

EMC[®] Symmetrix[®]
VMAX[®] Family with Enginuity
Product Guide

Symmetrix VMAX[®] 10K (SN xxx987xxxx), VMAX[®] 20K, VMAX[®] 40K

Rev 02

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

Published October, 2013

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

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

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

For the most up-to-date regulatory document for your product line, go to EMC Online Support.

CONTENTS

Preface

Audience.....	11
Related documentation.....	11

Chapter 1

Symmetrix VMAX Family Overview

Symmetrix VMAX family comparison.....	16
VMAX 10K.....	17
VMAX 10K File.....	17
VMAX 20K.....	17
VMAX 40K.....	18
eLicensing.....	18
System capacities.....	19
VMAX 10K.....	19
VMAX 20K.....	20
VMAX 40K.....	20
Host and External Connectivity.....	21
Usable system ports.....	21
VMAX 10K host port configurations.....	21
VMAX 20K and 40K host port configurations.....	22
External connections.....	24

Chapter 2

Symmetrix VMAX Architecture

The Engine.....	25
Engine redundancy.....	25
VMAX 20K and VMAX 40K power subsystem.....	25
VMAX 10K power subsystem.....	26
Physical memory.....	26
Back-end connections.....	26
Environmental monitoring.....	26
The Virtual Matrix Architecture.....	26
Disk array enclosures.....	27
Monitoring and control.....	27
The Service Processor.....	28
Environmental monitoring.....	28
Supported disk drives.....	28
Flash drives.....	29
Fibre Channel drives.....	29
Serial Attached SCSI drives.....	29
Serial Advanced Technology Attachment II drives.....	30
eMLC Enterprise Flash Drives (EFD).....	30

Chapter 3	Engenuity Operating Environment	
	Auto-provisioning Groups.....	32
	Dynamic configuration changes.....	33
	Concurrent configuration changes	33
	Virtual Provisioning	33
	VP compression	34
	Thin devices.....	34
	Data devices	34
	Thin pool.....	34
	Virtual Provisioning pool rebalancing	35
	Space reclamation	36
	Zero space reclamation	36
	Support for draining of data devices.....	36
	Fully Automated Storage Tiering	37
	FAST configuration	38
	Fully Automated Storage Tiering for Virtual Pools	39
	Thin tiers.....	40
	Cascaded storage groups	41
	FAST VP configuration	41
	Allocation by policy.....	42
	SRDF coordination.....	42
	FAST VP for FTS.....	43
	FAST VP compression	43
	Enhanced Virtual LUN technology.....	43
	Configuration considerations	44
	Virtual LUN Virtual Pool (VP) mobility.....	44
	Direct sparing	44
	Permanent sparing.....	45
	Permanent sparing example.....	45
	Configuration considerations	46
	Full spare coverage	47
	Data replication and Recovery software	47
	Symmetrix Remote Data Facility	48
	TimeFinder	48
	TimeFinder VP Snap	49
	RecoverPoint.....	49
	SRDF and RecoverPoint	52
	Performance considerations.....	52
	Performance planning considerations	52
	Dynamic Cache Partitioning.....	53
	Symmetrix Priority Controls	53
	Host I/O limits.....	54
	Cache performance	54
	Disk optimization	55
	Dynamic Mirror Service Policy.....	55
	Rotational positioning ordering	55
	100 percent fast writes.....	55
	Tiered storage optimization.....	56
	EMC XtremSW Cache	56

Chapter 4	Open Systems and Mainframe Support	
	Mainframe support on VMAX 20K and VMAX 40K	57
	FICON configuration	58
	IBM 2107 support	58
	IBM i 512-byte D910 support	58
	Logical control unit capacities	58
	Disk drive emulations	59
	Cascading configurations	59
	Error reporting to the Mainframe host.....	59
	SIM Severity Reporting	60
	Environmental errors (Enginuity 5874 and 5875).....	61
	Environmental errors (Enginuity 5876)	64
	Operator messages	66
	Open systems support on VMAX 10K, 20K, and 40K.....	67
	Open systems intermix configurations.....	67
Chapter 5	Data Integrity and Protection	
	Data Integrity Checking	69
	Physical memory error correction and error verification	69
	Block CRC	69
	Drive integrity monitoring and correction.....	70
	RAID options	70
	Vaulting	71
	The vault operation	71
	Vault configuration considerations.....	72
Chapter 6	Data Migration Tools	
	Federated Tiered Storage	73
	eDisks.....	75
	Modes of operation	75
	Federated Live Migration	76
	Open Replicator for Symmetrix	76
	Open Migrator/Live Migration	77
	PowerPath Migration Enabler	77
	SRDF/Data Mobility.....	77
	Unisphere for VMAX	78
Chapter 7	Storage management tools	
	Unisphere for VMAX	80
	EMC z/OS Storage Manager	80
	EMC ControlCenter and ProSphere	80
	Solutions Enabler SYMCLI	81
	SMI-S Provider	81
	Mainframe Enablers	81
	EMC Virtual Storage Integrator.....	82
	Geographically Dispersed Disaster Restart	82

Chapter 8	Security	
	Audit logs	85
	Symmetrix audit log	85
	Audit Integration with RSA enVision	86
	RSA enVision configuration considerations	86
	Data erasure	86
	IPv6 and IPsec support	87
	Symmetrix Access Controls.....	87
	Symmetrix Data at Rest Encryption	88
	External Enterprise Key Manager	88
	Symmetrix Service Credential, Secured by RSA.....	89
	User authorization controls	90

FIGURES

	Title	Page
1	Auto-provisioning Groups.....	32
2	Thin pool (balanced)	35
3	Thin pool (unbalanced)	35
4	Thin pool (rebalanced)	36
5	FAST configuration example	39
6	FAST VP Configuration example	42
7	The permanent sparing process	46
8	RecoverPoint configurations.....	50
9	z/OS IEA480E service alert error message format (AC power failure)	66
10	z/OS IEA480E service alert error message format (mirror-1 volume in “not ready” state) ... 66	
11	z/OS IEA480E service alert error message format (mirror-2 resynchronization)	66
12	z/OS IEA480E service alert error message format (mirror-1 resynchronization)	67
13	FTS configuration - fabric.....	74
14	FTS configuration - arbitrated loops.....	75

TABLES

	Title	Page
1	Document Change History	14
2	Comparison of Symmetrix VMAX systems	16
3	VMAX 10K system capacities.....	19
4	VMAX 20K system capacities.....	20
5	VMAX 40K system capacities.....	20
6	Front End I/O module for host connectivity per engine.....	21
7	Symmetrix VMAX 10K typical and mixed configurations.....	21
8	Symmetrix VMAX 20K and 40K typical and mixed configurations	22
9	Specifications - per engine	25
10	Supported disk drives	28
11	FAST version differences	37
12	Logical control unit maximums.....	58
13	SIM severity alerts.....	60
14	Environmental Errors Reported as SIM messages	61
15	Environmental errors reported as SIM messages	64
16	RAID options	70
17	Space dedication	72

PREFACE

As part of an effort to improve its product lines, 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 EMC representative if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. New versions of this document might be released on EMC Online Support <https://support.EMC.com>. Check to ensure that you are using the latest version of this document.

Audience

This manual provides information about EMC® Symmetrix® VMAX® 10K (SN xxx987xxxx), VMAX® 20K, and VMAX® 40K systems running Enginuity™ operating environment. This document is intended for use by customers or EMC representatives who wish to understand the offerings for Symmetrix systems.

Note: The VMAX® 10K (SN xxx987xxxx) is referred to in the rest of the guide as VMAX 10K.

Related documentation

The following documentation portfolios contain documents related to the hardware platform and manuals needed to manage your Symmetrix software and storage system configuration. Also listed are documents for external components which interact with your Symmetrix system.

General platform documentation

- ◆ *EMC Symmetrix VMAX Family Documentation Set* — Contains the hardware platform product guide, the physical planning guide, TimeFinder product guide, SRDF product guide, and the Front-end port layout diagrams for the Symmetrix VMAX Family (10K, 20K, and 40K).
- ◆ *EMC Symmetrix System Viewer for Desktop and iPad®* — Illustrates VMAX 10K, VMAX 20K, and VMAX 40K system hardware, incrementally scalable system configurations, and available host connectivity offered for Symmetrix systems.

General host software and Mainframe documentation

- ◆ *EMC Solutions Enabler x.x Complete Documentation Set* — Contains all the product guides and installation manuals needed to manage your Symmetrix using the Solutions Enabler SYMCLI mechanisms.
- ◆ *Mainframe Enablers Vx.x Complete Documentation Set* — Contains all the product guides needed to manage your Symmetrix using the Mainframe Enablers software.

RecoverPoint

- ◆ *RecoverPoint x.x Documentation Set* — Contains the RecoverPoint documentation set, including: Release Notes, Administrator's Guide, CLI Reference Guide, Security Configuration Guide, Glossary of Terms, Deployment Manager Product Guide, and Deployment Manager Release Notes.
- ◆ *EMC® RecoverPoint Deploying with Symmetrix Arrays and Splitter Technical Notes* — Provides information, required procedures, and best practices for deploying RecoverPoint with Symmetrix arrays and splitter.

Unisphere for VMAX

- ◆ *Unisphere for VMAX Online Help* — Explains how to use EMC Unisphere for VMAX for storage system configuration, management, and monitoring. This is the stand-alone help file.
- ◆ *Unisphere for VMAX Release Notes* — Describe the contents of your kit and how to prepare for an installation. These release notes identify any known functionality restrictions and performance issues that may exist with the current version and your specific storage environment.
- ◆ *Unisphere for VMAX Installation Guide* — Provides installation instructions for EMC Unisphere for VMAX.

Typographical conventions

EMC uses the following type style conventions in this guide.

Normal font	<p>In running text:</p> <ul style="list-style-type: none"> • Interface elements (for example, button names, dialog box names) outside of procedures • Items that user selects outside of procedures • Java classes and interface names • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, filenames, functions, menu names, utilities • Pathnames, URLs, filenames, directory names, computer names, links, groups, service keys, file systems, environment variables, notifications
Bold	<p>In procedures:</p> <ul style="list-style-type: none"> • Names of dialog boxes, buttons, icons, menus, fields • Selections from the user interface, including menu items and field entries • Key names • Window names <p>In running text:</p> <ul style="list-style-type: none"> • Command names, daemons, options, programs, processes, notifications, system calls, man pages, services, applications, utilities, kernels
<i>Italic</i>	<p>Used for:</p> <ul style="list-style-type: none"> • Full publications titles referenced in text • Unique word usage in text

Where to get help

EMC support, product, and licensing information can be obtained on the EMC Online Support site as described next.

Note: To open a service request through the EMC Online Support site, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Product information

For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to EMC Online Support (registration required) at:

<http://support.EMC.com>

Technical support

EMC offers a variety of support options.

Support by Product — EMC offers consolidated, product-specific information on the Web at:

<https://support.EMC.com/products>

The Support by Product web pages offer quick links to Documentation, White Papers, Advisories (such as frequently used Knowledgebase articles), and Downloads, as well as more dynamic content, such as presentations, discussion, relevant Customer Support Forum entries, and a link to EMC Live Chat.

EMC Live Chat — Open a Chat or instant message session with an EMC Customer Engineer (CE).

eLicensing support

To activate your entitlements and obtain your Symmetrix license files, visit the Service Center on <http://support.EMC.com>, as directed on your License Authorization Code (LAC) letter e-mailed to you.

For help with missing or incorrect entitlements after activation (that is, expected functionality remains unavailable because it is not licensed), contact your EMC Account Representative or Authorized Reseller.

For help with any errors applying license files through Solutions Enabler, contact the EMC Customer Support Center.

If you are missing a LAC letter, or require further instructions on activating your licenses through the Online Support site, contact EMC's worldwide Licensing team at licensing@emc.com or call:

- ◆ North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.
- ◆ EMEA: +353 (0) 21 4879862 and follow the voice prompts.

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

Document Change History

[Table 1 on page 14](#) provides a description of document changes based on the operating environment release. The *EMC® Symmetrix® VMAX® Family Arrays with Enginuity™ Operating Environment Release Notes* contain supplemental information about release features.

Table 1 Document Change History

Version	Description	Enginuity Operating System
01	Initial release of combined Symmetrix VMAX 10K, 20K and 40K systems running Enginuity Operating System.	5876 to 5876 Q2 2013 SR
02	Addition of VMAX 10K File and 16 Gb/s Front End I/O module.	5876 to 5876 Q2 2013 SR

CHAPTER 1

Symmetrix VMAX Family Overview

The EMC® VMAX® Family of enterprise storage systems offers the world's most comprehensive range of storage solutions for enterprises of all sizes. The VMAX family includes:

- ◆ The EMC® Symmetrix® VMAX® 10K scales from a single VMAX 10K engine in one system bay with no storage bays to a four VMAX 10K engines in four system bays and two additional storage bays. The high density configuration supports only four system bays.
- ◆ The EMC® Symmetrix® VMAX® 20K scales from a single VMAX 20K engine system with one standard storage bay to a large eight engine system with a maximum of ten standard storage bays for standard configurations and a maximum of one system bay and eight high density storage bays for high density configurations.
- ◆ The EMC® Symmetrix® VMAX® 40K scales from one to eight VMAX 40K engines with one system bay and up to ten standard storage bays for standard configurations and a maximum configuration of one system bay and eight high density storage bays for high density configurations. The VMAX 40K also offers the ability to mix standard and high density drives within a system.

This chapter introduces the EMC Symmetrix VMAX 10K, VMAX 20K, and VMAX 40K systems with Enginuity™ operating environment.

◆ Symmetrix VMAX family comparison.....	16
◆ eLicensing	18
◆ System capacities	19
◆ Host and External Connectivity	21
◆ External connections.....	24

Symmetrix VMAX family comparison

When planning the installation site for your Symmetrix system, meet with your EMC Systems Engineer and EMC Customer Engineer to complete the installation planning task sheet and the presite survey. All physical planning details are shown and described in the *EMC Symmetrix VMAX Family Physical Planning Guide*.

For full descriptions of the components and available system configurations, use the *EMC Symmetrix System Viewer for Desktop and iPad®* available on [EMC Online Support](#). The system viewer illustrates the system hardware, incrementally scalable system configurations, and available host connectivity offered for Symmetrix systems.

Table 2 Comparison of Symmetrix VMAX systems

Feature	VMAX 10K	VMAX 20K	VMAX 40K
Architecture	<ul style="list-style-type: none"> VMAX 10K Engine Dual Virtual Matrix Architecture Remote support 	<ul style="list-style-type: none"> VMAX 20K Engine Dual Virtual Matrix Architecture Remote support 	<ul style="list-style-type: none"> VMAX 40K Engine Quad Virtual Matrix Architecture Remote support
Connectivity	FC, 16 Gb FC, 1GbE, 10GbE, 10 Gb/s, FCoE, iSCSI	FC, FICON, 1GbE, 10GbE, 10 Gb/s, FCoE, iSCSI	FC, 16 Gb FC, FICON, 1GbE, 10GbE, 10 Gb/s, FCoE, iSCSI
Dispersion	Up to 32 feet (10 meters)	N/A	Up to 82 feet (25 meters)
Drives (Maximum)	1,560	3,200	
Front-end ports	64	128	
Hardware options	System securing kits, Topside cable routing kit	System securing kits, Symmetrix silencer, Topside cable routing kit	
Power	Single-phase	Single-phase or Three-phase	
Open Systems support	Yes		
Mainframe support	No	Yes	
IBM i	IBM i 512-byte D910	IBM i 512-byte D910, IBM 2107	
Protocols	IPsec (1 Gb/s I/O modules) for SRDF or iSCSI		
Data Migration	Federated Live Migration, Federated Tiered Storage, EMC Open Replicator, EMC Open Migrator/LM, PowerPath, PowerPath Migration Enabler, PowerPath SE		
Local Replication	TimeFinder, RecoverPoint		
Remote Replication	SRDF, RecoverPoint		
Management tools	Solutions Enabler, Unisphere for VMAX	MainFrame Enabler, Solutions Enabler, Unisphere for VMAX	
RAID protections	RAID 1, RAID 5 (3+1), RAID 5 (7+1), RAID 6 (6+2), RAID 6 (14+2)	RAID 1, RAID 5 (3+1), RAID 5 (7+1), RAID 6 (6+2), RAID 6 (14+2), RAID 10	
Security	Audit Logs, Data Erasure, IPsec security features, Symmetrix Access Controls, Symmetrix Data at Rest Encryption, Symmetrix Service Credential, Secure by RSA, User Authorization Controls		

VMAX 10K

The EMC Symmetrix VMAX 10K with Enginuity is a multi-controller, scale-out architecture with consolidation and efficiency for the enterprise that can non-disruptively scale from one to four VMAX 10K Engines. The maximum standard configuration consists of four system bays and two storage bays while the maximum high density configuration consists of four system bays only.

Symmetrix VMAX 10K systems are available with standard 3.5” drives or high-density 2.5” mixed within a storage bay.

Symmetrix VMAX 10K dispersed system bays allow a single VMAX 10K to be separated across two locations by up to 32 feet 10 inches (10 meters) enabling deployment in dense data center environments where floor loading and other physical limitations and obstructions would otherwise limit configurations.

Note: Depending on your geographical area, the VMAX 10K may support different hardware configurations. Contact your EMC representative for more information.

3rd party racking

The Symmetrix VMAX 10K system supports the installation of EMC components into a customer-supplied rack that conforms to the National Electrical Manufacturers Association (NEMA) standard for 19-inch cabinets allowing an installation into a standard rack that conforms to your existing computer room and infrastructure.

Components that are installed in a customer rack are shipped in a fully tested EMC system or storage bay (shipping rack), and are installed by EMC support personnel.

VMAX 10K File

The VMAX 10K File is a solution that provides block and file system storage in a single VMAX 10K storage array. The array includes additional hardware to support file-level access (data movers and control stations) and pre-installed storage management software.

The block storage (Unisphere for VMAX) and the file system (Unisphere) management software are linked using the context sensitive “Link & Launch” feature.

VMAX 20K

The EMC Symmetrix VMAX 20K with Enginuity is a distributed multi-engine storage system that can non-disruptively scale from one to eight VMAX 20K Engines. The maximum standard configuration consists of one system bay and ten storage bays while the maximum high density configuration consists of one system bay and eight storage bays with both configurations allowing for the integration of either standard 3.5” drives or high-density 2.5” drives for space or power constrained data centers.

With Enginuity 5876 Q1 2013 SR, the VMAX 20K supports the ability to mix standard and high density drives within a system.

VMAX 40K

The EMC Symmetrix VMAX 40K with Enginuity is a distributed multi-engine storage system that can non-disruptively scale from one to eight VMAX 40K Engines. The maximum standard configuration consists of one system bay and ten standard storage bays while the maximum high density configuration consists of one system bay and eight high density storage bays with both configurations allowing for the integration of either standard 3.5” drives or high-density 2.5” drives.

With Enginuity 5876 Q4 2012 SR, the VMAX 40K supports the ability to mix standard and high density drives within a system but not within the same storage bay

Symmetrix VMAX 40K dispersed system bays allow a single VMAX 40K to be separated across two locations by up to 82 feet (25 meters) enabling deployment in dense data center environments where floor loading and other physical limitations and obstructions would otherwise limit configurations.

eLicensing

Each Symmetrix system requires a license file that is installed during system installation. The license file contains all the licensed entitlements purchased for the systems, as well as entitled capacity associated with those entitlements. A License Activation Code (LAC) is created and sent as part of the ordering and shipping process.

Once the installation is complete, licensed entitlements can be viewed, managed or activated by using the Unisphere for VMAX, Solutions Enabler CLI, or Mainframe Enabler CLI.

System capacities

Symmetrix VMAX systems are expanded by adding engines, disk array enclosures, and physical memory.

VMAX 10K

[Table 3 on page 19](#) provides VMAX 10K system capacity details.

