part of the page.

The **Unmount Confirmation** dialog displays all the copies of other file systems that were mounted along with the selected copy to be unmounted.

3. Click **Yes** to confirm the unmount of all the copies shown in the dialog.

The **Unmount** page displays the progress of the unmount operation. All copies associated with the selected copy will be unmounted.

Override mount settings in a service plan

If there are multiple file systems subscribed from different hosts to the same plan, you can select different mount settings for each file system, overriding the generic mount settings.

Before you begin

This operation requires the Data Administrator role in AppSync.

Note

Mount overrides are not supported for multiple file systems on the same host.

Procedure

1. Browse to **Service Plans > FileSystems** and click one of the plans from the list.
2. From the **Settings** tab, select the **Mount copy** phase.
3. On the right pane, select the **Mount Copy Overrides** tab.

The list of file systems includes all file systems subscribed to this plan. The mount settings display the default settings.

4. Select the file system whose settings you want to override and click **Set Overrides**.

Press and hold the Shift or Ctrl keys to select multiple file systems.

5. On the **Mount Copy Overrides** dialog box, select options only for those mount settings that you want to override.

Fields that do not have a selection retain their default settings.

6. Click **OK**.

A pencil icon appears in the first column of the file system's row whose default mount settings you changed.

7. To revert to default settings, select the file systems and click **Use Default Settings**.

Mounting a file system after reboot

When an AppSync protected file system copy is mounted on to the mount host, AppSync automatically modifies the mount point related entries in the `/etc/fstab` file for Linux hosts. The AppSync agent modified entries in the `/etc/fstab` file is followed by a comment - # line added by AppSync Agent. The file system remains mounted even if the mount host reboots after a file system mount operation.

The following is an example of the `/etc/fstab` entry on a Linux host:

```
/dev/aps_1datavg_51fd06a41290/datalv /appsync-mounts/
oracle_mounts/d1 ext4 rw,nofail,_netdev 0 0 # line
added by AppSync Agent
```

On an AIX host, when a file system copy is unmounted, the entries are removed from the `/etc/filesystem` file.

The following is an example of the `/etc/filesystem` entry on an AIX host:

```
/appsync-mounts/symm1:
dev = /dev/APMlv02
vfs = jfs
log = /dev/APMloglv27
mount = true
check = false
options = nodev,rw
account = false
```

Note

To automatically mount a file system after a production host reboot, ensure that the value of `mount` is set to `TRUE` for the production file system. For file systems created using `crfs` or `smitty`, the automatic mount option is turned to `TRUE` by default. You can check `/etc/filesystems` and verify if the `mount` flag is set to `TRUE`, which allows a reboot.

Nested mount support for File systems

If a file system is mounted at a mount point `/a` and another file system (from the same volume group or a different volume group) is mounted at a mount point below `/a` such as `/a/b`, then `/a` and `/a/b` are considered as nested file systems because its mount point is a sub-directory of the parent file system.

You can mount and unmount nested file systems even if they are spread across volume groups.

It is recommended that you avoid creating circular nested file systems across volume groups such as the following because `mount`, `umount`, and `restore` operations fail even if protection succeeds.

- `/a` in volume group 1
- `/a/b` in volume group 2
- `/a/b/c` in volume group1

Instead, do one of the following:

- Create the file system layout in the following manner:
 - /a in volume group 1
 - /a/b in volume group 2
 - /a/b/c in volume group 3
- Create all nested file systems in one volume group

Note

- Supported only on UNIX platforms.
 - NFS file systems are not supported in nested layouts.
 - All the file systems from a single volume group must be mounted and unmounted together. Otherwise, the second mount of a file system fails, if another file system is already mounted from that volume group. Even in the case of non-nested file systems in a volume group such as /a and /b, if /a is mounted, then the second mount of /b fails.
-

Path mapping

The path mapping option mounts the copy to a host using a path mapping table set to user-defined locations. When you use a path mapping table, you have more control over where data is located.

You must specify the path where you want to mount a specific file system. You must provide a path map where the source file system and the target mount point is specified.

The following is a sample path mapping table for UNIX.

Source file system	Target mount path
/opt/emc/db1	/opt/emc/copies/db1
/opt/emc/db2	/opt/emc/test/db2
/opt/emc/db3	/opt/emc/reserved/db3
/prd/db1	/tst/db1

The following is a sample path mapping table for Windows.

The first two target paths, G:\ and H:\ drives must already be available on the mount host. That is, the root drive for the mount path must pre-exist before attempting a mount.

Source file system	Target mount path
D:\Test1	G:\Test1
E:\	H:\Test2
F:\Test3	I:\
L:\	N:\

Note

- If a target path is not provided for a source path, then it is mounted to a path same as the source path on the mount host.
 - Ensure that you type in the absolute mount path on the target host. If the path is invalid, mount fails.
 - Mount copy overrides is unavailable, if you select the mount path as Mapped path.
 - For Windows, if one of the entered path is invalid, VSS import fails. Therefore, the entire mount fails. Partial failed scenarios are not supported for Windows mount.
 - For Windows and NFS file systems on Unix, nested target mount points are not supported.
 - Path mapping is not supported for Oracle ASM disk groups.
 - Path mapping is not supported for Oracle NFS RAC database.
 - Path Mapping is not applicable to metadata paths for Microsoft Exchange and Microsoft SQL Server.
-

Specify path mapping settings

You can specify the path where you want to mount a specific copy.

Procedure

1. On the AppSync console, go to the **Copies** page to display available copies.
2. Select a copy to mount, and then click **Mount**.
The Select Copy to Mount page of the Mount wizard launches.
3. Select the copy to mount, and click **Next**.
The Specify Mount Settings page appears.
4. On the **Specify Mount Settings** page:
 - a. Select the mount host.
 - b. From the **Mount Path** list, select **Mapped Path**.
The Path Mapping Settings link appears.
5. Click on the link to open the Path Mapping Settings window.
6. From the **Select Source Host** list, select a host.
All the file systems on the selected host are displayed in the source path column.
7. Specify the target path.
8. Click **Save** to save your settings.
If you want to set the target path for a file system on another source host, repeat steps 6 to 8.
9. Click **Reset**, to clear all the entered target paths for the selected source host.
10. Click **OK** to exit the Path Mapping window.

Note

If you change the path mapping settings, the earlier saved path mapping settings is not valid and the new path mapping settings takes precedence. Therefore, ensure that you save the path mapping settings for all the hosts before changing it.

Restoring a file system

Use the File System Restore wizard to restore an existing file system copy.

Procedure

1. On the AppSync console go to **Copy Management > File systems**.

The file system Hosts page launches displaying available servers.

2. Select the server which has the file systems for restore, to launch the file systems page.
3. Select the file system that you want to restore (**Select All** option appears in lower left of page), and then click **Restore** from the Recover menu in the lower left area of the file system page.

This action launches the File System Restore wizard. All protected copies are listed in the wizard. You can only select one file system copy at a time for restore. The copy date, file system name, server name, service plan, and copy type appear for each copy.

4. Select the file system copy for restore and click **Next**.

The Restore Warnings page launches. The Restore Warnings page lists the file system which belongs to the same consistency group or volume group where the copy is restored. If this file system is protected as well as the selected file system copy, this file system is overwritten.

The restore warning page also lists any application that is installed in the file system that is being restored.

5. Read the warnings and click the **I have read warnings** checkbox.
6. Click **Next**.

The **Configure Storage Options** screen appears.

- **Wait for mirror rebuild to complete** - This option is applicable for VPLEX Snap copies whose production data resides on local or distributed RAID-1 volumes. By default, this option is selected.
- **Disable VMWare SRM** - Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later.
- **Perform device restore in background** - Allows you to optimize restore of VMAX devices. If you select this option, AppSync restore operation does not wait for VMAX track synchronization to complete. The production application is available instantly.

7. If the Summary page looks correct, click **Finish**.

The Results page loads where you can view the restore results.

CHAPTER 9

Protect VMware Datacenters

This chapter includes the following topics:

- [Configuration prerequisites](#) 226
- [Discovering datacenters](#) 228
- [Considerations when mounting a VMFS copy](#) 237
- [Restoring a datastore from a copy](#) 239
- [Restoring a virtual machine from a copy](#) 242
- [File or folder restore with VMFS or NFS datastores](#) 245

Configuration prerequisites

AppSync can create, mount, and restore copies in VMware vStorage VMFS and NFS data store configurations. Configuration prerequisites are required to integrate AppSync with VMware vStorage VMFS protection. Configure RecoverPoint and VMware according to the product documentation.

VMware configuration prerequisites

- VMware vCenter Server must be used in the environment.
- AppSync supports VMware's use of VSS with VM snapshots when a supported version of vSphere is installed and the VMware Tools facility is present on the virtual machine on the VMFS you are replicating. Refer to VMware documentation for information on the VSS-related characteristics in an AppSync copy. Contact VMware regarding considerations that are related to VSS in this configuration.
- When there is a configuration change in the vCenter Server, perform a discovery of data centers in the vCenter Server from the AppSync console before you protect a data store. Ensure that the VMFS UUID is unique in the virtual center inventory across all data centers.

Note

You can also create and manage copies of VMware data stores with the VSI for VMware vSphere web Client. The plug-in is available as a separate download from the AppSync Online Support page at support.EMC.com.

RecoverPoint configuration prerequisites

- Configure RecoverPoint protection (Local/Remote/Local and Remote) for the production LUNs before deploying AppSync. Refer to RecoverPoint documentation to create consistency groups and define replication sets.
- In an ESX cluster, target LUNs should be made visible to all the ESX hosts in the cluster.
- The AppSync server must connect to the RPA through the network.

VMware vStorage VMFS requirements

Some considerations apply when AppSync is introduced into a VMware environment for protecting VMware data stores.

All VMware specific operations occur through the VMware vCenter Server.

AppSync can be configured to require vCenter Server login credentials to allow protection of a certain VMFS for security purposes. Unless you instruct AppSync to omit this feature, AppSync takes a VMware Snapshot for each virtual machine that is online and residing on the VMFS before protection. This action ensures operating system consistency for the resulting replica. The following user roles for a Virtual Center ESX cluster are allowed with AppSync:

- Administrator
- VM power user
- VM user
- Resource pool Administrator
- VMware consolidated backup user

- Data store consumer
- Network Administrator

The following privileges must be assigned to the VC role that you plan to use in AppSync:

Datastore

Folder > Create Folder

Host > CIM

Host > Configuration > Storage partition configuration

Resource > Assign virtual machine to resource pool

Resource > Migrate powered off virtual machine

Resource > Migrate powered on virtual machine

Sessions > Validate session

Virtual Machine > Configuration > Add existing disk

Virtual Machine > Configuration > Add new disk

Virtual Machine > Configuration > Add or remove device

Virtual Machine > Configuration > Advanced

Virtual Machine > Configuration > Modify device settings

Virtual Machine > Configuration > Raw device

Virtual Machine > Configuration > Reload from path

Virtual Machine > Configuration > Remove device

Virtual Machine > Configuration > Rename

Virtual Machine > Configuration > Settings

Virtual Machine > Guest Operations > Guest Operation Program Execution

Virtual Machine > Interaction > Guest operating system management by VIX API

Virtual Machine > Interaction > Power Off

Virtual Machine > Interaction > Power On

Virtual Machine > Inventory

Virtual Machine > Provisioning > Allow disk access

Virtual Machine > Provisioning > Allow read-only disk access

Virtual Machine > Provisioning > Clone virtual machine

Virtual Machine > Snapshot management

AppSync supports VMware's use of VSS with VM snapshots when a supported version of vSphere is installed and VMware Tools are present on the virtual machine on the VMFS you are protecting. Refer to VMware documentation for use of the VSS-related characteristics in the AppSync copy and contact VMware regarding considerations that are related to VSS in this configuration.

If virtual machines in the data store have RDMs or iSCSI LUNs visible to them, the resulting copy does not contain those LUNs. If the virtual machine has virtual disks other than the boot drive located in other data stores, it is possible to capture these disks by configuring the service plan to include virtual machine disks.

Discovering datacenters

To keep AppSync up to date, discover datacenters on the VMware vCenter Server when there is a change in the configuration of the vCenter Server.

Before you begin

- This operation requires the Data Administrator role in AppSync.
- At least one vCenter server must be added to AppSync.

Procedure

1. Browse to **Copy Management > VMware Datacenters**.
2. On the VMware Datacenters page, click **Discover Datacenters > On Virtual Center** and select a vCenter server to discover its datacenters.

Optionally, you can also add a vCenter server by clicking **Add vCenter Server**. See [Add a vCenter Server](#).

List of datacenters

The top level of the VMware Datacenters page shows all datacenters registered with AppSync.

Column	Description
Protection status of datacenter	<ul style="list-style-type: none"> • Green: Latest copies of all datastores on the datacenter protected successfully • Yellow: One or more of the latest datastore copies on the datacenter completed with errors • Red: One or more of the latest datastore copies on the datacenter failed to complete • "i" symbol: One or more datastores on the datacenter are either not subscribed to service plans or do not have copies associated with them
Name	Name of the datacenter on the vCenter server.
vCenter Server	Name of the vCenter server that hosts the datacenter.
Last Discovery	Time when a discovery was last performed on the vCenter server.
Alert Recipients	List of email aliases to receive email alerts.

Clicking on a datacenter name shows the datastores.

Adding a VMware vCenter Server

Add a VMware vCenter Server to AppSync when a virtual machine is used as a mount host.

Before you begin

- This operation requires the Resource Administrator role in AppSync.
- Ensure that you know the credentials of an account with Administrator privileges on the vCenter Server.

Procedure

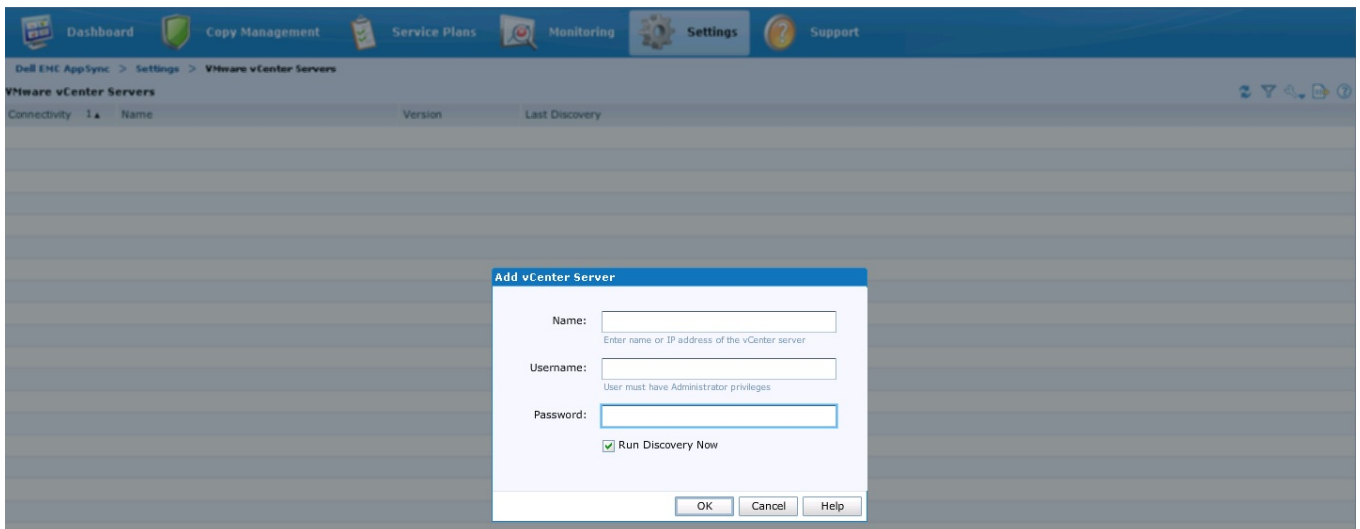
1. Select **Settings > VMware vCenter Servers**.
2. Click **Add**.
3. Type the vCenter Server name.
4. Type the credentials for an account that has Administrator privileges on the vCenter Server.

Note

AppSync allows you to mount a file system or a database (that is, the underlying storage LUN on which they reside) from a physical Windows or Linux environment to a VMware virtual environment as a RDM device. Ensure that you add the vCenter managing that virtual machine to AppSync before performing a mount.

5. Click **OK**.

Figure 6 Add vCenter Server



List of VMware datastores

The list contains VMware datastores that have been discovered and stored in the AppSync database.

Clicking on the datastore name displays the copies of the datastore.

The Service Plan column shows the plans that the datastore is subscribed to. Other details include the type of datastore (VMFS or NFS), and name of the ESX server.

Protect a VMware datastore

Protect a VMware datastore by subscribing it to an AppSync VMware service plan.

AppSync's protection mechanism for datastores is by means of service plans. You subscribe a datastore to a service plan and run the service plan immediately, or schedule the service plan to run at a later time.

- Choose **Subscribe to Plan and Run** when you want to protect selected datastores immediately. The service plan is executed for the datastores alone.

- Choose **Subscribe to Plan** when you want to schedule the protection for later. Protection for datastores that are part of the service plan are executed at the scheduled time.
- Choose an appropriate service plan from **Create a copy using** in the datastore **Copies** page.
- Choose **Run** from the VMware Datacenters Service Plan page to run the whole plan immediately.

Subscribing VMware datastores to a service plan

The **Subscribe to Plan** operation schedules the protection for later. Protection for all datastores that are part of the service plan are executed at the scheduled time.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > VMware Datacenters**.
2. Click a datacenter to display its datastores.
3. From this list, select the datastore to protect.
Select multiple datastores by holding down the Shift or Ctrl keys on your keyboard.
4. From the **Protect** list, select the appropriate service plan from **Subscribe to Plan**.

The selected plan appears in the Service Plan column for the datastore.

Protecting VMware datastores immediately

The **Subscribe to Plan and Run** operation adds datastores to an existing service plan and runs the service plan immediately for the selected datastores only.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > VMware Datacenters**.
2. Click a datacenter to display its datastores.
3. From this list, select a datastore to protect.
Select multiple datastores by holding down the Shift or Ctrl keys on your keyboard.
4. From the **Protect** list, select the appropriate service plan from **Subscribe to Plan and Run**.

The **Subscribe to Plan and Run** dialog appears displaying the progress through the different phases.

List of protected virtual machines

The list contains virtual machines belonging to datastores that are protected as part of a service plan run.

Click on the virtual machine name to display copies of the virtual machine. To perform a restore operation, select a virtual machine and click **Restore**.

Other details include the OS platform on the virtual machine, the version of the virtual machine, the ESX host on which the virtual machine resides, as well as the path to the

virtual machine file. In the path, the name of the datastore that the virtual machine resides on is within the [] parentheses.

Unsubscribing datastores from a service plan

You can unsubscribe datastores from a service plan or from all service plans that they are subscribed to.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > VMware Datacenters**.
2. Click a datacenter to display its datastores.
3. From this list, select the datastore to remove from a service plan.

You can select multiple datastores if they are subscribed to the same service plan. Select multiple datastores by holding down the Shift or Ctrl keys on your keyboard.

4. From the **Protect** list, select a service plan from **Unsubscribe from Plan**, or select **All** to remove the datastore(s) from all plans.

The service plan name is removed from the **Service Plan** column for the datastore(s).

VMware snapshots

When the VM consistency option is selected, AppSync creates snapshots of all the virtual machines that are in powered on state while the datastore is being replicated.

AppSync creates a Quiesced snapshot of the virtual machines that are in powered on state. VMware Tools is used to quiesce the file system in the virtual machine. Quiescing a file system is a process of bringing the on-disk data of a physical or virtual computer into a state suitable for backups. This process might include operations such as flushing dirty buffers from the operating system's in-memory cache to disk, or other higher-level application-specific tasks. If the VM consistency option is not set, AppSync skips the process of creating the virtual machine snapshots.

Viewing datastore copies

The copies are named with the time at which they were made.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > VMware Datacenters**.
2. Select a datacenter, then a datastore to view its copies.

You can see other copy details such as copy type, mount status, the VM consistency status and the service plan associated with the copy.

3. Select a copy to see more details in the **Details** pane.
 - **Virtual Machines** tab: lists the virtual machines that are part of the selected datastore copy.
 - **Virtual Disks** tab: lists the virtual disks that are part of the selected datastore copy.

- **Events** tab: lists the events that occurred when the datastore copy was created.

Viewing virtual machine copies

View the list of copies for a virtual machine based on the time they were made.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > VMware Datacenter**.
2. Select a datacenter, then the **Protected Virtual Machines** tab.
3. Select a virtual machine from the list to view its copies.

You can see copy details such as the size of the copy, storage on the disk, mode of the disk, and so on.

4. Select a copy to see additional details in the **Details** pane.
 - **Virtual Disks** tab: lists the virtual disks that are part of the selected virtual machine copy.
 - **Events** tab: lists the events that occurred when the virtual machine copy was created.

Creating a datastore copy from the Copies page

Create a copy of a datastore by subscribing it to an AppSync VMware service plan from the **Copies** page.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Copy Management > VMware Datacenters**.
2. Click a datacenter to display its datastores.
3. From this list, click the datastore to view its copies.
4. From the **Create a copy using** list, select the appropriate service plan.

