



CATE Storage and Backup Engineering

EMC Symmetrix VMAX Storage Array

**CATE STORAGE AND BACKUP ENGINEERING
SG-09065-SG**

Solutions Guide

Version 1.5

1. Version Control

Version	Date	Author	Reviewed	Comments
0.1	Oct 08,2009	Guy Shimabuku	Quang Le	
0.2	Oct 08,2009	Quang le	Guy Shimabuku	Added section 32 – Anti-virus Software Management
1.0	Oct 30, 2009	Guy Shimabuku Quang Le	Chris Clayton	
1.1	Nov 23, 2009	Guy Shimabuku	Quang Le	Appendix C – ACLX database and device
1.2	Aug 02, 2010	Guy Shimabuku		Added section 38 – Appendix D: Host Front-end Port Distribution Edited section 17.1.4 - TDEV metavolume recommendation
1.3	Jan 12, 2011	Guy Shimabuku		Updated section 16 & 35 with the new SCSI flag settings
1.4	Aug 1, 2011	Guy Shimabuku		Sections edited/added to support: • 5875 Enginuity • EFD & SATA II • Loop Population • Extended Loops • Pool Expansion • Thin pool tiers • eLicensing • Symdisk configuration • Port segment capacity
1.5	Oct 4, 2011	Chris Clayton		Symmetrix Performance Analyzer

2. Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, CITIGROUP and the Arc Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

3. Contact

Please address any comments or questions regarding the information contained in this document via email to the “*CATE Storage and Backup Engineering” distribution list.

4. Latest Versions

It is recommended that the latest version of the document be used from the [CSBE website](#). Newer versions may include expanded scope, or comments on issues discovered after deployment.

5. Table of Contents

1. VERSION CONTROL	2
2. LEGAL NOTICE.....	4
3. CONTACT.....	4
4. LATEST VERSIONS	4
5. TABLE OF CONTENTS	5
6. INTRODUCTION.....	8
7. REFERENCE INFORMATION	8
8. SYMMETRIX VMAX INTRODUCTION.....	9
9. VMAX SYSTEM BAY	11
10. STORAGE BAY.....	14
11. AC POWER SPECIFICATIONS.....	15
11.1. POWER AND COOLING REQUIREMENTS	15
12. FLOOR-LOAD BEARING SPECIFICATION	16
13. SYSTEM SPACE REQUIREMENTS	17
14. DISK DRIVES AND CAPACITIES	17
15. COMPONENT REPLACEMENT	18
16. CAPACITY GUIDELINES	18
16.1. VMAX CAPACITY CONFIGURATIONS	19
16.2. CAPACITY SUMMARY TABLE	30
16.3. VMAX EXPANSION FEATURES.....	30
16.4. SEPARATED STORAGE BAYS	30
17. VMAX AND ENGINUITY 5874/5 CONFIGURATION GUIDELINES:.....	31
17.1. FRONT-END PORT CONFIGURATION	33
17.2. DRIVE TYPE CONFIGURATION AND PROTECTION	34
17.3. THICK DEVICES:	35
17.4. GATEKEEPER DEVICE CONFIGURATION GUIDELINES:.....	36
17.5. ELICENSING GUIDELINES	36
18. THIN PROVISIONING GUIDELINES	37
19. AUTO-PROVISIONING GUIDELINES	41
20. DRIVE SPARING.....	42
21. VMAX VIRTUAL PROVISIONING	44
21.1. I/O FUNCTIONALITY	45
21.2. VOLUME TABLE OF CONTENTS	46
21.3. THIN POOL ALERTING.....	46

21.4.	AUDIT LOGGING AND ALERTING	47
21.5.	THIN POOL EXPANSION.....	49
22.	AUTO-PROVISIONING INTRODUCTION	51
23.	PHYSICAL INSTALLATION	53
24.	STANDARD HARDWARE CONFIGURATION	53
25.	DYNAMIC LUN ADDRESSING	53
26.	HARDWARE REPLICATION	53
27.	SYMMETRIX VMAX MANAGEMENT.....	54
28.	VENDOR REMOTE ACCESS TO SYMMETRIX.....	54
29.	AUDIT LOGGING	54
30.	SYMMETRIX PRIORITY CONTROLS.....	54
31.	ENABLING ECC SUPPORT	55
32.	SYMACL SYMMETRIX SETUP	55
33.	FUNCTIONAL IDENTIFICATION	55
34.	VMAX SERVICE PROCESSOR ANTI-VIRUS SOFTWARE MANAGEMENT	55
35.	CHANNEL ADAPTER FLAG SETTINGS	57
36.	APPENDIX A – PORT AND INITIATOR FLAG SETTINGS AND PROCEDURES	59
36.1.	E-LAB NAVIGATOR - VMAX CHANNEL ADAPTER FLAG SETTINGS	59
36.2.	CONFIGURING SPC-2 AND OS2007 FLAGS PER FA PORT	61
36.3.	CONFIGURING SPC-2 AND OS2007 FLAGS PER INITIATOR	63
36.4.	HOST VERIFICATION OF THE SPC-2 AND OS2007 FLAG SETTING.....	65
37.	APPENDIX B – GATEKEEPER CONFIGURATION	66
37.1.	GATEKEEPER DEVICE CONFIGURATION GUIDELINES:.....	67
37.2.	GATEKEEPER CLI COMMANDS:	67
38.	APPENDIX C - ACLX DATABASE/DEVICE	68
38.1.	ACLX DATABASE	68
38.2.	ACLX DEVICE.....	69
38.3.	ACLX DEVICE AND PORT FLAG SPECIFICATIONS:.....	69
38.4.	ACLX DEVICE HOST VISIBILITY	69
38.5.	ACLX DATABASE MAINTENANCE.....	72
38.6.	ACLX DATABASE EXAMPLE:	73
39.	APPENDIX D: HOST FRONT-END PORT DISTRIBUTION	74
39.1.	2-ENGINE HOST PORT DISTRIBUTION.....	75
39.2.	4-ENGINE HOST PORT DISTRIBUTION	76
39.3.	6-ENGINE HOST PORT DISTRIBUTION.....	77

39.4.	8-ENGINE HOST PORT DISTRIBUTION	79
40.	APPENDIX F: METAVOLUMES – STRIPED VS. CONCATENATED	80
41.	APPENDIX G: ELICENSING MANAGEMENT	81
41.1.	eLICENSE INTRODUCTION	81
41.2.	eLICENSE GLOSSARY	83
41.3.	GUIDELINES	83
41.4.	VIEW ACTIVATED ENTITLEMENT FILE:.....	84
41.5.	VIEW ACTIVATION CERTIFICATE	87
41.6.	ENTITLEMENT ACTIVATION	89
41.7.	ACCESSING ENTERPRISE LICENSE AGREEMENTS	94
41.8.	COORDINATOR (USER) MANAGEMENT	95
41.9.	VIEW LICENSES THROUGH SOLUTIONS ENABLER.....	98
41.10.	ENTITLEMENT FILE INSTALLATION	100
42.	APPENDIX H: CAPACITY TO SEGMENT RATIO	101
42.1.	FIBER ADAPTER PORT MONITORING:.....	103
43.	APPENDIX I: BCV CACHE UTILIZATION	107

6. Introduction

The goal of the Symmetrix VMAX Solution Guide is to provide standards and recommended practices for Citigroup storage teams. Implementation of these practices should create an efficient support model with consistent configurations enabling common controls and policies to be implemented for the distributed storage environment.

This guide defines the requirements for successful hardware installation and configuration of an EMC Symmetrix VMAX storage array. This is a living document and as such, it will be updated as new functionality or services are certified.

This document is intended for all Citigroup technology groups responsible for the deployment and support of enterprise storage. The information compiled in this document has been tested and certified by Global Distributed Storage Engineering ("CATE STORAGE AND BACKUP ENGINEERING").

The best practices outlined in this document include Standards and Recommendations. Standards should be adhered to for effective common controls. Recommendations are applicable to most environments and should be employed where possible.

7. Reference Information

Document	Location
CATE Storage Engineering Website	http://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/default.aspx
CATE STORAGE AND BACKUP ENGINEERING Symmetrix Documentation	http://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/techstdshmpgs/symmetrix.aspx
EMC Corporation. <i>EMC Symmetrix Virtual Provisioning Implementation and Best Practices</i>	http://powerlink.emc.com

Best Practices and Engineering Papers

- *EMC Solutions Enabler Symmetrix Array Management CLI Product Guide*
- *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*

- *EMC Solutions Enabler Symmetrix CLI Command Reference*
- *Best Practices for Fast, Simple Capacity Allocation with EMC Symmetrix Virtual Provisioning* (This document provides step-by-step instructions on configuring, managing and monitoring of thin devices)
- *EMC Symmetrix VMAX series Product Guide*
- *EMC Symmetrix VMAX series Physical Planning Guide*
- *Storage Provisioning With EMC Symmetrix Auto provisioning*
- *EMC Symmetrix Virtual Provisioning*

8. Symmetrix VMAX Introduction

The VMAX system bay contains from 1 to 8 VMAX Engines, where each engine contains two VMAX directors operating independently. Each Engine can support up to 240 dual-ported fibre channel drives and up to 360 drives in certain high capacity configurations. Each Director is the primary initiator for half the drives, and the secondary (backup) for the other half.

VMAX is EMC's the first enterprise-class array built on industry standard components using Intel® Xeon® 4 core processors and support logic, PCIe internal interconnects, and DDR2 fully-buffered ECC SDRAM DIMMs. VMAX is built from off-the-shelf commodities; however, the design is still enterprise class (e.g. every I/O controller, processor boards, blower and power supply is hot-swappable). All the I/O controllers are designed with power ramps to protect components when they are removed or inserted.

The VMAX RapidIO-based Virtual Matrix Interconnect represents the major architectural change over previous Symmetrix arrays. This in-house designed matrix contains the distributed lock management and coordinates all references to local and remote global memory.



Enginuity

Enginuity, the operating system for the frame, has also been enhanced to support the Virtual Matrix Architecture. Several of the most significant changes are depicted below:

- Memory management
- Block CRC
- RAID Virtual Architecture, Virtual LUNs and Mirror Positions
- Dynamic and concurrent configuration changes
- Expanded configuration limits

Memory Management

VMAX employs new techniques that reduces the amount of memory required per usable GB of capacity, improves memory reference performance and is dynamically extensible to support future metadata expansion. Block-level CRC is now stored alongside every block of data with VMAX. Each 512 byte data block is incorporated with an 8-byte data integrity field using an algorithm based on the T10-DIF proposal. The integrity field is appended to each 512 byte block on writes as the data enters the array and verified at each instance as the data is moved within the system. In addition to the CRC, referential information that specifies the Logical Block Addresses (LBA) is included within the integrity field. Enginuity uses this info to validate that the content of the block hasn't been corrupted, and that it is indeed the block of data from the requested LBA. Also, Enginuity stores check bits on track tables that are stored separately from the

data and the block CRC. According to EMC, the VMAX memory enhancements supports 60% more usable capacity than a similarly configured DMX-4. In a full configuration, the VMAX supports 1TB of global memory and over 2PB of usable capacity, which is 3x more than the DMX-4.

RAID Virtual Architecture

The RAID Virtual Architecture configures all RAID types (i.e. RAID-10, RAID-5, RAID-6) as a single mirror position on VMAX. These changes enable several key customer benefits:

- Drive rebuilds will no longer be delayed until replication or volume relocation is completed.
- Support device relocations between different RAID protection types
- "Virtual LUN" relocation without disrupting local AND remote replication sessions or relationships

Bin File


VMAX still uses a bin file consistent with the Symmetrix product line. EMC has made significant strides in reducing the time required to implement bin file changes. Dynamic and concurrent provisioning enables more operations to be performed simultaneously; whereby, the latency to perform operations such as device relocations, hot sparing and changing port flags has been reduced.

Device Limits

VMAX will support 512 disk hypers, which is twice as many hypers over the previous versions of Engenuity code. The maximum single volume size is 262,668 cylinders (239GB), an increased of 4X over the DMX-3/4. The maximum number of volumes remains at 64,000; however, the metadata structure allows EMC to easily extend these limits in the future.

9. VMAX System Bay

- VMAX Engines
- Matrix Interface Board Enclosure (MIBE)
- Power subsystem components
- Standby power supplies (SPS)
- Service Processor, keyboard/video/mouse, UPS



VMAX Engine

Symmetrix VMAX system supports a maximum of eight VMAX Engines within a single system bay. Each VMAX Engine includes:

- Two directors that support front-end, back-end and SRDF connections.
- Each director has:
 - 16, 32, or 64 GB of physical memory
 - One System Interface board (SIB) that connects the director and the Matrix Interface Board Enclosure (MIBE)
 - Two Back End I/O Modules (4 ports, 4 Gb/s) that connect to storage bay drives.
- Two I/ O Module carriers that provide connectivity between the director and the front-end and (optional) SRDF I/O ports. Front End I/O Modules support:
 - Fibre Channel host connectivity (4 ports, 2 or 4 Gb/s)
 - Fibre Channel SRDF connectivity
 - FICON host connectivity (2 ports, 2 or 4 Gb/s)
 - iSCSI host connectivity (2 ports, 1 Gb/s)
- Two management modules that provide environmental monitoring.
- Two enclosure power supplies.
- Four cooling fans.
- Matrix Interface Board Enclosure (MIBE) contains a matrix interface board and two power/fan modules and provides redundant connections for system components.


9.1.1. Power subsystem:

Two PDPs, one for each power zone, connect to the PDUs and provide a centralized bay interface and distribution control of the AC power input lines. The PDPs also contain the manual on/off power switches. The PDUs, one for each

power zone, provide the main interface between the input AC from the PDPs and the various components contained within the system bays.

Two Standby Power Supply (SPS) modules for each VMAX Engine provide backup power in the event of an AC power interruption. SPS modules maintain power for two 5 minute periods of AC loss while the Symmetrix system shuts down.

10. Storage Bay



Each storage bay holds from 8 to 16 disk array enclosures (40–240 drives). Each disk array enclosure contains:

- Two redundant disk array enclosure power /cooling fans
- Two link control card (LCC) modules
- From 5 to 15 drives per direct attached enclosure
- From 4 to 15 drives per daisy chained enclosure

10.1.1. Power subsystem:

Two PDPs, one for each power zone, connect to the system bay PDUs and provide a centralized bay interface and distribution control of the AC power input lines. The PDPs contain the manual on/off power switches. The PDUs, one for each power zone, provide the main interface between the input AC from the PDPs and the various components contained within the storage bay.

Two Standby Power Supply (SPS) modules that provide backup power for each row of four disk array enclosures. If an AC power source is interrupted or fails, the SPS modules can maintain power for two 5-minute periods of AC loss while the Symmetrix storage bay shuts down.

11. AC Power Specifications

North American AC power specifications for three-phase (Delta)
and single-phase power Specification North America three-phase (Delta)

Input voltage:	200–240 VAC L-L nom (240 VAC optimal)
Range:	200–240 VAC $\pm$ 10%
Frequency:	50–60Hz
Circuit breakers:	50 A
Power zones:	Two
Power requirements at customer site (min):	Two 50 A, three-phase drops per bay

International AC power specifications for three-phase, five-wire
WYE power and single-phase power Specification International three-phase (WYE) International single-phase

Input voltage:	220–240 VAC L-N nom (240 VAC optimal)
Range:	200–240 VAC + 10%
Frequency:	50–60Hz
Circuit breakers:	32 A
Power zones:	Two
Power requirements at customer site (min):	Two 32 A, three-phase drops

11.1. Power and Cooling Requirements

System configuration description	Total power consumption (kVA)	Heat dissipation (Btu/Hr)
System bay	7.8	26,300

Storage bay	6.1	19,800
system bay and one storage bay (two VMAX Engines)	8.3	27,200
system bay and two storage bays (four VMAX engines)	16.5	53,300
system bay with four storage bays (six VMAX engines)	30.4	99,200
system bay and six storage bays (eight VMAX Engines)	44.4	145,100
system bay and eight storage bays (eight VMAX Engines)	56.6	184,700
system bay and nine storage bays (eight VMAX Engines)	62.7	204,500
system bay and ten storage bays (eight VMAX Engines)	68.8	224,300

12. Floor-load bearing specification

System bay weights are calculated based on the number of VMAX Engines and storage bay weights assuming a fully loaded configuration. Use the power and weight calculator, at <http://powercalculator.emc.com>, to refine the weights shown here to more closely match your Symmetrix VMAX system.

System configuration description	Weights lb (kg)
Symmetrix VMAX system bay (four VMAX Engines)	1,830 lb (830.0 kg)
Symmetrix VMAX system bay (eight VMAX Engines)	2,774 lb (1258.3 kg)
Storage bay	2,144 lb (972.5 kg)
Symmetrix VMAX system bay and one storage bay (two VMAX Engines)	3,502 lb (1588.5 kg)
Symmetrix VMAX system bay and two storage bays (four VMAX Engines)	6,118 lb (2775.0 kg)
Symmetrix VMAX system bay and three storage bays (six VMAX Engines)	8,734 lb (3961.7 kg)
Symmetrix VMAX system bay and four storage bays (eight VMAX Engines)	11,350 lb (5148.3 kg)
Symmetrix VMAX system bay and five storage bays (eight VMAX Engines)	13,494 lb (6120.8 kg)
Symmetrix VMAX system bay and six storage bays (eight VMAX Engines)	15,638 lb (7093.3 kg)
Symmetrix VMAX system bay and seven storage bays (eight VMAX Engines)	17,782 lb (8065.8 kg)
Symmetrix VMAX system bay and eight storage bays (eight VMAX Engines)	19,926 lbs (9038.3 kg)
Symmetrix VMAX system bay and nine storage bays (eight VMAX Engines)	22,070lbs (9557.2 kg)
Symmetrix VMAX system bay and ten storage bays (eight VMAX Engines)	24,214 lbs (10983.3 kg)

EMC recommends the use of 24 X 24 inch heavy-duty, concrete-filled steel floor tiles. If a different size or type of tile is used, the customer must ensure that the tiles have a minimum load rating that is sufficient for supporting the system weight. To ensure proper physical support of the system, the following requirements are based on the use of 24 x 24 in. (61 x 61 cm) heavy-duty, concrete-filled steel floor tiles. Raised floors must meet the following requirements:

- Floor tiles and stringers must be rated to withstand concentrated loads of two casters, each weighing up to 1,000 lb (453.6 kg).
- Floor tiles and stringers must be rated for a minimum static ultimate load of 3,000lbs (1,360.8 kg) per tile.
- Floor tiles must be rated for a minimum of 1,000 lb. (453.6 kg) rolling load.
- Use or create no more than one floor tile cutout that is no more than 8 in. (20.3 cm) wide by 6 in. (15.3) deep in each 24 x 24 in. (61 x 61 cm) floor tile.

13. System Space requirements

The system bay is placed in the center and the storage bays on each side of the system bay. The table below describes physical space requirements for system and storage bays:

Description	Symmetrix VMAX system bay measurements	Storage bay measurements
Height	76.66in(194.7cm)	76.66in(194.7cm)
Width	30.2in(76.7cm)	30.2in(76.7cm)
Depth	41.88in(106.4cm)	41.88in(106.4cm)
Clearance for service, airflow		
Front	42in(106.7cm)	42in(106.7cm)
Rear	42in(106.7cm)	42in(106.7cm)
Top	18in(45.7cm)	18in(45.7cm)

14. Disk Drives and Capacities

The table below describes supported Symmetrix VMAX disk types, rotational speeds and formatted capacities. The information is updated whenever EMC modifies drive support for Symmetrix VMAX systems.

Supported Disk Drive Types				
Disk Name	Formatted Capacity (GB)	Array Size (Cylinders)	Speed RPM	EMC Nomenclature
Fibre Channel 300GB	288.19	305175	15000	T300155
Fibre Channel 450GB	432.29	457763	15000	H450156, HUC4515
Fibre Channel 600GB	576.39	586337	15000	EGC6015
SATA II 2TB	1882.72	1915210	7200	N02THJK
EFD (Flash) 400GB (note1)	393.84	400640	N/A	STC0400
EFD (flash) 200GB (note 1)	186.34	200320	N/A	CTSG200

Note 1: For use on three Tier FAST VP frames only


15. Component Replacement

The VMAX modular design allows for non-disruptive component replacement that minimizes the number of potential failure points. Non-disruptive component replacement allows the customer support personnel to install a new component, initialize it if necessary, and bring it online without disrupting access to unaffected volumes. VMAX supports the non-disruptive replacement of all major components as follows:

- Physical memory
- Directors
- Front End I/O Modules
- Disk drives
- Fan modules
- Link control cards (LCCs)
- System power supply (SPS) modules
- Power supplies
- VMAX Engines
- Serial Interface Boards (SIBs)
- Matrix Interface Board Enclosure (MIBE)
- Power supplies
- Power distribution units (PDUs)
- Power distribution panels (PDPs)
- Fibre Channel, Ethernet, RS-232 and power cables
- Service Processor — keyboard, video display, mouse (KVM), uninterruptible power supply (UPS)

16. Capacity Guidelines

VMAX capacities are configured and expanded by adding VMAX Engines and disk array enclosures. Each VMAX Engine connects directly to eight disk array enclosures. Additional disk array enclosures can be daisy-chained into 30 or 45-drive loops. The number of allowable drive loops depends on the number of VMAX Engines and the number and location of the disk array enclosures within the storage bays. VMAX Engines are installed in a specific order, and must reside in specific cabinet locations (see diagram below). Engines 4 and 5 represents the starting configuration and engine expansions move outward from these engines. For example, a 2 to 4-engine expansion incorporates engines 3 through 6.




