



CATE Storage and Backup Engineering
Citigroup Architecture & Technology Engineering

EMC VMAX3

CSBE-15013-SG

Solutions Guide

Version 0.9

©Citigroup Inc. 2015, All Rights Reserved.

This document contains material internal to Citigroup. The contents are protected by specific employee and contractor agreements regarding confidentiality. Reproduction for distribution outside of Citigroup is prohibited.

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	03-Mar-2015	Quang Le	-	Initial Draft
0.2	02-Apr-2015	Quang Le	Ludwig Derodel	Peer Review
0.3	03-Apr-2015	Quang Le	Global Stakeholders	Release for Testing
0.4	07-May-2015	Quang Le	-	Added additional McAfee ePO instructions
0.5	13-May-2015	Quang Le	-	Modified standard hardware configurations
0.6	24-July-2015	Quang Le	-	See Appendix A – Document Version Control
0.7	4-Aug-2015	Quang Le	Gary Johnson Jeff Hernandez	Peer Review
0.8	3-Sep-2015	Quang Le	-	See Appendix A – Document Version Control
0.9	30-Sep-2015	Quang Le	-	Release for Production

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 [*CATE Storage and Backup Engineering](#).

4 Latest Versions

It is recommended that the latest version of the document be used from the [CATE Storage and Backup Engineering website](#). Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5 Regional Reviewers

The purpose of this section is to list all the stakeholders who reviewed this document during the Request for Comment process.

Region	Current Version Reviewed by	Previous Version Reviewed
Asia Pacific	No feedback	
EMEA	No feedback	
Japan	No feedback	
Latin America / Mexico	No feedback	
North America	No feedback	

Table 1 - Regional stakeholder review

6 Table of Contents

1	VERSION CONTROL.....	2
2	LEGAL NOTICE	2
3	CONTACT.....	2
4	LATEST VERSIONS.....	2
5	REGIONAL REVIEWERS	3
6	TABLE OF CONTENTS	4
7	INTRODUCTION	9
8	ASSUMPTIONS AND REFERENCES.....	10
8.1	ASSUMPTIONS	10
8.2	HOW TO USE THIS DOCUMENT.....	10
8.3	REFERENCES	10
9	GENERAL DESIGN	12
9.1	OVERVIEW	12
9.2	VMAX3 HARDWARE	13
9.2.1	VMAX3 Hardware Components	13
9.3	CITI STANDARD VMAX3 STANDALONE HARDWARE CONFIGURATIONS	15
9.4	CITI STANDARD VAMX3 SRDF HARDWARE CONFIGURATIONS	16
9.4.1	VMAX3 Supported SRDF/S Configurations.....	16
9.4.2	VMAX3 Supported SRDF/A Configurations	17
9.4.3	SRDF Network Hardware	17
10	VMAX3 INSTALLATION	18
10.1	PREREQUISITES FOR A NEW VMAX3 ENVIRONMENT	19
10.2	PHYSICAL VMAX3 STORAGE ARRAY INSTALLATION.....	19
10.3	VMAX3 HYPERMAX OS CODE	19
10.4	VMAX3 MMCS ETHERNET CONNECTIONS.....	19
10.5	VMAX3 SAN CONNECTIONS	19
11	ARRAY MANAGEMENT	20
11.1	SOLUTIONS ENABLER.....	20
11.1.1	Solutions Enabler Management Server Setup.....	21
11.1.2	SE Server System Requirements.....	21
11.1.3	New SE Installation	21
11.1.3.1	Post-Installation Procedure	22
11.1.4	SE Uninstall Procedure on RHEL.....	23
11.1.5	Upgrading Solutions Enabler	23
11.1.5.1	Pre-Upgrade Procedure	24
11.1.5.2	Upgrade Procedure.....	24
11.1.6	Backing Out of an Upgrade on RHEL	26
11.1.7	Enabling the SYMAPI Server.....	27
11.1.8	Solutions Enabler Daemon Management	29
11.1.9	ECOM Configuration	30
11.1.10	ViPR SRM	37
11.1.11	HPNA Setup.....	38
11.1.12	Group Name Services (“GNS”) Configuration	38

11.2	UNISPHERE FOR VMAX	42
11.2.1	<i>Unisphere for VMAX Architecture</i>	42
11.2.2	<i>Unisphere for VMAX Components</i>	42
11.2.2.1	U4V Server	42
11.2.2.2	SYMAPI Server	44
11.2.2.3	Unisphere Client Requirements	44
11.2.3	<i>Solutions Enabler Installation/Configuration on U4V Server</i>	44
11.2.4	<i>Solutions Enabler Upgrade for U4V Server</i>	46
11.2.5	<i>Solutions Enabler SYMAPI Configuration for U4V</i>	47
11.2.6	<i>Unisphere for VMAX Installation</i>	47
11.2.7	<i>Upgrading Unisphere for VMAX</i>	49
11.2.8	<i>Configuring Unisphere Security</i>	50
11.2.9	<i>Unisphere User Authentication</i>	51
11.2.10	<i>Unisphere User Authorization</i>	56
11.2.11	<i>User Access to Unisphere</i>	58
11.2.12	<i>Unisphere Login Banner Configuration</i>	58
11.2.13	<i>Filtering Array Discovery in Unisphere</i>	59
11.2.14	<i>McAfee Virus Scan – Addition of processes to Low Risk Process Policy</i>	63
11.2.15	<i>Web Server Certificates</i>	63
11.2.16	<i>Registering arrays for Data Collection</i>	67
11.2.17	<i>Unisphere Server backup</i>	68
11.2.18	<i>Unisphere Performance Database Backup</i>	68
11.2.19	<i>Configure Idle User Timeout Session</i>	69
11.2.20	<i>Remove authorization via local directory</i>	70
11.2.21	<i>Unisphere Alerting</i>	71
11.2.21.1	Default Alerts	71
11.2.21.2	U4V System Runtime Checks/Alerts	71
11.2.21.3	Target Enginuity level	72
11.2.21.4	Unisphere Alert Policy Configuration	73
11.2.22	<i>Unisphere Performance Thresholds and Alerts</i>	75
11.2.22.1	Recommended Thresholds	76
11.2.23	<i>Unisphere Alert Notification</i>	77
11.2.23.1	Optional Email Notification	77
11.2.24	<i>Unisphere Reports</i>	78
11.2.24.1	Dashboard Reports	78
11.2.24.2	Customized reports	80
11.2.24.3	Export Unisphere Reports	84
12	VMAX3 LICENSING	86
13	VMAX3 STORAGE ARRAY CONFIGURATION	86
13.1	SRP RESERVE CAPACITY	86
14	STORAGE PROVISIONING	87
14.1	SAN ZONING	87
14.2	PROVISION STORAGE TO HOSTS	88
15	VMAX3 SECURITY CONFIGURATION	91
15.1	SYMACL HOST ACCESS CONTROL	91
15.1.1	<i>Symacl Setup</i>	91
15.1.2	<i>SYMACL Backup and Restore</i>	95
15.2	SYMAUTH ROLE-BASED USER ACCESS CONTROL	95
15.2.1	<i>To create symauth host based user entries on the array:</i>	96
15.2.2	<i>To create symauth domain based group entries on the array:</i>	97
15.2.3	<i>LDAP/Kerberos Linux Authentication</i>	102
15.2.4	<i>Backing Up and Restoring the Symauth Database</i>	103

15.3	SECURITY COMPENSATING CONTROLS	103
15.4	TESTSMIPROVIDER REMOVAL	103
16	MONITORING AND ALERTING.....	103
16.1	CONFIGURING THE EVENT DAEMON (STOREVNTD)	103
16.2	CONFIGURING SYMMETRIX ALERTING	105
16.3	AUDIT LOG RECONCILIATION VIA SEMS.....	107
16.3.1	<i>Performance</i>	<i>108</i>
16.4	REMOTE VENDOR SUPPORT.....	108
16.4.1	<i>EMC Secure Remote Support (ESRS)</i>	<i>108</i>
17	SUPPORTED VMAX3 FEATURES	109
17.1	SRDF REMOTE REPLICATION	109
17.1.1	<i>SRDF Limitations</i>	<i>109</i>
17.1.2	<i>SRDF/A Over Very Long Distances</i>	<i>110</i>
17.1.3	<i>Role of SRDF in Disaster Recovery</i>	<i>110</i>
17.1.4	<i>SRDF Configuration.....</i>	<i>111</i>
17.1.4.1	Collection of Basic Application Requirements	112
17.1.4.2	Network Equipment Configuration.....	112
17.1.4.3	SRDF Management server requirements.....	113
17.1.4.4	RA Configuration.....	114
17.1.4.5	Mapping RDF Groups to RAs.....	116
17.1.4.6	Creating dynamic RDF device pairs.....	117
17.1.4.7	Deleting RDF Device Pairs	118
17.1.4.8	Checking the status of dynamic RDF device pairs.....	118
17.1.4.9	Deleting RDF Groups.....	119
17.1.4.10	Enabling SRDF/A replication	119
17.1.4.11	SRDF/A Cycle Time Configuration.....	120
17.1.4.12	SRDF/A Delta Set Extension (DSE) Configuration.....	120
17.1.4.13	Checking DSE Status.....	121
17.1.4.14	SRDF/A Cache Utilization	122
17.1.4.15	SRDF/A Transmit Idle Configuration	122
17.1.4.16	Checking Transmit idle status	123
17.1.5	<i>SRDF Replication Monitoring</i>	<i>123</i>
17.1.5.1	Fault alerting via storevntd log file and Tivoli Endpoint Agent	123
17.1.5.2	Monitoring with scripts (legacy or backup solution)	124
17.1.5.3	Monitoring via Unisphere for VMAX.....	124
17.1.6	<i>Target site Point-In-Time copy for SRDF/A Resynchronization</i>	<i>127</i>
17.1.6.1	Are R2 Point-in-time copies mandatory?.....	127
17.1.6.2	Recovery Process Using SnapVX Snapshots as the Gold Copy	128
17.1.6.3	Recovery Process Using BCVs (Legacy) as the Gold Copy	131
17.2	ZERO SPACE RECLAMATION	134
18	STORAGE PROVISIONING BEST PRACTICES	135
18.1	VIRTUAL PROVISIONING	135
18.2	HOST CONNECTIVITY	135
18.2.1	<i>Capacity Monitoring with SRM.....</i>	<i>135</i>
19	DATA PROTECTION BEST PRACTICES.....	136
19.1	LOCAL PROTECTION	136
19.2	REMOTE PROTECTION	136
19.3	BACKUP STANDARDS	136
20	SECURITY BEST PRACTICES.....	137
20.1	SECURITY STANDARDS	137
20.2	CONFIGURING THE EVENT DAEMON (STOREVNTD) ON THE RHEL SE SERVER	137
20.3	AUDIT LOGGING	138

20.4	MMCS ANTI-VIRUS	140
20.5	DATA SECURITY	141
20.5.1	<i>Disk Sanitization</i>	141
20.5.2	<i>Disposal of failed disks</i>	141
20.6	CITMP INFORMATION SECURITY POLICY CONVERGENCE	141
20.6.1	<i>Information Ownership</i>	141
20.7	ROLE BASED ACCESS CONTROLS (RBAC)	142
20.7.1	<i>Symauth Role-Based Access Controls</i>	142
20.8	LOCAL FID SECURITY GUIDELINES	142
20.9	STANDARD BUILD CONFIGURATION CHECKLIST (SBCC) WITH HPNA	143
21	APPENDIX A – DOCUMENT VERSION CONTROL	144
22	APPENDIX B - LIST OF SYMMETRIX EVENT CODES	144
23	APPENDIX C – 2-SITE MULTI-SESSION CONSISTENCY	149
23.1	OVERVIEW OF MSC	149
23.2	CHOICE OF CONSISTENCY MODELS	150
23.3	RUNNING MSC BETWEEN FRAMES	152
23.4	HIGH-AVAILABILITY OF MSC REPLICATION	152
23.4.1	<i>Overview</i>	152
23.5	CONFIGURING MSC SUPPORT ON MANAGEMENT SERVERS	153
23.5.1	<i>Gatekeeper requirements</i>	153
23.5.2	<i>High-availability of RDF daemon</i>	154
23.6	USE OF COMPOSITE GROUPS	154
23.6.1	<i>Scalability restrictions</i>	155
23.6.2	<i>Creating the Composite Group</i>	155
23.6.3	<i>Synchronizing Composite Group definitions across management server</i>	155
23.6.4	<i>Enabling MSC replication</i>	156
23.7	2-SITE MULTI-SESSION CONSISTENCY OPERATIONS	158
23.7.1	<i>Independent RDF group failover/failback</i>	158
23.7.2	<i>Composite group failover/failback</i>	159
23.8	CYCLE TIME UNDER MSC	160
24	APPENDIX D – TEMPLATE OPERATION SCRIPTS SUMMARY	161
24.1	SUMMARY OF SCRIPTS	162
24.2	SCRIPT MON_SRDFA_INACTIVE	162
24.2.1	<i>Overview</i>	162
24.2.2	<i>Invocation and required arguments</i>	163
24.2.3	<i>Usage notes</i>	163
24.3	SCRIPT MON_SRDFA_CACHE	164
24.3.1	<i>Overview</i>	164
24.3.2	<i>Invocation and required arguments</i>	164
24.3.3	<i>Usage notes</i>	164
24.4	SCRIPT MON_SRDFA_CYCLE	165
24.4.1	<i>Overview</i>	165
24.4.2	<i>Invocation and required arguments</i>	165
24.4.3	<i>Usage notes</i>	165
24.5	SCRIPT SRDFA_HEALTH	166
24.5.1	<i>Overview</i>	166
24.5.2	<i>Invocation and required arguments</i>	166
24.5.3	<i>Usage notes</i>	166
24.6	SCRIPT MON_SRDFA_DSE_POOL	167
24.6.1	<i>Overview</i>	167
24.6.2	<i>Invocation and required arguments</i>	167
24.6.3	<i>Usage notes</i>	167

24.7	SCRIPT SEQ_SRDFA	168
24.7.1	Overview	168
24.7.2	Invocation and required arguments	168
24.7.3	Usage notes	168
24.8	SCRIPT 2SITE_ADD_VMAX_DEVICE.....	169
24.8.1	Overview	169
24.8.2	Invocation and required arguments	169
24.8.3	Usage notes	169
24.9	SCRIPT RESUME_SRDFA.....	171
24.9.1	Overview	171
24.9.2	Invocation and required arguments	171
24.9.3	Usage notes	171
24.10	SCRIPT 2SITE_ASYNC_FAILOVER	172
24.10.1	Overview	172
24.10.2	Invocation and required arguments	172
24.10.3	Usage notes	172
24.11	SCRIPT 2SITE_ASYNC_FAILBACK.....	173
24.11.1	Overview	173
24.11.2	Invocation and required arguments	173
24.11.3	Usage notes	173
24.12	SCRIPT 2SITE_ASYNC_SWAP	174
24.12.1	Overview	174
24.12.2	Invocation and required arguments	174
24.12.3	Usage notes	174
24.13	SCRIPT 2SITE_MSC_FAILOVER.....	175
24.13.1	Overview	175
24.13.2	Invocation and required arguments	175
24.13.3	Usage notes	175
24.14	SCRIPT 2SITE_MSC_FAILBACK.....	176
24.14.1	Overview	176
24.14.2	Invocation and required arguments	176
24.14.3	Usage notes	176
24.15	SCRIPT 2SITE_CG_FAILOVER	177
24.15.1	Overview	177
24.15.2	Invocation and required arguments	177
24.15.3	Usage notes	177
24.16	SCRIPT 2SITE_CG_FAILBACK.....	178
24.16.1	Overview	178
24.16.2	Invocation and required arguments	178
24.16.3	Usage notes	178

7 Introduction

This document provides detailed installation and configuration instructions for use by all operational and regional engineering teams for EMC VMAX3 deployment within Citi's global environment, to ensure adherence to configuration, support, and management standards as well as compliance to Citigroup Information Security Standards.

This is a living document and will be updated accordingly as new products and code versions are certified in the Citi environment. Its main purpose is to provide the reader with a comprehensive view of the VMAX3 storage arrays and their usage at Citi. However, the document does not cover the internals of the VMAX3.

This document is intended for use by all offices within Citi who intend to purchase and implement a VMAX3 Storage Infrastructure for block storage purposes.

8 Assumptions and References

8.1 Assumptions

This document assumes that the reader is familiar with the basics of block storage, including: hardware platforms, data protection, LUNs and storage operations. Although experience of a previous version is not a requirement, familiarization with the product is strongly recommended before continuing to read this document.

8.2 How to use this document

This document is not intended to be read cover to cover, it is purely a reference document designed to augment the product documentation provided by EMC; provide focused and detailed insights of storage management and to consolidate important reference materials required by Citi's Information Technology Management Policies (CITMP).

8.3 References

Item	Location / Value
EMC VMAX Family (100K, 200K, 400K) Documentation Set	http://support.emc.com
EMC VMAX3 Family with HYPERMAX OS 5977 Release Notes	http://support.emc.com
Solutions Enabler Symmetrix Array Management CLI Product Guide	http://support.emc.com
EMC VMAX3 Family New Features – A Detailed Review for Open Systems Environments	http://support.emc.com
EMC Simple Support Matrix (EMC VMAX 400K, VMAX 200K, and VMAX 100k) Director Bit Settings	http://support.emc.com
EMC VMAX Family (VMAX 100K, 200K, 400K) Planning Guide	http://support.emc.com
EMC VMAX Family Security Configuration Guide	http://support.emc.com
EMC Solutions Enabler Timefinder Family CLI User Guide	http://support.emc.com
EMC Solutions Enabler SRDF Family CLI User Guide	http://support.emc.com
EMC VMAX Reliability, Availability, and Serviceability Technical Notes	http://support.emc.com
Drive Sparing in EMC Symmetrix VMAX Family Systems	http://support.emc.com
EMC VMAX3 Service Level Provisioning with Fully Automated Storage Tiering (FAST)	http://support.emc.com
VMAX Standard Build Configuration Checklist (SBCC)	CSBE Symmetrix Product Page
CSBE Solutions Enabler 8.0.2 Solution Guide	CSBE Symmetrix Product Page
VMAX3 SRDF and SnapVX Release Notes	CSBE Symmetrix Product Page


9 General Design

This section explains the general design for the implementation of an EMC VMAX3 for block-based storage, utilizing a fibre channel Storage Area Network (SAN).

9.1 Overview

The VMAX3 storage array provides block-based storage to host servers via a fibre channel SAN. The SAN is also used for management communication from the management hosts running the VMAX3 management software (e.g. Solutions Enabler and Unisphere). The management hosts will also be used for Tivoli alerting, audit log reconciliation, and monitoring. Communication with operational tools such as Tivoli and Arcsight will utilize the IP network connections on the management host. The VMAX3 storage array also has its own IP interfaces used for vendor remote access managed by the ESRS infrastructure. Below is a diagram of the network topology for the VMAX3 solution.

The diagram below depicts an overview of the infrastructure implemented for a standard Citi deployment, highlighting data and management communication.



9.2 VMAX3 Hardware

9.2.1 VMAX3 Hardware Components

The main hardware components for the VMAX3 storage array are described in this section.

System Bay

The VMAX3 system bay is a cabinet that hold all the hardware components of the VMAX3. There are 2 types of system bays, a single-engine system bay and a dual-engine system bay. . The system bay also holds the disk drive enclosures (DAE) that house the disk drives. Single-engine system bays can contain 2-6 disk array enclosures (DAE). Dual-engine system bays can contain 2-4 DAEs. The system bay also contains power supplies with LI-ON battery backup, Virtual matrix infiniband interconnect (for VMAXs with more than one engine), and component cabling.

A picture of a dual-engine system bay is shown below. If the VMAX3 has more than one dual-engine system bay, subsequent system bays will have a similar layout with the exception of the Infiniband switches and KVM. The Infiniband switches and KVM are always installed in System Bay 1.




A picture of a VMAX3 configured with single-engine system bays. The Infiniband switches and KVM are always installed in System Bay 1 just as they are in arrays configured with dual-engine system bays.



Engine

The VMAX3 engine comprises of two director boards which house the MMCSs, I/O modules for front-end ports, back-end ports, vault flash drives, and Infiniband ports for connection to the Virtual Matrix. There are a total of 11 I/O module slots numbered 0-10 from left to right on each director. Slots 0, 1, 6, and 7 are used for vault flash drives. Slots 2, 3, 8, and 9 are used for front-end connectivity (8Gbps FC, 16Gbps FC, 1Gbps Ethernet, or 10Gbps Ethernet). Citi currently only supports fibre channel for VMAX3 front-end connectivity¹. Slots 4 and 5 are used for back-end connectivity to the disk drives. Slot 10 is used for the Infiniband connections to the Virtual Matrix if the storage array has more than one engine. The picture below depicts the components of a VMAX3 engine along with the logical port numbers for the front-end connections.



¹ Only 8Gbps FC ports can be used for SRDF (RF ports). Replication (using SRDF) will be certified by CATE in Q2 of 2015.


Disk Drive Enclosure (DAE)

The DAEs come in two sizes, 120 drives (for 2.5" drives) and 60 drives (for 3.5" drives). The Citi standard configurations will use the 120 drive DAEs (DAE120). The DAE120 comprises of 4 power zones. Drives will be installed so that a single power zone failure will not cause a data unavailability scenario. For example, for a RAID 1 raid group, members of the same RAID pair will be installed in slots that are in different power zones. For RAID 5 configurations, no two members in the RAID group will be installed in the same power zone.



Infiniband Switches

There are two Infiniband switches that are installed in VMAX3 storage arrays with more than one engine. The Infiniband switches form the Virtual Matrix which is used for inter-engine communication. The 100K and 200K models use the 12 port (56Gbps/port) Infiniband switch shown below:



9.3 Citi Standard VMAX3 Standalone Hardware Configurations

The following VMAX3 standalone hardware configurations have been approved based on Citi's performance and capacity requirement for enterprise tier storage:

Tier	Model	# of Engines	System Cache (Raw)	# of Drives/Drive type/RAID	Usable Capacity (TB)
Bronze	100K	2	1TB	32 x 800GB/EFD/R5(3+1), 232 x 1200GB/SAS/R6(6+2)	204
Silver	100K	2	1TB	32 x 800GB/EFD/R5(3+1), 432 x 600GB/SAS/R1	130
Silver	200K	2	2TB	64 x 800GB/EFD/R5(3+1), 670 x 600GB/SAS/R1	210
Silver	200K	2	2TB	64 x 800GB/EFD/R5(3+1), 1042 x 600GB/SAS/R1	307

A single 2-engine system bay is used for the 100K configuration and two single-engine system bays are used for the 200K configurations above.

If a configuration with less capacity than those listed above is required; take the larger configuration that is closest to the capacity you need and decrease the number of SAS drives in that configuration to come close to the capacity you need while keeping the number of EFD (flash) drives and cache the same to maintain performance. The drive counts for RAID 5 (3+1) configurations should be divisible by 8 and RAID 1 should be divisible by 2. This ensures that the RAID groups are balanced in the engine for these RAID configurations are balanced between the directors in the engine.

The current code level for VMAX3 does not support expansions (e.g. adding engines, drives, DAEs). The ability to expand capacity and processing power non-disruptively and without data loss will be supported in later code releases. For this reason, it is important to size the capacity of the array appropriately upfront so that expansion will not be needed for at least a year out.

The drive counts in the table above do not include hot spares. At a minimum, 1 hot spare is needed for every 50 drives per disk group per engine.

9.4 Citi Standard VAMX3 SRDF Hardware Configurations

This section covers the supported SRDF hardware configurations based on modelling done by EMC that will meet Citi's performance and CoB requirements.

9.4.1 VMAX3 Supported SRDF/S Configurations

SRDF/S is supported for replication between sites that are less than 200 km apart. SRDF/S configurations do not require point-in-time copies of the data since data at both sites are always synchronized during normal operations; therefore, both the source and target site should have the same hardware configuration. Below are the Citi standard VMAX3 hardware configurations that should also be used for SRDF/S configurations.

Tier	Model	# of Engines	System Cache (Raw)	# of Drives/Drive type/RAID	Usable Capacity (TB)
Bronze	100K	2	1TB	32 x 800GB/EFD/R5(3+1), 232 x 1200GB/SAS/R6(6+2)	204
Silver	100K	2	1TB	32 x 800GB/EFD/R5(3+1),	130

				432 x 600GB/SAS/R1	
Silver	200K	2	2TB	64 x 800GB/EFD/R5(3+1), 670 x 600GB/SAS/R1	210
Silver	200K	2	2TB	64 x 800GB/EFD/R5(3+1), 1042 x 600GB/SAS/R1	307


9.4.2 VMAX3 Supported SRDF/A Configurations

The SRDF/A configurations have been sized with approximately 30% additional capacity to the standard VMAX3 configurations to support point-in-time (PIT) copies (gold copy) of the data during SRDF resynchronization operations. EMC has done some analysis on the data change rate in Citi's Symmetrix environment and results showed that the peak change rate was 3% of the capacity per day. A snapshot buffer of 30% will allow for approximately 10 days of snap space according to their analysis. The SRDF/A configurations below are based off of Citi's standard VMAX3 hardware configurations with added capacity for the snapshot buffer, delta set extension (DSE), and added cache to handle mounting a single snapshot if needed. These hardware configurations can be used at both the source and target sites to provide equivalent COB protection at both sites. These configurations are not meant to support mounting snapshots for purposes other than recovery from a failure in a replicated configuration.

Tier	Model	# of Engines	System Cache (Raw)	# of Drives/Drive type/RAID	Usable Capacity (TB)
Bronze	100K	2	1TB	48 x 800GB/EFD/R5(3+1), 304 x 1200GB/SAS/R6(6+2)	270
Silver	100K	2	2TB	48 x 800GB/EFD/R5(3+1), 500 x 600GB/SAS/R1	170
Silver	200K	2	2TB	80 x 800GB/EFD/R5(3+1), 886 x 600GB/SAS/R1	275
Silver	200K	3	3TB	112 x 800GB/EFD/R5(3+1), 1290 x 600GB/SAS/R1	402

9.4.3 SRDF Network Hardware

The following diagram shows an example 2-Site SRDF/A network configuration.



Cisco McData MPS cards are used for the FC to IP conversion. Guidance on how to plan a Cisco implementation for SRDF/A usage can be found in the GDSE document “**Cisco MDS FCIP**” Solutions Guide (see references in Section 1) and is not discussed further here.

However, it must be noted that this is a very significant piece of work and the above document should be reviewed as soon as an SRDF/A implementation is considered.

The above example shows a source Symmetrix to the left. 4 RAs are logically connected to Cisco MPS cards in Cisco 9506 Directors. The MPS cards convert SRDF traffic running over Fibre-Channel into IP traffic. The SRDF IP traffic flows across the network to the remote site, where it is converted back into FC and into the target Symmetrix.

The above example is for a 2-site SRDF/A configuration. The number of MPS cards required would depend on the number of RAs to be used on the SRDF/A leg, which in turn depends on the number of RDF groups to be supported. Determination of the number of RAs required is discussed in a later sub-section. Advice on sizing the infrastructure is also covered in the “**Cisco MDS FCIP and “WAN Throughput Optimization for SRDF/A Solutions Guide with Cisco MDS**” Solutions Guides referenced previously.

Each member of pair of RAs supporting a given set of RDF groups should be connected to a different MPS card. Port Channelling is used to aggregate the GigE ports together.

Note: Only 8 Gbps FC ports on the VMAX3 can be used as RAs. 16 Gbps FC ports cannot be used as RAs.

The links via the IP network that are used for SRDF/A replication:

1. Must be provided by your Networks team with dedicated or QoS guaranteed bandwidth
2. Must be redundant (primary/secondary configuration)
3. Must be set up so that under normal running, the links have a single fixed latency i.e. there is no load balancing across two different routes with differing latencies.
4. The data transfer must be via the Citigroup network; else it would need to be encrypted.

10 VMAX3 Installation

This section provides the steps needed to install the array into an enterprise datacenter and connect to its infrastructure.