The service plan runs immediately for the datastore.

Expiring a datastore copy on demand

Expiring a copy removes it from the AppSync database and can free up storage, depending on the replication technology and the copy state.

Before you begin

This operation requires the Data Administrator role in AppSync.

Expiring a copy that was made with RecoverPoint does not remove the corresponding bookmark from RecoverPoint itself.

Procedure

1. Select **Copy Management > VMware Datacenters**.
2. Click a datacenter to display its datastores.

3. Click the datastore whose copies you want to expire.
4. From the **Copies** page, select one or more copies to expire.
You can also perform this action from the Service Plan's **Copies** tab.
5. Select **Expire**.
6. Verify that you want to expire the copy you selected and any associated copies listed and confirm.

Service plan schedule

The schedule of a service plan is set in the **Plan Startup** phase.

The **Startup Type** (scheduled or on demand) determines whether the plan is run manually, or configured to run on a schedule. Options for scheduling when a service plan starts are:

- Specify a recovery point objective (RPO)
 - Set an RPO of 30 minutes or 1, 2, 3, 4, 6, 8, 12, or 24 hours
 - Minutes after the hour are set in 5 minute intervals
 - Default RPO is 24 hours
- Run every day at certain times
 - Select up to two different times during the day
 - Minutes after the hour is in 5 minute intervals
 - There is no default selected
- Run at a certain time on selected days of the week
 - One or more days of the week (up to all seven days) can be selected
 - There is no default day of the week selected. Default time of day is 12:00 AM.
- Run at a certain time on selected days of the month
 - Select one or more days of the month (up to all days)
 - Select one time of day. Available times are at 15 minute intervals.
 - Default is the first day of the month

Overriding service plan schedules

You can set different schedules for individual datastores subscribed to a service plan, overriding the generic recurrence setting.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can override only the settings of the recurrence type already selected for the service plan.

Procedure

1. Navigate to **Service Plans** and select one of the plans from the list.
2. From the **Settings** tab, select the **Plan Startup** phase.
3. In the **Plan Startup Defaults** pane on the right, note the **Recurrence Type** selected for the plan.

A recurrence type can be set only if **Scheduled** is selected as the **Startup Type**.

4. Click the **Plan Startup Overrides** tab.
You can see the list of all datastores subscribed to the plan.
5. Select one or more datastores and click **Override Schedule**.
6. In the **Override Schedule** dialog, set the schedule based on your requirement and click **OK**.
For example, if the default recurrence type is **On specified days of the month**, and the rule setting is to **Run at 12:00 AM on the 1st day of every month**, you can override the time and the day for individual datastores.
A Pencil icon indicates that default settings have been overridden.

Application discovery

Before creating the copy, AppSync performs discovery on the selected datastores and updates the AppSync database if there is any change in configuration of the vCenter server.

There are no user settings associated with this phase and it cannot be disabled.

Application mapping

After discovering the application, AppSync maps it to array storage, and protection services such as RecoverPoint.

There are no user settings associated with this phase and it cannot be disabled.

Create copy

The Create Copy phase creates a copy based on the preferred storage type specified by the user.

This phase specifies the type of datastore copy to make, and the storage settings for the copies. The copy phase creates a local copy, remote copy, or a local and remote copy based on whether you have chosen the bronze, silver, or gold service plan.

Review [Overview: Service Plan](#) for more service plan copy information.

Datastore copy options

Select the copy type, the virtual machines to ignore for snaps, storage preferences, and the number of snapshot copies to retain.

- **Copy Consistency**

`VM Consistent` creates a copy of the datastores in the service plan including running programs, processes, and even windows that were open at the time of the snapshot. `Maximum Simultaneous VM Snapshots` is the number of simultaneous snapshots of all VMs present. The default value is four snapshots.

`Crash Consistent` creates a copy of the datastores in the service plan.

Crash consistent copies have everything except data from the memory at the time of taking the snapshot.

`Configure VM Snapshots for VMs link` allows you to select virtual machines from the datastores added to the service plan. By default, the `Exclude VMs for Snapshot` option is enabled. This means that the selected VMs are ignored while taking VMware snapshots during the service plan run. If you select the `Include VMs for Snapshot`, only the selected VMs are considered for VMware snapshot creation during the service plan run.

`Include Virtual Machine Disk` includes all the datastores that are associated with the virtual machines running on the datastores being protected. For example, Datastore DS1 is subscribed to the service plan. Virtual Machine VM1 which is a part of DS1 has virtual disks in Datastores DS2 and

DS3. When the service plan runs, datastores DS2 and DS3 are protected along with DS1. However, datastores DS2 and DS3 are not subscribed to the service plan.

- **Storage Ordered Preference**- the preferred order of storage technology to use while creating copies. You can order, select, or clear storage preferences. Copies are created using the first technology preference when possible. If the first technology cannot be used, the remaining copies are processed using the next selected preference instead. For example, if the first preference was a bookmark but not all the application data in the service plan was mapped to RecoverPoint, then AppSync uses VNX snapshots instead. If you want AppSync to skip using a particular replication technology, deselect that preference from the storage ordered preference list.

Note

A single service plan can contain a mix of VNX block, VNX file, and RecoverPoint replication objects. For example, if you have a Bronze service plan for VMware, the datastores can be a mix of RecoverPoint, VNX file, and VNX block replication.

- **Expiration** - the maximum desired number of array snapshot copies that can exist simultaneously.

Automatic expiration of array snapshot copies

The automatic expiration value in a service plan's Create Copy phase specifies the maximum number of snapshot copies that can exist simultaneously.

When the "Always keep *x* copies" value is reached, older copies are expired to free storage for the next copy in the rotation. Failed copies are not counted. AppSync does not expire the oldest copy until its replacement has been successfully created. For instance, if the number of copies to keep is 3, AppSync does not expire the oldest copy until the fourth copy is created successfully.

This setting is independent of the VNX pool policy settings in Unisphere for automatic deletion of oldest snapshots. The service plan administrator should work with the storage administrator to ensure that the VNX pool policy settings enable the support of the specified number of snapshot copies for the application residing in that pool.

AppSync does not expire copies under the following circumstances:

- Mounted copies are not expired.
- A copy that contains the only replica of a datastore is not expired.

Include RecoverPoint copies in expiration rotation policy: Check this option to include RecoverPoint copies when calculating rotations.

Note

If this option is not selected, then RecoverPoint copies accumulates, and will remain until the bookmarks expire from the RecoverPoint appliance.

Unmount previous copy

The service plan unmounts a previously mounted copy after creating the new copy. The exception is a copy that was mounted on-demand as opposed to by the service plan. The on-demand mounted copy is not unmounted.

There are no user settings associated with this phase and it can be enabled or disabled.

Mount copy

The Mount copy phase mounts all the datastore copies created by that service plan run.

The **Mount Copy Defaults** settings for the copy to mount depends on the service plan. Other mount settings determine the mount host, access mode and mount signature.

This phase can either be enabled or disabled.

General Settings:

- **Mount on host:** lists all the ESX servers discovered on the registered vCenter servers.
- **Mount Signature:** lists **Use original signature** and **Use new signature** to select from. When **Use new signature** is selected, AppSync resignatures the VMFS volume on mount. Applicable only for VMware VMFS datastores.
- **Cluster Mount:** Select Yes or No .

RecoverPoint Settings:

- **Image access mode** (during RecoverPoint mount):
- **Logged Access:**
Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes.
- **Virtual Access with Roll:**
Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing.
- **Virtual Access:**
Provides nearly instant access to the image; it is not intended for heavy processing.
- **Desired Service Level Objective (SLO):** Specifies the required VMAX 3 Service Level Objectives. SLO defines the service time operating range of a storage group.

VNX File Settings: This option is available only for VMware VNXFile datastores.

- **Mount Copy with access:** Select the type of access the copy should be mounted with - Read-only or Read-Write.

Overriding mount settings in a service plan

If there are multiple VMware datastores subscribed to the same plan, you can select different mount settings for each datastore, overriding the generic mount settings.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Navigate to **Service Plans > VMware Datacenters** and click one of the plans from the list.
2. From the **Settings** tab, select the **Mount copy** phase.
3. On the right pane, select the **Mount Copy Overrides** tab.

The list of datacenters includes all vCenter datacenters whose datastores are subscribed to this plan. The mount settings display the default settings.

Additionally, for VMAX v3 Datastores, SLO Service Level Objective appears as another option.

4. Select the datastore whose settings you want to override and click **Set Overrides**.

Select multiple datastores by holding down the Shift or Ctrl keys on your keyboard.

5. On the **Mount Copy Overrides** dialog, select options only for those mount settings that you wish to override.

For example, if you want to mount a copy to the production host, you would select **Use new signature** from the **Mount Signature** drop-down.

Fields that do not have a selection retain their default settings.

6. Click **OK**.

A pencil icon appears in the first column of the datastore's row whose default mount settings you changed.

7. To revert back to default settings, select the datastore(s) and click **Use Default Settings**.

Unmount copy

The final phase in the service plan unmounts the copy.

This phase is disabled if the **Unmount previous copy** phase is enabled. There are no user settings associated with this phase.

Considerations when mounting a VMFS copy

When you mount a VMFS copy to an alternate ESX Server, AppSync performs all tasks necessary to make the VMFS visible to the ESX Server.

- After these tasks complete, further administration tasks such as restarting the virtual machines and the applications must be completed by scripts or manual intervention.
- For datastore and virtual disk mounts on ESXi 5.x and RecoverPoint 4.0 environments, disable hardware acceleration to ensure successful virtual access type mounts. For more details, refer VMware Knowledge Base article 2006858.

Mounting a datastore copy on-demand

You can initiate an on-demand mount of a datastore copy from the datastore's **Copies** page, service plan's **Copies** tab or from a datacenter's **Datastore** page.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. From the **Recover** popup button, select **Mount a Copy** in the **Datastore** or **Copies** page or Service Plan **Copies** tab.
2. Use the **Copies** or **Service Plan** filters to select the appropriate copy to mount.
The copies list is refreshed based on the filters selected.
3. Select the copy to mount.

For a RecoverPoint copy, you also have the option to select a bookmark based on a specific time. However, there should be a copy available in AppSync prior to the time you select.

Click **Select a point in time** to select a copy with a specific time stamp and select the location of the copy (remote or local). Select **remote** to mount remote copy or **local** to mount local copy. The time shown here is the AppSync console's time. If the console is in a different time zone from the RecoverPoint Appliance (RPA), specify the time as per the server's time zone to mount the copy.

4. In the **Mount Additional Copies** page, select one or more additional copies to mount. The copies listed here are of other datastores that were protected at the same time and on the same datacenter as the copy you selected in the previous step.
5. On the **Select Mount Settings** page:
 - a. From the **Mount on Host** list, select the host on which to mount the copy.

All ESX Servers under the vCenter Server registered with AppSync are listed as possible mount host candidates. If a selected ESX is part of an ESX cluster, another field, Clutser Mount, is displayed, which can be set to Yes or No. If the ESX Server is part of a cluster you also have the option of mounting to all nodes of the cluster or just to the mount host you choose.
 - b. From the **Mount Signature** list, select from one of these options: **Use new signature**, **Use original signature**. Applicable only for VMware VMFS datastores.
 - c. In case of a RecoverPoint copy, from the **Image Access Options** list, select from one of these options:
 - **Logged access:** Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. Logged access is the only option available when you mount to the production host.
 - **Virtual access with roll:** Provides nearly instant access to the copy, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing. With RP VMAX, and RP XtremIO, virtual access with roll is not supported.
 - **Virtual access:** Provides nearly instant access to the image. Virtual access is not intended for heavy processing. Virtual access with RP VMAX and RP XtremIO is not supported.
 - d. From the **Mount copy with access** drop-down list, select the type of access the copy should be mounted with: **Read-only** or **Read-write**. Applicable only for VMware VNXFile datastores.
 - e. For VMAX 3 arrays, select the desired Service Level Objective (SLO) for the mount copy.

Note

The SLO values are dynamically fetched from the VMAX 3 arrays, and only the unique values are displayed.

- f. For VMAX V2 arrays, select the desired FAST VP policy for the mount copy.
- g. From the **VPLEX Mount option** list, select one of the following:
 - **Native array:** Use this option if you want to mount the copy as native array volumes.
 - **VPLEX virtual volume mount:** Use this option if you want to mount the copy as VPLEX virtual volumes.
 - **Enable VMware cluster mount:** Clear this option if you do not want to perform an ESX cluster mount. By default, this option is enabled.
6. From the **Summary** page, review the copy and mount settings that you chose in the previous pages and click **Finish** to mount the copy.
7. In the **Results** page, click **Details** link to see the progress of the different phases that are part of mounting a copy.

The last phase completed is displayed at the bottom of the list of phases.

Unmounting a VMware datastore copy

When you select a copy to unmount, other copies that were mounted along with the selected copy are unmounted.

Before you begin

This operation requires the Data Administrator role in AppSync.

You can unmount a copy only from a list of copies made for a datastore.

Procedure

1. Navigate to the **Copies** page from the **Protection** or **Service Plan** pages:
 - **Copy Management > VMware Datacenters** > select the VMware datacenter that hosts the datastore, then select the datastore with the copy to unmount.
 - **Service Plans > VMware Datacenters** > select a service plan, then select the **Copies** tab.
2. From the list of copies, select the datastore copy and click **Unmount** from the button below.

The **Unmount Confirmation** dialog displays all the copies of other datastores that were mounted along with the selected copy to be unmounted.

3. Click **Yes** to confirm the unmount of all the copies shown in the dialog.

The **Unmount** window displays the progress of the unmount operation. All copies that were mounted along with the selected copy will be unmounted.

Restoring a datastore from a copy

You can perform a restore of a datastore copy from the datastore's Copies page, the Service Plan's Copies tab, or from a datacenter's Datastore page.

Before you begin

- This operation requires the Data Administrator role in AppSync.
- Prior to restoring a datastore, it is recommended that you power off the VMs in the datastore.

Procedure

1. Select **Recover > Restore** from Datastore or Copies page or Service Plan Copies tab.

The Datastore Restore wizard launches.

2. Select the copy to restore.

Use the **Copies** or **Service Plan** filters to select the appropriate copy to restore. The copies list is refreshed based on the filters selected.

For a RecoverPoint copy, you also have the option to select a bookmark based on a specific time. However, there should be a copy available in AppSync prior to the time you select.

Click **Select a point in time** to select a copy with a specific time stamp. The time shown here is the console's time. If the console is in a different time zone from the RPA, specify the time as per the server's time zone to restore the copy.

3. Click **Next**.

If the selected copy has affected entities, the **Restore Warnings** page is displayed.

4. Read the warning messages for the affected datastores. Select the checkbox to indicate your agreement to restore other entities along with the selected copy.

You can manually unmount the datastores that will be overwritten prior to restore.

Only RecoverPoint copies have affected entities.

5. In the **Virtual Machine Operations** step, select the appropriate actions that you want AppSync to perform before and after restore. See [Virtual Machine Operations](#) for details.

6. Click **Next**.

The **Configure Storage Options** screen appears.

- **Wait for mirror rebuild to complete** - This option is applicable for VPLEX Snap copies whose production data resides on local or distributed RAID-1 volumes. By default, this option is selected.
- **Disable VMware SRM** - Allows you to manage consistency groups, if the SRM flag is enabled on the RecoverPoint consistency group. This is only applicable for RecoverPoint 4.1 and later.
- **Perform device restore in background** - Allows you to optimize restore of VMAX devices. If you select this option, AppSync restore operation does not wait for VMAX track synchronization to complete. The production application is available instantly.

7. In the **Summary** page, review the settings that you selected in the previous pages and click **Finish** to perform the restore.

8. In the **Results** page, click **View Details** to see progress of the different phases that are part of restoring a copy.

The last phase completed is displayed at the bottom of the list.

Virtual Machine Operations during restore

AppSync can perform operations on the virtual machines associated with the datastores selected for restore.

Table 32 Virtual Machine operations

Virtual Machine Operation	Description
VMs present at start of restore	Power down VMs at start of restore: If the virtual machines are present at the start of restore, AppSync shuts them down prior to beginning the restore operation. This is important for a successful restore.
Perform VM operations after restore	<ul style="list-style-type: none"> • Return VMs back to state found at start of restore: After restore, AppSync powers on the virtual machines and returns them to the same state they were at prior to restore. • Register all virtual machines: After restore, AppSync registers all virtual machines to the vCenter inventory. • Register and power up all virtual machines: After restore, AppSync powers on the virtual machines and registers them to the vCenter inventory. <p>You can select only one of these options.</p>
VMs not present at start of restore. Perform VM operations after restore.	<ul style="list-style-type: none"> • Register all virtual machines: After restore, AppSync registers all virtual machines to the vCenter inventory. • Register and power up all virtual machines: AppSync powers on the virtual machines and registers them to the vCenter inventory.

Datastore affected entities during restore

When you restore a datastore, AppSync calculates affected entities for other datastores that share the same storage.

An affected entity is data that resides on your ESX server that unintentionally becomes part of a replica because of its proximity to the data you intend to protect. You can prevent affected entity situations by properly planning your data layout.

In case of RecoverPoint or ViPR Controller, the granularity is at the consistency group (CG) level. If the CG is selected for restore, AppSync identifies other datastores residing on the same CG that were also protected alongside, and restores them. If the affected entity was not protected, AppSync will not be able to restore it properly. This is displayed as a warning in the Restore wizard.

There are no affected entities for VNX because multiple datastores cannot span the same LUN and multiple datastores cannot be hosted on the same File System.

If there are affected entities in your underlying storage configuration, the Restore Wizard notifies you of these items requiring you to acknowledge that additional items will be restored.

Restoring a virtual machine from a copy

You can perform a restore of a virtual machine from the **Protected Virtual Machines** tab or the virtual machine's **Copies** page.

Before you begin

- This operation requires the Data Administrator role in AppSync.
- You must be using vSphere Enterprise Edition.
- All datastores used by the virtual machine must be protected by selecting the **Include Virtual Machine Disk** option in the **Create copy** phase of the service plan.
- The virtual machine should not have any pre-existing snapshots.
- Virtual machines with RDMS cannot be restored.

Procedure

1. Select **Restore** from the **Protected Virtual Machines** tab or the virtual machine's **Copies** page.

The virtual machine Restore Wizard launches.

2. Select the copy to restore.

Use the **Copies** or **Service Plan** filters to select the copy to restore. The copies list is refreshed based on the filters selected.

For a RecoverPoint copy, you also have the option to select a bookmark based on a specific time. However, there should be a copy available in AppSync before the time you select.

Click **Select a point in time** to select a copy with a specific timestamp. The time shown here is the console's time. If the console is in a different time zone from the RPA, specify the time according to the server's time zone to restore the copy.

3. Click **Next**.

If other VMs were also protected along with the selected virtual machine, the **Multiple VM Restore** page is displayed.

Select one of the following options:

- Continue to restore only one virtual machine
 - View and/or select the other VMs for restore
4. In the **Select Restore Location** page, make the appropriate selections. See [Virtual Machine restore options](#) for details.
 5. In the **Select Mount Host** page:
 - a. Select the mount ESX.

If production data resides on RecoverPoint storage, the target devices should be visible to the selected mount host. For VNX, the mount host

should be registered to the VNX. For all other storage, the mount host should be registered to the storage array where the copy resides.

- b. For a RecoverPoint bookmark copy, select the RecoverPoint image access mode from the list - Logged Access, Virtual Access, or Virtual Access with Roll.
- c. For a VPLEX Snap copy, select the VPLEX mount option - VPLEX virtual volume or native array volume. For VPLEX virtual volumes, the mount host needs to be added to VPLEX storage view; for native array volumes, it needs to be zoned to the VPLEX backend array where the snapshot is created.
- d. Click **Next**.

Note

- AppSync employs VMware vMotion technology to move the virtual machine from mount host to restore location. Therefore, the mount host and host at the restore location should satisfy the VMware vMotion prerequisites such as network requirement.
 - In the case of a VPLEX Snap copy, if the ESX which is selected to mount the datastore for VM restore is part of an ESX cluster, the datastore is mounted only on that ESX and not on all the ESXs of that cluster. You must select the same ESX under **Select Restore Location** and **Select Mount host**, if you do not want VM files to be copied over the network.
-

In the **Choose Instant Restore** page, you can make a selection only if one of the following conditions is met:

- The mount and restore hosts are the same.
 - The mount and restore hosts are different but are nodes of the same ESX cluster.
6. In the **Choose Instant Restore** page, select **Yes** or **No** for the **Do you want to perform an instant restore option** option, based on whether you want to perform an instant restore.

During instant restore, you can continue to use the virtual machine. Though the virtual machine is powered on, the VMs are restored in the background.

If you select **No**, and if you had chosen to restore multiple virtual machines in Step 2 of this wizard, specify a number in the **Maximum number of simultaneous virtual machines to be restored** box. By default, the number is 2.

Note

If you are restoring multiple virtual machines belonging to a vApp, set **Maximum number of simultaneous virtual machines to be restored** to 1.