Table 3 VMAX 10K system capacities

				Standard 3.5"		Dense 2.5"		3.5" and 2.5" ^a	
Engines	FC Back End Ports	Min Drives ^b	Min TBu ^c	Max Drives	Max TBu ^d	Max Drives	Max TBu ^c	Max Drives	Max TBu ^c
1	16	24	2.8	240	579.0	300	248.1	360	392.9
2	32	44	5.7	480	1199.4	600	496.3	760	730.7
3	48	64	8.6	720	1489.0	900	758.3	1160	1068.5
4	64	84	11.5	960	1489.0	1200	1006.4	1560	1406.3

- Calculations are based on either 1 TB (dense) drives or 3 TB (standard) drives RAID 6 (14+2).
- Drive quantities include minimum number of required spares; spare drives do not contribute to usable capacity.
- Calculations are based on 80 mirrored drives. (RAID 1).
- Mixed drive configurations must have a minimum quantity of four 3.5" DAEs. The maximum drive count is achieved with a quantity of four 3.5" DAEs and 60 2.5" DAEs. The maximum usable capacities shown are based on this unique configuration.

VMAX 20K

[Table 4 on page 20](#) provides VMAX 20K system capacity details.

Table 4 VMAX 20K system capacities

				Standard 3.5"		Dense 2.5"		Extended Drive
Engines	FC Back End Ports	Min Drives ^a	Min TBu ^b	Max Drives	Max TBu ^c	Max Drives	Max TBu ^c	Max TBu ^c
1	16	48	2.8	240	579.0	400	330.8	n/a
2	32	88	5.7	480	1199.4	800	675.5	2026.7
3	48	128	8.6	720	1819.9	1200	1006.4	2026.7
4	64	168	11.4	960	2026.7	1600	1351.1	2026.7
5	80	272	14.3	1320	2026.7	2000	1682.0	n/a
6	96	248	17.2	1680	2026.7	2400	2026.7	n/a
7	112	418	20.0	2040	2026.7	2800	2026.7	n/a
8	128	328	22.9	2400	2026.7	3200	2026.7	n/a

- Drive quantities include minimum number of required spares; spare drives do not contribute to usable capacity.
- Calculations are based on 146 mirrored drives, (RAID 1).
- Calculations are based on 1 TB or 3 TB drives RAID 5 (7+1).

VMAX 40K

[Table 5 on page 20](#) provides VMAX 40K system capacity details.

Table 5 VMAX 40K system capacities

				Standard 3.5"		Dense 2.5"	
Engines	FC Back End Ports	Min Drives ^a	Min TBu ^b	Max Drives	Max TBu ^c	Max Drives	Max TBu ^c
1	16	48	5.7	240	579.0	400	330.8
2	32	88	11.5	480	1199.4	800	675.5
3	48	128	17.2	720	1819.9	1200	1006.4
4	64	168	23.0	960	2398.9	1600	1351.1
5	80	272	38.0	1320	3329.6	2000	1682.0
6	96	248	34.5	1680	3874.6	2400	2026.7
7	112	418	58.7	2040	3874.6	2800	2357.6
8	128	328	46.1	2400	3874.6	3200	2702.3

- Drive quantities include minimum number of required spares; spare drives do not contribute to usable capacity.
- Calculations are based on 300 GB mirrored drives, (RAID 1).
- Calculations are based on either 1 TB (dense) drives or 3 TB (standard) drives RAID 6 (14+2).

Host and External Connectivity

The Symmetrix systems supports dual connections to the host or server by allowing multiple connections from a host (direct attached) or a Fibre Channel switch (SAN) to a Symmetrix system. Multiple connections are distributed across Front End I/O Modules and engines to ensure uninterrupted access in the event of an I/O failure. A minimum of two connections per server or SAN is required to provide full redundancy.

Front End I/O modules are ordered in pairs and installed in engine slots 4 and 5. To provide redundancy, Front End I/O modules connect to alternate directors. Each engine supports four I/O modules (two pairs).

The *EMC Symmetrix VMAX Family Documentation Portfolio* provides front-end port layout diagrams for the Symmetrix VMAX 10K, VMAX 20K and VMAX 40K systems.

Usable system ports

[Table 6 on page 21](#) describes the supported front end I/O modules available for host connections per engine.

Table 6 Front End I/O module for host connectivity per engine

VMAX 10K ^a	VMAX 20K	VMAX 40K
<ul style="list-style-type: none"> • FC (up to 8 ports, 16 Gb/s) • FC (up to 16 ports, 2, 4, 8 Gb/s) • FC SRDF (up to 8 ports, 2, 4, 8 Gb/s) • FCoE (up to 8 ports, 10 Gb/s) • GbE SRDF (up to 8 ports, 1 Gb/s) • GbE SRDF (up to 8 ports, 10 Gb/s) • iSCSI (up to 8 ports, 1 Gb/s) • iSCSI (up to 8 ports, 10 Gb/s) 	<ul style="list-style-type: none"> • FC SRDF (up to 8 ports, 2, 4, 8 Gb/s) • FC (up to 16 ports, 2, 4, 8 Gb/s) • FCoE (up to 8 ports, 10 Gb/s) • FICON (up to 8 ports, 8 Gb/s) • GbE SRDF (up to 8 ports, 1 Gb/s) • GbE SRDF (up to 8 ports, 10 Gb/s) • iSCSI (up to 8 ports, 1 Gb/s) • iSCSI (up to 8 ports, 10 Gb/s) 	<ul style="list-style-type: none"> • FC (up to 8 ports, 16 Gb/s) • FC (up to 16 ports, 2, 4, 8 Gb/s) • FC SRDF (up to 8 ports, 2, 4, 8 Gb/s) • FCoE (up to 8 ports, 10 Gb/s) • FICON (up to 8 ports, 8 Gb/s) • GbE SRDF (up to 8 ports, 1 Gb/s) • GbE SRDF (up to 8 ports, 10 Gb/s) • iSCSI (up to 8 ports, 1 Gb/s) • iSCSI (up to 8 ports, 10 Gb/s)

a. The VMAX 10K can support either two or four front end I/O modules depending on your geographical area. Contact your EMC representative for more information.

VMAX 10K host port configurations

[Table 7 on page 21](#) provide typical and mixed usable host port types by Front End Module types.

Table 7 Symmetrix VMAX 10K typical and mixed configurations

Module Types ^b	Total Ports	Hosts Ports ^a					
		16 Gb FC	4 Gb FC	8 Gb FC	10 Gb FCoE	10 Gb iSCSI	1 Gb iSCSI
4 Gb FC	16		16				
8 Gb FC	16			16			
16 Gb FC	8	8					
1 GbE	8						8

Table 7 Symmetrix VMAX 10K typical and mixed configurations

Module Types ^b	Total Ports	Hosts Ports ^a					
		16 Gb FC	4 Gb FC	8 Gb FC	10 Gb FCoE	10 Gb iSCSI	1 Gb iSCSI
10 GbE	8					8	
					4	4	
					8		
4 Gb FC/1 GbE	12		8				4
4 Gb FC/ 10 GbE	12		8			4	
	12		8		4		
4 Gb FC/ 8 Gb FC	16		8	8			
8 Gb FC/1 GbE	12			8			4
8 Gb FC/ 10 GbE	12			8		4	
	12			8	4		
16 Gb FC/ 1GbE	8	4					4
16 Gb FC/ 10 GbE	8	4		4		4	
	8	4		4	4		
1 GbE /10 GbE	8					4	4
	8				4		4

a. Fibre Channel is available with single mode and multimode ports. SRDF is supported over Fibre Channel single mode and multimode ports.

b. Availability may vary depending on the geographical area. Contact your EMC representative for more information.

VMAX 20K and 40K host port configurations

[Table 8 on page 22](#) provide typical and mixed usable host port types by Front End Module types.

Table 8 Symmetrix VMAX 20K and 40K typical and mixed configurations

Front End I/O Module Type	Total Ports	Host Port Type ^{ab}					SRDF ^c		
		FC	FICON	iSCSI 1 Gb	iSCSI 10 Gb	FCoE 10 Gb	FC	1 Gb	10 Gb
Typical configurations									
8 Gb FC	16	16							
	14	12					2		
	12	8					4		
16 Gb FC ^d	8	8							
FICON	8		8						

Table 8 Symmetrix VMAX 20K and 40K typical and mixed configurations

Front End I/O Module Type	Total Ports	Host Port Type ^{ab}					SRDF ^c		
		FC	FICON	iSCSI 1 Gb	iSCSI 10 Gb	FCoE 10 Gb	FC	1 Gb	10 Gb
1 GbE	8			8					
	8			6				2	
	8			4				4	
10 GbE	8					8			
	8					8			
	8				4	4			
	8					6			2
	8					4			4

Mixed configurations	Total Ports	FC	FICON	iSCSI 1 Gb	iSCSI 10 Gb	FCoE 10 Gb	FC	1 Gb	10 Gb
FC 8 Gb/FICON	12	8	4						
	10	4	4				2		
	8		4				4		
FC/1 GbE	12	8		4					
	12	8		2				2	
	8			4			4		
FC/10 GbE	12	8			4				
	12	8				4			
FC 16 Gb/FICON ^d	8	4	4						
FC 16 Gb/1 GbE ^d	8	4						4	
FC 16Gb/10 GbE ^d	8	4				4		2	
FC 16 Gb/1 GbE ^d	8	4						4	
FICON/1 GbE	8		4	4					
	8		4	2				2	
	8		4						
FICON/10 GbE	8		4		4				
	8		4			4			

Table 8 Symmetrix VMAX 20K and 40K typical and mixed configurations

Front End I/O Module Type	Total Ports	Host Port Type ^{ab}					SRDF ^c		
		FC	FICON	iSCSI 1 Gb	iSCSI 10 Gb	FCoE 10 Gb	FC	1 Gb	10 Gb
iSCSI/1 GbE	8			4	4				
	8			4					
	8			4					4
iSCSI/10 GbE	8			4	4				
	8			4		4			
	8			4					4

- a. FCoE is supported on 10 GbE Front End I/O Modules only. It is not supported on 1 GbE Front End I/O Modules.
b. IPsec is available on 1 GbE Front End I/O Modules.
c. SRDF is supported on both Fibre Channel (8 Gb) and Ethernet Front End I/O Modules.
d. Supported in VMAX 40K only.

External connections

Symmetrix systems support Federated Tiered Storage (FTS) which provides the ability to attach external storage to a Symmetrix system through Fibre Channel Front End I/O modules and ports. FTS requires Enginuity emulation type (DX). DX emulations are configured in pairs, with each DX emulation associated to two front-end ports, thus consuming four front-end ports per DX pairs. DX emulations are back-end connections and are not used to connect to hosts. [“Federated Tiered Storage” on page 73](#) outlines how FTS works.

Note: FTS is only supported on 8 Gb Fibre Channel I/O modules.

The *EMC Symmetrix VMAX Family Documentation Portfolio* provides front-end port layout diagrams outlining FTS and DX emulation information.

CHAPTER 2

Symmetrix VMAX Architecture

The Symmetrix VMAX architecture provides a storage subsystem that can scale beyond the confines of a single system footprint. The core element of the VMAX architecture is the Symmetrix VMAX 10K, VMAX 20K, or VMAX 40K Engine, which includes high-availability (HA) directors with either Dual or Quad Virtual Matrix™ Architecture.

This chapter discusses the VMAX architecture. Topics include:

- ◆ [The Engine](#) 25
- ◆ [The Virtual Matrix Architecture](#) 26
- ◆ [Disk array enclosures](#) 27
- ◆ [The Service Processor](#) 28
- ◆ [Supported disk drives](#) 28

The Engine

The Symmetrix VMAX 10K, VMAX 20K, or VMAX 40K engine is a system bay component that provides physical memory, back-end connectivity, front-end host connectivity, and connectivity to other engines. The engine contains two integrated, highly-available Symmetrix directors. Each Symmetrix director contains a CPU complex, protected global memory, and dual-interfaces to either a Dual or Quad Virtual Matrix Architecture™.

[Table 9](#) lists the details per engine supported in each Symmetrix system.

Table 9 Specifications - per engine

System	Intel® XEON® processor	Physical memory per engine
VMAX 10K ^a	Two six-core 2.8 GHz Intel XEON processors	24, 96 or 128 GB (512 GB maximum)
VMAX 20K	Four quad-core 2.33 GHz Intel XEON processors	32, 64, or 128 GB (1 TB maximum)
VMAX 40K	Four six-core 2.8 GHz Intel XEON processor	48, 96, 192 or 256 GB (2 TB maximum)

a. Depending on your geographical area, the VMAX 10K may support different hardware configurations. Contact your EMC representative for more information.

Engine redundancy

The engines contain redundant management modules which provide environmental monitoring. The power and cooling subsystems are also redundant allowing for continuous availability.

VMAX 20K and VMAX 40K power subsystem

For VMAX 20K and 40K, two SPS modules for each engine provide backup power in the event of an AC power interruption. SPS modules maintain power for two five-minute periods of AC loss while the Symmetrix system shuts down.

VMAX 10K power subsystem

For VMAX 10K, two SPS modules for each engine provide backup power in the event of an AC power interruption. SPS modules maintain power for two five-minute periods of AC loss while the system shuts down. A second SPS assembly provides backup power for the vault drives in each system bay.

Physical memory

Memory is accessible by any director within the system:

- ◆ If a system has a single engine, physical memory mirrors are internal to the enclosure.
- ◆ If a system has multiple engines, physical memory mirrors are provided between enclosures.

Dual write technology is maintained by the system. In the event of a director and memory failure, data is obtained from the redundant copy.

Back-end connections

Back End I/O Modules provide dual access across directors to each drive. One module connects to one physical path to the drive and the other connects to a second physical path. The I/O modules are responsible for moving data between physical memory and the disks and, as such, are each connected to the physical memory through redundant internal paths to eliminate any single points of failure.

Environmental monitoring

Under normal conditions, the monitoring and control for each engine is accomplished by individual Management Modules by using an RS-232 and redundant Ethernet communications. In the event of an Ethernet failure, the Management Modules use the alternate Ethernet connection to transmit management traffic between directors to the Service Processor.

The Virtual Matrix Architecture

Symmetrix systems include a high bandwidth, low-latency Dual or Quad Virtual Matrix Architecture that allows engines to communicate with each other. The VMAX 10K and VMAX 20K dual virtual matrix architecture and the VMAX 40K quad virtual matrix architecture use redundant star topology to connect eight engines via four matrix interface boards (MIBEs) per system. Each MIBE contains the Matrix Interface Board (MIB) and two power supply/cooling modules.

Each director in the engine continually monitors the virtual matrix and collects events that measure performance and provide information for early fault detection and isolation. System monitoring includes the Ethernet and RS-232 cabling, virtual matrix communication and hardware components, and link and I/O data errors. Error messaging between the directors and the virtual matrix, and communication between the virtual matrix components provide virtual matrix health reporting, fault management, and failover information in the event of an unrecoverable error condition.

The dual or quad virtual matrix:

- ◆ Provides connectivity among engines
- ◆ Ensures fully redundant, fault tolerant connections that allow directors to access distributed physical memory.

Disk array enclosures

Disk array enclosures (DAEs) are storage modules that contain drives, link control cards (LCCs), power supplies/ cooling components. Symmetrix systems support either a standard disk array enclosure or a high density disk array enclosure.

Each disk array enclosure provides redundant connection to two separate directors. One director connects over one physical path to the disk array enclosure, and the other connects over a second physical path. The disk array enclosure supports 4 Gb/s back-end Fibre Channel interface (FC-AL). The dual-loop configuration provides redundancy, port bypass capability, and disk access at a maximum of 4 Gb/s per loop.

In VMAX 20K and 40K systems, disk array enclosures are contained in one or more storage bays and are configured and expanded in half bay (eight DAEs) or full bay (sixteen DAEs) increments.

In VMAX 10K systems, system bays contain a maximum of 10 DAEs in standard and a maximum of 12 DAEs in high density configurations.

Monitoring and control

Each disk enclosure contains a pair of link control cards (LCC A and B) that provide redundant connections to the disk. The LCCs monitor the Fibre Channel layer for connection and protocol errors. If errors occur, the status is reported and, if necessary, bypass circuits are used to switch out of the drive.

Symmetrix systems also include point-to-point back-end switches that enhance fault detection and isolation. The back-end has an independent relationship with each drive on the loop that optimizes problem isolation and improves serviceability. This dedicated relationship between a back-end controller and each disk drive also allows Symmetrix systems to analyze drive health prior to adding new disks to an existing configuration. Symmetrix systems will not add a faulty drive to an existing loop.

To ensure the highest level of availability, the communications link to the monitoring and control functions is not carried by the Fibre Channel loop. It is implemented through a separate path. And, if the Fibre Channel loop is not operating, the director can also use the out-of-band control path to access the LCC and reconfigure the loop into a working state.

The Service Processor

Through the integrated Service Processor, the Enginuity operating environment monitors all end-to-end I/O operations for errors and faults. By tracking these errors during normal operation, Enginuity can recognize patterns of error activity and predicts a potential hard failure before it occurs. This capability often prevents component failures by fencing off (removing from service) a suspect component before a failure occurs.

The integrated Service Processor provides remote notification and remote support capabilities, and allows EMC personnel to access the system locally or remotely. The Service Processor automatically notifies the Customer Support Center whenever a component failure or environmental violation is detected. This allows an EMC customer engineer to remotely run diagnostics to determine the source of the problem, often resolving it before it becomes critical.

In systems which are dispersed, the service processor located in System bay 2 is used only as a remote service terminal.

Environmental monitoring

Remote notification enables EMC to monitor all environmental statuses reported by the subsystems and for statuses to be collected by the Service Processor. If a failure occurs, the Symmetrix system automatically activates redundant elements and the Service Processor initiates a call-home request for service.

Supported disk drives

[Table 10 on page 28](#) provides a list of supported disk drives by Symmetrix systems. [“RAID options” on page 70](#) provides detailed information about available RAID options.

Note: EMC adds new and higher capacity drives on an ongoing basis, which may occur after the release of published documentation.

Table 10 Supported disk drives

Drive type	System
Flash	VMAX 10K, 20K, 40K
Fibre Channel (FC)	VMAX 10K, 20K, 40K
Serial Attached SCSI (SAS)	VMAX 10K, 20K, 40K
Serial Advanced Technology Attachment (SATA) II	VMAX 20K, 40K

The specification sheet for each system available on [EMC ONE](#) provides detailed information on drive types, rotational speeds, and formatted capacities for Open Systems, Mainframe, and IBM i series.

Flash drives

Flash drives greatly benefit applications with high transaction rates, such as trading systems, that require the fastest data retrieval and storage. Flash drives also improve performance in mixed workload environments, because they improve response time for I/O activity even when placed on only a small, I/O-intensive portion of the data set.

Flash drives are contained in the same 3.5 inch drive or 2.5 inch drive carriers as Fibre Channel drives. However, because there are no mechanical components, Flash drives use less energy and weigh less than traditional rotating disk drives. Flash drives appear as Fibre Channel drives to existing Symmetrix system management tools. This allows the management of Tier 0 storage without special tools.

Flash drives have the same RAID options as other supported drive types and can be used as source and target devices for TimeFinder. You can migrate devices to and from Flash drives.

On Symmetrix VMAX 20K and VMAX 40K systems you can configure metavolumes on Flash drives as long as all of the logicals in the metagroup are on Flash drives.

Symmetrix systems provide software that helps optimize high-performance environments that include Flash drives.

Flash configuration considerations

The following list describes planning and configuration considerations for Flash drives:

- ◆ Sparing is supported using Flash drives. However, Symmetrix will block sparing between Flash drives and magnetic disk drives to avoid potential performance issues.
- ◆ There is a minimum of one spare for one to 32 Flash drives per drive type or two spares per 100 drives if there are greater than 32 Flash drives per drive type.
- ◆ Flash drive spares do not count towards the minimum number of hard disk drive spare drives.