10.1 Prerequisites for a new VMAX3 Environment

The following items are required in order to implement a VMAX3 storage array according to CSBE standards:

- A RHEL 5.x or 6.x server for Solutions Enabler
- A Windows server for Unisphere for VMAX and Solutions Enabler (Windows version)
- 2 IP connections to the CEP network for each VMAX3 storage array
- Citi standard monitoring infrastructure environment
- Citi standard SEMS environment
- Appropriate datacenter requirements (e.g. physical floor space, cooling, power)
- Fibre channel SAN
- ESRS infrastructure for vendor remote access and monitoring

10.2 Physical VMAX3 Storage Array Installation

Reference [EMC VMAX Family Planning Guide](#) for hardware specifications, datacenter requirements (e.g. power, cooling, floor-load-bearing capacity, etc.), and shipping/delivery instructions and logistics. EMC will power-on and run health-checks on the array once the VMAX3 is physically installed in the datacenter. EMC should also ensure that the 'Disk Service' attribute is set to 'Normal' so that a request to replace the drive occurs immediately after a drive failure.

The following command can be issued to verify the 'Disk Service' setting:

```
# symcfg -sid <SymmID> list -v
```

Example: `symcfg -sid 973 list -v | grep "Disks Service"`

```
Disks Service          : Normal
```

10.3 VMAX3 Hypermax OS Code


Refer to the [CSBE Symmetrix Configuration Matrix](#) to determine what array code is supported. EMC should install/upgrade the VMAX3 to the most recently certified Hypermax OS version specified in the matrix.

10.4 VMAX3 MMCS Ethernet Connections

Each VMAX3 storage array will need two Ethernet connections to the Citi CEP network for vendor remote management and support via ESRS. The Ethernet connections will be connected to the VMAX3's primary and secondary MMCS Ethernet management interfaces. Once connected, EMC will configure the management interfaces and add the VMAX3 storage array to the ESRS infrastructure. EMC will verify that ESRS discovers the VMAX3 successfully. Refer to the [ESRS Solutions Guide](#) for ESRS setup.

10.5 VMAX3 SAN Connections

The VMAX3 will be connected to the core switches of dual-fabric fibre channel SAN to provide front-end connectivity to the array to host servers. FA ports on the odd numbered directors should be connected to Fabric A. FA ports on the even numbered directors should be connected to Fabric B. Connecting the VMAX3 this way will ensure that each engine has connectivity to each fabric in the SAN to guarantee redundancy and resiliency. The diagram below depicts the standard SAN connectivity for a VMAX3 populated with 4 engines:



11 Array Management

This section covers the management infrastructure that is needed to manage the VMAX3 storage array.

11.1 Solutions Enabler

The Solutions Enabler (SE) software application from EMC allows configuration and management of the EMC Symmetrix storage arrays. It consists of commands that can be invoked via the command line, or within scripts. The Citigroup standard configuration of Solutions Enabler is installed on a SOE Linux host that is connected to the Storage Area Network (SAN). The application utilizes the fibre channel transport to send and receive commands from the array.

11.1.1 Solutions Enabler Management Server Setup

The RHEL 5.x or 6.x Solutions Enabler (SE) management server(s) will need to be setup in order to manage and administer the VMAX3 storage array. An SE management server can manage multiple arrays. The following high-level steps will need to be completed in order for the SE server to manage the VMAX3:

- The SE management server(s) will need to be connected to both fabrics in the SAN and zoned to the FA port with the ACLX device visible (FA port 1d4 by default) for initial discovery.
- Symacl will need to be configured entry to allow the SE management server to manage the VMAX3.
- 6 Gatekeepers per array will need to be provisioned each SE management server.
- Symauth will need to be configured to control/allow users on the SE management server to manage the array with the appropriate rights.
- Setup the SE management server to monitor events/logs on the VMAX3 for Tivoli Alerting and SEMS audit log reconciliation.

Once gatekeepers are discovered by the SE management server(s) and the server(s) can successfully discover the array(s), the 'Show_ACLX_Device' flag should be disabled on all ports so that the ACLX device does not show up when application servers are zoned to those ports.

For redundancy, it is recommended that two hosts be used to manage the same set of Symmetrix storage arrays. These hosts do NOT require shared storage (clustering configuration) from the SAN; each host should have the Group Name Services enabled to share device group information. Please see [Group Name Services \("GNS"\) Configuration](#) on how the Group Name Service (GNS) is to be configured.

11.1.2 SE Server System Requirements

- RHEL 5.x or 6.x server
- Minimum amount of disk space needed: 352 MB

11.1.3 New SE Installation

Steps to install SE on a RHEL server:

1. The SE package for RHEL is located in the following directory

```
# cd /net/stealth/export1/home1.localhost/sw/Linux
```

2. Copy the SE package to SE server and expand file

```
# cp se8030-Linux-x86_64-ni.tar.gz <target location>
# gunzip se8030-Linux-x86_64-ni.tar.gz
# tar -xvf se8030-Linux-x86_64-ni.tar
```

3. Run the following command to install the software:

```
# ./se8030_install.sh -install -silent
```

- Run the following to install add-on software:

```
# ./se8030_install.sh -install -silent -srm -smis
```

- Verify that the SE packages are at 8.0.1.xxx:

```
# ./se8030_install.sh -check

#-----
#                               EMC Installation Manager
#-----
Copyright (c) [1997-2015] EMC Corporation. All Rights Reserved.

This software contains the intellectual property of EMC Corporation or is
licensed to EMC Corporation from third parties. Use of this software and the
intellectual property contained therein is expressly limited to the terms and
conditions of the License Agreement under which it is provided by or on behalf
of EMC.

Checking for Solutions Enabler Native Installer kit Installation.....

Sl No RPM                               Version
-----
 1  symcli-base                          8.0.3.2026-0.3.x86_64
 2  symcli-cert                          8.0.3.2026-0.3.x86_64
 3  symcli-data                          8.0.3.2026-0.3.x86_64
 4  symcli-jni                           8.0.3.2026-0.3.x86_64
 5  symcli-smi64                         8.0.3.2026-0.3.x86_64
 6  symcli-srm                           8.0.3.2026-0.3.x86_64
 7  symcli-symcli                        8.0.3.2026-0.3.x86_64
 8  symcli-symrecover                    8.0.3.2026-0.3.x86_64
 9  symcli-thincore                      8.0.3.2026-0.3.x86_64

All the Solutions Enabler components are installed!
```

11.1.3.1 Post-Installation Procedure

- Set the following in the **.bash_profile** file:

```
PATH=$PATH:/usr/storapi/bin:/usr/storapi/storbin:
MANPATH=$MANPATH:/usr/storapi/man:/usr/storapi/storman
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/storapi/shlib:
export PATH MANPATH LD_LIBRARY_PATH
```

- A minimum of 6 gatekeepers should be assigned to the server if that hasn't already been done. Additional gatekeepers can be added at a later time if commands get delayed or rejected for lack of gatekeeper resources.
- Run the following command to build the configuration (SYMAPI) database:

```
# symcfg discover
```

- Start the stordaemons if they are not running (storapi, storgns, storrdfd, storevntd, storwatchd, storsrmd, storstd and storsrvd):

```
# stordaeomon start <daemon name>
```

5. Run the following command to start the stordaemons automatically after a system reboot.

```
# stordaemon install storapid -autostart
# stordaemon install storgnsd -autostart
# stordaemon install storrdfd -autostart
# stordaemon install storevntd -autostart
# stordaemon install storwatchd -autostart
# stordaemon install storsrmd -autostart
# stordaemon install storstd -autostart
# stordaemon install storsrvd -autostart
```

6. Start ECC Agents:

```
# <ECC Agent install dir>/ecclexec/eccmad start
```

7. For DMX arrays running 5773 code, the VCM device is used for the initial discovery of the storage arrays by a Solutions Enabler host. Once the DMX is visible to the host, a sufficient number of gatekeepers should be assigned to the host. After the Solutions Enabler host has successfully discovered the assigned gatekeepers, the VCM device should be unmapped from all FAs on the DMX so that it will not be visible to other hosts.

The VMAX/VMAX3 contains an ACLX device that is used for the initial discovery of the storage arrays similar to the VCM device on the DMX. By default, the ACLX device is mapped to every front-end adapter port on the VMAX and to FA 1d4 on VMAX3. Once hosts are zoned to the front-end ports, the ACLX device will automatically be visible to those hosts for the initial discovery of the array and will function as the communications medium until gatekeepers on the array are provisioned to the host. As a best practice, once gatekeepers have been provisioned and discovered on the Solutions Enabler management server(s), the ACLX device should be unmapped from all FAs so that it is not visible to any host. This will alleviate any attempt to use the device as a storage device and stop any system messages from flooding the logs due to an unlabeled device.

Warning: If this is the initial management server for the Symmetrix storage array, ensure that gatekeepers have been assigned and discovered on the server before unmapping the ACLX device from the FAs.

11.1.4 SE Uninstall Procedure on RHEL

To completely remove SE on the server, do the following:

```
# stordaemon shutdown all -immediate
# <ECC Agent install dir>/ecclexec/eccmad stopall
<Locate SE install shell and run with uninstall option>
# ./se8030_install.sh -uninstall
```

11.1.5 Upgrading Solutions Enabler

Follow the steps in this section to upgrade to the most current certified version of SE on RHEL.

11.1.5.1 Pre-Upgrade Procedure

1. Run the following command to update the configuration database:

```
# symcfg discover
```

2. Create a backup of the masking database of each Symmetrix local to the SE server being upgraded:

For VMAX models, issue the command:

```
# symaccess -sid <SymmID> backup -file <BackupFilename>
```

For DMX4 models and earlier, issue the command:

```
# symmaskdb -sid <SymmID> backup -file <BackupFilename>
```

3. As a precautionary step, export all device and composite groups to a file as a backup:

```
# symdg exportall -f <filename>
```

```
# symcg exportall -f <filename>
```

4. Record the current symacl unique ID for the server in case the upgrade process changes the unique ID of the server:

```
# symacl -unique > /var/tmp/namCSBE14_uniqueID_preupgrade
```

5. Stop all Solutions Enabler daemons:

```
# stordaeomon shutdown all -immediate
```

6. Stop all ECC agents:

```
# <ECC Agent install dir>/ecc/exec/eccmad stopall
```

7. Make a backup copy of the `/var/symapi/db` directory:

```
# mkdir /var/tmp/symapi_bk
```

```
# mkdir /var/tmp/symapi_bk/db_old
```

```
# cd /var/symapi/db
```

```
# find . -print|cpio -dumpv /var/tmp/symapi_bk/db_old
```

8. Make a backup copy of the `/var/symapi/config` directory:

```
# mkdir /var/tmp/symapi_bk/config_old
```

```
# cd /var/symapi/config
```

```
# find . -print|cpio -dumpv /var/tmp/symapi_bk/config_old
```

9. Make a backup copy of the `/opt/emc/ECIM/ECOM/` directory:

```
# mkdir /var/tmp/ECOM_bk
```

```
# cd /opt/emc/ECIM/ECOM
```

```
# find -print|cpio -dumpv /var/tmp/ECOM_bk
```

10. Perform a backup of the system if possible.

11.1.5.2 Upgrade Procedure

1. Complete the pre-upgrade procedure in the previous section.

2. The SE package for RHEL is located in the following directory

```
$ cd /net/stealth/export1/home1.localhost/sw/Linux
```

3. Copy the SE package to SE server and expand file

```
$ cp se8030-Linux-x86_64-ni.tar.gz <target location>  
$ gunzip se8030-Linux-x86_64-ni.tar.gz  
$ tar -xvf se8030-Linux-x86_64-ni.tar
```

4. Run the following command to install the software:

```
$ ./se8030_install.sh -install -silent
```

5. Run the following command to install add-on software:

```
# ./se8030_install.sh -install -silent -srm -smis
```

6. Verify that the Solutions Enabler packages are at version 8.X:

```
# ./se8030_install.sh -check

#-----
#                               EMC Installation Manager
#-----
Copyright (c) [1997-2015] EMC Corporation. All Rights Reserved.

This software contains the intellectual property of EMC Corporation or is
licensed to EMC Corporation from third parties. Use of this software and the
intellectual property contained therein is expressly limited to the terms and
conditions of the License Agreement under which it is provided by or on behalf
of EMC.

Checking for Solutions Enabler Native Installer kit Installation.....

SI No RPM                                     Version
-----
1  symcli-base                               8.0.3.2026-0.3.x86_64
2  symcli-cert                               8.0.3.2026-0.3.x86_64
3  symcli-data                               8.0.3.2026-0.3.x86_64
4  symcli-jni                                8.0.3.2026-0.3.x86_64
5  symcli-smi64                              8.0.3.2026-0.3.x86_64
6  symcli-srm                                8.0.3.2026-0.3.x86_64
7  symcli-symcli                             8.0.3.2026-0.3.x86_64
8  symcli-symrecover                         8.0.3.2026-0.3.x86_64
9  symcli-thincore                           8.0.3.2026-0.3.x86_64

All the Solutions Enabler components are installed!
```

7. Run the 'symacl -unique' to check that the unique ID of the server hasn't changed from the unique ID recorded during the pre-upgrade procedure. If the unique ID has changed, use another SE server that has symacl permissions to the arrays and create a new entry for the server with its new unique ID on each array it's going to manage.

8. Rediscover the arrays:

```
# symcfg discover
```

9. Verify that prior settings in the SE configuration files have persisted after the upgrade: symapi_licenses.dat, daemon_options, netcnfg, options. If for some reason the settings did not persist, they can be restored from the backup copy.

10. If GNS was enabled, the device and composite groups on the arrays before the upgrade should still exist. If the device and composite groups were stored locally on the server and were exported before the upgrade, they can be restored by issuing the following commands:

```
# symdg importall -f <filename of file specified in the exportall command>
```

11. Start the stordaemons if they are not running (storapi, storgns, storrdfd, storevntd, storwatchd, storsrmd, storstd and storsrvd):

```
# stordaemon start <daemon name>
```

12. Run the following command to start the stordaemons automatically after a system reboot.

```
# stordaemon install storapid -autostart
# stordaemon install storgnsd -autostart
# stordaemon install storrdfd -autostart
# stordaemon install storevntd -autostart
# stordaemon install storwatchd -autostart
# stordaemon install storsrmd -autostart
# stordaemon install storstd -autostart
# stordaemon install storsrvd -autostart
```

13. Start ECC Agents:

```
# /export/opt/ecc/exec/eccmad start
```

11.1.6 Backing Out of an Upgrade on RHEL

If it is needed to back out of a SE upgrade, do the following:

1. Verify that there is a backup copy of the configuration files from the previous version which should have been created during the upgrade process in the /backup/se8.X_rhel/symapi_bk directories.

2. Stop all SE and ECC daemons:

```
# stordaemon shutdown all -immediate
# <ECC Install Dir>/ecc/exec/eccmad stopall
```

3. Uninstall the current version of the SE:

```
# ./se8XXX_install.sh -uninstall
```

4. Install the prior version of SE that was installed on the server before the upgrade.

```
$ cd /net/stealth/export1/home1.localhost/sw/PKG/SE
$ cp se76XX-Linux-i386-ni.tar.gz <target location>
$ gunzip se76XX-Linux-i386-ni.tar.gz
$ tar -xvf se76XX-Linux-i386-ni.tar
# ./se7X_install.sh -install -silent
```

5. Check to see if the accid has changed for the server after the rollback by comparing the value against the backup copy.

```
# symacl -unique
```

If the value has changed, the symacl entry will need to be removed and re-added with the new acid for this server. You will need to do this from another SE server that has access to the arrays. If there are no other SE servers, EMC will need to re-add the symacl entries.

6. If GNS was not enabled on the server, copy the backup copy of the SYMAPI database into the **/var/symapi/db**

```
# cd /backup/se7.X_rhel/symapi_bk/db_old
# find .| cpio -dumpv /var/symapi/db
```

7. Copy the old **symapi_licenses.dat** file back to the **/var/symapi/db** directory:

```
# cp /backup/se7.X_rhel/symapi_bk/config_old/symapi_licenses.dat
/var/symapi/config
```

8. Copy the backup copy of the **options** and **daemon_options** file back to the **/var/symapi/config** directory.

9. Run the following command to rebuild the configuration (SYMAPI) database:

```
# symcfg discover
```

10. Start the stordemons if they are not running (storapi, storgns, storrdfd, storevntd, storwatchd, storsrmd, storstd and storsrvd):

```
# stordemon start <daemon name>
```

11. Run the following command to start the stordemons automatically after a system reboot.

```
# stordemon install storapid -autostart
# stordemon install storgnsd -autostart
# stordemon install storrdfd -autostart
# stordemon install storevntd -autostart
# stordemon install storwatchd -autostart
# stordemon install storsrmd -autostart
# stordemon install storstd -autostart
# stordemon install storsrvd -autostart
```

12. Start ECC Agents:

```
# <ECC Install Dir>/ecc/exec/eccmad start
```

11.1.7 Enabling the SYMAPI Server

The SYMAPI server that runs on the SE server and provides a method for a client server that doesn't have fibre channel access to the storage array to manage the array. The client server connects to the SE server's SYMAPI server via the TCP/IP network to issue SYMCLI commands. These commands are then executed via fibre channel from the SE server to the array.

Enabling SYMAPI server functionality

On the client server (e.g. SANscreen server), add the following to the `/var/symapi/config/netcnfg` file:

```
SYMAPI_SERVER - TCPIP - <SE server's IP address> 2707 -
```

For example:

```
SYMAPI_SERVER - TCPIP - 169.193.70.169 2707 -
```

This instructs SYMCLI commands to connect to the SE server's SYMAPI server.

To specify which hosts on the TCP/IP network can use the SYMAPI server, the "nethost" file has to be configured on the SE server. If it does not already exist create and edit the `/var/symapi/config/nethost` file with the host address:

```
<this-host-IP-address or FQDN> <username>
```

Where <username> is the login ID of an authorized member of the SAN management team.

For example:

```
169.193.70.169 SYSTEM #SANscreen requires the SYSTEM user.
```

The SYMAPI server should now be started with:

```
# stordaeomon start storsrvd
```

The following command should be executed to start the storsrvd automatically after a system reboot:

```
# stordaeomon install storsrvd -autostart
```

This creates the `S99emc_storsrvd` startup script in the `/etc/rc3.d` directory.

Only authorized storage operators should be configured to log on to the host². Each operator should be assigned their unique user ID, such as SOE ID, to logon to the host. The user IDs so assigned should be unprivileged (non-root)³.

Each storage operator's account should be profiled with the following:

The following should be added to each user's profile file on the client server:

```
PATH=$PATH:/usr/storapi/bin;usr/storapi/storbin
MANPATH=$MANPATH:/usr/storapi/man:/usr/storapi/storman
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/storapi/shlib
```

² CSBE does not recommend sharing hosts that are used to manage storage with other applications for security reasons.

³ These are information security requirements. This will allow for tracking of storage operator activity on the server.

```

SYMCLI_CONNECT=SYMAPI_SERVER
SYMCLI_CONNECT_TYPE=REMOTE
export PATH MANPATH LD_LIBRARY_PATH SYMCLI_CONNECT SYMCLI_CONNECT_TYPE

```

The “SYMCLI_CONNECT” environment variable should only be set to “SYMAPI_SERVER” on the SE client server and NOT on the SE server. This will force the SYMCLI commands to use the SYMAPI server for executing the commands, instead of directly using the SYMCLI libraries⁴.

The command ‘**symcli -def**’ will show if this environment variable has been set.

11.1.8 Solutions Enabler Daemon Management

The following command should be run for the storevntd, storgnsd, storwatchd, storrdfd, storstpd, and storsrvd to ensure that these daemons are started after a reboot of the SE server:

```
# stordaeon install <daemon name> -autostart
```

If there is a monitoring system available (e.g. Tivoli ITM), it should be used to monitor the storevntd, storwatchd, storrdfd (for arrays with srdf configured), and storsrvd to ensure that these daemons are running on the SE management station. The storwatchd daemon will try to automatically restart any daemon that has crashed. If a daemon was started while the storwatchd was down, that daemon will not be monitored by the storwatchd. To verify that a daemon is being watched by storwatchd, issue the following command:

```
# stordaeon show storwatchd
```

```

Daemon State           : Running
Daemon Start Time      : Mon Mar 9 16:05:13 2015
Version                : V8.0-1953 (1.0)

Total Number of Connections : 11
Number of Active Connections : 6
Total Number of Requests    : 9
IPC Authentication         : Passcred


Daemon Restarts Performed   : 0
Daemons Being Monitored

storapid                 : Normal Autorestart
storgnsd                 : Normal Autorestart
storrdfd                 : Normal Autorestart
storstpd                 : Normal Autorestart
storevntd                : Normal Autorestart
storsrvd                 : Normal Autorestart

```

11.1.9 ECOM Configuration

ECOM is needed to provide applications such as EMC ViPR to discover the VMAX3 via SMI-S. A default account named 'admin' with password '#1Password' can be used to access the ECOM administration page (<https://<SE server hostname>:5989/ecomconfig>). The default password for ECOM must be changed via the ECOM administration page. The options in the ECOM administration is show in the screenshot below:



An alternate user account can be created for applications such as ViPR SRM that will be using SMI-S to discover the Symmetrix storage arrays. This can be done by selecting **Add User** from the ECOM administration page and entering the user information:

ECOM Administration

ECOM Version 2.8.3.0.0.109 (Build Date & Time: Apr 28 2015, 16:32:29)

Logged in as admin

Logging:

[Display Log File](#)

[Logging Options](#)

Security:

[Add User](#)

[Modify User](#)

[Change Password](#)

[Set Password quality](#)

[Delete User](#)

[List Users](#)

[Display Security Log File](#)

[Client IP Filtering](#)

[Local IP Filtering](#)

[SSL Certificate Management](#)

[LDAP Configuration](#)

[OSLogin Configuration](#)

[Dynamic Settings](#)

[List Running Providers](#)

[Simple CIM Browser](#)

Add User:

User Name:

Password:

Re-Enter Password:

Role: ▼

Scope: ▼

Account settings:

- ☐ User must change password at next login
- ☐ User cannot change password
- ☒ Password never expires
- ☐ User is disabled

[Back to ECOMConfig Page](#)

All HTTP access to ECOM should be changed to use HTTPS. To do this, set the <secure> port settings to "true" in the Port_settings.xml file located in /opt/emc/ECIM/ECOM/conf. The settings should look like the following:

```
<ECOMSettings>
<ECOMSetting Name="Port0">
<!--
  <portRange>5988</portRange>
-->
  <port>5988</port>
  <secure>true</secure>
  <slp>true</slp>
</ECOMSetting>

<ECOMSetting Name="Port1">
<!--
  <portRange>5989</portRange>
-->
  <port>5989</port>
  <secure>true</secure>
  <slp>true</slp>
</ECOMSetting>

<ECOMSetting Name="Port2">
<!--
  <portRange>5985</portRange>
-->
  <port>5985</port>
  <secure>true</secure>
  <slp>true</slp>
</ECOMSetting>

<ECOMSetting Name="Port3">
<!--
  <portRange>5986</portRange>
-->
  <port>5986</port>
  <secure>true</secure>
  <slp>true</slp>
</ECOMSetting>

<!-- CIM-RS ports -->
<ECOMSetting Name="Port4">
<!--
  <portRange>5993</portRange>
-->
  <port>5993</port>
  <secure>true</secure>
  <slp>true</slp>
</ECOMSetting>

<ECOMSetting Name="Port5">
<!--
  <portRange>5994</portRange>
-->
  <port>5994</port>
  <secure>true</secure>
  <slp>true</slp>
</ECOMSetting>
```

In order to access this page, the ecomtls.crt certificate in the /opt/emc/ECIM/ECOM/conf/ssl directory will need to be imported in the client browser.

A Citi signed certificate must be imported for the ECOM administration page by doing the following:

1. From the ECOM administration page, select **SSL Certificate Management**:

ECOM Administration

ECOM Version 2.8.3.0.0.109 (Build Date & Time: Apr 28 2015, 16:32:29)
Logged in as admin

Logging:

[Display Log File](#)

[Logging Options](#)

Security:

[Add User](#)

[Modify User](#)

[Change Password](#)

[Set Password quality](#)

[Delete User](#)

[List Users](#)

[Display Security Log File](#)

[Client IP Filtering](#)

[Local IP Filtering](#)

[SSL Certificate Management](#)

[LDAP Configuration](#)

[OSLogin Configuration](#)

[Dynamic Settings](#)

[List Running Providers](#)

[Simple CIM Browser](#)

[Logout](#)

2. Fill out the Certificate Signing Request form:

SSL Certificate Management

Option #1: Signed Certificate

Sign a ECOM generated Certificate Signing Request (CSR).

1. Generate a Certificate Signing Request.
2. Export the PKCS#10 formatted Certificate Signing Request.
3. Get Certificate Signing Request signed by your Certificate Authority (CA).
4. Import the signed certificate as PEM data.

See Wikipedia for more information about [Certificate Signing Requests](#).

Generate a Certificate Signing Request

Import Signed Certificate

Certificate Signing Request

Option #1: Signed Certificate

Sign a ECOM generated Certificate Signing Request (CSR).

1. [Generate a Certificate Signing Request](#).
2. Export the PKCS#10 formatted Certificate Signing Request.
3. Get Certificate Signing Request signed by your Certificate Authority (CA).
4. Import the signed certificate as PEM data.

See Wikipedia for more information about [X.509 Subject fields](#).

Common Name (<i>semicolon separated list</i>):	namgds14.nam.nsroot.net
CountryName (2 letter code):	US
State or Province Name:	NJ
LocalityName:	Warren
Organization Name:	Citigroup Inc.
Organizational Unit Name:	149141
Email Address:	quang.le@citi.com
SAN IP (<i>semicolon separated list</i>)	
SAN URI (<i>semicolon separated list</i>)	
KeyUsage (<i>semicolon separated list</i>)	DigitalSignature;NonRepudiation

Generate a Certificate Signing Request

Back

3. Copy the CSR and submit a request for a signed certificate at https://cmp.nj.smb.com/marketplace/control/product/~productId=47119_29952_GLOBAL:

Certificate Signing Request

Option #1: Signed Certificate

Sign a ECOM generated Certificate Signing Request (CSR).

1. Generate a Certificate Signing Request.
2. Export the PKCS#10 formatted Certificate Signing Request.
3. Get Certificate Signing Request signed by your Certificate Authority (CA).
4. Import the signed certificate as PEM data.