The Instant restore option is not available for:

- VMAX copies if the source devices are thick
- VMAX copies
- ViPR snap copies

- VPLEX snap copies
7. In the **Summary** page, review the settings that you selected in the previous pages, and then click **Finish** to perform the restore.
 8. In the **Results** page, click **View Details** to see progress of the different phases that are part of restoring a virtual machine.

The last phase completed is displayed at the bottom of the list.

Virtual Machine Restore options

You can select the restore location as well as restore operations.

Table 33 Virtual machine restore options

Restore Option	Description
Original location	<p>Restores to the location where the virtual machine was present at the time of protection.</p> <hr/> <p>Note</p> <p>For a RecoverPoint copy, restoring to the original location is not recommended. AppSync displays an appropriate warning when you select this option.</p>
Alternate location	<p>Restores to a location selected from the following options. All are mandatory.</p> <ul style="list-style-type: none"> • vCenter Server: You can select either the same vCenter Server where the datastore with the virtual machine was at the time of protection or a different server. • Datacenter • Host • Datastore
Options if the VM being restored already exists in the restore location	<ul style="list-style-type: none"> • Fail the restore: AppSync checks for the existence of the virtual machines in the restore location. For those virtual machines that exist in the restore location, the restore operation is aborted. For the rest, the restore operation continues. This is a precautionary option. • Create a new virtual machine: AppSync creates a new virtual machine before restoring. • Unregister the virtual machine: If the virtual machines selected for restore exist in the restore location, AppSync unregisters them from the inventory before restoring.

Table 33 Virtual machine restore options (continued)

Restore Option	Description
	<ul style="list-style-type: none"> • Delete from disk before performing restore: If the virtual machines being restored exist in the restore location, AppSync deletes them before restoring. <hr/> <p>Note</p> <p>It is recommended you take a backup of the virtual machine before proceeding with the restore operation.</p> <hr/> <ul style="list-style-type: none"> • Delete from disk after performing restore: If the virtual machines being restored exist in the restore location, AppSync deletes them after restoring.

File or folder restore with VMFS or NFS datastores

Files or folders stored on virtual disks on a virtual machine in VMFS and NFS datastores can be restored through AppSync.

The virtual disks stored in a VMFS or NFS datastore that are protected by an AppSync service plan can be used for file or folder level restore by specifying the location for mounting the virtual disk copy.

Within AppSync, file or folder level restore is a three phase process: To complete the restore, the final step is performed manually outside of AppSync. You must copy the files or folders from the location where the virtual disk is mounted to a location of your choice.

1. AppSync mounts the datastore snapshots to the ESX server on which the virtual machine with the AppSync agent resides.
2. The vCenter server adds the virtual disks from the datastore snapshots to the mount VM without powering off the VM.
3. AppSync agent performs a filesystem mount to the mount VM.

Restore of files or folders from virtual disks with multiple partitions is supported.

If the ESX server version is 5.0 and higher, the original VM can also be the mount VM.

Restrictions

- File or folder level restore is not possible on dynamic disks.
- Virtual disks belonging to same phase pit cannot be mounted to the same virtual machine even when created using the Gold service plan. In other words, if a virtual disk from a local copy is mounted; then the same virtual disk from a remote copy cannot be mounted.
- To perform an Any Point in Time (APIT) file restore, you must first perform an APIT mount of the datastore and then launch the Granular File Restore wizard from the APIT copy.

Restoring a file or folder from a virtual disk

You can perform the restore of a file or folder of a virtual disk from the **Protected Virtual Machines** tab or the virtual machine's **Copies** page.

Before you begin

- This operation requires the Data Administrator role in AppSync.
- You must be using vSphere Enterprise Edition.
- The virtual machine on which the copy is mounted and restored must be 64-bit with Windows 2008 or Windows 2012 as the operating system. The AppSync host plugin must be installed on it and it should be registered with the AppSync server.

Procedure

1. Select **Restore** > **File** from the **Protected Virtual Machines** tab or the virtual machine's **Copies** page.

The Granular File Restore Wizard launches.

2. Select the copy to restore.

Use the **Copies** or **Service Plan** filters to select the appropriate copy to restore. The copies list is refreshed based on the filters selected.

3. Click **Next**.

The **Select Virtual Disk** page appears displaying the virtual disks in the selected virtual machine.

4. Select the virtual disks whose files or folders must be restored and click **Next**.

5. In the **Select Host** page:

- a. Select the virtual machine on which the copy must be mounted.

In addition, specify the location in the selected virtual machine where the disk must be restored to. By default, the files are restored to the following location: `%system drive%\AppSyncMounts\ where:`

- `%system drive%` is system drive of the selected virtual machine on which the copy is to be mounted
- `<VM_name>` is the name of the virtual machine whose virtual disks are being restored
- `<copy_id>` is an AppSync generated ID
- `Hard disk#` is the number of the hard disk in the virtual machine. This number is the same as on the original virtual machine.

- b. Select the RecoverPoint image access mode from the list - Logged Access, Virtual Access or Virtual Access with Roll.

6. In the **Summary** page, review the settings that you selected in the previous pages and click **Finish** to start the restore of the disk.
7. In the **Results** page, click **View Details** to see progress of the different phases that are part of restoring a virtual disk.

The last phase completed is displayed at the bottom of the list.

8. Next, perform the manual step of copying the required files or folders from the mount location to a location of your choice.

9. Optionally, unmount the datastore. [Unmount a datastore copy.](#)

CHAPTER 10

Repurposing

- [Repurposing overview](#).....250
- [Using the Repurpose wizard](#).....254

Repurposing overview

This topic explains how to use the AppSync repurposing feature for database and file systems.

AppSync allows you to create copies of your database and file systems for application testing and validation, test and development, reporting, data masking, and data analytics. AppSync identifies copies that are created from a repurpose action as first generation and second generation copies. The source of a second generation copy is a first generation copy. You can create multiple second generation copies from a first generation copy.

AppSync supports repurposing on File systems, SQL Server and Oracle databases.

There are two types of repurposing:

- Native array repurposing - The first generation copy is a copy of the source database. For example, in the case of an XtremIO array, snapshot of the source is the first generation copy.
- RecoverPoint bookmark repurposing - The first generation copy is a copy of the LUNs at the local or remote replication sites in the RecoverPoint consistency groups.

Note

To create a snap of a bookmark on a remote site (remote RecoverPoint repurposing), add both the local and remote native array to AppSync.


Review the following additional information about repurposing:

- A first generation copy creates a copy that can be used as source for other copies.
- Repurpose copies do not figure in RPO calculations.
- You can create first generation and second generation repurpose copies on-demand or schedule it.
- Restore is not supported for second generation copies.
- Restore of a first generation copy is not supported in the case of RecoverPoint bookmark repurposing.
- Restore from a first generation copy is not allowed in the case of VMAX, if it is created on the remote site.
- The first generation copy of a database creates an application consistent copy. It includes application discovery, mapping, and database freeze/thaw. For File systems, you can configure freeze/thaw operation using callout scripts. [Create copy phase features freeze and thaw callout scripts](#) provides more information on configuring and running freeze and thaw callout scripts.

Note

For a first generation copy of SQL, you can configure a VSS retry count and retry interval for freeze/thaw operation using the Repurpose Wizard. VSS retry options are not applicable for Crash Consistent SQL copies.

- Second generation copies are created using the first generation copy as the source without impacting the application. They do not include application discovery, mapping, and application freeze/thaw. If a first generation copy is mounted with recovery, and if the second generation copy is refreshed, the second generation copy might not be recoverable after the mount.



Additional Notes

- SQL
 - Log backup is not supported as part of repurposing.
 - Repurposing of multiple SQL databases is not supported.
- Oracle - RMAN options are not available in the Repurposing wizard.
- File system
 - Repurposing of NFS file system is not supported.
 - When repurposing multiple file systems together, failing to protect one or more file systems fails the repurpose operation completely.
 - You cannot simultaneously protect two file systems residing on two separate storage arrays.
 - File systems that are repurposed together are mounted together.
 - File systems that are protected together are repurposed together for second generation copies.

Repurpose schedule

- If you attempt to create both the first generation and second generation copies simultaneously using the Repurpose wizard from the Database page, the second generation copy is created automatically after the first generation copy is created. This is applicable for **Run Now**, **Run Recurrently As Per Schedule**, and the **Run Only Once At later time** options.
- If you create a schedule for the second generation copy, the second generation copy is not triggered after the first generation copy is created. The second generation copy runs according to the schedule.
- If you create a second generation copy using the Repurpose wizard from the Copies page, the second generation copy is not triggered even though the first generation copy runs according to the schedule. However, if the second generation copy is scheduled, it runs according to the schedule.
- On the first schedule, a repurposed copy is created, and on subsequent schedules, the copy is refreshed.
- Multiple file systems can be scheduled together as part of the same repurpose plan.
- If multiple file systems are scheduled to be repurposed together as part of the same schedule, selecting any of the file system and viewing the schedule displays

the schedule. If you delete the schedule, the schedule for all the file systems that were scheduled to be repurposed together are deleted.

Modifying a repurpose plan

Each copy is associated with a unique repurposing plan. To modify a repurpose plan:

1. On the AppSync console, select **Copy Management** > *application*.
2. Select a copy.
3. Click on the **Repurpose** link in the service plan column of the Copies page.

Note

- The options that you cannot modify are disabled.
 - If you modify the label, the updated label is reflected in the Copies page only when you refresh the copy.
-

Repurpose refresh

Refresh means to discard the current copy data and recreate the copy data using its parent. When you refresh a copy, changes from the source are reflected in the copy. This is normally done by creating a new point-in-time copy and expiring the old copy.

- First generation and second generation copies can be refreshed.
- Refreshing a first generation copy creates an application consistent copy with a new time.
- Second generation copies are not modified if you refresh the first generation copy.
- Refresh of a second generation copy recreates the second generation copy with the first generation parent. (Used for discarding changes of second generation copy and starting over.)
- The timestamp on the second generation copy is the same as first generation copy. If the first generation copy is refreshed, then the timestamp differs.
- When you refresh a mounted repurpose copy, AppSync unmounts the copy, expires the copy and creates a new copy and the refreshed copy is mounted back with the same options as previous mount operation.
- When you refresh a mounted RecoverPoint repurposed copy, AppSync unmounts the copy and refreshes the copy. It is not mounted back by the end of the refresh operation.
- When a repurposed copy is refreshed because of a scheduled repurpose service plan run, the refreshed copy is mounted with the mount options as specified in the mount copy phase of the repurpose plan.

You can refresh a repurposed copy at any time. To start the refresh:

1. From the **Copies** page of the AppSync console, select the repurposed copy that you want to refresh.
2. Click **Refresh**.

Native array refresh

When storage arrays provide copy refresh capability natively (that is, the ability to refresh the same copy instead of creating a new one), AppSync uses it instead of expiring the copy and recreating a new copy to refresh the copy data.

Native array refresh benefits

- Reduces the time required to complete a refresh workflow.
- If the copy being refreshed is mounted to a host, then the storage LUNs are not removed from the host. Only the application is stopped, the file system is unmounted, data is refreshed on the existing copy, file system is mounted back, and the application is started. In a VMWare virtual environment where application resides on virtual disks, the virtual disks are removed from the virtual machine, but LUNs are not removed from ESX.
- Eliminates the need for a rescan in a VMWare virtual environment because storage LUNs are not removed from the mount host.
- The WWN, volume, and device details of the mounted LUN remains the same because storage LUNs are not removed from the mount host. This is beneficial in the case of some backup scenarios, where external post-mount scripts depend on the copy WWN, volume, or device details for any action.

Native array refresh usage

- On Unity, first generation copies are refreshed using the native array refresh. However, for a second generation copy, native array refresh is used only if the second generation snapshot is refreshed when the first generation snapshot is ATTACHED.

Note

This is only applicable for Unix based applications.

- On XtremIO, if all the LUNs to be protected are in a single XtremIO consistency group, native array refresh is used.
- In the case of VPLEX virtual devices on XtremIO storage devices, native array refresh is used, if all XtremIO LUNs belong to the same consistency group.
- For Windows based applications, the application must be fully part of a single XtremIO consistency group. For example, if a SQL database resides on LUNs from two different XtremIO consistency groups, then native array refresh is not used.
- For Oracle databases, the data, redo, and control files must reside on one consistency group and archive log, and FRA must be on a separate consistency group.

Repurpose expire

You can expire a repurposed copy when you no longer need the copy.

Note

In the case of VMAX 2 arrays, the session still persists on the array even after you expire a copy.

Data masking using scripts

You can use the AppSync repurposing feature to mask sensitive data.

To mask data using scripts:

1. Mount and recover the first generation copy of a database.

2. Apply data masking on the same copy using the post mount script phase and unmount the database. The post mount script phase provides a way to execute a script.
3. Unmount the first generation copy.
4. Create a second generation copy from the first generation copy using the Repurpose wizard.

Using the Repurpose wizard

Use the Repurpose wizard to schedule or immediately create first generation or second generation copies as required.

Before you begin

This operation requires the Data Administrator role in AppSync.

Procedure

1. Log in to the AppSync console and go to **Copy Management**.

Option	Description
To Repurpose an SQL copy:	<ol style="list-style-type: none"> a. Select Microsoft SQL Server. A list of available SQL server instances appears. b. Select an instance. c. On the Database folders page, select User databases. The list of available database appears. d. Select the database that you want to repurpose.
To Repurpose an Oracle copy:	<ol style="list-style-type: none"> a. Select Oracle. A list of available databases appears. b. Select the database that you want to repurpose.
To Repurpose a File system copy:	<ol style="list-style-type: none"> a. Select File system. A list of available hosts appears. b. Select a host. A list of available file systems appears. c. Select one or more file systems that you want to repurpose.

2. Select **Repurpose** from the drop-down list in the lower left of the console screen.

This action launches the Repurpose wizard, and leads you to the Intentions page where you can select the action that you want to perform:

- **Create First Generation copy**
- **Create First Generation copy and a Second Generation copy**

3. Select the desired copy type, and select **Local** or **Remote** in the Site drop-down to continue creating the copy.

Note

Creating first generation and second generation copies on the remote site is local to the remote array.

4. Select **Use Bookmark as an intermediate step**, and then click **Next**.

The option is available only if RecoverPoint is registered.

5. Click **Next** to launch the **Settings** screen.

From the Settings screen, you can define the specific options for first generation and second generation copies. Specifically you can

- Define labels for each copy to help identify the copy purpose.
-

Note

In the case of file systems, first generation and second generation labels determine the callout script names. Freeze and thaw callout scripts can be used and is only applicable for first generation copies. Unmount callout scripts can be used for both first generation and second generation copies. The name of the freeze/thaw/unmount callout scripts uses the first generation label. [Create copy phase features freeze and thaw callout scripts](#) provides more information on configuring and running freeze and thaw callout scripts.

- Select application-specific copy options for the first generation copy only.
 - Configure storage options (for first generation copy only).
 - Choose appropriate copy type (wizard fails if incorrect type is chosen).
6. Select the desired options for the copy, and then click **Next**.
The schedule page of the wizard appears allowing you to identify when the first generation copy must be created. Create the copy now or schedule the copy.
 7. Select one of the following scheduling options:
 - Run Now - Creates a copy when you click **Finish** on this wizard.
 - Run Recurrently As Per Schedule - Creates a copy based on the specified recurrence type. On the first schedule, a repurposed copy is created, and on subsequent schedules, it refreshes the copy.
 - Run Only Once At later time - Creates a copy only once on the specified date and time.
 8. Click **Next** to complete the wizard.

The AppSync Service Plan In-Progress Monitor

The AppSync Service Plan In-Progress Monitor allows you to view all currently running repurpose activities, and monitor their progress. The AppSync Service Plan In-Progress Monitor shows the item being repurposed (source) and the label of the item being created or refreshed along with the application type.

View or delete repurpose copy schedules

You can view or delete a repurpose copy schedule.

Procedure

1. From the Appsync console, navigate to **Copy Management > Application**, and then select the appropriate database or file system.
2. Select **Repurpose > View Repurpose Copy Schedule** to view all scheduled repurpose actions.
3. Select one or more schedules, and click **Delete** to delete the repurpose schedule.

Note

In the case of file systems, deleting the schedule by selecting one file system also removes the schedule for other file systems.

View repurposed copies

You can view repurposed copies for File systems, Oracle databases and SQL application instance.

Procedure

1. From the Appsync console, navigate to **Copy Management**, and then select the appropriate database or application instance.
2. Click **Repurposed Copies** to view the repurposed copies for all the databases, or application instance.

The **Repurposed Copies** window appears. It lists all the first generation and second generation copies.

You can perform the following operations:

- Refresh
- Repurpose
- Mount
- Unmount
- Restore

CHAPTER 11

Monitor AppSync

This chapter includes the following topics:

- [RPO concepts and best practices](#)..... 258
- [Alerts and associated events](#).....259
- [Email alerts](#)..... 260
- [The AppSync Service Plan In-Progress Monitor](#)..... 261

RPO concepts and best practices

A recovery point objective (RPO) is one of several scheduling options that can be selected as part of a service plan's Plan Startup phase.

When you subscribe an object (such as an Exchange database) to a service plan that uses RPO as its recurrence type, the object acquires the recovery point objective specified in the service plan.

Since you can subscribe an object to more than one service plan, it is possible for an object to have more than one recovery point objective. When an object has more than one RPO, the service plan with the highest RPO frequency (that is, the lowest RPO hours value) is used for calculation and reports.

As a best practice, you should subscribe an object to only one RPO-enabled service plan. If you subscribe an object to additional service plans, they should not use the RPO-based recurrence type.

Note

RPO is an alerting mechanism only, no copies are initiated based on this setting.

Recovery point compliance report

The recovery point compliance report shows the recoverability for all objects that are subscribed to service plans with an RPO recurrence type. The report is at **Monitoring > Recovery Point Compliance Report**.

Table 34 Recovery Point Compliance Report

Column	Description
Server	Host level object, such as a Microsoft Exchange Mailbox Server
Application	Name of the protected object, such as a Microsoft Exchange database. Click the name to go to the list of copies for the object.
Recovery Point Objective	The recovery point objective as defined in the Start phase of the associated service plan.
Time Since Last Recovery Point	Amount of time since the last copy or bookmark created by the associated service plan. A green icon indicates the copy is RPO compliant. A red icon indicates non-compliance.
Service Plan	Name of the service plan. Click the name to go to the service plan definition.

Exporting an RPO compliance report to CSV

You can create a recovery point objective (RPO) compliance report in comma-separated value format.

Before you begin

No particular AppSync role is required for this operation.

Procedure

1. Navigate to **Monitoring > Recovery Point Compliance Report**.
2. Sort and arrange columns as desired for the report.
3. Click **Export** to run the Export wizard.

You have the option to include table headers and export only selected rows. The default file name is "Recovery Point Objectives_*mm_dd_yyyy_hh_mm_ss*.csv".

Summary of RPO compliance

The Recovery Point Objectives (RPO) summary on the dashboard shows the percentage of RPOs met across all objects that are subscribed to RPO-enabled service plans.

Alerts and associated events

AppSync generates an alert when a service plan phase fails, when a recovery point objective (RPO) is not met, or when a mount or restore fails.

Service plan failure alerts are generated immediately on failure of a service plan phase. When an application goes out of RPO compliance, the associated alert is generated within 1 hour. Refer to [Acknowledging alert icons for database, file system, and data store service plan runs](#), for information.

AppSync displays alerts in the console at **Monitoring > Alerts**.

Table 35 Details of alerts

Column	Description
Alert (!)	Level of alert
Time	Date and time of the alert.
Server	Application server, such as a Microsoft Exchange server.
Application	Replicated object, such as a Microsoft Exchange database.
Category	RPO, Phase Failure, License, and Other.
Service Plan	Service plan that was running when the alert was generated, or the service plan that created the copy that failed a mount or restore.
Message	Describes the cause of the alert.
Alert Acknowledged	Indicates if the alert has been acknowledged. Note that acknowledged alerts will not display in the AppSync Dashboard.

You can filter alerts by the time they were generated and by the associated service plan.

View the associated events that led up to the alert by clicking the alert. Expand the top-level events to see additional details. You can filter associated events by any column.

Acknowledging alerts

With AppSync version 2.2.2 and later you can acknowledge alerts. Refer to [Acknowledging alerts](#), and [Acknowledging alert icons for database, file system, and data store service plan runs](#) for more information.

Acknowledging alerts

These steps show you how to acknowledge alerts for monitoring.

You can choose to acknowledge alerts that are shown in the console, **Monitoring > Alerts**. The alerts page display shows a column for **Alert Acknowledged**.

A value of No is the default. When you acknowledge an alert, the value of the alert changes to YES from the default value NO.

Procedure

1. Go to **Monitoring > Alerts** and select an alert from the alerts table with a current value of **NO** in the Alert Acknowledged column.
2. Click **Acknowledge Alert**.

Results

The alert displays a value of **YES** in the Alert Acknowledged column of the Alert table.

Acknowledging alert icons for database, file system, and data store service plan runs

You can acknowledge an alert icon within the AppSync console for Exchange, Oracle and SQL databases, File systems, and VMWare datastores.

An alert icon indicates the status of the most recent service plan run. The icon appears beside a database, file system, or datastore after the run.