Fibre Channel drives

Symmetrix systems use industry-standard 4 Gb/s Fibre Channel disk drives and support Fibre Channel loops with all RAID options supported.

- Symmetrix VMAX 10K support from 5 to 75 drives per loop.
- Symmetrix VMAX 20K support from 15 to 75 drives per loop. With mixed DAEs, 40 drives per loop are supported.
- Symmetrix VMAX 40K support 50 drives per loop. With mixed DAEs, 40 drives per loop are supported.

Serial Attached SCSI drives

The Symmetrix system uses industry-standard 2.5 inch Serial Attached SCSI (SAS) disk drives, contained in the same 2.5 (high density DAEs) or 3.5 (standard DAEs) inch drive carriers as Fibre Channel drives. RAID 1, 5, and 6 options are supported.

Serial Advanced Technology Attachment II drives

Serial Advanced Technology Attachment (SATA) II drives are a high-density storage option.

SATA II drives and Fibre Channel drives can coexist in the same disk array enclosure, but because of performance differences they should not be mixed in volume protection strategies. For example, SATA II and Fibre Channel drives should not be mixed together in RAID groups.

eMLC Enterprise Flash Drives (EFD)

eMLC (Multi-Level Cell) Enterprise Flash Drives (EFD) that deliver comparable performance, reliability and life expectancy of SLC-based EFDs. eMLC EFDs allow Fully Automated Storage Tiering for Virtual Pool (FAST™ VP) configurations to be built to utilize more flash and more nearline capacity to improve both GB and IOPS.

Chapter 3

Enginuity Operating Environment

The EMC® Enginuity™ operating environment provides the intelligence that controls the components in Symmetrix systems. Enginuity is a storage operating environment that controls storage data flow.

Enginuity optimizes service levels required by high-end environments and is specialized and specifically optimized for storage-based functions. It is driven by realtime events related to the input and output of data. It applies self-optimizing intelligence to deliver the ultimate performance, availability, and data integrity required in a platform for advanced storage functionality.

This chapter provides an overview of features for simplified storage management and provisioning. Topics include:

◆ Auto-provisioning Groups.....	32
◆ Dynamic configuration changes.....	33
◆ Concurrent configuration changes	33
◆ Virtual Provisioning.....	33
◆ Virtual Provisioning pool rebalancing	35
◆ Fully Automated Storage Tiering	37
◆ Fully Automated Storage Tiering for Virtual Pools.....	39
◆ Enhanced Virtual LUN technology.....	43
◆ Virtual LUN Virtual Pool (VP) mobility.....	44
◆ Direct sparing	44
◆ Permanent sparing.....	45
◆ Full spare coverage	47
◆ Data replication and Recovery software.....	47
◆ Performance considerations.....	52
◆ Cache performance	54
◆ Disk optimization.....	55
◆ Tiered storage optimization.....	56
◆ EMC XtremSW Cache.....	56

Auto-provisioning Groups

Auto-provisioning Groups on open systems make provisioning operations faster and easier by allowing grouping of initiators, front-end ports, and devices together, and to build masking views that associate the devices with the ports and initiators. When a masking view is created, the necessary mapping and masking operations are performed automatically to provision storage. Once a masking view has been created, any changes to the grouping of initiators, ports, or storage devices are propagated throughout the view and the mapping and masking are automatically updated as required. Auto-provisioning Groups reduce complexity, execution time labor cost, and the risk of error.

Figure 1 on page 32 depicts the concept of Auto-provisioning Groups.

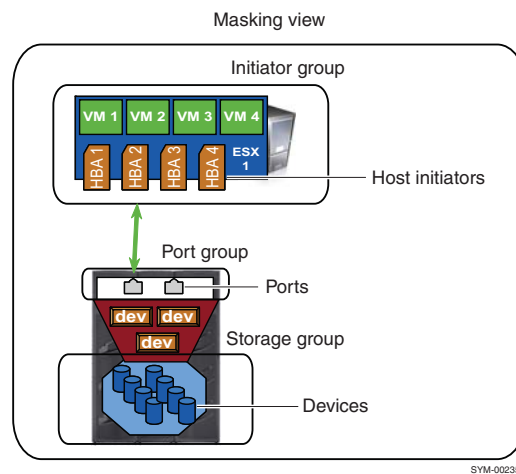


Figure 1 Auto-provisioning Groups

- ◆ An **initiator group** is a logical grouping of up to 32 Fibre Channel initiators or eight iSCSI names or a combination of both. An initiator group may also contain the name of another initiator group to allow the groups to be cascaded to a depth of one.
- ◆ A **port group** is a logical grouping of Fibre Channel and/or iSCSI front-end director ports. The only limit on the number of ports in a port group is the number of ports in the Symmetrix system; however it is likely that a port group contains a subset of the available ports in order to isolate workloads to specific ports.
- ◆ A **storage group** is a logical grouping of up to 4,096 Symmetrix devices. LUN addresses are assigned to the devices in the cascaded storage group when the masking view is created using the dynamic LUN addressing feature.
- ◆ A **masking view** defines an association between one initiator group, one port group, and one storage group. When a masking view is created, the devices in the storage group are mapped to the ports in the port group and masked to the initiators in the initiator group. Depending on the server and application requirements, each server or group of servers may have one or more masking views that associate a set of Symmetrix devices to an application, server, or cluster of servers.

Dynamic configuration changes

The dynamic configuration change framework reduces the impact to hosts while making changes to various configuration parameters.

Enginuity 5876 includes new configuration objects in this framework, including physical disk groups, physical drives, RAID groups, logical volumes and static RDF device assignments. The new dynamic configuration change mechanism also includes operations related to VLUN migrations, permanent sparing, direct sparing and FTS.

Concurrent configuration changes

Concurrent configuration changes provide the ability to run scripts concurrently instead of serially.

Uses for concurrent configuration changes include parallel device mapping, unmapping, or addition, as well as parallel metavolume configuration from different hosts. Concurrent configuration changes are also available during drive repair.

Virtual Provisioning

Symmetrix Virtual Provisioning improves storage capacity utilization and simplifies storage management by allowing storage to be allocated and accessed on demand from a pool of storage that services one or many applications. This type of storage has multiple benefits:

- ◆ Enables LUNs to be “grown” into over time with no impact to the host or application as space is added to the thin pool.
- ◆ Delivers space from the thin pool on demand.
- ◆ Provides wide striping for a thin pool.
- ◆ Relieves the storage administrator’s efforts of physical device/LUN configuration.

The Virtual Provisioning feature introduces three new concepts: thin devices, data devices, and thin pools. Thin devices can be created with an inflated capacity, because the actual storage space for the data written to the thin devices is on the data devices. In this way, when additional storage is needed, more data devices can be created in the thin pool.

Virtual Provisioning simplifies data layout, with automated wide striping that provides equivalent or potentially better performance than standard provisioning. Virtual Provisioning is appropriate for all storage types in a tiered storage environment and supports both local and remote replication with SRDF and TimeFinder.

Virtual provisioning provides the ability to persistently preallocate space. Extents that are persistently preallocated are not reclaimed by a standard reclaim operation.

Note: Enginuity 5876 introduces the support for thin CKD 3390 and thin IBM i 512-byte D910 devices. Enginuity 5876 Q2 2013 SR provides support of thin IBM i 512-byte D910 devices on VMAX 10K systems.

VP compression

5876 Q4 2012 SR introduces VP compression, allowing thin device data to be compressed within a thin pool. Data may be compressed manually for an individual device or group of devices, via Solutions Enabler or Unisphere for VMAX. Alternatively, inactive data may be compressed automatically for thin devices that are managed by FAST VP.

In order for data to be compressed, the thin pool containing the data must be enabled for compression. Only allocated extents are compressed. VP compression is supported on FBA and CKD 3390 devices.

Thin devices

Thin devices, also known as VP devices (FBA and CKD), have no storage allocated to them when they are created; rather storage is allocated on-demand from a “bound” thin pool. The first write to a location in a thin device results in space being allocated on a data device from the bound pool.

Data devices

Data devices are grouped in a thin pool and are dedicated to the purpose of providing the actual physical storage used by thin devices. As with thin pools, data devices must have identical device emulation type, reside on identical drive technologies, and use the same RAID protection types, and drive technology.

Thin pool

A thin pool, also known as a VP pool, contains thin devices of identical emulation and protection type, all of which reside on disks of the same technology type and speed.

When a write is performed to a portion of the thin device, the Symmetrix system allocates a minimum allotment of physical storage from the pool and maps that storage to a region of the thin device including the area targeted by the write. These storage allocation operations are performed in small units of storage called thin device extents.

Symmetrix systems balance the allocation of extents across all the data devices in the pool that are enabled and that have remaining unused capacity.

When a read is performed on a thin device, the data being read is retrieved from the appropriate data device in the thin pool to which the thin device is bound. Reads directed to an area of a thin device that has not been mapped do not trigger allocation operations. The result of reading an unmapped block is that a block in which each byte is equal to zero will be returned. When more storage is required to service existing or future thin devices, data devices can be added to existing thin storage pools. New thin devices can also be created and associated with existing thin pools.

A thin device can be presented for host use before all of the reported capacity of the device has been mapped. If the reported capacities of the thin devices using a given pool exceeds the pool available storage capacity, the thin device configuration is said to be oversubscribed.


Oversubscribed thin pools

Oversubscribing of thin pools allows the presenting of larger than needed devices to hosts and applications without having enough physical drives to fully allocate all of the space represented by the thin devices.

Virtual Provisioning pool rebalancing

Virtual Provisioning pool rebalancing, once enabled, is a feature that scans and rebalances thin pools, minimizing the performance issues caused by imbalances. Pool imbalances are usually caused by the addition of devices to the pool. Up to eight thin pools may be rebalanced at any one time.


[Figure 2 on page 35](#) shows load balancing in a thin pool. Because Engenuity software writes data in a round robin fashion, and because the devices in this pool were created at the same time, the devices fill at the same rate and maintain a balanced capacity.



SYM-002362

Figure 2 Thin pool (balanced)


The capacity becomes unbalanced when a new device is added to the pool, as shown in [Figure 3 on page 35](#). The older devices, which contain more data, no longer participate equally when accepting new tracks.



SYM-002363

Figure 3 Thin pool (unbalanced)

[Figure 4 on page 36](#) shows the thin pool after rebalancing. By moving data to the new devices, I/O activity is redistributed across the thin pool, allowing all devices to participate equally.



SYM-002364

Figure 4 Thin pool (rebalanced)

Space reclamation

Space reclamation allows the user to reclaim space already on a drive by deallocating data chunks that contain all zeros on FBA devices. This is most effective when migrating from standard, fully provisioned devices to thin devices.

With Enginuity 5876, space is reclaimed from CKD devices using a utility which processes the VTOC and determines the location of the unused space. The Symmetrix system is then instructed to deallocate the associated extents.

Zero space reclamation

Zero space reclamation eliminates separate steps to reclaim “zero” space when moving data into Symmetrix systems. As data arrives, all zero buckets are stripped out and only chunks that contain user data are stored to the physical drives in the storage pool bound to the corresponding thin device.

Open Replicator for Symmetrix and Federated Live Migration support zero space reclamation when migrating from Symmetrix DMX, CLARiiON and third-party systems to Symmetrix systems (Open Replicator for Symmetrix) and third-party systems to Symmetrix systems (Federated Live Migration).

SRDF supports zero space reclamation. The *EMC Symmetrix Remote Data Facility (SRDF) Product Guide* related to your system provides more details about full SRDF and SRDF migration support for thin devices.

z/OS systems using thin devices that were not defined using the PREALLOCATE and PERSIST attributes, are able to reclaim free space using Mainframe Enablers allowing users to reconcile free space in the VTOC with allocated but unused space in thin pools.

Support for draining of data devices

Support for draining of data devices allows the nondisruptive removal of one or more data devices from a thin device pool, without losing the data that belongs to the thin devices.

Virtual Provisioning support for draining of data devices allows the correction of over-provisioned configurations. Thin pools can be shrunk nondisruptively, improving efficiency by freeing up space for reuse outside the pool. Data devices in a thin pool can be disabled and “drained,” a process by which allocated tracks (containing data) from the draining devices are moved to other data devices in the pool. The disabled data devices can then be removed from the pool and reused for other purposes.

Fully Automated Storage Tiering

Fully Automated Storage Tiering (FAST™) analyzes user workloads and moves volumes from one performance tier to another to optimize performance and/or economy. Symmetrix systems use sophisticated algorithms to monitor activity and recommend strategies for moving the most-used data to the fastest (and most expensive) storage, such as Enterprise Flash Drives (EFD), the least-used data to the slowest (and least expensive) storage, such as SATA, while maintaining the remaining data on Fibre Channel (FC) drives,

There are two FAST products: FAST and FAST for Virtual Pools (FAST VP). [Table 11 on page 37](#) shows the differences.

Table 11 FAST version differences

FAST on Symmetrix VMAX 20K and 40K	FAST VP on Symmetrix VMAX 10K, 20K and 40K
Requires Solutions Enabler 7.1 and higher	Requires Solutions Enabler 7.2 and higher
Requires Engenuity 5874 and higher	Requires Engenuity 5875 and higher
Supports standard devices	Supports thin devices
Supports FBA and CKD device emulations	<ul style="list-style-type: none"> Supports FBA device emulation Engenuity 5876 supports thin CKD 3390 and thin IBM i 512-byte D910 devices Engenuity 5876 Q2 2013 SR provides support of thin IBM i 512-byte D910 devices on VMAX 10K systems
Disk group provisioning (DP) tiers: contain disk groups	Virtual pool (VP) tiers: contain thin pools
DP modes: Auto Approve and User Approve	VP modes: Auto Approve or None
User visible data movement plans and history	No plans or history generated
Federated Tiered Storage (eDisks) not supported	Supports Federated Tiered Storage (eDisks) with Engenuity 5876
FAST VP compression not supported	Supports FAST VP compression with Engenuity 5876 Q4 2012 SR and Solutions Enabler V7.5 and higher
Thin device/thin pool compression not supported	Supports compression for thin devices and thin (VP) pools with Engenuity 5876 Q4 2012 SR and Solutions Enabler V7.5 and higher
Three tiers per policy supported	Four tiers per policy supported with Engenuity 5876 Q4 2012 SR and Solutions Enabler V7.5 and higher

Systems that use FAST can:

- ◆ Reduce acquisition costs by maximizing utilization of less expensive drives (SATA) for infrequently used data
- ◆ Improve performance by optimizing data placement for high access requirements data
- ◆ Reduce overall operating costs and simplify management by consolidating systems, using fewer drives, and reducing power, cooling, and floor space requirements
- ◆ Reduce the total cost of ownership by optimizing the performance of a tiered Symmetrix system.

FAST configuration

FAST is available on systems with Flash, Fibre Channel, SAS, or SATA drives and is supported on systems with two or three of these drive types. To configure the Symmetrix system for FAST, the following actions are taken:

- ◆ Define Symmetrix tiers — A Symmetrix tier is a type of storage, and a set of resources (disk groups/thin pools) from which the storage is selected.
- ◆ Define a FAST policy — A FAST policy groups Symmetrix tiers and assigns an upper limit for each tier. The upper limit specifies how much of an associated storage group can reside on the tier.
- ◆ Define storage groups — A storage group is a set of devices. A storage group is associated with a FAST policy and assigned a priority. A storage group can only be associated with one policy, however, one policy may be associated with many storage groups.

The *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide* provides detailed information about FAST commands and configuration.

Sample configuration

[Figure 5 on page 39](#) shows a Symmetrix system after FAST configuration. The Symmetrix system contains three storage tiers:

- ◆ PrimeTier contains drives that deliver the best performance
- ◆ WorkTier contains one high performance Fibre Channel disk group
- ◆ ArchiveTier contains high capacity SATA drives.

A percentage of the high performance Flash drives on PrimeTier is allocated to the storage groups accessed by three applications:

- ◆ up to 50% of Exchange storage groups can be allocated to PrimeTier
- ◆ up to 50% of Oracle Finance storage groups can be allocated to PrimeTier
- ◆ up to 20% of Oracle Sales storage groups can be allocated to PrimeTier

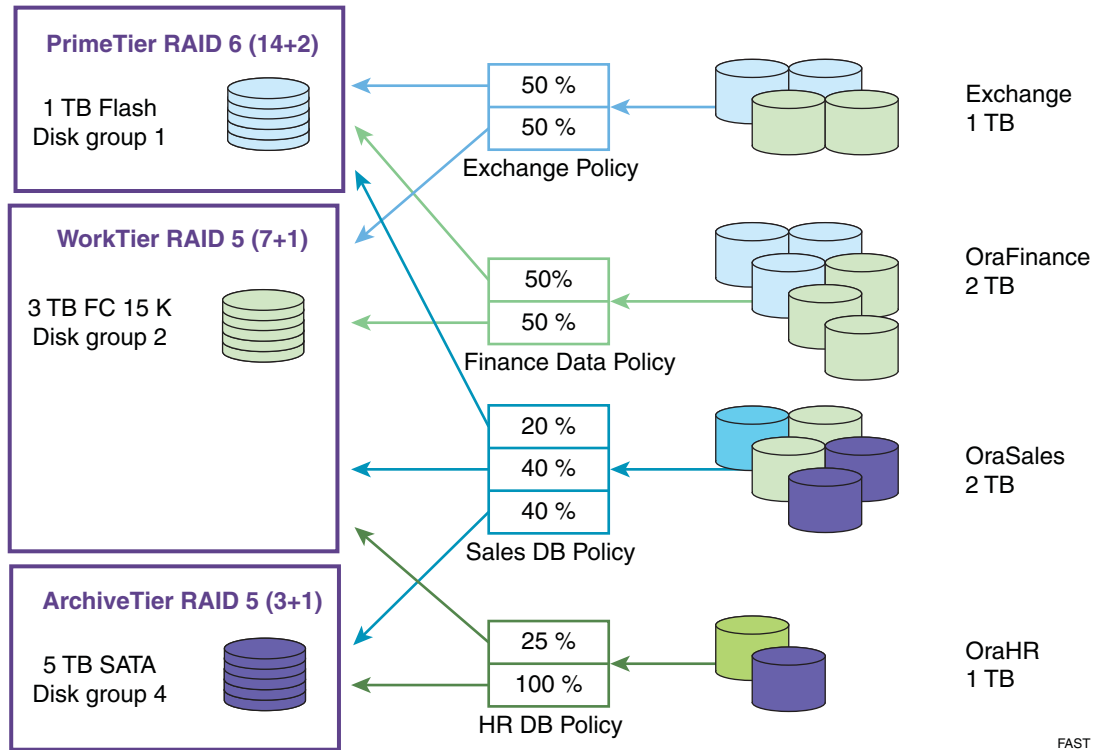


Figure 5 FAST configuration example

Configuration considerations

The following configuration considerations apply to systems that use FAST software:

- ◆ Symmetrix Optimizer is required.
- ◆ FAST systems must configure Dynamic Reallocation Volumes (DRVs) to ensure sufficient space during swap operations. Symmetrix Optimizer uses DRVs to hold data internally during the data swap. Systems without DRV devices that have insufficient free space are degraded, and may be unable to complete a data swap.

Systems that use FAST software must follow the same DRV configuration requirements as Symmetrix Optimizer.

- ◆ FAST operations are not supported on thin devices.

Fully Automated Storage Tiering for Virtual Pools

Fully Automated Storage Tiering for Virtual Pools (FAST VP) automates the identification of data volumes allowing for relocation of application data across different performance/capacity tiers within a Symmetrix system. FAST VP monitors workloads at both the LUN and sub-LUN (FBA) or sub-volume (CKD) level in order to identify data that would benefit from being moved to higher-performing drives. FAST VP identifies data that could be moved to higher-capacity drives, without existing performance being affected based on policies that associate a storage group to multiple drive technologies, or RAID protection schemes, via thin pools, as well as the performance requirements of the application contained within the storage group. Data movement executed during this activity is performed nondisruptively, without affecting business continuity and data

availability. With FAST VP, the data for a thin device may reside in its bound pool and potentially in one or more other pools. The role of thin pools in FAST VP is analogous to the role of disk groups in FAST. Both comprise the back-end storage available to devices under FAST control.