See Wikipedia for more information about [PKCS#10](#).

```
-----BEGIN CERTIFICATE REQUEST-----
MIIDBTCCAeOCAQAwDEPMA0GA1UECzMGMTCQ5MTQxMRcwFQYDVQQKEw5DaXRpZ3Jv
dXAqSW5jLjEPMMA0GA1UEBhMGMGV2FycmVudQswCQYDVQQIEwJOSjEELMAkGA1UEBhMC
VVMxITAfBgNVBAMTGG5hbWdke2UxNC5uYW0ubnNyb290Lm5ldDCCASIwDQYJKoZI
hvcNAQEBBQADggEPADCCAQoCggEBBAK8kFJ1LPtTpGHGANCjfiFgEgQBvK2zc5yW
o8smvGGfEmIVCjuurUf2J7m6Wt9UFVotpACqGidH1RLJ6VomNo9okSoutyywRES
6ZaWBMTT6wsFkSp1TPx/n0CMnph21AYV1z+n0QqzGwulM8cPMKwddDi48cmv3256
ggvJ8kz1taGLw7AeRQkJFN2rCw5j0aF7JDw8thWiWgTRauNxJxNFEonxjw1Ri3M
4rHnTrhDMky03fawxjkAoxbVq+/Aapky7yic+kQCJ5fUdby6gaspapPvkXkEeLoA
NI18LvRM2fDqfqt7aJveRo9xeFe0pQR65NjeYYmEaYpWJQ00JEMCAwEAAABIMEYg
CSqGSIb3DQEJJDjE5MDcwHAyDVR0RBBUwE4ERcXVhbmducGVAY210aS5jb20wCQYD
VR0TBAlwADAMBgNVHQ8EBQMD8/+AMA0GCSqGSIb3DQEBwUAA4IBAQC8cYBPe6Lh
9aR1L+qk2sqC5QEyh0a05cu5mGvOM9x3pYwJBr3kO/TpFC7SxnTJzreGNiz8Aq7t
jFfJH3m64kr2jp9zPQCy/HmoXZIoG5SSTK+KMV6pzMMA4nltRADOfozCN7+A8S270
33bSRMer+0FX92KEz3FtALXt/V2jhv4/EnBAQ/1rtu/OGqy4eI8yM1LKIF3fx3M
qpND1vFhFwK9cdOfFNRp/GkH9q9/vM2CwFRVEVgY8v0zvA6NvVm1k8v1IerKcUWS
XYzuoV5xkIo22522ARMjUZsd9Noc6zYTwhQ8ewhI3+P/NF/OjgCv8gfvpH2ze96
OyI7Yf3eUiY+
-----END CERTIFICATE REQUEST-----
```

[Continue submit the signed certificate](#)

Back

Sample CMP request for a signed certificate:

require:	
Which Certificate Authority do you require for the certificate?:	DEVELOPMENT
Environment:	DEV
Target Hostname - DEV:	namgdse14.nam.nsroot.net
What Type of Certificate is Required:	MS CA Web Server
Port Number (default 443):	5993
Service Generated or User provided CSR:	CSR
Enter CSR:	-----BEGIN CERTIFICATE REQUEST----- MIIDBTCCAe0CAQAwDEPMA0GA1UECzMGMtQ5MTQxMRcwFQYDVQKew5DaXRpZ3Jv dXAgSW5jLjEPMA0GA1UEBxMGV2YcmVuMQswCQYDVQKIEwJOSjELMAkGA1UEBhMC VVMxITAFBgNVBAMTGG5hbWdkc2UxNC5uYW0ubnNyb290Lm5ldDCCASlwDQYJKoZI hvcNAQEBBQADggEPADCCAQoCggEBBAK8kFJILPtTPGHGANCjflfgEgQBuVk2zc5yW o8smvGGfEmIVCjuurUfZJ7m6Wt9UFVotpACqGiDhHJRLJ6VomNo9okSoutyywRES 6ZaWBMtT6wsFkSp1TPr/n0CMnph2IAyVlz+n0QqzGwulM8cPMKwddDi48cmv3256 ggvJ8kzltGLw7AeRQkJFN2rCw5j0aF7JDw8thWiWgTRauNxJxNF Eonxjw1Ri3M 4rHnTrhDMky03fawxjkAoxbVq+/AapKY7yic+kQCJ5fUdby6gzpatPvkXkEcLoA NII8LvRM2fdqfT7aJveRo9xeFe0pQR65NjeYYmEaYpWJQ0JEMCAwEAAaBIMEYg CSqGSIb3DQEJDDjE5MDcwHAYDVR0RBBUwE4ERcXVhbm cubGVAY2l0aS5jb20wCQYD VR0TBAlwADAMBgNVHQ8EBQMDB/+AMA0GC SqGSIb3DQEBCwUAA4IBAQCScYBP6Lh 9aR1L+qk2sqC5QEyh0a05cu5mGvOM9x3pYwjBr3kO/TpfC7SxnTJsreGNiz8Aq7t jffJH3m64kr2jp8zPQCy/HmoXZIoG5STK+KMyV8pzMA4nltRADOfosCN7+A85Z70 33bSKMer+0FX92KEz3FtALXt/V2jhw4/KnBAQ/1rtu/OGqy4el8yMILKIF3fx3M qpND1vFhWk9cdOffNRp/GkH9q9/vM2CwfrVEVgY8v0zvA6NvVm1k8villerKcUWS XYzuaoV5xklo2Z5Z2ARMjUZsd9NOc6zYTwHQSeWhl3+P/NF/OjgCv8gfvPHZze96 Oyl7Yf3eUiy+ -----END CERTIFICATE REQUEST-----
Generate the output in Base64 (DER/PKCS #12 coming soon):	Base64

- Import the signed certificate once it is received:

Submit CA Certificate

Option #1: Signed Certificate

Sign a ECOM generated Certificate Signing Request (CSR).

1. Generate a Certificate Signing Request.
2. Export the PKCS#10 formatted Certificate Signing Request.
3. Get Certificate Signing Request signed by your Certificate Authority (CA).
4. Import the signed certificate as PEM data.

See Wikipedia for more information about [PEM format](#).

```

ZHBwMSxDTj1DRFA=Q049UHVibG1jJTIwS2V5JTIwU2VydmljZXM=Q049U2Vydmlj
ZXM=Q049Q29uZmlndXJhdG1vbixEQ=1uc3Jvb3Q=REM9bmV0P2N1cnRpZmljYXR1
UmV2b2NhdG1vbkkpc3Q/YmFzZT9vYmplY3RDG1vY3RDbGFzc31jUkxkaXN0cmli
dXRpb25Qb21udDCCAVGgCCsGAQUFBwEBBIIIBSjCCAUYwfwQYIKwYBBQUHMAKGcW
h0dHA6Ly9jZKJ0aWZpY2F0ZKJ1dm9jYXRpb25saXN0LmN1cnRpZmljYXR1cy50a
S5jaXRpb29Jv dXAubmV0L1BL3VJvb3QvQ210aSUyME1udGVybmFzJTIwRGV2aW
N1Q0E1MjBVbnRy dXN0ZWQwY3J0MIEHBggrBgEFBQcwAoaBt2xkYXA6Ly8vY249
Q210aSUyME1udGVybmFzJTIwRGV2aW N1Q0E1MjBVbnRydXN0ZWQ=Q049QU1B
LENOFVB1YmwpYyUyMEt1 eSUyMFN1cnZpY2VzLENOPVN1cnZpY2VzLENOPUNvbm
ZpZ3VYXRpb24sREM9bnNy b290LERDPW51dD9jQU1cnRpZmljYXR1P2Jhc2U/b2
JqZWNOQ2xhc3M9Y2VydGlm aW NhdG1vbkkF1dGhvcml0eTA8BgkrBgEEAYI3FQc
ELzAtBiUrBgEEAYI3FQcEgN8s h/7g4eFjRqE1pxOk/1CG4OsmQqE5YtbAgFkAgE
IMBMGA1UdJQQMMaGgCCsGAQUF BwMBMBsGCSsGAQQBgjcVCgQOMAwCgYIKwY
BBQUHAAwEwRAYJKoZIhvcNAQkPBDcw NTA0BgqqghkiG9w0DAgICAIAwDgYIKo
ZIhvcNAwQCAgCAMAcGBSsOAwIHMAoGCCqG S1b3DQMhMAoGCSqGSIb3DQEBB
QUAA4IBAQC2311QpR2Dz5Co6k3jwU+L91VwWax4 bkYrQCz2Pria+530DbCS13F
57gq1YQuDkVgiP8D20UPongYN4GM2jkmnjEBEWNzM eFpu7gVzuEEnOqPFD37
U82QMdyys2JJBkoGRngO3wg9mUAVjG18+IiW+eNe11S zZLOsmzOqja4Yu0sh4
5118aF2B9UYunnLusoe2eopa8HjtUQn3xxa1NgUUVGVLGd Aw7jt2aCzDeyt3Qv
339fTiIdmwsx4ERIEjEq88I2bnoau0M3qiv9UzYFuxOVHYkBTu04s68NmMQ8XGob
LKyb9oG6Sxt8weycZtapWOWREh06vgtvBCzapFF1 -----END CERTIFICATE-----

```

Submit the Certificate

Back

5. Stop and restart the ECOM service

```

# ps -eaf | grep ECOM
# kill -9 <ECOM PIDs>
# /opt/emc/ECIM/ECOM/bin/ECOM -d

```

11.1.10 ViPR SRM

Once certified by Citi (expected end of Q2 2015), ViPR SRM will be used for chargeback reporting for VMAX3. ViPR SRM will utilize SMI-S on the SE management server and Unisphere for VMAX to discover and report on the VMAX3 storage array. A functional ID will need to be created for ViPR SRM and assigned the symauth role of Auditor for SRM to gather the data.

11.1.11 HPNA Setup

Reference [GNAST Design and Standard for SAN Automated Configuration Assurance](#) for instructions for installing the HPNA client on the RHEL SE server. Once installed, a user account called 'hpnaro' is created on the OS. A symauth entry for the 'hpnaro' user will need to be created and assigned the symauth role of 'Auditor' for each Symmetrix array the SE server is managing. The hpnaro user will also need r-x permissions to the /var/symapi/db and /var/symapi/config directories on the SE server.

11.1.12 Group Name Services ("GNS") Configuration

CSBE recommends enabling GNS. GNS provides the ability to store device and composite group definitions in a shared repository located on each Symmetrix array, which would then automatically become visible to all hosts that are configured to run Solutions Enabler. This allows all GNS-enabled hosts to see the same group definitions across the Symmetrix environment, while sharing real-time updates to group definitions and configurations made by other hosts.

As of SE 6.0, the GNS daemon (storgnsd) can now create remote RDF mirrors for composite groups as well as device groups. For composite groups, only the standard devices are mirrored. In addition, this daemon synchronizes changes to SRDF/MSR consistency groups as it works in companion with the RDF daemon (storrdtd).

Also with SE 6.0 or above, you can now create and restore backup (shadow) files for GNS manually from the command line for attached storage array.

The present document only covers basic GNS configuration. However, the EMC Engineering White Paper **"Using SYMCLI and GNS to propagate Group Definition to Multiple Hosts"** gives an excellent summary of GNS operations, with examples, and is recommended reading when enabling GNS.

GNS User Authentication

GNS is intended to share group definitions across multiple users and hosts in a Symmetrix environment via the host installed GNS daemons. As such, limiting permission to the GNS daemon controls access to GNS functionality. However, this type of permission control is **NOT** required in the Citigroup environment because all users logging on to the SOE host have the same storage management responsibilities (and therefore need the same privileges).

Configuring GNS

If the daemon_options file does not already exist, copy the template README.daemon_options file to /var/symapi/config/daemon_options and edit. The following options should be set:

Option	Description
storgnsd:autorestart=enable	Make use of the watchdog mechanism to automatically restart the daemon if it crashes.
storgnsd:gns_device_poll_interval=15	The interval, in seconds, at which to poll locally-attached Symmetrix arrays for group changes.
storgnsd:gns_device_remote_poll_interval=60	The interval, in seconds, at which to poll

	remotely-attached Symmetrix arrays for group changes.
storgnsd:gns_rmtarr_update_interval=60	Interval in seconds to use to check for group changes that need to be propagated to remote Symmetrix arrays. The default is 60 seconds.
storgnsd:gns_symlsdiscover_interval=600	The interval, in seconds, at which to check for Symmetrix configuration changes.
storgnsd:gns_ppath_poll_interval=60	The interval, in seconds, at which to check whether PowerPath definitions need to be updated.
storgnsd:gns_shadow_file=enable	If set to enable, the daemon automatically maintains backup copies of GNS data for locally attached Symmetrix arrays. These backups, maintained on the local disk, are automatically restored to the GNS repository on an array if the daemon detects that a repository has been erased (for example, in the event of a power failure under Enginuity 5x67 or 5x66).

Using GNS

To instruct SYMCLI to use GNS instead of the local SYMAPI database, edit the **/var/symapi/config/options** file, and uncomment the following line:

```
SYMAPI_USE_GNS=ENABLE
```

GNS daemon storgnsd will need restarting.

Stop the daemon with the command

```
# stordaeon shutdown storgnsd
```

The GNS daemon is started automatically when a SYMCLI command that manipulates groups, such as symdg, is issued.

Alternatively, the daemon can be started manually via the stordaeon command line utility as follows:

```
# stordaeon start storgnsd
```

In addition, the daemon can be set to start automatically every time the local host is booted using the following command line:

```
# stordaeon install storgnsd -autostart
```

Pre-starting the daemon, either manually or via the automatic option is useful because the daemon may take a while to initially construct its cache — depending on the number of groups and Symmetrix arrays it has to load.

If the daemon crashes the Solutions Enabler watchdog daemon (storwatchd) will automatically try to restart the daemon. A combination of the watchdog mechanism and the auto-start option described above can be used to ensure that the daemon is always running.

Any existing group definitions should now be imported to GNS with the following commands:

```
# symdg activateall
# symcg activateall
```

To see if a group has been imported successfully, you can use the `symdg show DgName` or `symcg show CgName` commands. The output should have “Device Group in GNS : Yes” added to it.

GNS is now ready for use. Group definitions can now be created, edited or deleted as before with `symdg` and `symcg` SYMCLI command. All modifications will automatically be applied to the GNS database on the Symmetrix. The local database will not be used.

The GNS daemon (`storgnsd`) can now create remote RDF mirrors for composite groups as well as device groups. For composite groups, only the standard devices are mirrored. In addition, this daemon synchronizes changes to SRDF/MSR consistency groups as it works in companion with the RDF daemon (`storrdfd`). This is covered in a separate subsection.

Group Name Collisions

Each GNS daemon attempts to ensure that group names are unique (relative to a group type). It is possible, however, for duplicate named groups to occur. For example, it is possible for two different groups to be created with the same name on different Symmetrix arrays. A subsequent reconfiguration could then allow the GNS daemon to see both groups. If a GNS daemon detects duplicate group names, it returns a state flag informing clients of this fact and modifies the group names to be unique by attaching a numeric suffix. For example, duplicate device groups with the name `MyDG` could be modified in the following way:

```
MyDG#12345678
MyDG#87654321
```

While in this state (duplicate names), the only operations permitted on them are those that can be used to resolve the name conflict: group rename and delete (with the `-force` flag). For example:

```
# symdg delete MyDG#12345678 -force
# symdg rename MyDG#87654321 MyDG
```

GNS Backup and Restore

GNS group definitions are stored within the each individual Symmetrix array. A GNS daemon can be configured to automatically backup and restore GNS data from individual arrays. When backup is enabled, the daemon will write an image of the GNS group information to the host’s local disk each time it makes a change or it notices a change made from another locally attached host. These backups are maintained as separate per-Symmetrix files and named `SymmID.shadow` within the `/var/symapi/gns/` directory on the host machine⁵.

Automatic backup of the GNS database should be enabled by setting the `storgnsd:gns_shadow_file=enable` option in the `/var/symapi/config/daemon_options` file.

The creation and restoration from backup (shadow) files maybe done manually with the following commands:

⁵ Because the maintenance of numerous backup images can impact performance, it is recommended that you do not enable it on all GNS daemons. Typically, it is recommended that you configure a maximum of 2 GNS daemons to maintain backups for each Symmetrix array that may hold GNS data.

```
# /usr/symcli/daemons/storgnsd backup -sid <SID>
# /usr/symcli/daemons/storgnsd restore -sid <SID>
```

Note, automatic backups and manual backups write to the same shadow file. Consequently, if you want to preserve a manual backup made at a particular time, you should make a copy of the shadow file (/var/symapi/gns/SID.shadow) after the manual backup is completed.

GNS remote Mirroring

CSBE supports mirroring of RDF Group Definitions in a Remote Symmetrix Array.

With remote mirroring, GNS will automatically maintain a mirrored group definition on the remote Symmetrix array. This mirroring of groups' converts a local RDF1 type group to a remote RDF2 type group, and vice versa. Also local BCVs are converted to remote BCVs and vice versa.

When changes are made to the local group definition, GNS automatically propagates corresponding change to the remote group definition, keeping the two groups synchronized. The remote mirrored group has the same name as the local group and has a mirror image of its contents, bar those elements that are swapped such as the RDF type.

To enable mirroring of an RDF group definition, set the following option in the GNS daemon options file.

```
storgnsd:gns_remote_mirror=enable
```

By default, mirrored groups are read-only and cannot be modified or renamed. The default behaviour is recommended at this time.

Note that propagation of group definitions is controlled by the daemon_options file entry

```
storgnsd:gns_device_remote_poll_interval=60
```

Hence, it may take a minute or so before the definition is visible at the remote site.

GNS remote mirroring is not supported for device groups containing Concurrent or Cascaded SRDF devices. It must not be used on any SE server that is used to manage a 3-Site SRDF configuration. This applies to both the production and DR site.

If remote mirroring must be enable on an SE server that is also managing a 3-Site SRDF configuration, the remote Symmetrix (remote to the SE server) in the 3-site configuration MUST be excluded from GNS management using the gns_symavoid option for the storgnsd. For example, if the remote Symmetrix on a 3-site configuration are 1380 & 217:

```
# vi daemon_options
storgnsd:gns_symavoid = 000192601380, 000194900217
```

For further information on remote mirroring, refer to the EMC Engineering White Paper “**Using SYMCLI and GNS to propagate Group Definition to Multiple Hosts**”.

11.2 Unisphere for VMAX


Unisphere for VMAX (U4V) is a web browser based user interface designed to provide storage management for EMC Symmetrix products under a common framework and offers features such as provisioning, managing, monitoring, alerting and performance analysis. Unisphere for VMAX (U4V) is a web browser based user interface designed to provide storage management for EMC Symmetrix products under a common framework and offers features such as provisioning, managing, monitoring, alerting and performance analysis.

U4V provides a Performance Real Time view for examining high frequency data points, Historical View for all Symmetrix array Key Performance Indicators (KPI), and Diagnostic View for root cause analysis. U4V allows data to be collected and stored for analysis and reports to be produced on workload, resource usage trends for service level management, diagnostics and capacity planning.

The following tables list the required components and versions for the deployment of Unisphere for VMAX

Component Name	Capability Description
Unisphere for VMAX	V8.0.3 and above is the certified for deployment
Solutions Enabler	V8.0.3 and above is certified for deployment with U4V 8.0.3
Unisphere for VMAX server	Unisphere for VMAX software should only be installed on these O.S. types for this certification: Windows 2008R2/2012 (SOE Built)

11.2.1 Unisphere for VMAX Architecture



11.2.2 Unisphere for VMAX Components

11.2.2.1 U4V Server

Unisphere for VMAX does have some scalability limits due to:

- The potentially large amount of data that it needs to collect, store and manipulate.
- For each U4V instance, EMC only support visibility + data collection of a maximum of 256,000 devices, EMC provide a guideline limit of up to 20 Symmetrix arrays per U4V instance – with caveats on CPU, Memory and Disk.

The current hard limit is 256,000 devices per U4V server pair, and hence the maximum number of Symmetrix systems that can be managed may be less than 20 (or potentially more depending on device count per array).

The 256,000 device count per U4V server pair includes ANY array/devices the symapi server can see (locally or remotely) – **Not just the devices on the arrays registered for data collection.**

In larger estates it is highly likely the SYMAPI server will discover/have access to more arrays than a single U4V server will be configured to monitor – **in this scenario certain arrays may need to be excluded from U4V in order to ensure the setup does not exceed the products limitation and impact data collection and general usage** – please see the [related appendix](#) for further details on how to apply a filter for U4V.

The following table can be used to **estimate** how many U4V and SYMAPI servers will be needed. It assumes that each VMAX arrays no more than an average of 10000 volumes. Of course if the average device amount increases then the amount of arrays will decrease.

A minimum of 2 SE servers is assumed for all sites for CLI management.

Unisphere for VMAX server should be run on a VMware virtual machine with:

- Windows Server 2008 R2 Standard Edition (64-bit), 2012 or 2012R2
- Disk & Memory – The space required will depend on the amount of devices/arrays – EMC Advise the following scaling based on number of arrays/volume count:

Maximum Volume count	Maximum managed storage system count	Minimum Processor Specification	Minimum Available memory	Minimum available disk space
0-64K	10	2 Core 1.8GHz	12GB	100GB
64-128K	20	2 Core 1.8GHz	16GB	140GB
128-256K	20	2 Core 1.8GHz	20GB	180GB
Deployments of more than 256K devices will typically consist of multiple symapi and Unisphere for VMAX virtual machines, potentially with a mix of the above specifications depending on the amount devices.				
CTI should validate deployments larger than the above table with CATE and their local EMC account team.				
Memory: The above stated is the minimal dedicated amount of memory (not shared memory) required to run Unisphere for VMAX with performance option <u>this is in addition to OS requirement.</u>				
Disk: The above stated is the minimal dedicated disk recommended to run Unisphere for VMAX <u>this is in addition to OS requirement.</u>				

- Solutions Enabler will need to be installed as a prerequisite for Unisphere.

Due to the large memory requirements of Unisphere for VMAX, **it is mandatory that no other applications (over and above standard build items) be installed on the same server.**

11.2.2.2 SYMAPI Server

The Solutions Enabler configuration guide should be referenced for the setup and configuration of the physical SYMAPI server.

Unisphere for VMAX will interact with a SYMAPI server using the Solutions Enabler application that is set up in a Client /Server configuration. The SYMAPI server will have physical connectivity to the storage array via Fibre channel. The following restrictions apply to client/server SYMAPI connection:

- U4V server can only connect to a single SYMAPI server.
- U4V must only be used to collect performance data from Symmetrix that are local to the SYMAPI server. Data collection over RDF links to a remote Symmetrix is prohibited.

Typically existing Linux servers running Solutions Enabler (SE) should be used where possible to collect the STP data. In relation to U4V & SE a maximum of 256,000 devices can be collected by a single STP daemon i.e. a single SE server.

- An SE server can collect STP data on a defined subset of the arrays that it can see
- The primary or secondary SE server can manage (via CLI) all of the Symmetrix visible to both the primary or secondary SE server

As SE servers are deployed in redundant pairs. Each SE server in the pair could act as a separate SYMAPI server for multiple U4V instances.

11.2.2.3 Unisphere Client Requirements

Browser	Version	Memory
Internet Explorer	9.0 →11.0	600MB
Chrome	21.0.1180 or later	


11.2.3 Solutions Enabler Installation/Configuration on U4V Server

Solutions Enabler Installation on the Unisphere for VMAX Server is a pre-requisite for U4V. The provision of this package should be via TPM. It will also detect and upgrade the previous version automatically.


This is an **“interactive”** install package, so that regional deployment differences can be accommodated and hence screenshots are provided below.

TPM Package Name: CITI_WIN_SolutionsEnabler_8030_00


- Request the SA to push the package from TPM
 - Note: TMP will not install the program, but drop the binaries on the VM for the SA/Storage admin to carry out movement (if required) and installation.
- Run se8030-WINDOWS-x64.exe in the C:\server_packages\Solutions Enabler 8.0.3.0 directory and follow the installation instructions
 - Install to appropriate location (typically default)




- Select Typical Set up
- Select the SE Base Daemon. Other daemons are not needed if the U4V server is going to use a remote SYMAPI connection.



- Confirm default selection and proceed with the installation









11.2.4 Solutions Enabler Upgrade for U4V Server

Before upgrading, backup the <install dir>\Program Files\EMC\SYMAPI\config and C:\Program Files\EMC\SYMAPI\config\db directories.

- Run se8030-WINDOWS-x64.exe in the C:\server_packages\Solutions Enabler 8.0.3.0 directory
- Select the SE Base Daemon



11.2.5 Solutions Enabler SYMAPI Configuration for U4V

Unisphere for VMAX server and remote SYMAPI server should be configured to utilize SECURE SSL sessions.

From U4V server, edit the netcnfg file "<installpath>\emc\symapi\config\netcnfg" with the details of the SYMAPI server for example:

```
SYMAPI_SERVER - TCPIP symapi.domain.net 111.222.333.444 2707 SECURE
```

See Enabling the SYMAPI Server for instructions on setting up the SYMAPI server.

Restart SMAS service for the changes to take effect.


11.2.6 Unisphere for VMAX Installation

Installation of this package should be via TPM. It will also detect and upgrade the previous version automatically.


This is an **"interactive"** install package and hence screenshots are provided below.

TPM Package Name: CITI_WIN_U4V_8039_00


- Request the SA to push the package from TPM
- Run the installation package UNIVMAX_V8.0.3.9_WINDOWS_X86_64.exe in the C:\server_packages\Unisphere 8039 and follow the installation instructions
 - Select Next
 - Install to appropriate location (typically default)




- You should select the "Remote" connection type, providing the fully qualified name of the remote SYMAPI server that you will be using, leaving the net port as 2707 default. > Next



- Leave the checkbox unchecked for Certificate-based Client Authentication



- Run the installation, and select 'Done' on completion



11.2.7 Upgrading Unisphere for VMAX

Do the following before upgrading:

- If this is an upgrade, take a screen shot of the LDAP binding information. The upgrade should persist LDAP binding information but this will ensure that you have the information if for some reason the LDAP binding information is lost.





Authentication Authority States	
Local Directory	Enabled
Windows OS/AD	Enabled
LDAP-SSL	Enabled

LDAP Details	
Server	namdev.nsrootdev.net
Bind DN	CN=srv_emc_unisphere,OU=Svc_Accts,OU=StorageEng,DC=namdev,DC=nsrootd...
Port	3269
User Search Path	DC=nsrootdev,DC=net
User Object Class	user
User ID Attribute	sAMAccountName
Group Search Path	DC=nsrootdev,DC=net
Group Object Class	group
Group Name Attribute	cn
Group Member Attribute	member
Limit Authentication to Group(s)	per_csbe_univmax_authentication

- Take a backup of the Unisphere database and performance database within Unispehere.

To upgrade the Unisphere software:

- Request the SA to push the package from TPM
- Run the UNIVMAX_V8.0.3.9_WINDOWS_X86_64.exe in the C:\server_packages\Unisphere 8039 directory




11.2.8 Configuring Unisphere Security

The following table & Diagram summarizes the security configuration that has been adopted in Citi relating to Unisphere deployments.

Requirement	Implementation for Unisphere for VMAX
Centralized Authentication	User Authentication via LDAP using Active Directory as the source

Restricting activity when in Unisphere for VMAX	Authorization based on AD Group/User mapped to symauth roles
User access to U4V	Users will utilize desktop web browser for access (rather than direct RDP to a server)



11.2.9 Unisphere User Authentication

From the 7.6 release of Solutions Enabler CATE have provided instruction on using group based authentication based on Kerberos groups. However U4V cannot utilize these groups for authentication, and there have continued to be issues with U4V using the native AD authentication option.

Using AD based authentication U4V cannot:

- Look up and pass user group information to symauth.
- Use AD group based symauth rules – therefore CTI would still need to maintain individual rules for each user
- Look up and pass nested AD group information
- Enforce a group based restriction for user access to U4V

To work-around the restrictions found with U4V using native AD, U4V does allow authentication via LDAP where you specify the authentication source. This has successfully been tested utilizing the AD domain controllers as the authentication source.

Prerequisites:

- A Windows AD group will need to be set up to restrict access to U4V application to only the members of the AD groups that will be nested within it.
 - The Nested groups will be the ones that have been mapped to symauth roles ([Reference previous section on symauth group set up](#)).
- A Functional ID/service account will need to be set up to carry out the user look up for authentication queries controlling access to Unisphere for VMAX.


FID/Service account and active directory groups can be requested and membership edited via Citi Marketplace (search for Active Directory), alternatively an existing FID could be leveraged. e.g. DCNM/HPNA

- The default U4V 'smc' account has been given admin role on the arrays via symauth to allow initial set up.

Process:

To set up the Authentication:

- Access Unisphere for VMAX via URL <https://<UnisphereforVmax-server-name>:8443>
Login into U4V with the initial default 'smc' account.
- Select Home → 'Administration' → 'Security' → 'Authentication' which will display a summary of the current configuration.
- Select Edit and the Authentication Wizard will launch.
- Select LDAP-SSL



1. Click Next
2. LDAP-SSL configuration window will open Input details.

Configure Authentication

1 Authentication Authority

2 Configure LDAP-SSL

3 Summary

* LDAP Bind

* Server (IP or Hostname) * Port

* Bind DN

* Bind Password

* LDAP Users

* User Search Path

* User Object Class

* User ID Attribute

* LDAP Groups

* Group Search Path

* Group Object Class

* Group Name Attribute

* Group Member Attribute

SSL Certificate

Certificate File

<< Hide Advanced

☒ Limit authentication to members of LDAP group(s):

Group Name(s)

3. Click Next once complete

Please see the table below providing further reference of the input details.