After you acknowledge the icon in the Acknowledge column, AppSync changes the icon to an information icon. Also, you can acknowledge the alert after every Service Plan run. You are not restricted to acknowledge an alert only once.

Note

If a delete is pending, then the Acknowledge button becomes disabled for the database, file system, or data store.

The following procedure shows you how to view and acknowledge the alert.

Procedure

1. Select an alert from the database, file system, or data store table that has an alert icon that is associated with the last service plan run.
2. Click the **Acknowledge Alert** button that is located below the database, file system, or data store table.

The alert icon changes to an information icon.

3. If required, re-run the service plan on the same database, file system, or data store to display the alert icon for this run.

Email alerts

You can configure AppSync to send alerts via email to a list of recipients. By default, only failure alerts are emailed.

You enable alert emails and add recipients per application instance, such as an Exchange mailbox server.

To configure AppSync to also send an email after successful completion of a scheduled service plan:

1. Select **Settings > AppSync Server Settings**.
2. Select **Notify service plan success**.
3. Click **Apply**.

Configure server settings for email alerts

Configure SMTP services on a machine that the AppSync server can access.

Before you begin

This operation requires the Resource Administrator role in AppSync. Refer to SMTP documentation for configuration procedures.

Procedure

1. Select **Settings > AppSync Server Settings**.
2. Under **Send Mail** settings, enter values for SMTP server, port, sender, and recipient.

Note

License non-compliant alerts are sent to the specified recipient address.

Use **Test** to validate the settings. Enter a recipient's address and then check the recipient's inbox. You can also select **Insert Default Values**.

3. Select **Notify service plan success** to receive notifications on the successful completion of service plans.
4. Click **Apply**.

Specify email alert recipients

Configure email alerts per application instance level.

Before you begin

- This operation requires the Data Administrator role in AppSync.
- SMTP services must be configured on a machine accessible by the AppSync server. See [Configure server settings for email alerts](#) on page 261.

Procedure

1. Select **Copy Management > Application (select application)**.
2. Select the application object for which to set alert settings.
3. Select **Alert Settings**.
4. Enter one or more email recipients, separated by commas (,) and click **OK**.
A test email is sent to recipients for verification.

The AppSync Service Plan In-Progress Monitor

This topic describes the AppSync Service Plan In-Progress Monitor.

The AppSync Service Plan In-Progress Monitor shows all in-progress protection, repurpose, and refresh processes. It enables you to view currently running service plan and repurpose activities and to monitor their progress. It shows the current copy you

are repurposing (making a snap of a snap copy), and related copy labels such as 1st Gen (generation) or 2nd Gen.

You typically use the AppSync Service Plan In-Progress Monitor for long-running repurpose, protection activities (for example, VMAX Clone synchronization). When you select this monitor, it lists all the repurpose, protection activities in progress. Select any protection or repurpose activity, and then select **Monitor** to launch a progress dialog to directly view the progress of the selected activity.

If two datasets are protected, there are two separate entries in the AppSync Service Plan In-Progress Monitor. Click **Monitor** to fetch the respective events.

To launch the AppSync Service Plan In-Progress Monitor:

1. On the AppSync console select **Monitoring**.
2. Select **AppSync Service Plan In-Progress** on the Monitoring page.

CHAPTER 12

Storage considerations

- [VNX Block](#)264
- [VNX file](#)266
- [VNXe](#)..... 268
- [VMAX](#)270
- [VMAX 3 and VMAX All Flash](#) 274
- [ViPR Controller](#) 275
- [XtremIO](#)277
- [RecoverPoint](#) 280
- [Unity](#) 285
- [VPLEX](#).....287

VNX Block

AppSync supports the creation and management of application copies using VNX Snap copy technology. Consider best practices for VNX array setup before deploying AppSync.

Connectivity

Consider the following information for VNX connectivity with AppSync.

- AppSync supports Fibre, iSCSI, and FCoE connectivity between the host running the AppSync software and the VNX array. Network connectivity is required between the array and the AppSync server.
- Storage control occurs only on the AppSync server. No zoning is necessary from the VNX array to the AppSync server. It is not necessary to pre-expose any LUNs.
- Configurations with multiple AppSync servers per VNX array are supported.

VNX LUN support

Consider the following information for VNX LUN support.

- AppSync only supports LUNs that are in a pool. If you run VNX Snapshot replication you must use pool LUNs, not RAID LUNs. For RecoverPoint, use RAID storage.
- AppSync cannot create snapshots on LUNs with compression enabled.
- For VNX Snapshots, AppSync supports primary LUNs of any size.

VNX consistency groups

Consider the following information when using VNX consistency groups:

- All limitations that apply to VNX consistency groups also apply to AppSync.
- You can have multiple consistency groups within a single service plan.
- If you are using VNX consistency groups, all file systems that are related to an application in the snapshot set should reside in the same service plan. If not, you can encounter problems with mount and restore.

VNX mount

The following considerations apply:

- Mount hosts require SAN visibility to the VNX array.
- Changes made to a VNX Snapshot while it is mounted are persistent.
- If you accidentally delete hardware copies on the array, you cannot mount those copies with AppSync.
- Do not change the name of the storage group for the mount host when the copy is mounted. If you change the name of the storage group, ensure that you revert to the original name before unmounting the copy from AppSync.

Avoiding inadvertent overwrites

When you use AppSync to create a copy of one set of data that shares a LUN with other data, the copy contains all the data on that LUN. During restore, you may unintentionally write older data over newer data. The entities that are overwritten are called *affected entities*. Always configure data so that affected entities are reduced or eliminated.

Each LUN should contain a single file system or database file. If you are certain that the file system and/or database table residing on that LUN is always be backed up or restored as a unit, exceptions apply.

Service plan considerations for applications on VNX Block storage

After you register VNX storage, you can subscribe the application to a service plan to create and manage copies.

Bronze plans are supported. AppSync supports VNX Snapshot as the copy technology. Subscribe to the Bronze service plan to create and manage local copies for operation recovery, backup acceleration, or repurposing (create copies for test/dev). AppSync supports pooled LUNs (TLU/DLU) if the Snapshot technology supports these LUNs.

The maximum number of copies that AppSync can create and manage for VNX Block is dictated by the limits of the VNX Snap technology. The maximum number of VNX snaps per source is 256. This allows a maximum AppSync service plan rotation of 255.

Dynamic mounts

With proper zoning, AppSync automatically presents storage to the host when a copy is mounted.

Physical host

When AppSync mounts a copy, it dynamically assigns a snapshot to the host. The physical host must be zoned to the VNX array.

Virtual machine

Dynamic mounts happen as raw device mapping (RDM) or through native iSCSI on the virtual machine.

- For RDM, the ESX server where the virtual machine resides must be zoned to the VNX array.
- For RDM and virtual disks, virtual center (which manages the ESX server that the virtual machine mount host resides on) must be registered with the AppSync server.
- For native iSCSI, the virtual machine must be zoned to the VNX array.
- For virtual disks, virtual center of ESXi server (where mount host resides) must be registered with AppSync. Register the virtual machine with `disk.EnableUUID` flag enabled. AppSync installs the host plug-in on the virtual machine during registration for virtual disk and application level protection.

Microsoft Cluster Server mounts for SQL Server

Microsoft Cluster Server (MSCS) mounts for SQL Server can be done on production or alternate clusters.

When you mount to a cluster node using VNX storage:

- The storage group configuration applies only to physical hosts or virtual machines with NPIV or iSCSI that are directly connected to the VNX. For clusters configured using virtual machines with RDM or virtual disk, the copy that is mounted is only visible to the selected node (usually passive node).
- When you mount to a cluster node for backup purposes, create a dedicated storage group for one of the nodes of the cluster, preferably the passive node.

- If PowerPath 5.7 is installed, the host IP in the VNX storage group changes to the IP that corresponds to Microsoft failover cluster virtual adapter. PowerPath 5.7 has an auto-host registration feature that intercepts host agent operation and overwrites the IP with its own selection. This feature cannot be turned off. Upgrade to PowerPath 5.7.2 and above to correct this issue.

SAN policy on Windows Server Standard Edition

On Windows Server, the SAN policy determines whether a disk comes in online or offline when it is surfaced on the system. For Enterprise Edition systems, the default policy is offline. On Standard Edition the default policy is online. You need to set the policy to offlineshared to prevent mount failures.

To set the SAN policy to offline on a Windows Server 2008 Standard Edition host, open a command line window and run the following commands:

```
C:\>diskpart
Microsoft DiskPart version 6.0.6001
Copyright (C) 1999-2007 Microsoft
Corporation.
On computer: abcxyz
DISKPART> san policy=offlineshared
DiskPart successfully changed the SAN policy
for the current operating system.
```

VNX file

AppSync supports the creation and management of application copies using VNX File SnapSure copy technology. AppSync-managed copies can be local, remote (off the VNX Replicator target) or identical point-in-time local and remote copies.

Consider best practices for VNX file setup before deploying AppSync.

VNX SnapSure

VNX SnapSure creates a point-in-time copy of all the data on the network file system (NFS). For the initial snapshot, this method creates a full copy of the original file system, therefore requiring the same amount of space on the file system. Subsequent snapshots space usage depends on how much the data has changed since the last snapshot was taken.

SnapSure has the following characteristics:

- Storage Service — VNX File Server
- Source — VNX LUN
- Target — VNX SnapSure local snapshot
- Storage Requirements — The following storage requirements apply:
 - The source data must reside on VNX file systems.
 - Storage must include enough space for the snapshots on the VNX.
 - Storage pools cannot be defined for VNX jobs.
- Mount and Recovery — You can mount the replica on a target host and/or perform direct recovery from target to source.

VNX Replicator

VNX Replicator creates a point-in-time copy of all the data on the network file system (NFS). VNX Replicator maintains consistency between the source and target file systems that are based on the Time Out of Sync policy settings.

VNX Replicator has the following characteristics:

- Storage Service — VNX File Server
- Source — VNX NFS
- Target — Replicator remote snapshot
- Storage Requirements — The source data must reside on network file systems.
- Mount and Recovery — Can mount the copy on a target host and perform recovery from the copy if required.

VNX remote protection

Protection occurs between a local Data Mover and a Data Mover on a remote VNX system.

Both VNX for file cabinets must be configured to communicate with one another by using a common pass phrase, and both Data Movers must be configured to communicate with one another by using a Data Mover interconnect. After communication is established, a remote session can be set up to create and periodically update a source object at a remote destination site. The initial copy of the source file system can either be done over an IP network or by using the tape transport method.

Some recommendations for the session include:

- The session must be created with the Time Out of Sync update policy instead of a manual refresh.
- The Time Out of Sync value should be set to lowest value possible for the network configuration.

After the initial copy, changes made to the local source object are transferred to a remote destination object over the IP network. These transfers are automatic and are based on definable protection session properties and update policy.

One-to-many replication configurations are not supported in AppSync.

Protecting data on VNX network file systems

For service plans configured for remote protection, the NFS copy is created as a SnapSure Snapshot on the local and/or remote file system. Copies of NFS data stores can be created from service plans configured for local, remote, and local and remote protection.

During restore of an NFS copy, AppSync creates a roll back snapshot for every file system that has been restored. The name of each roll back snapshot can be found in the restore details.

You can manually delete the roll back snapshot after verifying the contents of the restore. Retaining these snapshots beyond their useful life can fill the VNX snap cache and cause resource issues.

Service plan considerations for an application on VNX File storage

Once you register VNX storage, you can subscribe the application to a service plan to create and manage copies.

Bronze, Silver, and Gold plans are supported for copies of applications (NFS data store, Oracle NFS) residing on VNX File.

The limits of VNX SnapSure technology determine the maximum number of copies that AppSync can create and manage for VNX File.

For Local SnapSure copies you can have a maximum of 96 RO (read-only) snaps. AppSync service plan rotation for VNX NFS file system is a maximum of 95.

For Remote SnapSure copies (across Remote Replicator), you can have a maximum of 95 RO snaps. AppSync service plan rotation for VNX NFS file system is a maximum of 94.

For RW (read/write) mounts, SnapSure allows for up to 16 RW snaps off existing RO snaps. A maximum of 16 snapshots for a given source can be mounted RW at any specified time. The service plans, by default, unmount the provision copy before mounting the new copy so this limit has no consequences. However, if the implementation requires simultaneously mounting multiple copies for the same source RW, the limit of 16 must be considered.

VNX file mount

You can mount any VNX File Snapshot copy created in the service plan at any time, independent of other copies created on the same service plan.

The following considerations apply:

- ESX mount hosts must belong to a vCenter server.
- ESX mount hosts require visibility on the network.
- NFS.MaxVolumes, an advanced setting on the ESX server, should be set to the number of NFS datastores that will be mounted to each ESX on the network.
- When mounting to an ESX server, AppSync uses the lowest number interface that has connectivity.

Note

For VNX File Snapshot mount, alias of export is not used, if the source file system is NFSv4.X. Instead, in such a scenario, the checkpoints are exported and mounted using the checkpoint's full name. This is because of a limitation from VNX file storage.

VNXe

Learn about VNXe and AppSync copy management, application support, and allowed storage and replication types.

VNXe arrays support all applications within AppSync. VNXe does not support the following configuration with AppSync:

- NFS file systems on UNIX platforms
- Repurposing copies on VNXe arrays

VNXe copy management

This section describes a typical AppSync workflow where you can create and manage application-consistent copies on VNXe storage. AppSync manages VNXe arrays with the Management Interface instead of the Service Processor interface.

Perform resource registration for VNXe when you start AppSync after installation. Register hosts as well as vCenter and storage systems so that AppSync can perform various operations that are required to create and manage copies of applications. Typically, registration of an entity includes identifying the system using name/IP address and providing the necessary credentials (username/password) for AppSync to discover and operate on the registered system.

Note

AppSync rediscovers file systems, interfaces, and file shares every hour. Therefore, any new file, file system, or file share added on the array is not identified by AppSync until the next refresh cycle.

Storage and replication type

AppSync supports Unified Snapshot replication technology to create and manage local copies of applications that reside on VNXe block or file storage.

Source block storage LUNs can be pool LUNs that are either thick or thin. You can provision the required source block devices using the VNXe LUN or VMware data store wizards within the Unisphere UI. When creating basic LUNs using the LUN wizard select one of the following options:

- Create a LUN – Creates an individual LUN from a desired storage pool.
- Create a LUN group – Creates a grouping of LUNs from a desired storage pool. The advantage of using a LUN group is that all LUNs within the group are snapped together guaranteeing consistency on the array.

Best practice states that in Microsoft environments you should use LUN groups in their storage layout to help the application consistent creation of VNXe snapshots within the Microsoft VSS Service time window.

You can provision the required source file devices using the VNXe file system or VMware data stores wizard. When provisioning from the file system wizard, only NFS share file systems are supported.

Service plan considerations with VNXe

Before you add a service plan to create VNXe copies, review these considerations.

After you register VNXe storage, subscribe to the Bronze service plan to create and manage local copies for operation recovery and backup acceleration. After you register VNXe storage, AppSync selects snap for Bronze plans by default.

For copies across RecoverPoint Continuous Remote Replication or CLR, subscribe to Silver or Gold service plan respectively. You can change snap to Bookmark for RecoverPoint copies.

Mount and unmount copy considerations for VNXe

Mount/unmount operations on VNXe arrays involve attaching/detaching snapshots to SMPs, and granting/removing (masking/unmasking) snapshot access to SMP LUNs or a set of SMPs to a host.

VNXe only allows one snapshot at a time of a set of LUNs or LUN group for attachment to SMP LUNs. LUNs contained in a LUN group are attached and detached together, there is no partial attach/detach.

Before performing a mount and unmount, zone the mount host to the VNXe array, and register the host name with its initiators.

The first step AppSync performs when mounting a snapshot on VNXe, is host initiator discovery for the mount host. Based on the snapshot information, AppSync maps to the appropriate source LUN/LUN group(s) to determine host access for the mount host. AppSync verifies that no other snapshots are attached to the SMP LUNs. Next AppSync modifies host access to either grant snapshot access to perform a mount, or remove host access to unmount.

For RDM or vDISK mount/unmount, AppSync identifies host access based on the host initiator for the ESX server.

LUN Groups have an all or nothing approach towards mounts. On Microsoft products, all devices within or outside of a LUN group need to be initially exposed to the mount host so that a VSS import can be performed correctly. Some devices can remain visible but not connected to a mount point on the mount host.

Mount/unmount VNXe NFS datastore considerations

Review the following information for NFS datastore mount/unmount.

AppSync creates a share based on the VNXe file snap that you want to mount. The share will be visible to the mount host. During unmount, the share created during the mount will be deleted. Share name appears in the following format:

AS-Share-

lastFourDigitOfVNXeSerialNum-ProductionFilesystemId-

TimeOfShareCreated

Regarding export IPs, AppSync creates a list of export IP interfaces from the VNXe array. Production export IP is a priority.

VMAX

To create and manage copies of your applications, AppSync supports TimeFinder Clone and TimeFinder VP Snap replication technology. AppSync also supports remote copy management off of an R2 in a SRDF/S or SRDF/A configuration.

Review the following sections before adding your VMAX storage.

Service plan considerations for applications on VMAX storage

Once you register VMAX storage you can subscribe your application to a service plan to create and manage copies.

Bronze and Silver plans are supported. TimeFinder VP Snap is the default replication technology used for service plans. You can change your preference to clone if TimeFinder Clone copies are desired.

The recommended maximum number of copies to keep before expiration is 6 for Timefinder Clones and up to 31 for VP Snap. The number of Timefinder Clone and VP Snap copies that can be created and managed is influenced by other copy and replication technologies used on the source LUNs. Refer to section on Copy Session Limits" for your planning. Refer to [VMAX copy session limits](#).

Bronze plan

You can subscribe to the Bronze service plan to create and manage local copies for operation recovery, backup acceleration or repurposing (create copies for test/dev).

For RAID LUNs AppSync chooses TimeFinder Clones for the Bronze plan. If the source is a RAID LUN or a mix of RAID and thin LUNs, then AppSync defaults to clone even if you select TimeFinder VP Snap as your preference.

Silver plan

For copies across SRDF/S or SRDF/A subscribe applications to the Silver service plan.

Note

Creation of remote copies in an SRDF/A configuration is not supported with Microsoft applications.

SRDF/A caveats: Creating a TimeFinder VP Snap or TimeFinder Clone of the R2 device is not allowed if either of the following is true:

- SRDF/A device-level write pacing is not activated and supported on the SRDF/A session.
- The SRDF pair is the R21-> R2 of a cascaded configuration, and any of the following apply:
 - The R21 Symmetrix array is running an Enginuity level lower than 5876.159.102
 - The R2 Symmetrix array is running an Enginuity level lower than 5875.
 - The R21 device is not pace-capable.
- Restore from SRDF/A is not supported

Source storage LUNs can be traditional RAID LUNs or thin LUNs (TDEVs). TimeFinder VP Snap support is only for thin LUNs. Consider the following recommendations:

- R1 > R2 should be in Synchronized state (for SRDF/S) and Consistent state (for SRDF/A)
 - For Silver plan Create copy: Affinitizer splits the applications based on the RDF Group (RA Group) to which the source devices belong.
 - Put all application LUNs in the same RDF group.
-

Note

One SMI-S provider that discovers the local array and the remote array is required when making copies of R2 devices. Refer to the *AppSync VMAX Array Support Guide* on <https://support.emc.com> for additional information.

Copy session limits

Symmetrix VMAX series arrays support up to 16 differential sessions per source device, which can be used for TimeFinder/Clone, TimeFinder/Snap, TimeFinder VP Snap, SRDF/Star, Solutions Enabler Open Replicator (ORS), or Symmetrix Differential Data Facility (SDDF) operations.

This limits the number of available copies that can be created.

TimeFinder VP Snap allows an additional 32 sessions per Symmetrix device which includes availability of one session of the traditional 16 sessions available. If you want to perform a restore, an additional session is required from the 16.

For example, if you use VP Snap for a source LUN and then desire a restore operation, this action leaves 14 sessions available for other copy technologies (TF Clone, ORS, Timefinder/Snap, SRDF/Star, on so on).

Additionally, if you want to create and manage TF Clone copies for the same source, you can create no more than 7 TF Clone copies using AppSync. (AppSync creates differential TimeFinder Clone copies which take up 2 differential sessions per copy $14/2 = 7$.) Since AppSync does not delete or expire a copy prior to creating a new one, the source of the AppSync rotation for the TimeFinder Clone copies can be no more than 6. This allows for an additional copy to be created prior to delete/expire of the oldest copy.

Note

For additional TimeFinder session limits, refer to *Solutions Enabler Symmetrix TimeFinder Family CLI Product Guide*.

Mount and unmount VMAX copies

Mount/unmount operations on VMAX involve masking/unmasking LUNs or set of LUNs to a host.

AppSync relies on the VMAX Auto-Provisioning capability. AppSync requires the mount host to be zoned to the VMAX array. You should create a masking view with the appropriate initiator group, port group and storage group.

When AppSync performs a mount operation on VMAX, it discovers the host initiator for the mount host first. Based on the host initiator, it maps to the appropriate masking view to determine the Storage Group to or from which the target LUNs are masked/unmasked to perform a mount/unmount operation.