Note: FAST VP with Enginuity 5875 supports FBA devices only. FAST VP with Enginuity 5876 supports thin CKD 3390, thin IBM i 512-byte D910 devices, and FTS.

Enginuity 5876 introduces Federated Tiered Storage (FTS) which allows the virtualization of external storage as an external disk (eDisk). Adding the eDisk to the Symmetrix system makes its capacity available to the system as an external spindle. The order for fastest to slowest tiers is: EFD, FC, SATA.

Enginuity 5876 Q4 2012 SR allows the distinguishing between a local tier and an external tier (where the thin pool resides on external storage). You can define an external tier as EFD, FC, and SATA.

Enginuity 5876 Q4 2012 SR also allows FTS tiers to be associated with a technology type. The technology associated with the FTS tier indicates to the FAST VP controller the expected performance from the tier. This feature enables you to place the FTS tier at the right location for the expected performance of the external tier.

FAST VP adds finer granularities of performance measurement and data movement as small as 7.5 MB or ten 768K track groups. The data from a single thin device under FAST control can be spread across multiple tiers. The FAST controller is free to relocate individual chunks of a thin device, based on performance data gathered at the sub-LUN or sub-volume level.

Systems that use FAST VP software can:

- ◆ Reduce acquisition costs by maximizing utilization of less expensive drives (SATA) for infrequently used data
- ◆ Improve performance by optimizing data placement for high access requirements data
- ◆ Reduce overall operating costs and simplify management by consolidating systems, using fewer drives, and reducing power, cooling, and floor space requirements
- ◆ Reduce the total cost of ownership by optimizing the performance of a tiered Symmetrix system.

[“Virtual Provisioning” on page 33](#) provides detailed information on thin devices and thin pools.

Thin tiers

FAST VP adds a new type of tier, the thin tier which is a set of thin pools. A thin tier consists of thin pools of the same type. To be a member of a thin tier, a thin pool must contain only data devices that reside on tier technology type and match the tier protection type.

Cascaded storage groups

Introduced in 5876, cascaded storage groups allow you to simplify the management of your FAST VP controlled storage and eliminates the need to duplicate masking views for applications with the same masking views but different FAST VP policies and storage groups.

Cascaded storage groups allow you to create a *parent* storage group comprised of multiple storage groups (parent storage group members) and *child* storage groups comprised of devices. By assigning child storage groups to the parent storage group members and applying the masking view to the parent storage group, the masking view inherits all devices in the corresponding child storage groups.

FAST VP configuration

FAST VP is available on systems with Flash, Fibre Channel, SAS, or SATA drives and is supported on systems with two or three of these drive types. To configure the Symmetrix system for FAST VP, the following actions are taken:

- ◆ Define Symmetrix tiers — A Symmetrix tier is a type of storage and a set of resources (disk groups/thin pools) from which the storage is selected. Tiers are a shared resource.

Enginuity 5876 Q4 2012 SR supports up to four tiers, with the fourth tier being FTS, per FAST VP policy.
- ◆ Define a FAST VP policy — A FAST VP policy groups Symmetrix tiers and assigns an upper limit for each tier. The upper limit specifies how much of an associated storage group can reside on the tier.
- ◆ Define storage groups — A storage group is a set of devices. A storage group is associated with a FAST VP policy and assigned a priority. A storage group can only be associated with one policy, however, one policy may be associated with many storage groups.

In order to have different performance levels between storage groups, multiple policies would share a tier. Policies would not be shared for one type of performance level. If the performance level varies, two policies can be created, one which gives more percentage and one with a smaller percentage to a Flash tier.

The *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide* provides detailed information about FAST commands and configuration.

Sample configuration

[Figure 6 on page 42](#) shows a Symmetrix system after FAST VP configuration. The Symmetrix system contains three storage tiers:

- ◆ PrimeTier contains devices that deliver the best performance
- ◆ WorkTier contains one high performance Fibre Channel disk group
- ◆ ArchiveTier contains high capacity SATA drives.

A percentage of the high performance Flash drives on PrimeTier is allocated to the storage groups accessed by three applications:

- ◆ up to 50% of Exchange storage groups can be allocated to PrimeTier

- ◆ up to 50% of Oracle Finance storage groups can be allocated to PrimeTier
- ◆ up to 20% of Oracle Sales storage groups can be allocated to PrimeTier

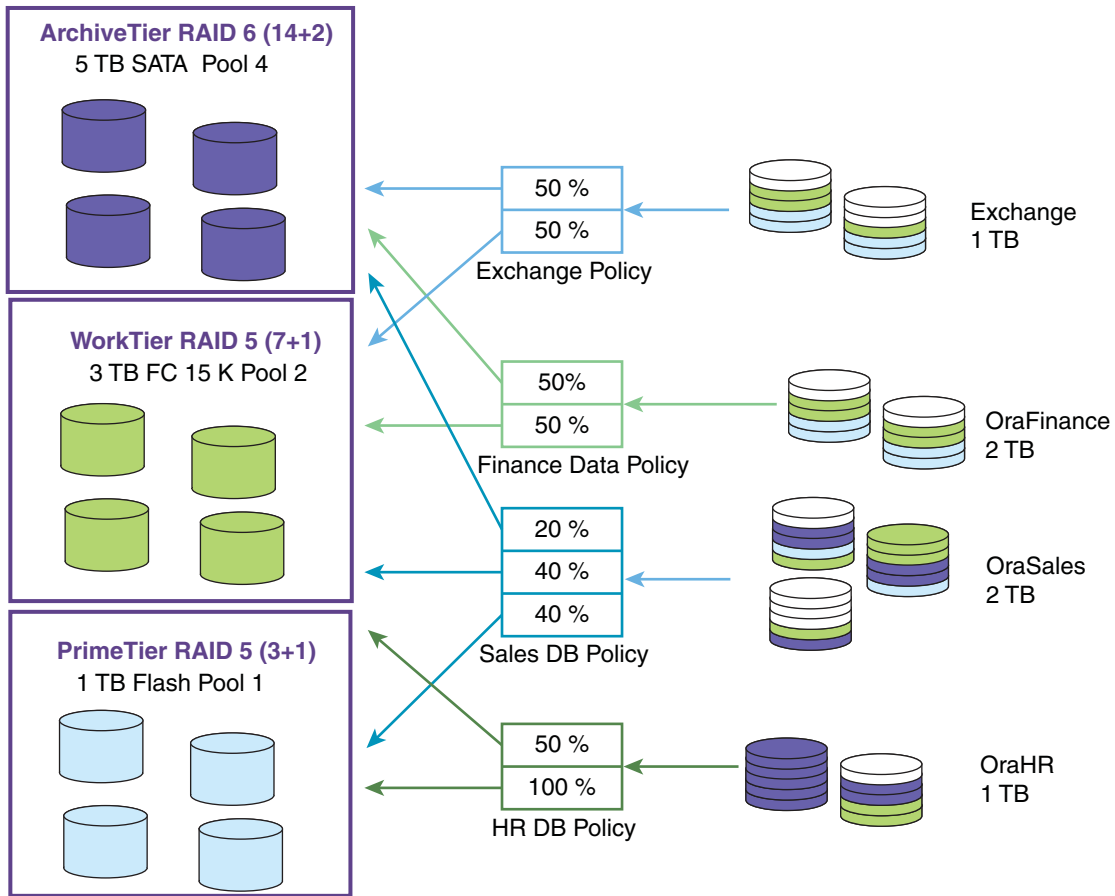


Figure 6 FAST VP Configuration example

Allocation by policy

Introduced in Enginuity 5876, FAST VP will allow for allocations to come from any pool contained within the tiers of the associated FAST VP policy. FAST VP will attempt to allocate on an appropriate tier based on performance metrics and policies.

SRDF coordination

FAST VP is interoperable with SRDF solutions. The FAST VP software can act independently on both the local and remote VMAX Family systems. With Enginuity version 5876 and EMC host-based SRDF software, you can coordinate data movements on both sides of the SRDF links in a two-site SRDF topology. With Enginuity 5876 Q2 2013 SR, this capability is extended to multi-site SRDF topology.

FAST VP for FTS

FAST VP supports tiers of externally provisioned VP pools. Encapsulated devices are not supported. The support for FAST VP is explained below:

- ◆ Tiers — You can create thin tiers, which can contain externally provisioned pools. The external tiers can only contain thin pools configured with externally provisioned data devices. FAST VP for FTS introduces an additional tier.

In Enginuity 5876 Q4 2012 SR, FAST VP no longer treats FTS tiers as the lowest tier. For each FTS tier, you are able to define the tiers type of storage (SATA, FC, EFD) which sets the expected performance from the external tier. Enginuity 5876 Q4 2012 SR supports up to four tiers, with the fourth tier being FTS, per FAST VP policy.

- ◆ Policy — You can add external tiers to a FAST VP policy.
- ◆ Association — Once a policy with an external tier is part of the association, data from the associated storage group can move to/from the external tier without any restriction.

The *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide* provides additional details about FAST and FTS.

FAST VP compression

FAST VP Compression allows for the automated, non-disruptive compression of inactive thin device data for devices associated with a FAST VP policy. FAST VP is enabled at the system level by setting the FAST VP Time to Compress control parameter. The time to compress may be set between 40 and 400 days, or never – never means FAST VP will not compress data. Data that is seen to be inactive for a period of time longer than the time to compress will be considered eligible for automatic compression.

FAST VP Compression performs compression at the subLUN level. It requires that the FAST VP policy contain a pool that has been enabled for compression.

Enhanced Virtual LUN technology

Enhanced Virtual LUN technology provides the ability to change the protection and/or type of a Symmetrix device, thereby providing the means for quick, manual migration of an application's data to other storage within the Symmetrix system. Data migrations are typically initiated to:

- ◆ Move data to a lower storage tier as it ages and becomes less critical to business operations
- ◆ Optimize storage performance or respond to cyclical business changes by allowing the user to move data between storage tiers nondisruptively.

The Enhanced Virtual technology feature can utilize either configured or unconfigured storage to migrate the data of the Symmetrix devices to the desired RAID protection and/or disk type.

- ◆ When configured storage is used — The storage of an existing device, or target, with the desired protection and/or disk type is used to perform the migration. The storage of this target device is unassociated with it and associated with the Symmetrix device.

The data of the Symmetrix device is then copied from the existing to the newly associated storage. The storage originally associated with the Symmetrix device is unassociated with it and associated with the target. All data on the target is cleared.

- ◆ When unconfigured storage is used — The system creates storage of the requested protection and disk type and associates it with the Symmetrix device. The data is then copied from the existing to the newly associated storage. The storage originally associated with the Symmetrix device is deleted, thereby returning it to unconfigured storage.

Configuration considerations

The following requirements apply to configurations that use Virtual LUN technology:

- ◆ Enginuity supports the migration of metadevices. However, an individual meta member may not be specified as a volume for migration.
- ◆ Physical disks that contain the devices for migration (source) may not be specified as a physical destination disk.
- ◆ Source and target devices for Virtual LUN technology migrations can be:
 - Of any emulation type (FBA, CKD and IBM i Symmetrix device)
 - On any supported drive type (Flash, Fibre Channel, SAS or SATA)
 - Configured using RAID 1, RAID 5, or RAID 6.
- ◆ The source devices must be standard Symmetrix devices.

Virtual LUN Virtual Pool (VP) mobility

Virtual LUN Virtual Pool (VP) mobility, introduced in Enginuity 5875, extends the capabilities of Enhanced Virtual LUN technology by nondisruptively relocating all of a thin device's allocated extents from one thin pool to a target thin pool. In the case where FAST VP is in use and a thin device's allocated extents have been spread across multiple thin tiers (pools), Virtual LUN VP mobility enables the consolidation of all of the device's allocated extents in the target pool.

Direct sparing

Introduced in Enginuity 5876, Direct sparing allows for the invoked spare drive to be added as another member of the RAID group. When rebuilding the data, the option to directly copy the data from the failing drive to the invoked spare drive is supported. The failing drive is removed only when the copy process is finished.

Direct sparing is used with protection types, RAID 1, RAID 5, and RAID 6. RAID 6 (14+2) is not supported. RAID 6 (14+2) protection will be achieved via permanent sparing. Volumes with SRDF connectivity may also take advantage of Direct sparing.

Direct sparing provides the following benefits:

- ◆ The Symmetrix system can copy the data from the failing RAID member (if available), removing the need to read the data from all of the members and doing the rebuild. Copying to the new RAID member is less CPU intensive.

- ◆ If a failure occurs in another member, the Symmetrix system can still recover the data automatically from the failing member (if available).
- ◆ More than one spare for a RAID group is supported at the same time.

Permanent sparing

Symmetrix systems support permanent sparing, a process that automatically replaces a failing drive with a spare drive. A Symmetrix system initiates sparing when certain errors are detected, thereby reducing the amount of time that a failed or failing drive remains active in the system.

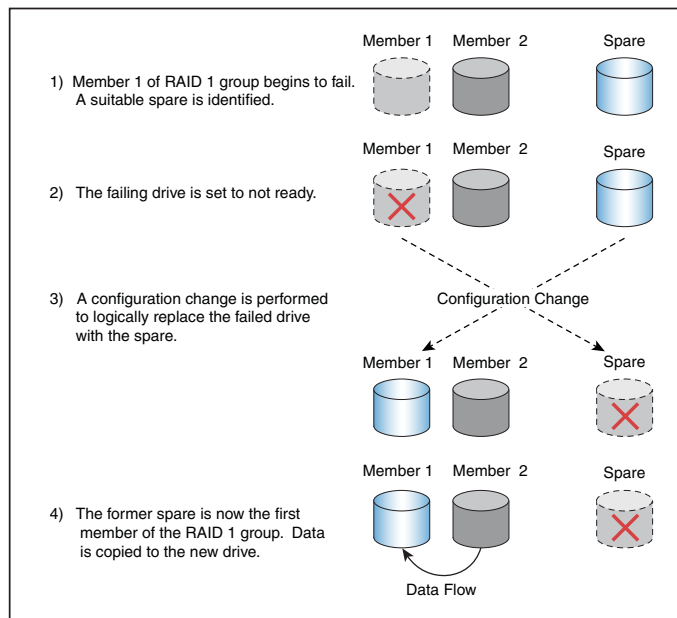
Permanent sparing permanently replaces a failing drive with a spare drive through a configuration change. Permanent sparing is used in conjunction with RAID 6 (14+2).

During permanent sparing, the Symmetrix system looks for a spare drive of the same block size, capacity, and speed, in a good location, to permanently replace the failing drive. Once a spare is identified, the permanent sparing process loads a new configuration file. This reconfigures the logical devices from the failing drive onto the selected spare. The configuration is locked during this change, which can take a few seconds to complete. Data is then rebuilt onto the drive. The configuration is not locked during the rebuild process, and the Symmetrix system continues to process I/Os as the highest priority.

The failed drive becomes a “not ready” spare in the spare pool. The “not ready” designation ensures the drive will not be used for any subsequent sparing activity. The “not ready” designation is removed when the drive is physically replaced, at which time it is again available as a spare.

Permanent sparing example

[Figure 7 on page 46](#) shows an example of permanent sparing on a RAID 1 group. Permanent sparing of RAID 5 and RAID 6 devices is identical in that the configuration change logically replaces the failed data drive with the spare drive.



SYM-001874

Figure 7 The permanent sparing process

Configuration considerations

Spare drives are required for every Symmetrix system configuration. The required number of drives needed for spares is determined automatically when the EMC service representative configures a new Symmetrix system or adds disk array enclosures to an existing configuration. Spare drives are marked automatically and distributed among the available directors. The EMC service representative can override automatic spare placement to meet special configuration needs, or to manually configure an individual spare or spares.

The amount and type of spare drives are calculated as follows:

- ◆ The required number of spares is based on drive type and is automatically configured on ordering. The minimum number of spares required are:
 - Four spare drives for VMAX 10K
 - Eight spare drives for VMAX 20K and VMAX 40K
- ◆ Two spare drives for every 100 drives of a specific type are required.
- ◆ EMC determines the rules for the quantities and types of spares after careful analysis of field data. However, these calculations are considered the minimum and additional drives can be configured. To provide the best possible protection, consider configuring spare drives on every loop and for every type of data drive in the loop.

Configurations with Flash drives must meet Flash drive requirements. [“Flash drives” on page 29](#) provides detailed information about configurations that use Flash drives as spares.

Full spare coverage

Introduced in 5876 Q4 2012 SR, full spare coverage is initiated on initial configurations for Symmetrix systems. Direct sparing does not permanently relocate a drive to a suboptimal location in the Symmetrix system, for example, another disk director slice or loop.

Data replication and Recovery software

Symmetrix system local data replication software provides replicated datasets for business continuance, backup, testing, data restoration, and data migration. Symmetrix system remote replication software extends these protections to include protection for primary site disasters and extended power outages.

Symmetrix system data replication software products include:

- ◆ [“Symmetrix Remote Data Facility” on page 48](#) — Provides remote replication (remote mirroring), disaster recovery and data migration services.
- ◆ [“TimeFinder” on page 48](#) — Provides local data replication services.
- ◆ [“RecoverPoint” on page 49](#) — Provides integrated continuous data protection and continuous remote replication.

Symmetrix Remote Data Facility

The EMC Symmetrix Remote Data Facility (SRDF) family of products offers a range of Symmetrix-based disaster recovery, parallel processing, and data migration solutions.

SRDF disaster recovery solutions are based on active remote mirroring and dependent-write consistent copies of data maintained at one or more remote locations. A dependent-write is a write operation that cannot be issued by an application until a prior, related write I/O operation is completed. Dependent-write consistency is required to ensure transactional consistency once the applications are restarted at the remote location.

SRDF configurations require at least two Symmetrix systems. These systems are also known as the primary and the secondary system. Both sites can be located in the same room, in different buildings within the same campus, or hundreds to thousands of kilometers apart.

The SRDF family of products require that Enginuity run on every Symmetrix system that is part of an SRDF solution. Different versions of Enginuity running on different Symmetrix hardware models can be used in the SRDF solution. The *EMC Symmetrix Remote Data Facility (SRDF) Product Guide* outlines SRDF features dependant on the version of Enginuity code your system is running.

The *SRDF Two-site Interfamily Connectivity* application on EMC Online Support provides details about SRDF two-site interfamily connectivity options.

TimeFinder

The TimeFinder family of products are Symmetrix local replication solutions designed to nondisruptively create point-in-time copies of critical data. You can configure backup sessions, initiate copies, and terminate TimeFinder operations from mainframe and open systems controlling hosts using EMC Symmetrix host-based control software.

The TimeFinder local replication solutions include TimeFinder/Clone, TimeFinder/Snap, and TimeFinder VP Snap. TimeFinder/Clone creates full-device and extent-level point-in-time copies. TimeFinder/Snap creates pointer-based logical copies that consume less storage space on physical drives. TimeFinder VP Snap provides the efficiency of Snap technology with improved cache utilization and simplified pool management.

Note: VMAX 10K does not support TimeFinder/Snap, mainframe host environments, or extent-level point-in-time copies.

Each solution guarantees high data availability. The source device is always available to production applications. The target device becomes read/write enabled as soon as you initiate the point-in-time copy. Host applications can therefore immediately access the point-in-time image of critical data from the target device while TimeFinder copies data in the background.

TimeFinder/Clone is suitable if:

- ◆ Full-volume copies are intended for recovery scenarios.
- ◆ Full-volume or extent-level point-in-time copies of production data have to be immediately available to applications for activities such as reporting and testing.

- ◆ The majority of data on the production volumes changes between subsequent backup sessions.
- ◆ Multiple copies of production data are needed, and you want to reduce disk contention and improve data access speed to the production data.