16.1. VMAX Capacity Configurations


EMC supports +50 VMAX capacity configurations; therefore, in an effort to simplify the ordering process CATE STORAGE AND BACKUP ENGINEERING has narrowed down the possible configurations to eleven. Each of the eleven configurations is based on 2, 4, 6, and 8 VMAX engines. The direct connect bays are numbered 1, and the daisy chained bays are numbered 2 and 3. The A, B, C, and D are used to identify the storage bays that share Symmetrix VMAX Engine resources.

Extended loop configuration are presented for 2 and 4-engine frames. Extension allows up to 45 drives per loop increases a 2-engine frame to a maximum of 720 drives and 1440 drives for a 4-engine frame. The expanded loop option pertains to frames with low to medium I/O load that need additional capacity to accommodate existing or a small number of new hosts.


16.1.1. Two Engine Configuration

Configuration Number	# of Engines (# dir)	Total Storage Bay	Full Storage Bays	Half-filled Storage Bays	Direct Connected Storage Bays	Daisy Chained Storage Bays	Total Disk Enclosures	# of loops with 15 drives	# of loops with 30 drives	Maximum # of Drives
SB-CONFIG05	2 (4)	1	1	0	1	0	16	16		240
SB-CONFIG07	2 (4)	2	2	0	1	1	32		16	480






Two Engine 480 Drive Configuration



16.1.2. Two Engine Configuration – Extended Loop



Configuration Number	# of Engines (# dir)	Total Storage Bay	Full Storage Bays	Half-filled Storage Bays	Direct Connected Storage Bays	Daisy Chained Storage Bays	Total Disk Enclosures	# of loops with 45 drives	# of loops with 30 drives	Maximum # of Drives
----------------------	----------------------	-------------------	-------------------	--------------------------	-------------------------------	----------------------------	-----------------------	---------------------------	---------------------------	---------------------


SB-CONFIG02E	2 (4)	3	3	0	1	2	48	16		720
--------------	-------	---	---	---	---	---	----	----	--	-----



16.1.3. Four Engine Configuration


Configuration Number	# of Engines (# dir)	Total Storage Bay	Full Storage Bays	Half-filled Storage Bays	Direct Connected Storage Bays	Daisy Chained Storage Bays	Total Disk Enclosures	# of loops with 15 drives	# of loops with 30 drives	Maximum # of Drives
SB-CONFIG12	4 (8)	2	2	0	2	0	32	32		480
SB-CONFIG 14	4 (8)	3	3	0	2	1	48	16	16	720
SB-CONFIG16	4 (8)	4	4	0	2	2	64		32	960







16.1.4. Four Engine Configuration – Extended Loop


Configuration Number	# of Engines (# dir)	Total Storage Bay	Full Storage Bays	Half-filled Storage Bays	Direct Connected Storage Bays	Daisy Chained Storage Bays	Total Disk Enclosures	# of loops with 45 drives	# of loops with 60 drives	Maximum # of Drives
SB-CONFIG10E	4 (8)	6	6	0	2	4	96	32		1440



16.1.5. Six Engine Configuration



Configuration Number	# of Engines (# dir)	Total Storage Bay	Full Storage Bays	Half-filled Storage Bays	Direct Connected Storage Bays	Daisy Chained Storage Bays	Total Disk Enclosures	# of loops with 15 drives	# of loops with 30 drives	Maximum # of Drives
SB-CONFIG 22	6 (12)	4	4	0	3	1	64	32	16	960
SB-CONFIG24	6 (12)	5	5	0	3	2	80	16	32	1200
SB-CONFIG26	6 (12)	6	6	0	3	3	96		48	1440







16.1.6. Eight Engine Configuration

Configuration Number	# of Engines (# dir)	Total Storage Bay	Full Storage Bays	Half-filled Storage Bays	Direct Connected Storage Bays	Daisy Chained Storage Bays	Total Disk Enclosures	# of loops with 15 drives	# of loops with 30 drives	# of loops with 45 drives	Maximum # of Drives
SB-CONFIG36	8 (16)	6	6	0	4	2	96	32	32		1440
SB-CONFIG38	8 (16)	7	7	0	4	3	112	16	48		1680
See Notes	8 (16)	8	8	0	4	4	128		64		1920
SB-CONFIG44	8 (16)	10	10	0	4	6	160		32	32	2400





Notes:

1. 8-engine 30 drives per loop configuration have not been included in the EMC ordering guide. This configuration is technically valid and will be in the EMC ordering guide 1Q10. An RFQ will be used to order this configuration.

2. SB-CONFIG44 eight engine 2400 drive configuration is restricted to upgrade expansions only. The additional two storage bays over the 1920 drive configuration provides another 480 drives or approximately 189TB (450GB FC drives) at RAID-6. The additional two storage bays increases the loop populations for engine 1, 2, 7 and 8 to 45 drives, making this option one of the least attractive options for expanding a frame. The primary recommendation is to purchase another array if additional capacity over the 1920 drive frame is required; however, if space restrictions and/or time frames prevent the purchase of another frame, this expansion upgrade may be ordered on an as needed basis.

16.2. Capacity Summary table

Component	2-engine	4-engine	6-engine	8-engine
Number directors	4	8	12	16
Total backend loops	16	32	48	64
Minimum number of drives	80	160	240	320
Maximum number of drives	480	960	1440	1920
Min/Max global memory	64/256 GB	128/512 GB	192/768 GB	256/1024 GB
FA Ports	32	64	96	128

16.3. VMAX Expansion Features

The VMAX architecture is inherently scalable and provides many options for online hardware upgrades. As depicted in the capacity configuration, [section 15.1](#), drive count overlap between the engine counts were purposely inserted to allow for the pre-placement of engines to satisfy future growth through expanding the loop population. The sheer capacity of the VMAX array will make engine and drive expansions a common occurrence. Therefore, EMC has incorporated several expansion features into the VMAX architecture:

- Each storage bay is divided into an upper and lower half. VMAX Engines in the upper half of the system bay are cabled to the upper half storage bays and VMAX Engines in the lower half of the system bay are cabled to the lower half of the storage bays. In this way cabling is controlled, consistent and tidy through all configurations.
- Symmetrix VMAX Engines in the system bay are populated from the center position outward.

16.4. Separated Storage Bays

Physically separating VMAX storage bays provides storage bay placement flexibility and addresses floor loading limitations within a data center. The following items provide a basic overview of a VMAX array with a separation of two or more storage bays from an adjacent storage bay:

- The four storage bays that connect directly to the system bay (1A, 1B, 1C, 1D) cannot be separated from the system bay.
- Storage bays 2 and 3 can be located up to 12 feet away from storage bay 1. Actual separation distance will be determined on a case-by-case basis and will depend on the customer's configuration, cable routing and data center design.
- A separated storage bay is cabled to the adjacent storage bay using 10-meter cables (32 cables— 2 bundles of 16 cables).

17. VMAX and Enginuity 5874/5 Configuration Guidelines:

- Vaulting requires the use of five disk drives on each backend loop. The first five disk drives on each drive loop in the direct-attached Storage Bay must be Fibre Channel or SATA II to support the vault devices. Each vault device is 5120 cylinders in size. Example, a 2-engine frame will require 80 vault devices.
- Replicated and non-replicated devices should not reside on the same frame. For operational simplicity, it is recommended that replicated (RDF) devices be segregated by frame from non-replicated devices. Exceptions are allowed for replicated devices to augment a frame dedicated to non-replicated devices, due to a change in operational requirements from the original intent of the frame.
- VMAX engine counts of 2, 4, 6 and 8 are supported, accommodating from 80 to 1920 drives.
- Enginuity 587X utilizes only one mirror position for all RAID protection types.
- Thick and thin devices may co-exist on the same frame.
- Each director on a frame shall contain a like amount of cache with a minimum of 32GB each. Any upgrade to global cache shall be equal distributed between each director. (See [Appendix I](#))
- A maximum of 30 disks per loop is supported. As a restricted configuration, 2 and 4-engine frames maybe upgraded to a maximum of 45 drives per loop. A detailed performance evaluation at 30 drives per loop is needed to determine that the probability of performance

bottlenecks will be low with the additional drives. A Loop population of 45 drives expands the 2-engine frame to 720 drives and the 4-engine frame to 1440 drives. The loop extension upgrade is non-disruptive. (See sections [16.1.2](#) and [16.1.4](#))

- Symmetrix disk (symdisk) groups shall categorize the same physical disk types and specific data protection types. Symdisk groups shall utilize working physicals equally distributed across each backend loop. For example, a 2-engine frame contains 16 backend loops; therefore, each symdisk group shall be created or expanded in increments equal to 16 working physicals.
- The ACLX device is the initial communication medium between the Solutions Enabler management server and the VMAX array. Once gatekeepers are assigned to the management server, the ACLX device need not be mapped to each FA port. (See [Appendix C](#) - ACLX Database/Device)
- The Symmetrix File System (SFS) uses four 8738 cylinder devices for the internal storage.
- The VMAX backend subsystem supports the industry-standard 4Gbps Fibre Channel (FC) and SATA II drives. Placing 2Gbps drives into the VMAX is not supported.
- Dynamic Cache Partitioning is not recommended due to the possibility of misappropriating cache resources. Dynamic Cache Partitioning provides logical partitioning of cache into eight groups (seven new ones with one default).
- VMAX 1150 devices created at 1150 cylinders (1.053GB) in size with all certified protection schemes supported. This increase in quorum device size is pursuant to supporting newer cluster management applications. Recommend that quorum devices be 100% preallocated.
- The VMAX frame supports 8192 storage groups with a maximum of 4096 devices per group.
- Each daisy chained disk enclosure requires a minimum of four drives.
- Priority Service Time (PST) prioritizes a device or device group by increasing or decreasing the potential I/O delay of the request. PST utilizes a performance metric of 1 through 16, where level 1 is the highest level of throughput and the default value for each device within the array. Each PST level is relative to each other; therefore, throughput differences between PST levels are dependent upon the load on the array. It is

recommended that PST only be used to lower a level of performance on a device or device group. PST levels should be lowered in small increments (1-3 levels) until an acceptable level of performance is achieved.

17.1. Front-end Port Configuration

- Channel Adapters (front-end) shall have the following SCSI and Fibre Channel flags enable (see [Section 34](#) for explanation of SCSI flags and implementation guidelines):

SCSI Flags Enabled
Common_Serial_Number(C)
Disable_Q_Reset_on_UA(D)
SCSI_3(SC3)
SPC2_Protocol_Version(SPC2)
SCSI_Support1(OS2007)
Fibre Channel Flags Enabled
Init_Point_to_Point(PP)
Unique_WWN(UWN)
Access_Logix(ACLX)
Auto_Negotiate(EAN)

- Setting port flags overrides the front-end port settings and replaces the heterogeneous configuration flag setting used with symmask. Symaccess shall be used to configure the port flags on a per initiator basis. For example, V-bit override, for volume set-addressing on HP-UX servers, can be enabled on an as needed basis per WWN. (See [Appendix A](#), for examples on setting the SCSI flags).
- Host front-end port access shall be distribution between the available engines. The rule-of-17 or even-odd technique is the recommended method to distribute host FA port access. (see [section 38](#) for additional information on host port distribution)
- An FA port shall be limited to a maximum of 50TB of mapped capacity. A maximum of 4096 devices, including metavolume members, are support on a single FA slice. (See [Appendix H](#))
- A host shall be zoned to a minimum of one segment; pair of FA ports. If the host application is deemed to be a high performing, additional segments up to a maximum of four may be zoned to the host. Additional

segment appearances per host increase the LUN capacity that must be mapped to each front-end port. [Appendix H: Front-end Port Performance an Analysis](#)

- VMAX Fiber Adapter (FA) ports should be set with auto-negotiation enabled (default) setting.

17.2. Drive Type Configuration and Protection

- All supported drive types (see table below) may co-exist on the same frame:

Supported Disk Drive Types				
Disk Name	Formatted Capacity (GB)	Array Size (Cylinders)	Speed RPM	EMC Nomenclature
Fibre Channel 300GB	288.19	305175	15000	T300155
Fibre Channel 450GB	432.29	457763	15000	H450156, HUC4515
Fibre Channel 600GB	576.39	586337	15000	EGC6015
SATA II 2TB	1882.72	1915210	7200	N02THJK
EFD (Flash) 400GB (note1)	393.84	400640	N/A	STC0400
EFD (flash) 200GB (note 1)	186.34	200320	N/A	CTSG200
Note 1: For use on three Tier FAST VP frames only				

- A minimum of 8 spare disk drives per array. Allocate 2 spares for every 100 disk drives (FC and SATA) types.
- Enterprise Flash drives (EDF) require 1 spare per 32 drives or two per 100 drives greater than 32 drives. EDF drive spares are counted separately from disk drives. EDF do not count toward the minimum disk drive spare count.
- Permanent sparing is the **only** sparing mechanism supported on VMAX (see [section 20](#)). Be aware that permanent sparing may invoke the following configuration anomalies:
 - Locate two RAID 6 device members on the same disk director, loop, or drive enclosure. Having RAID group members share the same backend elements is not a recommended practice.
 - Vault devices can be spared to any disk drive within the same loop, does not need to be one of the first 5 drives. Flash drives cannot contain vault devices.

The Samsung CTSG200 flash drive is the preferred EDF choice over the STEC (STC0X00) flash drives. The Samsung flash drives has shown greater throughput potential as compared to the 200/400GB STEC drives.

17.3. Thick Devices:

- RAID-6 14+2 are supported for each Fibre Channel SATA II drive types. RAID-6 provides the highest level of data protection. Gatekeepers utilize RAID-1 protection.
- Each drive and RAID type must reside in separate symdisk groups.
- Thick device sizes of 36828 to 262668 cylinders are supported in size increments of 36828 cylinders. Recommended the use of the largest devices possible, whereby limiting the use of metavolumes. The larger devices will reduce the number of individual splits per physical drive, thereby reducing the resources required to manage the backend subsystem:

units	size (cyl)	size (bytes)	size (gigabytes)
1	36828	36203397120	33.7170
2	73656	72406794240	67.4341
3	110484	1.0861E+11	101.1511
4	147312	1.44814E+11	134.8682
5	184140	1.81017E+11	168.5852
6	220968	2.1722E+11	202.3022
7	262668	2.5821E+11	240.646

- The one cylinder striping of hypervolumes is the only supported metavolume configuration for thick devices.
- Metavolumes are supported for device sizes of 262668 cyl and larger.
- Metavolumes may contain 255 members with a total capacity of ~60TB. Six meta-members encompassing 1.4TB device is the recommended maximum.
- Engenuity 5784 allows a maximum of 512 hypervolumes on each physical drive. Engenuity 5785 allows a maximum of 1024 hypervolumes on each physical drive.

17.4. Gatekeeper Device Configuration Guidelines:

See Appendix B for Gatekeeper description

- Solutions Enabler host to contain 16 Gatekeepers from each frame for basic operation:
 - 10 for simple Solutions Enabler commands and configuration manager
 - 1 for each daemon. There are usually 6 daemons initiated per SE server.
 - 6 for SMC and Ionix Control Center (if needed)
- A gatekeeper shall only be masked to only one host.
- Gatekeepers are typically 3 (2888KB) or 6 (5760KB) cylinders in size. The maximum size should not extend beyond 10 cylinders, as the inquiry Gatekeeper flag will be avoided. The minimum Gatekeeper size should be large enough to contain a host volume label/signature.
- All gatekeeper type devices should use RAID-1 protection.
- Gatekeepers support multipathing through EMC PowerPath.

17.5. eLicensing Guidelines

- Enginuity 5875 and Solutions Enabler +7.2 require frame and host based feature sets to be licensed through EMC's PowerLink License Management Application. The License Management Application activates new license keys by creating entitlements files that pertain to a specific frame. License Management Application also manages user access; entitlement queries and feature set updates. (See [Appendix G](#)). Initial and updates to the entitlement file is initiated through the sale order process.
- The "symlmf" command in Solutions Enabler 7.2 and later is required to install the entitlement file into an array or host.
- Moving a license from a DMX to a VMAX is not supported by the software license agreement. Customers do not receive credit for old licenses.
- EMC License Operations is authorized to update sales orders from the old to the new product numbers. However, keys for unlicensed products cannot be created.

- SRDF commands spanning between an Enginuity 5875 and earlier array (i.e. Symmetrix DMX, VMAX) releases requires the SRDF and SYMCLI keys in the 5875 to command both boxes. The SYMCLI key must be on the host issuing the commands even if connected only to the 5875 array.
- VMAX arrays purchased before 2011 may appear to lack feature authorization in Enginuity 5875, which is the result of the sales order containing old product model numbers. This requires update to the sales transaction for new model numbers, which triggers an entitlement creation. If this issue is suspected, contact EMC License Operations at Licensekeys@emc.com or call 800-782-4362 option 4.
- For support with missing or incorrect entitlements after activation contact the EMC License Key Group at LicenseKeys@emc.com or call 1-800-782-3797 (Americas), or + 353 (0) 21 4879862 (EMEA).

18. Thin Provisioning Guidelines

Thin Provisioning Solutions Guide located at the following link:

https://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/platforms/Symmetrix/09058-sg.pdf

- Thin devices may co-exist on the same frame with thick devices.
- A thin pool capacity shall consist of DATADEVs constructed from a single symdisk group.
- Each symdisk group shall include disk drives from the same family and rotational speed. The disk drives shall be configured with the widest possible backend loop distribution.
- A symdisk group capacity shall not be shared between thin pools.
- The minimum number of Fibre Channel disk drives assigned to a symdisk group for a thin pool is equal to 4 times the number of backend loops. This minimum threshold assures proper I/O distribution required to maintain a high level of throughput performance. Example, a two engine VMAX has 16 backend loops supporting a minimum thin pool size of 64 physicals.
- The minimum number of SATA II disk drives assigned to a symdisk group for a thin pool is equal to 2 times the number of backend loops. This minimum threshold assures proper I/O distribution required to maintain a

high level of throughput performance. Example, a 2-engine VMAX has 16 backend loops supporting a minimum thin pool size of 32 physicals.

- EFD (flash) drives are supported for Fully Automated Storage Tiering (FAST) only. The number of EFD drives assigned to a symdisk group is equal to 1-4 times the number of backend loops. Each frame supports one EFD thin pool. For example, a 2-engine VMAX will support one EFD thin pool consisting of 16 to 64 drives equally distributed on all available loops.
- A maximum of 8 thin pools per array. Strive for a lesser number of pools with greater capacity per loop; providing higher backend I/O distribution.
- Thin pool capacity expansions with Fibre Channel and SATA II disk drives shall be accomplished by adding physical disks to the associated symdisk group. Expansions shall be accomplished in disk drive increments equal to the number of backend loops; all drives are equally distributed across all loops. Example, a 2-engine VMAX with 16 backend loops shall add minimum of 16 disk drives, one drive per loop to the associated symdisk group. Recommend that pool rebalancing be initiated after each capacity expansion.
- EFD thin pool expansion shall be accomplished in increments equal to the loop count. The maximum number of EFD per frame is equal to 4 times the loop count. For example, a 2-engine VMAX with one EFD per loop can be expanded up to 64 drives in increments of 16 drives.
- Management and administration of the VP environment can be accomplished through SYMCLI commands. Symmetrix Management Console (SMC) and/or Control Center are not required for deployment.
- **Target no higher than 70% capacity utilization per thin pool.** TDEVs allocations should be suspended when a thin pool reaches 70% utilization of the available capacity. This provides sufficient buffer for unexpected growth or “runaway” application(s) on the existing TDEVs. See section [20.5](#) for details on thin pool expansions.
- DATADEV and standard devices should not share the same symdisk group. Exceptions are made for quorum and gatekeeper devices.
- DATADEV are not directly host accessible.
- DATADEV does not support metavolumes.
- DATADEV constructed from Fibre Channel or SATA disk drives shall use RAID-6 14+2 data protection.

- DATADEV constructed from EFD (flash) drives shall use RAID-5 7+1. EFD drives are only supported for automated storage tiering.
- A thin pool can only contain DATADEV with one emulation type. The first DATADEV added to the thin pool defines the emulation type.
- DATADEV shall be created at 262,668 cylinders in size or as large as feasibly possible.
- The removal of DATADEV from an existing thin pool should be avoided. Removing DATADEV requires draining the existing data to the remaining DATADEV in the thin pool. The pace at which the data is moved from the source DATADEV is extremely slow as to not affect the performance of the pool.
- Dynamic Cache partitioning shall not be used for thin devices. Splitting of TDEV and DATADEV between cache partitions causes the Write Pending (WP) counts to double, which may lead to global cache threshold violations.
- Rebalance variance for an individual thin pool shall be no higher than 3%.