You can select the desired FAST VP policy for the target LUN in the mount phase of the service plan. If there is a storage group for the mount host with the desired FAST VP, AppSync adds the LUN to the storage group. If this storage group does not exist, AppSync adds the LUN to any storage group that is masked to the host.

If a storage group is configured to pick target devices, AppSync removes the devices from the storage group at the time of mount and adds them to the storage group for the mount host. The devices are added to the original storage group when the copy is expired.

Note

- When you select FAST VP policy, ensure that the storage pool of the storage group (FAST VP policy's storage group) and the storage pool of the copy devices are of the same storage pool type (that is, they must be on the same storage tiers). If the copy devices and FAST VP policy storage pools are on different storage tiers, the copy devices cannot be moved between different storage tiers and mount operation fails. For example, If a VMAX V2 source device is created on a Flash Drive Pool and a TimeFinder Snap is also created from the same pool in AppSync, to mount the copy to a desired FAST VP policy, ensure that you select a FAST VP policy associated with the Flash Drive Pool because LUNs cannot be moved from one tier to another having pools with mismatched disk drives.
- To use the FAST VP policy feature after an upgrade from AppSync 2.2.3 or earlier versions to 3.0, rediscover the array.
- Do not change the name of the storage group for the mount host when the copy is mounted. If you change the name of the storage group, ensure that you revert to the original name before unmounting the copy from AppSync.

For RDM or Vdisk mount/unmounts, AppSync identifies the Masking view based on the host initiator for the ESX server.

Microsoft Cluster Server mounts for SQL Server

Microsoft Cluster Server (MSCS) mounts for SQL Server can be done on production or alternate clusters.

When you mount to a cluster node using VMAX storage:

- The storage group configuration applies only to physical hosts or virtual machines with iSCSI that are directly connected to the VMAX. For clusters configured using virtual machines with RDM or virtual disk, the copy mounted is only visible to the selected node (usually the passive node).
- When you mount to a cluster node for backup purposes, create a dedicated storage group for one of the nodes of the cluster, preferably the passive node.
- AppSync does not support mount to a cluster as a clustered resource. To mount to a cluster as a clustered resource (in a physical/iSCSI environment), deselect the default setting **Use dedicated storage group** on the AppSync mount dialog . During mount AppSync will make the copy visible to multiple nodes in the cluster by using a Storage Group with multiple assigned nodes. AppSync also mounts the filesystems to the selected mount host. After the completion of AppSync mount, manually add the mounted devices under cluster management to avoid the possibility of any data corruption.

Repurpose copies on VMAX

Consider this information when repurposing Oracle and SQL Server database copies residing on a VMAX.

You can repurpose a VMAX source copy where the database resides or the source can be the target device in an SRDF session.

You can only repurpose a 1st Gen copy from the source, or a 2nd Gen copy (copy of the copy).

You can repurpose a RecoverPoint bookmark copy of an Oracle or SQL Server database.

When considering repurposing, review the following information:

- VMAX does not support a mix of thick and thin devices in cascading.
- A first generation copy must be a clone. The second generation copy can be a clone or a TimeFinder VP Snap.
- If the source device is thick, then the first and second generation targets are thick. If the source device is thick, AppSync only supports a clone of a clone for the source device.
- If the source is thin, AppSync supports both a TimeFinder VPSnap of clone and a clone of a clone.
- The first generation copy can be a remote copy or local copy in an SRDF session. But the second generation is local only.
- If you have thick source devices, configure the storage group with thick devices, otherwise the first generation copy creation fails.

VMAX restore

VMAX restricts the maximum number of hops in cascading to two.

If source A has the following sessions such as A > D and A > B > C (when created by a service plan or using Repurposing workflow), then during the restore from D the number of hops changes to 3 as the restore session leads to D > A > B > C. Therefore, Appsync provides an option to terminate the session B > C if it is a clone.

During consecutive runs if AppSync chooses C as a target for B, then it will be a full sync instead of a resync since the session B > C will be terminated during restore.

For example:

A > B > C

A > D > E

Restore from B will terminate D > E. Restore from D will terminate B > C. Necessary sessions will be terminated only if the you select the option.

If the second generation copy is a TimeFinder VP Snap, you must expire second generation snaps manually. The restore fails if a snap of clone exists for source and displays all the copies that need to be expired in the progress window.

Note

Refer to the *AppSync VMAX Array Support Guide* on <https://support.emc.com> for additional information.

VMAX 3 and VMAX All Flash

VMAX 3 and VMAX All Flash arrays are supported with AppSync. This section describes supported features, and service plan considerations including mount/unmount and restore of VMAX 3 and VMAX All Flash copies.

To create and manage copies of applications, AppSync supports SnapVX snapshot replication technology in VMAX 3 and VMAX All Flash arrays.

Service plan considerations for applications on VMAX 3 and VMAX All Flash storage

Review these considerations for service plan support with VMAX V3 and VMAX All Flash.

Overview

After you register (add) VMAX 3 and VMAX All Flash storage, subscribe an application to a service plan to create and manage copies.

Bronze, Silver, and Gold service plans are supported with VMAX 3 and VMAX All Flash. Snap in a service plan with VMAX 3 and VMAX All Flash is equivalent to a SnapVX snapshot linked in `no copy` mode. Clone in a service plan is equivalent to SnapVX snapshot linked in `copy` mode. The default preference for a service plan is Snap.

SnapVXClone optimization

The SnapVX clone is created and linked with the target device during copy creation. AppSync optimizes the SnapVX clone linking using the relinking strategy. While linking the SnapVX clone with the target device, AppSync attempts to find a suitable target device, which is already linked with a previous SnapVX of the same source device, and can be used for relinking. If you relink the new SnapVX with the same target device, it only synchronizes the delta data, that is the data that changed from the earlier SnapVX to the new SnapVX. This minimizes the linking time for larger devices. The non-repurposing service plan employs the relinking feature during copy creation. The repurposing service plans can refresh the repurposed copies and use the relink feature.

Mount/unmount VMAX 3 and VMAX All Flash copies

Mount/unmount operations on VMAX 3 and VMAX All Flash include masking/unmasking LUNs or a set of LUNs to a host. AppSync relies on the VMAX 3 and VMAX All Flash Auto-Provisioning capability.

The mount host must be zoned to the VMAX 3/VMAX All Flash array. Next, you can create a masking view with the initiator group, port group, and storage group.

When AppSync performs a mount operation on VMAX 3 or VMAX All Flash, it discovers the host initiator for the mount host first, then based on this host initiator, AppSync maps to (or from) the masking view. This operation determines the storage group where the target LUNs are masked/unmasked. For RDM or Vdisk mount/unmount, AppSync identifies the masking view that is based on the host initiator for the ESX server.

You can select the wanted Service Level Objective (SLO) for the target LUN in the mount phase of the service plan. If there is a storage group for the mount host with the wanted SLO, AppSync adds the LUN to the storage group. If this storage group does not exist, AppSync adds the LUN to any storage group that is masked to the host.

If a storage group is configured to pick target devices, AppSync removes the devices from the storage group at the time of mount and adds them to the storage group for the mount host. The devices are added to the original storage group when the copy is expired.

Note

Do not change the name of the storage group for the mount host when the copy is mounted. If you change the name of the storage group, ensure that you revert to the original name before unmounting the copy from AppSync.

VMAX 3 and VMAX All Flash repurpose overview

Review VMAX 3 and VMAX All Flash support for Repurposing.

AppSync supports local and remote repurposing of VMAX 3 and VMAX All Flash SnapVX copies. Refresh of an existing copy will create a new snapshot of the source LUN and link it to the original target of the copy in the required mode. The old snapshot is then expired.

Use the Repurpose wizard to create a local or remote repurposed copy with VMAX 3 or VMAX All Flash.

Note

Refer to the *AppSync VMAX Array Support Guide* on <https://support.emc.com> for additional information.

ViPR Controller

Learn about support for ViPR Controller with AppSync, including configuration considerations, supported service plans and applications.

With this release of AppSync, ViPR Controller operates as another storage system. This means that, ViPR Controller needs to be registered and discovered into AppSync. Refer to the section on adding ViPR Controller to AppSync.

AppSync supports ViPR Controller snapshots as copy technology for applications provisioned using ViPR Controller. .

Note

An application with a mix of LUNs for protection is not supported. For example, AppSync does not support a mix of native (AppSync-based) LUNs on VMAX and ViPR Controller LUNs.

AppSync only supports applications that are provisioned by ViPR Controller with block virtual pools backed by VMAX/VPLEX with VMAX and XtremIO storage. AppSync does not support applications provisioned on file and object virtual pools.

Note

AppSync does not support ISCSI with AIX for either XtremIO native storage or XtremIO on ViPR Controller.

Migration from native array to ViPR controller

AppSync supports ViPR Controller copy management functionality such that migration of applications from native to ViPR platform occurs without any disruption to existing copies and service plans and without manual intervention. For example:

1. Assume you have an Oracle database running on a VMAX, and you are using AppSync to protect this database using a Bronze service plan for local protection.
2. AppSync creates a daily snapshot of the database by directly interacting with VMAX storage natively using an SMI-S provider.
3. At some time in the future, you decide to use a ViPR Controller platform to automate your storage infrastructure. As part of this automation, you move all existing volumes into ViPR Controller. ViPR Controller requires all storage management tasks such as volume provisioning for snapshot creation, deletion and so on to be performed using ViPR Controller.
4. You still need local protection of your Oracle database on VMAX storage but you want to migrate to ViPR Controller for all snapshot operations.
5. With AppSync you can seamlessly migrate to the ViPR Controller platform without any manual intervention or modification to your existing plans.

Service plan considerations with ViPR Controller

Before you add a service plan for ViPR Controller copies, review these considerations.

Copy management functionality for ViPR Controller follows the same workflow as AppSync for native storage array support. For example, you must configure a host, a ViPR Controller resource inside AppSync, and then perform discovery of the application. Once discovered, you can subscribe the application to an existing service plan within AppSync. For this release, Bronze service plans are supported.

Mount and unmount ViPR Controller copies

Review this information regarding mounting and unmounting ViPR Controller copies with AppSync.

AppSync allows mounting of application copies created using ViPR snapshot to hosts registered with ViPR Controller. AppSync makes use of ViPR export groups while mounting copies of an application.

As part of mount, AppSync determines if a mount host already has a provisioned export group in the project and ViPR array where snapshots are stored. If AppSync finds an existing export group, it updates the export group by adding snapshots to the export group. If AppSync cannot find an export group for the mount host and ViPR array where the snapshot is stored, AppSync makes use of ViPR APIs to provision a new export group for the host in the required project for the ViPR arrays where the snapshots are located.

If the applications resides on VPLEX virtual volumes, the mount host must be zoned to the back-end native array on which the VPLEX virtual volumes reside. This is required

because AppSync supports exporting VPLEX volume snapshots only as native array volumes and not as VPLEX virtual volumes.

Automatic zoning troubleshooting tip: When a Fabric Manager account has read-only access and a ViPR array has SAN zoning set to **Automatic > Create Export Group**, this action fails if the host is not zoned to the array. Automatic zoning cannot be performed since the Fabric Manager account has read-only access to fabric. To work around this issue, set SAN Zoning to **Manual** and then manually create a zone between the host and physical array.

ViPR Controller copy restore

Review these scenarios when considering a restore of a ViPR Controller copy.

AppSync does not support restore from application copies which are created using ViPR Controller where the volume was provisioned by ViPR Controller on an XtremIO 3.x back-end array (either natively or by using VPLEX in front-end).

LUN level restore from ViPR snapshot

While restoring from a ViPR Controller snapshot, AppSync checks for an affected entity in case of consistency groups and warns you about other applications which may become impacted due to the restore.

Granular Restore

AppSync supports granular restore for virtual machines and file with ViPR Controller as storage platform. Since ViPR Controller does not have capability to create a copy of a copy, granular restore for virtual machines always uses cloning of virtual machines to recover the virtual machines from a replica.

Instant restore functionality is disabled while using ViPR Controller.

XtremIO

Review the supported applications, replication technology, configuration requirements, and restrictions for XtremIO arrays with AppSync before you begin the installation.

Application support

AppSync creates write-consistent snapshots on the XtremIO array for each application you add to a service plan. XtremIO supports the following applications:

- Oracles databases
- SQL Server databases
- Exchange databases and DAG
- File systems
- VMware data stores
- RecoverPoint (4.1.2 minimum version) is supported with XtremIO 4.0 and later and AppSync 2.2.2 and later.

Replication technology

AppSync creates XtremIO Snapshots of the application data and places the snapshots in the following common folder: `/AppSyncSnapshots/APPSYNCSERVERNAME/PRODUCTIONHOSTNAME/`. Move XtremIO Snapshots from the common folder only if you use the AppSync copy Remove feature, which removes a copy from AppSync, but does not remove the XtremIO snapshots.

AppSync names snapshots by using the following naming convention: `TIMESTAMP.snap.ORIGINALVOLNAME`. Due to the XtremIO limit of 64 characters in a snapshot name, the `ORIGINALVOLNAME` should be 40 characters or less.

Review XtremIO limits regarding the number of LUNs that can be snapped at one time. Do not let Oracle database configurations exceed the number of volumes limited by maximum XtremIO snapshot capabilities.

You can adjust the maximum number of outstanding disk requests with the `Disk.SchedNumReqOutstanding` parameter. The default value for `Disk.SchedNumReqOutstanding` is 32. XtremIO recommends this to be 256 for any device presented to a host. For instructions on setting the maximum number of outstanding disk requests for XtremIO snapshot devices mounted on a host, see the relevant VMware knowledge base article.

AppSync supports the use of XtremIO consistency group APIs to create and refresh snapshots, allowing for the fastest possible operation time. However, support is limited to the repurposing workflow. The minimum XtremIO version required is 4.0.2.

To use XtremIO consistency group APIs, the following conditions must be met:

- All source LUNs must have the same consistency group
- All source LUNs must be part of only one consistency group
- All snapshots must be part of a single Snapshot-set (consistency group level refresh or restore)

Note

A single source LUN in a consistency group is also supported.

For Windows based applications, all the LUNs must be in one consistency group.

For Oracle, all archive log LUNs must be part of one consistency group and the database LUNs must be part of one consistency group.

Note

- XtremIO volumes are always deleted at a rate of 1 volume per every 60 seconds. Therefore, the expiration of an application copy spread across 5 volumes takes 5 minutes to delete.
 - When using the XtremIO consistency group APIs, leftover volumes as part of the operation are also deleted at a rate of 1 volume per every 60 seconds.
 - For information on XtremIO limitations, particularly limitations regarding XtremIO Snapshots and folders, review the *XtremIO Release Notes* available on <https://support.emc.com>.
-

Restrictions

Consider the following restrictions for XtremIO with AppSync:

- XtremIO Initiator Groups must be defined in XtremIO for all mount hosts to which AppSync mounts XtremIO copies.
- AppSync does not support XtremIO with iSCSI connectivity for AIX hosts.

Configuration considerations

- The XtremIO Management Server (XMS) should be configured on a SAN with at least one XtremIO array.

- Zone XtremIO arrays to production and mount hosts (physical) or ESX servers (virtual).
- For mount and unmount of copies:
 - Ensure that you configure Oracle or SQL Server databases on XtremIO arrays for data and logs.
 - Fibre Channel and iSCSI are supported.

Considerations before adding an array:

To add and configure an XtremIO array to work with AppSync, you need at least one XtremIO Management Server (XMS) configured for that XtremIO array. Review the following considerations before adding an array:

- Administrator privileges are required to add the XtremIO array.
- Ensure XtremIO storage is zoned to production hosts (physical) or ESX servers (virtual). RDM and virtual disk are supported on VMware virtual machines. iSCSI is supported for Windows and Linux hosts, allowing you to see XtremIO storage over an iSCSI LAN. iSCSI is supported for physical or virtual hosts, and also ESX servers.
- Oracle, file systems and VMFS data stores on Linux/AIX are supported. File systems and virtual disks are supported on Windows.
- You need the XMS name/IP address and credentials.

Note

If you change the array credentials, ensure that you update the same in AppSync before attempting any operation.

The *AppSync Installation and Configuration Guide* provides instructions to add an XtremIO array.

Restore options with XtremIO storage

Learn about restore options for application copies on XtremIO arrays when planning the installation.

AppSync 2.2.2 and later supports automated restore of XtremIO 4.0 and later copies. The following applications are supported:

- SQL Server databases
- Exchange standalone databases and Exchange Data Availability Groups (DAG)
- VMware data stores
- File systems
- RecoverPoint

AppSync uses the Restore wizard for automated restore on XtremIO storage. Click the Restore button to launch Restore wizard for respective applications. During restore, AppSync creates another XtremIO-generated snapshot, stored under the tag `/volumes/AppSyncSnapshots/RestoredSnapshots`. **An Administrator must clean up these snapshots manually.**

RecoverPoint

Consider best practices for RecoverPoint setup before deploying AppSync. For example, be sure to observe RecoverPoint consistency group granularity best practices.

Service plan considerations for applications with RecoverPoint protection

AppSync supports different RecoverPoint replication options.

Three types of replication options:

Local (Continuous Data Protection)

In Local protection, RecoverPoint replicates to a storage array at the same site. In a RecoverPoint installation that is used exclusively for local protection, you install RPAs at only one site and do not specify a WAN interface. The Bronze service plan protects application replication.

Remote (Continuous Remote Replication)

In Remote replication, RecoverPoint replicates over a WAN to a remote site. There is no limit to the replication distance. The Silver service plan protects application replication.

Local and Remote (Concurrent Local and Remote)

In Local and Remote replication, RecoverPoint protects production LUNs locally using local protection and remotely using remote replication. Both copies have different protection windows and RPO policies. The Gold service plan protects application replication. RecoverPoint multi-site (multiple remote sites) is not supported at this time.

Note

For RecoverPoint bookmarks:

- Source VNX volume, target VMAX volume—virtual and virtual with roll access modes are not supported.
 - Source VMAX volume, target VNX volume—virtual and virtual with roll access modes are supported.
-

RecoverPoint prerequisites

Verify that the RecoverPoint configuration meets the prerequisites necessary for use with AppSync.

- Install and configure RecoverPoint according to the RecoverPoint documentation.
- Use RecoverPoint to create consistency groups.
- Ensure that the splitters for all mount hosts are attached to the RecoverPoint target volumes they are going to use.
- Synchronize time on all systems. Follow the steps in the operating system documentation to configure the AppSync server and all production and mount hosts to be synchronized with a time server. This includes all hosts, VNX, VNXe, Unity storage, and RecoverPoint appliances.

- For failover preparation, keep in mind that AppSync requires that RecoverPoint Local and Remote consistency groups have both local and remote copies, even in a failover situation. This may require RecoverPoint administrator configuration steps after failover to configure a local copy on the remote site.
- During AppSync configuration, the RecoverPoint site is added as a resource. In a Local and Remote configuration, AppSync discovers all sites in the RecoverPoint system configuration. Credentials for an account that has RecoverPoint admin privileges is required when adding the site.

Dynamic or static mounts

RecoverPoint copies can be mounted in two ways, statically or dynamically.

AppSync supported static mounts of RecoverPoint targets. Using static mounts, the RecoverPoint target LUNs (Local or Remote) had to be pre-exposed (masked) to the mount host before you could mount the RecoverPoint copies. If you are using static mounts in a virtual machine environment, the RecoverPoint target LUNs must be masked to the ESX server, and added as RDMS to the virtual machines prior to mounting the copy.

RecoverPoint targets may also be dynamically mounted. RecoverPoint target LUNs are mapped at mount time to identify the LUNs, and the LUNs are masked (moved to the mount host storage group) and surfaced prior to mounting. When the target LUNs for dynamic mount are on VNX storage, the VNX must be registered with AppSync. This is also applicable for VNXe and Unity storage.

AppSync does not have a prerequisite that replica devices must be made visible to the mount host. AppSync can dynamically expose devices across all storage technologies. For VMAX2 and VMAX3, AppSync does not support static mounts. This is also applicable for RecoverPoint environments involving VMAX.

For VMAX dynamic mounts, follow the VMAX auto provisioning instructions so that masking succeeds.

If you are using dynamic mounts in a virtual environment, do not mask the target LUNs to the ESX server. AppSync will mask the LUN to the ESX server, and then add the LUN as an RDM to the mount host. Refer to [Mount and unmount VMAX copies](#).

When unmounting:

- LUNs which were dynamically mounted are dynamically unmounted, that is, the LUNs are removed from the storage group.
- LUNs which were statically mounted remain in the storage group after the unmount completes.
- For application copies with LUNs that are mixed (both statically and dynamically mounted), the LUNs will be dynamically unmounted. All mounted LUNs are removed from the storage group.

Given proper zoning, AppSync presents storage to the host automatically when a copy is mounted.

Physical host

AppSync dynamically assigns a snapshot to the host when the copy is mounted. The physical host must be zoned to the VNX, VNXe, or the Unity array of the RecoverPoint target LUNs (Local or Remote).