TimeFinder/Snap is suitable if:

- ◆ Only a fraction of data on the production volumes changes between subsequent backup sessions.
- ◆ Only a fraction of data on the production volumes frequently changes during the peak I/O activity window when multiple point-in-time copies are required.

TimeFinder VP Snap is suitable if:

- ◆ You want to create space-efficient snaps for thin devices
- ◆ You want multiple sessions to share capacity allocations within a thin pool, thus reducing the storage required for saved tracks.

TimeFinder includes the following features:

- ◆ Supports RAID 1, RAID 5, and RAID 6, and RAID 10 protection schemes.

Note: VMAX 10K does not support RAID 10.

- ◆ Provides restore capabilities.
- ◆ Provides incremental re-synchronization between the source and the target.
- ◆ Supports Virtual Provisioning.
- ◆ Tightly integrated with SRDF.

The *EMC Symmetrix TimeFinder Product Guide* outlines features dependent on the version of Engenuity code your system is running.

TimeFinder VP Snap

Engenuity 5876 supports TimeFinder VP Snap. TimeFinder VP Snap leverages TimeFinder/Clone technology to create space-efficient snaps for thin FBA devices by allowing multiple sessions to share capacity allocations within a thin pool. VP Snap provides the efficiency of Snap technology with improved cache utilization and simplified pool management. With VP Snap, tracks can be stored in the same thin pool as the source, or in another pool of your choice. Both the source device and the target device must be thin devices. VP Snap is not supported on CKD devices.

Note: VP Snap source devices work with FAST VP, but shared data will not be moved.

RecoverPoint

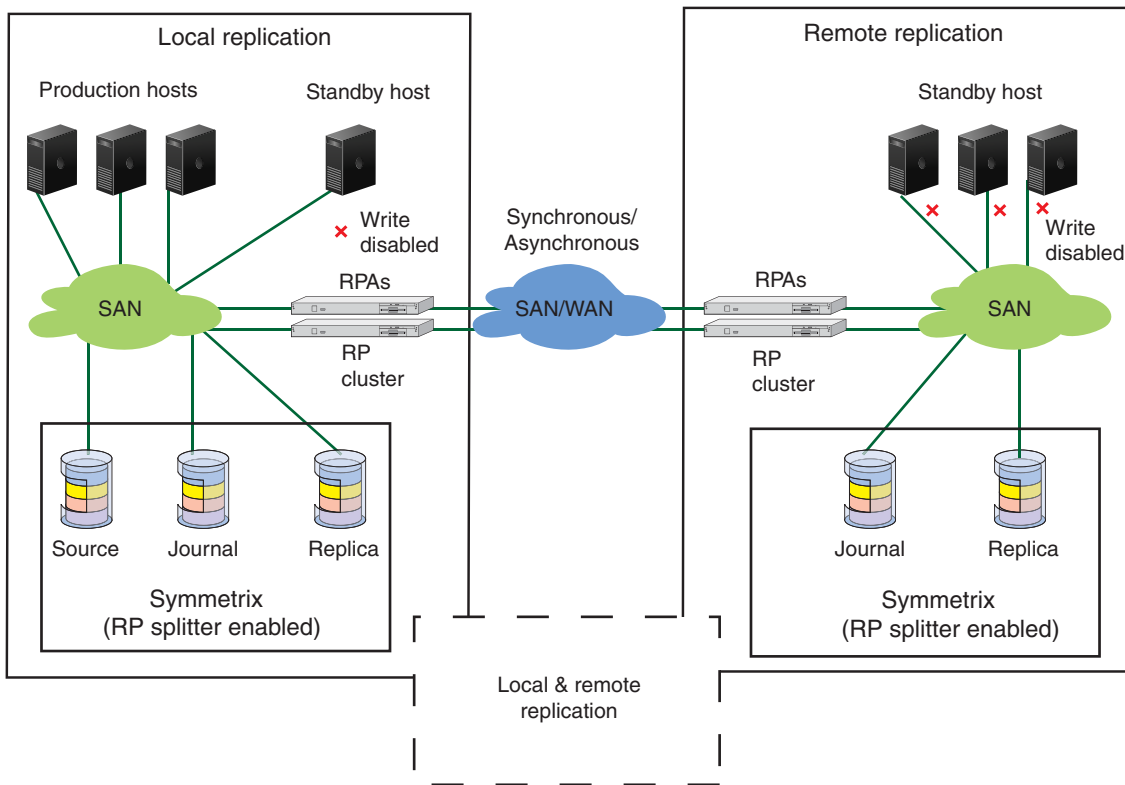
Introduced in 5875 for VMAX 10K, and 5876 for VMAX 20K and VMAX 40K, EMC RecoverPoint is a comprehensive data protection solution providing integrated continuous data protection and remote replication to recover applications at any point in time and provides disaster recovery. The integrated RecoverPoint splitter, which is enabled by default, supports 4,096 LUNs and is an enhanced implementation of Open Replicator for

Symmetrix which has been designed to operate as a RecoverPoint write I/O splitter. This RecoverPoint array-based splitter further enhances RecoverPoint’s existing array-based splitting options for Symmetrix systems. The RecoverPoint array-based splitter requires either:

- ◆ RecoverPoint/EX simplifies continuous data protection and replication using array-based write splitting
- ◆ RecoverPoint/CL is the full-featured offering that adds support for heterogeneous server and storage platforms

Local Data Corruption Protection, Disaster Recovery, Secondary Device Repurposing and Data Migrations are primary instances for using RecoverPoint, which provides:

- ◆ Continuous data protection (CDP) — Provides block-level local replication between LUNs in the same SAN using CDP technology that journals every write for later recovery to any point in time
- ◆ Continuous remote replication (CRR) — Provides dynamic synchronous and asynchronous block-level remote replication between LUNs in two different SANs using near-CDP technology that journals groups of writes for later recovery to significant points in time
- ◆ Concurrent local and remote data protection (CLR) — Provides both local replication (CDP) and remote (CRR) block level replication.



ICO-IMG-000969

Figure 8 RecoverPoint configurations

RecoverPoint systems support local replication of data that applications are writing over Fibre Channel to SAN-attached storage. The systems use existing Fibre Channel infrastructure to integrate seamlessly with existing host applications and data storage subsystems. For remote replication, the systems use existing IP connections to send the replicated data over a WAN, or use Fibre Channel infrastructure to replicate data asynchronously or synchronously. The systems provide failover of operations to a secondary site in the event of a disaster at the primary site.

Introduced in 5876, Symmetrix systems provides support for smaller to larger device size RecoverPoint replication enabling replication in heterogeneous environments where size may be different between devices.

Enhanced RecoverPoint thin LUN awareness supports only the allocated data to be synchronized during a full sweep. For example, if there is 10% of allocated data, RecoverPoint synchronizes only that amount of data.

RecoverPoint Appliance

The RecoverPoint Appliance contains four Fibre Channel ports. A minimum of two ports are connected to the fibre channel switch within your network with the connection completing on the front-end I/O module ports on the engine.

The *RecoverPoint Administrator's Guide* provides detailed information related to this feature.

Replication Manager

Replication Manager's support for EMC RecoverPoint provides continuous protection and near-instantaneous restoration of critical business data. It allows rapid recovery of the data on a volume to any point in time (crash-consistent recovery) or to a significant point in time (application-consistent recovery).

Replication Manager with RecoverPoint addresses the problem of lengthy recoveries of files and databases made by traditional backup software. It allows rapid reconstruction of a volume image that may be only seconds or minutes old.

Support for RecoverPoint complements Replication Manager's mirror- and snapshot-based recovery support by offering:

- ◆ Near-instantaneous restore time, providing full access to data made unavailable due to physical loss of media, data corruption, data pollution, or database or file system integrity problems
- ◆ Unlimited number of recovery points
- ◆ Practically no limit on the number of application-consistent checkpoints (as opposed to snapshot replicas)
- ◆ Application-consistent protection
- ◆ Instantaneous access to recovery points

SRDF and RecoverPoint

With 5876 Q2 2013 SR, SRDF and RecoverPoint CDP can co-exist on the same source device in a two-site SRDF solution, which provides you the ability to perform remote data protection with SRDF and logical data protection with RecoverPoint at the same time. The SRDF R1 devices are allowed to be tagged for RecoverPoint, and the RecoverPoint tagged devices are allowed to be configured as SRDF R1 devices.

This feature is supported for both static and dynamic SRDF devices, and is available in the following SRDF modes of operations:

- ◆ Synchronous
- ◆ Asynchronous
- ◆ Adaptive copy

The *EMC Solutions Enabler Symmetrix SRDF CLI Product Guide* provides details about controlling operations. The *EMC® Symmetrix® Remote Data Facility (SRDF®) Product Guide* provides additional details on requirements.

The *EMC® RecoverPoint Deploying with Symmetrix Arrays and Splitter Technical Notes* provides information, required procedures, and best practices for deploying RecoverPoint with Symmetrix arrays and splitter.

Performance considerations

Symmetrix systems use patented intelligent algorithms to manage data flow across host I/O channels, memory, and disks. Symmetrix systems also provide configuration tools that help consolidate storage resources while also providing the flexibility you need to manage mission-critical data and applications.

Performance planning considerations

To maximize Symmetrix system performance, review the following planning considerations:

- ◆ In multi-host environments, rank distributed workloads from the busiest to the least busy.
- ◆ Analyze workload and application requirements across storage tiers.
- ◆ Determine the data size requirements for each host connected to the Symmetrix system.
- ◆ Determine the number and type of front-end I/O connections (FICON, Fibre Channel, GbE, iSCSI), as well any requirements for multipathing.
- ◆ Analyze the nature of the applications executed on each host connected to the Symmetrix system.
- ◆ Determine the availability of a logical volume manager (LVM) on the hosts and the use of data striping.
- ◆ Determine and record the distribution of hypervolumes across the Symmetrix back-end, including:
 - Hypervolume size

- Allocation of hypervolumes between different hosts and applications
- Applications that use Virtual Provisioning
- Virtual Provisioning requirement.
- ◆ Determine the maximum drive and file system sizes supported by each host connected to the Symmetrix system.
- ◆ Define local and remote replication requirements.
- ◆ Analyze the requirements for device sharing.
- ◆ Define the number of Front End I/O Modules used and the number and type of ports on each engine.
- ◆ Define any special considerations for host-level mirroring for device distribution in the Symmetrix system.
- ◆ Plan for the possibility of upgrading the Symmetrix system with additional drives in the future, and the implications if the Symmetrix system installed is not at maximum capacity.

EMC Professional Services provides a suite of storage modeling applications to help you manage configurations and plan capacities and workloads.

Dynamic Cache Partitioning

Dynamic Cache Partitioning helps optimize storage because it allows the allocation of portions of cache to specific device groups.

Dynamic Cache Partitioning allows the definition of a maximum of eight cache partitioned groups, including the default group, to which all devices initially belong. With cache partitioning enabled, portions of cache can be allocated to a specific cache partition group.

Cache partition groups are monitored by the Symmetrix system to ensure that they do not consume more cache than they are allotted. Cache partitions can be:

- ◆ Dynamic — Allow the temporary donation of unused cache to other partitions after a specified donation time.
- ◆ Static partitions — Remain fixed in size, the donation parameter is ignored. These are often used for charge-back control.

Symmetrix Priority Controls

Symmetrix Priority Controls enhance tiered storage management by allowing prioritization of host application read I/O and SRDF/S transfers by assigning a priority level to device groups.

Symmetrix Priority Controls allow the definition of a maximum of 16 priority levels. Each user-defined priority level has a potential delay time that is determined by the chosen priority hierarchy, I/Os per second, and device groupings. Symmetrix systems implement prioritization of the I/O queue by using the Back End I/O Modules and Front End I/O Modules configured for SRDF. The Back End I/O Modules order read and write requests in the queue to prioritize what is sent to the disk. The Front End I/O Modules examine and prioritize SRDF/S transfer requests.

The task's priority level determines the task's position in the queue. During non-peak periods and periods of decreased utilization, all queued requests are satisfied in a timely manner even if they are assigned a low priority. It is only when the disk or SRDF/S transfer is "in demand" that service differentiation occurs.

Note: In mainframe environments, if workload manager (WLM) is active, each I/O has its own priority assigned by WLM. In such cases, Symmetrix Priority Controls honor the priorities assigned by WLM.

Host I/O limits

Enginuity 5876 Q4 2012 SR introduce host I/O limits (quotas) to be defined by the user and associated with storage groups. The quota definitions contain the operating parameters of the input/output (I/O) per second and/or bandwidth limitations. A defined quota is equally divided amongst the total number of directors included in the masking view associated with the storage group for which the quota is defined. All devices in that storage group share that quota.

When applications are configured, you can associate the limits to storage groups that contain a list of devices. A single storage group can only be associated to one limit and a device can only be in one storage group that has limits enabled. If the device is part of a hierarchy of storage groups and its parent already has an assigned quota, then the device cannot have a quota.

Enginuity 5876 Q2 2013 SR supports a quota increase to 2048 and also supports the following enhancements:

- ◆ Cascaded host I/O limits controlling parent and child storage groups limits in a cascaded storage group configuration.
- ◆ Support for front end host I/O limits on FCoE ports.
- ◆ Offline and failed director redistribution of quota that supports all available quota to be available instead of losing quota allocations from offline and failed directors.
- ◆ Dynamic host I/O limits support for dynamic redistribution of steady state unused director quota.

Cache performance

Symmetrix systems dramatically improve cache performance because they combine statistical prefetch algorithms, based on comprehensive analysis from thousands of installed Symmetrix systems, with dynamic intelligence that adjusts to immediate conditions on the Symmetrix system. EMC intelligent algorithms automatically adjust to the workload by constantly monitoring, evaluating, and optimizing cache decisions.

Storage systems use statistical prefetch algorithms to predict which information to place in cache before the host actually requests it. Inefficient prefetch algorithms have a significant impact on performance because they populate memory space with erroneous information, or fail to retrieve high-demand data that is required.

Disk optimization

Engenuity disk performance algorithms minimize the mechanical latencies and limitations of traditional hard disks. Disk optimization features include:

- ◆ Dynamic Mirror Service Policy — Reduces disk latencies in mirrored environments by reading from both mirrored pairs.
- ◆ Rotational position ordering — Reduces rotational latencies by scheduling I/Os based on physical drive location.
- ◆ 100 percent fast writes — Improves the performance of applications initially by writing to cache instead of disk.

Dynamic Mirror Service Policy

Symmetrix Dynamic Mirror Service Policy (DMSP) reduces rotational latency in mirrored environments by collecting latency statistics on both disks in the mirrored pair. Engenuity uses this information to improve performance by selecting a disk reading pattern that automatically and dynamically balances workload. The Dynamic Mirror Service Policy can choose:

- ◆ A split disk read pattern — Keeps each disk head over only half of the disk by splitting the disk into two halves, the inner cylinders and the outer cylinders. Disk one (primary) might service I/O corresponding to the inner part of the mirrored disk. Disk two (secondary) might handle I/O corresponding to the outer part of the mirrored disks.
- ◆ An interleaved read pattern — Allows both disks to alternately fetch tracks by reading logically assigned tracks. For example, even-numbered tracks could be serviced by disk one (primary) and odd-numbered tracks could be serviced by disk two (secondary).

Rotational positioning ordering

Symmetrix system rotational position ordering code schedules I/Os based on their physical location on the drive. Rotational position ordering is especially useful with large capacity drives, because large capacity drives are more likely to have several I/O queues.

100 percent fast writes

Symmetrix systems reduce latencies for mission-critical, write-intensive applications by performing all write operations in cache. When writing to cache in fast mode, Symmetrix systems send a “write complete” acknowledgment to the host as soon as the data is written and verified in cache, thereby releasing the application from the I/O operation without participating in the slow process of accessing the disk. The Symmetrix system then coalesces multiple writes to the same disk location, and safely destages the data to disk.

Tiered storage optimization

Tiered storage is a way to categorize the performance, functions, availability, and cost requirements necessary to support specific business applications. Tiered storage falls into four categories:

- ◆ Tier 3 — For applications like backup and archive, that do not require the performance, availability and functionality of higher performance tiers.
- ◆ Tier 2 — For important information, like order fulfillment and batch processing, that requires a high level of performance, availability or functionality, but may benefit from a lower cost solution.
- ◆ Tier 1— For mission-critical applications, like online ordering, that require high performance, high availability, and robust information safeguards.
- ◆ Tier 0 — For applications that require the ultimate performance and the lowest possible latency, as well as the same high availability and information safeguards that are also provided in tier 1.

Tier 0 storage is built on a generation of Flash drive technology that surpasses the limits imposed by magnetic disk drives. Symmetrix systems combine Flash drive technology with state-of-the-industry software, and provide high performance, flexible tier 0 storage options.

Note: Enginuity 5876 Q4 2012 SR supports up to four tiers, with the fourth tier being FTS, per FAST VP policy.

EMC XtremSW Cache

EMC® XtremSW Cache™ is intelligent caching software that leverages server-based flash technology to reduce latency and accelerate throughput for dramatic application performance improvement. XtremSWCache accelerates reads and protects data by using a write-through cache to the networked storage to deliver persistent high availability, integrity, and disaster recovery. XtremSWCache coupled with system-based EMC FAST software creates the most efficient and intelligent I/O path from the application to the data store. The result is a networked infrastructure that is dynamically optimized for performance, intelligence, and protection for both physical and virtual environments.

CHAPTER 4

Open Systems and Mainframe Support

Symmetrix systems support connectivity to environments with both mainframe and open systems hosts.

The Symmetrix Enterprise Storage Platform (ESP) option is required when both mainframe hosts (FICON) and open systems hosts (FCoE, Fibre Channel, iSCSI) connect to the same Symmetrix system. Contact your EMC representative for the current list of supported hosts, models, operating systems, and open systems support policies.

This chapter introduces the mainframe and open systems features supported on Symmetrix systems. Topics include:

- ◆ [Mainframe support on VMAX 20K and VMAX 40K](#) 57
- ◆ [Error reporting to the Mainframe host](#)..... 59
- ◆ [SIM Severity Reporting](#)..... 60
- ◆ [Environmental errors \(Engenuity 5876\)](#) 64
- ◆ [Open systems support on VMAX 10K, 20K, and 40K](#)..... 67

Mainframe support on VMAX 20K and VMAX 40K

Symmetrix VMAX 20K and VMAX 40K systems support mainframe features, including:

Note: The availability of certain features may be dependent on Engenuity versions.

• Compatible Native Flash (Flash Copy)	• MultiSubsystem Imaging
• Concurrent Copy	• Parallel Access Volumes
• Dynamic Channel Management (DCM)	• PDS Search Assist
• Dynamic Parallel Access Volumes/Multiple Allegiance (PAV/MA)	• Peer-to-Peer Remote Copy (PPRC)
• Extended Address Volumes (EAV)	• Persistent IU Pacing (Extended Distance FICON)
• High Performance FICON (zHPF)	• Parallel Access Volumes
• HyperPAV	• PDS Search Assist
• Modified Indirect Data Address Word (MIDAW)	• Peer-to-Peer Remote Copy (PPRC)
• Multiple Allegiance (MA)	• Sequential Data Striping
• Multitrack High Performance FICON (zHPF multitrack)	• z/OS Global Mirror
• Multi-Path Lock Facility/Concurrent Access (MPLF/CA)	• Virtual Provisioning
• Basic HyperSwap	

FICON configuration

FICON is a high-speed I/O interface that provides native point-to-point or switched point-to-point single or multimode connections from mainframes to storage devices. FICON connections support multiple concurrent I/O channels, I/O program multiplexing, and better link utilization than IBM's previous Fibre Channel standard.

Symmetrix systems auto-detect 2, 4 or 8 Gb/s FICON link speeds at switch or host port login. Symmetrix systems also provide support for Extended Distance FICON.