18.1.1. Data Migrations

- EMC does not currently support a frame based data aware migration tool. “Data aware” refers to the ability of an application to distinguish between legitimate (data written by an application) and illegitimate data (i.e. padding bytes, placeholder words or non-application written chunks). Host based data aware migration tools are currently available and provides a fairly efficient data migration mechanism.
- Block-level migration tools for thick to thin device migrations are supported (e.g. EMC Open Replicator). Block-level migration tools do not distinguish between used and unused application blocks; consequently, the entire source thick device capacity is transferred to the thin target device. Recommend running the Zero Space Reclamation process following a thick to thin device migration (see [Section 13](#)).

18.1.2. Data Protection

- All DATADEV created from Fibre Channel or SATA drives shall use RAID-6 14+2 data protection. RAID-6 provides the highest level of protection among supported RAID protection types.

- EFD DATADEVs, supported for automatic tiering only, shall use RAID-5 7+1 data protection.
- TDEVs do not utilize a protection scheme; it inherits the protection type from the DATADEV in the pool to which it is bound.

18.1.3. Pre-Allocation and Oversubscription

- All TDEV shall utilize a preallocation value of 0%. Preallocation does not remove the first write penalty, but allocation latency is slightly decreased. Exceptions are made for quorum, boot and swap devices that maybe 100% preallocated.
- **Oversubscription will be limited to 200%** of the actual thin pool capacity. Progressive steps in reaching the 200% oversubscription are supported. Example, configuring the thin pool at 125% oversubscription and moving to 150% and/or 175% is permissible.
- Capacity utilization of oversubscribed pools must be monitored to prevent an out-of-space condition. Thin pool utilization monitoring is accomplished through Solutions Enabler CLI, symevent log and Tivoli NetCool.

18.1.4. Thin Devices

- TDEV sizes of 36,828 through 262,668 cylinders are supported in increments of 36828 cylinders:

units	size (cyl)	size (bytes)	size (gigabytes)
1	36828	36203397120	33.7170
2	73656	72406794240	67.4341
3	110484	1.0861E+11	101.1511
4	147312	1.44814E+11	134.8682
5	184140	1.81017E+11	168.5852
6	220968	2.1722E+11	202.3022
7	262668	2.5821E+11	240.646

- CATE STORAGE AND BACKUP ENGINEERING recommends the use of fewer larger devices over many smaller TDEVs, thus reducing the cache resources and processing cycles to manage TDEVs.
 - TDEVs are cache entities that utilize 143KB per device plus 8KB for each gigabyte of device size. For example, a 36828 cylinder thin

device uses 422.5KB of cache space versus a 257,796 cylinder device uses 2120KB of cache space. The larger TDEV uses 8.224Bytes of cache per cylinder versus 11.47B per cylinder for the smaller TDEV, which represents a 28% cache space savings per cylinder.

- TDEV stripe metavolumes are recommended over concatenated volumes, due to its higher level of throughput performance. Initial EMC documentation recommended concatenating TDEV, due to the intrinsic ability to expand online. However, the superior performance obtained with plaiding TDEV negates the minimally used online expansion capabilities of concatenating TDEV. [Appendix F: Metavolumes – Striped vs. Concatenated](#)
- Metavolumes cannot be formed by combining thin and thick devices.
- TDEV Metavolumes shall be created with members from the same thin pool.

18.1.5. Thin-Hostile Environments

- Ability to recognize operating systems and applications that are thin hostile is essential to maintaining the overall health of the thin pool:
 - File systems that do not reuse the space associated with deleted files resulting in unused space and reducing or eliminating the capacity utilization benefits of VP. (See [section 11.3](#))
 - Applications or data managers that pre-write dense patterns of specific fill words to reserve the available capacity.
 - Workloads that generate high levels of write requests for prolonged durations. This type of workload causes the accumulation of write pendings within global cache, which may cause SRDF and Timefinder operations on other devices to halt, occurrence of system disconnects, and the dramatic reduction in I/O throughput. (See [Section 12.3.3](#)).

19. Auto-provisioning Guidelines

Auto-Provisioning Solutions Guide located at the following link:
https://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/platforms/Symmetrix/09034-sg.pdf

- **Symaccess will be the only method used to provisioning VMAX devices.** Symmask/db will continue to be used on DMX-4 earlier frames.
- Auto-provisioning requires that the ACLX port flag must be enabled on all FA ports. See [Appendix C](#).
- Each VMAX supports the creation of 8192 individual storage groups, port groups, initiator groups and masking views.
- A World Wide Name (WWN) can only belong to a single initiator group.
- Each initiator group can contain from 0 to 32 entries. Empty initiator groups allow the masking view for a host to be created without knowing the HBA information. The HBA information can be added to the initiator group when available that will promptly bring the storage online to the host.
- Initiator groups may be cascaded one level to another initiator group. Cascading of initiator groups supports the clustering of servers, where multiple initiator groups share the same storage devices.
- Volume dynamic addressing is enabled by default, where the LUN address is automatically determined when the masking view is created. A LUN address or start of a range of addresses may be manually specified with an additional command argument. (see [section 22](#))
- LUN offset on the array is no longer supported.
- Storage group names are limited to 64 characters and are not case sensitive.
- A storage group can contain up to 4096 devices.
- Devices can belong to multiple storage groups.

20. Drive Sparing

Spare drives are required for every Symmetrix VMAX system. The bin file configuration program prevents creation of Symmetrix VMAX configurations that does not contain the required amount and type of spare drives. All drive types must be considered where capacity and speed constitute different drive types. The rules for the amount and types of HDD spares are as follows:

- Two spare drives for every 100 disk drive type
- One spare drive for every 100 solid state drive type

- Minimum of eight spare drives for the entire system

20.1.1. Permanent Sparing


The Symmetrix array implements permanent drive sparing to increased data availability with little effect on performance and functionality. The array uses self-diagnostic information and active error checking to make drive sparing decisions. If excessive errors are detected indicating that a hard failure is imminent, the Symmetrix initiates a search for a suitable spare and the EMC Customer Support Center is alerted of the impending drive failure. The VMAX only supports permanent sparing to replace the failed resource with a suitable alternative.

Permanent sparing is often referred to as “permanent replacement” where the failing or failed drive is replaced with an online spare. When a failing drive is detected, the Symmetrix array locates a suitable spare that satisfies the following criteria:

- No two RAID-1 or RAID-5 members can share the same loop
- No more than 2 RAID-6 members on the same loop
- Vault devices must be spared on the same disk adapter as the failed drive
- Vault devices may use any disk drive type or combination of disk drive type

Permanent sparing does not require a mirror position to transfer data to the spare drive, which avoids potential conflicts with other Enginuity features (e.g. BCV). When the Symmetrix frame identifies a suitable spare, the data on the failing drive is copied or rebuilt to the new drive. The failed drive takes the profile of a not ready spare until replaced. The diagram below illustrates the RAID-1 drive sparing procedure.

Diagram 3: Symmetrix Permanent Drive Sparing



21. VMAX Virtual Provisioning

EMC Virtual Provisioning (VP) product is based on traditional thin provisioning functionality, which uses thin devices presented to the host while consuming physical storage from a shared storage pool. VP is a feature set supported on the VMAX array with Enginuity 587X Microcode and Solution Enabler +7. VP provides several distinct advantages over standard provisioning methods:

- Improves storage capacity utilization by presenting an application with sufficient storage and allowing the capacity to increase dynamically.
- Support oversubscription is where the cumulatively advertised capacity exceeds the physical capacity of storage available from the storage pool.
- Increased performance through the reduction in hotspots with greater backend I/O distribution.

VP utilizes a storage pool to group virtual devices called Thin Devices (TDEV) with the actual storage volumes called Data Devices (DATADEV), (see diagram 1). The TDEV is a cache entity that is viewed as a typical Symmetrix device, but functions like a pointer to DATADEV. A DATADEV is created from physical storage that adheres to the same configuration requirements as standard Symmetrix devices. The VP storage pool, often referred to as the thin pool, associates and manages the relationship between the TDEV and DATADEV. The thin pool is where the oversubscription value is set and utilization specifications are displayed. The existing thin pool capacity can be augmented online by adding or removing DATADEV; however, the removal of DATADEV is considered hazardous and time consuming.




Diagram 1:

21.1. I/O Functionality

VP widely stripes the data across all available DATADEV within the thin pool in a round-robin write scheme, spanning a greater number of physical spindles than standard devices. This distribution of workload across the entire thin pool provides for a consistent level of performance by reducing the probability of hot spots.

In a write request, VP allocates a minimum amount of physical storage from the thin pool, which is allocated sequentially regardless if the data is written sequentially or randomly to the TDEV. This allocation is performed in units of storage called device extents that are 12 tracks (768 KB) in size. A round-robin

mechanism is used to balance the allocation of device extents across all the DATADEV in the pool.

A read request retrieves the appropriate data from the DATADEV utilizing maps to locate the requested information in the storage pool. If read requests are generated to an area that has not been mapped, a block where each byte is equal to zero is returned. Reading from an unmapped area will not result in any new storage allocations.

21.2. Volume Table of Contents

Instant Volume Table of Contents (IVTOC) is the low-level formatting of tracks and is the primary underlining processes for the implementation of VP. IVTOC utilizes a logical and physical process to prepare and format the allocated capacity. When storage is allocated to a thin device, the logical process creates a table with the TDEV and DATADEV relational information. When a write request is submitted to the TDEV, the frame formats the needed capacity prior to writing the new data and updates the table of contents.

The initial thin device write request may have a slightly higher latency than standard devices due to the partial-track first write penalty. The first write process is where a write is performed to an area of a thin device that has not yet been mapped, at which time the logical VTOC process allocates an extent and updates the information table (IVTOC state). The partial-track first write penalty occurs when the first write fails to completely fill the first track (64KB) in the extent as required by VP, thus the backend disk adapter pads with zeros the remainder of the track.

When a TDEV is unbound from the storage pool, the relational information created during the IVTOC process is removed and the storage capacity is returned to the pool as available storage. The original data remains on the DATADEV following the unbind process; however, the low-level format is required prior to accessing the same disk locations.

21.3. Thin Pool Alerting

Virtual provisioning provides five alerting variables to inform the users of state changes and current status of virtual devices and associated pools:

- *Device Pool Configuration Change* – Alerts the user when the status of the pool has changed as a result of adding or removing data devices or binding or unbinding thin devices.
- *Device Pool Free Space* – Alerts the user when the amount of free space remaining in a device pool meets or exceeds a specific percentage.

- *Device Pool Status* – Alerts the user when the state of a pool changes (i.e. hardware failure).
- *Thin Device Allocation* – Alerts the user when the amount of allocated space on a thin device meets or exceeds a specific percentage.
- *Thin Device Usage* – Alerts the user when the amount of used space on a thin device meets or exceeds a specific percentage.

21.3.1. Capacity Alerting

Of primary importance for the proper operation of the thin provisioning is the ability to monitor the consumption level of the thin pool. Recommend that storage pools be capped (no new allocations) at 70% of the available pool capacity which provides a 30% buffer to handle existing host applications assigned to the pool. The Symmetrix places pool utilization level alert messages into the “syમેvent” log at the following fixed levels:

- 60% - issue a warning alert
- 65% - issue a minor alert
- 70% - issue a major alert
- 80% - issue a critical alert
- 100% - issue a fatal alert

According to EMC, pool utilization alert messages do **not** warrant a call home event. Therefore, the initiative to monitor and respond appropriately in the event of a pool fill condition is the total responsibility of Citi.

Depicted below are thin pool threshold alert messages taken from syમેvent log on a Solution Enabler server:

```
[fmt=evt] [evtid=1208] [date=2008-08-15T15:36:54] [symid=000190103110]
[TPDataPool=glab_ssd] [sev=warning] = Data Pool utilization is now 60 percent.
```

```
[fmt=evt] [evtid=1208] [date=2008-08-15T18:17:41] [symid=000190103110]
[TPDataPool=glab_ssd] [sev=minor] = Data Pool utilization is now 65 percent
```

```
[fmt=evt] [evtid=1208] [date=2008-08-15T18:20:56] [symid=000190103110]
[TPDataPool=glab_ssd] [sev=major] = Data Pool utilization is now 70 percent.
```

```
[fmt=evt] [evtid=1208] [date=2008-08-15T18:30:39] [symid=000190103110]
[TPDataPool=glab_ssd] [sev=critical] = Data Pool utilization is now 80 percent.
```

```
[fmt=evt] [evtid=1208] [date=2008-08-15T19:52:37] [symid=000190103110]
[TPDataPool=glab_ssd] [sev=fatal] = Data Pool utilization is now 100 percent.
```

21.4. Audit Logging and Alerting

The symaudit log entries can be viewed by issuing the 'symaudit' command. If symauth is enabled on the Symmetrix, the user issuing the symaudit command must have the role of SecurityAdmin, Auditor, or Admin.

21.4.1. Configuring Symmetrix Alerting

The storevntd daemon should be configured to record Symmetrix events to a file on the Solaris 10 Solutions Enabler management server. The following entries (in bold red font) in the **daemon_options** file should be created to configure the EVP SNMP alerts:

```
storevntd:log_event_targets = file

storevntd:log_event_file_name = symevents

storevntd:log_symmetrix_events = \
    sid=000190100919, audit;\
    sid=000190103110, audit;\
    sid=000190300717, audit;\
    status;\
events, tgt=file
```

Once these entries have been created in the daemon_options file, the storevntd should be stopped and restarted with the new entries by following these steps:

```
# stordaeon action storevntd -cmd reload
# stordaeon shutdown storevntd -immediate
# stordaeon start storevntd
```

21.4.2. Configuring Tivoli Agent

The Tivoli Endpoint agent will then be used to scrape the event file used by storevntd and exported to the Tivoli monitoring system for alert notification. (Solutions Enabler Guide) The following is a sample of Symmetrix events as captured by the Tivoli system. Note that the red highlighted text is the portion of the notification message that will be sent to the storage support team responsible for supporting that Symmetrix.

```
Citi_Hdw_San_Generic;source=APPLLOG;severity=CRITICAL;citi_probearg='1208
, [symid=000194900217] [TPDataPool=qtpool]
';hostname=sangs1;citi_profilename=CCC-CCC-CCC-EMC.D.app.prc-
logadapter;sub_source=EMC;msg='Data Pool utilization is now 100
percent. [symid=000194900217] [TPDataPool=qtpool]
';origin=10.40.42.33;END
Citi_Hdw_San_Generic;source=APPLLOG;severity=CRITICAL;citi_probearg='1111
, [symid=000194900217] ';hostname=sangs1;citi_profilename=CCC-CCC-CCC-
EMC.D.app.prc-logadapter;sub_source=EMC;msg='Save or data device pool is
full. [symid=000194900217] ';origin=10.40.42.33;END
Citi_Hdw_San_Generic;source=APPLLOG;severity=CRITICAL;citi_probearg='1115
, [symid=000194900217] ';hostname=sangs1;citi_profilename=CCC-CCC-CCC-
EMC.D.app.prc-logadapter;sub_source=EMC;msg='Save or data device pool is
almost full. [symid=000194900217] ';origin=10.40.42.33;END
Citi_Hdw_San_Generic;source=APPLLOG;severity=CRITICAL;citi_probearg='1208
, [symid=000194900217] [TPDataPool=qtpool]
';hostname=sangs1;citi_profilename=CCC-CCC-CCC-EMC.D.app.prc-
logadapter;sub_source=EMC;msg='Data Pool utilization is now 20 percent.
[symid=000194900217] [TPDataPool=qtpool] ';origin=10.40.42.33;END
```

21.5. Thin Pool Expansion

21.5.1. Determining When to Expand the Thin Pool

The expansion of an existing thin pool should only be selected after all other options are exhausted. Prior to the thin pool reaching 70% of capacity, a determination is required if the pool should be capped (prohibiting new allocations) and/or expanded, whereby reducing the probability of an out-of-space condition. To properly determine if a thin pool should be expanded, careful monitoring and data collection should be done prior to reaching the 70% utilization threshold. Thin pools should routinely be monitored once the 50 - 60% utilization is realized with the objective in determining the rate of increase (%utilization/time). The thin pools rate of increase will determine the steps needed to reduce the possibility of an out-of-space condition. The following three scenarios functions as guidelines available to the storage administrator (SA):

- If the pool allocation increases from 50 to 60% in 3 months, the pool should be immediately capped with all new TDEV allocations routed to new pools or to existing pool that are under 50% utilization (include pools from another array). With the thin pool capped and allocations continues to rapidly increase the migration of applications to new or lesser utilized thin pool(s) and/or the associated symdisk group augmented with additional disks may be necessary.

- If the pool allocation increases from 50% to 60% in 12 months the pool should be periodically monitoring for any out-of-ordinary increase in utilization. If no dramatic increase occurs, the pool should be capped at 70% and new allocations configured on lesser utilized thin pools.
- If the pool allocation increases from 50% to 60% in 24 months the pool should be periodically monitoring for any out-of-ordinary increase in utilization. If no dramatic increase occurs the pool should be capped when 75-80% utilization is reached.
- If the frame is at the maximum physical disk count and all thin pools approaching 50% fill with a fairly aggressive rate of increase the following options may be employed:
 - If the frame contains six engines or less and maximum drive count per that configuration is reached, an additional two engines and associated storage bays can be added.
 - If an 8-engine frame with 1920 drives is running low on capacity, recommend that another frame be purchased. If this is not possible based on the timeframe, floor space and/or cost, the existing 1920 drive count can be expanded to 2400 drives by extending certain loops to 45 drives per loop.
 - In a 2 or 4-engine frame the loop population may expand to 45 drives without additional engines added.

The thin pools rate of allocation is the primary factor in determining when a thin pool should be capped or expanded. However, other variables that make-up the VP environment such as lifecycle of the array, lifecycle of the applications and remaining storage resources are also determining factors. There is no one formula to determine when a thin pool should be capped and/or expanded; therefore, the array management team should consider all variables including loop expansion and/or new frame(s) to handle capacity expansions.

21.5.2. Thin Pool Expansion Guidelines:

VP does contain the ability to rebalance existing extents across the thin pool. Adding or removing DATADEV from an existing pool creates a utilization imbalance between the existing and new DATADEVs. The thin pool will utilize the least exhausted DATADEVs for new writes until utilization balance among all DATADEVs are restored. The process of adding DATADEVs to the thin pool should be accomplished in sufficient quantity and distribution to maintain the throughput performance of the thin pool. Once the DATADEV has been enabled

with a thin pool, rebalancing of the written extents is recommended. Thin pool rebalancing redistributes utilized extents among all available DATADEV.

Thin Provisioning Solutions Guide located at the following link:
[https://catecollaboration.citigroup.net/domains/platstor/CATE Storage and Backup Engineering/products/platforms/Symmetrix/09058-sg.pdf](https://catecollaboration.citigroup.net/domains/platstor/CATE%20Storage%20and%20Backup%20Engineering/products/platforms/Symmetrix/09058-sg.pdf)

If thin pool expansion is deemed necessary, the following requirements for adding new DATADEVs should be satisfied:

- Symdisk group associated with the thin pool shall be augmented with physical disks of the same profile.
- DATADEVs shall have the same protection type as the original devices in the thin pool.
- All DATADEVs in a pool should be of the same size. Using different size devices could result in uneven data distribution.
- The thin pool expansions shall be accomplished by increasing the number of physical drives within the associated symdisk group with a minimum physical disk increase equal to the loop count.
- Expansion of an existing thin pool can be accomplished non-disruptively; however this action should be done within the approved green zone.
- Thin pool rebalance process shall be run following each DATADEV addition.

22. Auto-provisioning Introduction

VMAX Auto-provisioning (AP), introduced with Enginuity 5874 and Solutions Enabler V7.X, reduces the complexity and streamlines the steps required to provision devices. AP utilizes the symaccess command set to manage the provisioning process, configuring host initiator flags, setting port security and monitoring host connectivity. Symaccess increases provisioning efficiency by utilizing containers to group together multiple elements to complete the mapping and masking process. EMC has positioned symaccess to be the sole provisioning mechanism on the VMAX frame; therefore, **CATE STORAGE AND BACKUP ENGINEERING mandates that all VMAX device provisioning utilize the symaccess command set.**

Auto-provisioning Elements	
Container	Description
Storage group	contains the devices (i.e. thick, thin, RDF) required by the host
Port group	contains the front-end ports that the devices are mapped
Initiator group	contains the initiator (e.g. WWN) of the host HBA
Masking View	associates the storage, port and initiator group together. Once the masking view is created the devices are automatically masked and mapped.