Virtual machine

Dynamic mounts happen as a raw device mapping (RDM) or through native iSCSI on the VM.

- For RDM, the ESX server where the VM resides must be zoned to the VNX, VNXe, or the Unity array of the RecoverPoint target LUNs (Local or Remote).
- For RDM and virtual disks, virtual center (which manages the ESX server where the VM mount host resides) must be registered with the AppSync server.
- For native iSCSI, the virtual machine must be logged into the array (VNX, VNXe, or Unity) initiators of the RecoverPoint target LUNs (Local or Remote). The VNX must have a storage group defined for the host.

Repurpose RecoverPoint Bookmark copies of Oracle or SQL Server databases

AppSync supports the ability to repurpose RecoverPoint Bookmark copies for Oracle or SQL Server databases.

Use AppSync to repurpose a RecoverPoint Bookmark on a VMAX target (VMAX 3 is not supported) and create a first generation (1st Gen) copy, which leverages TimeFinder Clone or TimeFinder VPSnap replication technology. You can repurpose the clone copy further (not for VPSnap) to create a 2nd Gen copy that leverages TimeFinder Clone or TimeFinder Clone VPSnap.

- **Bookmark (hidden) > Clone**
- **Bookmark (hidden) > Clone > Snap**
- **Bookmark (hidden) > VPSnap**
- **Bookmark (hidden) > Clone > Clone**

To copy Bookmarks, use the RecoverPoint repurpose wizard. The RecoverPoint Appliance and SMI-S provider or VNX must be registered in AppSync.

Supported configurations include:

- Application: Oracle and SQL Server
- Storage: VNX and VMAX (1st Gen copy is a VMAX copy of
- Bookmark: The 2nd Gen copy is a copy of the 1st Gen copy).

In the repurpose wizard, select **Use Bookmark as an intermediate step** to perform RecoverPoint repurposing. If you do not select this option, AppSync begins native repurposing. The drop-down list lists **create a 1st Gen copy from site**. This option determines if the system uses RecoverPoint Continuous Data Protection or RecoverPoint Continuous Remote Replication Bookmark repurposing.

Considerations

- If you refresh the 1st Gen copy, AppSync takes a new copy of the database.
- 1st Gen and 2nd Gen copies are always local.
- Manual expire of 1st Gen the Bookmark copy.
- Refresh a 2nd Gen to create a copy of the 2nd Gen from the 1st Gen.
- If the 1st Gen copy is a VMAX clone, the same LUN is used during refresh (instead of rotation).

Repurpose (create) a 1st Gen copy of a RecoverPoint Bookmark

Learn to repurpose a RecoverPoint Bookmark on a VMAX target and create a 1st Gen (clone) copy of an Oracle or SQL Server database, which leverages TimeFinder Clone or TimeFinder VPSnap replication technology. You can repurpose the clone copy further (not for VPSnap) to create a second generation copy that leverages TimeFinder Clone or TimeFinder Clone VPSnap.

Before you begin

The RecoverPoint Appliance and SMI-S provider or VNX must be registered in AppSync.

To copy Bookmarks, use the RecoverPoint repurpose wizard.

Procedure

1. Log in to the AppSync console and select **Copy Management**.
2. In the repurpose wizard, select **Use Bookmark as an intermediate step** to perform RecoverPoint repurposing.

Recovery using RecoverPoint system created Bookmark

For a Recover point copy, you can specify a time which corresponds to the RecoverPoint system created bookmark (any point in time copy) during recovery (mount and/or restore) operations. However, there must be a bookmark copy available for the application in AppSync before you specify the time. The resultant copies are crash-consistent. Therefore, application recovery might not be successful.

To mount an any point in time (APIT) copy:

Procedure

1. In the Mount wizard, select a bookmark copy.
2. In the **Select a copy** screen, select **Select a point in time** to specify a time which corresponds to the RecoverPoint system created bookmark you want to mount.

Note

- The time shown is the console's time. If the console is in a different time zone from the RPA, specify the time as per the AppSync server's time zone.
- You must specify a time which is later than the oldest available Bookmark copy for that application in AppSync server.

3. Specify the time, and click **Next**.
4. Select the mount settings, and click **Next**.
5. Click **Finish**.

After the mount operation completes, a new APIT bookmark copy with the specified timestamp is listed in AppSync.

Note

- In the case of a copy created using the Gold service plan (local and remote), you can specify the location where you want to mount the copy. It can either be a local or a remote site.
 - If there is no RecoverPoint system created bookmark available for the time specified, AppSync picks the nearest bookmark available (below the specified time) in RecoverPoint.
 - The APIT bookmark copy created during mount can be re-used, if you want to mount the same RecoverPoint system created bookmark again (all applications supported by AppSync).
-

Restore from an APIT copy

Procedure

1. In the Restore wizard, select a bookmark copy.
 2. In the **Select a copy** screen, select **Select a point in time** to specify a time which corresponds to the RecoverPoint system created bookmark you want to mount.
-

Note

- The time shown is the console's time. If the console is in a different time zone from the RPA, specify the time as per the AppSync server's time zone.
 - You must specify a time which is later than the oldest available Bookmark copy for that application in AppSync server.
-

3. Select the restore options, and click **Next**.
 4. Click **Finish**.
-

Note

- In the case of a copy created using the Gold service plan (local and remote), you can specify the location where you want to restore the copy. It can either be a local or a remote site.
 - If there is no RecoverPoint system created bookmark available for the time specified, AppSync picks the nearest bookmark available (below the specified time) in RecoverPoint.
 - Unlike APIT mount, no APIT bookmark copy is created at the end of the restore operation.
 - APIT restore operation is only supported for VMware datastores and file systems.
 - Restoring from a mounted copy is only supported for VMware datastores.
-

Unity

This section describes Unity support with AppSync. It includes information on configuration considerations, supported service plan and application details.

Unity arrays support all applications within AppSync.

AppSync does not support the following configuration with Unity:

- Unity File storage on Windows platform
- Remote copies of unified snapshots

Unity copy management

This section describes a typical AppSync workflow where you can create and manage application-consistent copies on Unity storage. AppSync manages Unity arrays with the Management Interface instead of the Service Processor interface.

Perform resource registration for Unity when you start AppSync after installation. Register hosts as well as vCenter and storage systems so that AppSync can perform various operations that are required to create and manage copies of applications. Typically, registration of an entity includes identifying the system using name/IP address and providing the necessary credentials (username/password) for AppSync to discover and operate on the registered system.

Note

AppSync rediscovers file systems, interfaces, and file shares every hour. Therefore, any new file, file system, or file share added on the array is not identified by AppSync until the next refresh cycle.

Storage and replication type

AppSync supports Unified Snapshot replication technology to create and manage local copies of applications that reside on Unity block or file storage.

Source block storage LUNs can be pool LUNs that are either thick or thin. You can provision the required source block devices using the Unity LUN or VMware data store wizards within the Unisphere UI. When creating basic LUNs using the LUN wizard select one of the following options:

- Create a LUN - Creates an individual LUN from a desired storage pool.
- Create a Consistency group - Creates a grouping of LUNs from a desired storage pool. The advantage of using a consistency group is that all LUNs within the group are snapped together guaranteeing consistency on the array.

Best practice states that in Microsoft environments you must use consistency groups in their storage layout to help the application consistent creation of Unity snapshots within the Microsoft VSS Service time window.

You can provision the required source file devices using the Unity file system or VMware data stores wizard. When provisioning from the file system wizard, only NFS share file systems are supported.

Service plan considerations with Unity

Before you add a service plan to create Unity copies, review these considerations.

After you register Unity storage, you can subscribe to the Bronze service plan to create and manage local copies for operation recovery and backup acceleration. After you register Unity storage, AppSync selects snap for Bronze plans by default.

For copies across RecoverPoint Remote replication or Local and Remote replication, subscribe to Silver or Gold service plan respectively. You can change snap to Bookmark for RecoverPoint copies.

Mounting and unmounting Unity NFS datastore copies

Review the following information for NFS datastore mount/unmount.

AppSync creates a share based on the Unity file snap that you want to mount. The share is visible to the mount host. During unmount, the share created during the mount is deleted. The share name appears in the following format:

```
AS-Share-
lastFourDigitOfUnitySerialNum-ProductionFilesystemId-
TimeOfShareCreated
```

AppSync creates a list of export IP interfaces from the Unity array. Production export IP is a priority.

Mounting and unmounting Unity copies

Review this information before you mount and unmount Unity copies with AppSync.

Mount and unmount operations on Unity arrays involve attaching and detaching snapshots to SMPs, and granting and removing (masking/unmasking) snapshot access to SMP LUNs or a set of SMPs to a host. Unity only allows one snapshot at a time for a set of LUNs or consistency group for attachment to SMP LUNs. LUNs contained in a consistency group are attached and detached together, there is no partial attach or detach.

Before performing a mount or unmount, zone the mount host to the Unity array, and register the host name with its initiators.

The first step AppSync performs when mounting a snapshot on Unity, is host initiator discovery for the mount host. Based on the snapshot information, AppSync maps to the appropriate source LUN/consistency group(s) to determine host access for the mount host.

AppSync verifies that no other snapshots are attached to the SMP LUNs. Next AppSync modifies host access to either grant snapshot access to perform a mount, or remove host access to unmount.

For RDM or vDISK mount/unmount, AppSync identifies host access based on the host initiator for the ESX server.

Mounting and unmounting Unity NFS File system copies

Review the following information for NFS File system mount and unmount.

AppSync creates a share based on the Unity unified snapshot for file that you want to mount. The share is made visible to the mount host by exporting the NFS file system's unified snapshots, and the file system is created on that NFS exports. During unmount, the file system is unmounted and the NFS share created during mount is deleted. The share name appears in the following format:

```
AS-SharelastFourDigitOfUnitySerialNum-ProductionFilesystemId-
TimeOfShareCreated
```

AppSync creates a list of export IP interfaces from the Unity array. Production export IP is a priority.

Oracle database on Unity NFS file system is supported. The unified snapshots for file is created for Oracle data and logs. During mount of an Oracle database, use one of the following options:

- Mount and Recovery - This option mounts the file system on the mount host and recovers the database.
- Mount the file system - This option only mounts the file system on the mount host.

Repurposing copies on Unity

AppSync supports repurposing application copies on Unity arrays. Both snapshot and thin clone copy technologies of Unity are supported. The minimum Unity version required is 4.2.

[Repurposing overview](#) provides additional information on repurposing.

- The first generation copy is always a snapshot.
- The second generation copy can either be a snapshot or a thin clone. By default, the second generation copy is a snapshot. You can change the default option.

For additional information on Unity snapshot and thin clone copy technologies, see Unity documentation.

VPLEX

AppSync can create application consistent and crash consistent Snapshot (VPLEX Snap and VPLEX Clone) copies on the underlying managed array hosting VPLEX virtual volumes. AppSync supports the following applications on VPLEX storage:

- Oracle databases
- SQL Server databases
- File systems
- Microsoft Exchange
- VMware data stores

The following VPLEX device configurations are supported:

- VPLEX Local
 - RAID 0
 - RAID 1
- VPLEX Metro
 - Distributed devices

Note

- VPLEX virtual volumes must be mapped 1:1 to an array volume.
 - Concatenated devices (RAID-C) are not supported.
 - Nested devices are not supported.
 - Remote volumes (local device with global visibility by setting remote access) is not supported.
 - If there is a mobility job in progress, the device cannot be protected until the mobility job completes.
-

Service plan considerations for applications on VPLEX storage

After you register VPLEX storage, you can subscribe your application to a service plan to create and manage copies.

- AppSync supports the Bronze service plan for applications on VPLEX storage. This means that you can only create application specific local copies.
- When you select the storage preference as Snapshot in a service plan, AppSync creates a snapshot on the back-end storage array.
- During mapping, AppSync queries VPLEX about the virtual volume details such as device components, extents, and storage volume. It then communicates to the back-end storage array and maps the corresponding storage LUN to be protected.
- If applications on the same hosts are from different VPLEX clusters, the applications are grouped separately during protection.
- If applications are on the same host and on the same VPLEX virtual volumes, they are grouped together during protection.
- If the underlying VPLEX storage device is a RAID1, or a distributed device, you must configure storage options in the create copy phase under service plan settings.
- A VPLEX virtual volume on a RAID 1 device has two legs. The two legs can be from two different back-end storage arrays. The leg to be protected is determined by the array that is selected in configure storage options. A VPLEX distributed virtual volume has storage devices on both the clusters. The leg to be protected is determined by the cluster that is selected in configure storage options.
- In the case of RAID 0 devices, the VPLEX back-end array's storage LUN which maps to the VPLEX virtual volume is protected.
- In the case of RAID-1 devices, the VPLEX back-end array's storage LUN which maps to the selected leg of the RAID-1 device is protected.

Mount and unmount VPLEX copies

Review this information before you mount and unmount VPLEX copies with AppSync.

AppSync provides two mount options:

- Mount as VPLEX virtual volumes
- Mount as native array volumes

Mount as VPLEX virtual volumes

When you select this option, the snapshot on the back-end array is made visible to VPLEX. AppSync creates VPLEX virtual volumes on the underlying native array snapshots and provisions these virtual volumes to the mount host. The provisioned virtual volumes are added to the storage view of the mount host. During unmount, the virtual volumes are removed from the storage view of the mount host, and it tears down the created VPLEX virtual volume on the underlying native array snapshots. The snapshot on the back-end array is de-provisioned from VPLEX.

Consider the following:

- The mount host must be zoned to the VPLEX cluster where the production copy is created.
- The mount host must be zoned to VPLEX, but does not have to be zoned to the native array where the snapshot is created.

- Copy of production volumes on VPLEX RAID 0 devices are mounted as local RAID 0 volumes on the same cluster.
- Copy of production volumes on VPLEX RAID 1 devices are mounted as local RAID 1 devices with a single leg on the same cluster. If you manually add a mirror leg, ensure that you manually remove that leg before unmount.
- Copy of production volumes on VPLEX distributed RAID 0 devices are mounted as local RAID 0 volumes on the same cluster.
- Copy of production volumes on VPLEX distributed RAID 1 devices are mounted as local RAID 1 volumes on the same cluster. If you manually add a mirror leg, ensure that you manually remove that leg before unmount.

Mount as native array volumes

When you select this option, AppSync provisions the native array snapshots to the mount host. The mount host must be zoned to the native array where the snapshot is created. All other mount considerations of the native array are applicable.

Enable VMware cluster mount

If the mount host is a VMware virtual machine residing on an ESX cluster, the target LUN is made visible to all the nodes of the ESX cluster during mount. By default, this is enabled. If you do not want to perform an ESX cluster mount, you can clear this option. Then the target LUN is made visible only to the ESX cluster on which the mount host resides. This is applicable for both RDM and vDisk device types.

VPLEX restore considerations

Consider the following when restoring data from VPLEX Snap copies:

- The VPLEX production virtual volume layout must be the same as it was when the copy was created. If there is any change in the production virtual volume layout, AppSync detects it and the restore fails.
- In the case of a RAID 1 and distributed devices, AppSync restores one leg of the mirror for which the copy was created. The other leg is rebuilt and synchronized after restore is complete from the native array snapshots. If you do not want to wait for mirror synchronization, ensure that you clear the **Wait for mirror rebuild to complete** option in the Restore wizard.
- During restore, AppSync removes VPLEX virtual volumes from the consistency group, restores from native array snapshot, and adds the virtual volumes back to the consistency group. It also invalidates cache of all the VPLEX virtual volumes.
- AppSync does not support restore of VPLEX production virtual volumes, which are protected by RecoverPoint.
- When restoring from VPLEX Snap copies, ensure that no other operation is performed on the device being restored.

Storage considerations

CHAPTER 13

Troubleshooting AppSync

This section provides information on the common problems encountered while using AppSync.

• Reboot required after installing the AppSync host plug-in.....	292
• User account does not have the correct permissions.....	292
• AppSync Exchange Interface service is partially registered.....	293
• VSS timeout issue.....	293
• Host installation and deployment issue.....	294
• Oracle ASM disk groups cannot be mounted after a host reboot.....	294
• Mount of ASM disk groups fail on RHEL 6.x and 7.x MPIO configurations.....	294
• AppSync fails to mount Oracle ASM disk groups (Event - ORCL_000043)	295
• AppSync fails to unmount Oracle ASM disk groups (Event - ORCL_000044) ..	295
• Oracle database discovery failure.....	296
• Oracle database discovery failure - / file system full.....	296
• Oracle database fails to start after a reboot.....	296
• Oracle restart script not removed for UNIX hosts registered using a SUDO user on agent uninstallation.....	297
• AppSync fails to create symlinks.....	297
• Checking system logs for Oracle restart.....	297
• Database fails to start with the created SPFile	298
• Oracle recovery failure.....	298
• AppSync fails to freeze the SQL Server database in a timely manner (Event - SQL_000018).....	299
• SQL Cluster - second generation copy mount failure.....	299
• Timeout error during SQL database discovery.....	300
• <AppSync>\jboss\standalone\tmp\vfs\ folder disk usage.....	300
• XtremIO copy creation takes time.....	300
• Changing an XMS IP.....	300
• Error during datastore or virtual disk mount.....	301
• Virtual disk mapping failure.....	301
• Protection of File systems or Oracle applications on MPIO devices fail during mapping.....	301
• Repurposing file systems on multiple LUNs fail.....	302
• Changed file system type not updated after host discovery.....	302
• Mount of a file system snapshot to RHEL7 fails.....	303
• CST lockbox restore failure.....	303
• Protection and repurposing failure on VMAX.....	303
• AppSync services do not start after reboot.....	304
• Windows server configuration issue.....	304
• Flash on Windows Server 2012 R2.....	304
• Browser refresh.....	305
• Scheduled service plan fails.....	305
• Event logging.....	306
• Error handling.....	306

Reboot required after installing the AppSync host plug-in

Problem

If you are installing the AppSync host plug-in for the first time, you might have to reboot your Exchange server after the install completes. In some environments, the install of the Visual C++ 2010 runtimes require a reboot. In previous releases of AppSync, the install rebooted automatically without warning. However, automatic reboots have been removed from AppSync 2.0 and later.

Resolution

To verify if a reboot is required, look in `Windows\temp\appsync_host_plugin_setup.log` for the following:

```
>: "C:\Program Files\EMC\AppSync Host
Plug-in\msredist\vcredist-10.0.40219.1-x64.exe" /q
Command.run(): starting stdout monitor...
Command.run(): starting stderr monitor...
Command.complete(): waiting for process to complete.
Command.complete(): process completed with exit
code: 3010
```

The exit code of 3010 indicates that a reboot is required.

User account does not have the correct permissions

Problem

If the EMC AppSync Exchange Interface service fails to register properly, check the `ExchangeInterfaceInstall.log` file in the AppSync host plug-in\logs directory. A common problem is that the user account for running the service was not granted the `Log on as a batch job` permission.

If AppSync fails to discover databases, verify the EMC AppSync Exchange Interface service user account has been granted the correct Exchange permissions.

Resolution

To grant the useraccount the correct permissions, and manually register the EMC AppSync Exchange Interface service:

1. Grant the user account that will run the EMC AppSync Exchange Interface service `Log on as a batch job` and `Log on as a service user rights`.
2. Open a command prompt and navigate to the directory where the AppSync Host Plug-in is installed. The default location is `C:\Program Files\EMC\AppSync Host Plug-in`.
3. Run the following command to register the service and the DCOM component:

```
awExchangeInterface /service /user <"domain\username"> /
password <"password"> /nopriv For example: awExchangeInterface /
service /user mydomain /appsyncechuser /password
mYp@55W0rd.
```

4. To configure the password for the DCOM component, run `DCOMCNFG`.
5. Expand **Component Services > Computers > My Computer > DCOM Config**.

6. Right click on **EMC AppSync Exchange Interface** and select **Properties**.
7. Click on the **Identity** tab.
8. Select **This user** and enter the user account and password from step 3.
9. Click **OK**.
10. Verify that you can start the EMC AppSync Exchange Interface service by running: `net start appsyncexchangeinterface`.
11. Use the AppSync console to rediscover the server. Go to **Settings > Servers**, select the server, and then click **Rediscover**.
12. Discover the Exchange mailbox databases. Go to **Copy Management > Exchange** and click on the Exchange server. You may have to re-enter the credentials.

AppSync Exchange Interface service is partially registered

Problem

If the rights and permissions are not granted properly to the user account, or if conflicting software is installed, the AppSync Exchange Interface service does not register correctly. You might have to perform a manual cleanup.

Resolution

Do the following:

1. Open a command prompt and navigate to the directory where the AppSync host plugin is installed. The default location is `C:\Program Files\EMC\AppSync Host Plug-in`.
2. Run the following command to remove the service and delete the DCOM component: `awExchangeInterface /unregserver`
3. Using the Services console (`service.msc`), verify that the EMC AppSync Exchange Interface service is removed. If it persists, run: `sc delete AppSyncExchangeInterface`
4. Using the Component Services console (DCOMCNFG), verify that the EMC AppSync Exchange Interface DCOM component was removed. **Expand Component Services > Computers > My Computer > DCOM Config**.
5. If the component persists, click **DCOM Config**, then in the center pane, click **EMC AppSync Exchange Interface**, and then click **Delete**.
6. Using **REGEDIT**, verify that all the stale entries related to AppSync Exchange Interface are deleted.