IBM 2107 support

Symmetrix systems provide support for the IBM 2107 command and feature set. When Symmetrix systems emulate an IBM 2107, they externally represent the serial number as an alphanumeric serial number in order to be compatible with IBM command output. Internally, Symmetrix systems retain a numeric serial number for IBM 2107 emulations. Correlation between the numeric and alphanumeric serial numbers is handled within Engenuity.

IBM i 512-byte D910 support

Symmetrix systems provide support for IBM i 512-byte 910 devices to provide a standard 512 byte device format that can natively attached to the IBM i host. The D910 is similar to the existing 2107 device type which also enables the EMC Virtual Provisioning features.

Note: Engenuity 5876 Q2 2013 SR provides support of thin IBM i 512-byte D910 devices on VMAX 10K systems.

Logical control unit capacities

[Table 12 on page 58](#) shows logical control unit (LCUs) capacities.

Table 12 Logical control unit maximums

Description	Maximum
LCUs per director slice (or port)	255 (within the range of 00 to FE)
LCUs per Symmetrix split ^a	255
Splits per Symmetrix system	16 (0 to 15)
Devices per split	65,280
LCUs per Symmetrix system	255
Devices per LCU	256
Logical paths per port	2,048
Logical paths per LCU per port	128
Symmetrix system host addresses per Symmetrix system (base and alias)	65,280
I/O host connections per Symmetrix system engine	8

- a. A Symmetrix split is a logical partition of the Symmetrix system, identified by unique Symmetrix devices, SSIDs, and host serial number.

Disk drive emulations

When Symmetrix systems are configured to mainframe hosts, the Symmetrix system drives emulate IBM ECKD Direct Access Storage Device (DASD).

Cascading configurations

Cascading configurations greatly enhance FICON connectivity between local and remote sites by using switch-to-switch extensions of the CPU to the FICON network. These cascaded switches communicate over long distances using a small number of high speed lines called interswitch links (ISLs). A maximum of two switches may be connected together within a path between the CPU and the Symmetrix system.

Same switch vendors are required for a cascaded configuration. To support cascading, each switch vendor requires specific models, hardware features, software features, configuration settings, and restrictions. Specific IBM CPU models, MVS release levels, host hardware, and Enginuity levels are also required.

For the most up-to-date information about switch support, consult the EMC Support Matrix (ESM), available through E-Lab™ Interoperability Navigator (ELN) at <http://elabnavigator.emc.com> under the PDFs and Guides tab.

Error reporting to the Mainframe host

The Symmetrix Enginuity operating environment can detect the following error types to the mainframe host in the Symmetrix storage systems:

- ◆ **Data Check** — Enginuity has detected an error in the bit pattern read from the disk. Data checks are due to hardware problems when writing or reading data, media defects, or random events.
- ◆ **System or Program Check** — Enginuity has rejected the command. This type of error is indicated to the processor and is always returned to the requesting program.
- ◆ **Overrun** — Enginuity cannot receive data at the rate it is transmitted from the host. This error indicates a timing problem. Resubmitting the I/O operation usually corrects this error.
- ◆ **Equipment Check** — Enginuity has detected an error in hardware operation.
- ◆ **Environmental** — Enginuity internal test has detected an environmental error. Internal environmental tests monitor, check, and report failures of the critical hardware components. They run at the initial system power-up, upon every software reset event, and at least once every 24 hours during regular operations.

If an environmental test detects an error condition, it sets a flag to indicate a pending error and presents a unit check status to the host on the next I/O operation. The test that detected the error condition is then scheduled to run more frequently. If a device-level problem is detected, it is reported across all logical paths to the device experiencing the error. Subsequent failures of that device are not reported until the failure is fixed.

If a second failure is detected for a device while there is a pending error-reporting condition in effect, Enginuity reports the pending error on the next I/O and then the second error.

Engenuity reports error conditions to the host and to the EMC Customer Support Center. When reporting to the host, Engenuity presents a unit check status in the status byte to the channel whenever it detects an error condition such as a data check, a command reject, an overrun, an equipment check, or an environmental error.

When presented with a unit check status, the host retrieves the sense data from the Symmetrix system and, if logging action has been requested, places it in the Error Recording Data Set (ERDS). The EREP (Environment Recording, Editing, and Printing) program prints the error information. The sense data identifies the condition that caused the interruption and indicates the type of error and its origin. The sense data format depends on the mainframe operating system. For 2105, 2107, or 3990 controller emulations, the sense data is returned in the SIM format.

SIM Severity Reporting

Engenuity supports SIM Severity Reporting that enables you to filter SIM severity alerts reported to the multiple virtual storage (MVS) console. [Table 13 on page 60](#) lists the default settings for SIM Severity Reporting:

- ◆ SIM **ACUTE**, **SERIOUS**, and **MODERATE** Alerts are by default reported to the MVS console.
- ◆ All SIM severity alerts are by default reported to the EREP.

Table 13 SIM severity alerts

Severity	Description
SERVICE	No system or application performance degradation is expected. No system or application outage has occurred.
MODERATE	Performance degradation is possible in a heavily loaded environment. No system or application outage has occurred.
SERIOUS	A primary I/O subsystem resource is disabled. Significant performance degradation is possible. System or application outage may have occurred.
ACUTE	A major I/O subsystem resource is disabled, or damage to the product is possible. Performance may be severely degraded. System or application outage may have occurred.
REMOTE SERVICE	EMC Customer Support Center is performing service/maintenance operations on the system.
REMOTE FAILED	Service Processor cannot communicate with the EMC Customer Support Center.

Environmental errors (Enginuity 5874 and 5875)

Table 14 on page 61 lists Enginuity 5874 and 5875 environmental errors in SIM format.

Table 14 Environmental Errors Reported as SIM messages

Hex code	Severity level	Description	SIM reference code
042F	MODERATE	Send a SIM alert for error code 052F (a sync SRDF write failure occurred). Note: Epack required for 5874 only; fix 58335 and fix 20011029.	E42F
043E	MODERATE Note: Severity level can be modified via SymmWin.	An SRDF Consistency Group was suspended.	E43E
0454	SERVICE	SRDF switched to Adaptive Copy write pending mode due too numerous suspend or halt chains. This condition happens in SRDF configurations when Enginuity detects a mainframe memory paging I/Os to the SRDF devices. This error triggers a call home to the EMC Customer Support Center.	E454
0461	SERVICE	The M2 is resynchronized with the M1 device. This event occurs once the M2 device is brought back to a Ready state. ENGINEERING RECOMMENDATION is to leave disabled.	E461
0462	SERVICE	The M1 is resynchronized with the M2 device. This event occurs once the M1 device is brought back to a Ready state. ENGINEERING RECOMMENDATION is to leave disabled.	E462
0463	SERIOUS	One of the back-end directors failed into the IMPL Monitor state.	2463
0465	SERVICE	Device resynchronization process has started. ENGINEERING RECOMMENDATION is to leave disabled.	E465
0467	MODERATE Note: Severity level can be modified via SymmWin.	The remote Symmetrix system reported an SRDF error across the SRDF links.	E467
046B	SERVICE	The Event Trace utility has been running for more than 30 days.	E46B
046D	MODERATE Note: Severity level can be modified via SymmWin.	An SRDF group is lost. This event happens, for example, when all SRDF links fail.	E46D
046E	SERVICE Note: Severity level can be modified via SymmWin.	An SRDF group is up and operational.	E46E
0471	SERIOUS	The DSE, DATA, or SAVE device is too small or the session is too long.	2471

Table 14 Environmental Errors Reported as SIM messages (continued)

Hex code	Severity level	Description	SIM reference code
0473	SERIOUS <hr/> Note: Severity level can be modified via SymmWin. <hr/>	A periodic environmental test (env_test9) detected the mirrored device in a Not Ready state.	E473
0474	SERIOUS <hr/> Note: Severity level can be modified via SymmWin. <hr/>	A periodic environmental est (env_test9) detected the mirrored device in a Write Disabled (WD) state.	E474
0475	SERIOUS <hr/> Note: Severity level can be modified via SymmWin. <hr/>	An SRDF R1 remote mirror is in a Not Ready state.	E475
0476	SERVICE	Service Processor has been reset.	2476
0477	REMOTE FAILED	The Service Processor could not call the EMC Customer Support Center (failed to call home) due to communication problems.	1477
0478	MODERATE	Check if a FLASH programming lock is held.	2478
047C	ACUTE	A SAVE device with user data is Not Ready.	247C
047D	MODERATE <hr/> Note: Severity level can be modified via SymmWin. <hr/>	Either the SRDF group lost an SRDF link or the SRDF group is lost locally.	E47D
047E	SERVICE <hr/> Note: Severity level can be modified via SymmWin. <hr/>	An SRDF link recovered from failure. The SRDF link is operational.	E47E
047F	REMOTE SERVICE	The Service Processor successfully called the EMC Customer Support Center (called home) to report an error.	147F
0492	SERVICE	The available disk space on the Service Processor is low.	2492
01BA 02BA 03BA 04BA	ACUTE	A problem with the Engine power supply or the Engine SPS module.	24BA

Table 14 Environmental Errors Reported as SIM messages (continued)

Hex code	Severity level	Description	SIM reference code
04CA	MODERATE <hr/> Note: Severity level can be modified via SymmWin. <hr/>	An SRDF/A session dropped due to a non-user request. Possible reasons include fatal errors, SRDF link loss, or reaching the maximum SRDF/A host-response delay time.	E4CA
04D1	REMOTE SERVICE	Remote connection established or remote control connected.	14D1
04D1	REMOTE SERVICE	Remote connection established. Remote control connected.	14D1
04D2	REMOTE SERVICE	Remote connection closed. Remote control rejected.	14D2
04D3	SERVICE	Flex filter problems.	24D3
04D4	REMOTE SERVICE	Remote connection closed. Remote control disconnected.	14D4
04DA	SERVICE	Problems with task/threads.	24DA
04DB	SERVICE	SYMPL script generated error.	24DB
04DC	SERVICE	Service Processor related problems.	24DC
04E0	REMOTE FAILED	Communications problems.	14E0
04E1	SERVICE	Problems in error polling.	24E1
04F9	MODERATE	An SRDF Consistency Group was suspended.	E4F9

Environmental errors (Enginuity 5876)

Table 15 on page 64 list the Enginuity 5876 or higher environmental errors in SIM format.

Note: All Severity levels listed can be modified via SymmWin.

Table 15 Environmental errors reported as SIM messages

Hex code	Severity level	Description	SIM reference code
042F	NONE	Send a SIM Alert for error code 052F (a sync SRDF write failure occurred). EMC recommendation: NONE.	E42F
043E	MODERATE	An SRDF Consistency Group was suspended.	E43E
0454	SERVICE	SRDF switched to Adaptive Copy write pending mode due to numerous suspend or halt chains. This condition happens in SRDF configurations when Enginuity detects a mainframe memory paging I/Os to the SRDF devices. This error triggers a call home to the EMC Customer Support Center.	E454
0461	NONE	The M2 is resynchronized with the M1 device. This event occurs once the M2 device is brought back to a Ready state. EMC recommendation: NONE	E461
0462	NONE	The M1 is resynchronized with the M2 device. This event occurs once the M1 device is brought back to a Ready state. EMC recommendation: NONE.	E462
0463	SERIOUS	One of the back-end directors failed into the IMPL Monitor state.	2463
0465	NONE	Device resynchronization process has started. EMC recommendation: NONE.	E465
0467	MODERATE	The remote Symmetrix system reported an SRDF error across the SRDF links.	E467
046B	SERVICE	The Event Trace utility has been running for more than 30 days.	E46B
046D	MODERATE	An SRDF group is lost. This event happens, for example, when all SRDF links fail.	E46D
046E	SERVICE	An SRDF group is up and operational.	E46E
0470	ACUTE	OverTemp condition based on memory module temperature.	2470
0471	SERIOUS Note: Severity changes from SERIOUS to ACUTE at 98% full capacity.	The DSE, DATA, or SAVE device is too small or the session is too long.	2471
0473	SERIOUS	A periodic environmental test (env_test9) detected the mirrored device in a Not Ready state.	E473
0474	SERIOUS	A periodic environmental est (env_test9) detected the mirrored device in a Write Disabled (WD) state.	E474
0475	SERIOUS	An SRDF R1 remote mirror is in a Not Ready state.	E475

Table 15 Environmental errors reported as SIM messages

Hex code	Severity level	Description	SIM reference code
0476	SERVICE	Service Processor has been reset.	2476
0477	REMOTE FAILED	The Service Processor could not call the EMC Customer Support Center (failed to call home) due to communication problems.	1477
047A	ACUTE	AC Power lost to Power Zone A or B.	247A
01BA 02BA 03BA 04BA	ACUTE	A problem with the engine power supply or the engine SPS module.	24BA
047C	ACUTE	A SAVE device with user data is Not Ready.	247C
047D	MODERATE	Either the SRDF group lost an SRDF link or the SRDF group is lost locally.	E47D
047E	SERVICE	An SRDF link recovered from failure. The SRDF link is operational.	E47E
047F	REMOTE SERVICE	The Service Processor successfully called the EMC Customer Support Center (called home) to report an error.	147F
0492	SERVICE	The available disk space on the Service Processor is low.	2492
04CA	MODERATE	An SRDF/A session dropped due to a non-user request. Possible reasons include fatal errors, SRDF link loss, or reaching the maximum SRDF/A host-response delay time.	E4CA
04D1	REMOTE SERVICE	Remote connection established. Remote control connected.	14D1
04D2	REMOTE SERVICE	Remote connection closed. Remote control rejected.	14D2
04D3	SERVICE	Flex filter problems.	24D3
04D4	REMOTE SERVICE	Remote connection closed. Remote control disconnected.	14D4
04DA	SERVICE	Problems with task/threads.	24DA
04DB	SERVICE	SYMPL script generated error.	24DB
04DC	SERVICE	Service Processor related problems.	24DC
04E0	REMOTE FAILED	Communications problems.	14E0
04E1	SERVICE	Problems in error polling.	24E1
04F9	MODERATE	An SRDF Consistency Group was suspended.	E4F9

Operator messages

On z/OS, SIM messages will be displayed as IEA480E Service Alert Error messages. They have the format shown in [Figure 9 on page 66](#) and [Figure 10 on page 66](#).

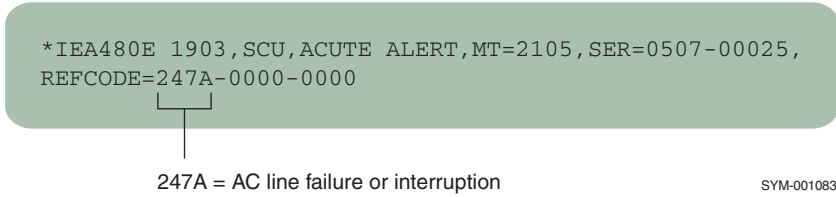


Figure 9 z/OS IEA480E service alert error message format (AC power failure)

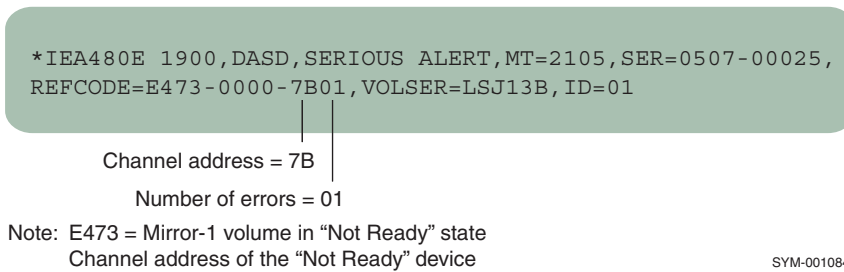


Figure 10 z/OS IEA480E service alert error message format (mirror-1 volume in “not ready” state)

Note: All host channel paths to that device (target volume) will report this error message. Therefore, this message may appear several times.

The Symmetrix system also reports events to the host and to the service processor. These events are:

- ◆ The mirror-2 volume has synchronized with the source volume.
- ◆ The mirror-1 volume has synchronized with the target volume.
- ◆ Device resynchronization process has begun.

On z/OS, these events are displayed as IEA480E Service Alert Error messages. They have the format shown in [Figure 11 on page 66](#) and [Figure 12 on page 67](#).

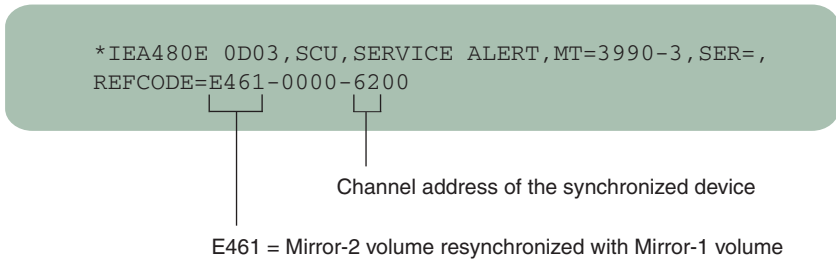


Figure 11 z/OS IEA480E service alert error message format (mirror-2 resynchronization)

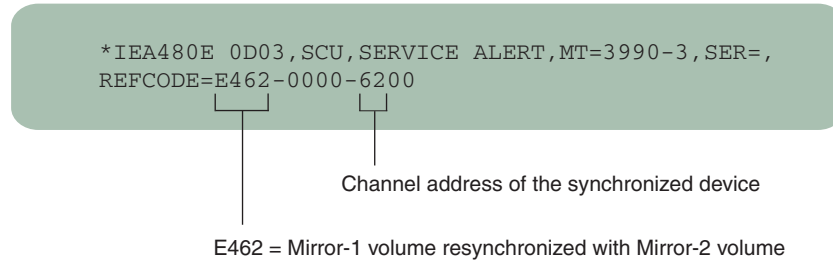


Figure 12 z/OS IEA480E service alert error message format (mirror-1 resynchronization)

Open systems support on VMAX 10K, 20K, and 40K

Symmetrix VMAX 10K, VMAX 20K and VMAX 40K systems provide the drive emulation and volume manager compatibility required for open systems environments. On open systems hosts, Symmetrix system logical disk devices appear as standard SCSI drives with data stored in fixed-block architecture format. Any logical device manager software installed on a host can be used with Symmetrix system disk devices.

Symmetrix systems provide Fibre Channel connections to IBM i hosts. Fibre Channel provides up to 8 Gb/s full duplex with IBM i.

EMC Symmetrix systems support open systems features, including:

- EMC Open Migrator/Live Migration (LM)
- FAST VP
- Federated Live Migration
- Federated Tiered Storage
- IBM i host connections
- Open Replicator for Symmetrix
- RecoverPoint
- Virtual Provisioning

Open systems intermix configurations

Intermix configurations on open systems allow the definition of separate FICON zones on new or existing open systems switches. These switches can also be cascaded to enhance connectivity and facilitate remote backup and recovery. The EMC and IBM branded McDATA and INRANGE switches are supported.

To support open systems intermix configurations, each switch vendor requires specific models, hardware and software features, configuration settings, and restrictions. Specific IBM CPU models, MVS release levels, channel hardware, and Engenuity levels are also required.

For the most up-to-date information about switch support, consult the EMC Support Matrix (ESM), available through E-Lab Interoperability Navigator (ELN) at <http://elabnavigator.emc.com>, under the PDFs and Guides tab.

CHAPTER 5

Data Integrity and Protection

Symmetrix systems provide a suite of integrity checks, RAID options, and vaulting capabilities to ensure data integrity and to protect data in the event of a system failure or power outage.

This chapter describes Symmetrix system data integrity protections. Topics include:

- ◆ [Data Integrity Checking](#) 69
- ◆ [Drive integrity monitoring and correction](#)..... 70
- ◆ [RAID options](#)..... 70
- ◆ [Vaulting](#) 71

Data Integrity Checking

Symmetrix systems validate the integrity of data it holds at every possible point during the lifetime of the data. From the point at which data enters a Symmetrix system, the data is continuously protected by error detection metadata. This protection metadata is checked by hardware and software mechanisms any time data is moved within the Symmetrix subsystem, allowing the Symmetrix system to provide true end-to-end integrity checking and protection against hardware or software faults.