The use of AP for device provisioning will constitute the majority of the day-to-day operations. The table above depicts the symaccess building blocks necessary to complete the mapping and masking process:

The following table compares the traditional provisioning against the auto-provisioning process:

Operation	DMX symmask Approach	VMAX Auto-provisioning Approach
Initial Storage Provisioning to a Server	1. Identify next available channel address for each front-end port 2. Map devices to each front-end port 3. For each initiator to front-end port connection, create masking entry	1. Create Storage, port, and Initiator groups 2. Create masking view – all required mapping and masking is performed automatically
Adding Additional Capacity to an Existing Server	1. Identify next available channel address for each front-end port 2. Map devices to each front-end port 3. For each initiator to front-end port connection, create masking entry	1. Add additional devices to storage group – all required mapping and masking is performed automatically
Storage Reclamation when a Server or Application is Retired	1. For each initiator to front-end port connection, delete masking entry 2. For each front-end port, unmap each device	• Delete masking view – all required unmapping and unmasking is performed automatically

23. Physical Installation

Physical and environmental setup of all enterprise storage arrays should conform to vendor specifications and facilities provided by the local data center planning group. Please review “*EMC Symmetrix VMAX series Physical Planning Guide*” on EMC Powerlink site.

24. Standard Hardware Configuration

The certified Enginuity code level for each Symmetrix array is available on the CATE STORAGE AND BACKUP ENGINEERING website:

[http://catecollaboration.citigroup.net/domains/platstor/CATE Storage and Backup Engineering/products/techstdshmpgs/symmetrix.aspx](http://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/techstdshmpgs/symmetrix.aspx)

25. Dynamic LUN Addressing

On the Symmetrix VMAX with Auto-provisioning Groups, dynamic LUN addressing is an integral part of the provisioning process. With the DMX masking approach, a separate masking entry is configured for each path and potentially different LUN addresses could be applied to different paths to the same device. With Auto-provisioning groups, the masking view defines all paths between the initiators and front-end ports and the dynamic LUN address associated with the view. This ensures the same address is assigned to the device for all paths. Many cluster applications such as VMware ESX require the same LUN address for all paths. The dynamic LUN address feature eliminates the potential for inconsistent LUN addresses for the same device on different paths.

26. Hardware Replication

VMAX support replication with SRDF (Symmetrix Remote Data Facility) for remote replication and TimeFinder for replication within a frame. Please see the CATE STORAGE AND BACKUP ENGINEERING Solutions Guides for replication operations:

[https://catecollaboration.citigroup.net/domains/platstor/CATE Storage and Backup Engineering/products/platforms/Symmetrix/09039-SG.pdf](https://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/platforms/Symmetrix/09039-SG.pdf)

[https://catecollaboration.citigroup.net/domains/platstor/CATE Storage and Backup Engineering/products/platforms/Symmetrix/09038-SG.pdf](https://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/platforms/Symmetrix/09038-SG.pdf)

27. Symmetrix VMAX Management

Symmetrix VMAX Array is managed with the Solutions Enabler application. Please see the CATE STORAGE AND BACKUP ENGINEERING Solutions Guide for Solutions Enabler:

https://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/platforms/Symmetrix/07060-sg-v2.8.pdf

28. Vendor Remote Access to Symmetrix

Communication between the EMC Customer Support Center and the Symmetrix VMAX system occurs through EMC Secure Remote Support (ESRS). Connectivity is achieved through a dedicated server as a conduit for all communications between EMC and the site information structure. EMC provides Secure Remote Support Gateway software at no-charge to any customer with a maintenance agreement. Policy Manager software is also provided that allows customers to manage, control and log gateway actions. Please reference the following ESRS Solutions Guide for additional information:

https://catecollaboration.citigroup.net/domains/platstor/CATE_Storage_and_Backup_Engineering/products/services/ESRS/07006-sg_v1.95.pdf

29. Audit Logging

See the [Solutions Enabler Solutions Guide](#) for instructions on setting up Solutions Enabler to gather the symaudit logs and exporting them to eTrust.

30. Symmetrix Priority Controls

Symmetrix Priority Controls (SPC) is a Quality-of-Service (QoS) parameter that allows the performance metric per device or group of devices to be managed. Using the Solutions Enabler symqos command set, SPC provides 16 levels (1 highest – 16 lowest priority) of device prioritization through setting the Priority Service Time (PST). PST prioritizes a device or device group by increasing or decreasing the potential I/O delay of the request.

EMC defaults the PST levels for all applicable devices to the highest level of 1 for the entire frame; therefore, utilize PST to set the performance metric between 2 and 16. Since PST levels are relative to each other and dependent upon the load on the array, recommend that PST levels be lowered in 1-3 level increments until an acceptable level of performance is achieved.

PST levels are **not** supported on SFS, VAULT, SAVE, DRV and VDEV.

31. Enabling ECC Support

Please refer to [Solutions Enabler Solutions Guide](#) for instruction on configuring the Symmetrix to work with ECC.

32. Symacl Symmetrix Setup

Refer the [Solutions Enabler Solutions Guide](#) for instructions on symacl setup.

33. Functional Identification

Refer the [Solutions Enabler Solutions Guide](#) for available functional identification.

34. VMAX Service Processor Anti-Virus Software Management

The VMAX service processor comes with an OEM version of McAfee VirusScan. Traditionally, the virus definitions in VirusScan only get updated manually by the EMC CE during a Symmetrix microcode upgrade. To automate the upgrade of virus definitions and VirusScan software, it is recommended that the McAfee ePolicy (ePO) Common Management Agent (CMA) be installed on the service processor and configured to use Citi's ePO infrastructure for VirusScan software and virus definition updates. By integrating the VMAX service processor with Citi's ePO infrastructure, the following can be accomplished:

- The McAfee VirusScan software can be upgraded to Citi's supported version.
- Virus definitions are automatically updated and current.

The following steps outline the procedure for integrating the VMAX service processor with Citi's ePO infrastructure:

- 1) Create firewall rules for ePO communication.

Since the VMAX service processors are connected to a CEP IP network, firewall rules will need to be created to allow IP communication between the VMAX service processors and the ePO server. The following rules will need to be created in the firewall:

Source	Destination	Protocol/Port Numbers
VMAX service processor	ePO server	TCP/18080 TCP/18443

ePO server	VMAX service processor	TCP/18081 UDP/18082
------------	------------------------	------------------------

- 2) The same procedure used to install the ePO agent on the VMAX SP can be used to install the ePO agent on the VMAX. The procedure for ePO agent installation is contain in the document [Configuring VirusScan Exclusions for EMC Symmetrix VMAX Storage Arrays via ePO](http://catecollaboration.citigroup.net/domains/security/isp_standards/hsestds/Antivirus%20%20McAfee%20ePolicy%20Orchestrator/AV%20Documentation/ePOusefulinfo_contacts.aspx) created by Host Security Engineering for creating the policies in the ePO server for the VMAX storage platform. The regional storage operations teams should coordinate with the local ePO administrator for ePO server configuration and with their local EMC support team for the ePO agent installation on the VMAX service processor. The ePO contact for each region can be found at http://catecollaboration.citigroup.net/domains/security/isp_standards/hsestds/Antivirus%20%20McAfee%20ePolicy%20Orchestrator/AV%20Documentation/ePOusefulinfo_contacts.aspx. The following should be done after following the instructions in the document:
 - The pre-configured policies created by HSE for the VMAX storage array should be imported and installed in the local ePO server.
 - The FramePkg.exe file obtained from the ePO administrator for the site should be given to EMC. EMC will use these files to install the ePO agent software on the service processor. EMC has created an internal procedure for the ePO agent install and should obtain that procedure through their internal resources.
- 3) Once the ePO policies have been installed on the ePO server and the ePO agent has been installed and configured on the VMAX service processor, the VMAX's ePO agent report to the ePO server and should show as an end node and can then be assigned to the container that contains the ePO policies for the VMAX storage array.
- 4) The VirusScan software on the VMAX service processor should be upgraded to a Citi supported version via the ePO agent. This procedure is the same as the procedure for the VMAX model. Instructions for upgrading the VirusScan software are included in the document [Configuring VirusScan Exclusions for EMC Symmetrix VMAX Storage Arrays via ePO](http://catecollaboration.citigroup.net/domains/security/isp_standards/hsestds/Antivirus%20%20McAfee%20ePolicy%20Orchestrator/AV%20Documentation/ePOusefulinfo_contacts.aspx).
- 5) EMC should verify that the ePO agent can successfully receive virus definition updates and that the VirusScan software has been successfully upgraded to the Citi supported version. The ePO policies should also be validated by the ePO administrator.

35. Channel Adapter Flag Settings

To ensure proper host to array device communications a common protocol is required between the initiator (host) and target (array). Incompatibility between initiator and target protocol may lead to incorrect volume identifiers, limited device range, and erratic/complete communication failure. To further complicate initiator and target protocol compatibility, most SAN environments utilize heterogeneous mode, where paths from different types of host operating systems share the same physical port.

Currently, there are 42 flag variables on each VMAX front-end port, of which 22 are modifiable by the user. See complete list of modifiable and non-modifiable flags below:

SCSI Flags	
Modifiable	Non-Modifiable
Negotiate_Reset(N)	Tagged_Commands(T)
Soft_Reset(S)	Linked_Commands(L)
Environ_Set(E)	Sync_Transfer(Y)
HP3000_Mode(B)	Wide_Transfer(W)
Common_Serial_Number(C)	Cyl_Count_In_Name(P)
Disable_Q_Reset_on_UA(D)	PBAY_Monitor(Q)
Sunapee(SCL)	Command_Reordering(R)
Siemens(S)	Set_Qerr(Z)
Sequent(SEQ)	Disable_Ultra(U)
Avoid_Reset_Broadcast(ARB)	Disable_False_Disconnect(
Server_On_AS400(A4S)	Disable_Mini_Q(DMQ)
SCSI_3(SC3)	Avoid_Force_Negotiate(AFN)
SPC2_Protocol_Version(SPC2)	Auto_Busy(AB)
SCSI_Support1(OS2007)	Disable_Interleaved_Cmds
Fibre Channel Flags	
Modifiable	Non-Modifiable
Volume_Set_Addressing(V)	Disk_Array(A)
Non_Participating(NP)	Hard_Addressing(H)
Init_Point_to_Point(PP)	Global_3rdParty_Logout(TP)
Unique_WWN(UWN)	Generic_VSA(GVSA)
Access_Logix(ACLX)	Class_2_Service
OpenVMS(OVMS)	
AS400(AS4)	
Auto_Negotiate(EAN)	

Operating systems and infrastructure software continues to evolve leading the standard committees to release new flags for the array vendors to implement. To maintain compatibility between manufactures and difference versions, the “American National Standard of Information systems” defines behaviors that are common to all SCSI devices. Implementation of these SCSI behaviors requires certain flags to be defined between the initiator and target to distinguish the request-response model and descriptor blocks supported. Correlating the flag requirements between the various manufactures’, EMC created the E-Lab Navigator tool available through the PowerLink website. E-Lab allows the user to determine to proper flag setting by specifying the operating system, infrastructure software and Engenuity environment. (See [section 36](#), Appendix A for the flag compatibility table)

Based upon the compatibility tests conducted by Citi and EMC, the following table contains the flags supported by the majority of the modern operating systems certified within Citi. The enabling of these nine flags depicted in the following table represents the default FA flag settings for all new VMAX array installs:

SCSI Flags Enabled
Common_Serial_Number(C)
Disable_Q_Reset_on_UA(D)
SCSI_3(SC3)
SPC2_Protocol_Version(SPC2)
SCSI_Support1(OS2007)
Fibre Channel Flags Enabled
Init_Point_to_Point(PP)
Unique_WWN(UWN)
Access_Logix(ACLX)
Auto_Negotiate(EAN)

Enabling of the C, D, SC3, PP, UWN, ACLX and EAN flags has been the recommended default settings for several generations of Symmetrix arrays. Enabling of the SCSI Primary Commands 2nd generation (SPC-2) and SCSI compliance 2007 (OS2007) flags has been added with the VMAX array certification:

- The InterNational Committee for Informational Technology Standards (INCITS) T10 Technical Committee defined the SPC-2 standard in 2001. This standard defines the device model and commands that are basic to all SCSI devices. Therefore, all current Citi certified operating systems supports enabling the SPC-2 flag as strict requirement or as an optional setting.
- The OS2007 flag implements specific feature sets defined in SCSI Primary Commands 3rd generation (SPC-3); circa 2005. The OS2007 flag are required by the Windows and AIX operating systems, an optional for Linux,

HP-UX and ESX and Solaris operating systems. The SPC-2 flag must be enabled before enabling the OS2007 flag.

Note: Solaris OS and OS2007 Compatibility

An issue concerning Solaris operating system and the OS2007 flag has been exposed by EMC. The Sun SFS driver initiates the Unit Attention (UA) sense key when a new device is presented. This causes an “I_T Nexus loss” alert message to be generated in the log file and possible increase in I/O response time. This UA process is mandated by the T10 SCSI device model for host rescans following device inventory changes. The OS2007 flag moves toward the industry goal of dynamically scanning and configuring LUNs without interaction.

Enabling the SPC-2 flag on existing VMAX arrays requires a momentary shutdown of the FA port. This shutdown causes all paths accessing the FA port to be temporarily offline for several minutes. Due to the downtime involve and/or the sheer volume of existing FA ports, the process of enabling the SPC-2 flag on existing VMAX maybe prohibitive; therefore, enabling the SPC-2 flag on through the initiator group override feature maybe the only option available . The initiator override allows all WWN’s in an initiator group to be set contrary to the FA port flag settings. For instance, the initiator group override capability allows the SPC-2 and OS2007 flags to be set contrary to the setting on the FA port. The advantage of the initiator group override removes the necessity of shutting the down the FA port to enable the SPC-2 flag. The disadvantage is the need to reboot the host that requires flag settings contrary to the FA port.

36. Appendix A – Port and Initiator Flag Settings and Procedures

36.1. E-Lab Navigator - VMAX Channel Adapter Flag Settings

Operating System -Director Flag Settings				
Operating System	Topology	Required Director Bit Setting	Optional Director Bit Setting	Storage Operating Environment
Microsoft Windows 2003 [x86] : Standard Edition, Enterprise Edition	FC-AL; FC-SW	(C) Common Serial Number; (OS2007) Host SCSI Compliance 2007; (SC3) SCSI-3 compliance; (SPC-2) SPC-2 compliance		EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Microsoft Windows 2003 [x64] : Standard Edition, Enterprise Edition				
Microsoft Windows 2008 [x86]: Standard Edition, Enterprise Edition				

Microsoft Windows 2008 [x64]: Standard Edition, Enterprise Edition				
Operating System	Topology	Required Director Bit Setting	Optional Director Bit Setting	Storage Operating Environment
Sun Solaris: 10, 8, x86 1059	FC-AL; FC-SW	(C) Common Serial Number; (D) Disable Queue Reset on Unit Attention; (UWN) Unique World Wide Name	(ACLX) Initiator Group; (OS2007) Host SCSI Compliance 200758; (SC3) SCSI-3 compliance; (SPC-2) SPC-2 compliance	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Operating System	Topology	Required Director Bit Setting	Optional Director Bit Setting	Storage Operating Environment
VMware ESX: 3.5, 4.0 (vSphere 4)	FC-AL; FC-SW	(C) Common Serial Number; (SPC-2) SPC-2 compliance	(OS2007) Host SCSI Compliance 2007; (SC3) SCSI-3 compliance	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
VMware ESXi: 3.5, 4.0 (vSphere 4),4.1				
Operating System	Topology	Required Director Bit Setting	Optional Director Bit Setting	Storage Operating Environment
Hewlett Packard HP-UX 11i v3 (HP-UX 11.31)	FC-SW	(C) Common Serial Number; (PP) Enable Point-to-point; (SPC-2) SPC-2 compliance; (V) Enable Volume Set Addressing	(ACLX) Initiator Group; (OS2007) Host SCSI Compliance 200760	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Hewlett Packard HP-UX 11i v3 (HP-UX 11.31)	FC-SW	(C) Common Serial Number; (V) Enable Volume Set Addressing		EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Hewlett Packard HP-UX 11i: v1.0 (HP-UX 11.11), v2 (HP-UX 11.23)	FC-SW	(C) Common Serial Number; (PP) Enable Point-to-point; (V) Enable Volume Set Addressing		EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Hewlett Packard HP-UX 11i: v1.0 (HP-UX 11.11), v2 (HP-UX 11.23)	FC-SW	(C) Common Serial Number; (PP) Enable Point-to-point; (V) Enable Volume Set Addressing	(ACLX) Initiator Group	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183

Hewlett Packard Open VMS: V7.3-2, V8.2, V8.2-1, V8.3, V8.3-1H1	FC-SW	(OVMS) Open VMS; (PP) Enable Point-to-point; (SC3) SCSI-3 compliance; (UWN) Unique World Wide Name	(OS2007) Host SCSI Compliance 2007; (SPC-2) SPC-2 compliance	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Operating System	Topology	Required Director Bit Setting	Optional Director Bit Setting	Storage Operating Environment
Red Hat RHEL [32-Bit]: 5.0, 5.1, 5.2, 5.3, 5.4; Red Hat RHEL: [32-bit] 5.5 ⁷ , [x86_64] 4.7 AS/ES, [x86_64] 4.8 AS/ES ⁶ , [x86_64] 5.0, [x86_64] 5.1, [x86_64] 5.2, [x86_64] 5.3 ¹⁷ , [x86_64] 5.4 ¹ , [x86_64] 5.5 ¹ ;	FC-AL; FC-SW	(C) Common Serial Number; (D) Disable Queue Reset on Unit Attention; (SC3) SCSI-3 compliance; (SPC-2) SPC-2 compliance	(OS2007) Host SCSI Compliance 2007	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183
Operating System	Topology	Required Director Bit Setting	Optional Director Bit Setting	Storage Operating Environment
IBM AIX: 5.3, 6.1	FC-AL; FC-SW	(ACLX) Initiator Group; (FA LOOP ID) FA Loop ID12; (OS2007) Host SCSI Compliance 2007; (PP) Enable Point-to-point; (SC3) SCSI-3 compliance; (SPC-2) SPC-2 compliance; (UWN) Unique World Wide Name	(C) Common Serial Number13	EMC Enginuity 5874: .121.102, .122.103, .123.104, .157.129, .207.166, .210.168, .229.182, .230.183

36.2. Configuring SPC-2 and OS2007 Flags per FA Port

The following procedures depict how to enable the SPC-2 and OS2007 flags on an FA port through configuration manager. The SPC-2 flag requires removing the port from service prior to changing the setting.

36.2.1. Setting the SPC-2 Flag:

Turn off FA port: <pre>symcfg -sid 1380 -sa 8f -p 1 offline -v</pre> Execute a port 'Offline' operation for director 8f, port 1? in Symmetrix unit '000192601380' (y/[n]) ? y A port 'Offline' operation execution is in progress for Symmetrix unit '000192601380'. Please wait... Offline FE director(s) on local Symmetrix..... The 'OFFLINE' control operation SUCCEEDED. The port 'Offline' operation successfully executed for Symmetrix Unit '000192601380'.
--

Set FA flag for SPC-2 Enable:

A Configuration Change operation is in progress. Please wait...

```
Establishing a configuration change session.....Established.
Processing symmetrix 000192601380
{
  set port 8F:1 SPC2_Protocol_Version(SPC2)=Enable;
  set port 9E:1 SPC2_Protocol_Version(SPC2)=Enable;
}

Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 019 of 047 steps.....Executing.
Step 056 of 080 steps.....Executing.
Step 056 of 080 steps.....Executing.
Step 061 of 080 steps.....Executing.
Step 066 of 080 steps.....Executing.
Step 066 of 080 steps.....Executing.
Local: COMMIT.....Done.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

Turn of FA port:

```
symcfg -sid 1380 -sa 8f -p 1 online -v
```

```
in Symmetrix unit '000192601380' (y/[n]) ? y
```

```
A port 'Online' operation execution is
in progress for Symmetrix unit '000192601380'. Please wait...
  Online FE director port(s) on local Symmetrix.....Done.

The 'ONLINE' control operation SUCCEEDED.
```

36.2.2. Setting the OS2007 flag:

A Configuration Change operation is in progress. Please wait...