Configuration Item	Description	Citi entry
Server	IP address or hostname of the server to use for authentication	AD Domain xxxxxx.nsroot.net Example: eur.nsroot.net
Port	Port number of the LDAP/S service (Global Catalogue)	3269
Administration Bind DN	Distinguished Name of the service account used to perform operations, such as searching users and groups	Request this information from AD operations if this is not known. CN="FID sAMAccountName",OU="XXX",OU="YYY",DC="ZZZ",DC=nsrootdev,DC=net Example: CN=srv_emc_unisphere,OU=Svc_Accts,OU=StorageEng,DC=namdev,DC=nsrootdev,DC=net OR sAMAccountName@domain Example srv_account@xxxxxx.nsroot.net

		srv_emc_unisphere,@namdev.nsrootdev.net
Bind Password	Password of the service account	Password that has been set for service account
User Search Path	Distinguished name of the node at which to begin user searches	DC=nsrootdev,DC=net
User Object Class	Object class identifying users in the LDAP hierarchy	User
User ID Attribute	Attribute identifying the user login ID within the user object	cn
Group Search Path	Distinguished name of the node at which to begin group searches	OU=XXXX,OU=YYYY, DC=nsrootdev,DC=net Example: OU=Svc_Accts,OU=StorageEng, DC=nsrootdev,DC=net
Group Object Path	Object class identifying groups in the LDAP hierarchy	Group
Group Name Attribute	Attribute identifying the group name	Cn
Group Member Attribute	Attribute indicating group membership for a user within the group object	Member
Limit Authentication to members of LDAP group	Specifies to limit authentication to only members of a specified group	This is the AD group created for containing the nested groups NOT the groups mapped to symauth rules e.g. per_univmax_authentication

If any settings are incorrect you will get an error banner, click close and you will return to the previous configuration page

If all settings are entered correctly a summary will be displayed

Configure Authentication

1 Authentication Authority

2 Configure LDAP-SSL

3 Summary

The following authentication authorities will be enabled:

- Local Directory
- Windows OS/AD
- LDAP-SSL

LDAP-SSL Configuration:

Server: namdev.nsrootdev.net Port: 3269

Bind DN:

CN=srv_emc_unisphere,OU=Svc_Accts,OU=StorageEng,DC=namdev,DC=nsrootdev,DC=net

User Search Path: DC=nsrootdev,DC=net

User Object Class: user

User ID Attribute: sAMAccountName

Group Search Path: DC=nsrootdev,DC=net

Group Object Class: Group

Group Name Attribute: cn

Group Member Attribute: member

LDAP-SSL Configuration (Advanced)


Limit authentication to members of LDAP group(s):

per_csbe_univmax_authentication

< Back Next > Finish Cancel

4. Click Finish

Confirmation box will appear – click close



The Authentication Summary will update once you apply the change:

LDAP Details	
Server	namdev.nsrootdev.net
Bind DN	CN=svr_emc_unisphere,OU=Svc_Accts,OU=StorageEng,DC=namdev,DC=nsrootd...
Port	3269
User Search Path	DC=nsrootdev,DC=net
User Object Class	user
User ID Attribute	sAMAccountName
Group Search Path	DC=nsrootdev,DC=net
Group Object Class	Group
Group Name Attribute	cn
Group Member Attribute	member
Limit Authentication to Group(s)	per_csbe_univmax_authentication

For initial setup, Local Directory should be left enabled as a temporary measure until all setup is complete. Remove “Local directory” only when there is a user set as “Administrator” or “SecurityAdmin” in symauth **and** the user has confirmed authentication access via LDAP.

If you disable Local access prematurely, you may lock yourself out of Unisphere for VMAX. You have been warned.

11.2.10 Unisphere User Authorization

The use of symauth feature as a user-based access control mechanism to a Symmetrix is **mandatory** when using Unisphere for VMAX. See Symauth Role-Based User Access Control for more details on symauth.

Symauth meets Citi audit requirements since all user changes are logged with fmt=symaud and reviewed accordingly, if audit logging is correctly set up as per the above Solutions Guide.

'PerfMonitor' is a role that limits access to managing performance related activities on Symmetrix systems specifically via Unisphere for VMAX. Where applicable the role should be assigned in order to limit certain users/AD groups to performance monitoring related and read-only operations, but not operational configuration via U4V.

The permissions Unisphere for VMAX users gain will be based on the SYMAUTH roles they or their AD groups are mapped to on each array.

Most SAN administration team members will already be assigned to at least the 'StorageAdmin' role in order to carry out day to day CLI manipulation.

Note: The 'Administrator' role should not be assigned for providing rights for standard storage administration activities. It is possible to assign more than one role to a user (directly or indirectly via AD Group membership) in U4V and the combined privilege will be granted to the user.

The symauth provides user-based authorization that assigns a user or group to a role to limit the management operations they can perform on a Symmetrix array.

Solutions Enabler represents user or group identity as a string that comprises of the User/Group Name and the authentication source. Although Symauth entries for U4V will be based on active directory groups, the authentication source however will **NOT** be specified as 'D:' (Windows Domain).

The Authentication source will need to be specified as 'L:' (LDAP)

New active directory groups can be requested and membership edited via Citi Marketplace (search for Active Directory) or you can utilize existing AD groups if they contain the appropriate users.

Once the AD groups have been set up (or you have identified the existing groups you wish to use) appropriate symauth rules will need to be applied, for example:

```
# symauth -sid 1234 -f filename commit
```

Where the contents of example file is

```
Add group L:namdev.nsrootdev.net\per_csbe_unispherevmax_admins to role
StorageAdmin;
Add group L:namdev.nsrootdev.net\per_csbe_unispherevmax_users to role
PerfMonitor;
```

Therefore In this **example**:

- Authentication source **LDAP**: (Using Active Directory Domain controllers)
namdev.nsrootdev.net ** (Windows AD group) 'per_csbe_unispherevmax_admin**' is mapped to the **StorageAdmin role**
- Authentication source **LDAP**: (Using Active Directory Domain controllers)
namdev.nsrootdev.net ** (Windows AD group) 'per_csbe_unispherevmax_users**' is mapped

to the **PerfMonitor** role

Once applied you can check/confirm the entry - Symauth list below has also been edited to just show these entries relevant to U4V.

```
# symauth -sid 4261 list -users -v
Role name           : StorageAdmin
Group name          :
L:namdev.nsrootdev.net\per_csbe_unispherevmax_admins

Role name           : PerfMonitor
Group name          :
L:namdev.nsrootdev.net\per_csbe_unispherevmax_users
```

If changes are made to a user's symauth role outside of U4V (e.g. a change to the membership of an AD group that is used to define which users have a specific role), then the change will not immediately be visible to U4V.

This is because it can take some time for the authorization rule to update on SFS and be sent back to SE and U4V. Authorization rules are only *immediately* recognized within the process in which they were created. So if you were to create a rule through U4V, then that rule would be immediately enforced for that instance. It will take a few minutes to propagate to other instances of U4V/SE.

The bottom line is that you may have to wait several minutes (depending on the size of your environment, although there is no way to determine exact time) for the change to be visible to and used by U4V.

11.2.11 User Access to Unisphere

U4V is not limited to only being accessible via the window management workstation it resides on. Users do not need to RDP to a specific server before launching a browser on the server and logging on or have need to have access to the server.

Users will continue to access U4V from their desktop web browser after set up is complete
Access Unisphere for VMAX via URL


<https://<UnisphereforVmax-server-name>:8443>

1. Enter **JUST** your SOE ID & Password
2. Click Login


11.2.12 Unisphere Login Banner Configuration

The login message should be set accordingly:

You are authorized to use this System for approved business purposes only. Use for any other purpose is prohibited. Authorized and unauthorized activities may be monitored. All Data residing on this system remains the property of Citi.



This is configured Home → Administration → Preferences → Client session log in message



11.2.13 Filtering Array Discovery in Unisphere

As highlighted in earlier sections, U4V scalability is limited to a maximum of 256,000 devices – Within larger estates a SYMAPI server may well see 40+ arrays.

By default U4V will discover **whatever** the SYMAPI server has discovered, which includes local and remote arrays, and is not limited to only arrays that are registered in U4V for collection.


In previous U4V versions, within some of the larger regional data centers, there have been incidents where U4V has not collected data for the arrays registered, with the cause being attributed to the U4V instances having more than the recommended devices allowed.

– In order to stay within the device limit U4V can be configured to exclude specified arrays from its discovery. The update requires making a change to an xml file on the U4V server, and example is shown below

It is recommended to limit U4V to only discover/monitor/view


- 1. Arrays Registered for Data collection**
- 2. Local Arrays ONLY**
 - **Remote arrays should be excluded to eliminate un-needed traffic on RDF links**

All Arrays visible on the SYMAPI server are visible on U4V instance



- Remove the deploy# folders in **<Install Dir>\EMC\SMAS\jboss\domain**. However, DO NOT delete the main “deploy” folder. When Unisphere starts, it will recreate the deploy0, deploy1, ... folders based on the “deploy” folder.
- Locate and open the ejb-jar.xml file.
Note: you will require appropriate rights to carry out changes to this file.

<Install Dir>\EMC\SMAS\jboss\domain\deploy\domain-symm.ear\domain-symm-ejb.jar\META-INF/




- In the ejb-jar.xml file locate the section as shown below

```
<!-- Comma-separated symm-id collection -->
<env-entry>
  <env-entry-name>filterSymms</env-entry-name>
  <env-entry-type>java.lang.String</env-entry-type>
  <env-entry-value></env-entry-value>
</env-entry>
```

- Update the file to include all SYM IDs to be filtered OUT, depicted in the example below. The example depicts one SYM IDs that will be excluded, multiple entries should be comma separated.

```
<!-- Comma-separated symm-id collection -->
<env-entry>
  <env-entry-name>filterSymms</env-entry-name>
  <env-entry-type>java.lang.String</env-entry-type>
  <env-entry-value>000195700524</env-entry-value>
</env-entry>
```

- Save the changes, Close the file
- Stop & start the SMAS services
Specified array should then no longer be visible to U4V.



An RPQ was submitted by Peter Kemp (EMEA TC) to support the use of this exclusion workaround – This was rejected however EMC have advised they will support Citi – If any support issues are raised by the use of this exclusion, please inform the author of this document.

11.2.14 McAfee Virus Scan – Addition of processes to Low Risk Process Policy

EMC recommend the exclusion of postgres and stp directories and files from antivirus scan **to avoid performance database corruption.**

As confirmed with EMC (SR 68912574) – the stp directory exclusion is only applicable when Unisphere is deployed in Local configuration hence why it is not included.

Host & Security Engineering (HSE) team advise that Citi's best practice is to avoid scanning on the processes rather than excluding files from virus scanning.

Therefore 'pg_ctl.exe' and 'postgres.exe' processes needs to be added into the 'Low Risk Processes' (LRP) policy whereby the antivirus will ignore files accessed by the 'postgres.exe' process, which will avoid virus scanning on the database files.

In order to request the processes are moved in to the low risk process policy the CTI teams carrying out the production deployments/upgrade will need to raise a ResolveIT Ticket to the HSE team.


- **From the CATE website download a copy of the Security Deviation Request Form zip file – Which includes the SDRF, SMT and TISO email approval**
- **Raise a ResolveIT Ticket to: CTI GL AD APPLICATIONS-Information Assurance Ops (IAO) detailing the requirement to move the pg_ctl.exe and 'postgres.exe' processes in to the low risk process policy for the specific deployments.**
 - o **Reference/Attach the approved SDRF form**

Please note the following advice from the BISO:

"Deviations tied to a CTC version ID are good for the life of the product as long as the CTC Product ID is 'Recommended or Restricted', and the CTC Version ID is 'Aging or Current'. Otherwise, deviations are good for one year and are the responsibility of the BISO and Application owner to file for an extension prior to expiration. "

11.2.15 Web Server Certificates

When you first try to connect to the Unisphere for VMAX via web browser the error message for security certificate as shown below:



As a workaround you can click 'Continue to this Website' to access Unisphere.

During CIVA testing for U4V 1.5 it was advised it is required to apply a valid CA-signed digital certificate instead of using the self-signed certificate from U4V.

High level steps on how to replace the certificate for U4V are:

1. Retrieve Keystore password
2. Generate CSR
3. Request signed certificate from CA
4. Retrieve Server and CA certificate
5. Import CA certificate chain into the U4V server keystore
6. Import CA-signed certificate into the U4V server keystore
7. Verify the certificate chain for the keystore

Step 1: Retrieve Keystore password

- Open `<install_directory>\emc\smas\jboss\domain\configuration\domain.xml.init`
- Search for 'password='
- The keystore password is displayed inside quotation marks in clear text

For example

```
<ssl key-alias="tomcat" password="ic1-unisp-vm01.namdev.nsrootdev.net@Keystore-2" certificate-key-
```

Step2:Generate CSR

```
<install_directory>\emc\smas\jre\bin\keytool -certreq -alias tomcat -keyalg RSA -file tomcatcert.csr -keystore keystore
```

You will be prompted for the keystore password, once input a csr file will be generated.

Step 3: Request signed certificate from CA

Submit request from Citi Market Place (CMP) by searching for 'Venafi MS-PKI On Demand Device Certificate - 163513 - Create/Revoke Certificate'

https://cmp.nj.ssm.com/marketplace/control/product/~productId=47119_29952_GLOBAL

Complete the questionnaire according to your local requirements, adding the CSR generated in the previous step.

Also refer to: [CMP instructions for SSL Cert request](#) & [MSPKI - SSL Guide](#)

NOTE: If required, certificate revocation can also be requested via marketplace –against the same market place product

Step 4: Retrieve CA signed certificate (Requester will be sent this via email)

CA signed certificate will be sent directly to the requester

From the email - Please copy everything between the ---- delimiters, including the delimiter lines themselves and save to a file for use.

```
-----BEGIN CERTIFICATE-----
and include everything through
-----END CERTIFICATE -----
```

And to obtain the environment appropriate Device, Policy and Root CA chains, please visit the CATE MS PKI Sharepoint site (linked within the email) and select from the right sidebar. It is important to obtain all intermediate and root CA certificates for the signed certificate.

Step 5: Import CA certificate chain to the keystore

It is required to import all certificates in the certificate chain starting with the root certificate, intermediate certificates (in descending order from the root certificate), and then the CA signed certificate last.

- Copy the root, intermediate, and the signed certificate to the
`<InstallDirectory>\SMAS\jboss\standalone\configuration\` directory.
- Go to the following directory from command prompt:

```
cd <InstallDirectory>\SMAS\jboss\standalone\configuration\
```

- Run the following command to import the root CA certificate in trusted chain:

```
<install_directory>\emc\smas\jre\bin\keytool -import -alias
root -trustcacerts -file <root certificate filename> -keystore
keystore
```

Example:

```
# keytool -import -alias root -trustcacerts -file
CitiInternalRootCAUntrusted_b64.cer -keystore keystore
```

- When prompted, enter the keystore password previously retrieved (step 1).
- Import the intermediate certificates in the chain in descending order from the root certificate:

Example:

```
# keytool -import -alias intermiediate1 -trustcacerts -file
CitiInternalPolicyCAUntrusted_b64.cer -keystore keystore

# keytool -import -alias intermiediate2 -trustcacerts -file
CitiInternalDeviceCAUntrusted_b64.cer -keystore keystore
```

Step 6: Import the CA-signed certificate into the keystore

- From the U4V server, stop the SMAS service

```
C:\>net stop "EMC Symmetrix Management Application Server"
```

- Go to the following directory from command prompt:

```
cd <InstallDirectory>\SMAS\jboss\standalone\configuration\
```


- Run the following command on the command prompt:

```
<install_directory>\jre\bin\keytool -import -alias tomcat -
trustcacerts -file <signed CA certificate filename> -keystore
keystore
```

Example:

```
# keytool -import -alias tomcat -trustcacerts -file icl-unisp-
vm01.cer -keystore keystore
```

- When prompted, enter the keystore password previously created (step1).
- Start the SMAS service
- On next attempted launch, you can install the Citi/CA signed certificate



Step 7: Verify the certificate chain for the keystore

Issue the following command to list the certificates for the keystore:

```
<install_directory>\jre\bin\keytool -list -keystore
keystore
```

Sample Output:

```
C:\PROGRA~1\EMC\SMAS\jboss\standalone\configuration>C:\PROGRA~1\EMC\SMAS\jre\bin
\keytool -list -keystore keystore
Enter keystore password:

Keystore type: JKS
Keystore provider: SUN

Your keystore contains 5 entries

root, Apr 22, 2015, trustedCertEntry,
Certificate fingerprint (SHA1): 5B:74:CA:A1:68:FE:9B:D5:4C:C8:DF:39:AF:50:1D:C9:
E6:8C:3E:0A
intermediate2, Apr 22, 2015, trustedCertEntry,
Certificate fingerprint (SHA1): 20:B6:21:27:14:E5:B9:42:F3:3A:A0:74:0E:A7:B3:F1:
4E:6B:8B:6F
intermediate1, Apr 22, 2015, trustedCertEntry,
Certificate fingerprint (SHA1): F8:C0:D3:76:4C:55:10:C3:14:8B:0A:0F:B7:05:F6:BC:
01:BF:4D:B2
tomcat, Apr 22, 2015, PrivateKeyEntry,
Certificate fingerprint (SHA1): 0A:DE:81:DF:E0:E2:DC:82:FF:35:4C:87:C0:61:41:FA:
A0:BF:DC:43
univmaxrestforwardingclient, Apr 22, 2015, PrivateKeyEntry,
Certificate fingerprint (SHA1): DB:89:D5:FA:29:31:64:30:9A:89:66:69:B4:18:DD:2E:
DE:CE:EC:B7

C:\PROGRA~1\EMC\SMAS\jboss\standalone\configuration>
```

Notice that there is also a private key entry named `univmaxrestforwardingclient` in the keystore. If that private key entry is not in the keystore, issue the following command to create it in the keystore:

```
<install_directory>\jre\bin\service.bat generatekeystore
<keystore password>
```

11.2.16 Registering arrays for Data Collection

The Performance component of Unisphere for VMAX only collects performance data after you register your Symmetrix systems. Note that **ONLY local Symmetrix systems** can be registered for Performance data collection.

Ensure that you have completed configuration of [symauth](#) to give the right privileges to set up data collection.


Login into Unisphere for VMAX using an AD account/Group with symauth roles 'PerfMonitor' or 'StorageAdmin' assigned on the Symmetrix array from which data is to be collected.

Once logged in, follow the below to register a Symmetrix system for performance data collection:

- **Select the Symmetrix system**

- Click 'Performance' → 'Settings' → 'Systems Registrations' to open the Systems Registrations view.
- Click Register to open the Register dialog box. The Symmetrix ID automatically populates the Select Symmetrix list.
- Select both the type of performance data to collect, 'Real Time' and 'Diagnostic'.
- Set Diagnostic Interval to 10 Minutes (Minimum)
- Click OK.

View the system registrations by selecting 'Performance' → 'Settings' → 'System Registrations'. Click on the 'View Details' button from bottom left panel to view further details or make changes (e.g. enable/disable Real time collections) to the system registrations on the Symmetrix array.



Properties	
Symmetrix ID	000196700973
Real Time	☒
Root Cause Analysis	☒
Collection Recovery (Hours)	24
Root Cause Analysis Interval (Minutes)	5
Server IP	10.40.52.45
SE Server IP	10.40.42.90

11.2.17 Unisphere Server backup

With Unisphere for VMAX 8.0.3 and above you are able to manually backup **all** the data currently on the database server, including Database Storage Analyzer, Workload Planner, performance, and infrastructure data. Database backups will enable you to recover from system crashes.

Noting that you can only restore the database to the same Unisphere for VMAX version and same operating system. For example, a V8.0.3 database on Windows, can only be restored to a V8.0.3 on Windows

At the point of testing the 8.0.3 GUI does **not** include a restore button! This was flagged to EMC who acknowledged this was missing – but have provided instructions for restore if this function is used (CTI should contact CATE If this function is needed) – this section will be updated once the GUI has the required function.

11.2.18 Unisphere Performance Database Backup

Unisphere for VMAX performance database should be backed up regularly to protect against accidental corruption. On Symmetrix arrays with many devices, the disk space required eventually may be high.

To backup U4V performance database,

- Select Performance → Settings → Databases to open the Databases view.

- Select a database from the list and click 'Backup' to open the Database Backup wizard.
- Select the Symmetrix system.
- Enter a backup File Name. The file name will be started with Symmetrix system ID and timestamp of backup taken follow by the file name as defined here. For example:
 - 000192604261_201402020543_test_backup.dat

- Optional: Select any of the following Data to Save:
 - *Historical*⁶
 - **Last day of Diagnostics*⁷
 - **Named Real Time Traces*
- Set the "Execution Time":
 - If 'Now' – Click Finish.
 - If selecting 'Schedule' - Click 'Next'. Select the 'Start Date and time'. Select 'Yes' to configure for recurring database backup and select the day to run on weekly basis.

- Click Finish.

The backup will be written to the directory <InstallDirectory>\Program Files\EMC\SMAS\backup\SPA

11.2.19 Configure Idle User Timeout Session

⁶ "Historical" is always selected and cannot be unselected

⁷ The two * options are not recommended for recurring backups

EMC Unisphere for VMAX Timeout

Due to inactivity you have been automatically logged out of Unisphere for VMAX.
To return to Unisphere for VMAX please use your web browsers Back button.

By default, the Unisphere for VMAX initiate automatic logout for user session after 480 minutes (8 hours).

Edit the **web.xml** file - Changing the setting from so users will be automatically logout after inactive session for more than 30 minutes.

<i>Web.xml</i>	<i>located</i>	<i>in</i>
<code><installdrive>\Program Files\EMC\SMAS\jboss\domain\deploy0\smc0.ear\smc.war\WEB-INF</code>		

Edit the parameter from 480 to 30

<pre><param-name>sageIdleUserTimeoutInMinutes</param-name> <param-value>30</param-value></pre>
--

Restart SMAS service from the U4V server to take effective for this setting.

11.2.20 Remove authorization via local directory

Once setup has been completed and tested and you are 100% satisfied it works, local directory account assess must be disabled.

Users with appropriate Administrator or SecurityAdmin rights will need to login into U4V.
Select 'Administration' from the right bottom common tasks panel and select 'Authentication'
Unselect 'Local Directory' from the list



11.2.21 Unisphere Alerting

Unisphere for VMAX can be configured to monitor Symmetrix systems for specific events, error conditions and breaches of thresholds. When U4V has been enabled to capture the specified event, error or threshold breach it will **display** a reactive 'alert'.

An alert in U4V can be considered to be a visible representation of an entry in the audit log or Symevent file – the U4V Alert is **not** the same as an alert from an event monitoring system.

11.2.21.1 Default Alerts

As part of the Solutions Enabler configuration a number of essential events/alerts are already captured and forwarded to Citi Event Monitoring systems (see [Configuring Symmetrix Alerting](#)) – Some of the same events included as part of the recommended solutions enabler configuration are also enabled as default in U4V.

These pertain to various Pool & FAST Policy Utilization levels, providing Normal, Warning, Critical & Fatal alerts based on % Utilization.

The default threshold alerts provided in U4V cannot be disabled or edited. When a condition occurs triggering the alert, they are logged/displayed in the System Alerts Pane.

The screenshot shows the EMC Unisphere for VMAX V8.0.1.7 interface. The top navigation bar includes icons for Storage, Hosts, Data Protection, Performance, Databases, System, and Support. The breadcrumb trail is 000192604261 > System > Alerts. A table displays a single alert with the following details:

State	Severity	Type	Object	Description	Created	Acknowledged
NEW	✖	ARRAY	small_pool	Data Pool utilization is now 80 percent. - Object is: 000192604261:small_pool	Wed 02/04/2015 01:36:52 PM GMT+0000	

Below the table are three buttons: View Details, Acknowledge, and Delete.

They can be acknowledged and/or deleted which will reduce/clear the alert count on the Home screen adjacent to the array.

11.2.21.2 U4V System Runtime Checks/Alerts

Unisphere for VMAX performs the following runtime checks. If a check returns a result outside the expected range, Unisphere for VMAX will raise the appropriate alert.

Checks are run on 10 minute intervals and alerts are raised on 24 hour intervals from the time the server was last started.

The checks and alerting is provided by default and cannot be disabled or edited.

Runtime check	Number of volumes	Threshold	Alert
Total memory on the Unisphere server	0 - 64,000	12 GB	System memory <# GB> is below the minimum requirement of <# GB>
	64,000 - 128,000	16 GB	
	128,000 - 256,000	20 GB	
Free disk space on the Unisphere installed directory	0 - 64,000	100 GB	Free disk space <# GB> is below the minimum requirement of <# GB>
	64,000 - 128,000	140 GB	
	128,000 - 256,000	180 GB	
Number of managed storage systems	0 - 64,000	10	Number of managed arrays <#> is over the maximum supported number of #
	64,000 - 128,000	20	
	128,000 - 256,000	20	
Number of managed volumes	256,000		Number of managed volumes <#> is over the maximum supported number of <#>. Note that Solutions Enabler may indicate a slightly different number of volumes than indicated in this alert.
Number of gatekeepers	6		Number of gatekeepers <#> on Symmetrix (<i>SymmID</i>) is below the minimum requirement of 6.

11.2.21.3 Target Enginuity level

By default U4V will log an alert if the symmetrix are not at EMC's target Enginuity level. It is unlikely Citi would be at the same Enginuity levels as the latest vendor release – to avoid unneeded alert entries we would recommend this alert is disabled.

Home → Administration → Alert Settings → Alert Policies Select the Target Enginuity Warning listed against each array discovered, click Disable

The screenshot shows the EMC Unisphere for VMAX V8.0.2.6 interface. The top navigation bar includes 'All Symmetrix', 'Home', 'Performance', 'Search', and 'Support'. The breadcrumb trail is 'All Symmetrix > Home > Administration > Alert Settings > Alert Policies'. The main content area displays a table of alert policies.


Symmetrix ID	Type	Policy Name	State	Notification
000192604261	Array	SRDF/A Session recovered fro...	disabled	
000195700524	Array	SRDF/A Session recovered fro...	disabled	
000196700973	Array	Target Enginuity Warning	disabled	
000196700974	Array	Target Enginuity Warning	disabled	
000192604261	Array	Target Enginuity Warning	disabled	
000195700524	Array	Target Enginuity Warning	disabled	
000192604261	Array	Thin Device Allocation	disabled	

11.2.21.4 Unisphere Alert Policy Configuration

Unisphere alerts can be setup by enabling and enabling pre-defined alert policies. The policies can be enabled/disabled in Unisphere under the Alert Settings page in Unisphere:

Symmetrix ID	1 ▲ Type	2 ▲ Policy Name	3 ▲ State	Notification
000196700973	Array	Array Component Events	disabled	
000196700973	Array	Array Events	disabled	
000196700973	Array	CG Tripped	disabled	
000196700973	Array	DB Checksum Triggered	disabled	
000196700973	Array	Deferred Service Threshold Alert	disabled	
000196700973	Array	Device Config Change	disabled	
000196700973	Array	Device Pool Config Change	disabled	
000196700973	Array	Device Pool Status	disabled	
000196700973	Array	Device Status	disabled	
000196700973	Array	Director Status	disabled	
000196700973	Array	Disk Status	disabled	
000196700973	Array	Environmental Alert	disabled	
000196700973	Array	Event Lost Alert	disabled	
000196700973	Array	Event Overflow Alert	disabled	
000196700973	Array	External disk group fully allocated	enabled	
000196700973	Array	FAST Controller switched state	disabled	
000196700973	Array	FAST FTS Performance	enabled	
000196700973	Array	GK Timeout	disabled	
000196700973	Array	GK Utilization	disabled	
000196700973	Array	Hotspare Invoked	disabled	
000196700973	Array	Migration Complete Alert	disabled	
000196700973	Array	Optimizer switched mode	disabled	
000196700973	Array	Port Link Status	disabled	
000196700973	Array	Port Status	disabled	
000196700973	Array	RVA Spare Coverage	disabled	
000196700973	Array	SP Alerts	disabled	
000196700973	Array	SRDF Alerts	disabled	
000196700973	Array	SRDF Job Flow Control Change	disabled	
000196700973	Array	SRDF Link Status	disabled	
000196700973	Array	SRDF/A No Cycle Switch Alert	disabled	
000196700973	Array	SRDF/A Session	disabled	
000196700973	Array	SRDF/A Session dropped, transi...	disabled	
000196700973	Array	SRDF/A Session entering trans...	disabled	
000196700973	Array	SRDF/A Session recovered from...	disabled	
000196700973	Array	Target Engineity Warning	enabled	
000196700973	Array	Thin Device Allocation	disabled	
000196700973	Array	Thin Device Usage	disabled	
000196700973	Array	Thin Pool Rebalancing Complete...	disabled	
000196700973	Array	User approval required for Opti...	disabled	
000196700974	Array	Array Component Events	disabled	
000196700974	Array	Array Events	disabled	
000196700974	Array	CG Tripped	disabled	
000196700974	Array	DB Checksum Triggered	disabled	

A notification method can also be configured for each alert policy to define what mechanism to use to send out the alerts. The alerts can be sent out via email, snmp, and/or syslog:



The alert notification settings need to be configured for the selected alert method. See Unisphere Alert Notification for more details on configuring notifications.

11.2.22 Unisphere Performance Thresholds and Alerts

U4V provides the ability to create performance thresholds. There are many default system values for these thresholds which can be edited or new ones can be created.