The protection metadata is appended to the data stream, and contains information describing the expected data location as well as cyclic redundancy code (CRC) representation of the actual data contents. The expected values to be found in protection metadata are stored persistently in an area separate from the data stream. The protection metadata is used to validate the logical correctness of data being moved within the Symmetrix system any time the data transitions between protocol chips, internal buffers, internal data fabric endpoints, system cache, and system disks.

Physical memory error correction and error verification

Symmetrix systems correct single-bit errors and report an error code once the single-bit errors reach a predefined threshold. When a multibit error occurs, the Symmetrix system fences the physical memory segment (removes it from service) and retrieves the data from mirrored memory (if it was unwritten) from the drive. In the unlikely even that physical memory replacement is required, the Symmetrix system notifies EMC support, and a replacement is ordered.

Block CRC

Symmetrix systems support industry standard block CRC based on T10 Data Integrity Field (DIF) that enables host generated CRC's to be stored and used with the data. Symmetrix systems provide additional protections for address/control fault modes that provide increased levels of protection against address/control faults and provides this in user definable blocks provided by the T10 standard.

Symmetrix systems provide address information and write status information in the extra bytes in the application tag and reference tag portion of the block CRC.

Drive integrity monitoring and correction

Symmetrix systems monitor medium defects by both examining the result of each disk data transfer and proactively scanning the entire disk during idle time. If a block on the disk is determined to be bad, the director:

1. Rebuilds the data in the physical memory if necessary.
2. Remaps the defect block to another area on the disk set aside for this purpose.
3. Rewrites the data from physical memory back to the remapped block on the disk.

The director maps around any bad block(s) detected, thereby avoiding defects in the media. The director also keeps track of each bad block detected on a disk. If the number of bad blocks exceeds a predefined threshold, the Service Processor invokes a sparing operation to replace the defective drive and then automatically alerts EMC Customer Support to arrange for corrective action.

RAID options

Symmetrix systems support varying levels of protection including RAID 1, RAID 10, RAID 5 (3+1 and 7+1), and RAID 6 (6 + 2 and 14 + 2), as shown in [Table 16 on page 70](#). RAID protection options are configured at the physical drive level. Different levels of RAID protection are easily configured with different datasets within a Symmetrix system. EMC strongly recommends that you use one or more of the RAID data protection schemes for all data devices.

Table 16 RAID options (page 1 of 2)

RAID	System	Provides the following	Configuration considerations
RAID 1	All	<p>The highest level of performance for all mission-critical and business-critical applications. Maintains a duplicate copy of a device on two drives:</p> <ul style="list-style-type: none"> • If a drive in the mirrored pair fails, the Symmetrix system automatically uses the mirrored partner without interruption of data availability. • When the drive is (nondisruptively) replaced by a sparing operation, the Symmetrix system re-establishes the mirrored pair and automatically re-synchronizes the data with the drive. 	<p>Withstands failure of a single drive within the mirrored pair.</p> <p>A drive rebuild is a simple copy from the remaining drive to the replaced drive.</p> <p>The number of required drives is twice the amount required to store data (usable storage capacity of a mirrored system is 50%).</p>

Table 16 RAID options (page 2 of 2)

RAID	System	Provides the following	Configuration considerations
RAID 10	VMAX 20K and VMAX 40K	A combination of RAID 1 (mirrored) and RAID 0 (striping with no data protection) used for mainframe environments. RAID 10 allows four Symmetrix system devices (each one-fourth the size of the original IBM device) to appear as one IBM device to the host, accessible by way of one host address. Any four devices can be chosen to define a group provided they are equally sized, the same type (for example, all 3390), and have the same mirror configuration.	Withstands failure of multiple drives within one of the mirrored pairs. A drive rebuild is a simple copy from the remaining drive to the replaced drive. The number of required disks is twice the amount required to store data (usable data storage capacity of a mirrored system is 50%).
RAID 5	All	Distributed parity and striped data across all drives in the RAID group. Options include: <ul style="list-style-type: none"> RAID 5 (3 + 1) — Consists of four drives with parity and data striped across each device. RAID 5 (7 + 1) — Consists of eight drives with data and parity striped across each device. 	RAID 5 (3 + 1) provides 75% data storage capacity. RAID 5 (7 + 1) provides 87.5% data storage capacity. Withstands failure of a single drive within the RAID 5 group.
RAID 6	All	Striped drives with double distributed parity (horizontal and diagonal). The highest level of availability options include: <ul style="list-style-type: none"> RAID 6 (6 + 2) — Consists of eight drives with dual parity and data striped across each device. RAID 6 (14 + 2) — Consists of 16 drives with dual parity and data striped across each device. 	RAID 6 (6 + 2) provides 75% data storage capacity. RAID 6 (14 + 2) provides 87.5% data storage capacity. Withstands failure of two drives within the RAID 6 group.

Vaulting

Symmetrix system vault operations protect data during a manual power down or an unexpected power outage.

Each Symmetrix system comes with Standby Power Supply (SPS) modules. If power is lost, the Symmetrix system uses SPS power to write the data in physical memory onto designated disk storage called Power Vault devices. Vaulted images are fully redundant; the contents of physical memory are saved twice to independent disks.

The vault operation

The Symmetrix systems initiate a vault operation if the system is powered down, transitions offline, or when environmental conditions, such as the loss of a data center due to an air conditioning failure occur. A vault operation is initiated:

- ◆ During the first (save) part of the vault operation, the Symmetrix system stops all I/O. When physical memory reaches a consistent state, directors write the contents to the vault devices, saving two copies of the data. The Symmetrix system then completes the power down, or, if power down is not required, remains in the offline state.
- ◆ During the second (restore) part of the operation, the Symmetrix system startup program initializes the hardware and the environmental system, restores physical memory contents from the saved data (while checking data integrity) and re-initializes the physical memory.

The system resumes normal operation when the SPSeS are sufficiently recharged to support another vault. If any condition is not safe, the system will not resume operation and will notify Customer Support for diagnosis and repair. This allows Customer Support to communicate with the Symmetrix system and restore normal system operations.

Vault configuration considerations

[Table 17](#) lists the amount of dedicated vault space and the dedicated devices for vault space per engine.

Table 17 Space dedication

System	Dedicated vault space	Dedicated devices for vault space per engine
VMAX 10K	9 GB	20, total of 180 GB of vault space
VMAX 20K	5 GB	40, total of 200 GB of vault space
VMAX 40K	9 GB	40, total of 360 GB of vault space

The following configuration rules apply to the vault drives and directors:

- ◆ Vault devices are designated disk drives that reserve a dedicated amount space for vault data, including metadata.
- ◆ All drives types can be used for vault.
- ◆ The vault space is for internal use only. No other device can reside in this space.
- ◆ Five vault drives per loop are required to enable sparing. Vault drives are eligible for permanent sparing and direct sparing.
- ◆ The total capacity of all vault devices in the system will be at least sufficient to keep two logical copies of the persistent portion of physical memory.

CHAPTER 6

Data Migration Tools

Symmetrix systems provide the ability to perform the one-time movement of data from a source to a target using various data migration software. Once the migration operation is performed, applications will now access the data from the new location. When migrating data, one should consider the potential business impact and the type of data to be migrated, the site location(s), the number of systems and applications, and the total amount of data, as well as time considerations and schedules.

Symmetrix systems support several tools and services for data migration:

- ◆ [“Federated Tiered Storage” on page 73](#) — Allows the creation of external storage using EMC or third-party storage systems.
- ◆ [“Federated Live Migration” on page 76](#) — Allows data movement from source Symmetrix systems to a target Symmetrix system.
- ◆ [“Open Replicator for Symmetrix” on page 76](#) — Creates copies of data for migration between Symmetrix systems and other qualified storage systems.
- ◆ [“Open Migrator/Live Migration” on page 77](#) — Provides online data migration between the source and the target storage systems. The migration process is non-disruptive to production applications.
- ◆ [“PowerPath Migration Enabler” on page 77](#) — Enables non-disruptive or minimally disruptive data migration between storage systems or between logical units within a single storage system.
- ◆ SRDF/Data Mobility (SRDF/DM) — Enables fast data transfer from R1 to R2 devices over extended distances.
- ◆ [“Unisphere for VMAX” on page 78](#) — Supports browser-based user interface for the configuration and management of Symmetrix systems.

Federated Tiered Storage

Federated Tiered Storage (FTS) allows the creation of external storage using EMC or third-party storage systems. FTS enables data mobility between the Symmetrix system and the external systems, and data mobility between two external systems, using existing VLUN migration technology. FTS does not require any new hardware and uses available SAN ports (also usable as FAs and RFs) and is based on new Enginuity emulation functionality, known as DX (Disk Adaptor eXternal). [“External connections” on page 24](#) provides additional information.

Federated Tiered Storage (FTS) gives you the ability to attach supported, SAN-attached external systems to a Symmetrix system, as show in [Figure 13 on page 74](#). Enginuity 5876 Q4 2012 SR supports direct attach from your external storage system to a Symmetrix system, as shown in [Figure 14 on page 75](#), using Fibre Channel arbitrated loops between the external port and the target port.

Note: The maximum external capacity is determined by the VMAX cache.

FTS allows LUNs which exist on external systems to be used to provide physical disk space on existing systems while gaining access to Symmetrix features such as local replication, remote replication, storage tiering, data management, and data migration. These external LUNs can be used as raw storage space for the creation of Symmetrix devices in the same way internal Symmetrix physical drives are used. These devices are known as eDisks.

Figure 13 on page 74 shows the configuration of the network, storage systems, and application hosts using FTS.

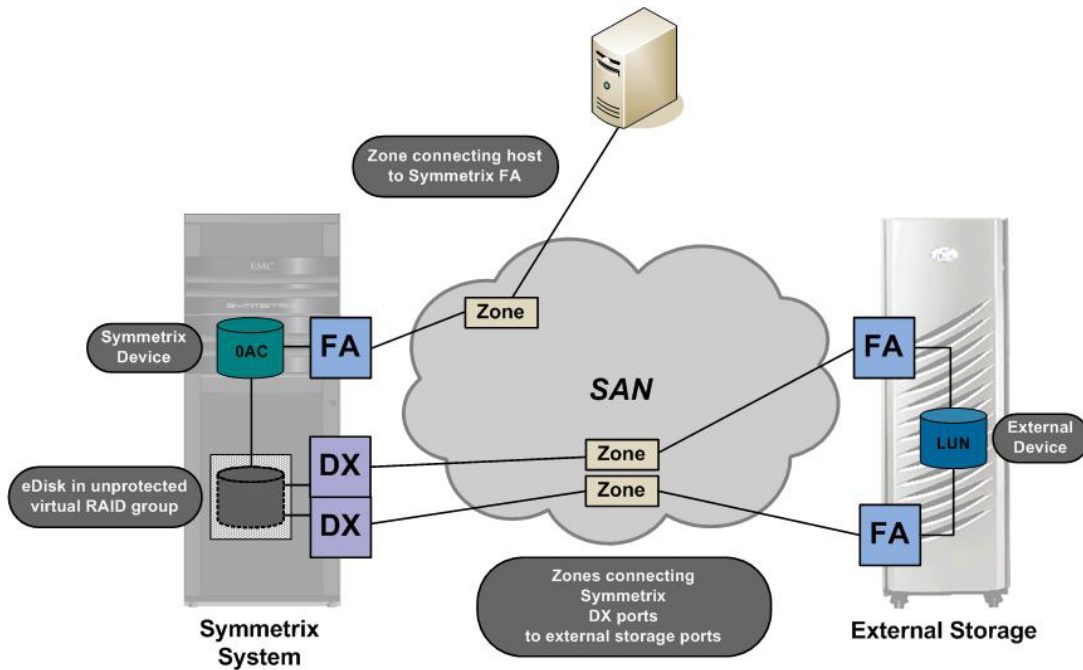


Figure 13 FTS configuration - fabric

Figure 14 shows direct attach from your external storage system to a Symmetrix system.

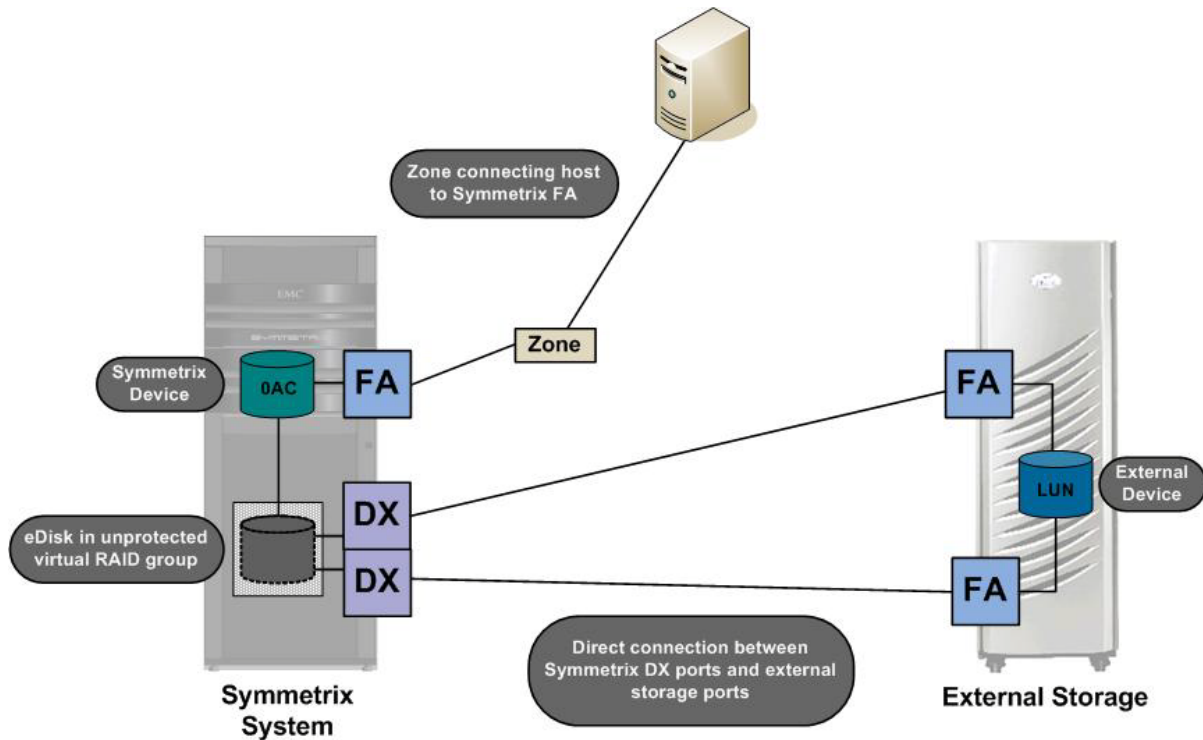


Figure 14 FTS configuration - arbitrated loops

eDisks

When you attach external storage to a Symmetrix system, FTS virtualizes an external system's SCSI logical units as Symmetrix disks called eDisks.

The following restrictions apply to eDisks:

- Can only be unprotected devices. The RAID protection scheme of eDisks is dependent on the external system.
- Cannot be IBM i, CKD, or gatekeeper devices.
- Cannot be used as VAULT, SFS, or ACLX devices.

Modes of operation

FTS has two modes of operation, depending on whether the external LUN is used as raw storage space or has data that must be preserved and accessed through a Symmetrix device.

- ◆ External Provisioning that allows you to access LUNs existing on external storage as raw capacity for new Symmetrix devices. These devices are called externally provisioned devices. When you virtualize an eDisk for external provisioning, you can then create Symmetrix devices from the external disk group and present the storage to users. You can also use this storage to create a new FAST VP tier. See [“Fully Automated Storage Tiering for Virtual Pools” on page 39](#) for additional details.

- ◆ Encapsulation which allows you to preserve existing data on external LUNs and access it through Symmetrix volumes. These devices are called encapsulated devices. There are two different options with encapsulation:
 - **Standard encapsulation** - The eDisk is encapsulated and exported from the Symmetrix system as disk group provisioned devices.
 - **Virtual provisioning encapsulation** - Data devices are created and added to a specified thin pool. Fully allocated thin devices are also created and bound to the pool. Extents are allocated to the external LUN through the data devices.

In either case, Enginuity automatically creates the necessary Symmetrix devices. If the eDisk is larger than the maximum Symmetrix device size or the configured minimum auto meta size, Enginuity creates multiple Symmetrix devices to account for the full size of the eDisk. These Symmetrix devices are concatenated into a single concatenated metadvice to allow access to the complete volume of data available from the eDisk.

FTS Interoperability Matrix at https://elabnavigator.emc.com/jsp/legacy_support.jsp and the *EMC Simple Support Matrix EMC Federated Tiered Storage (FTS)* provide greater detail on FTS support of external systems. The *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide* provides additional details.

Federated Live Migration

Federated Live Migration (FLM) allows volumes to be migrated non-disruptively, without loading software on any connected host and without host interruption to load virtualization software. This allows a new Symmetrix system to displace existing systems without needing to coordinate downtimes with host or application administrators.

Open Replicator for Symmetrix

EMC Open Replicator for Symmetrix (ORS) provides a method for copying data from a qualified storage system within a storage area network (SAN) to or from a another system. Open Replicator for Symmetrix provides the ability to migrate data from EMC systems, CLARiiON systems, and qualified third party systems to/from a Symmetrix system, as well as the ability to migrate data from a Symmetrix system to other storage systems within the SAN. When combined with PowerPath Migration Enabler, Open Replicator for Symmetrix provides a nondisruptive migration solution.

Open Replicator for Symmetrix features include:

- ◆ Introduced in Enginuity 5876, ORS is supported over FCoE where FCoE ports are enabled to act as a SCSI initiator.
- ◆ Heterogeneous data mobility — Migrates and copies data between Symmetrix systems and other supported third-party systems.
- ◆ Online migrations — Migrates data with minimum disruption to host applications.
- ◆ Integration with EMC replication software — Allows Open Replicator for Symmetrix to protect lower-tier applications at remote locations.
- ◆ SAN/WAN capabilities — Allows data migration to qualified storage systems over SAN/WAN topologies.

- ◆ Consistency support — Ensures immediate data availability to host applications using Symmetrix system consistency technology.
- ◆ Full or incremental copying — Accelerates backups by pushing a live source Symmetrix system volume to a target volume on a qualified storage system.

Open Migrator/Live Migration

EMC Open Migrator/Live Migration (LM) provides online data migration of Microsoft Windows, UNIX, or Linux devices between any source and EMC storage. Migration is performed from the production host.

With Open Migration/LM, critical applications are online and available during operations like server consolidation, storage upgrades, and performance tuning. Open Migrator/LM provides the mirroring and background copy functions that are used to synchronize data images on one or more source and target devices, LUNs, or LUN partitions. Data migration is supported between source and target devices of any device block type. During migration the source volume is available for I/O to production applications on the host.

PowerPath Migration Enabler

EMC PowerPath Migration Enabler (PPME) is a migration tool that enables non-disruptive or minimally disruptive data migration between storage systems or between logical units within a single storage system. Migration Enabler resides on the host and allows applications continued data access throughout the migration process. Migration Enabler works independently of PowerPath multipathing, but PowerPath must be fully installed.

PowerPath Migration Enabler integrates with other technologies to minimize or eliminate application downtime while migrating data.

SRDF/Data Mobility

SRDF/Data Mobility (SRDF/DM) permits operation in SRDF adaptive copy mode only and is designed for data replication or migration between two or more Symmetrix systems. SRDF/DM transfers data from primary volumes to secondary volumes permitting information to be shared, content to be distributed, and access to be local to additional processing environments. Adaptive copy mode enables applications using that volume to avoid propagation delays while data is transferred to the remote site. SRDF/DM supports all Symmetrix systems and all Enginuity levels that support SRDF, and can be used for local or remote transfers.