```
Establishing a configuration change session.....Established.
Processing symmetrix 000192601380
{
  set port 8F:1 SCSI_Support1(OS2007)=Enable;
  set port 9E:1 SCSI_Support1(OS2007)=Enable;
}

Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 014 of 047 steps.....Executing.
Step 048 of 080 steps.....Executing.
Step 056 of 080 steps.....Executing.
Step 059 of 080 steps.....Executing.
Step 061 of 080 steps.....Executing.
Step 062 of 080 steps.....Executing.
Step 066 of 080 steps.....Executing.
Step 075 of 080 steps.....Executing.
Local: COMMIT.....Done.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

36.2.3. Verify the Port Flag Settings:

```
symcfg -sid 1380 list -FA 7e -p 1 -v
```

Director Identification: FA-7E

Director Type : FibreChannel (563)
Director Status : Online

Number of Director Ports : 2
Director Ports Status : [ON,ON,N/A,N/A]
Director Connection Status : [Yes,Yes,N/A,N/A]

Director Symbolic Number : 07E
Director Numeric Number : 71
Director Slot Number : 7

Director Port: 1

WWN Node Name : 5000097208159000
WWN Port Name : 5000097208159119

Fibre Channel Loop ID : N/A

Fibre Adapter Type : N/A
Negotiated Speed (GB/Second) : 4

SCSI Flags

```
{
  Negotiate_Reset(N)      : Disabled
  Soft_Reset(S)           : Disabled
  Environ_Set(E)          : Disabled
  HP3000_Mode(B)          : Disabled
  Common_Serial_Number(C) : Enabled
  Disable_Q_Reset_on_UA(D) : Enabled
  Sunapee(SCL)            : Disabled
  Siemens(S)              : Disabled
  Sequent(SEQ)            : Disabled
  Avoid_Reset_Broadcast(ARB) : Disabled
  Server_On_AS400(A4S)    : Disabled
  SCSI_3(SC3)             : Enabled
  SPC2_Protocol_Version(SPC2) : Enabled
  SCSI_Support1(OS2007)   : Enabled
}
```

Fibre Specific Flags

```
{
  Volume_Set_Addresssing(V) : Disabled
  Non_Participating(NP)    : Disabled
  Init_Point_to_Point(PP)  : Enabled
  Unique_WWN(UWN)          : Enabled
  Access_Logix(ACLX)       : Enabled
  OpenVMS(OVMS)            : Disabled
  AS400(AS4)               : Disabled
  Auto_Negotiate(EAN)      : Enabled
}
```

36.3. Configuring SPC-2 and OS2007 Flags per Initiator

36.3.1. Enabling SPC-2 flag: per initiator group

```
# symaccess -sid 1380 -name sangs11_ig -type init set ig_flag on spc2 -enable
```

View HBA Flag Settings:

```
# symaccess -sid 1380 show sangs11_ig -type init -detail
```

```
Symmetrix ID           : 000192601380

Initiator Group Name    : sangs11_ig
Last updated at        : 10:57:44 AM on Mon Apr 26,2010

Port Flag Overrides    : Yes
  Enabled               : SPC2_Protocol_Version(SPC2)
  Disabled              : N/A
Consistent Lun         : No

  Originator Port wwn   : 10000000c9517ff5
  User-generated Name   : sangs11/port_A
  FCID Lockdown         : No
  Heterogeneous Host    : No
  Port Flag Overrides   : Yes
    Enabled              : SPC2_Protocol_Version(SPC2)
    Disabled             : N/A
  CHAP Enabled          : N/A
  Type                  : Fibre

  Originator Port wwn   : 10000000c95172cf
  User-generated Name   : sangs11/port_B
  FCID Lockdown         : No
  Heterogeneous Host    : No
  Port Flag Overrides   : Yes
    Enabled              : SPC2_Protocol_Version(SPC2)
    Disabled             : N/A
  CHAP Enabled          : N/A
  Type                  : Fibre
```

36.3.2. Enabling SPC-2 Flag per WWN

```
symaccess -sid 1380 -wwn 10000000c9517ff5 set hba_flags on SPC2 -enable
```

View HBA Flag Settings:

```
symaccess -sid 1380 show ig_sangs11 -type initiator -detail
```

```
Symmetrix ID           : 000192601380

Last updated at        : 08:50:05 PM on Wed May 13,2009

Initiator Group Name    : ig_sangs11
```

```

Originator Port wwn : 10000000c95172cf
User-generated Name : 10000000c9517ff5/10000000c9517ff5
FCID Lockdown      : No
Heterogeneous Host : No
Port Flag Overrides : No
CHAP Enabled       : N/A
Type               : Fibre

Originator Port wwn : 10000000c9517ff5
User-generated Name : 10000000c9517ff5/10000000c9517ff5
FCID Lockdown      : No
Heterogeneous Host : No
Port Flag Overrides : Yes
  Enabled          : SPC2_Protocol_Version(SPC2)
  Disabled         : N/A
CHAP Enabled       : N/A
Type               : Fibre

```

36.3.3. Enabling OS2007 Flag per WWN

```
symaccess -sid xxx -wwn 10000735c942edfe set hba_flag on os2007 -enable
```

Note - enable the SPC2 prior to enabling the OS2007.

36.4. Host Verification of the SPC-2 and OS2007 Flag Setting

The vital product data query through the “inq” command displays the SPC setting in byte 2 and OS2007 in byte 6. For the SPC-2 in byte 2, ANSI Vers=0x04 indicates the SPC-2 flag is enabled, and ANSI Vers = 0x03 indicates that the older version SPC is enabled. For OS2007 in byte 6, Reserve=0x10 indicates the OS2007 is enabled, and 0x00 means that it is disabled:

```
# inq -page0 -dev /dev/vx/rmdp/emcpower8s2
Inquiry utility, Version V7.3-956 (Rev 1.0)      (SIL Version V7.0.1.0 (Edit Level 956)
Copyright (C) by EMC Corporation, all rights reserved.
For help type inq -h.
```

```
STANDARD INQUIRY PAGE ---- RAW DATA FORMAT ---- STANDARD INQUIRY PAGE
```

```

-----
      00 01 02 03 04 05 06 07  08 09 0A 0B 0C 0D 0E 0F
-----
0000: 00 00 04 32 8f 00 10 0a  45 4d 43 20 20 20 20 20  ...2....EMC
0010: 53 59 4d 4d 45 54 52 49  58 20 20 20 20 20 20 20  SYMMETRIX
0020: 35 37 37 33 39 32 31 35  45 30 30 30 51 00 00 08  57739215E000Q...
0030: d1 00 00 00 00 00 00 00  00 00 00 30 0d 9c 09 00  .....0....
0040: 02 60 01 9c 00 00 00 00  00 00 00 00 00 00 00 00  .`.....
0050: 00 00 00 00 00 00 00 00  57 73 01 63 00 05 17 11  .....Ws.c....
0060: 00 00 00 00 30 00 00 00  13 00 00 00 54 10 c2 32  ...0.....T..2
0070: ff 80 00 00 00 00 00 00  51 00 2f 00 01 90 30 12  .....Q./...0.
0080: 92 00 65 00 30 31 35 45  44 30 00 00 00 00 00  ..e.015ED0....
-----

```

```

-----
BYTE 00: Peripheral Qualifer = 0x00      Peripheral Device Type      = 0x00

```

BYTE 01:	Removeable Media	= 0x00	Device Type Modifier	= 0x00				
BYTE 02:	ISO Ver	= 0x00	ECMA Ver	= 0x00	ANSI Vers	= 0x04		
BYTE 03:	Async event Notify	= 0x00	Terminate I/O Process	= 0x00				
BYTE 03:	Reserved	= 0x03	Response data format	= 0x02				
BYTE 04:	Additional Length		= 0x8f					
BYTE 05:	Reserved (CCS-Vendor Unique)	= 0x00						
BYTE 06:	Reserved		= 0x10					
BYTE 07:	RelAdr	WBus32	WBus16	Sync	Linked	R	CmdQue	StRe
	0	0	0	0	1	0	1	0

37. Appendix B – Gatekeeper Configuration

A Gatekeeper is a small Symmetrix device that simply functions as a SCSI target for syscalls and obeys the rules for standard disk I/O operations. The Solutions Enabler software considers any device at 10 cylinders or less to be a preferred gatekeeper device. Syscalls are encapsulated in Command Descriptor Blocks (CDB); therefore, no data is written onto or read from the Gatekeeper device.

Symmetrix syscalls are defined as a set of instructions to be interpreted by the Enginuity microcode for array control functions. Syscalls access and modify the Symmetrix internal states, statistics, and services from external clients, such as a standard Unix server. Syscalls fall into two categories:

- Passive Syscalls - request data from the Symmetrix
- Active Syscalls - request that the Symmetrix perform an action

Syscalls use the existing host channel/paths and operational codes to communicate with the Symmetrix array. Any disk device can be the target of a syscall, and Enginuity will decode the command string and perform the embedded array control instructions. However, using just any disk device is not optimal. Syscalls, like all I/O, use a locking mechanism to signal the device is in use. If syscalls are directed to actual data volumes, there is potential for a bottleneck between syscall I/O and actual I/O traffic simultaneously deploying locks on the in-band communication link. To avoid these types of locking conflicts, special devices known as gatekeepers are identified to be the target of all syscall communications.

The multipathing of gatekeepers is supported with EMC PowerPath. PowerPath is a host resident software application that manages host-to-storage I/O paths to maximize the performance and availability of server applications. PowerPath is able to distinguish between syscalls to a frame and standard disk I/O communications. The ability to distinguishing syscalls communications allows PowerPath to reserve the same path for active syscalls. Path reservation eliminates syscall failures due to movement of initial syscall path to the secondary path for load balancing. The frame will acknowledge the syscall down the initial path even if the initial syscall was moved to a secondary path. Having the syscall initiator and return acknowledgement on different paths causes the

embedded command in the syscall to fail. The path reservation technique is not supported on MPIO, MPXIO and DMP.

As of Solutions Enabler 6.0, the base daemon creates a pool of devices equal to 50% of the available gatekeepers. Predefining a pool of gatekeepers reduces the latency in establishing the in-band communications link. The pool assigns Gatekeeper(s) upon syscall requests. If more gatekeepers are needed, the base daemon will temporally grab gatekeeper(s) in the remaining 50% to satisfy the syscall requirements.

37.1. Gatekeeper Device Configuration Guidelines:

- Solutions Enabler host to contain 16 Gatekeepers from each frame for basic operation. The number of Gatekeepers dictate the logical paths between host and frame:
 - 10 for Solutions Enabler commands and configuration manager
 - 1 for each daemon. There are usually 6 daemons initiated per SE server.
 - 6 for SMC and Ionix Control Center (if needed)
- A gatekeeper should only be mapped/masked to one host.
- Gatekeepers are typically 6 (5760KB) cylinders in size. The maximum size should not extend beyond 10 cylinders, as the inquiry Gatekeeper flag will be avoided. The minimum Gatekeeper size should be large enough to contain a host volume label or signature.
- All gatekeeper type devices should use RAID-1 protection.
- EMC PowerPath supports gatekeeper multipathing.

37.2. Gatekeeper CLI Commands:

The following depicts how to display the Gatekeeper candidates and which devices are locked due to Gatekeeper operations:

List available Gatekeeper candidates:

```
# stord daemon action storapid -cmd list_gks <symid>
```

S Y M M E T R I X G A T E K E E P E R S

Symmetrix ID: 000192601380

GK Select : Disabled

Priority	State	Flags	PdevName
-----	-----	-----	-----
0	RU		/dev/rdisk/emcpower64c
1	RU		/dev/rdisk/emcpower65c
2	OP		/dev/rdisk/emcpower66c

```

3    OP    /dev/rdisk/emcpower67c
4    CL    /dev/rdisk/emcpower68c
5    CL    /dev/rdisk/emcpower69c
6    CL    /dev/rdisk/emcpower70c
7    CL    /dev/rdisk/emcpower71c
8    CL    /dev/rdisk/c1t5000097208159119d1s2
9    CL    /dev/rdisk/c1t5000097208159119d2s2

```

Legend for State:

CL = closed
 OP = opened as a pooled gatekeeper
 RU = opened as a repeat-use gatekeeper
 IN = device is invalid

Legend for Flags:

A = device is in the gkavoid file
 S = device is in the gkselect file
 D = device will be deleted when closed
 F = device failed to open
 > = device is selected for use by another application
 ~ = device was abandoned during open

Display Gatekeeper (GK) semaphores:

```
# symcfg list -semaphores
```

S Y M A P I S E M A P H O R E S

ID	Key	State	Lock Proc. Type Wait	Lock Full Path Specification
-----	-----	-----	-----	-----
0	0x45001168	Unlocked	GK	0 /dev/rdisk/c1t5006048AD5F11B00d0s2
6	0x45000158	Locked	GK	0 /dev/rdisk/emcpower21c
7	0x45000168	Locked	GK	0 /dev/rdisk/emcpower22c
8	0x45000178	Locked	GK	0 /dev/rdisk/emcpower23c
10	0x45000188	Unlocked	GK	0 /dev/rdisk/emcpower24c
2	0x45000408	Locked	GK	0 /dev/rdisk/emcpower64c
3	0x45000418	Locked	GK	0 /dev/rdisk/emcpower65c
4	0x45000428	Locked	GK	0 /dev/rdisk/emcpower66c
5	0x45000438	Locked	GK	0 /dev/rdisk/emcpower67c
9	0x45000448	Unlocked	GK	0 /dev/rdisk/emcpower68c

38. Appendix C - ACLX Database/Device

38.1. ACLX Database

The symaccess auto-provisioning module provides the ability to assign and mask access privileges of devices through the creation of masking views. A masking view defines the host to Symmetrix device connectivity, where this information is stored in the ACLX database. The ACLX database resides in the Symmetrix File System (SFS) of the VMAX array and is not directly accessible.

38.2. ACLX Device

The ACLX device functions as the initial communications medium upon initial turn-up of the VMAX frame. Once gatekeepers are provisioned to the Solutions Enabler management servers, the ACLX device no longer needs to be visible to any host. Auto-provisioning (symaccess) requires an ACLX device to be present on the frame and ACLX port flag enabled on each channel adapter (FA port) for it to be added to a symaccess port group. Note, the ACLX device and ACLX port flag are different entities, where the ACLX device does not need to be mapped to a port for the ACLX port flag to be enabled on that port. The following list describes the attributes of the ACLX device and ACLX port flag:

38.3. ACLX Device and Port Flag Specifications:

1. ACLX device usually assigned device identifier between 0008 and 0054
2. ACLX device is 6 cylinders (5760KB) in size.
3. Auto-provisioning requires an ACLX device to be present on the frame.
4. ACLX device will function in either write disabled or read/write enabled configuration.
5. ACLX port flag needs to be enabled on each FA port utilizing auto-provisioning.
6. ACLX device does not need to be visible to the application servers.
7. ACLX device shall not be used as a storage device.
8. PowerPath does not use the ACLX device.

38.4. ACLX Device Host Visibility

By default the ACLX device is mapped on each FA port. Once gatekeeper devices are provisioned to the Solutions Enabler management server(s), the appearance of the ACLX device on each FA port is no longer required. To minimize the confusion between the ACLX device and standard storage devices, the ACLX device should be unmapped from each FA port. The ACLX device maybe unmapped from all FA ports using the following procedure:

1. Verify gatekeepers are properly provisioned to the Solutions Enabler management server:

```
symcfg list -semaphores
```

S Y M A P I S E M A P H O R E S						
ID	Key	State	Lock Type	Proc. Wait	Lock Full Path Specification	
0	0x45001168	Unlocked	GK	0	/dev/rdisk/c1t5006048AD5F11B00d0s2	
3	0x45000158	Locked	GK	0	/dev/rdisk/emcpower21c	
5	0x45000168	Locked	GK	0	/dev/rdisk/emcpower22c	
6	0x45000178	Locked	GK	0	/dev/rdisk/emcpower23c	
9	0x45000188	Locked	GK	0	/dev/rdisk/emcpower24c	
2	0x45000408	Locked	GK	0	/dev/rdisk/emcpower64c	
4	0x45000418	Locked	GK	0	/dev/rdisk/emcpower65c	
7	0x45000428	Locked	GK	0	/dev/rdisk/emcpower66c	
8	0x45000438	Locked	GK	0	/dev/rdisk/emcpower67c	
11	0x45000448	Locked	GK	0	/dev/rdisk/emcpower68c	
13	0x45000458	Unlocked	GK	0	/dev/rdisk/emcpower69c	
14	0x45000468	Unlocked	GK	0	/dev/rdisk/emcpower70c	
15	0x45000478	Unlocked	GK	0	/dev/rdisk/emcpower71c	

2. Verify the device identifier of the ACLX device:

```
symdev -sid 1380 list -aclx
```

Symmetrix ID: 000192601380

Device Name		Directors		Device		
Sym	Physical	SA :P DA :IT	Config	Attribute	Sts	Cap (MB)
0008	Not Visible	???:? 08D:CE	2-Way Mir	N/Grp'd ACLX RW		6

3. Verify the ACLX port flag is enabled on each FA port. Enable the ACLX for each port depicted as “NO” under the “ACLX Enabled” column:

```
symcfg -sid 1380 list -fa all
```

Symmetrix ID: 000192601380

S Y M M E T R I X			F I B R E	D I R E C T O R S	
Dir	Port	WWN	ACLX Enabled	Volume Set Addressing	Pnt to Pnt
FA-7E	0	5000097208159118	Yes	No	Yes
FA-7E	1	5000097208159119	Yes	No	Yes
FA-8E	0	500009720815911C	Yes	No	Yes
FA-8E	1	500009720815911D	Yes	No	Yes
FA-9E	0	5000097208159120	Yes	No	Yes
FA-9E	1	5000097208159121	Yes	No	Yes
FA-10E	0	5000097208159124	No	No	Yes
FA-10E	1	5000097208159125	Yes	No	Yes
FA-7F	0	5000097208159158	Yes	No	Yes
FA-7F	1	5000097208159159	Yes	No	Yes
FA-8F	0	500009720815915C	Yes	No	Yes
FA-8F	1	500009720815915D	Yes	No	Yes
FA-9F	0	5000097208159160	Yes	No	Yes
FA-9F	1	5000097208159161	Yes	No	Yes
FA-10F	0	5000097208159164	Yes	No	Yes
FA-10F	1	5000097208159165	Yes	No	Yes

```
symconfigure -sid 1380 -cmd "set port 10E:0 ACLX=enable;" commit -v -nop
```

A Configuration Change operation is in progress. Please wait...

```
Establishing a configuration change session.....Established.
Processing symmetrix 000192601380
{
  set port 10E:0 Access_Logix(ACLX)=Enable;
}

Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 023 of 047 steps.....Executing.
Step 055 of 079 steps.....Executing.
Step 055 of 079 steps.....Executing.
Step 060 of 079 steps.....Executing.
Step 061 of 079 steps.....Executing.
Step 065 of 079 steps.....Executing.
Step 073 of 079 steps.....Executing.
Local: COMMIT.....Done.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

```
symcfg -sid 1380 list -fa all
```

Symmetrix ID: 000192601380

S Y M M E T R I X F I B R E D I R E C T O R S					
Dir	Port	WWN	ACLX Enabled	Volume Set Addressing	Pnt to Pnt
FA-7E	0	5000097208159118	Yes	No	Yes
FA-7E	1	5000097208159119	Yes	No	Yes
FA-8E	0	500009720815911C	Yes	No	Yes
FA-8E	1	500009720815911D	Yes	No	Yes
FA-9E	0	5000097208159120	Yes	No	Yes
FA-9E	1	5000097208159121	Yes	No	Yes
FA-10E	0	5000097208159124	Yes	No	Yes
FA-10E	1	5000097208159125	Yes	No	Yes
FA-7F	0	5000097208159158	Yes	No	Yes
FA-7F	1	5000097208159159	Yes	No	Yes
FA-8F	0	500009720815915C	Yes	No	Yes
FA-8F	1	500009720815915D	Yes	No	Yes
FA-9F	0	5000097208159160	Yes	No	Yes
FA-9F	1	5000097208159161	Yes	No	Yes
FA-10F	0	5000097208159164	Yes	No	Yes
FA-10F	1	5000097208159165	Yes	No	Yes

4. Write disable the ACLX device:

```
symdev -sid 1380 write_disable 008
```

```
Write Disable device '008' on
ALL Directors, ALL Ports (y/[n]) ? y

Write Disable operation successfully completed for the device.
```

5. Unmap the ACLX device from each FA port:

```
symconfigure -sid 1380 -cmd "unmap dev 008 from dir all:all;" commit -v -nop

A Configuration Change operation is in progress. Please wait...

  Establishing a configuration change session.....Established.
  Processing symmetrix 000192601380
  {
    unmap dev 0008 from dir ALL:ALL;
  }

  Performing Access checks.....Allowed.
  Checking Device Reservations.....Allowed.
  Locking devices.....Locked.
  Initiating COMMIT of configuration changes.....Queued.
  COMMIT requesting required resources.....Obtained.
  Step 020 of 020 steps.....Executing.
  Local: COMMIT.....Done.
  Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

38.5. ACLX Database Maintenance

It is recommended to back up the ACLX database periodically to provide protection against corruption or accidental or malicious deletion of masking information. Using the following symaccess command, a backup of the ACLX database to a specific file on the local management server can be created:

```
symaccess -sid 1234 backup -f backup_090909
```

The backup file can be used to restore the masking information. A restore will overwrite the current masking information with the masking views contained in the backup file; therefore, caution must be applied since legitimate configurations could be overwritten/removed. The ACLX database restore process is limited by the following:

- Ports will be returned to port groups only if they have the ACLX attribute set.
- Devices will be returned to storage groups if they are devices that can be masked.
- Views will be restored only if the port and storage groups contain valid entries
- LUN information will be restored if valid for port settings.

The backup file can also be used with the `symaccess` commands to query masking information by using the `-file` option instead of the `-sid`.