The threshold values can be viewed from the charts in the Analysis or Diagnostic view.

And if enabled, U4V Alerts can be generated against threshold breaches

To view, edit or create new performance thresholds and alerts select Home → Administration → Alert Settings → Performance Thresholds and Alerts to view and manage thresholds and alerts.

The screenshot shows the 'Performance Thresholds and Alerts' page in the Unisphere web interface. The top navigation bar includes 'All Symmetrix', 'Home', 'Performance', 'Search', and 'Support'. The breadcrumb trail is 'All Symmetrix > Home > Administration > Alert Settings > Performance Thresholds and Alerts'. On the left, there's a sidebar with a list of Symmetrix IDs (000192604261, 000195700524, 000196700973) and a 'Category' dropdown set to 'System'. Below the sidebar, there are expandable sections for 'Storage', 'Data Protection', 'FAST', and 'Hosts'. The main area displays a table of performance metrics with columns for 'Metric', 'Custom', 'First Thresh...', and 'Second Thr...'. The table lists various metrics like '% Cache WP', '% Hit', '% Reads', '% Writes', 'Allocated Capacity', 'BE IOs/sec', 'Cache Balance', 'DA Balance', 'DX Balance', 'EFD Balance', 'FC Balance', 'FE Balance', 'Host IOs/sec', 'Host MBs/sec', 'Number of R1 Cache Slots in Use', 'Number of R2 Cache Slots in Use', 'RDF Balance', 'RDFA WP Count', 'Read RT (ms)', 'SATA Balance', 'System Max WP Limit', 'System WP Count', and 'System WP Events/sec'. At the bottom, there are buttons for 'Create', 'Edit', 'Delete', and 'Notifications'.

Click the 'Create' button from bottom left to create performance thresholds and alerts. Select the Symm ID, Category, Instance, and Metric to set the thresholds on, and enter the value for first/second threshold. Alerts can be enabled based on severity and occurrence.

Create Performance Threshold and Alert

Symmetrix ID: 000192604261

Category: FE Director

Instance: All

Metric: % Busy

First Threshold: 50

Second Threshold: 70

☒ Enable Alert


First Threshold Severity: Information

Second Threshold Severity: Warning

Occurrence: 3 / 5 samples

OK Cancel Help

Thresholds can be seen from the chart in View too. Click on the menu bar on the chart in select 'Display Thresholds'. Threshold values will then be displayed in the chart.



11.2.22.1 Recommended Thresholds

The following metrics were formulated with EMC performance engineering and are recommended as the basic performance related threshold alerts that should be set. Any defaults thresholds can be edited.

	UNISPHERE PERSPECTIVE	ALERT	RECOMMENDATIONS	FROM	PERFORMANCE	
	Threshold 1	Data Points	Per Interval	Threshold 2	Data Points	Per Interval
SWP	50	2	6	80	1	1
%Busy DA	70	3	6	80	3	6
%Busy FA	70	3	6	90	3	6
%Busy Port 1	80	3	6	95	3	6
%Busy Port 0	80	3	6	95	3	6
Warning			Critical			

11.2.23 Unisphere Alert Notification

As previously highlighted the alerts that are captured and displayed in Unisphere for VMAX are not the same as an alert in an event monitoring system.

However the 'Notification' feature in U4V can provide Email, Syslog and SNMP Notifications against the alerts U4V receives.

Citi's strategic target is to move to a single platform for monitor, alerting and notification – hence why these are currently included as optional.


11.2.23.1 Optional Email Notification

It is acknowledged and understood that the receipt of U4V notifications via email would not likely be a welcome addition to the volume of emails the operations teams already receive.

We have included it as an option for the case that certain operations teams or even individuals did wish to receive set up and receive U4V notifications via email (even on a short term basis for a specific issue), as a supplementary addition to any existing alerting. Or just to be able to receive emailed reports.

The appropriate email DL needs to be set up with an external address – e.g. Unisphere-Alerts@citi.com this can be requested via CMP.

Once the DL has been set up you will need to configure the email server in Administration → Alert Setting → Notifications → Config Email Server
Enter the DL details in the Sender field, and enter appropriate mail-hub in the IP address/Host field, leave the port as 25 → Click Ok



Configure Email

From

Sender Email

Unisphere-Email@citi.com

IP Address/Host Server Port

mailhub-na.wlb.nam.nsroot.net 25

To


Recipient Email

cti-san-admin@citi.com

 Changing will override individual email addresses.

Test OK Close Help

Once configured you are able to select the level of Alerts that you wish to receive an email notification against.



Notifications

TYPES

☒ Syslog Enable

☒ Email

☒ SNMP Enable

Alerts Reports

Symmetrix	System Level	Performance Level	System Email	Performance Email
000192604261	   	  	sanadministrator@citi	sanadministrator@citi
000196700974	   	  	sanadministrator@citi	sanadministrator@citi
000195700524	   	  	sanadministrator@citi	sanadministrator@citi


11.2.24 Unisphere Reports

11.2.24.1 Dashboard Reports

The easiest way to set up a report is to use the report buttons at the bottom of the Performance → Monitor pane.

Once the dashboard or template is displayed, select either 'Save As, giving you the option to save the displayed data as a PDF to a user specified location or 'Schedule the report' up to 7 days ahead and set to reoccur, save locally and if configured email.

Note scheduled/reoccurring reports will be saved to Install Drive:\Program Files\EMC\SMAS\report\SPA\querydata



Dashboard Report Example

Dashboard Report Example

Report Name	Storage Group			
Name	Utilization			
Array ID	000192604261	Start	Thu Feb 5 2015 12:32:08	
		End	Thu Feb 5 2015 16:32:08	
Utilization				
Table Summary				
Category	% Busy	Utilization	Port Utilization	Queue Utilization
FE Directors (2)	0.4			
BE Directors (DA) (16)	13.5			
Cache Partitions (1)	0.3			
FC Disks (96)	9.6			
FE Directors				
ID	% Busy	Utilization	Port Utilization	Queue Utilization
FA-7F	0.4			
FA-7F:0	0.0			
FA-10F	0.4			
FA-10F:0	0.0			

11.2.24.2 Customized reports

Alternatively customized reports can be created and provided as CSV, XML, BTP
Go to Performance' → 'Settings' → 'Reports' to create report.

Add required queries to the report – the wizard will prompt for Categories, Format, and then you select the metrics required

Created/Scheduled reports will be visible in the summary

All Symmetrix > Performance > Reports						
Name	Symmetrix ID	Format	Recurring	Run Date	Email Enabled	Description
Dashboard Report	Multiple	PDF	●	Fri Feb 6 2015 04:15:00 PM	●	
U4V test	000192604261	CSV	●	Thu Feb 5 2015 04:48:00 PM	●	
Test	000192604261	CSV	●	Thu Feb 5 2015 08:00:00 PM		
Create Copy Edit Schedule Cancel Schedule Run Now Delete Navigate to Notifications						

- Reports ran on demand will be saved to user specified location.
- Scheduled reports will save to:
Scheduled XML, CSV, and BTP reports are saved to this location:
install_dir\SMAS\report\SPA\querydata
- Scheduled UPV reports are saved to this location:
SMAS\backup\SPA

Example CSV

SYMM_I	DATE	TIME(m)	INTERV	CATEG	VALUE	INSTAN	% Busy
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-7E	5.49
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-8E	0.43
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-9E	3.72
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-10E	5.26
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-7F	0.45
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-8F	0.41
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-9F	0.42
192604261	2/5/2015 12:20	1.42E+12	5	FE_DIR	MAX	FA-10F	0.47

Citi has experienced an issue with the U4V services crashing when running a large number of custom reports concurrently. As a result EMC has made modifications to the U4V code (just for Citi V8.0.3.9) to allow custom reports to be run on one day but request data from the previous 24 hours. The following operational procedure was created to rectify the issue:

Custom reporting procedure

1. When adding a new set of reports for an array we need to bundle them and schedule them inside U4V to execute every 15 minutes starting at 12:15am. This has been working on all U4V instances without crashing U4V (runs out of memory).

2. There is a script that runs on the U4V server (old windows AT scheduler) to collect the reports created at 15 minute intervals and move them all to a share drive for processing into various software tools. You will need to allow 30-60 minutes **“after”** the last scheduled custom report to insure the last set of reports have time to complete. This script normally runs around 2am but can be moved to later in the morning based on the number of arrays to process.
3. There is a second script normally run a 3am to check U4V status and restart U4V if custom reports caused U4V to crash. You can move this script to execute later in the morning (60 mins after last custom report script) if needed.

Sample showing array report bundles are schedule at 15 minute intervals (12:15, 12:30, 12:45 etc.)

EMC Unisphere for VMAX V8.0.3.9

All Symmetrix

Home

Performance

Search

Support

[All Symmetrix](#) > [Performance](#) > [Reports](#)

Name	Symmetrix... 1▲	Format	Recurring	Run Date	Email Enabled	Description
LPPERFPORTS1369	000195701369	CSV	●	Sat Aug 29 2015 12:15:00 AM		
LPPERFHOSTIO1369	000195701369	CSV	●	Sat Aug 29 2015 12:15:00 AM		
LPPERFBE1369	000195701369	CSV	●	Sat Aug 29 2015 12:15:00 AM		
LPPERFFE1369	000195701369	CSV	●	Sat Aug 29 2015 12:15:00 AM		
LPPERFWP1369	000195701369	CSV	●	Sat Aug 29 2015 12:15:00 AM		
LPPERFWP1703	000195701703	CSV	●	Sat Aug 29 2015 12:30:00 AM		
LPPERFPORTS1703	000195701703	CSV	●	Sat Aug 29 2015 12:30:00 AM		
LPPERFHOSTIO1703	000195701703	CSV	●	Sat Aug 29 2015 12:30:00 AM		
LPPERFBE1703	000195701703	CSV	●	Sat Aug 29 2015 12:30:00 AM		
LPPERFFE1703	000195701703	CSV	●	Sat Aug 29 2015 12:30:00 AM		
LPPERFHOSTIO2186	000195702186	CSV	●	Sat Aug 29 2015 12:45:00 AM		
LPPERFBE2186	000195702186	CSV	●	Sat Aug 29 2015 12:45:00 AM		
LPPERFFE2186	000195702186	CSV	●	Sat Aug 29 2015 12:45:00 AM		
LPPERFWP2186	000195702186	CSV	●	Sat Aug 29 2015 12:45:00 AM		
LPPERFPORTS2186	000195702186	CSV	●	Sat Aug 29 2015 12:45:00 AM		

Create

Copy

Edit

Schedule

Cancel Schedule

Run Now

Delete

Navigate to Notifications

When adding/editing a new set of reports check to see when the last array report bundle was run and schedule the new report bundle to run at the next available 15 minute interval.

Edit Report

General **Schedule** Email

☒ **Schedule**

Select the initial run time and the days of the week to run the report

First Run Time 12 : 45 am

Day(s) to run ☒ Sunday ☒ Monday ☒ Tuesday ☒ Wednesday
☒ Thursday ☒ Friday ☒ Saturday

Retention (Days) 7

OK Cancel Help

The second part of adding a new report is selecting the data extraction range. In U4V V8+ select “custom” from the “Time Range” line and setup as follows.

- Insure start “day” is the same day (actual date does not matter just that the day is the same)
- The start time should be 12:00AM
- The end time should be 11:55PM (the last 5 minute interval that can be collected by U4V)
- After you click OK to save the entry please double check the times again as sometimes the AM/PM do not take. Use the up/down arrows to change AM/PM and click ok/double check again.

Edit Report

General Schedule Email

Name LPPERFPORTS2186

Description

Symmetrix ID 000195702186

Format CSV

Time Range Custom

Generated Time Zone (UTC) Coordinated Universal Time

Queries

Title	Query
PM Ports Pct Busy	

Add Delete Edit

Custom Time Selection

Select a Start and End Time

Start Time July 30 2015 12:00 am

End Time July 30 2015 11:55 pm

OK Cancel Help

OK Cancel Help

Scripts running via AT schedule to check

Run command “at” from a windows command prompt. This will display the scripts scheduled via windows “at” scheduler

Sample “at” command showing 3 scripts scheduled

```
C:\>at
Status ID      Day                Time                Command Line
-----
1      Each Su                3:00 PM             D:\scripts\U4V-cleanup.cmd
15     Each M T W Th F S Su    3:00 AM             D:\scripts\start-u4v.cmd
1>D:\scripts\u4v-restart.log 2>&1
3      Each M T W Th F S Su    2:07 AM             D:\scripts\move-uv4-querydata-
to-nas-share.cmd 1>D:\scripts\u4v-move.log 2>&1
```

If you need to change the run time of the move or U4V restart scripts you will need to first delete the existing scheduled script then re-add it back in with a new run time. Get the status ID number and issue the delete command (sample below)

```
C:\>at 3 /delete
```

The commands to add the U4V report extract move to share or check/restart U4V are below. Just change the run time and submit the entire command as displayed below.

U4V move to share command

```
at      02:00      /every:M,T,W,Th,F,S,Su    D:\scripts\move-uv4-querydata-to-nas-share.cmd
1^>D:\scripts\u4v-move.log 2^>^&1
```

U4V restart command

11.2.24.3 Export Unisphere Reports

U4V has an excellent export feature that will allow most tables/charts & data to be exported to various formats.

On the pane where you have the metric you wish to export - click the export button




A wizard launches, where you specify the Item you wish to export, and depending on what you are exporting options are provided as to what data to export and what format



Exported Data is saved locally to your desktop in the chosen format

Read and Write Response Times



Date	SG_SAMPLED Avg	SG_SAMPLED Avg
Thu Feb 5 11:30:00 GMT+0000 2015	0	0
Thu Feb 5 11:35:00 GMT+0000 2015	0	0
Thu Feb 5 11:40:00 GMT+0000 2015	0	0
Thu Feb 5 11:45:00 GMT+0000 2015	0	0
Thu Feb 5 11:50:00 GMT+0000 2015	0	0
Thu Feb 5 11:55:00 GMT+0000 2015	0	0
Thu Feb 5 12:00:00 GMT+0000 2015	0	0
Thu Feb 5 12:05:00 GMT+0000 2015	0	0
Thu Feb 5 12:10:00 GMT+0000 2015	0	0

12 VMAX3 Licensing

VMAX3 licenses are installed on the VMAX3 storage array. A license file that contains all the license keys for features purchased for the array can be obtained from [EMC Online Support](#). The license file can then be copied onto the management server running Solutions Enabler and pushed out to the VMAX3 storage array by issuing the following command:

```
# symmlmf add -type emclm -sid <SymmID> -file <license filename> -v
```

The following command to list the licenses on the VMAX3:

```
# symmlmf -sid <SymmID> list -type emclm
```

Refer to the [EMC Solutions Enabler Installation Guide](#) for descriptions of the VMAX3 licenses.

13 VMAX3 Storage Array Configuration

Fully Automated Storage Tiering (FAST) on VMAX3 is not user configurable. Disk groups, data volumes, thin pools, and storage resource pools (SRP) are configured by EMC in the factory or via a bin file change. Online changes to SRPs without re-initialization is currently not supported. A maximum of 2 SRPs can be configured on the VMAX3. Only one SRP should be configured containing all drives in the system for best utilization and efficiency. EMC will create a disk group for each drive type (technology, capacity, and speed) and RAID configuration. Each disk group will have one thin pool associated with it. EMC will also create the TDATs on the disk groups based on drive capacity and RAID type. Once the SRP is created, the usable capacity of the SRP is available to be provisioned to the host servers.

13.1 SRP Reserve Capacity

The reserved capacity of an SRP is the percentage of the SRP's total capacity that is reserved strictly for host writes. Once the SRP's free capacity reaches the percentage specified by the reserved capacity parameter, only host writes will be allowed to use up the remaining free space. Local replication (e.g. SnapVX) and remote replication (e.g. SRDF/A DSE) features will not be able to consume additional capacity if the SRP's free capacity goes below the reserved capacity. Any existing snapshot sessions will be invalid once the reserved capacity is reached.

Since none of Citi's storage arrays should go over 85% allocated capacity, the reserved capacity should be no less than 15%. The default reserved capacity for an SRP is 10%. The reserved capacity can be changed on an SRP using symconfigure with the following line in the symconfigure command file:

```
set srp SRP_1, resv_cap=15;
```

To view the SRPs reserve capacity, use the following command:


```
# symcfg -sid 973 list -srp -v
```

Symmetrix ID	: 000196700973
Name	: SRP_1
Description	:
Default SRP	: FBA
Usable Capacity (GB)	: 87570.0
Allocated Capacity (GB)	: 7141.3
Free Capacity (GB)	: 80428.7
Subscribed Capacity (GB)	: 15994.7
Subscribed Capacity (%)	: 18
Reserved Capacity (%)	: 15
Usable by RDFA DSE	: Yes

14 Storage Provisioning

14.1 SAN Zoning

Host initiators must be zoned to the storage array front-end ports on the SAN before storage can be provisioned to the host. The figure below shows how a host should be connected to the SAN:



The host server should have at least 2 initiators connected to the SAN, one to fabric A and one to fabric B. If more than 2 host connections are available, split the host connections evenly between each fabric. Each host initiator should be zoned to a port on 2 different directors. One zone should contain one host initiator and 2 FA ports on the VMAX3 (single initiator, dual target zone). Zoning using this method will ensure that each host will have at least 4 paths to the storage array, each path going to a different director (assuming the VMAX3 has 2 or more engines).

14.2 Provision Storage to Hosts

All devices (LUNs) on the VMAX3 are thin (or virtual) LUNs. Auto-provisioning is still used to provision storage to host servers. Auto-provisioning uses masking views to provision storage to hosts. A masking view comprises of 3 components: initiator group, storage group, and port group.

Initiator Groups

Host initiators (WWPNs) for a host are grouped into one initiator group. Port flag settings specific to that host OS platform are also set at the initiator group level. Reference the [EMC ESSM VMAX3 Director Bit Settings](#) document to determine which port flags should be set according to the host's OS platform.

The following command creates an initiator group:

```
# symaccess -sid <SymmID> create -name <ig name> -type initiator
```

Example: `symaccess -sid 973 create -name namCSBE14_ig -type initiator`

The following command adds host initiators to an initiator group. Run this command for each host initiator in the host:

```
# symaccess -sid <SymmID> -name <ig name> -type initiator add -wwn <initiator wwpn>
```

Example: `symaccess -sid 973 -name namCSBE14_ig -type initiator add -wwn 10000000c94cae48`

The following command sets port flags for an initiator group:

```
# symaccess -sid <SymmID> -name <ig name> -type initiator set ig_flags <on <Flag> <-enable | -disable> | off [Flag]>
```

Example: `symaccess -sid 973 -name namCSBE14_ig -type init set ig_flags on spc2 -enable`

A cascaded initiator group is an initiator group that contains multiple initiator groups. This is helpful when creating shared storage for nodes in a cluster. In this case, the parent initiator group will contain the initiator groups for the nodes in the cluster. It is important that port group that is used for the cluster's masking view contains all ports that each node is zoned to.

Port Groups

Port groups should comprise of 4 FA ports, each on a different director. For simplicity, use the same port number on each director in the port group. Hosts should be spread evenly among all available port groups to provide an even distribution of load on the FAs. Host initiators need to be zoned to the FA ports in the port group that will be eventually be used in that host's masking view(s).

The following command lists host initiators that are logged into the FA ports:

```
# symaccess -sid <SymmID> list logins
```

The following command creates a port group:

```
# symaccess -sid <SymmID> create -name <pg name> -type port -dirport <dir:port#>
```

Example: `symaccess -sid 973 create -name 1d7-2d7-3d7-4d7_pg -type port -dirport 1d:7,2d:7,3d:7,4d7`

Storage Groups

Storage groups specify which devices (TDEVs) are going to be provisioned hosts. Citi's [LUN Sizing Guidelines](#) should be followed when creating VMAX3 devices. All devices on VMAX3 are TDEVs.

The following command creates a storage group:

```
# symmsg -sid <SymmID> create <SG name>
```

Example: symmsg -sid 973 create namCSBE14_sg

Use symconfigure with the following command to create TDEVs and add them to a storage group:
create dev count=<# of devices>, size=<device size in MB/GB/cyl>, emulation=FBA, config=TDEV, sg=<SG Name>;

Example: create dev count=2, size=34 GB, emulation=FBA, config=TDEV, sg=pp-dmp-migration_sg;

A Service Level Objective (SLOs) and workload can be associated with the storage group to provide it with an appropriate level of performance. There can be up to 5 SLOs (diamond, platinum, gold, silver, and bronze) to choose from depending on what type of drives are installed in the storage array. If no SLO is associated with the storage group, the default setting of Optimized will be used which will place the most active data on the highest performing storage and least active data on the lowest performing storage. The default setting of Optimized should be used if the performance requirements are not known for the application. The following table shows what drives can be used for device allocation for each SLO.

FAST Service Level Objective Compliance				
	Flash	15k	10k	7k
Diamond	Alloc			
Platinum		Alloc		
Gold		Alloc		
Silver			Alloc	
Bronze			Alloc	
Optimized		Alloc		

(EMC, 2015)

A workload (OLTP, OLTP_Rep, DSS, or DSS_Rep) can also be defined with the storage group to further fine tune the performance.

The following command lists all the SLOs and their associated workloads along with the approximate response time for each:

```
# symcfg -sid 973 list -slo -detail -all
```

SERVICE LEVEL OBJECTIVE

S

Symmetrix ID : 000196700973

Name	Workload	Approx Resp Time (ms)
-----	-----	-----
Optimized	N/A	N/A
Diamond	OLTP	0.8
Diamond	OLTP_REP	2.3
Diamond	DSS	2.3
Diamond	DSS_REP	3.7
Diamond	<none>	0.8
Platinum	OLTP	3.0
Platinum	OLTP_REP	4.4
Platinum	DSS	4.4
Platinum	DSS_REP	5.9
Platinum	<none>	3.0
Gold	OLTP	5.0
Gold	OLTP_REP	6.5
Gold	DSS	6.5
Gold	DSS_REP	7.9
Gold	<none>	5.0
Silver	OLTP	8.0
Silver	OLTP_REP	9.5
Silver	DSS	9.5
Silver	DSS_REP	10.9
Silver	<none>	8.0
Bronze	OLTP	14.0
Bronze	OLTP_REP	15.5
Bronze	DSS	15.5
Bronze	DSS_REP	16.9
Bronze	<none>	14.0

A cascaded storage group is a storage group (the parent SG) that contains multiple storage groups (the child SGs). Each child SG can have its own SLO/workload associated with it. This is useful for hosts that have devices with different performance requirements.

The following command associates a storage group with an SLO and workload:

```
# symmsg -sid <SymmID> -sg <sg name> set -slo <slo> -wl <workload>
```

Example: `symmsg -sid 973 -sg pp-dmp-migration_sg set -slo diamond -wl oltp`

Masking View

A masking view will need to be created once the initiator group, port group, and storage group has been created for the host. Once the masking view is created, the devices in the storage group will

be accessible to all ports in the port group and the host will now be able to access the devices through those ports.

The following command creates a masking view:

```
# symaccess -sid <SymmID> create view -name <mv name> -sg <sg name> -pg <pg name> -ig <ig name>
Example: symaccess -sid 973 create view -name namCSBE14_mv -sg namCSBE14_sg -pg 1d7-1d7-3d7-4d7_pg -ig namCSBE14_ig
```

15 VMAX3 Security Configuration

15.1 Symacl Host Access Control

Symmetrix Host Access Control management (SYMACL) sets up or updates the level of access control for a Symmetrix array. Before a Solutions Enabler host can issue SYMCLI commands, it requires permissions in each Symmetrix storage array's SYMACL database. All information regarding access control is stored within directly on the Symmetrix array.

The symacl database uses the following access control mechanisms to govern access to the Symmetrix device pools:

- *Access control groups* - Unique *access IDs* and *names* are assigned (together) to hosts and users and then sorted into access control *groups* according to similar needs (determined by an Administrator). Access groups are allowed to act on device pools based on *rights (access types)* granted by the Administrator.

The unique host ID for open systems can be viewed by running “symacl –unique”

- *Device pools* - Access rights (types), such as BCV, SRDF, ADMIN, give permission to a host or user to perform certain Solutions Enabler functionality on a specified device pool. Device pools are sometimes referred to as access pools or accpool.
- *Access control entries (ACEs)* - Once the group and device pool mechanisms are established, the ACEs are created, which grant access rights to these pools. The ACEs of the various access groups (along with groups and pools) are managed and stored in the access control database.
- *Access control list (ACL)* - is a group of ACEs that are associated with the same “accgroup”. For example, the ACEs that are associated to a control group comprise an ACL in the database.

15.1.1 Symacl Setup

SYMACL is initialized by EMC. The SE server's unique ID must be provided to EMC before SYMACL is enabled on the Symmetrix to ensure that it has management access to array after SYMACL is enabled. The SYMAPI_ALTERNATE_ACCESS_ID parameter in the /var/syampi/config/options file should be set to ENABLE on the SE management server and all SE clients connecting to this server via a symapi server/client connection. This provides a randomly generated unique ID for the host that is not tied to the server's hardware. This is the recommended setting for all intel servers. To set this parameter, enter the following line in the /var/syampi/config/options file:

SYMAPI_ALTERNATE_ACCESS_ID = ENABLE

After editing the options file on the SE servers/clients, issue the 'symacl -unique' command to generate the unique ID (the unique ID returned in the command will be the same on the SE server and client).

```
# symacl -unique
The unique id for this host is: 2A5B04A6-54408FCE-9B397178
#
```

Note: all host identifiers in the examples are invalid and for example only

Once EMC has enabled SYMACL, the initial symacl list should have at least the following for VMAX and earlier Symmetrix models:

```
# symacl -sid 919 list -acl

Symmetrix ID: 000190100919
```

Group Name	Pool Name	Access Type
AdminGrp	ALL_DEVS	ALL
AdminGrp	ALL_DEVS	ADMIN
UnknwGrp	ALL_DEVS	Base
UnknwGrp	!INPOOLS	ALL

The initial symacl list will have the following for VMAX3:

```
# symacl -sid 973 list -acl

Symmetrix ID: 000196700973
```

Group Name	Pool Name	Access Type
AdminGrp	ALL_DEVS	ADMIN
AdminGrp	ALL_DEVS	ALL
EmcIntrnlGrp	ALL_DEVS	ALL
EmcIntrnlGrp	ALL_DEVS	ADMINRD
EmcNasGrp	ALL_DEVS	BASE
EmcNasGrp	EmcNasPool	BASECTRL
UnknwGrp	!INPOOLS	ALL
UnknwGrp	ALL_DEVS	BASE

The initial SE host should already be a member of the AdminGrp accgroup. All SE hosts should be made a member of the SYMACL group called "AdminGrp" on all local and remote Symmetrix arrays that it will manage. The entry in the command file used to add a host to an accgroup is:

```
add host accid F963D5A6-7FECA3CD-5FE83C25 name CSBEse01 to accgroup AdminGrp;
```

Note: There is an 8 character limit for the hostname used to name the acid on VMAX and earlier Symmetrix arrays. The limit is 31 characters on VMAX3.

When executing the command file, the user will be prompted to enter a PIN (the PIN is the last four digits of the array serial number):

```
# symacl -sid 973 -f addhost.cmd -v commit

Enter Access PIN:

Contents of command file: addhost.cmd
.....
add host accid F963D5A6-7FECA3CD-5FE83C25 name CSBEse01 to accgroup AdminGrp;
.....

Command file: (addhost.cmd)
PREVIEW.....Started.
PREVIEW.....Done.

PREPARE.....Started.

  Adding Host access id CSBEse01 to group AdminGrp.....Done.

PREPARE.....Done.

Starting COMMIT.....Done.
```

The accgroup UnknwGrp shall be deleted from the symacl list to prevent unauthorized servers from accessing the storage array. The entry in the command file to remove the UnknwGrp is:

```
delete accgroup UnknwGrp remove_aces=true;
```

The Symmetrix service processor will also need symacl entries created for it. The unique ID for the Symmetrix SP can be obtained from EMC. An accgroup called EMC_SP should be created with VLOGIX, ADMINRD, OPTMZR, and BASE rights to ALL_DEVS. The Symmetrix SP will need to be added to the accgroup EMC_SP. On VMAX3, the symacl entries are created for the MMCSS by default and are contained in the accgroup EmcIntrnlGrp.