VSS timeout issue

Problem

During protection of applications which reside on Unity, XtremIO, or VPLEX on XtremIO, protection fails with the VSS timeout error.

Resolution

Add the IP address and the FQDN of the XMS in the host file located at `C:\Windows\System32\drivers\etc\hosts`. For example, `10.247.169.71 lrmb071`.

Note

You can configure VSS retry settings in the create copy phase of a service plan for Windows applications such as File system, Microsoft SQL, and Microsoft Exchange. For SQL Server databases, you can consider taking Crash Consistent backup. The Protect SQL Server chapter provides more information on the Crash Consistent backup type.

Host installation and deployment issue

Problem

When installing the agent plug-in, SUDO user installation might fail.

Resolution

Ensure the following:

1. SSH service is configured on the Unix/Linux systems.
2. BZip2, OpenSSH, and OpenSSL packages are available for AIX.
3. The sg3_utils package is available for Linux.
4. SSH port 22 is unblocked.
5. There is sufficient space available at the install location.
6. The home/install directory of the SUDO user has the "write" privilege.

The *AppSync Installation and Configuration Guide* provides additional information on installing the SUDO user.

Oracle ASM disk groups cannot be mounted after a host reboot

Problem

Production Oracle ASM disk groups cannot be mounted after a host reboot because of conflicting ASMLIB disks. The udev rules that mask the devices of an AppSync mounted copy does not get loaded after a reboot leading to conflict between the production ASMLIB devices and the mounted copy's devices. If udev rules are not loaded, then the mounted copy's devices are exposed through their ASMLIB header because that information is present on the replicated device and it is not hidden by the udev rules. Therefore, the ASM instance sees two ASMLIB disks with the same name and gets confused.

Resolution

Do one of the following:

- Unmount the copy in AppSync.
- Manually reload the udev rules according to the Linux platform version.

Mount of ASM disk groups fail on RHEL 6.x and 7.x MPIO configurations

Problem

If you set the disk string to `/dev/mapper/*` on the mount host, it can lead to a conflict because AppSync attempts to mask devices using the disk string `/dev/emc-`

appsync-*. The /dev/emc-appsync-* paths are UDEV rules based NAME parameter (in the case of RHEL 6.x) or UDEV rules based SYMLINK+ parameter (in the case of RHEL 7.x), and it is like an alias over the /dev/mapper/* devices. The conflict occurs because the same target device is masked using two paths - /dev/mapper/* path and /dev/emc-appsync-* path, and ASM does not accept duplicate paths for candidate disks.

Resolution

Remove /dev/mapper/* from the asm_diskstring parameter using the following command:

```
alter system set asm_diskstring= '<paths without /dev/mapper/*>'
scope=both
```

For example, if existing ASM disks have paths with MPIO aliases such as /dev/mapper/asm_disk<n>, change /dev/mapper/* to /dev/mapper/asm_disk*.

AppSync fails to mount Oracle ASM disk groups (Event - ORCL_000043)

Problem

AppSync fails to mount Oracle ASM disk groups.

Resolution

1. Check the previous agent log for mountASMFilesystems operation to confirm if all the related devices have surfaced correctly.
2. If MPIO on Linux 6.x and 7.x is used, ensure that no duplicate paths are presented to ASM through the existing asm_diskstring parameter. [Mount of ASM disk groups fail on RHEL 6.x and 7.x MPIO configurations](#) provides more information.
3. This issue might occur if the asm_diskstring parameter is empty or if it is set to nested paths such as /dev/*, /dev/asm-disk*. Ensure that a proper value is assigned to the asm_diskstring parameter.
4. For Linux flavors, this issue might occur if there is any spurious udev rules file present under /etc/udev/rules.d/ directory masking the same target devices with some other NAME/SYMLINK parameter. Ensure that no such file exists and remove the files, if any.
5. Ensure that there is enough space in /tmp.

AppSync fails to unmount Oracle ASM disk groups (Event - ORCL_000044)

Problem

AppSync fails to unmount Oracle ASM disk groups during a restore operation.

Resolution

1. Check if there are any affected databases that must be shutdown manually before restore. AppSync reports unprotected affected databases before restore and you must shut them down manually.

2. If the failure occurs during unmount of a mounted copy, connect to the ASM instance on the mount host and manually dismount the mounted disk group using the sqlplus `alter diskgroup <diskgroup name> dismount` command . Before executing this command, ensure that the mounted and recovered database is shutdown.

Oracle database discovery failure

Problem

If a symbolic link is used for `ORACLE_HOME` in `/etc/oratab` for a database, AppSync fails to discover the database, if you start the database with the actual path.

Resolution

Ensure that you start the database with the `ORACLE_HOME` variable set to the same symbolic link path.

Oracle database discovery failure - / file system full

Problem

Oracle database discovery fails with the following error:

```
No valid instance available when / file system is full
```

Resolution

Ensure that either disk space or inode usage in root (`/`) file system is not full.

Oracle database fails to start after a reboot

Problem

The AppSync mounted Oracle databases does not start after a host reboot.

Resolution

If the mounted and recovered databases do not come up post a reboot, then the most common reasons can be:

- The file systems cannot be mounted before the database restart is triggered. AppSync waits for a maximum of 1 minute for the file systems to be mounted.
- The ASM instance cannot be started (if database resides on ASM disk groups). AppSync waits for a maximum of 2 minutes for the ASM instance to start.
- After reconnecting to the rebooted host, if you notice that the ASM disk groups and/or the file systems on which the mounted and recovered database resides are mounted, do the following:
 1. Navigate to the AppSync agent install path (usually `/opt/emc/appsync`).
 2. If root is the owner of AppSync installation, run `./acp -b` .
 3. If a SUDO user is the owner of AppSync installation, then run `su - <sudouser> and sudo ./acp -b`.
 4. Check whether all the databases that were mounted and recovered using AppSync are opened in the desired open mode using sqlplus.

When you run `acp -b`, check for the events.

To check logs after a host reboot, refer to the following locations:

On Linux:

- Check `/var/log/messages`.
- Check `/var/log/boot.log`.

On AIX:

1. Add a line in `/etc/inittab` to direct the logs to boot log and console logs.

```
asdbora:2:once:/etc/asdbora 2>&1 | alog -tboot > /dev/console
#oracle restart
```

2. Enable syslog, if it is not enabled, and check for errors.
3. Check the boot log and the console log.

Oracle restart script not removed for UNIX hosts registered using a SUDO user on agent uninstallation

Problem

The Oracle restart script created by AppSync, that is, `/etc/asdbora` for AIX and `/etc/init.d/asdbora` for Linux might not be removed, if a SUDO user is used for registering the host. There is no impact on functionality, it is only a cleanup issue.

Resolution

After the host agent is removed and uninstalled from the AppSync server:

1. Manually remove the script.
2. Remove the symlinks pointing to the `asdbora` file, if they are not removed from `/etc/rc.d/[rc0.d,rc2.d, rc3.d, rc5.d]`.

AppSync fails to create symlinks

Problem

AppSync fails to create symlinks to the `asdbora` file when you select the restart database post reboot option.

Resolution

On Linux, run the following commands:

- `ln -s /etc/init.d/asdbora /etc/rc.d/rc0.d/K01asdbora`
- `ln -s /etc/init.d/asdbora /etc/rc.d/rc3.d/S99asdbora`
- `ln -s /etc/init.d/asdbora /etc/rc.d/rc5.d/S99asdbora`

On AIX, run the following commands:

- `ln -s /etc/asdbora /etc/rc.d/rc0.d/K01asdbora`
- `ln -s /etc/asdbora /etc/rc.d/rc3.d/S99asdbora`

Checking system logs for Oracle restart

To check system logs for Oracle restart:

On Linux, see:

- `/var/log/messages`
- `/var/log/boot.log`

On AIX:

1. Enable `syslog` (Note: Be careful while editing `/etc/syslog.conf`)
2. Create an empty file such as `touch /tmp/syslog.out`
3. Copy the following lines and paste it in `/etc/syslog.conf`

```
*.debug          /tmp/syslog.out    rotate size 100k
files 4
*.crit          /tmp/syslog.out    rotate time 1d
*.info         /tmp/syslog.out    rotate time 1d
*.*            /tmp/syslog.out    rotate size 100k
files 4
```

4. Type the following command to restart the syslog service:

```
refresh -s syslogd
```

5. Check whether `syslog.out` is populated with the logs.
6. Confirm that an entry in `/etc/inittab` exists for: `asdbora:2:once:/etc/asdbora 2>&1 | alog -tboot > /dev/console #oracle restart`
7. Type the following command to view the boot log:

```
alog -o -t boot
```

8. Type the following command to view the console log:

```
alog -o -t console
```

Database fails to start with the created SPFile

Problem

When you mount a copy, the database does not start with the created SPFile.

Resolution

To start the database with the created SPFile, connect to the database and run the following commands:

- `shutdown immediate`
- `startup mount`

Oracle recovery failure

Problem

Oracle database recovery might fail during mount, if the database name starts with a numeric character.

Resolution

Ensure that the database name does not start with a numeric character because Oracle does not allow an `ORACLE_SID` to begin with a numeric character.

AppSync fails to freeze the SQL Server database in a timely manner (Event - SQL_000018)

Problem

AppSync might fail to freeze the SQL Server database in a timely manner, if there is heavy IOPs on the database or due to other database performance issues.

Resolution

1. Add a registry key `CC_AGENT_THREAD_WAIT_TIME` of type `REG_DWORD` with value of 1200 (See step 3 to determine the actual time taken during a Microsoft VDI backup). Also, add another key of type `REG_DWORD` with the same value for VDI timeout as `CC_SQL_VDI_TIMEOUT`.
2. Consider taking a non-VDI backup. Refer the SQL server mount and restore considerations for limitations with non-VDI backups.
3. Contact the SQL Server database administrator to address the performance issues. Microsoft also provides VDI backup diagnostic tools that can be leveraged to check the time taken by Microsoft VDI backups for a database. Contact Microsoft or Dell EMC support for more information.
4. You can consider taking Crash Consistent backup. The Protect SQL Server chapter provides more information on Crash Consistent backup.

Note

For Crash Consistent SQL backup, AppSync does not use VSS or VDI framework. Therefore, there is no VSS timeout or Freeze/Thaw failure issues.

Note

This does not address the VSS timeout issues that occur due to 10 second limitation from Microsoft. This resolution is applicable only if AppSync fails to quiesce the SQL Server database before VSS comes into the picture.

SQL Cluster - second generation copy mount failure

Problem

When you mount both the first generation and the second generation copies sequentially under the same mount drive in a mount cluster, the mount of the second generation copy fails.

Resolution

Ensure that you place the existing mounted drives in maintenance mode. After all the clustered mount points are placed in the maintenance mode, the second generation mount succeeds.

AppSync places only the root of the mount point in maintenance mode. You must place the other clustered drives mounted within the same root drive in maintenance mode. Before unmounting the copies, you must take out all such drives out of the maintenance mode.

Timeout error during SQL database discovery

Problem

During SQL database discovery, you might receive a request timed out message even after configuring the SQL server settings with the appropriate credentials.

Resolution

On the SQL Server, navigate to **Control Panel > Administrative Tools > ODBC application** to start the ODBC application and ensure that the application is functioning.

Alternatively, you can consider running the `sfc` or the `scannow` command in DOS as Administrator to resolve possible issues with the ODBC application. However, check with your SQL or Windows administrator before running the command, or work with the owner of the ODBC application to resolve issues.

<AppSync>\jboss\standalone\tmp\vfs\ folder disk usage

Problem

The files in <AppSync>\jboss\standalone\tmp\vfs\ directory are temporary files. If the files accumulate in the folder, disk usage can be very high.

Resolution

To free up disk space:

1. Stop the AppSync server service.
2. Stop the AppSync datasource service.
3. Delete all the files from the `C:\EMC\AppSync\jboss\standalone\tmp\vfs` directory.
This cleans up the old `tmp\vfs` files that can impact swapping.
4. Start the AppSync datasource service.
5. Start the AppSync server service.

XtremIO copy creation takes time

Problem

Snapshot creation on XtremIO takes significant time (more than 2 seconds) irrespective of the host (Windows or UNIX).

Resolution

Add the IP address and the FQDN of the XMS in the host file located at `C:\Windows\System32\drivers\etc\hosts`. For example, `10.247.169.71 lrmb071`.

Changing an XMS IP

Problem

Changing an XMS IP when a mount operation is in progress can lead to unmount failure because AppSync looks for the old XMS IP.

Resolution

Before you change your XMS IP address, ensure that you unmount all the XtremIO copies mounted in AppSync.

Error during datastore or virtual disk mount

Problem

If copy target LUNs are exposed to ESX, but they are not mounted (unmounted state) as datastores to ESX, you might encounter the following error:

```
Host Platform Config fault
```

Resolution

Ensure that more than one copy of the same datastore is not left in an unmounted state on ESX.

Virtual disk mapping failure

Problem

Virtual disk mapping fails even after VMware vCenter Server and appropriate storage array is added.

Resolution

You must set the `disk.EnableUUID` value to true.

1. Power off the virtual machine.
2. Log into vCenter Server or the ESXi/ESX host through the vSphere Client.
3. Right-click the virtual machine, and click **Edit settings**.
4. Click the **Options** tab.
5. Go to **Advanced > General > Configuration Parameters**.
6. Add or modify the row `disk.EnableUUID` with the value TRUE.
7. Click **OK** to save
8. Click **OK** to exit.
9. Right-click the virtual machine and click **Remove from Inventory** to unregister the virtual machine from the vCenter Server inventory.

Note

If you perform this change using the command line, use the `vim-cmd` command to reload the vmx file. For more information, see the relevant VMware knowledge base article.

10. Power on the virtual machine.

Protection of File systems or Oracle applications on MPIO devices fail during mapping

Problem

The protection of File systems or Oracle applications on MPIO devices on RHEL fails during mapping with the following error: `disk.EnableUUID=TRUE` is not enabled for the virtual device.

The reason might be one of the following:

- File systems are created on `/dev/mpath/` MPIO devices.
- The `lvm.conf` file might read `preferred_names = ["^/dev/mpath/", "^/dev/mapper/mpath", "^/dev/[hs]d"]`.

Resolution

- AppSync supports devices created only on `/dev/mapper/` devices, because they are persistent and created early during boot.
- The `lvm.conf` file must be modified to read `preferred_names = ["^/dev/mapper/mpath", "^/dev/[hs]d"]`.

Repurposing file systems on multiple LUNs fail

Problem

Repurposing file systems on multiple LUNs might fail, if the file systems being repurposed are from different storage types.

Resolution

Ensure that you repurpose file systems on the same array together.

Problem

Repurposing file systems on multiple LUNs might fail, if the number of LUNs to be protected exceeds the number of storage units allowed by the AppSync server.

Resolution

- Try repurposing fewer file systems.
- By default, the number of storage units allowed is 12. If you want to repurpose additional file systems, configure the server settings in the following manner:
 - For VMAX arrays - `max.vmax.block.affinity.storage.units`
 - For VNX arrays - `max.vnx.block.affinity.storage.units`
 - For VPLEX and XtremIO arrays - `max.common.block.affinity.storage.units`

For more information on configuring the server settings, contact Dell EMC Support.

Problem

Repurposing file systems on multiple LUNs might fail, if the subset of file systems to be repurposed does not exist on the host.

Resolution

Create repurpose copies and schedules again.

Changed file system type not updated after host discovery

Problem

AppSync does not automatically update the Windows file system type when a file system is changed from NTFS to ReFS or ReFS to NTFS.

Resolution

Do the following:

1. Remove the file system on the host.

2. Rediscover file systems in AppSync.
3. Add the new file system with the changed type on the host.
4. Rediscover file systems in AppSync again.

Mount of a file system snapshot to RHEL7 fails

Problem

Mount of a file system snapshot to RHEL 7 fails.

Resolution

Ensure that you set the `auto_activation_volume_list` parameter in the `lvm.conf` file to read the following: `auto_activation_volume_list = ["vplexvg1", "vplexvg2", "ol"]`

CST lockbox restore failure

Problem

Login to AppSync fails with a lockbox issue. Check logs at `C:\EMC\AppSync\apache-tomcat\logs` to verify failure.

Resolution

1. Stop all the three AppSync services, that is EMC APPSYNC Datastore Service, EMC AppSync Security Server Service, and the EMC AppSync Server Service.
2. Delete all the files under `C:\EMC\AppSync\cst\xml`.
3. Copy the files from `C:\EMC\cstBackup\<timestamp>\xml` to `C:\EMC\AppSync\cst\xml`. Ensure that you copy the files with the latest working timestamp.
4. Start all the three AppSync services.
5. Login to AppSync.

Protection and repurposing failure on VMAX

Problem

When you protect or repurpose applications on VMAX, it fails with the following error:

```
the device is already in requested state
```

Resolution

Protection or repurposing must have failed before activating the clone session, which AppSync had placed in the PreCopy state. To verify, fetch the state of the clone session for each of the source devices. Do the following:

1. Login to the host that has SMI-S installed.
2. Open a command prompt window. Type the following commands:

```
cd "\Program Files\EMC\SYMCLI\bin"
symclone -sid <symmetrix_id> list
```

3. If the clone sessions of the source are in the PreCopy state, add the source and target device information in a file (each line of the file must be in the <SOURCE

DEVICE> <TARGET DEVICE> format) and activate the clone session using the `symcli` command.

4. Type the following command to activate the session and change the clone session from the PreCopy state to CopyInProg state:

```
symclone -sid <symmetrix_id> -file <full_path_to_filename>
activate -noprompt
```

AppSync services do not start after reboot

Problem

If AppSync Server is installed on a Windows host, the AppSync security server and the AppSync server services do not start after a reboot.

Resolution

Run a repair of the AppSync server on the host to resolve this issue.

Windows server configuration issue

Problem

When configuring Windows server, it fails with the following error: Host plug-in registered to another AppSync server

Resolution

Do the following:

1. Select **Start > Run**.
2. Type `regedit` and click **OK**.
3. Navigate to **HKEY_LOCAL_MACHINE > Software > EMC > AppSync**.
4. Delete the following registry key entries:
 - CC_AUTH_PEER_CERT
 - CC_AUTH_SELF_CERT
 - CC_AUTH_SELF_KEY
5. Restart the AppSync host plug-in service before attempting to add the Windows agent to another AppSync server.

Flash on Windows Server 2012 R2

Problem

There is an issue with running Flash on Windows Server 2012 R2. Adobe Flash is only installed and enabled when you enable the Desktop Experience. It is not necessary to enable Flash Player on Windows 2008 R2.

Resolution

For Windows 2008 R2, go to the Adobe website to download and install the latest version of Flash Player.

For Windows Server 2008 R2 or Windows Server 2012 R2, you must turn off Internet Explorer Enhanced Security Configuration. Close any open Internet Explorer windows. Review these scenarios:

1. • If you are running Windows Server 2008 R2, browse to the Security Information section of Server Summary, and then click **Configure IE ESC** to open the **Internet Explorer Enhanced Security Configuration** dialog.
 - If you are running Windows Server 2012, click **Configure this local server** to open the **Local Server** configuration page. Browse to Properties, next to IE Enhanced Security Configuration, and then click **On** to open the Internet Explorer Enhanced Security Configuration dialog.
2. To edit Internet Explorer Enhanced Security Configuration when members of the local Administrator group are logged on, browse to **Administrators**, and then click **Off**.
3. To edit Internet Explorer Enhanced Security Configuration when all other users are logged on, browse to **Users**, and then click **Off** (Recommended).
4. Click **OK** to apply your changes.

Additionally, enable the Desktop Experience on Windows Server 2012. Follow these steps:

1. Open Server Manager and click **Add Roles and Features**.
2. When the Add Roles and Features Wizard appears, specify the values on the Installation Type, Server Selection, and Server Roles pages.
3. On the Features page, expand User Interfaces and Infrastructure and select **Desktop Experience**.
4. On the Confirmation page, select **Restart the destination server automatically if required** and click **Install**.

Browser refresh

Problem

If you refresh the browser using the browser's **Refresh** icon or pressing F5 on your keyboard, content in the AppSync console are not displayed.

Resolution

To refresh content in the AppSync console, use the **Refresh** icon on the console's top right corner.

Scheduled service plan fails

Problem

If you have scheduled a recurring service plan for a database, for example, everyday at 3 PM, and a backup tool is also scheduled to run at the same time on a particular day (for example, Friday at 3PM), AppSync protection might fail because of resource conflicts.

Resolution

To schedule a service plan to run every day, excluding a particular day (for example, Friday at 3PM), you must create two schedules:

- Set a schedule to run "Every day at..." at > 12 AM, 3 AM, 6 AM, 9 AM, 12 PM, 6PM, 9PM
- Set a schedule to run "On selected days..." at > 3PM on Sunday, Monday, Tuesday, Wednesday, Thursday, Saturday

Event logging

- The `events.log` file captures the generated events for every run of a service plan. It is located in the same place as the `server.log` file.
- The rotation count of the log file and the size of log file is dependent on **Settings > AppSync Server Settings > Trace Logs**.