38.6. ACLX Database Example:

- ACLX database backup:

```
symaccess -sid 1380 backup -file symaccessDB-111609
```

```
Backup all masking data on symmetrix 000192601380 to backup file symaccessDB-111609
(y/[n])? y
```

```
The masking data on symmetrix 000192601380 was backed up to file symaccessDB-111609.
```

- Displaying database information from a backup file:

```
symaccess -f ACLX-DB-112309 list view -v
```

```
Backup file      : ACLX-DB-112309
```

```
Symmetrix ID in Backup: 000192601380
```

```
Masking View Name : sangs1_gatekeepers
```

```
Initiator Group Name : ig_sangs1
```

```
Port Group Name : sangs1_gatekeepers
```

```
Storage Group Name : sangs1_gatekeepers
```

```
Masking View Name : c14_p560_viocl4
```

```
Initiator Group Name : c14_p560_viocl4
```

```
Port Group Name : Seg1_FabA_7e0_FabB_10F0
```

```
Storage Group Name : c14_p560_viocl4
```

```
Masking View Name : view_Thin_RDF_sangs1
```

```
Initiator Group Name : ig_sangs1
```

```
Port Group Name : sangs1_gatekeepers
```

```
Storage Group Name : Thin_RDF_sangs1
```

```
Masking View Name : view_Fat_RDF_sangs1
```

```
Initiator Group Name : ig_sangs1
```

```
Port Group Name : sangs1_gatekeepers
```

```
Storage Group Name : Fat_RDF_sangs1
```

```
Masking View Name : intux28
```

```
Initiator Group Name : intux28
```

```
Port Group Name : sangs1_gatekeepers
```

```
Storage Group Name : intux28_29
```

```
Masking View Name : intux29
```

```
Initiator Group Name : intux29
```

```
Port Group Name : sangs1_gatekeepers
```

```
Storage Group Name : intux28_29
```

```
Masking View Name : vw_sangs4_DB
```

```
Initiator Group Name : ig_sangs4_DB
```

```
Port Group Name : pg_sangs4_DB
```

```
Storage Group Name : sg_sangs4_DB
```

```
Masking View Name : or_dbtest1
```


```
Initiator Group Name : or_dbtest1
```

```
Port Group Name : or_dbtest1
```

```
Storage Group Name : or_dbtest1
```

39. Appendix D: Host Front-end Port Distribution

The rule-of-17 is used to distribute front-end port connections to avoid a Single Point of Failures (SPOF) on the channel directors and adapters. As depicted in the diagram below, rule-of-17 constitutes zoning a host to an adapter pair to avoid losing total communication with a Symmetrix array. For example, the host would be zoned to an adapter pair such as FE3 and FE14, FE4/FE13, FE7/FE10 or FE8/FE9. Note, the addition of each adapter pairing totals 17, hence the rule-of-17.



* Note: These slots can be configured with either front-end or back-end adapters.

SYM-000078

EMC implemented rule-of-17 to avoid SPOF due to the hardware design configuration within the frame. In the Symm4/5 (8000 series), the even/odd memory and power buses were divided into two domains, where an issue on one bus could shutdown half the frame. In the Symm6 (DMX-1/2/3000), the shelves are divided into two power domains, where a failure could shutdown the half the directors. None of the memory or power bus issues applicable in Symm4, 5, 6 applies to the DMX-3/4 and VMAX. **The DMX-3/4 direct matrix and VMAX engine based design eliminates the rule-17 requirement.** Provisioning host connections to different engines provides sufficient port distribution.

39.1. 2-engine Host Port Distribution

No direct means to distribute host ports is required other than different engines; however, a means to simply to track the port distribution is recommended. Citi certification limits VMAX configuration to 2, 4, 6 and 8 engines; which allows rule-17 to be continually utilized. For example, in a 2-engine VMAX (w/o RDF) rule-17 can be utilized by pairing ports on director 07 with ports on director 10. Front-end #1 in fabric “A” consists of all ports in director 07 (engine 4) and all ports in director 09 (engine 5). Front-end #2 in fabric “B” consists of all ports in director 10 (engine 5) and all ports in director 08 (engine 4).

2-Engine			
Rule-17			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	10H1	Engine 5
	07H0	10H0	
	07G1	10G1	
	07G0	10G0	
	07F1	10F1	
	07F0	10F0	
	07E1	10E1	
	07E0	10E0	
Engine 5	09H1	08H1	Engine 4
	09H0	08H0	
	09G1	08G1	
	09G0	08G0	
	09F1	08F1	
	09F0	08F0	
	09E1	08E1	
	09E0	08E0	

With an even number of VMAX engines dividing the ports between even and odd can also be using in place of rule-of-17. This principal is called the rule-even/odd, where odd port numbers are paired with odd ports from its paired engine. For example, in a 2-engine VMAX rule-even/odd pairs director ports 07 and 09 together and even number director ports 08 and 10 together. Front-end #1 in fabric “A” consists of all ports in director 07 (engine 4) and all ports in director 10 (engine 5). Front-end #2 in fabric “B” consists of all ports in director 08 (engine 4) and all ports in director 09 (engine 5). Mixing rule-of-17 and rule-

even/odd leads to an uneven distribution; therefore, only one method should be used per frame.

2-Engine			
Rule-even/odd			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	09H1	Engine 5
	07H0	09H0	
	07G1	09G1	
	07G0	09G0	
	07F1	09F1	
	07F0	09F0	
	07E1	09E1	
	07E0	09E0	
Engine 5	10H1	08H1	Engine 4
	10H0	08H0	
	10G1	08G1	
	10G0	08G0	
	10F1	08F1	
	10F0	08F0	
	10E1	08E1	
	10E0	08E0	

39.2. 4-engine Host Port Distribution

4-Engine			
Rule-even/odd			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	09H1	Engine 5
	07H0	09H0	
	07G1	09G1	
	07G0	09G0	
	07F1	09F1	
	07F0	09F0	
	07E1	09E1	

4-Engine			
Rule-17			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	10H1	Engine 5
	07H0	10H0	
	07G1	10G1	
	07G0	10G0	
	07F1	10F1	
	07F0	10F0	
	07E1	10E1	

	07E0	09E0	
Engine 5	10H1	08H1	Engine 4
	10H0	08H0	
	10G1	08G1	
	10G0	08G0	
	10F1	08F1	
	10F0	08F0	
	10E1	08E1	
	10E0	08E0	
Engine 3	05H1	11H1	Engine 6
	05H0	11H0	
	05G1	11G1	
	05G0	11G0	
	05F1	11F1	
	05F0	11F0	
	05E1	11E1	
	05E0	11E0	
Engine 6	12H1	06H1	Engine 3
	12H0	06H0	
	12G1	06G1	
	12G0	06G0	
	12F1	06F1	
	12F0	06F0	
	12E1	06E1	
	12E0	06E0	

	07E0	10E0	
Engine 5	09H1	08H1	Engine 4
	09H0	08H0	
	09G1	08G1	
	09G0	08G0	
	09F1	08F1	
	09F0	08F0	
	09E1	08E1	
	09E0	08E0	
Engine 3	05H1	12H1	Engine 6
	05H0	12H0	
	05G1	12G1	
	05G0	12G0	
	05F1	12F1	
	05F0	12F0	
	05E1	12E1	
	05E0	12E0	
Engine 6	11H1	06H1	Engine 3
	11H0	06H0	
	11G1	06G1	
	11G0	06G0	
	11F1	06F1	
	11F0	06F0	
	11E1	06E1	
	11E0	06E0	

39.3. 6-engine Host Port Distribution

6-Engine			
Rule-even/odd			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	09H1	Engine 5
	07H0	09H0	
	07G1	09G1	
	07G0	09G0	
	07F1	09F1	

6-Engine			
Rule-17			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	10H1	Engine 5
	07H0	10H0	
	07G1	10G1	
	07G0	10G0	
	07F1	10F1	

	07F0	09F0	
	07E1	09E1	
	07E0	09E0	
Engine 5	10H1	08H1	Engine 4
	10H0	08H0	
	10G1	08G1	
	10G0	08G0	
	10F1	08F1	
	10F0	08F0	
	10E1	08E1	
	10E0	08E0	
Engine 3	05H1	11H1	Engine 6
	05H0	11H0	
	05G1	11G1	
	05G0	11G0	
	05F1	11F1	
	05F0	11F0	
	05E1	11E1	
	05E0	11E0	
Engine 6	12H1	06H1	Engine 3
	12H0	06H0	
	12G1	06G1	
	12G0	06G0	
	12F1	06F1	
	12F0	06F0	
	12E1	06E1	
	12E0	06E0	
Engine 2	03H1	13H1	Engine 7
	03H0	13H0	
	03G1	13G1	
	03G0	13G0	
	03F1	13F1	
	03F0	13F0	
	03E1	13E1	
	03E0	13E0	
Engine 7	14H1	04H1	Engine 2
	14H0	04H0	
	14G1	04G1	
	14G0	04G0	
	14F1	04F1	
	14F0	04F0	

	07F0	10F0	
	07E1	10E1	
	07E0	10E0	
Engine 5	09H1	08H1	Engine 4
	09H0	08H0	
	09G1	08G1	
	09G0	08G0	
	09F1	08F1	
	09F0	08F0	
	09E1	08E1	
	09E0	08E0	
Engine 3	05H1	12H1	Engine 6
	05H0	12H0	
	05G1	12G1	
	05G0	12G0	
	05F1	12F1	
	05F0	12F0	
	05E1	12E1	
	05E0	12E0	
Engine 6	11H1	06H1	Engine 3
	11H0	06H0	
	11G1	06G1	
	11G0	06G0	
	11F1	06F1	
	11F0	06F0	
	11E1	06E1	
	11E0	06E0	
Engine 2	03H1	14H1	Engine 7
	03H0	14H0	
	03G1	14G1	
	03G0	14G0	
	03F1	14F1	
	03F0	14F0	
	03E1	14E1	
	03E0	14E0	
Engine 7	13H1	04H1	Engine 2
	13H0	04H0	
	13G1	04G1	
	13G0	04G0	
	13F1	04F1	
	13F0	04F0	

	14E1	04E1	
	14E0	04E0	

	13E1	04E1	
	13E0	04E0	

39.4. 8-engine Host Port Distribution

8-Engine			
Rule-even/odd			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	09H1	Engine 5
	07H0	09H0	
	07G1	09G1	
	07G0	09G0	
	07F1	09F1	
	07F0	09F0	
	07E1	09E1	
	07E0	09E0	
Engine 5	10H1	08H1	Engine 4
	10H0	08H0	
	10G1	08G1	
	10G0	08G0	
	10F1	08F1	
	10F0	08F0	
	10E1	08E1	
	10E0	08E0	
Engine 3	05H1	11H1	Engine 6
	05H0	11H0	
	05G1	11G1	
	05G0	11G0	
	05F1	11F1	
	05F0	11F0	
	05E1	11E1	
	05E0	11E0	
Engine 6	12H1	06H1	Engine 3
	12H0	06H0	
	12G1	06G1	
	12G0	06G0	

8-Engine			
Rule-17			
Front-end #1 (fabric A)		Front-end #2 (fabric B)	
Engine 4	07H1	10H1	Engine 5
	07H0	10H0	
	07G1	10G1	
	07G0	10G0	
	07F1	10F1	
	07F0	10F0	
	07E1	10E1	
	07E0	10E0	
Engine 5	09H1	08H1	Engine 4
	09H0	08H0	
	09G1	08G1	
	09G0	08G0	
	09F1	08F1	
	09F0	08F0	
	09E1	08E1	
	09E0	08E0	
Engine 3	05H1	12H1	Engine 6
	05H0	12H0	
	05G1	12G1	
	05G0	12G0	
	05F1	12F1	
	05F0	12F0	
	05E1	12E1	
	05E0	12E0	
Engine 6	11H1	06H1	Engine 3
	11H0	06H0	
	11G1	06G1	
	11G0	06G0	

	12F1	06F1	
	12F0	06F0	
	12E1	06E1	
	12E0	06E0	
Engine 2	03H1	13H1	Engine 7
	03H0	13H0	
	03G1	13G1	
	03G0	13G0	
	03F1	13F1	
	03F0	13F0	
	03E1	13E1	
	03E0	13E0	
Engine 7	14H1	04H1	Engine 2
	14H0	04H0	
	14G1	04G1	
	14G0	04G0	
	14F1	04F1	
	14F0	04F0	
	14E1	04E1	
	14E0	04E0	
Engine 1	01H1	15H1	Engine 8
	01H0	15H0	
	01G1	15G1	
	01G0	15G0	
	01F1	15F1	
	01F0	15F0	
	01E1	15E1	
	01E0	15E0	
Engine 8	16H1	02H1	Engine 1
	16H0	02H0	
	16G1	02G1	
	16G0	02G0	
	16F1	02F1	
	16F0	02F0	
	16E1	02E1	
	16E0	02E0	

	11F1	06F1	
	11F0	06F0	
	11E1	06E1	
	11E0	06E0	
Engine 2	03H1	14H1	Engine 7
	03H0	14H0	
	03G1	14G1	
	03G0	14G0	
	03F1	14F1	
	03F0	14F0	
	03E1	14E1	
	03E0	14E0	
Engine 7	13H1	04H1	Engine 2
	13H0	04H0	
	13G1	04G1	
	13G0	04G0	
	13F1	04F1	
	13F0	04F0	
	13E1	04E1	
	13E0	04E0	
Engine 1	01H1	16H1	Engine 8
	01H0	16H0	
	01G1	16G1	
	01G0	16G0	
	01F1	16F1	
	01F0	16F0	
	01E1	16E1	
	01E0	16E0	
Engine 8	15H1	02H1	Engine 1
	15H0	02H0	
	15G1	02G1	
	15G0	02G0	
	15F1	02F1	
	15F0	02F0	
	15E1	02E1	
	15E0	02E0	

40. Appendix F: Metavolumes – Striped vs. Concatenated

Striped metavolumes has a higher throughput performance potential than concatenated metavolumes. The reason is two fold: 1st, the metavolume

configuration and the backend stripe of virtual provisioning provide two cuts at striping, effectively spreading I/O over more backend resources. 2nd, is that a single FA is only actively working on a single SCSI command to a specific device at a time. With a metavolume it is possible to get one command actively running per meta member. In a purely random workload chances are that each meta member is being accessed, resulting in similar performance between stripe and concatenated. In a highly sequential workload the concatenated metavolume will likely execute one I/O at a time. In a stripe model the sequential workload spans multiple devices allowing multiple commands to be executed in parallel, one per meta member.

41. Appendix G: eLicensing Management

41.1. eLicense Introduction

eLicensing is an EMC license enforcement solution for Enginuity 5875 on Symmetrix VMAX Series arrays. The eLicensing process allows VMAX customers to manage, track, and comply with software license keys for each frame. VMAX frames with Enginuity 5875 or existing systems being upgraded to 5875 will require an entitlement file obtained through the eLicensing process.

The eLicensing process requires access to EMC's PowerLink website and a License Authorization Code (LAC) for each array. The LAC is used to retrieve license files from PowerLink, where the license files entitle users to feature sets on specific arrays. PowerLink License Management allows customers to manage initial product selection, product upgrades, Enterprise License Agreements, and re-hosting license keys. Products without the proper license keys are blocked from activation, including new installations and upgrade to 5875 that have not obtained the proper keys. The following table depicts the license required for each feature:

Enginuity Feature Activated	License Needed
Enginuity	Enginuity
FAST (full volume)	FAST and Optimizer
FAST (thin devices)	FAST VP
SRDF Asynchronous	SRDF/A
SRDF Synchronous	SRDF/S
SRDF/STAR	
Site A VMAX >>>>	SRDF/S, SRDF/A, and STAR
Site B VMAX >>>>	SRDF/S, SRDF/A, and STAR
Site C VMAX >>>>	SRDF/A and STAR
TF/Clone	TF/Clone
TF/Snap	TF/Snap

Dynamic Cache Partitioning	Dynamic Cache Partitioning
Symmetrix Priority Controls	Symmetrix Priority Controls

EMC will adhere to the “do no harm” policy for the foreseeable future, where existing software products/features are automatically activated as part of the upgrade and continue to run regardless of the entitlement discovery process. If necessary, EMC will also provide emergency overrides to avoid disruptions to customer operations.

Caution:

This is the first iteration of eLicensing issues have been uncovered in user security, entitlement searches, field validation, help screens and entitlement file creation. EMC has been made aware of all issues Citi considers detrimental.

Customers upgrading to Enginuity 5875 will receive a LAC letter, as seen below, from EMC License Operations as part of the Legacy Customer Transition (LCT) process. The customer will be allowed to activate and download the entitlement file for a specific VMAX frame, where it can then be install using SYMCLI “symlmf” command on Solutions Enabler 7.2 or newer.

Dear EMC Software User,

Thank you for choosing EMC software. Your EMC Software License Authorization Code (LAC) is R1WJ914MHH1XXXXXXXXXX. You must redeem this LAC for license keys to activate your software. Please protect your LAC like you would any other license key to prevent anyone from improperly activating your software.

Activating Your Software

1. Log into Powerlink at <http://powerlink.emc.com> (New users follow the new member registration steps)
2. Navigate to Support > Software Downloads and Licensing > License Management and click on the product you wish to activate.
3. Enter your assigned LAC R1WJ914MHH1XXXXXXXXXX via the Activation menu.
4. Click the Activate button and follow the steps to obtain your license key(s).

For help generating a license key see the Powerlink License Management page or help links in Powerlink Licensing.

Downloading Your Software

1. Log into Powerlink at <http://powerlink.emc.com> (New users follow the new member registration steps)
2. Navigate to Support > Software Downloads and Licensing>navigate to the appropriate product download.
3. Locate your product and download the appropriate software packages

If you have any questions about your order please contact your EMC Account Representative or your Authorized Reseller.

If you have any questions about EMC Software Licensing please contact our world wide 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

Access to the online activation and download licensing modules for a specific locale or frame requires the user to be defined as an eLicense coordinator.

There are two levels of coordinators: 1) fulfillment coordinator – single user that is "supposed" to receive any software updates, 2) site coordinator - user(s) that need to have access to specific sites to manage licenses. EMC permits any coordinator to add or delete other coordinators without verification or notification; therefore, the defined coordinator list should be verified on a regular basis.

41.2. eLicense Glossary

License Authorization Code (LAC)	unique 11-20 character code provided with each order. At part of the software order a LAC will be provided with directions of entitlement activation and software downloads.
Rehost	Rehost is the process of moving activated entitlements from one device or machine to a different device or machine.
Site	A site represents some meaningful way for a company to manage its entitlements. For example, a site might represent a business division or a geographic location.
Coordinator	A coordinator is a point of contact at a customer site that is responsible for the licensing aspect of the customer's relationship with EMC.
Entitlement	An entitlement is a record of the type, version, and quantity of software a customer has purchased the right to use.
License File	The output of activation used to determine authorization for software usage.
Locking ID	Locking ID is a unique identifier (fingerprint) on a device.
Machine name	User defined unique name to reference where licenses has been activated

41.3. Guidelines

- Engenuity 5875 and Solutions Enabler +7.2 require frame and host based feature sets to be licensed through EMC's PowerLink License Management Application. The License Management Application activates new license keys by creating entitlements files that pertain to a specific frame. License Management Application also manages user access; entitlement queries and feature set updates. (See [Appendix G](#)). Initial and updates to the entitlement file is initiated through the sale order process.
- The "symlmf" command in Solutions Enabler 7.2 and later is required to install the entitlement file into an array or host.

- Moving a license from a DMX to a VMAX is not supported by the software license agreement. Customers do not receive credit for old licenses.
- EMC License Operations is authorized to update sales orders from the old to the new product numbers. However, keys for unlicensed products cannot be created.
- SRDF commands spanning between an Enginuity 5875 and earlier array (i.e. Symmetrix DMX, VMAX) releases requires the SRDF and SYMCLI keys in the 5875 to command both boxes. The SYMCLI key must be on the host issuing the commands even if connected only to the 5875 array.
- VMAX arrays purchased before 2011 may appear to lack feature authorization in Enginuity 5875, which is the result of the sales order containing old product model numbers. This requires update to the sales transaction for new model numbers, which triggers an entitlement creation. If this issue is suspected, contact EMC License Operations at Licensekeys@emc.com or call 800-782-4362 option 4.
- For support with missing or incorrect entitlements after activation contact the EMC License Key Group at LicenseKeys@emc.com or call 1-800-782-3797 (Americas), or + 353 (0) 21 4879862 (EMEA).

41.4. View Activated Entitlement File:

The following procedure determines if an entitlement file exists and provides a view of the entitlements associated with the LAC.

- Procedure:
 - From the PowerLink home page, choose Support > Software Downloads and Licensing > License Management.
 - Select Symmetrix VMAX Series.
 - Select Manage My Licenses.
 - Enter the LAC or search entitlements.
 - Click View entitlements

EMC²

Powerlink Licensing

[Home](#)
[Log Out](#)
[Help](#)
[Feedback](#)

[Home](#)
[Entitlements](#)
[Activation](#)
[Reports](#)
[Help](#)
[Log out](#)

Search Entitlements

[Return](#)

To display a list of entitlements, provide search criteria, then click Search Entitlements. To display additional search criteria, click Advanced Search. [Click here for additional Search Tips](#)

%indicates wildcard search

%LAC:

Product line:

Family: (Cannot display the entire list. To filter the list, type the value.)