The ECC storage agent server will also need symacl entries created for it. Create a group called ECC_HOST and grant it BASE, ECC, and ALL rights to ALL_DEVS. Add the ECC storage agent server to the ECC_HOST accgroup.

The final symacl list should look like the following for VMAX and earlier Symmetrix models:

``` # symacl -sid 380 list -acl ```

Symmetrix ID: 000192601380

Group Name	Pool Name	Access Type
AdminGrp	ALL_DEVS	ADMIN
AdminGrp	ALL_DEVS	ALL
EMC_SP	ALL_DEVS	ADMINRD
EMC_SP	ALL_DEVS	OPTMZR
EMC_SP	ALL_DEVS	BASE
EMC_SP	ALL_DEVS	VLOGIX
ECC_HOST	ALL_DEVS	BASE
ECC_HOST	ALL_DEVS	ECC
ECC_HOST	ALL_DEVS	ALL

The final symacl list should look like the following for VMAX3:

``` # symacl -sid 973 list -acl ```

Symmetrix ID: 000196700973

Group Name	Pool Name	Access Type
AdminGrp	ALL_DEVS	ADMIN
AdminGrp	ALL_DEVS	ALL
EmcIntrnlGrp	ALL_DEVS	ALL
EmcIntrnlGrp	ALL_DEVS	ADMINRD
EmcNasGrp	ALL_DEVS	BASE
EmcNasGrp	EmcNasPool	BASECTRL

Note: If symauth is enabled on the storage array, only users with the Admin or SecurityAdmin role will be able to make symacl changes.

The following should be done if application hosts need access to their devices for Timefinder/SRDF control:

1. Create an accgroup for the application host.
2. Add the application host to the accgroup.
3. Create a device pool comprised of only the devices that the application host owns and needs to control.
4. Grant BCV or RDF access to the accgroup for the device pool.

-
5. Grant BASE access to the accgroup for ALL devices.

15.1.2 SYMACL Backup and Restore

A backup copy of the symacl database should be created before and after any symacl configurations are made.

The backup operation saves the contents of the access control database in the file specified by the file option. The file must not previously exist. The backup file created is a text command file compatible for use with the symacl utility.

To create a backup of the symacl database, use the following command:

```
symacl -sid SymmID backup -file Filename
```

Note: The backup file contains encrypted versions of the unique IDs; therefore, if you compare the values in the backup file to the original file used to create the database, they will be different.

```
$ symacl -sid 973 backup -f 0973_acl_2015-02-23.bak

$ cat 0973_acl_2015-02-23.bak
add restored accid 00FD3198-E16B13EC-DE471DA2 name nmCSBE14 to accgroup
AdminGrp;
add restored accid 62514033-E9816427-DEEF42EC name admin to accgroup AdminGrp;
add restored accid AB1723BA-B173858A-C16489F0 name CSBEse01 to accgroup
AdminGrp;
grant access=ADMIN to accgroup AdminGrp for ALL devs;
grant access=ALL to accgroup AdminGrp for ALL devs;
```

The restore operation replaces the contents of the access control database with the contents of the file specified by the file option. It eliminates whatever currently exists in the Symmetrix configuration database and replaces it with the contents of the backup file.

To restore a backup using symacl you use the `-restore` option as shown below:

```
symacl -sid SymmID commit [-v|-noecho] -restore -file Filename
```

15.2 Symauth Role-Based User Access Control

The symauth feature provides role-based user access control for Symmetrix storage arrays. The `symauth symCLI` command provides user-based authorization that assigns a user or group to a role to limit the management operations they can perform on a Symmetrix array. Solutions Enabler user ID is assigned to one role only. The following depicts the available roles:

- **Admin** - Performs all operations on a Symmetrix system, including security operations in addition to all StorageAdmin and Monitor operations.
- **SecurityAdmin** - Performs security operations on a Symmetrix system in addition to all monitor operations.
- **StorageAdmin** - Performs all management (active or control) operations on a Symmetrix system in addition to all Monitor operations. This role does not allow users to perform security operations.

- **Monitor** - Performs read-only (passive) operations on a Symmetrix system excluding the ability to read the audit log or access control definitions.
- **Auditor** - Grants the ability to view, but not modify, security settings for a Symmetrix system (including reading the audit log, symacl list, and symauth) in addition to all Monitor operations. This is the minimum role required to view the Symmetrix audit log.
- **PerfMonitor** - Performs the same operations as a monitor, with the addition of being able to set performance alerts and thresholds.

Internally, Solutions Enabler represents a user/group identity as a string that comprises the user's name and authentication source. The possible encodings are displayed below:

Authentication String	Description
H:HostName\UserName	A user authenticated by the local host operating system
D:DomainName\UserName	A user authenticated by a specific Domain on Windows
L:ServerName\UserName	A user authenticated by an LDAP Server
C:HostName\UserName	A user authenticated by the private SMC or Unisphere for VMAX 1.0 authentication service on some host
V:DomainName\UserName	A user authenticated by a Virtualization Domain

Before enabling symauth on the storage array, symauth entries should be created for users that will be managing the storage array. The symauth entries should be associated with local host authenticated users, where each user has direct access.

15.2.1 To create symauth host based user entries on the array:

1. Ensure that symauth is disabled on the array:

```
# symauth -sid 110 list
```

```
S Y M M E T R I X   A U T H O R I Z A T I O N   S T A T U S
```

```
Symmetrix ID: 000190103110
```

```
Authorization Control      : Disabled
```

```
Time Enabled              : Fri May 15 10:13:51 2009
```

```
Time Disabled             : N/A
```

```
Time Updated              : Mon Sep 21 20:07:26 2009
```

```
Enforcement Mode          : Enforce
```

```
Server Policy              : Trust clients
```

2. Create a command file called **adduser.cmd** with the following entries:

```
assign user H:sangs7lsw54321 to role SecurityAdmin;
assign user H:sangs7lcp12345 to role StorageAdmin;
```

3. After the command file is create, execute the command file with the **symauth** command:

```
# symauth -sid 110 preview -f adduser.cmd

PREVIEW.....Started.
PREVIEW.....Done.
# symauth -sid 110 commit -f adduser.cmd

Perform an Authorization change on Symmetrix unit 000190103110 (y/[n]) ? y

PREVIEW.....Started.
PREVIEW.....Done.

COMMIT.....Started.
assign user H:sangs7\sw54321 to role SecurityAdmin.....Done.
assign user H:sangs7\cp12345 to role StorageAdmin.....Done.
COMMIT.....Done.
#
```

4. Verify the user entries:

```
# symauth -sid 110 list -users

SYMMETRIX AUTHORIZATION USERS

Symmetrix ID: 000190103110

Role name      Username
-----
SecurityAdmin   H:sangs7\sw54321
StorageAdmin    H:sangs7\cp12345
#
```

On RHEL, the users with the StorageAdmin role will need to have **rwX** access to the files in the **/var/symapi/config** and **/var/symapi/db** directory.

15.2.2 To create symauth domain based group entries on the array:

User authentication can be accomplished though Windows Domain groups. This simplifies access management by assigning users to groups then managing the groups. Groups can belong to multiple SE servers, removing the redundancy of managing users for each SE server independently. The following example uses the “D” symauth nomenclature representing domain identifier of an AD group:

1. Create a command file called **assign-grp** with the following entries:

```
assign group D:NAMDEV\per_csbe_symm_admin to role StorageAdmin
```

2. After the command file is create, execute the command file with the **symauth** command:

```
# symauth -sid 4261 -file assign-grp commit -v -nop
```

```
PREVIEW.....Started.
assign group D:NAMDEV\per_csbe_symm_admin to role StorageAdmin

PREVIEW.....Done.

COMMIT.....Started.
assign group D:NAMDEV\per_csbe_symm_admin to role StorageAdmin.....Done.

COMMIT.....Done.
```

3. Verify group entry

```
# symauth -sid 4261 list -users
```

```
S Y M M E T R I X   A U T H O R I Z A T I O N   U S E R S
```

```
Symmetrix ID: 000192604261
```

```
Flags
```

Role	User/Group name	Component	E
Admin	User Administrator	N/A	.
Admin	User root	N/A	.
Admin	User smc	N/A	.
Admin	Group L:namdev.nsrootdev.net*	N/A	.
StorageAdmin	User H:rd-2k8r2-507srm\SYSTEM	N/A	.
StorageAdmin	Group D:NAMDEV\per_csbe_symm_*	N/A	.
Auditor	User H:* \storevntd	N/A	.
Monitor	User H:CSBEEMCSMC01\SPA	N/A	.
None	User *	N/A	.
PerfMonitor	Group D:namdev.nsrootdev.net*	N/A	.

```
Legend for Flags:
```

```
(E) : N = Rule has no effect since a different rule grants greater rights.
      : . = Rule is active.
```

```
# symauth -sid 4261 list -user -v
```

```
S Y M M E T R I X   A U T H O R I Z A T I O N   U S E R S
```

```
Symmetrix ID: 000192604261
```

```
Role name          : Admin
```

```

User name      : Administrator
Component type :
Component name :
Flags         :

Role name      : StorageAdmin
Group name     : D:NAMDEV\per_csbe_symm_admin
Component type :
Component name :
Flags         :

```

4. Commands generated by users in an AD group are tracked by the individual user identifier:

Symaudit log:

```

Record Number      : 375635
  Records in Seq   : 1
  Offset in Seq    : 1
  Time             : 01/09/13 11:14:15
  Vendor ID        : EMC Corp
  Application ID    : SYMCONFIGURE
  Application Version : 7.5.0.7
  API Library       : SEK
  API Version       : V7.5.0.7 (Edit Level: 1604)
  Host Name        : CSBE-ECC03
  OS Name          : WinNT
  OS Revision       : 5.2.3790Se
  Client Host      :
  Process ID       : 00007152
  Task ID          : 00009664
  Function Class    : CfgChg
  Action Code       : Commit
  Text             : The local CFGCHG COMMIT operation SUCCEEDED
  Username         : D:NAMDEV\gs96899
  Activity ID      : SE92f7657591

```

Windows Event Log:

```

Event Type: Success Audit
Event Source: Security
Event Category: Object Access
Event ID: 560
Date: 1/9/2013
Time: 11:12:37 AM
User: NAMDEV\gs96899
Computer: CSBE-ECC03
Description:
Object Open:
  Object Server: Security
  Object Type: Key
  Object Name: \REGISTRY\MACHINE\SOFTWARE\Microsoft\Windows
NT\CurrentVersion\Winlogon
  Handle ID: 1560
  Operation ID: {0,21504116}
  Process ID: 7152

```

```

Image File Name:    C:\Program Files\EMC\SYMCLI\bin\symconfigure.exe
Primary User Name:  gs96899
Primary Domain:     NAMDEV
Primary Logon ID:   (0x0,0x3E65B7)
Client User Name:   -
Client Domain:      -
Client Logon ID:    -
Accesses:           DELETE
                   READ_CONTROL
                   WRITE_DAC
                   WRITE_OWNER
                   Query key value
                   Set key value
                   Create sub-key
                   Enumerate sub-keys
                   Notify about changes to keys
                   Create Link

Privileges:         -
Restricted Sid Count: 0
Access Mask: 0xF003F

```

For more information, see Help and Support Center at <http://go.microsoft.com/fwlink/events.asp>.

Symevent log:

```

[fmt=symaud] [date=2013-01-09T11:12:39] [symid=000192604261] [orig=SE]
[user=D:NAMDEV\gs96899] [host=CSBE-ECC03] [actid=SE92f7657591]
[appid=SYMCONFIGURE] [aud-cls=CfgChg] [aud-act=Commit] [aud-num=375629,1/2] =
STARTING a local CFGCHG COMMIT to create new symdevs
[fmt=symaud] [date=2013-01-09T11:12:39] [symid=000192604261] [orig=SE]
[user=D:NAMDEV\gs96899] [host=CSBE-ECC03] [actid=SE92f7657591]
[appid=SYMCONFIGURE] [aud-cls=CfgChg] [aud-act=Commit] [aud-num=375630,2/2] =
create dev count=2, size=36828 cyl, emulation=FBA, config=TDEV, mvs_ssid=0;

```

To find out what the current username is for the user currently logged into the system, type the following command:

#symauth show -username

For example:

```
#symauth show -username
```

Your current username is: H:sangs1\ql68074

Your current groupname is: H:sangs1\adm

symauth -sid 110 list -users

SYMMETRIX AUTHORIZATION USERS

Symmetrix ID: 000190103110

Role name	Username
-----	-----

SecurityAdmin	H:sangs7lsw54321
StorageAdmin	H:sangs7lcp12345

On RHEL, the users with the StorageAdmin role will need to have **rwX** access to the files in the **/var/symapi/config** and **/var/symapi/db** directory.

To set ACL entries for individual users, do the following:

```
# setfacl -R -m user:cp12345:rwX /var/symapi/config
# chmod 775 /var/symapi/config
# setfacl -R -m user:cp12345:rwX /var/symapi/db
# chmod 775 /var/symapi/db
```

If wanting to set the ACL entries for user groups, do the following:

```
# setfacl -R -m group:503:rwX /var/symapi/db
# getfacl /var/symapi/db
getfacl: Removing leading '/' from absolute path names
# file: db
# owner: root
# group: root
user::rwX
group::r-x
group:ql68074:rwX
mask::rwX
other::r-x
# setfacl -R -m group:503:rwX /var/symapi/config
# getfacl /var/symapi/config
getfacl: Removing leading '/' from absolute path names
# file: var/symapi/config
# owner: root
# group: root
user::rwX
group::r-x
group:ql68074:rwX
mask::rwX
other::r-x
```

Note: SE will only detect the user's primary group so when setting ACL's for groups, ensure that the ACL is set for the users' primary group.

1. The root user or a user with the proper file ACL will need to modify the **/var/symapi/config/daemon_users** file to allow the alternate users to use the daemons. For simplicity, the **<all> <all>** entry should be used so that this step does not have to be done for each modification to symauth. Ensure that the following line (all lower-case) is added to the daemon_users file to allow non-root users to access the Solutions Enabler daemons:

<all> <all>

2. Once all of the symauth user entries are entered on both the local and remote storage arrays (symauth entries must be also be created for the remote arrays for SRDF functionality), enable symauth on the Symmetrix:

symauth -sid 110 enable

Perform an Authorization change on Symmetrix unit 000190103110 (y/[n]) ? **y**

3. Verify that the **Enforcement Mode** is set to **Enforce**:

symauth -sid 110 list

```

S Y M M E T R I X   A U T H O R I Z A T I O N   S T A T U S

Symmetrix ID: 000190103110

Authorization Control      : Enabled
Time Enabled              : Thu Dec  6 11:03:05 2007
Time Disabled             : N/A
Time Updated              : Thu Dec  6 11:03:05 2007

Enforcement Mode          : Enforce
Server Policy             : Trust clients

```

If it is not set to **Enforce**, then set the mode by:

symauth -sid 110 enable

Perform an Authorization change on Symmetrix unit 000190103110 (y/[n]) ? **y**

15.2.3 LDAP/Kerberos Linux Authentication

LDAP/Kerberos Redhat Linux authentication utilizes group accounts in accessing the SE server. LDAP/Kerberos groups accounts simplify user management and will eventually become the default method for Linux SE server access.

Solutions Enabler (SE) provides access to EMC Symmetrix arrays for provisioning, troubleshooting, monitoring and basic management. Typical user access is performed by creating a local user account on each SE server required by the user. The creation and ongoing management of individual user accounts has become time consuming with the large number of SAN production personnel, Symmetrix arrays, SE servers and SAN operational support systems deployed by Citi.

Utilizing LDAP/Kerberos accounts reduce management complexity by combining users into a group, eliminating the need for creating and managing local host accounts. The LDAP/Kerberos group can be define into a SE authentication role, where all individual members within the group will have the same access profile as the group. The centralization of access management via LDAP/Kerberos coupled with role based SE symauth group functionality simplifies user administration while maintaining a sufficient level of security.

15.2.4 Backing Up and Restoring the Symauth Database

The symauth entries for each array should be backed up on a periodic basis; especially after changes are made to the entries. Use the following command to backup the symauth entries to a file on the SE management server:

```
symauth -sid <SymmID> backup -f <backup filename>
```

To restore the symauth entries from an earlier backup of the symauth database; use the following command:

```
symauth -sid <SymmID> commit -restore -f <backup filename>
```

Note: As of SE 7.2, the quotes are added around the username automatically when backing up the symauth entries.

For example: assign user "D:NT AUTHORITY\SYSTEM" to role Monitor;

15.3 Security Compensating Controls

This section covers compensating controls put in place to close security gaps found during the Citi ISRP process.

15.4 TestSMIProvider Removal

The TestSmiProvider script in the /opt/emc/ECIM/ECOM/bin directory on RHEL and the <install dir>\Program Files\EMC\ECIM\ECOM\bin on Windows management servers should be deleted due to passwords being logged in clear text when using the script. Storing passwords in clear text is a violation of Citi's security standards so the deletion of the TestSmiProvider script is mandatory. EMC is aware of this violation and is working on a fix for this issue in a later version of Solutions Enabler.

16 Monitoring and Alerting

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

16.1 Configuring the Event Daemon (storevntd)

The SE event daemon (storevntd) is a daemon that runs on the SE server and can be used to import the symaudit log as well as to send out Symmetrix alerts via SNMP, syslog, the OS's event log, or to a file. The following steps should be followed to configure the storevntd:

1. The following are the main entries should reside in the `/var/symapi/config/daemon_options` file on the SE server to specify the behaviour of the storevntd (There are some other parameters in the `daemon_options` file for the storevntd but the following are the basic settings to get storevntd to functioning):

```
storevntd:log_event_targets = file syslog      #Direct the logs and events to a local file and
syslog target.

storevntd:log_symmetrix_events = \
    sid=000192601380, audit, tgt=file|syslog;\ #Directs symaudit logs from the specified
Symmetrix to the local file and also send the audit event to the syslog server specified below.
    sid=000190301292, audit, tgt=file|syslog;\
    status, tgt=file;\ #Sends status messages to the file target
    events, tgt=file #Sends event messages to the file target

storevntd:log_event_file_name = symevents #Specifies the prefix of the local log file.

storevntd:log_event_syslog_host = 10.40.41.199 #IP address of the SEMS server specified by
the GSOC team.
```

2. Add a symauth entry for storevntd. In the example below, a command file called **adduser.cmd** was created with the entry **assign user H:sangs7\storevntd to role Auditor**; to create the symauth entry for storevntd:

```
# symauth -sid 110 commit -f adduser.cmd

Perform an Authorization change on Symmetrix unit 000190103110 (y/[n]) ? y

PREVIEW.....Started.
PREVIEW.....Done.

COMMIT.....Started.
assign user H:sangs7\storevntd to role Auditor.....Done.
COMMIT.....Done.

#
```

3. Start storevntd:

```
# stordaeon start storevntd
Waiting for daemon to start. This may take several seconds.
#
```

4. Run the command to have storevntd start automatically after a system reboot:

```
# stordaeon install storevntd -autostart
```

5. Reload the new options specified in the `daemon_options` file for the storevntd:

```
# stordaeomon action storevntd -cmd reload
storevntd      Action performed
#
```

16.2 Configuring Symmetrix Alerting

Symmetrix alert messages are sent to EMC Support Center by the call-home feature based in the service processor. EMC responds by initiating fixes and alerting Citi storage teams of the issue. In parallel, Citi storage teams should be seeing these alerts generated through the storage event daemon (storevntd) on the SE host and should follow up with EMC to make sure they have seen the alert and are acting on it. This redundant process mitigates the risk from a failure of the call-home subsystem transporting the alert to EMC Support Center. Citi problem management system can be used to internally document the fault and track the progress of resolution.

The storevntd daemon should be configured to record Symmetrix events to a file on the Solutions Enabler management server as specified in the previous section. Once the storevntd is configured to capture Symmetrix events, the management server will need to be configured to send events to be alerted on to Citi's Tivoli system. The Tivoli Endpoint agent will be used to scrape the event file used by storevntd and send the alerts to the Tivoli monitoring system for alert notification to the storage support teams. An event ID include list specifies the set of Symmetrix alert messages that CTI and CSBE considers critical for operation. The objective of the include list is to limit the amount of alerts deemed non-critical received by the global operations center. A complete list of event ID's is depicted in [Appendix C](#):

Do the following to setup Tivoli monitoring and alerting:

- 1) The ITM6 Redhat Linux OS Base agent need to be installed on the RHEL SE server if it hasn't already. Refer to the [RFP- Global ITM6 Linux\(i386\) OS Agent for Tivoli Monitoring](#) for instructions on installing the agent.
- 2) Submit Global Data Center Operations Automations Marketplace request to CTI DC Monitoring Services to monitor the storevntd file on the Solutions Enabler management server. Marketplace request can be found at:

https://cmp.nj.smb.com/marketplace/control/product/~productId=47307_253071_GLOBAL#

Notes on submitting request:

- Select **Tivoli** for the technology
- Select **Add or Delete Logfile Adapters** for the Tivoli option
- Select **Add Logfile Adapter** for the Logfile Adapter action
- Select **App Support** for Monitor Type
- Complete EMA_MonitoringSpc MS Excel questionnaire with the following parameters:

User Questionnaire	ESS.Tivoli.Add Logfile Adapter
Tivoli	Add logfile adapter
Hostname	<SE Server>

Full Path to Logfile	/var/symapi/log/symevents.log0 /var/symapi/log/symevents.log1
Case sensitive string to look for	evtid=1050 evtid=1074 evtid=1077 evtid=1080 evtid=1117 evtid=1119 evtid=1142 evtid=1144 evtid=1147 evtid=1148 evtid=1152 evtid=1201 evtid=1202 evtid=1203 evtid=1204 evtid=1208 evtid=1221 evtid=1512

Below are the descriptions of the event IDs:

Event ID	Description
1050	Diagnostic event trace triggered.
1074	Service processor could not complete a call for service.
1077	Power system AC line interruption detected.
1080	Single SRDF link in an RDF group is not operational.
1117	Director not responding.
1119	SRDF/A session dropped: write pending limit reached. Host throttling is disabled.
1142	Power Supply A multiple fan fault.
1144	Power Supply A faulted.
1147	Power Supply B single fan fault.
1148	Power Supply B faulted.
1152	Supplemental Power Supply internal fault.
1201	Array state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1202	Director state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1203	Port state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1204	Disk state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1208	[<NAME> Pool utilization is now <NN> percent.

1221	Disk state is now [Offline Offline Spare Online Online Spare Online Degraded] (was: [Offline Offline Spare Online Online Spare Online Degraded])
1512	The allocation for SRP has changed to <NN> %u percent.

3. Once the request is received and processed, the Tivoli administrator will push out the Tivoli KLO Agent to the SE server to start monitoring the SE symevents files for alerting.

16.3 Audit Log Reconciliation via SEMS

Citi's Security Event Monitoring System (SEMS) is used to centrally collect security audit events. Custom reports can then be generated from the events collected by SEMS and sent to the log review teams for reconciliation. The following steps need to be followed to integrate the Symmetrix with SEMS for audit log reconciliation:

1. storevntd needs to be configured on the SE management server to send Symmetrix syslog audit events to the SEMS server.
2. The SEMS GSOC team should be contacted to initiate the request for integrating the Symmetrix with SEMS for log reconciliation. The SE management station should be specified as the source device in the request. Use the following distribution: ***CTI Global SOC SDM**
 - a. If PowerBroker is used for user access control and audit logging, the existing log reconciliation process for PowerBroker should be followed for reconciling Symmetrix changes performed by Citi personnel. Only vendor initiated changes should be reconciled using the SEMS system for environments using PowerBroker since PowerBroker cannot track vendor changes which are initiated on the Symmetrix Service Processor. Please note in the request to the SEMS team whether PowerBroker is being used. If Powerbroker is not being used, ensure that symauth is implemented.
 - b. For environments that have implemented symauth for user access control and audit logging, all symaud events should be reported and reconciled with the exception of the events in the exclude list specified in the next step.
3. The GSOC team will then configure Arcsight to receive the Symmetrix syslog audit events and generate daily reports to the regional log review teams.

The GSOC team will apply a filter to the report to only include audit events (fmt=symaud) and an event description containing the following:

- 1) UNMAPPING
- 2) UNMAP
- 3) MAPPING
- 4) MAP
- 5) BASE CONTROL
- 6) CREATE
- 7) ADD
- 8) DELETE
- 9) APPLY
- 10) RENAME
- 11) INITIALIZE
- 12) MODIFY

-
- 13) REMOVE
 - 14) RWENABLE
 - 15) RWDISABLE
 - 16) CODELOAD
 - 17) HBA_PORT_FLAGS_ENABLE
 - 18) RENAME_WWN_ALIAS
 - 19) CLEAR_WWN_ALIAS
 - 20) HBA_PORT_FLAGS_DISABLE
 - 21) HBA_PORT_FLAG_OPTION_CLEAR
 - 22) REFRESH_DEVMASK_DB
 - 23) ADD_DEVICES
 - 24) REMOVE_DEVICES
 - 25) FAST
 - 26) FAST_VP
 - 27) FAST_VP
 - 28) DELETING
 - 29) CREATING
 - 30) LICENSE
 - 31) DEVMASK
 - 32) METADEVICE
 - 33) ATTRIBUTE
 - 34) BCV
 - 35) CLONE
 - 36) RDF
 - 37) RCoPY
 - 38) QoS
 - 39) CFGCHG
 - 40) RECOVERY

For environments still using PowerBroker, audit events showing up under the 'root' user of the SE management server should also be excluded from the report since PowerBroker has its own log review process. It is important to specify whether PowerBroker is being used in the request to the GSOC team so that they know which filter to apply to the reports for log reconciliation.

16.3.1 Performance

16.4 Remote Vendor Support

16.4.1 EMC Secure Remote Support (ESRS)

The VMAX3 will be integrated with Tivoli for alerting Citi Storage operations teams. Reference the Solutions Enabler Solutions Guide for instructions on setting up the SE management server for Tivoli alerting for VMAX3.

ESRS will be used for vendor (EMC) remote monitoring and access. Reference the [ESRS Solutions Guide](#) for ESRS setting up ESRS for VMAX3.

17 Supported VMAX3 Features

The table below lists the VMAX3 features and their current status in Citi.

Feature	Item	Function	Citi Status
Virtual Provisioning	Native thin provisioning	All VMAX3 volumes are thin devices.	Supported
FAST	Fully automated storage provisioning	Distributes data across tiers of storage based SLO and activity on the volume	Supported
SRDF	Remote replication	Replicate data between storage arrays. Only 2-site replication supported.	Supported
SnapVX	Local replication	Provides local point-in-time copies of the data.	Supported
ZSR	Zero Space Reclamation	De-allocates zero space in LUNs and puts the reclaimed space back into the SRP.	Supported

17.1 SRDF Remote Replication

SRDF is a license-enabled (Remote Replication Suite) feature on Symmetrix storage arrays used for replicating data between Symmetrix storage arrays. Citi only supports replication between 2-sites using either SRDF/S or SRDF/A. The local EMC account team should be contacted as soon as an SRDF implementation is proposed and instructed to engage EMC's Shared Services to properly size the SRDF configuration (e.g. RDF links, network bandwidth, DSE maximum capacity, etc.).

SRDF/S provides synchronous replication between arrays that are less than 200 km apart. With SRDF/S, writes on the source array are not acknowledged to the host until the data is written on the target array; thus, guaranteeing that the data on the replicated LUNs are always the same on the source and target array but with a negative performance impact due to the wait time for the write acknowledgement on the target side.

SRDF/A provides asynchronous replication between arrays without any distance limitation. Unlike SRDF/S, the source array does not have to wait for the write to complete on the target array before acknowledging the write to the host; therefore, doesn't have the negative performance impact that SRDF/S does. It uses a Delta Set architecture to maintain a consistent point-in-time target (R2) image with only a slight lag behind the source (R1) side, resulting in minimal data loss if the R1 is lost. SRDF/A allows one to configure an SRDF environment in which R1 devices transfer data to R2 devices in chunk cycles.

17.1.1 SRDF Limitations

The following table lists the limits for SRDF for all VMAX3 models⁸.

⁸ The following EMC limits are for information only as CATE would not expect Citi deployments to get close to these limits.

SRDF scalability	Max
RDF groups per array	250
RAs/Director	16
RDF groups per RA	250
RDF devices per RA	64K
RDF devices per RDF group	32K
RDF devices per array	64K

Note: All RDF groups on VMAX3 are dynamic. It is not possible to create static RDF groups on VMAX3.