Error handling

The AppSync logging format is fixed and any log monitoring tool can be tuned to match the appropriate expression to raise an alert in service desk. You can check the metadata part of a logged event to determine if an event is an error event or not. If you see a TYPE-ERROR, then it is an error event, the ID appears after the EVENT in [], and the text that appears (excluding the metadata information) after the event ID is the event message. Other details such as the time of error, from which AppSync server, from which Appsync user, and so on can also be tracked. The category of events are indicated in the EVENT ID (for example, ORCL, DPL, HST, SPP, and so on).

The following are some examples of AppSync generated events:

- 07-09-2016 14:06:57.489 INFO [Thread-58 (HornetQ-client-global-threads-2113097824)]
[com.emc.archway.service.eventservice.EventServiceBean] []
[] EVENT [ORCL_000104]: During discovery, AppSync detected that the following database(s) were offline: SymASM. As a result, they will not be available for protection.(METADATA: TYPE-ERROR, TIME-2016-07-09 14:06:57.469-0400NATIVETIME-2016-07-09 14:06:57.469-0400, HOST-lrmk096, PHASE-, THREAD=Thread-58 (HornetQ-client-global-threads-2113097824), USER-admin, CATEGORY-GENERIC, SESSIONID-5br0S4+nz3-n6SzORPWYGAcn.undefined)
- 07-21-2016 05:07:45.076 INFO [Thread-156 (HornetQ-client-global-threads-436170702)]
[com.emc.archway.service.eventservice.EventServiceBean] []
[] EVENT [SPP_000001]: Mount copy phase for RPdb1 beginning(METADATA: TYPE-INFO, TIME-2016-07-21 05:07:45.076-0400NATIVETIME-2016-07-21 05:07:45.076-0400, HOST-lrmk096, PHASE=Mount copy, THREAD=Thread-156 (HornetQ-client-global-threads-436170702), USER-admin, CATEGORY-GENERIC, SESSIONID-FPO4yGMevayAeRzNPJC8AAF0.undefined)
- 07-21-2016 05:04:26.027 INFO [Thread-158 (HornetQ-client-global-threads-436170702)]
[com.emc.archway.service.eventservice.EventServiceBean] []
[] EVENT [UNM_000001]: Skipping unmount phase. There were no previously mounted copies found for the applications under protection during this cycle.(METADATA: TYPE-INFO, TIME-2016-07-21 05:04:26.027-0400NATIVETIME-2016-07-21 05:04:26.027-0400, HOST-lrmk096, PHASE=Create CRR bookmark copy, THREAD=Thread-158 (HornetQ-client-global-threads-436170702), USER-admin, CATEGORY-GENERIC, SESSIONID-)

- 07-21-2016 05:04:10.178 INFO [Thread-158 (HornetQ-client-global-threads-436170702)]
 [com.emc.archway.service.eventservice.EventServiceBean] []
 [] EVENT [SPP_000001]: Application mapping phase for RPdb1 beginning(METADATA: TYPE-INFO, TIME-2016-07-21 05:04:10.178-0400NATIVETIME-2016-07-21 05:04:10.178-0400, HOST-lrmk096, PHASE-Application mapping, THREAD=Thread-158 (HornetQ-client-global-threads-436170702), USER-admin, CATEGORY-GENERIC, SESSIONID-)
- 07-21-2016 05:04:09.398 INFO [Thread-162 (HornetQ-client-global-threads-436170702)]
 [com.emc.archway.service.eventservice.EventServiceBean] []
 [] EVENT [MILE_000006]: Application discovery phase for RPdb1 completed successfully(METADATA: TYPE-INFO, TIME-2016-07-21 05:04:09.398-0400NATIVETIME-2016-07-21 05:04:09.398-0400, HOST-lrmk096, PHASE-Application discovery, THREAD=Thread-162 (HornetQ-client-global-threads-436170702), USER-admin, CATEGORY-MILESTONE, SESSIONID-)
- 07-21-2016 05:03:11.102 INFO [Thread-89 (HornetQ-client-global-threads-436170702)]
 [com.emc.archway.service.eventservice.EventServiceBean] []
 [] EVENT [VNX_000052]: Successfully created repurpose VNX snapshot copy AppSyncSnap-20160721_050227:80-545e3dbf-d238-4c8e-bd2e-883666512bb8-APPSYNC_TMP_CG_20160721_050228:484_400_0_488.CopySnap.oracle.autol_repurpose.2.1.20160721_050310:946 of source VNX snapshot AppSyncSnap-20160721_050227:80-545e3dbf-d238-4c8e-bd2e-883666512bb8-APPSYNC_TMP_CG_20160721_050228:484_400_0_488.(METADATA: TYPE-INFO, TIME-2016-07-21 05:03:11.102-0400NATIVETIME-2016-07-21 05:03:11.102-0400, HOST-lrmk096, PHASE-Create 2nd gen archLogs copy, THREAD=Thread-89 (HornetQ-client-global-threads-436170702), USER-admin, CATEGORY-GENERIC, SESSIONID-Hgg-ZRqpCV0ixTrPp-tXbJ+C.undefined)
- 07-21-2016 05:00:28.724 INFO [Thread-158 (HornetQ-client-global-threads-436170702)]
 [com.emc.archway.service.eventservice.EventServiceBean] []
 [] EVENT [LIC_000004]: Storage array APM00140431583 is not licensed for use with AppSync, but is within the 90 day trial period.(METADATA: TYPE-WARNING, TIME-2016-07-21 05:00:28.724-0400NATIVETIME-2016-07-21 05:00:28.724-0400, HOST-lrmk096, PHASE-Create CRR bookmark copy, THREAD=Thread-158 (HornetQ-client-global-threads-436170702), USER-admin, CATEGORY-GENERIC, SESSIONID-Hgg-ZRqpCV0ixTrPp-tXbJ+C.undefined)

GLOSSARY

A

- alternate path** An user defined alternate path that AppSync uses while mounting a copy.
- array** A collection of disk drives where user data and parity data may be stored. Devices can consist of some or all of the drives within an array.
- ASM (Automated Storage Management)** A disk volume manager used for storing Oracle files, ASM allows administrators to add and remove disks while the database is available. Data is automatically striped across all disks in a disk group.

C

- command line interface (CLI)** Method of operating system or application software by typing commands to perform specific tasks.
- continuous data protection** The method of data protection in which all changes to data are continuously captured and tracked, allowing for data recovery to any point in time.
- copy - SQL backup** SQL Server backup type that is used to protect the database and active part of the transaction logs without affecting the sequence of backups. The copy backup type allows you to take a backup without affecting other backup tools that might be creating full copies of the database.

D

- default path** AppSync's predefined alternate path used while mounting a copy. The format is %SystemDrive%\AppSyncMounts\%%ProdServerName%\.
- distributed device** A RAID 1 device whose mirrors are in different VPLEX clusters.

E

- event** A log message that results from a significant action initiated by a user or the system.

F

- failover** Automatically switching to a redundant or standby device, system, or data path upon the failure or abnormal termination of the currently active device, system, or data path.
- full - SQL backup** SQL Server backup type that is used to protect the database and active part of the transaction logs. The full backup type allows you to restore transaction logs, so that you can restore the database to a point-in-time that is newer than the copy.

H

host-plugin An AppSync agent software that is installed and used on production and mount hosts to perform certain functions in order to facilitate protection and recovery of applications.

L

logical unit number (LUN) Virtual storage to which a given server with a physical connection to the underlying storage device may be granted or denied access. LUNs are used to identify SCSI devices, such as external hard drives, connected to a computer. Each device is assigned a LUN number which serves as the device's unique address.

M

metadata Information about data, such as data quality, content, and condition.

mount host The system that AppSync uses to mount a copy. This can be different from the production system.

N

network System of computers, terminals, and databases connected by communication lines.

O

original path The mount path AppSync uses to mount copies. The same path as on the production host.

P

production host The production computer that contains the information system that manages the production data: a database server, Web server, application server, or file server.

R

RAC Real Application Clusters. Allows multiple Oracle instances on different nodes of a cluster to access a shared database on the cluster to facilitate load balancing.

RAID The use of two or more storage volumes to provide better performance, error recovery, and fault tolerance.

RAID 0 A performance-orientated striped or dispersed data mapping technique. Uniformly sized blocks of storage are assigned in regular sequence to all of the arrays disks. Provides high I/O performance at low inherent cost. No additional disks are required. The advantages of RAID 0 are a very simple design and an ease of implementation.

RAID 1 Also called mirroring, this has been used longer than any other form of RAID. It remains popular because of simplicity and a high level of data availability. A mirrored array consists of two or more disks. Each disk in a mirrored array holds an identical image of the user data. RAID 1 has no striping. Read performance is improved since either disk can be read at the same time. Write performance is lower than single disk storage. Writes must be performed on all disks, or mirrors, in the RAID 1. RAID 1 provides very good data reliability for read-intensive applications.

recover The additional operation performed on a protected application to bring the application online and running, after the copy has been mounted or restored.

RecoverPoint Appliance (RPA) Hardware that manages all aspects of data protection for a storage group, including capturing changes, maintaining the images in the journal volumes, and performing image recovery.

RecoverPoint site All RecoverPoint entities on one side of the replication.

restore The process performed on a copy of a protected application to bring the production data of the application to contain consistent data up to a point in time that is earlier than the current time.

S

service plan Defines the attributes of a copy that AppSync creates and manages.

SRDF A technology that allows two or more Symmetrix systems to maintain a remote mirror of data in more than one location. The systems can be located within the same facility, in a campus, or hundreds of miles apart using fibre or dedicated high-speed circuits. The SRDF family of replication software offers various levels of high-availability configurations, such as SRDF/Synchronous (SRDF/S) and SRDF/Asynchronous (SRDF/A).

storage area network (SAN) A high-speed special purpose network or subnetwork that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users.

storage volume A Logical Unit Number (LUN) or unit of storage presented by the back end array.

T

TimeFinder Symmetrix TimeFinder is a business continuance solution that allows you to use special Symmetrix devices called business continuance volumes (BCVs) to create mirrors of Symmetrix data.

TimeFinder Clone Copies of a source device on multiple target devices. The source and target devices can be either standard devices or BCV devices as long as they are all of the same size and emulation type. Clone copies of striped or concatenated meta devices can also be created, but the source and target meta devices must be completely identical in stripe count, stripe size, and capacity. Once activated, the copy can be instantly accessed by a target's host, even before the data is fully copied to the target device.

TimeFinder Snap A host-accessible device containing track-level location information (pointers), that indicates where the copy session data is located in the physical storage. TimeFinder Snap operations provide instant snapshot device copies, using virtual devices (VDEVs).

U

universal unique identifier (UUID) A 64-bit number used to uniquely identify an AppSync copy.

V

VDI Virtual Device Interface. SQL Server provides an API called Virtual Device Interface (VDI) that helps AppSync agent in providing support for backup and restore operations. These APIs provide maximum reliability and performance, and support the full range of SQL Server backup and restore functionality, including the full range of snapshot backup capabilities.

virtual volume Unit of storage presented by the VPLEX front end ports to hosts. A virtual volume looks like a contiguous volume, but can be distributed over two or more storage volumes.

VSS Volume Shadow Copy Service. A Windows service and architecture that coordinate various components to create consistent point-in-time copies of data called shadow copies.

INDEX

A

- agent, See AppSync host plug-in , See AppSync host plug-in
- alert 259–261
 - associated events 259
 - email 261
 - filtering 259
 - when generated 259
 - where displayed 259
- alerts 260
- AlwaysOn Availability Groups 110
- application discovery 234
- AppSync
 - architecture 16
 - console 17
 - host plug-in 16
 - overview 14
 - REST interface 17
 - server 16
 - summary of deployment steps 80
 - user interface 17
- assqlrestore 156
- assqlrestore commands 154
- audience 11

B

- browser refresh 305

C

- call out scripts 55
- CLI 17
- CLI actions 25
- cli utility 24
- cluster mount 265, 272
- coexistence with Replication Manager 17
- comments 11
- console
 - effect of user role on 20
 - overview 20
 - performing actions 20
 - preferences 20
 - starting 21
 - times shown 20
- conventions for publication 11
- copy
 - expire 83
- copyDetails 36
- creating database copy from Copies page 177
- CST lockbox 303

D

- data masking 253
- database

- remove from plan 124
- datacenter 228, 231
 - images 231
- datastore 228, 230, 231, 241
 - affected entities 241
 - configure protection 230
 - unsubscribe from plan 231
- datastore, vmware 229
- datastores 225
- definition 50
- delete a copy 83
- delete a datastore copy 232
- delete a SQL database copy 126
- disableSP 28
- discover 124, 228
- dynamic and statis mounts 281
- dynamic mount 265

E

- email 260
- email alerts 261
- enableSP 27
- event
 - alert-related 259
- event logging 306
- exchange 81
- Exchange 58, 77, 78, 80–82, 86, 87, 91, 93, 96, 99, 100, 102, 103, 237, 258
 - backup type 86
 - consistency check 91
 - DAG 81
 - database and log layout 87
 - deleted database restore 103
 - discover databases 80
 - event log errors 87
 - included plans 58
 - interaction with VSS 86, 93
 - mount 93, 237
 - overview of support 78
 - protect 81
 - protect immediately 82
 - remove mailbox server 81
 - requirements for partial restore 102
 - restore 93, 99, 100
 - subscribe database to a service plan 82, 258
 - VSS 86
- exchange interface service 293
- exchange service plan
 - table of default settings 58
- expire 30
- expiring database 177

F

- file system 210, 220, 302

File system 301
 filesystem 69, 200, 203, 212
 Filesystem 199, 201, 210, 216
 Mount Wizard 216
 Filesystem copies page 202
 Filesystem Mount copy 214
 Filesystem page 201
 Filesystem restore 223
 Filesystem, change mount point 218
 flash 304

H

HACMP cluster integration 165
 HACMP Restore 166, 203

I

installation 294

L

listCopies 34
 login 25
 logout 26

M

Microsoft Exchange, See Exchange , See Exchange
 monitor 260
 Monitor 257
 mount 36, 90, 91, 93, 96, 139, 146, 218, 237, 239, 265, 286, 288
 by server 90
 copy by server 91
 dynamic 265
 host override 91
 locations 93
 minimize log option 96
 on demand 93, 139, 237
 override 90
 throttle option 96
 unmount 96
 unmount datastore copy 239
 unmount SQL database 146, 218
 validation options 96
 virtual machine 265
 wizard 93, 139, 237
 Mount copy 183
 mount on cluster 191
 mount override
 VMware datastore 215, 219, 236
 mount phase
 datastore copy 236
 Mount VMFS copy 237
 MSCS 265, 272

N

NFS filesystems 200

O

on virtual disks 107

Oracel, Veritas Cluster Services 163
 oracle 178
 Oracle 64, 160, 183, 189, 190, 195, 294, 296–298
 affected entities in restore 195
 override mount settings 183
 summary of service plan settings 64
 Oracle copies page 175
 Oracle Data Guard 161
 Oracle mount copy 185
 Oracle on VMware virtual disks 170
 Oracle pre-mount script 182
 oracle service plans 178
 Oracle support 171
 Oracle, application discovery 179
 Oracle, application mapping 179
 oracle, copy expiration 181
 Oracle, mount 187
 oracle, post-copy script 181
 Oracle, pre-copy script 180
 oracle, prerequisites and supported configurations 166
 oracle, protecting 174
 oracle, restore 196
 Oracle, restore 194
 oracle, restoring RAC copy for affected entities 197
 oracle, storage preferences 179
 oracle, unmount copy 184
 oracle, unmount previous copy 182
 Oracle, Unsubscribe from service plan 175
 override schedule 84, 127, 208, 233
 override service plan schedules 178

P

path mapping 97, 98, 147, 148, 192, 193, 221, 222
 post mount script 184
 PowerHA cluster 202
 preface 11
 protect 229
 protect database 82
 Protected virtual machines list 230

R

RecoverPoint 99, 102, 103, 149, 280, 281, 283, 284
 affected entities in restore 99
 affected entities in SQL restore 149
 consistency groups 280
 prerequisites 280
 replication options 280
 restore any-point-in-time 103
 restore granularity 102
 setup 280
 recovery point objective
 compliance report 258, 259
 concepts 258
 setting 84, 127, 209, 233
 refresh 26
 related documentation 11
 report 29
 reports
 overview 16
 recovery point compliance 258, 259
 save to CSV 258

- repurpose 252, 273, 282
 - Repurpose
 - bookmark) 282
 - repurpose copies 256
 - repurpose expire 253
 - Repurpose Monitor 261
 - repurpose refresh 252
 - repurpose schedule 251
 - repurpose, view 256
 - repurposing 250, 254, 287, 302
 - REST interface 17
 - restore 99–103, 149, 195, 289
 - affected entities 99, 195
 - deleted database 103
 - Exchange 99, 100
 - manual option 101
 - partial 102
 - RecoverPoint any-point-in-time 103
 - SQL affected entities 149
 - wizard 100
 - restore folder from virtual disk 246
 - Restore VMFS or NFS datastores 245
 - RHEL 7 303
 - role
 - definition 15
 - effect on console 20
 - rotation 86, 130, 235
 - runSP 27
- S**
- SAN policy
 - prerequisite 266
 - script
 - post-copy 88, 131
 - post-mount 92, 137, 216
 - pre-copy 85, 128
 - service plan 15, 54, 55, 82–86, 88–93, 124, 127, 128, 131, 132, 137, 174, 208–210, 216, 233–235, 237, 285, 288
 - application discovery phase 85
 - application mapping phase 85, 128, 210, 234
 - create copy phase 86
 - create new 54
 - database schedule 84, 127, 208, 233
 - definition 15
 - delete 55
 - disable 54
 - enable 54
 - mount copy phase 89, 90
 - phase details 83
 - post-copy script phase 88, 131
 - post-mount script phase 92, 137, 216
 - pre-copy script phase 85, 128
 - remove database 82
 - run on demand 54
 - save copy as 54
 - schedule 84, 127, 208, 209, 233
 - override 84, 127, 208, 233
 - start 84, 127, 209, 233
 - subscribe database 82
 - unmount copy phase 93, 137, 237
 - unmount previous copy phase 89, 235
 - unmount previous SQL copy 132
 - unsubscribe database 124
 - validate copy phase 91
 - service plan schedule 178
 - service plan summary and details 178
 - service plans 50
 - sql 113
 - SQL 56, 299, 300
 - SQL database
 - copies page 115
 - SQL Database 123
 - SQL database copy
 - expire 126
 - sql log backups 114
 - SQL restore 148, 149
 - SQL server
 - mount and recover copy 143
 - mount copy 141
 - SQL Server 61, 105–109, 119, 122–124, 132, 138, 139, 146, 151, 218
 - configure protection 123
 - considerations in a cluster 119
 - database discovery 122
 - discover server instances 122
 - included plans 61
 - login credentials 109
 - mount 139
 - mount phase of service plan 132
 - permissions 108
 - restore copy 151
 - restore damaged database 151
 - rights 108
 - user databases folder 122
 - user privileges 108, 109
 - SQL server backup 129
 - sql server log backup expiration 118
 - SQL Server service plan
 - table of default settings 61
 - sql server transaction log 110
 - SQL Server:user privileges 108, 109
 - subscribe 31
 - subscribe to plan
 - copy 232
 - subscribe to planVMware
 - datastore 230
 - support information 11
- T**
- transaction log
 - configure 111
- U**
- Unity 285, 286
 - unmount 47, 89, 235
 - as part of service plan 89, 235
 - unmount a copy
 - as part of service plan 93, 137, 237
 - on demand 96
 - unmount a database copy
 - on demand 146, 218
 - unmount a datastore copy

- on demand 239
- unsubscribe 32

V

- vCenter Server
 - adding 228
- Vdisk restore 196
- VIO vSCSI 171
- ViPR C 275
- ViPR Controller
 - service plan 276
- virtual machine
 - copies page 232
- vm consistency 231
- VMAX 57, 270, 272, 273, 303
 - restore 273
- VMAX 3 274, 275
 - mount, unmount 274
- vmax service plans 181
- VMAX service plans 270
- VMAX, Symmetrix copy session arrays 271
- vmfs requirements 226
- vmware
 - restore options 241
- VMware
 - included plans 74
- VMware datastore 239
- VMware datastore copy
 - expire 232
- VMware service plan
 - table of default settings 74
- vmware snapshots 231
- VNX
 - setup 264
- VNX array
 - dynamic mount 265
- VNX file service plan 267
- VNX file setup 266
- VNX File Snapshot copy mount 268
- VNX remote protection 266
- VNX Replicator 266
- VNX service plan 265
- VNX Snapshot
 - affected entities in restore 99, 195
 - affected entities in SQL restore 149
 - expiration 86, 130, 235
 - partial restore 102
- VNX Snapshot, RP bookmarks
 - control replication storage utilization 84
- VNX SnapSure' 266
- VNXe 269
 - service plan 269
- VPLEX 56, 287
- VSS failure 88, 130, 212
- VSS timeout 293

W

- Windows server 304

X

- XtremIO 277, 279
 - restore 279