License type:

Expire date: to: (ex. Apr 18, 2011)

Status: Available or Active

Entitlement ID:

License Enabler:

License Key:

Pick Line ID:

Purchase Order #:

Sales Order #:

[Advanced Search](#)

© Copyright 2003 - 2010 EMC Corporation. All Rights Reserved | [Privacy Policy](#) | [Help](#) | [Feedback](#) | [FAQs](#)

EMC²

EMC Service Center

EMC Licensing Activation

Product ID:

This product has been registered previously and license keys were already generated. Please download the license files below and follow your product's activation instructions to activate your product.

An email containing the license keys has been sent to guy.shimabuku@citi.com

Name and Description	Key	Date
Product Line : Symmetrix Product Name : Symmetrix Download File	INCREMENT SYMM_VMAX_OPTIMIZER EMCLM 2.0 permanent uncounted \	Apr 06, 2011

Enter the additional email addresses to forward the certificate letter with the activation files.

The Certificate letter has been sent successfully to the email addresses: guy.shimabuku@citi.com

Emails:

Enter frame ID

Enter email address to receive certificate

Options	Entitlement	Parent Company	EndCo Pub ID	Site	Product Line	Title	Product #	Total Qty	Available Qty	Status	LAC	License Enabler ID	Pick Line Order #	Purchase Order #	Sales Order #	Sold-to Company	Dealer Company	Expiration Date	Entitlement Type
View	View Certificate	4944220	Citigroup Incorporated	G197	13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059- 138038548	Symmetrix TF-CLONE BASE NEW LICENSE	457- 100-	1 seats	0	Active	R1WJ9	828803	N/A	84100056	Citigroup Incorporated				Permanent
View	View Certificate	4944221	Citigroup Incorporated	G197	13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059- 138038548	Symmetrix TF-SNAP ADD ON NEW LICENSE	457- 100-	1 seats	0	Active	R1WJ9	828803	N/A	84100056	Citigroup Incorporated				Permanent
View	View Certificate	4944222	Citigroup Incorporated	G197	13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059- 138038548	Symmetrix ADV FAST SUITE BASE NEW LICENSE	457- 100-	1 seats	0	Active	R1WJ9	828803	N/A	84100056	Citigroup Incorporated				Permanent

41.5. View Activation Certificate

License File Retrieval for Enginuity 5875 installs.

- Procedure:
 - From the PowerLink home page, choose Support > Software Downloads and Licensing > License Management.
 - Select Symmetrix VMAX Series.
 - Select Register My Product.
 - Enter the Serial number of the Symmetrix in the Product ID field and insert email address to receive a copy of the listed entitlements.
 - A web page will open depicting the entitlements. The entitlements will automatically be emailed to the address inserted in the previous page. The entitlement file maybe save it to your desktop or media device with the "Serial_Num_Date.LIC" format.

EMC Powerlink

Search Knowledgebase | All Support Resources | SEARCH | Advanced Search | Saved Searches | Open a Service Request | Live Chat

Products | Support | Services | Training | Solutions

Home > Support > Software Downloads and Licensing > License Management

License Management

This area centralizes EMC's license management applications. Whether you need to register your software by site or activate a license key using a License Activation Code (LAC) you can locate the appropriate resource below. Simply select a product family from the alphabetical list to launch the corresponding license management application.

Additional resources to help you manage better:

- Frequently asked questions, including how to locate your Host ID
- Contact the Licensing Support team at Licensing@emc.com

Licensing A-C	Licensing D-Q	Licensing R-Z
Access Logic	Data Domain	RecoveryPoint
AlteraStor	Data Protection Advisor	Application Manager
ApplicationTender	Database Edition by VERITAS	Application Manager/SE
ArchivelTender	DatabaseUttender	Replicator
AS/400 SRDF and TimeFinder Suite	DiskTender	ReportTender
Ascent Capture	Documentum	ResourcePins
Automated Resource Manager	Documentum eRoom	Retrospect
AutoStart	Double Checksum	SAN Advisor
Backup Advisor	DS-M Connect	SAN Copy
Celerra iN3e	EmailMiner	SAN Manager
Centra Software Developers Kit	EmailTender	Smartx
CLARION AX100	EMC Data Manager (EDM)	SMI-E Provider
CLARION AX150	EMC Disk Library	SnapView
CLARION AX6	ERMTender	Solutions Enabler
CLARION Utility Kit	ESN Manager	SourceOne
Com-Pay	Extractor	StorageOne eDiscovery - Kaseya
Connectix Directors and Switches	File Management Appliance	SRDF
Connectix Manager	Foundation Suite by VERITAS	Storage Administrator for Exchange
Consistency Group	Intra	StorageScope
ControlCenter	InfoVault	StorageScope File Level Reporter
CopyCross	Inscape	Symmetrix Performance Analyzer 2.0 and above
CopyPoint	Invista	Symmetrix VMAX Series
	Ironix Data Center Insights	TimeFinder
	Ironix Service Manager	TimePoint for SAP
	MirrorView	VirtualSAN
	Navigator	VirtualSAN
	NDMP TapeLibrary	Volume Logic
	NetWorker	Voxense
	ONCourse	VSAH Space Manager
	Open File Manager	VSS Provider
	Open Migrator/VM	WebTender
	Open Replicator for Symmetrix	WorkflowTender
	OpenScale	z/OS Storage Manager (BSM)
	Performance Manager	
	PowerPath	

Select → [StorageScope File Level Reporter](#)

© Copyright 2011 EMC Corporation. All Rights Reserved. | Terms Of Use | Privacy | Feedback | EMC.com

EMC²
EMC Service Center

EMC Licensing Activation

Enter frame ID

Product ID: HK192601380 [FAQ](#)

This product has been registered previously and license keys were already generated.
Please download the license files below and follow your product's activation instructions to activate your product.

An email containing the license keys has been sent to guy.shimabuku@citigroup.com

Name and Description	Key	Date
Product Line : Symmetrix Product Name : Symmetrix Download File	INCREMENT SYMM_VMAX_OPTIMIZER EMCLM 2.0 permanent uncounted \	Apr 06, 2011

Enter the additional email addresses to forward the certificate letter with the activation files.
The Certificate letter has been sent successfully to the email addresses: guy.shimabuku@citigroup.com

Emails: [»](#)

Enter email address to receive certificate

Sample of entitlement email

From: licensing@emc.com
To:
Cc:
Subject: Your Activation Certificate

Message

Dear EMC Software Product User:
We are pleased to provide you with your EMC License Certificate:
Machine Name 00019260
Locking ID's HK19260

Current activation


Symmetrix		Entitlement ID	Product #	Description	Installed Qty	LAC	Expiration Date
4944	457-100	TF-CLONE BASE NEW LICENSE	1	R1WJ9			
4944	457-100	TF-SNAP ADD ON NEW LICENSE	1	R1WJ9			
4944	457-100	ADV FAST SUITE BASE NEW LICENSE	1	R1WJ9			
4944	456-102	TF-CLONE SATA 251-400TB (TCS0800)	200	R1WJ9			
4944	456-103	TF-SNAP ADD SATA 251-400TB (TCS0800)	200	R1WJ9			

41.6. Entitlement Activation

Activation of available entitlements and creation of the license file.

- Procedure:
 - From the PowerLink home page, choose Support > Software Downloads and Licensing > License Management.
 - Select Symmetrix VMAX Series.
 - Select Manage My Licenses.
 - Enter the LAC and select activate.
 - From the Activate Entitlements page verify address, parent company and user name > next
 - Verify individual entitlement an availability > next

- Confirm legal notice and add necessary email notifications > next
- Save license file



EMC² Powerlink Licensing
where information lives[®]

Home Entitlements Activation Reports Help Log out

» Home

Guy Shimabuku, welcome to the Powerlink Licensing integrated license management system.

Insert LAC

Quick Key(s)
To begin the activation process, provide the LAC or LACs to use, then click Activate. You can also provide the machine name and locking ID to use for the activation.

*LAC(s): R1WJ9..... Add more LACs

Machine name: ?

Locking ID: ?

Activate View Certificate Rehost Regenerate Key

Entitlement Management	Activation Management	Reports
Entitlement Dashboard	Machine Dashboard	Reports
View Entitlements	View Certificates	
Manage Coordinators	Activate Licenses	
	More...	

[Search by Enabler Code](#)

[Download Enabler Code](#)

Select activate

© Copyright 2003 - 2010 EMC Corporation. All Rights Reserved | [Privacy Policy](#) | [Help](#) | [Feedback](#) | [FAQs](#)

EMC² Powerlink Licensing

Home Entitlements Activation Reports Help Log out

Activate Entitlements

Quick Keys

To register a different site, select from the dropdown list or search for sites. Select the Next button to keep the entitlements assigned to their current site.

Move to site:

The company and user listed below are the registered company and user for this transaction.

Parent Company
Citigroup Incorporated - G197
399 PARK AVE
FL 3
NEW YORK, 10022
United States

Search Site

To display a list of sites, specify search criteria, then click Search Sites. [Click here for additional Search Tips](#)

% indicates wildcard search

%Site name:

%Site ID:

%City: warren

Country: United States

Search Sites Cancel

Showing: 1-9 of 9

Select	Site Name	Site ID	City	Country
Select	13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059	136036546	WARREN	United States
Select	4109954 CITIGROUP PTY LIMITED 283 KING GEORGE RD WARREN NJ US 07059	132152979	WARREN	United States
Select	4260241 CITIGROUP PTY LIMITED 283 KING GEORGE RD WARREN NJ US 07059	132220802	WARREN	United States
Select	4263552 CITIBANK DDD, BUILDING B /B4-C002 WARREN NJ US 07059	132224099	WARREN	United States
Select	5044605 CITIGROUP PTY LIMITED 283 KING GEORGE RD WARREN NJ US 07059	132977418	WARREN	United States
Select	5108640 CITIGROUP CT/CE INC. 283 KING GEORGE RD WARREN NJ US 07059	132991587	WARREN	United States
Select	5200174 CITIGROUP 283 KING GEORGE RD E2 C 073 WARREN NJ US 07059	133035487	WARREN	United States
Select	7351275 CITIGROUP 283 KING GEORGE RD WARREN NJ US 07059	133420829	WARREN	United States
Select	7610046 *EXPIRED*DSE NA DCTM -CITIGROUP 283 KING GEORGE RD WARREN NJ US 07059	133530135	WARREN	United States

Showing: 1-9 of 9

© Copyright 2003 - 2010 EMC Corporation

EMC² Powerlink Licensing

Home Entitlements Activation Reports Help Log out

Activate Entitlements

Quick Keys **Register** Activate Confirm Complete

Cancel Next >

To register a different site, select from the dropdown list or search for sites. Select the Next button to keep the entitlements assigned to their current site.

Move to site: 13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059 - 1:

The company and user listed below are the registered company and user for this transaction.

Parent Company	User Name
Citigroup Incorporated - G197329691	Prem Anand S.
399 PARK AVE	140 58TH ST
FL 3	BROOKLYN, 11220
NEW YORK, 10022	United States
United States	softenger.premnand.s@citigroup.com
	[Change]

Cancel Next >

© Copyright 2003 - 2010 EMC Corporation. All Rights Reserved | Privacy Policy | Help | Feedback | FAQs

Activate Entitlements - Activate - [ie]

https://powerlink.emc.com/nsepr/webapps/powerlinklicensing/poeticWeb/activation/activationDetails.htm?Q=44B2C3C8778483B24B182FCBC0955F5D09B34

Activate Entitlements - Activate

Home Entitlements Activation Reports Help Log out

Activate Entitlements

Quick Keys Register **Activate** Confirm Complete

Cancel < Back Next >

Please provide required information. Enter or validate Quantity then click Next.

Machine Name Machine Information

000192601380 * Symm Serial Number: HK000192601380 Locking ID Help

Symmetrix

Qty	Avail	Installed	Entitlement	
1	1 seats	0 seats	ADV FAST SUITE BASE NEW LICENSE Product # 457-100-485	Remove
128	128 seats	0 seats	ADV FAST SUITE NON-SATA 251-400TB (TCF0800) Product # 456-102-687	Remove
200	200 seats	0 seats	ADV FAST SUITE SATA 251-400TB (TCS0800) Product # 456-102-698	Remove
1	1 seats	0 seats	SRDF-A BASE NEW LICENSE Product # 457-100-519	Remove
128	128 seats	0 seats	SRDF-A NON-SATA 251-400TB (TCF0800) Product # 456-102-872	Remove
200	200 seats	0 seats	SRDF-A SATA 251-400TB (TCS0800) Product # 456-102-883	Remove
128	128 seats	0 seats	SRDF-S ADD NON-SATA 251-400TB (TCF0800) Product # 456-102-837	Remove
1	1 seats	0 seats	SRDF-S ADD ON NEW LICENSE Product # 457-100-513	Remove
200	200 seats	0 seats	SRDF-S ADD SATA 251-400TB (TCS0800) Product # 456-102-848	Remove
1	1 seats	0 seats	TF-CLONE BASE NEW LICENSE Product # 457-100-541	Remove
128	128 seats	0 seats	TF-CLONE NON-SATA 251-400TB (TCF0800) Product # 456-102-966	Remove
200	200 seats	0 seats	TF-CLONE SATA 251-400TB (TCS0800) Product # 456-102-977	Remove
128	128 seats	0 seats	TF-SNAP ADD NON-SATA 251-400TB (TCF0800) Product # 456-103-050	Remove
1	1 seats	0 seats	TF-SNAP ADD ON NEW LICENSE Product # 457-100-563	Remove
200	200 seats	0 seats	TF-SNAP ADD SATA 251-400TB (TCS0800) Product # 456-103-061	Remove

Discussions Discussions not available on https://powerlink.emc.com/

Internet | Protected Mode: Off 100%

Activate Entitlements - Confirm - [ie]

https://powerlink.emc.com/nsepn/webapps/powerlinklicensing/poeticWeb/activation/activationConfirmation.htm?Q=4482C3C8778483B248182FCBC0955F5

EMC² Powerlink Licensing

Home Entitlements Activation Reports Help Log out

Activate Entitlements

Quick Keys Register Activate **Confirm** Complete

Cancel < Back Finish

Legal Notice

You must agree to the following:

As your company's registered user, you warrant that:UPON CHECKING THIS CHECK BOX YOU ARE PROVIDING AFFIDAVIT THAT YOU ARE IN COMPLIANCE WITH THE FOLLOWING TERMS. You are in full compliance with your License Agreement and all the restrictions therein with the change you are about to commit. My company understands this restriction and agrees to safeguard its provisions. Nevertheless, my company wishes to perform this transfer or other change.

☒ To accept, check the check box

Confirmation Information

To send additional certificates to specified email addresses, enter the addresses in the field below. Use commas to separate multiple addresses.

Email to: (Multiple separated by commas)

Language: U.S. English

Comments:

You can add notes concerning this transaction.

Notes

Here is a summary of this activation transaction.

EMC² Powerlink Licensing

Home Entitlements Activation Reports Help Log out

Activate Entitlements

Return to Quick Keys Entitlement Dashboard Machine Dashboard

Quick Keys Register Activate Confirm **Complete**

The activation process is completed. To continue, select one of the links below.

The certificate is being sent automatically to the registered user.

Email certificate View certificate

Machine Name	Activation Information
000192601380 View summary	Activation date: Apr 6, 2011 Certificates sent to: softenger.premanand.s@citigroup.com, guy.shimabuku@cit.com

Save to File

Annotation: Symmetrix

Key: INCREMENT SYMM_VMAX_OPTIMIZER EMCLM 2.0 permanent uncounted \VENDOR_STRING=LICENSETYPE=IND.SYMMID=HK192601380;COUNT=128;TYPE=RAW-NON-SATA-TB \HOSTID=ANY dist_info="ACTIVATED TO 13043290 CITI 283 KING \GEORGE RD WARREN NJ US 07059" ISSUER=EMC IS (TRUNCATED) [View key](#)

Copyright 2003 - 2010 EMC Corporation. All Rights Reserved | Privacy Policy | Help | Feedback | FAQs

Download entitlement file

41.7. Accessing Enterprise License Agreements

An Enterprise Licensing Agreement (ELA) applies to multiple arrays and is tied to a Site ID or Customer ID per the ELA contract. Each array under an ELA still has

an array-specific license file. With Engenuity 5875 all licenses are permanent.

- Procedure:
 - From the PowerLink home page, choose Support > Software Downloads and Licensing > License Management.
 - Choose Symmetrix VMAX Series.
 - Choose Manage My Licenses.
 - Enter the LAC.
 - View entitlements

The screenshot displays the Symmetrix VMAX License Management interface. At the top, there are navigation links: Return to previous page, Print, Email, and Options. Below these are tabs for Overview, Ownership, Keys, and Transactions. The main content area is divided into two sections: Machine Summary and Activation Details.

Machine Summary:

- Machine name: 000192601380
- This machine includes the following locking IDs.
- SYMM ID: 000192601380
- MAC Address (SYMM): 08-00-20-c5-1b-41
- Symm Serial Number: HK192601380

Activation Details:

The Activation Details section shows two license families: ADV FAST Suite and SRDF-S AO. Each family has a table of entitlements.

ADV FAST Suite License Details:

Entitlement	Qty	Start Date	Expiration Date	Can Update Until	Move Count	Entitlement ID
ADV FAST SUITE NON-SATA 251-400TB (TCF0800)	128 seats	Apr 6, 2011			5	4944230
Product #.456-102-687						
ADV FAST SUITE SATA 251-400TB (TCS0800)	200 seats	Apr 6, 2011			5	4944229
Product #.456-102-698						
ADV FAST SUITE BASE NEW LICENSE	1 seats	Apr 6, 2011			5	4944222
Product #.457-100-485						

SRDF-S AO License Details:

Entitlement	Qty	Start Date	Expiration Date	Can Update Until	Move Count	Entitlement ID
SRDF-S ADD NON-SATA 251-400TB (TCF0800)	128 seats	Apr 6, 2011			5	4944234
Product #.456-102-837						
SRDF-S ADD SATA 251-400TB (TCS0800)	200 seats	Apr 6, 2011			5	4944226
Product #.456-102-848						
SRDF-S ADD ON NEW LICENSE	1 seats	Apr 6, 2011			5	4944228
Product #.457-100-513						

41.8. Coordinator (user) Management

eLicensing provide two levels of user identifiers; 1) fulfillment coordinator – only one per frame, is the user that "supposed" to receive any software updates, 2) site coordinator - users that need to have access to specific sites to manage licenses. Coordinators may only be selected from a database managed by EMC. Additions to the coordinator database is accomplished by an email to the license management team, which verifies the new users through the Powerlink profile. Removal of a site coordinator or replacement of a fulfillment coordinator can be accomplished by any coordinator without verification or notification.

- Procedure:
 - From the PowerLink home page, choose Support > Software Downloads and Licensing > License Management.
 - Select Symmetrix VMAX Series.
 - Select Manage My Licenses.

- Enter the LAC.
- Click View Certificate.
- From the Certificate page click the Ownership tab.
- From the ownership page select view coordinators.
- Add, replace or remove coordinators as necessary.

Click ownership

Return to previous page | Print | Email | Options

Overview | **Ownership** | Keys | Transactions

Machine Summary

Machine name: 000192601380

This machine includes the following locking IDs.

SYMM ID : 000192601380
 MAC Address (SYMM) : 08-00-20-c5-1b-41
 Symm Serial Number : HK192601380

Activation Details

Symmetrix	Family	License Version	Total Qty			
Symmetrix Permanent activation type Lowest move count: 5	ADV FAST Suite	All	329 seats	Details		
[Close]						
Details						
Entitlement	Qty	Start Date	Expiration Date	Can Update Until	Move Count	Entitlement ID
ADV FAST SUITE NON-SATA 251-400TB (TCF0800) Product # 456-102-687	128 seats	Apr 6, 2011			5	4944230
ADV FAST SUITE SATA 251-400TB (TCS0800) Product # 456-102-698	200 seats	Apr 6, 2011			5	4944229
ADV FAST SUITE BASE NEW LICENSE Product # 457-100-485	1 seats	Apr 6, 2011			5	4944222
SRDF-S AO						
All						
329 seats						
[Close]						
Details						
Entitlement	Qty	Start Date	Expiration Date	Can Update Until	Move Count	Entitlement ID
SRDF-S ADD NON-SATA 251-400TB (TCF0800) Product # 456-102-837	128 seats	Apr 6, 2011			5	4944234
SRDF-S ADD SATA 251-400TB (TCS0800) Product # 456-102-848	200 seats	Apr 6, 2011			5	4944226
SRDF-S ADD ON NEW LICENSE Product # 457-100-513	1 seats	Apr 6, 2011			5	4944228

Overview Ownership Keys Transactions

The following is the registered company for this certificate.

Parent Company
 Citigroup Incorporated - G197329691
 399 PARK AVE
 FL 3
 NEW YORK, 10022
 United States

The following list shows how the certificates are allocated among the company sites.