17.1.2 SRDF/A Over Very Long Distances

In principle, SRDF/A can run over an unlimited distance. In practice, there are in fact some practical limitations that need to be taken into account before deciding to go ahead with an implementation.

1. Large latencies can dramatically reduce the effective throughput on a network link. For example, the throughput on a nominal 100MB/sec network link with a 4K writes IO workload (typical for SRDF/A) has been measured to reduce dramatically as latency increases.

CSBE has developed Best Practice to optimize WAN throughput for SRDF/A replication using the Cisco FCIP product. See references section for further information.

2. SRDF/A becomes sensitive to packet loss with very high latencies (e.g. 250ms). The consequence are:
 - a. an increased probability that SRDF/A will abort (i.e. replication will cease) when IP packet loss occurs between the source and target sites.
 - b. a significant reduction in throughput when packet loss occurs on the IP network between the source and target sites.

CSBE has developed Best Practice to optimise WAN throughput for SRDF/A replication using the Cisco FCIP product. See references section for further information.

3. Cycle time increases with latency for replication sessions under Multi-Session Consistency control. This becomes significant over 40ms.
4. Much longer runtime for SRDF control and query operations. This becomes significant for latencies over 40ms.

17.1.3 Role of SRDF in Disaster Recovery

It must be stressed that SRDF will only give a business an I/O consistent copy of their data set at the DR site for a particular server, application or set of applications (depending at what level consistency is enforced).

In the case of a production failure that only impacts a server and not the storage infrastructure (e.g. a server hardware failure), there will be no data loss.

In the case where the storage infrastructure or indeed the whole production site fails, the DR copy will be from a point-in-time somewhere between 15-30 seconds (for the default cycle time) prior to the production failure and data loss will occur if SRDF/A is used. If SRDF/S is used, the DR copy will be at the time of failure.

In the case that SRDF/A is used:

- the storage infrastructure will not be able to define exactly how much data was lost (it will be possible to obtain an actual timestamp, indicating the point-in-time that the SRDF/A DR copy represents. This can be provided to application owners after a disaster to help with quantifying data loss in a disaster scenario).

- the storage infrastructure does NOT provide a mechanism for dealing with the data loss. This is the responsibility of the application.

- it is the application owners that need to deal with issues relating to recovery using different point-in-time data set for applications that were not bound together from a consistency perspective i.e. where they were not in the same replication group.

17.1.4 SRDF Configuration

The main steps in preparing for an SRDF implementation for existing applications are as follows. Additional steps required for an MSC implementation are covered in Appendix C.

1. Collect basic requirements information for target applications to determine SRDF configuration and capacity requirements, licensing requirements, etc.
2. Engage with EMC Shared Services to size the SRDF configuration (e.g. RDF links, network bandwidth, DSE max capacity, etc.).
3. Order additional network links or bandwidth (if necessary), other network infrastructure and other storage licenses and hardware consistent with the Solution Qualification.
4. Install VMAX storage, SAN management servers, Cisco MDS FCIP switches. Connect Cisco MDS FCIP switches units to IP network.
5. Configure VMAX ready for SRDF usage with 5875 (or later) microcode. Install latest certified version of Solutions Enabler on management servers. Configure Cisco MDS FCIP switches.
6. The SRDF emulation must be enabled on all directors and then director ports need to be assigned to the SRDF emulation; thus, creating RAs.
7. Map RDF groups to RAs. Map applications to RDF groups. Map applications to devices and hence devices to RDF groups.
8. Create required RDF groups and populate with application devices.
9. Configure SRDF mode (SRDF/S or SRDF/A) and settings.

17.1.4.1 Collection of Basic Application Requirements

Before any design work can be undertaken for a specific SRDF/A implementation, it is vital to collect some basic information on the requirements of the target applications to understand the scope of the project ready for initial discussions with EMC.

Basic information for each application in scope should include, but is not limited to:

- Number of servers
- Amount of data to protect
- Recovery Point Objective⁹
- If I/O consistency at the DR site between servers that comprise the application is to be enforced by the storage infrastructure. If so, they need to be in the same replication group.
- If I/O consistency at the DR site between applications is to be enforced by the storage infrastructure. If so, multi-session consistency will need to be deployed.

One of the most crucial decisions that must be made is which applications require point-in-time consistency with each other at the target site. This critical decision determines:

- How many RDF groups will be needed
- If multi-session consistency will be needed

Hence, these decisions must be made very early in the planning process.

Enforcing consistency between applications at the storage subsystem level does have negative points and hence it is essential that is undertaken only where such consistency enforcement is genuinely required and where the application owners have been informed of the “binding” together of these applications, and the negative consequences (see Appendix C).

See Appendix C for a more in-depth discussion on choice of consistency models.

Another key decision that must be made is whether to use SRDF/A or SRDF/S solution. This will depend on factors such as distance, network bandwidth, RPO, etc.

17.1.4.2 Network Equipment Configuration

Cisco MDS FCIP switches should be configured for SRDF as per the following detailed documents:

- [Cisco MDS FCIP Solutions Guide](#)
- [WAN Throughput Optimization for SRDF/A with Cisco MDS Solutions Guide](#)

⁹ Recovery Point Objective is the point-in-time to which data must be restored to successfully resume processing (often thought of as time between last backup and when outage occurred).

17.1.4.3 SRDF Management server requirements

The SE management servers will be used to:

1. Configure and manage the SRDF/A environment
2. For controlling replication if Multi-Session Consistency (MSC) is implemented.

There should be a minimum of 2 management servers at the source and target sites. Only one is actually needed for SRDF control, but a second should be provided for redundancy.

If MSC is implemented, the management server service continuously controls replication and hence must be highly available. Implementation of high availability of the replication service is described in detail later in this document.

A minimum of 3 additional gatekeepers is required on each management server for SRDF configuration, control and reporting. If Multi-Session Consistency (MSC) is to be used, then one additional gatekeeper is needed for every three Composite Groups on the source site management server.

The following options should be configured in the **/var/symapi/config/options** file on the SE management servers managing SRDF arrays.

```
SYMAPI_USE_RDFD           = ENABLE
SYMAPI_ALLOW_RDF_SYMFORCE = TRUE
SYMAPI_WAIT_ON_LOCKED_GK  = ENABLE
```

SYMAPI_USE_RDFD controls if the RDF Daemon is used. This will need to be enabled for MSC control.

SYMAPI_ALLOW_RDF_SYMFORCE is FALSE by default, as it is dangerous. However, this needs to be set to TRUE in an MSC environment, as symforce is needed with Composite Group failovers following an MSC-triggered suspension of SRDF/A.

SYMAPI_WAIT_ON_LOCKED_GK should be enabled to avoid command failure when due to Gatekeepers not being available.

Also, if GNS is to be used for mirroring of device and composite group definitions (recommended), you will need the following additional entries:

```
SYMAPI_USE_GNS           = ENABLE
```

The following options should be configured in the **/var/syampi/config/daemon_options** file on the SE management servers managing SRDF arrays. It may be necessary to first copy the template file and then edit.

```
# cp /var/symapi/config/README.daemon_options /var/symapi/config/ daemon_options
```

The logfile_type entries are not essential but do help with debugging issues with the daemons.

```
storrdfd:logfile_type=dated
storapid:logfile_type=dated
storgnsd:logfile_type=dated
storwatchd:logfile_type=dated
```

Also, if GNS is to be used for remote mirroring of device and composite group definitions, you will need the following additional entries:

```
storgnsd:autorestart=enable
storgnsd:gns_remote_mirror=enable
```

17.1.4.4 RA Configuration

In order to configure RA ports for SRDF, the SRDF emulation must be enabled on the directors. The SRDF emulation can only be enabled on a directors by EMC. Once the SRDF emulation is enabled, Solutions Enabler can be used to associate ports to the SRDF emulation; thus, making them RA ports. Only 8Gbps FC ports can be used for RA ports at this time. A maximum of 16 ports can be associated to the SRDF emulation per director. As a Citi standard, no more than 2 ports on each director should be configured as RAs. RAs should be distributed evenly between directors and engines. For example: a two-engine system with 4 RAs configured will have 1 RA port configured on each of its 4 directors.

Once EMC has enabled the SRDF emulation on the directors, the following steps can be followed to associate front-end ports to the SRDF director emulation:

1. Identify the director IDs for the SRDF emulation:

```
symcfg -sid <sid> list -dir ALL
```

Sample Output:

```
# symcfg -sid 973 list -dir ALL
```

Symmetrix ID: 000196700973 (Local)

SYMMETRIX DIRECTORS						
Ident	Type	Engine	Cores	Ports	Status	
-----	-----	-----	-----	-----	-----	-----
IM-1A	IM	1	4	0	Online	
IM-2A	IM	1	4	0	Online	
IM-3A	IM	2	4	0	Online	
IM-4A	IM	2	4	0	Online	
ED-1B	EDS	1	8	0	Online	
ED-2B	EDS	1	9	0	Online	
ED-3B	EDS	2	9	0	Online	
ED-4B	EDS	2	8	0	Online	
DF-1C	DISK	1	8	4	Online	
DF-2C	DISK	1	9	4	Online	
DF-3C	DISK	2	9	4	Online	
DF-4C	DISK	2	8	4	Online	
FA-1D	FibreChannel	1	8	16	Online	
FA-2D	FibreChannel	1	6	15	Online	
FA-3D	FibreChannel	2	6	15	Online	
FA-4D	FibreChannel	2	8	16	Online	
RF-1E	RDF-BI-DIR	1	4	1	Online	
RF-2E	RDF-BI-DIR	1	4	1	Online	
RF-3E	RDF-BI-DIR	2	4	1	Online	
RF-4E	RDF-BI-DIR	2	4	1	Online	

2. Determine which ports are available to associate to the SRDF emulation:

symcfg -sid <sid> -port -free

Sample Output:

```
# symcfg list -sid 973 -port -free

Symmetrix ID: 000196700973

      Flags Speed
Slot Port FCIS Gb/sec Status
-----
  4  10 Y...    8 Powered

Legend:
Flags:
(F)C   : Y = Yes, . = No
F(C)OE : Y = Yes, . = No
F(I)CON : Y = Yes, . = No
I(S)CSI : Y = Yes, . = No
```

Note: The Slot number in the output above is the director number, not the slot on the director board.

3. A symconfigure command file will need to be created to associate a port to the SRDF director emulation. Below is the sample output of the symconfigure command which contains the command in the command file:

Sample Output:

```
# symconfigure -sid 973 commit -f associate_port.cmd -v

Execute a symconfigure operation for symmetrix '000196700973' (y/[n]) ? y

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

Establishing a configuration change session.....Established.
Processing symmetrix 000196700973
{
  associate port 10 to dir 4E;
}

Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Committing configuration changes.....Started.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

4. Confirm the ports have been associated with the SRDF director:

symcfg -sid <sid> list -dir <RDF dir ID> -port

Sample Output:

```
# symcfg -sid 973 list -dir 4e -port

Symmetrix ID: 000196700973 (Local)

      SYMMETRIX DIRECTOR PORTS

      Speed
Ident Port WWN      Type  Gb/sec Status
```

RF-4E	8	50000973500F34C8	RDF-BI-DIR	8	Online
RF-4E	10	50000973500F34CA	RDF-BI-DIR	8	PendOn

17.1.4.5 Mapping RDF Groups to RAs

There can be a maximum of 250 RDF groups per VMAX3 storage array. There can be a maximum of 250 RDF groups defined on a single port but the aggregate number of groups for the entire array cannot exceed 250.

Analysis of the SRDF requirements should be done in conjunction with EMC Shared Services to determine the expected application workload and how many RAs are needed for replication. Typically, there is one RDF group created for each application.

As a best practice, the same RA ID on the source and target array should be used when creating RDF groups for ease of management. If this is not possible, the following should be carefully documented:

- Which RA ports on the source and target array correspond with each other. For example, port 1e:8 on the source VMAX3 may not correspond to port 1e:8 on the SRDF/A target.
- Which RDF groups are to run over each pair of ports.

This information will be needed when the RDF groups are created. Each RDF group should be given a sensible label. RDF groups cannot be defined with leading zeros. In many implementations, the source RA and target RA may not be the equivalent port.

As new applications are added to a Symmetrix, they should be spread over different RA pairs to minimize the risk of a given pair becoming overloaded.

Dynamic RDF groups are created on-the-fly with the symrdf command:

```
symrdf addgrp -label <rdrg name> -rdrg <src rdrg #> -sid <src SID> -dir
<src dir:port> -remote_rdfg <tgt rdrg #> -remote_dir <tgt dir:port> -
remote_sid <tgt SID> -nop
```

Example:

```
# symrdf addgrp -label DYN1 -rdrg 1 -sid 0973 -dir 1e:8,2e:8,3e:8 -
remote_rdfg 1 -remote_dir 1e:8,2e:8,3e:8 -remote_sid 000196700974 -nop
```

Note: Must use 12 digit SID for remote array in the command

RA pairs can be added and removed from an RDF group with the symrdf command:

```
symrdf modifygrp {-add | -remove} -rdrg <src rdrg #> -sid <src sid> -dir
<src dir:port> -remote_dir <tgt dir:port>
```

Examples:

```
# symrdf modifygrp -add -rdrg 1 -sid 973 -dir 4e:8 -remote_dir 4e:8
# symrdf modifygrp -remove -rdrg 1 -sid 973 -dir 4e:8 -remote_dir 4e:8
```

The mapping between RDF groups and RA ports of the dynamic RDF groups can be shown using the symcfg command.

```
# symcfg list -ra all -switched -sid <SID>
```

17.1.4.6 Creating dynamic RDF device pairs

RDF device pairs can be created either by a pairs file or by storage groups.

Creating RDF Device Pairs with Pairs File

The device pairs file syntax contains two columns (R1 and R2 designated devices in the first and second columns respectively). Each SRDF pair must be listed on a separate line in the device file.

The creation of the dynamic SRDF pair utilizes the `createpair` option of the `symrdf` command. Similarly, the deletion of a dynamic SRDF pair utilizes the `deletepair` option of the `symrdf` command.

If metas are used, the pairs file should only include meta heads. In the example below, each meta is made up of 4 hypers. The first device in the first column has the meta head of 0157 and 3 additional hypers (0158, 0159, 015A) making up the volume. Only the meta head is included in the file.

Device file "pairsfile_DYN2"

```
0157 0157
015B 015B
015F 015F
0163 0163
0167 0167
016B 016B
016F 016F
etc.etc
```

In this example, the R1 and corresponding R2 have the same volume number. It is strongly recommended that all using Dynamic RDF be configured in this way.

The following pair creation example uses a pairfile called "pairsfile_DYN2", which only contains the devices to be added to RDF group2. The `-type rdf1` option indicates that the devices in the first column (which are in the Symm specified by the `-sid` option) are R1 devices. If GNS is use with RDF group mirroring enabled, then the newly created device group will be mirrored to the target site (in a 2-site configuration only) and to the other local SAN management servers.

The following command invalidates the R2, merges track table, brings up the RDF links and starts the copy process from R1 to R2.

```
# symrdf -file pairsfile_DYN2 -sid 1380 -rdfg 2 createpair -type rdf1
      -establish -g DG2
```

The `-g` option creates an `rdf1`-type device group called DG2 that contains the devices in the pairs file and which can be used subsequently for SRDF control. Alternatively, local tools can be used to create the device group if preferred.

Full establish throughput is less for a fully allocated thin devices compared with thick devices. However, for newly created thin devices (with no pre-allocation), full establish synchronization is very quick as there is no data to copy (e.g. 4*33GB LUNs reached Synchronized state is less than

2 minutes). This is a notable one-off benefit with using thin provisioned rather than standard storage.

Creating RDF Device Pairs with Storage Groups

The following steps can be used to create RDF device pairs using storage groups if the RDF devices have not been created:

1. Create empty storage groups on the source and target arrays using the symmsg command:

```
# symmsg -sid 973 create qtest-rdf1_sg
# symmsg -sid 974 create qtest-rdf2_sg
```

2. Create and add devices to the SRDF Storage Groups using symconfigure. The following is the symconfigure command file entry:

```
create dev count=10, size=34GB, emulation=FBA, config=RDF1+TDEV,
remote_config=RDF2+TDEV, ra_group=1, sg=qtest-rdf1_sg,
remote_sg=qtest-rdf2_sg;
```

If the RDF devices have already been created on the source and target, use the following steps:

1. Create the source and target SGs and add the devices to the storage groups.
2. Create the SRDF device pair with the symrdf command

```
symrdf createpair -sid 123 -rdfg 456 -sg localSG -type r1 -remote_sg
remoteSG
```

17.1.4.7 Deleting RDF Device Pairs

Storage groups cannot be used to delete device pairs so a device file will be required. If a device file was used to create the device pairs, the same device file can be used to delete the device pairs. To delete the pairing set up in the previous section, one would issue the command

```
# symrdf -file pairsfile_DYN2 -sid 1380 -rdfg 2 deletepair
```

Note that the symrdf deletepair command is only allowable in one of the following states, where the links are not ready status:

- Suspended
- Split
- Failed over

17.1.4.8 Checking the status of dynamic RDF device pairs

There are two commands for querying the dynamic RDF device pair status.

```
# symdev -sid <sid> list -dynamic
```

will show the configuration of the dynamic RDF devices.

- Devices marked as either RDF1+TDEV or RDF2+TDEV have SRDF enabled.

```
# symrdf -sid <sid> list -dynamic
```

will show the pair state of the dynamic RDF devices (Synchronized, Failed Over etc).

Note: that 2-Way mirror devices are not shown with the symrdf command.

17.1.4.9 Deleting RDF Groups

Dynamic RDF groups can be deleted with the command of the form

```
# symrdf removegrp -sid 0973 -rdfg 2
```

Note: All RDF device pairs must be removed from the RDF group before it can be deleted

17.1.4.10 Enabling SRDF/A replication

RDF groups are by default created in synchronous mode. The replication mode must be explicitly set to asynchronous. The `-rdfg <RDFG num>` is optional for a 2-site SRDF/A configuration.

The enable command enables consistency between devices within the device group.

```
# symrdf -g <DGname> -rdfg <RDFG num> set mode async -nop
# symrdf -g <DGname> -rdfg <RDFG num> enable
```

Confirm that SRDF/A is running successfully and that consistency is enabled.

```
# symrdf -g <DGname> -rdfg <RDFG num> query
```

Under the MODES section, the M column for each device pair should show A (for async), and the “RDF Pair STATE” should be Consistent. Note that it may take up to two cycles before it appears as “Consistent.”

For example

```
# symrdf -g DG19 -rdfg 19 query
```

```
Device Group (DG) Name      : DG19
DG's Type                   : RDF1
DG's Symmetrix ID           : 000187701380
Remote Symmetrix ID         : 000187721497
RDF (RA) Group Number       : 19 (12)
```

Source (R1) View					Target (R2) View				MODES	
-----					-----				-----	
	ST				LI	ST				
Standard	A				N	A				
Logical	T	R1 Inv	R2 Inv	K	T	R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MDA
-----					-----				-----	
DEV001	00BD	RW	0	0	RW	00BD	WD	0	0	A..
DEV002	00C1	RW	0	0	RW	00C1	WD	0	0	A..
-----					-----				-----	
Total										
Track(s)			0	0				0	0	
MB(s)			0.0	0.0				0.0	0.0	

Legend for MODES:

M(mode of Operation): A = Async, S = Sync, E = Semi-sync, C = Adaptive Copy

```
D(omino)           : X = Enabled, . = Disabled
A(daptive Copy)    : D = Disk Mode, W = WP Mode, . = ACp off
```

Also, the command

```
# symstat -type cycle -retype rdfa -rdfg <RDFG num>
```

will show a non-zero “Last Switch” entry and the “Active Cycle #” will be increasing.

17.1.4.11 SRDF/A Cycle Time Configuration

The default SRDF/A cycle time for VMAX3 is 15 seconds and should not be changed unless absolutely necessary due to RDF bandwidth constraints. The SRDF cycle time is changed at the RDF group level with the `symrdf` command:

```
symrdf -sid <src sid> -rdfg <rdfg #> set rdfa -cycle_time <1-60>
```

Example:

```
# symrdf -sid 973 -rdfg 1 set rdfa -cycle_time 30
```

17.1.4.12 SRDF/A Delta Set Extension (DSE) Configuration

SRDF/A uses cache to buffer host I/Os and shield applications from the remote replication process in the array. These cache resources can be fully consumed if there are unplanned spikes in workloads or reduced bandwidth in the backend network pipeline, therefore causing SRDF/A to either drop or throttle.

With Delta Set Expansion (DSE), rather than storing all pending data in SRDF/A cycles in cache (Global Memory), this feature allocates a portion of disk storage in an existing SRP as temporary storage location. Since Citi standard is to only have one SRP configured on the storage array, the DSE space will come from the same SRP. This allows SRDF/A to be more resilient and flexible to handle the situations where host application throughput outpaces the RDF throughput.

Note that the function of the SRDF/A DSE feature is to temporarily extend the amount of cache space for SRDF/A by offloading some or all of the cycle data to free space in the SRP within the VMAX, to allow for I/O peaks. It will not help if the SRDF/A drop is a result of unbalanced configurations.

DSE cannot be used as a method to avoid properly assessing a workload and sizing the required network. DSE can only be used to deal with temporary workload peaks and network outages. It cannot make SRDF/A work with an undersized network i.e. one where the effective network bandwidth is less than the average SRDF/A workload.

The maximum amount of space in the SRP that can be used by DSE can be specified (default is unlimited). The main factors when sizing the DSE maximum capacity is:

- How much spare disk space you have in the SRP. This is the simplest approach, if you do not have a specific network outage duration that you want to be able to survive.
- Host workload. It will be necessary to capture the application host write profiles for time intervals that include the periods of heaviest write activity to establish base DSE sizing requirements.

- How long a network outage, or workload peak, that you want to be able to survive without SRDF/A dropping.

Note also that some “spare” network capacity will be required to allow “catch up” after an outage or peak. The more spare capacity there is, the quicker the system will return to normal RPO.

If you intend to use DSE, EMC Shared Services should recommend a DSE maximum capacity value for your workload and outage survival requirements as part of the sizing process. If you have WLA data for the workload with which DSE is to be used, EMC can use their Symmerge tool to size the DSE pool for you (Symmerge refers to this as “usable paging disk space”).

The DSE max capacity can be set using the following symconfigure command (the value is in GBs):

```
# symconfigure -sid <sid> commit -cmd "set symmetrix dse_max_cap = <1-100000>;"
```

The DSE max capacity value can be check by using the symcfg command:

Sample Output:

```
$ symcfg -sid 973 list -v

...
SRDF/A DSE Maximum Capacity (GB) : 100

Auto Meta : N/A
Minimum Auto Meta Size : N/A
Auto Meta Member Size : N/A
Auto Meta Configuration : N/A
```

17.1.4.13 Checking DSE Status

There are unfortunately many different commands for checking the status of DSE, each showing slightly different information. These are summarized below and examples of output are also shown.

Interactive DSE Checking

```
# symrdf -g <DGname> query -rdfa
```

This useful report shows %cache used and DSE tracks used (not percentages) by this RDF group, as well as other standard query output. It also shows the amount of time that Transmit Idle has been invoked for.

```
# symcfg list -rdfg <RDFG number> -sid <SymmID> -rdfa
```

This report shows flag information for specified RDF group, such as if DSE is active, if autostart is set and DSE threshold. It also shows if Transmit Idle is enabled for the group.

```
# symstat -type cache -reptype rdfa -sid <SymmID> -I <seconds>
```

This useful report shows SRDF/A cache and DSE usage for **all active SRDF/A groups**. The rdfa_dse threshold for each group is also shown.

```
# symcfg list -rdfg all -sid <SymmID> -rdfa
```

This useful report shows flag information for all RDF groups, such as if DSE is active, if autostart is set and DSE threshold. It also shows if Transmit Idle is enabled for each group.

17.1.4.14 SRDF/A Cache Utilization

During the SRDF sizing process, EMC Shared Services should recommend how much cache is required for the SRDF/A caching.

The maximum cache usage percentage represents the fraction of user cache that SRDF/A is allowed to use if the need arises. The default setting is 75% for VMAX3. It should be left at the 75% unless specified otherwise by EMC.

The following example shows how to change the SRDF/A cache percentage. In this example, the percentage is changed to 75%.

The “SRDF/A Maximum Cache Usage” value can be checked using the symcfg command:

Sample Output:

```
# symcfg -sid 973 list -v
.....
SRDF/A Maximum Cache Usage (Percent) : 75
```

The “SRDF/A Maximum Cache Usage” value can be changed using symconfigure with the following line in the symconfigure file:

```
set Symmetrix rdfa_cache_percent=<0-100>;
```

17.1.4.15 SRDF/A Transmit Idle Configuration

Transmit Idle allows SRDF/A to seamlessly ride out a transient link failure. This feature provides an extra level of protection from sessions dropping when the RDF link cannot transmit data. If the link recovers, SRDF/A can pick up operations where it left off.

With this feature enabled, if the link fails, the RDF pair state will change to `TransIdle` after the link limbo time expires (instead of dropping). When transmit idle is enabled for an SRDF/A group, the system will indicate when the `TransIdle` state is active for a specific RDF group. SRDF/A Transmit Idle support must be enabled on both the local and remote side for each RDF group.

SRDF/A Transmit Idle is not a panacea for all causes of SRDF/A drops. It is targeted to address only short link outages when the Symmetrix cache is not constrained. It is important to note that SRDF/A Transmit Idle does not replace SRDF Link Limbo. SRDF/A Transmit Idle will only invoke after the SRDF Link Limbo timer expires. SRDF/A Transmit Idle is of limited benefit if the Symmetrix system is cache constrained while running SRDF/A.

SRDF/A Transmit Idle cannot help you if the SRDF/A environment is not properly balanced between the frames and across the SRDF links. The following SRDF/A component configurations must be balanced:

- SAN control hosts must be running a minimum of SE 8.0 for VMAX3
- Symmetrix frames in use on each side of the link

- Cache sizes and allocation to SRDF/A
- RDF link bandwidth sized appropriately to the requirements of the applications and data being replicated
- Drive counts and sizes, and data protection schemes in use

SRDF/A Transmit Idle does not take the place of ensuring that the SRDF network between the sites is of sufficient quality to meet an application's replication objectives.

By default Transmit Idle is automatically enabled when an RDF group is created so no configuration is required. If Transmit Idle is not enabled (e.g. it has been disabled, or Transmit idle has been enabled whilst SRDF/A is already running), then it should be manually enabled for each RDF symcfg list -ra group on both the source and target symmetrix. For example:

```
# symrdf -sid 973 -rdfg 1 set rdffa -transmit_idle on -both_sides
```

The Link Limbo Timer (LLT) should be set to the default 10 seconds when Transmit Idle is used. If LLT is set to > 10 seconds on an existing implementation (not likely with VMAX), it should be changed back to 10 seconds when SRDF/A transmit idle is enabled.

```
# symrdf -sid 973 -rdfg 1 set rdffa -limbo 10 -both_sides
```

17.1.4.16 Checking Transmit idle status

Interactive Checking

```
# symcfg -sid <sid> list -rdfg <rdfg #> -v
```

Provides link limbo timer value.

```
# symrdf -g <DGname> query -rdffa
```

This useful report shows if Transmit Idle has been invoked (see RDF Pair State). It also shows the amount of time that Transmit Idle has been invoked for.