The *EMC® Symmetrix® Remote Data Facility (SRDF®) Product Guide* provides greater detail.

Unisphere for VMAX

EMC Unisphere for VMAX is a simple, intuitive, browser-based user interface for the configuration and management of Symmetrix systems running Enginuity 5876 and higher.

Unisphere for VMAX presents the functionality of the Solutions Enabler SYMCLI in a browser interface. You can use Unisphere to:

- ◆ Manage Symmetrix access controls, user accounts, and roles
- ◆ Discover Symmetrix systems
- ◆ Perform configuration operations (create volumes, map and mask volumes, set Symmetrix system attributes, set volume attributes, set port flags, create SAVE pools)
- ◆ Manage volumes (change volume configuration, set volume status, reserve volumes, duplicate volumes, create/dissolve meta volumes)
- ◆ Manage Fully Automated Storage Tiering (FAST), Virtual Provisioning, and Auto-provisioning Groups
- ◆ Perform and monitor replication operations
- ◆ Specify and monitor alerts
- ◆ Monitor an application's performance and analyze system performance

Unisphere for VMAX is preinstalled on the service processor (but without the performance section, and can also run on a data center host.

In Enginuity 5876 Federated Live Migration includes support for clustered host configurations. Enginuity 5876 Q4 2012 SR introduces new host environments and support of additional arrays. Federated Live Migration supports migrations with zero space reclamation.

The *EMC Simple Support Matrix EMC Federated Live Migration* provides up-to-date information on supported host operating environments.

CHAPTER 7

Storage management tools

Symmetrix systems can be managed with the following tools:

- ◆ [“Unisphere for VMAX” on page 80](#) — A simple, intuitive browser based user interface for the configuration and management of Symmetrix systems. Unisphere for VMAX also provides performance monitoring and analysis, and shares a common look and feel with other EMC management consoles.
- ◆ [“EMC z/OS Storage Manager” on page 80](#) — An Interactive System Productivity Facility (ISPF) interface that manages Symmetrix systems in mainframe environments.
- ◆ [“EMC ControlCenter and ProSphere” on page 80](#) — An intuitive, browser-based family of products that provides management of the overall storage environment, including multivendor storage reporting, monitoring, configuration, and control.
- ◆ [“Solutions Enabler SYMCLI” on page 81](#) — A library of commands that are entered from a command line or from a script.
- ◆ [“SMI-S Provider” on page 81](#) — An SMI-compliant interface for EMC Symmetrix and CLARiiON® systems.
- ◆ [“Mainframe Enablers” on page 81](#) — A suite of components that can be used to monitor and manage a Symmetrix system.
- ◆ [“EMC Virtual Storage Integrator” on page 82](#) — A view of the relationships between VMware virtual machines, disk files, datastores, and underlying EMC storage devices on the Symmetrix VMAX.
- ◆ [“Geographically Dispersed Disaster Restart” on page 82](#) — Automate disaster restart of applications and systems in mainframe environments in the event of a planned or an unplanned outage.

Note: Mainframe Enablers does not provide management for FAST VP and Federated Tiered Storage (FTS).

Unisphere for VMAX

EMC Unisphere for VMAX is a simple, intuitive, browser-based user interface for the configuration and management of Symmetrix systems running Engenuity 5876 and higher.

Unisphere for VMAX presents the functionality of the Solutions Enabler SYMCLI in a browser interface. You can use Unisphere to:

- ◆ Manage Symmetrix access controls, user accounts, and roles
- ◆ Discover Symmetrix systems
- ◆ Perform configuration operations (create volumes, map and mask volumes, set Symmetrix system attributes, set volume attributes, set port flags, create SAVE pools)
- ◆ Manage volumes (change volume configuration, set volume status, reserve volumes, duplicate volumes, create/dissolve meta volumes)
- ◆ Manage Fully Automated Storage Tiering (FAST), Virtual Provisioning, and Auto-provisioning Groups
- ◆ Perform and monitor replication operations
- ◆ Specify and monitor alerts
- ◆ Monitor an application's performance and analyze system performance

Unisphere for VMAX is preinstalled on the service processor (but without the performance section, and can also run on a data center host.

EMC z/OS Storage Manager

The EMC z/OS Storage Manager (EzSM) is a mainframe software product providing storage management in a Symmetrix environment. EzSM provides mainframe storage managers and operations that staff a flexible z/OS-centric view of storage that presents both Symmetrix-specific information and z/OS storage management data in a single easy-to-use 3270 interface.

With EzSM, you can discover and monitor the volumes in a Symmetrix VMAX 20K and 40K, set alerts for volumes, summarize Symmetrix configuration information, and much more. You can use EzSM to:

- ◆ Monitor and report on Symmetrix system mainframe storage
- ◆ Perform SRDF and TimeFinder operations
- ◆ Perform Group Name Services (GNS) operations.

EMC ControlCenter and ProSphere

EMC ControlCenter® and ProSphere™ storage management software provides an end-to-end solution for multivendor storage reporting, monitoring, configuration, and control.

EMC ControlCenter consists of core infrastructure components (providing basic scalability, usability, and information sharing) coupled with a set of licensed applications and license packages that ensure the appropriate level of management and control over your environment.

ProSphere is EMC's next generation Storage Resource Management product that will integrate much of the EMC ControlCenter functionality into a single product.

EMC ControlCenter users under maintenance contracts are entitled to ProSphere, ensuring a seamless path to its advanced capabilities for virtual and cloud storage infrastructures.

The following EMC ControlCenter and ProSphere licenses manage Symmetrix systems:

- ◆ SAN Manager™
- ◆ EMC Symmetrix Manager
- ◆ EMC Symmetrix Optimizer
- ◆ Performance Manager
- ◆ EMC StorageScope™

Solutions Enabler SYMCLI

Solutions Enabler SYMCLI is a specialized library consisting of commands that can be invoked on the command line or from within scripts. Solutions Enabler SYMCLI commands can be used to monitor device configuration and status and perform control operations on devices and data objects within your managed storage complex.

SMI-S Provider

EMC SMI-S Provider is an SMI-compliant interface for Symmetrix and CLARiiON systems that supports the SNIA Storage Management Initiative (SMI). The SMI-Specification (SMI-S) is an open storage management interface that enables the interoperability of multiple vendors' storage management technologies to monitor and control storage resources in multivendor SAN topologies.

Mainframe Enablers

The EMC Mainframe Enablers is a suite of components that you can use to monitor and manage your EMC Symmetrix system. The components listed below are distributed and installed as a single package. This combined packaging simplifies installation and maintenance, and provides assurance of component compatibility.

- ◆ ResourcePak® Base for z/OS facilitates communication between the Symmetrix system and mainframe-based applications provided by EMC or independent software vendors.
- ◆ SRDF® Host Component for z/OS monitors SRDF status and controls SRDF processes through the use of commands executed from a host. SRDF maintains a real-time copy of data at the logical volume level in multiple Symmetrix systems located in physically separate sites.
- ◆ TimeFinder®/Clone Mainframe Snap Facility is the software foundation for three functional products:
 - TimeFinder/Clone allows creating point-in-time copies of full volumes or individual datasets
 - TimeFinder/Snap allows creating pointer-based copies where only the pre-images of changed data are written to the save area.

- TimeFinder/Consistency Group allows performing snap and clone operations on volumes so that the target is dependent write consistent.
- ◆ TimeFinder/Mirror for z/OS allows you to create Business Continuance Volumes (BCVs) and gives you the ability to ESTABLISH, SPLIT, RE-ESTABLISH and RESTORE from the source logical volumes.
- ◆ Consistency Groups for z/OS is designed to ensure the consistency of data remotely copied by the Symmetrix SRDF feature in the event of a rolling disaster.
- ◆ TimeFinder Utility for z/OS is used in conditioning SPLIT BCVs by relabeling the volume and (optionally) renaming and recataloging datasets. This allows the BCV to be mounted and used.
- ◆ **AutoSwap™ for z/OS** can move (swap) workloads from volumes in one set of Symmetrix systems to volumes in other Symmetrix systems without operational interruption.

The *EMC Mainframe Enablers Installation and Customization Guide* provides detailed descriptions of components.

EMC Virtual Storage Integrator

The EMC Virtual Storage Integrator (VSI) has several features that enable the ease-of-integration of Symmetrix systems with virtual environment. The features include:

- ◆ The Storage Viewer features extends the vSphere Client to facilitate the discovery and identification of EMC Symmetrix storage devices that are allocated to VMware ESX/ESXi hosts and virtual machines. Storage viewer presents storage details to the virtual datacenter administrator, merging the data of several different storage mapping tools into a few seamless vSphere Client views.
- ◆ The EMC Storage Pool Management (SPM) features simplifies the task of provisioning Symmetrix storage in VMware environments.

Geographically Dispersed Disaster Restart

EMC Geographically Dispersed Disaster Restart (GDDR) is a mainframe software product for Symmetrix VMAX 20K and VMAX 40K systems that automates business recovery following both planned outages and disaster situations, including the total loss of a data center. EMC GDDR achieves this goal by providing monitoring, automation and quality controls to the functionality of many EMC and third-party hardware and software products required for business restart. As EMC GDDR restarts production systems following disasters, it does not reside on the same servers that it is seeking to protect. EMC GDDR resides on separate logical partitions (LPARs) from the host servers that run your application workloads.

EMC GDDR is installed on a control LPAR at each site. Each EMC GDDR node is aware of the other EMC GDDR nodes through network connections between each site. This awareness allows EMC GDDR to perform the following:

- ◆ Detect disasters
- ◆ Identify survivors
- ◆ Nominate the leader
- ◆ Recover business at one of the surviving sites.

EMC GDDR has no limitations on the number of EMC Symmetrix systems that can be managed. Any limitations are subject to restrictions in EMC hardware and software.

EMC GDDR can be used with certain SRDF configurations and the *EMC® Symmetrix® Remote Data Facility (SRDF®) Product Guide* provides greater detail.

CHAPTER 8

Security

Symmetrix system security features support the latest security protocols and also provide robust authentication, audit logging, and data erasure services.

This chapter provides information about Symmetrix system security features. Topics include:

◆ Audit logs	85
◆ Audit Integration with RSA enVision	86
◆ Data erasure	86
◆ IPv6 and IPsec support	87
◆ Symmetrix Access Controls.....	87
◆ Symmetrix Data at Rest Encryption	88
◆ Symmetrix Service Credential, Secured by RSA.....	89
◆ User authorization controls	90

Audit logs

Symmetrix systems record system activities in secure, tamper-proof audit logs. Audit log protections include:

- ◆ The Symmetrix audit log — Records major Symmetrix system activities.
- ◆ RSA enVision log security — A security management platform that consolidates log security information from multiple sources.

Symmetrix audit log

The Symmetrix audit log records major activities on the Symmetrix system, including host-initiated actions, physical component changes, actions on the Service Processor, and attempts blocked by security controls such as Symmetrix Access Controls.

The Audit Log is secure and tamper-proof - event contents cannot be altered. Additional protection is provided by Symmetrix Access Controls which restrict access to the Audit log by assigning a role (auditor) that allows a user to view, but not modify, the log. [“Symmetrix Access Controls” on page 87](#) provides additional information about Symmetrix Access Control features.

Audit log configuration considerations

The audit log resides on the Symmetrix File System (SFS) within the Symmetrix system. Once the 40 MB capacity limit is reached, the log file will overwrite. To retain audit log history beyond 40 MB capacity, you must regularly monitor and capture audit log contents. You can also export the audit log to a file or log server.

Audit Integration with RSA enVision

The RSA enVision® platform provides a centralized log management service that enables organizations to simplify their compliance programs and optimize their security incident management. The RSA enVision solution facilitates the automated collection, analysis, alerting, auditing, reporting, and secure storage of all logs. Organizations can simplify compliance by using regulation-specific, out-of-the-box reports, alerts and correlations rules. Reports can be scheduled to be delivered at a specific time or run on an ad-hoc basis. Alerts can be delivered through the intuitive user interface, via SMS, or email. Administrators don't have to be glued to the interface at all times. Auditors can even be granted read-only access to the enVision platform so that they can access the reports whenever they need them.

The Solutions Enabler event daemon can be configured to push Symmetrix audit log entries to RSA enVision automatically at scheduled intervals. This integration saves time and resources by simplifying log management. It also helps to track Symmetrix system events in the context of the larger environment. Refer to the *EMC Solutions Enabler Installation Guide* for detailed information about the event daemon.

RSA enVision configuration considerations

RSA enVision provides configured security alerts for the following:

- ◆ Denied login attempts on a host
- ◆ Denied login attempts on the Service Processor
- ◆ Attempts to escalate privileges on Solutions Enabler
- ◆ Attempts to escalate privileges on the Service Processor
- ◆ Symmetrix Access Controls violations
- ◆ Removal of physical drives
- ◆ Successful disk erasure events
- ◆ Unsuccessful disk erasure events
- ◆ Start of a remote file transfer between the customer environment and EMC Service personnel.

This integration is available with RSA enVision v3.5.

Data erasure

Symmetrix systems can be configured to perform data erasure on drives requiring replacement in the Symmetrix system. With the feature enabled, failed drives are erased natively within the system, following the permanent sparing or direct sparing activity and prior to the EMC customer engineer replacing the drive. The default erasure process consists of a multiple pass overwrite with complementary and random character patterns followed by a write verify. The overwrite activity can be modified to address specific customer requirements and is supported on Fibre Channel, SAS, SATA and Flash drives. An optional certification service is available to provide a certificate of erasure and includes the customer retention of drives that fail the erasure activity.

In addition to the Symmetrix in-frame failed drive erasure feature, EMC offers the following data protection features and services:

- ◆ Data at Rest Encryption — Data on all physical drives in the system is encrypted. When the drive is removed from the system, data on the drive is inaccessible.
- ◆ Full system erasure service — Data on all drives in the system is overwritten when replacing, retiring or repurposing a Symmetrix system.
- ◆ Data Erasure/Single Disk service — Multi-platform environments whose security policies require formal certification that failed drives are erased before leaving the premises. This appliance-based service supports Symmetrix platforms using Fibre Channel, SATA, SAS and Flash drives.

All erasure services are performed on-site in the security of the customer's data center and provide a Data Erasure Certificate and report of erasure results.

IPv6 and IPsec support

Internet Protocol Security is an open standard protocol that enhances security for communications over IP networks. IPsec provides encryption and authentication guidelines that allow two computers to trust each other's identity and establish an encrypted communication channel. IPsec prevents or mitigates many security attacks, like spoofing, tampering with data en route, connection hijacking, eavesdropping, and transmission replay.

IPsec is the same security technology used in many corporate VPN solutions to secure communication between a remote PC and corporate systems.

Symmetrix systems 1 Gb/s I/O modules support IPsec. The need for additional encryption hardware is eliminated because IPsec-compatible encryption is included with the 1 Gb/s module. Any combination of IPv4/v6, IPsec enabled, or IPsec disabled is allowed on each engine.

Note: IPsec is not supported for 10 Gb/s link speed.

You can configure IPsec for SRDF or iSCSI connections using Unisphere for VMAX. IPsec configuration support is also available from EMC. IPsec configuration and status is monitored and viewed with Solutions Enabler SYMCLI commands.

IPsec sessions are limited to 100 sessions per port depending on the policy configuration.

Symmetrix Access Controls

Enterprise storage systems share data among different groups and customers and often allow host systems and applications, such as EMC ControlCenter or TimeFinder, to issue management commands to any visible storage device. Data in open enterprise environments such as these may be vulnerable to accidental or deliberate tampering.

Symmetrix Access Controls help protect data in open environments by providing:

- ◆ Access control groups — Provide access IDs and names to hosts with similar access needs.

- ◆ Access pools — Provide permissions, such as ADMIN, assigned to a host, that allow specific functions on a specified set of devices.
- ◆ Access control entries (ACEs) — Grant permissions to access control groups and access pools.
- ◆ Access control lists (ACLs) — Provide a group of ACEs associated the same group.
- ◆ Restricted host access to a defined set of devices (access pools).

Symmetrix Data at Rest Encryption

Symmetrix VMAX 10K, VMAX 20K and VMAX 40K systems support Symmetrix Data at Rest Encryption (D@RE) that provides hardware-based, on-system, back-end encryption for Symmetrix systems by using Fibre Channel I/O modules, validated for FIPS 140-2 Level 1, that incorporate AES-XTS 256-bit data-at-rest encryption. These modules encrypt and decrypt data as it is being written to or read from disk, thus protecting your information from unauthorized access even when disk drives are removed from the system.

D@RE supports either an embedded key manager, or RSA Data Protection Manager client for integration with an external enterprise key manager. The embedded key manager provides internal redundancy for keys, which are never exposed, with recovery built into replacement scripts. When D@RE is installed, you can install the embedded key manager, and never have to revisit.

Note: D@RE must be enabled in the Symmetrix system configuration settings before you start the installation procedure.

For the most up-to-date information on qualified interoperability with the RSA Key Manager Appliance, consult the EMC Support Matrix (ESM), available through E-Lab Interoperability Navigator (ELN) at <http://elabnavigator.emc.com> under the PDFs and Guides tab.

External Enterprise Key Manager

Data at Rest Encryption supports an external Enterprise Key Server, which provides the following benefits:

- ◆ Highly-available and resilient environment for long-term key storage external to the Symmetrix system.
- ◆ Support for a large numbers of keys and clients across geographic and organizational boundaries.
- ◆ Centralized key management and auditing for security compliance objectives.
- ◆ Separation of encryption key storage from encrypted data.
- ◆ Ability to temporarily decommission systems for secure transport from one data center to another

Symmetrix Service Credential, Secured by RSA

The Symmetrix Service Credential (SSC), Secured by RSA, provides industry-leading RSA security credentials to protect you against unauthorized remote or local access to your system via the Service Processor. EMC service personnel, authorized service providers, and EMC service-enabled partners can access the Service Processor only via RSA-protected accounts.

RSA protections provided include:

- ◆ Authentication — Ensures the correct identity of any EMC service personnel who has access to your system.

Any individual who can initiate a service call has been individually authenticated and determined to be a member of the EMC support team. Even field-based personnel are authenticated before they are associated with an account. Also, the call into the Symmetrix Service Processor must originate from one of several specifically designed remote support networks at EMC. This ensures that the necessary networking and security infrastructure is in place to protect your call.

- ◆ Authorization — Grants appropriate privileges based on the EMC service personnel's job. For example, an EMC customer engineer is not able to log in as a senior product support engineer (PSE).
- ◆ Auditing — Captures the actions, access level, and connections of EMC service personnel in a tamper-proof, viewable audit log.

SSC also provides token types that allow you to select the level of protection you need. These include:

- ◆ EMCDefault — Allows use of the Service Processor by EMC Global Services (broadest access).
- ◆ Site protection — Allows EMC personnel to access systems at a particular site (more restrictive).
- ◆ Serial protection — Allows EMC personnel access to a specific Symmetrix system, based on the Symmetrix system serial number (most restrictive).
- ◆ SecureEMR — Provides emergency access to the Symmetrix system (recommended for all configurations).

Unless you specify otherwise, EMCDefault and SecureEMR protections are preinstalled on your Symmetrix system.

User authorization controls

User authorization controls provide a predefined set of roles assigned to users that determine what operations the user can perform. User authorization controls differ from Symmetrix Access Controls because users are assigned a single role for the entire Symmetrix system rather than for individual logical devices or classes of devices. The following roles are defined:

- ◆ None — No actions authorized.
- ◆ Monitor — Authorized for read only actions (audit logs or access control definitions excluded).
- ◆ Storage Admin — Authorized for storage management and modification of group name service group definitions.
- ◆ Admin — Authorized for all operations, including security and Monitor operations.
- ◆ Security Admin — Authorized for additional security operations and all Monitor operations.
- ◆ Auditor — Authorized to grant the ability to view (but not modify) security settings for the Symmetrix system, in addition to all Monitor operations.