Site	Product Line	Family	Qty
13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059 - 136036546	Symmetrix	ADV FAST Suite	329 seats
283 KING GEORGE RD	Symmetrix	SRDF-S AO	329 seats
WARREN, 07059	Symmetrix	SRDF-A	329 seats
United States	Symmetrix	TF-Snap AO	329 seats
	Symmetrix	TF-Clone	329 seats

[\[View Coordinators\]](#)

Below are the registered users associated with this machine.

[\[Remove or replace registered users\]](#)

Prem Anand S.
 140 58TH ST
 BROOKLYN, 11220
 United States
 Email: softenger.premanand.s@citigroup.com

Select View Coordinators

Use this page to search for the sites and coordinators that belong to a given company. From the results, you can create and edit sites, and change the coordinators assigned to each site. When searching by coordinator name, provide either the first name or last name of a user coordinator. Do not provide the full name of a user. [Click here for additional Search Tips](#)

% indicates wildcard search

%Site name: 13043290 CITI 283 KING GEORGE RD W
 %Site ID: 136036546

Coordinator type:
 Site status: Active

You can provide either the first or last name of a user coordinator, or the company name of a company coordinator. You cannot search for both a user coordinator name and a company coordinator name.

%User coordinator:
 %Company coordinator:
[Search Coordinators](#)

Here are the sites and coordinators that matched your search criteria.

[Replace checked coordinators](#) [Remove checked coordinators](#) [Get total count](#)

Site	Options	Coordinator Type	Coordinator
13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059 - 136036546	☐ Options	Fulfillment Coordinator	Guy Shimabuku
283 KING GEORGE RD	☐ Options	Site Coordinator	Robert Hartigan (Citigroup Incorporated -- Citigroup Incorporated; 399 PARK AVE; NEW YORK; United States)
WARREN, 07059	☐ Options	Site Coordinator	Guy Shimabuku (Citigroup Incorporated -- Citigroup Incorporated; 399 PARK AVE; NEW YORK; United States)
United States	☐ Options	Site Coordinator	Adrian Pengelly (Citigroup Incorporated -- Citigroup Incorporated; 399 PARK AVE; NEW YORK; United States)

[Add Coordinator](#)

[Replace checked coordinators](#) [Remove checked coordinators](#) [Get total count](#)

Select items to manage coordinators

Use this page to search for the sites and coordinators that belong to a given company. From the results, you can create and edit sites, and change the coordinators assigned to each site. When searching by coordinator name, provide either the first name or last name of a user coordinator. Do not provide the full name of a user. [Click here for additional Search Tips](#)

% indicates wildcard search

%Site name: 13043290 CITI 283 KING GEORGE RD W

%Site ID: 136036546

Coordinator type:

Site status:

You can provide either the first coordinator name.

%User coordinator:

%Company coordinator:

Here are the sites and coordinators:

[Replace checked coordinators](#)

Site
13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059 - 136036546
283 KING GEORGE RD WARREN, 07059
United States

[Add Coordinator](#)

[Replace checked coordinators](#) [Remove checked coordinators](#) [Get total count](#)

Copyright 2003 - 2010 EMC Corporation. All Rights Reserved | [Privacy Policy](#) | [Help](#) | [Feedback](#) | [FAQs](#)

Add Coordinator

Select the type(s) of coordinator to add.

The following types of coordinators only allow one coordinator per site (or per site family): Fulfillment Coordinator. If you add these types of coordinators, the new coordinator may replace the existing coordinator assigned to the site.

Coordinator is a:

Coordinator type: ☐ Site Coordinator ☒ Fulfillment Coordinator

Below is the site that the selected coordinator will be added to.

Site name: 13043290 CITI 283 KING GEORGE RD WARREN NJ US 07059 - 136036546

Select the new coordinator. Either type the coordinator information, or click the find icon to search for and select the new coordinator.

*New coordinator email:

[\[Validate Entered User\]](#)

41.9. View Licenses through Solutions Enabler

Users can perform frame and host license queries using the SYMCLI. The query output depicts how the feature has been licensed plus the licensed and utilized capacity based on drive type. Features that have been activated using an emergency override at the service processor will be reported as manual override and displayed as "MAN." Features that have been activated as part of an Enginuity upgrade is reported as "IN USE." Features that have been properly licensed through an entitlement is displayed as "ENT."

symlmf query -type emclm -sid xxxx - command displays the state of all licenses for a Symmetrix that were activated by the license file or were deemed "in-use" at the time of upgrade from Enginuity 5874. In addition, it displays the capacity usage by license and drive type.

```
# symlmf -sid 1234 query -type emclm
```

```
Symmetrix ID : 000192601234
Activation ID: 1747749
Issue Date   : 13/06/2011
```

Feature Name	Act	Type	Capacity	
			Licensed	Usage
SYMM_VMAX_DCP	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_ENGINUITY	USE	R-TB-Non-SATA	0	88.4
		R-TB-SATA	0	128.0

SYMM_VMAX_FAST	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_FAST_VP	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_OPTIMIZER	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_SPC	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_SRDF	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_SRDF_A	ENT	R-TB-Non-SATA	128	-
		R-TB-SATA	200	-
SYMM_VMAX_SRDF_S	ENT	R-TB-Non-SATA	128	-
		R-TB-SATA	200	-
SYMM_VMAX_TF_CLONE	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0
SYMM_VMAX_TF_SNAP	ENT	R-TB-Non-SATA	128	88.4
		R-TB-SATA	200	128.0

Legend:

Act(ivation Type):

ENT = Entitlement

MAN = Manual Override

USE = In Use

symlmf show -type emclm -sid xxxx - command displays the license file that was last used for installation of licenses.

```
# symlmf -sid 1380 show -type emclm
INCREMENT SYMM_VMAX_OPTIMIZER EMCLM 2.0 permanent uncounted \
  VENDOR_STRING=LICENSETYPE=IND;SYMMID=HK192601380;COUNT=128;TYPE=RAW-
NON-
SATA-TB \
  HOSTID=ANY dist_info="ACTIVATED TO 13043290 CITI 283 KING \
  GEORGE RD WARREN NJ US 07059" ISSUER=EMC ISSUED=06-Apr-2011 \
  NOTICE="ACTIVATED TO SiteID: 136036546" SN=1747749 TS_OK \
  SIGN="000C A5EB 4810 846F 21E3 D1F3 200E 4C00 2C38 C4C4 584F \
  CDC7 C6F3 A3E8 A19C"
INCREMENT SYMM_VMAX_DCP EMCLM 2.0 permanent uncounted \
  VENDOR_STRING=LICENSETYPE=IND;SYMMID=HK192601380;COUNT=128;TYPE=RAW-
NON-
SATA-TB \
  HOSTID=ANY dist_info="ACTIVATED TO 13043290 CITI 283 KING \
  GEORGE RD WARREN NJ US 07059" ISSUER=EMC ISSUED=06-Apr-2011 \
  NOTICE="ACTIVATED TO SiteID: 136036546" SN=1747749 TS_OK \
  SIGN="000A 6C9E 0EC2 6AF6 8FE0 7626 37C9 BC00 D327 BFB7 4679 \
  F86D 4A75 9EBF E8AE"
INCREMENT SYMM_VMAX_SPC EMCLM 2.0 permanent uncounted \
  VENDOR_STRING=LICENSETYPE=IND;SYMMID=HK192601380;COUNT=128;TYPE=RAW-
NON-
SATA-TB \
  HOSTID=ANY dist_info="ACTIVATED TO 13043290 CITI 283 KING \
  GEORGE RD WARREN NJ US 07059" ISSUER=EMC ISSUED=06-Apr-2011 \
```

```

NOTICE="ACTIVATED TO SiteID: 136036546" SN=1747749 TS_OK \
SIGN="00FB CBF7 9937 8887 CF27 F8A7 DC20 AF00 B51A D9D7 FEE1 \
FC1E 0A8A 95A5 353A"
INCREMENT SYMM_VMAX_FAST EMCLM 2.0 permanent uncoun ted \
VENDOR_STRING=LICENSETYPE=IND;SYMMID=HK192601380;COUNT=128;TYPE=RAW-
NON-
SATA-TB \
HOSTID=ANY dist_info="ACTIVATED TO 13043290 CITI 283 KING \
GEORGE RD WARREN NJ US 07059" ISSUER=EMC ISSUED=06-Apr-2011 \
NOTICE="ACTIVATED TO SiteID: 136036546" SN=1747749 TS_OK \
SIGN="0062 B986 9B9A 2E27 9827 A72D 2D4F EA00 1992 E469 2AB3 \
1B04 FFF1 7828 6796"

```

41.10. Entitlement File Installation

The entitlement file obtained through the eLicense process is installed on the array through Solutions Enabler. In certain instances the entitlement file will contain both frame (i.e. SRDF, FAST) and host based licenses (i.e. SPA).

Procedure:

- Copy the entitlement file onto Solutions Enabler management server.
- To load the host based licenses run command “*symlmf add -type emclm (EMC License Management) -file <entitlement_file>*” Note, the frame based licenses will not be processed.
- Verify host features with “*symlmf query*” command.
- To load the frame based licenses run command “*symlmf -sid <Frame_ID> add -type emclm (EMC License Management) -file <entitlement_file>*”
- Verify the frame features and usage values with “*symlmf -sid <Frame_ID> query*” command.

Host Based License Installation

```
# symlmf add -type emclm -file ./HK192604261_14-Jun-2011.lic
```

```
Total licenses Processed      :    1
Total licenses Not Processed:   12
```

```
# symlmf list -type emclm
```

Feature Name	SymmID	Host ID	Days Until Expr	Capacity Type	Units
SPA_BASE	000192601380	GDSEEMCSPA02	-	-	-
SPA_BASE	000192604261	-	-	-	-
SPA_DIAG	000192601380	GDSEEMCSPA02	-	-	-

Legend:

Capacity:

R-TB = Raw capacity in TB

Reg-TB = Configured capacity in TB

- = No capacity limit

Frame Based License Installation

```
# symmlmf -sid 4261 add -type emclm -file HK192604261_14-Jun-2011-latest.lic
```

```
Total licenses Processed      :    12
```

```
Total licenses Not Processed:     1
```

```
# symmlmf query -type emclm -sid 4261
```

```
Symmetrix ID : 000192604261
```

```
Activation ID: 1781086
```

```
Issue Date   : 06/15/2011
```

Feature Name	Act	Type	Capacity	
			Licensed	Usage
SYMM_VMAX_DCP	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_ENGINUITY	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_FAST	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_FAST_VP	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_OPTIMIZER	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_PROSPHERE	USE	R-TB-Non-SATA	0	83.2
		R-TB-SATA	0	128.0
SYMM_VMAX_SMC	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_SPC	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_SRDF	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_SRDF_A	ENT	R-TB-Non-SATA	90	-
		R-TB-SATA	128	-
SYMM_VMAX_SRDF_S	ENT	R-TB-Non-SATA	90	-
		R-TB-SATA	128	-
SYMM_VMAX_TF_CLONE	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0
SYMM_VMAX_TF_SNAP	ENT	R-TB-Non-SATA	90	83.2
		R-TB-SATA	128	128.0

Legend:

Act(ivation Type):

ENT = Entitlement

MAN = Manual Override

USE = In Use

42. Appendix H: Capacity to Segment Ratio

The certification of high capacity Fibre Channel (FC) and SATA disk drives plus storage virtualization tremendously increases the potential frame capacity of the Symmetrix VMAX array. Previous thick device deployments were based on Fibre Channel 146/300GB disk drives, whereby allotting a maximum of 10TB of capacity per segment (FA pair) as sufficient to provision the entire backend capacity of the frame. The advent of the 600GB FC, 2TB SATA II and thin provisioning oversubscription essentially quadruples or more the capacity of the frame within the same physical footprint, drive count and channel ports. This renders the 10TB per segment policy insufficient to provision the entire logical and virtual capacity of the frame. The current Citi imposed limit of 10TB/segment has been increased to 50TB/segment, allowing the VMAX array to operate with a high percentage of capacity from SATA drives and support 200% thin pool oversubscription.

The 50TB/segment is an extensive increase from the current standard; therefore, a step approach of increasing the capacity per segment ratio in increments of 5-10TB to the 50TB limit is feasible. At each step increase verify that front-end port is not being overloaded. If an individual port displays a fairly consistent 8000 IOps and/or 100MB/s for more than 60 minutes each in duration, then action should be considered. This action should start with simple symstat data collection followed by a STP report showing the load across all FA ports on the frame. If analysis shows that heavy load is consistent on certain ports, these ports should be distributed to more or less busy FA ports. The increasing or adding segments for device(s) will also distribute the load; however, the capacity to segment ratio will also increase.

Storage tiering allows access workloads to be aligned with the best suited drive technology. This alignment groups drive technologies into individual tiers consisting of Enterprise Flash Drives (EFD), Fibre Channel (FC) and SATA drives (see table below). The EFD and FC drives represent the high performance tiers and the SATA drive represents the high density less costly tier. The amount of SATA drives will have the highest impact in increasing the capacity to segment ratio of the frame.

Symmetrix VMAX Supported Disk Drive Types				
Disk Name	Formatted Capacity (GB)	Array Size (Cylinders)	Speed RPM	EMC Nomenclature
Fibre Channel 300GB	288.19	305175	15000	T300155
Fibre Channel 450GB	432.29	457763	15000	H450156, HUC4515
Fibre Channel 600GB	576.39	586337	15000	EGC6015
SATA II 2TB	1882.72	1915210	7200	N02THJK
STEC EFD 400GB (note1)	393.84	400640	N/A	STC0400
Samsung EFD 200GB (note 1)	196.7	200320	N/A	N/A

Note 1: For use on three Tier FAST VP frames only

The aligning of drive technology with access work loads is essential in our effort to reducing data storage. The table below shows a three tier 4-engine VMAX array configured for tiering. The thin pool size column shows a total capacity of 777,291GB, which is equal to 24,290GB/segment when divided equally among the available FA pairs; assume no RDF ports. When factoring in the thin pool 200% oversubscription the capacity per segment increases to 48,580GB. If the current imposed limit of 10TB/segment were implemented on this 4-engine configuration, the frame would support 320,000GB or 20.6% of the potential thin device capacity. The capacity per segment ratio will be constantly increasing, where larger SATA, SAS (Serial Attached SCSI) and EFD drives coupled with 8 and 16 Gbps channel adapters will increase the capacity per segment beyond 50TB limit in the near future.

FAST VP Thin Pools 1 EFD (minimum configuration)								
Tier	Disk Type	Disk Size (cyl)	Disk Count per Loop	Loop Count %	Protection Type	4-Engine		Tier Ratios
						Disk Count	Thin Pool Size (GB)	
1	STEC	400640	3	10%	7R5	96	30811	3.96%
2	EGC6015	586337	17	57%	14R6	544	255520	32.87%
3	N02THJK	1915210	10	33%	14R6	320	490960	63.16%
Totals			30			960	777291	
Capacity per segment (FA Pair) in GB							24290	

42.1. Fiber Adapter Port Monitoring:

There are several methods to monitor the load on Fiber Adapter ports: 1) STP metrics collection, 2) symstat command through Solutions Enabler. STP provides the greatest detail and the number of metrics collected; however, the metrics may take several hours or more to produce. Symstat is simple command run from the Solutions Enabler server and provides a concise and near real-time output. Recommend using that symstat be used in conjunction with STP in determining heavily loaded segments.

The following symstat and STP output shows how to determine if segments are heavily loaded. The symstat output depicts the throughput (IO/sec) and transaction rate (KB/sec) for FA ports in 60 second intervals. This output shows FA ports 9e1 and 8f1 with a heavy throughput load, as verified through STP front-end port performance graphs showing a constant +85% busy. The STP device metrics shows which devices are presenting the heavy load to the frame Based on the data gathered, the 16 device(s) causing the high load should be distributed to segments with less load, like ports 7e1 and 10f1 that shows little to no load, or to more segments dividing the load between four to six FA ports.


42.1.1. SYMSTAT Command


```
# symstat -sid 1380 -i 10 -type port -dir all
```

42.1.2. SYMSTAT output:


TIME	DIRECTOR	PORT	IO/sec	Kbytes/sec
14:46:18	FA-7E	1	15	0
14:46:18	FA-9E	1	9525	76204
14:46:18	FA-10E	1	36	0
14:46:18	FA-7F	1	6	18
14:46:18	FA-8F	1	9757	77972
14:46:18	FA-9F	1	838	98522
14:46:18	FA-10F	1	2	0
14:47:19	FA-7E	1	12	0
14:47:19	FA-9E	1	8866	70928
14:47:19	FA-10E	1	26	0
14:47:19	FA-7F	1	26	110
14:47:19	FA-8F	1	9007	71964
14:47:19	FA-9F	1	743	88320
14:48:19	FA-7E	1	12	0
14:48:19	FA-9E	1	9562	76500
14:48:19	FA-10E	1	33	0
14:48:19	FA-7F	1	4	12
14:48:19	FA-8F	1	9639	77014
14:48:19	FA-9F	1	820	98512
14:49:19	FA-7E	1	12	0
14:49:19	FA-9E	1	9225	73804
14:49:19	FA-10E	1	13	0
14:49:19	FA-7F	1	4	10
14:49:19	FA-8F	1	9386	75002
14:49:19	FA-9F	1	812	96622
14:50:19	FA-7E	1	9	0
14:50:19	FA-9E	1	8930	71442
14:50:19	FA-10E	1	30	0
14:50:19	FA-7F	1	3	8
14:50:19	FA-8F	1	9006	71936
14:50:19	FA-9F	1	783	92704

42.1.3. STP Front-end Port Metrics:






42.1.4. STP Device metrics



43. Appendix I: BCV Cache Utilization

When establishing a cloning session to a BCV device, protection slots (e.g. cache slots) are dynamically allocated from global cache. These protection slots are used to track what needs to be and what has already been copied during the cloning process. The number of cache slots allocated is dependant on the size of the source device, where two cache slots is consumed for every 2144 cylinders of the source device. As depicted in the graphs below, the amount of useable cached required to support BCV's can be substantial based the amount of global cache.



As a mode of self-protection, the frame limits the protection slot allocation to 7% of global cache with Enginuity +5874.207 (early code levels limited to 5%), which is equal to 8.4GB in a frame with 120GB of cache. If the 7% protection slot threshold is surpassed all further cloning sessions will fail with an error message containing “Partial failure: BCV lib error.”

The table below depicts an example frame configuration with 1200 BCVs of varying sizes and the calculated cache required to support the BCVs. If this were a typical 2-engine VMAX with 64GB of global cache, the 8.09GB required to support cloning would be surpassed. In this instance, a 2-engine frame with 128GB of cache would support the protection slot allocation to accommodate 1200 BCVs. It is imperative that enough cache be initially installed to minimize the insufficient cache cloning errors and cache upgrade procedures.

Cache Utilization to BCV Ratio - insert the number of BCV's by size

Source BCV Size	Number of BVC's	Size of BCV in cylinders	Total BCV Capacity	Number of Cache Slots Needed	Amount of Cache in GB Needed to Process BCVs
less than 67GB	500	73656	36828000	34354	2.10
68-101GB	300	110484	33145200	30919	1.89
102-134GB	200	147312	29462400	27484	1.68
135-167GB	100	184140	18414000	17177	1.05
168-202GB	50	220968	11048400	10306	0.63
203-239GB	50	262668	13133400	12251	0.75
Totals	1200		142031400	132492	8.09
Note: This table depicts the estimated amount of cache needed to support different clone sessions per frame					

There are two primary methods to alleviate this insufficient cache issue: 1) reduce the number of devices in the cloning operation, 2) add more cache to global memory. Reducing or terminating clone jobs to free up protection slots are usually not feasible, and the deallocation process following cloning termination is a low priority background task that depending on frame load may take a while to complete. The upgrading of global cache capacity requires the serial shutdown of each director; each engine has two directors. Caching is provided through DIMM's (Dual Inline Memory Modules) on each director, where all directors in the frame shall have the same amount of cache. Each director usually takes about 45-60 minutes to upgrade, at which time the director protection pairing will be in degraded mode. For a 2-engine frame with four directors, EMC recommends a five hour window to complete and run the necessary checks.

44. Appendix J: Symmetrix Performance Analyzer

Symmetrix Performance Analyzer (SPA) is a browser based user interface for historical trending and analysis of Symmetrix array performance data. SPA provides a historical view of all Symmetrix Key Performance Indicators (KPIs), a Diagnostic view for root cause analysis, and a Real Time view for examining high frequency data points as they happen.

SPA provides easy-to-use graphical data representations as well as export and print functions to save the data. SPA collects and stores performance data for you to analyze and report on workload and resource usage trends for service-level management, diagnostics, and capacity planning.

You can use SPA to:


- View graphs detailing system performance

- Monitor performance and resource utilization over time
- Create your own dashboard for critical, periodic metrics
- Monitor a specific application's performance
- Diagnose device-level issues
- Perform root cause analysis
- Drill-down through data to investigate issues


For further details, see the CATE SPA Solutions Guide.

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/CSBE-11028-SG.pdf>

Example Dashboard view



Example snapshot view



Example Trend view