```
# symcfg list -rdfg all -sid <SymmID> -rdffa
```

This useful report shows flag information for all RDF groups, including if Transmit Idle is enabled for each group.

17.1.5 SRDF Replication Monitoring

There are multiple methods for SRDF replication monitoring/alerting with VMAX:

- Via storevntd log file and Tivoli Endpoint Agent (Recommended solution for VMAX)
- Via Citi scripts (legacy or backup solution)
- Via Unisphere for VMAX for monitoring (Complimentary to above)

17.1.5.1 Fault alerting via storevntd log file and Tivoli Endpoint Agent

The Symmetrix Storage Event daemon (storevntd) should be used to monitor SRDF-related events. The following events can be monitored for (inexhaustive list):

- SRDF/A session
 - SRDF/A session inactive
 - SRDF/A session dropped; write pending limit reached. Host throttling is disabled
 - SRDF/A session dropped: write pending limit reached. Host throttling is enabled
 - SRDF/A session dropped: device not ready. Tolerance mode is off
 - SRDF/A session dropped: device not ready through consistency group
 - SRDF/A dropped: no SRDF links operational
 - SRDF/A dropped: timeout in MSC mode
 - R1 To R2 Lag Time
- SRDF CG
 - RDF CG trip event triggered
- SRDF Link
 - No SRDF links in an RDF group are operational
 - Single RDF link in an RDF group is not operational

See Configuring Symmetrix Alerting for instructions for setting up Tivoli Monitoring.

17.1.5.2 Monitoring with scripts (legacy or backup solution)

This section describes how the following key SRDF/A status and metric information can be monitored and alerts generated under exceptional circumstance:

- 1) Monitoring for inactive SRDF sessions
- 2) Monitoring for cache utilization above a preset threshold
- 3) Monitoring for cycle times above a preset threshold

The first three of the above tasks can be undertaken by running the script **srdfa_health** (see Appendix B for more detail) from cron at regular intervals on a SE server being used to manage SRDF/A. Individual scripts for each of the above monitoring tasks are also available (**mon_srdfa_inactive**, **mon_srdfa_cache**, **mon_srdfa_cycle** – see Appendix D – Template Operation Scripts Summary for further details).

All of the above scripts will write a message to /var/adm/messages using the logger command if one of their criteria for raising an alert is met. The message will of type local0.error be of the form:

```
Aug  2 06:24:07 sarprod2-ny SRDFA_ALERT: [ID 702911 local0.error] Please contact
SAN Support team. Current cycle time for RDF group 25 (66 sec) on Symmetrix 1380
has exceeded acceptable threshold (60 sec)
```

All messages will include the string SRDFA_ALERT. You will need to instruct your local Enterprise Monitoring (EMS) Team to monitor the /var/adm/messages file on the server running the SRDF/A monitoring script(s) and set their system to raise an alert to the appropriate support team should a message with the string SRDFA_ALERT be detected. The contents of the message should also be forwarded by EMS to the support team.

17.1.5.3 Monitoring via Unisphere for VMAX

Amongst many other events, Unisphere for VMAX will capture certain SRDF related events, which can be viewed via its web interface; this can be helpful to correlate issues logged via alerting systems.

See the Unisphere Solutions Guide for further details on capturing alerts and EMC's Unisphere documentation.

SRDF Alerts that can be captured are:

- SRDF Alerts
- SRDF Link Status

- SRDF/A Session
- SRDF/A Session entering transmit idle state
- SRDF/A Session dropped, transmit idle state timeout
- SRDF/A Session recovered from a transmit idle state

The following screenshot shows where the standard canned SRDF alerts are set in Unisphere.

Administration → Alert Settings → Alert Policies

Symmetrix ID 1 ▲	Type	Policy Name	2 ▲	State
000192604261	Array	Event Lost Alert		disabled
000192604261	Array	Event Overflow Alert		disabled
000192604261	Array	FAST Controller switched state		disabled
000192604261	Array	FAST FTS Performance		enabled
000192604261	Array	GK Timeout		disabled
000192604261	Array	GK Utilization		disabled
000192604261	Array	Hotspare Invoked		disabled
000192604261	Array	Migration Complete Alert		disabled
000192604261	Array	Optimizer switched mode		disabled
000192604261	Array	Port Link Status		disabled
000192604261	Array	Port Status		disabled
000192604261	Array	RVA Spare Coverage		disabled
000192604261	Array	SP Alerts		disabled
000192604261	Array	SRDF Alerts		disabled
000192604261	Array	SRDF Link Status		disabled
000192604261	Array	SRDF/A No Cycle Switch Alert		disabled
000192604261	Array	SRDF/A Session		disabled
000192604261	Array	SRDF/A Session dropped, tran...		disabled
000192604261	Array	SRDF/A Session entering trans...		disabled
000192604261	Array	SRDF/A Session recovered fro...		disabled
000192604261	Array	Thin Device Allocation		disabled
000192604261	Array	Thin Device Usage		disabled
000192604261	Array	Thin Pool Rebalancing Complet...		disabled
000192604261	Array	User approval required for Opti...		disabled

Enable Disable Notify

Bespoke Alerts/Events can also be set up: Administration → Alert Settings → Performance Thresholds & Alerts

The screenshot shows the 'Performance Thresholds and Alerts' configuration page in the Symmetrix management interface. The left sidebar contains a tree view of categories: System, Storage, Data Protection, FAST, and Hosts. Under 'System', the 'RDFA Director' category is selected. The main area displays a table of metrics with columns for Metric, Custom, First Threshold, and Second Threshold. The '% Busy' metric is highlighted, showing a first threshold of 50 and a second threshold of 70. Other metrics listed include Avg IO Service Time (ms), Copy IOs/sec, Copy MBs/sec, IOs/sec, MBs Received/sec, MBs Sent and Received/sec, MBs Sent/sec, Read Reqs/sec, Reqs/sec, Rewrites/sec, and Write Reqs/sec. At the bottom, there are buttons for 'Create', 'Edit', 'Delete', and 'Notifications'.

Select desired category, Instance, metric, values and sampling

The 'Edit Performance Threshold and Alert' dialog box is shown. It contains the following fields and settings:

- Symmetrix ID:** 000192604261
- Category:** RDFA Group
- Instance:** *
- Metric:** Avg Cycle Time
- First Threshold:**
 - Value: 15
 - Enable Alert: ☒
 - Severity: Information
 - Occurrence: 2 / 5 samples
- Second Threshold:**
 - Value: 30
 - Enable Alert: ☒
 - Severity: Warning
 - Occurrence: 2 / 5 samples

Buttons at the bottom: OK, Cancel, Help.

Events are captured & displayed, and can be filtered on rdf group, SID, etc.

All Symmetrix > Home > Alerts

Alerts

State	Severity	Symmetrix	Object 1	Description	Created
NEW	INFORMATION	000192604261	79	SRDF/A session active. - Object is: 000192604261:79	Fri Oct 25 16:57:07 GMT+0100 2013
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Fri Oct 25 16:57:07 GMT+0100 2013
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Fri Oct 25 16:49:48 GMT+0100 2013
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Fri Oct 25 16:41:14 GMT+0100 2013
NEW	INFORMATION	000192604261	79	SRDF/A session active. - Object is: 000192604261:79	Fri Oct 25 16:31:43 GMT+0100 2013
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Fri Oct 25 16:29:00 GMT+0100 2013
NEW	INFORMATION	000192604261	79	SRDF/A session active. - Object is: 000192604261:79	Thu Oct 24 16:51:46 GMT+0100 2...
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Thu Oct 24 16:51:46 GMT+0100 2...
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Thu Oct 24 16:45:16 GMT+0100 2...
NEW	WARNING	000192604261	79	Average cycle time (74.0) exceeds threshold (30.0)	Thu Oct 24 16:40:00 GMT+0100 2...
NEW	INFORMATION	000192604261	79	Average cycle time (74.0) exceeds threshold (15.0)	Thu Oct 24 16:40:00 GMT+0100 2...
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Thu Oct 24 16:38:10 GMT+0100 2...
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Thu Oct 24 16:29:58 GMT+0100 2...
NEW	INFORMATION	000192604261	79	SRDF/A session active. - Object is: 000192604261:79	Thu Oct 24 16:29:35 GMT+0100 2...
NEW	CRITICAL	000192604261	79	SRDF/A session drop at cycle boundary requested [host software initiated]. - Obj...	Thu Oct 24 16:22:09 GMT+0100 2...
NEW	INFORMATION	000192604261	79	SRDF/A session active. - Object is: 000192604261:79	Fri Oct 4 10:48:53 GMT+0100 2013
NEW	INFORMATION	000192604261	79	Single SRDF link in an RDF group is now operational. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	INFORMATION	000192604261	79	SRDF/A Session recovered from a transmit idle state. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	INFORMATION	000192604261	79	SRDF/A Session recovered from a transmit idle state. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	INFORMATION	000192604261	79	All SRDF links in an RDF group are operational. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	INFORMATION	000192604261	79	SRDF/A Session entering transmit idle state. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	INFORMATION	000192604261	79	SRDF/A Session entering transmit idle state. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	WARNING	000192604261	79	No SRDF links in an RDF group are operational. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013
NEW	WARNING	000192604261	79	Single SRDF link in an RDF group is not operational. - Object is: 000192604261:79	Fri Oct 4 10:34:04 GMT+0100 2013

17.1.6 Target site Point-In-Time copy for SRDF/A Resynchronization

During a resynchronization, the data on the R2 will be inconsistent and hence during this period there is no valid COB copy of the data at the remote site, unless a PiT copy is first taken. SnapVX should be used to take a point-in-time (PiT) copy (i.e. snapshot) of the SRDF/A target site (R2 devices) data prior to a resynchronization to ensure that there is a consistent copy of the data at the remote site at all times. The SRDF/A supported hardware configurations in [9.4.2 VMAX3 Supported SRDF/A Configurations](#) are sized to accommodate the snapshot space needed to withstand the duration of an SRDF/A failure.

Hence, the snapshots are not required during normal SRDF/A replication operations, but are there as a safeguard during resynchronization operations.

17.1.6.1 Are R2 Point-in-time copies mandatory?

R2 point-in-time copies are not required in SRDF/S environments as SRDF resynchronization virtually never takes place because:

- 1) The metro-distance networks links between the local and remote Symmetrix are very reliable.

- 2) SRDF/S does not usually drop when there are workload and/or link bandwidth issues (it throttles the host writes instead).

However, with SRDF/A, it is recommended that deployments should implement R2 point-in-time copy protection. Without the PiT copy you will be vulnerable to complete data loss should your R1 copy be lost during a resynchronization exercise that might occur due to:

1. a network failure
2. a scheduled resynchronization associated with a COB/DR test
3. The current capacity of the network link being exceeded due to growth in utilization and improper capacity planning
4. The capacity of the network link being exceeded temporarily (but sufficiently long to cause SRDF/A to drop) due to a workload peak.

17.1.6.2 Recovery Process Using SnapVX Snapshots as the Gold Copy

The following is the process for recovering from primary site failure on the R2 side using SnapVX snapshots as the gold copy during resynchronization.

- 1) Once a production site failure has been detected and the production application has stopped, a recovery on the R2 side is needed. Immediately take a snapshot of the R2 LUNs. In this example, a snapshot of a storage group containing the R2 LUNs for a specific application is taken:

```
# symsnapvx -sid 973 -sg qtest-rdf1_sg -name qtest-rdf1_snapshot establish
```

If another snapshot is taken using the same snapshot name, a new “generation” of the snapshot will be created which means it’s just another PIT copy of the data with the same name. The generations are numbered starting at 0 with the newest snapshot being 0. The older snapshots will increment their generation value as new snapshots are created. The following output lists all the snapshots (and generations) for a storage group:

```
# symsnapvx -sid 973 -sg qtest-rdf1_sg list -detail
```

Storage Group (SG) Name : qtest-rdf1_sg									
SG's Symmetrix ID : 000196700973 (Microcode Version: 5977)									
		Total		Non-Shared					
Sym	Snapshot Name	Flgs	Gen	FLRG	Deltas	Snapshot	Timestamp	(Track	s) (Tracks) Expiration Date
000AE	qtest-rdf1_snapshot		0		Fri Jun 26 16:19:32	2015	27	8432	75942 Sat Jun 27 16:19:32
2015	qtest-rdf1_snapshot		1		Fri Jun 26 16:17:59	2015	27	8441	68786 Sat Jun 27 16:43:35
2015	qtest-rdf1_snapshot		2		Fri Jun 26 16:07:33	2015	27	8495	251940 NA
000AF	qtest-rdf1_snapshot		0		Fri Jun 26 16:19:32	2015	27	8446	75457 Sat Jun 27 16:19:32
2015	qtest-rdf1_snapshot		1		Fri Jun 26 16:17:59	2015	27	8459	68378 Sat Jun 27 16:43:35
2015	qtest-rdf1_snapshot		2		Fri Jun 26 16:07:33	2015	27	8500	251981 NA
000B0	qtest-rdf1_snapshot		0		Fri Jun 26 16:19:32	2015	27	8441	75568 Sat Jun 27 16:19:32
2015	qtest-rdf1_snapshot		1		Fri Jun 26 16:17:59	2015	27	8456	68342 Sat Jun 27 16:43:35
2015	qtest-rdf1_snapshot		2		Fri Jun 26 16:07:33	2015	27	8498	251545 NA
000B1	qtest-rdf1_snapshot		0		Fri Jun 26 16:19:32	2015	27	8449	75783 Sat Jun 27 16:19:32
2015	qtest-rdf1_snapshot		1		Fri Jun 26 16:17:59	2015	27	8465	68581 Sat Jun 27 16:43:35
2015	qtest-rdf1_snapshot		2		Fri Jun 26 16:07:33	2015	27	8499	251786 NA

000B2 qtest-rdf1_snapshot 2015	0 Fri Jun 26 16:19:32 2015	27	8433	75435	Sat Jun 27 16:19:32
qtest-rdf1_snapshot 2015	1 Fri Jun 26 16:17:59 2015	27	8465	68062	Sat Jun 27 16:43:35
qtest-rdf1_snapshot 000B3 qtest-rdf1_snapshot 2015	2 Fri Jun 26 16:07:33 2015	27	8498	251579	NA
qtest-rdf1_snapshot 2015	0 Fri Jun 26 16:19:32 2015	27	8442	75686	Sat Jun 27 16:19:32
qtest-rdf1_snapshot 2015	1 Fri Jun 26 16:17:59 2015	27	8456	68476	Sat Jun 27 16:43:35
qtest-rdf1_snapshot 000B4 qtest-rdf1_snapshot 2015	2 Fri Jun 26 16:07:33 2015	27	8497	251783	NA
qtest-rdf1_snapshot 2015	0 Fri Jun 26 16:19:32 2015	27	8447	75296	Sat Jun 27 16:19:32
qtest-rdf1_snapshot 2015	1 Fri Jun 26 16:17:59 2015	27	8455	68013	Sat Jun 27 16:43:35
qtest-rdf1_snapshot 000B5 qtest-rdf1_snapshot 2015	2 Fri Jun 26 16:07:33 2015	27	8499	251372	NA
qtest-rdf1_snapshot 2015	0 Fri Jun 26 16:19:32 2015	27	8427	75410	Sat Jun 27 16:19:32
qtest-rdf1_snapshot 2015	1 Fri Jun 26 16:17:59 2015	27	8447	68220	Sat Jun 27 16:43:35
qtest-rdf1_snapshot 000B6 qtest-rdf1_snapshot 2015	2 Fri Jun 26 16:07:33 2015	27	8496	251551	NA
qtest-rdf1_snapshot 2015	0 Fri Jun 26 16:19:32 2015	27	8437	75516	Sat Jun 27 16:19:32
qtest-rdf1_snapshot 2015	1 Fri Jun 26 16:17:59 2015	27	8457	68013	Sat Jun 27 16:43:35
qtest-rdf1_snapshot 000B7 qtest-rdf1_snapshot 2015	2 Fri Jun 26 16:07:33 2015	27	8500	251478	NA
qtest-rdf1_snapshot 2015	0 Fri Jun 26 16:19:32 2015	27	8437	75314	Sat Jun 27 16:19:32
qtest-rdf1_snapshot 2015	1 Fri Jun 26 16:17:59 2015	27	8451	68132	Sat Jun 27 16:43:35
qtest-rdf1_snapshot	2 Fri Jun 26 16:07:33 2015	27	8495	251729	NA

		835	3920	3955154	

Flgs:

(F)ailed : X = Failed, . = No Failure
 (L)ink : X = Link Exists, . = No Link Exists
 (R)estore : X = Restore Active, . = No Restore Active
 (G)CM : X = GCM, . = Non-GCM

A time-to-live (TTL) valute can be set for the snapshot. Be sure the value is high enough to allow ample time for the application to fully recover. The TTL values are set to 1 day below since they are just for example.

Set ttl to 1 day at snapshot creation:

```
# symsnapvx -sid 973 -sg qtest-rdf1_sg -name qtest-rdf1_snapshot establish -ttl -delta 1
```

Set ttl to 1 day for a specific generation snapshot:

```
# symsnapvx -sid 973 -sg qtest-rdf1_sg -snapshot_name qtest-rdf1_snapshot -generation 1  
set ttl -delta 1
```

Once a snapshot of the R2 devices has been established, it will be considered the gold copy of the data before the failure until the application is fully recovered.

- 2) After a snapshot of the R2 devices is established, replication between Prod/Cob is failed over & swapped so that the R1 device now appears to reside on the COB array, and R2 on the Production array.
- 3) SA and application owners start the application on the COB host.
- 4) If there is a need to restore the snapshot to the original volumes due to some corruption during the recovery process, that can be done with the following command:

```
# symsnapvx -sid 973 -sg qtest-rdf1_sg -snapshot_name qtest-rdf1_snapshot -gen 0 restore
```

To verify that the snapshot was restored successfully:

```
# symsnapvx -sid 973 -sg qtest-rdf1_sg list -summary
```

Storage Group (SG) Name : qtest-rdf1_sg		
Snapshot State	Count	
-----	-----	
Restored	10	
RestoreInProgress	0	
Established	40	
EstablishInProgress	0	
Linked	0	
CopyInProgress	0	
Copied	0	
NoSnapshot	0	
Failed	0	
-----	-----	
Total	50	
Track(s)		

Total Remaining	0	

- 5) Once the application is fully recovered on the COB side and the SRDF/A links and/or production array is fully recovered, replication between COB and production can now be re-established. The initial resynchronization of the RDF devices should be done in Adaptive Copy mode to avoid high demand on the storage array cache if there is a large amount of invalid tracks. At this point, the personality of the production array and COB array should still be swapped (COB volumes appear as R1) so application updates on the COB array are being replicated to the production array.
- 6) Once replication is synchronized, the snapshots used for the gold copy for the original failure can be deleted or allowed to expire if a TTL was set and SRDF mode changed from Adaptive Copy to Asynchronous.
- 7) Before reverting the applications back to the production array, the applications should be stopped and another snapshot of the LUNs on the COB array should be created as a gold copy.
- 8) Replication between prod/COB is failed-over and then swapped so that the R1 device appears back on the production array and the R2 on the COB array.
- 9) SA and application owners start their applications on the production array.

Note: In a configuration where two production arrays are replicating to each other, it is recommended that the SRDF/A configurations sized for snapshot space is used at both sites so that both arrays have the capability to provide PiT copies of data (gold copies) during SRDF resynchronization following a failure.

17.1.6.3 Recovery Process Using BCVs (Legacy) as the Gold Copy

BCVs have been used for PiT copies for SRDF/A configurations for Symmetrix arrays prior to VMAX3. VMAX3 has the capability to emulate BCVs utilizing the same SE commands to control and manage BCVs. Though BCVs are much more expensive than SnapVX snapshots since BCVs are full copies of the source device and snapshots only consume space to store the original data for changed data. BCVs are considered legacy and SnapVX is the recommended technology. This section is included for cases where BCVs are still desired.

Note: The COB array should be properly sized to accommodate BCV space. The SRDF/A standard configurations are only sized with capacity that will accommodate SnapVX snapshot space, not BCVs.

The following steps should be followed to setup BCVs on the Cob array in an SRDF/A configuration. In the examples below, 973 is the production array and 974 is the COB array:

- 1) From the production SE management server create a disk group for containing the R1 source devices:

```
# symdg create -type RDF1 test-rdf1_dg
# symdg -g test-rdf1_dg -sid 973 addall -dev ae:b7
```

- 2) On the management server at the source site, the BCVs in the target VMAX should be associated as remote BCVs.

```
# symbcv -g test-rdf1_dg -sid 973 associateall -devs 008b:0094 -rdf
-rdfg 1
# symmir -g test-rdf1_dg establish -full -rdf -nop
```

Verify that the remote BCVs have been associated:

```
# symdg show test-rdf1_dg
```

```
Group Name: test-rdf1_dg
```

```

Group Type           : RDF1   (RDFS)
Device Group in GNS  : Yes
Valid                : Yes
Symmetrix ID         : 000196700973
Group Creation Time   : Wed Dec 31 19:00:00 1969
Vendor ID            : EMC Corp
Application ID        : N/A

```

```

Number of STD Devices in Group      : 10
Number of Associated GK's           : 0
Number of Locally-associated BCV's  : 0
Number of Locally-associated VDEV's : 0
Number of Locally-associated TGT's  : 0
Number of Remotely-associated VDEV's(STD RDF): 0
Number of Remotely-associated BCV's (STD RDF): 10
Number of Remotely-associated TGT's(TGT RDF) : 0
Number of Remotely-associated BCV's (BCV RDF): 0
Number of Remotely-assoc'd RBCV's (RBCV RDF) : 0
Number of Remotely-assoc'd BCV's (Hop-2 BCV) : 0
Number of Remotely-assoc'd VDEV's(Hop-2 VDEV): 0
Number of Remotely-assoc'd TGT's (Hop-2 TGT) : 0
Number of Composite Groups          : 0
Composite Group Names               : N/A

```

```
Standard (STD) Devices (10):
```

```
{
-----

```

LdevName	PdevName	Sym Device Dev Config	Att. S	Cap ts (MB)
DEV001	N/A	000AE RDF1+TDEV	R	W 34817
DEV002	N/A	000AF RDF1+TDEV	R	W 34817
DEV003	N/A	000B0 RDF1+TDEV	R	W 34817
DEV004	N/A	000B1 RDF1+TDEV	R	W 34817
DEV005	N/A	000B2 RDF1+TDEV	R	W 34817
DEV006	N/A	000B3 RDF1+TDEV	R	W 34817
DEV007	N/A	000B4 RDF1+TDEV	R	W 34817
DEV008	N/A	000B5 RDF1+TDEV	R	W 34817
DEV009	N/A	000B6 RDF1+TDEV	R	W 34817
DEV010	N/A	000B7 RDF1+TDEV	R	W 34817

```

}

```

```
BCV Devices Remotely-associated (STD RDF) (10):
```

```
{
-----

```

LdevName	PdevName	Sym Device Dev Config	Att. S	Cap ts (MB)
RBCV001	N/A	0008B BCV+TDEV	N	R 34817
RBCV002	N/A	0008C BCV+TDEV	N	R 34817
RBCV003	N/A	0008D BCV+TDEV	N	R 34817
RBCV004	N/A	0008E BCV+TDEV	N	R 34817
RBCV005	N/A	0008F BCV+TDEV	N	R 34817
RBCV006	N/A	00090 BCV+TDEV	N	R 34817

```

}

```

RBCV007	N/A	00091 BCV+TDEV	N	R	34817
RBCV008	N/A	00092 BCV+TDEV	N	R	34817
RBCV009	N/A	00093 BCV+TDEV	N	R	34817
RBCV010	N/A	00094 BCV+TDEV	N	R	34817
}					

- 3) On the management server at the target site, the BCVs in the target VMAX should be associated as local BCVs. If GNS remote mirror replication is in use, the device group definition changes made at the source site will be replicated to the remote site and the BCV type changed to local automatically.

The BCVs can now be managed using symmir commands:

```
# symmir -g test-rdf1_dg split -rdf -nop
# symmir -g test-rdf1_dg establish -rdf -nop
```

If the R2 was consistent at the time that the SRDF/A session was suspended, the procedure for recovering back to a replicating state, whilst still ensuring that a consistent point-in-time copy is available at the target site at all times is as follows:

1. Verify BCV is fully established.
2. Split PIT BCV off at SRDF/A target site
3. Wait until BCVs finished splitting in the background
4. Set mode to adaptive copy disk mode
5. Establish RDF replication for device group
6. Wait until RDF Pair Status synchronized
7. Set mode to async copy and re-enable device group consistency
8. Establish PIT BCV

The above steps can be executed automatically by using the symrecover command. The symrecover command that restarts SRDF/A whilst automatically creating a point-in-time copy. This standard program can be used in place of resume_async script. Full details can be found in the document “**EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide**”. However, the following Best Practice is recommended:

- Do not run in monitor mode by default as one symrecover is needed per RDF group which will not scale. Also symrecover will automatically restart a manually suspended RDF group, which is probably not desired functionality.
- Instead, it is recommended that symrecover be used manually as required to automate the complex restart process, including taking the point-in-time copy at the remote site before .

For example, run at the production site:

```
# symrecover start -g <Device-Group> -mode async -options <options-file>
```

where options file contains:

```
monitor_cycle_time=60
```

```
run_once=1
restart_group_on_startup=1
goldcopy_bcv_r2_mirror_state_post_restart=establish
```

This will:

- split off the remote BCV
- set the RDF group to acp_disk mode
- establish the RDF pair
- change mode to async when establish is complete
- re-establish the remote BCV once the RDF group is in consistent state

symrecover actions are logged in /var/symapi/log/symrecover-yyyymmdd.log.

17.2 Zero Space Reclamation

Zero Space Reclamation (ZSR) is a feature that frees up unwritten tracks or tracks written with zeros back to the SRP. ZSR can be run on device groups, storage groups, or a range of devices. ZSR should be run on a periodic basis during periods of low activity on the array.

Note: Running ZSR on replicated volumes will not initiate ZSR on the remote array. ZSR has to be run on both the source and target arrays to reclaim space on both arrays.

This example shows ZSR being run on a storage group:

```
# symmsg -sid 973 -nop -v reclaim -sg qtest-rdf1_sg
Requested devices: 000AE 000AF 000B0 000B1 000B2 000B3 000B4 000B5 000B6
000B7
```

```
STARTING a 'RECLAIM_START' control operation for 10 devices.
Reclaim Start device(s) at local Symmetrix.....Done.
```

```
The 'RECLAIM_START' control operation SUCCEEDED.
```

18 Storage Provisioning Best Practices

18.1 Virtual Provisioning

- Only one SRP should be configured per storage array.
- Maximum of 200% oversubscription of SRP capacity
- Event 1512 should be added to event include list for Tivoli Monitoring. This will send an alert each time the allocated capacity of the SRP reaches the following thresholds:

% SRP Capacity Allocated	Severity
>= 80%	Critical
>= 70%	Major
>= 65%	Minor
>= 60%	Warning

- All devices should be created without pre-allocating capacity unless the application specifically requires pre-allocated (i.e. thick) devices.
- The largest devices that can be created on VMAX3 is 16TB. Refer to the [CSBE LUN Sizing Guidelines](#) for recommended LUN sizes. Since VMAX3 no longer supports meta-devices, all LUNs should be configured as a single device.
- SLO should be set to 'Optimized' when provisioning storage.

18.2 Host Connectivity

- A minimum of two host connections to the SAN; one to Fabric A and one to Fabric B.
- Each host initiator should be zoned to a minimum of two FA ports, on different directors and on different engines.

18.2.1 Capacity Monitoring with SRM

Once certified by Citi (expected end of Q2 2015), ViPR SRM will be used for chargeback reporting for VMAX3. ViPR SRM will utilize SMI-S on the SE management server and Unisphere for VMAX to discover and report on the VMAX3 storage array. An Active Directory functional ID with need to be created for ViPR SRM (srm_readonly) and assigned the symauth role of 'Auditor' for SRM to gather the data.

19 Data Protection Best Practices

19.1 Local Protection

All VMAX hardware components are fully redundant and field replaceable. Disk RAID protection is used to protect data against drive failures.

19.2 Remote Protection

SRDF is used to replicate data between 2 sites to provide protection against a single site failure.

19.3 Backup Standards

Refer to [CSBE Backup Services](#) product page for available backup solutions.

20 Security Best Practices

20.1 Security Standards

Citi Information Security Standards (“CISS”) are available from [Information Security Services \(“ISS”\) website](#).

Citi’s SAN, NAS and Backup infrastructure stores and processes information across the full range of information security classifications.

For data classified as Public or Internal, no specific provision is required for data encryption either in-flight or at-rest.

For Confidential and Confidential PII support, all new applications are required to encrypt confidential PII data in transit (refer to Encryption Requirements Section of [CISS](#)). For Restricted or Authentication data, encryption of data is required both in-flight and at-rest.

The recommended approach for Restricted, Authentication and new applications storing Confidential or Confidential PII data would be to encrypt the data within the application prior to transmission over the network. This approach has the benefit of also meeting the requirement in all these cases for persistent encryption when non-Citi managed storage infrastructure is used. At the time of writing this document there are no tested or approved solutions at the storage infrastructure layer to encrypt data in-flight or to encrypt the data at rest, meaning that where required, this functionality needs to be applied by the application.

This classification, by definition, infers a “Medium Risk” status on all Citi’s SAN, NAS and Backup infrastructure.

20.2 Configuring the Event Daemon (storevntd) on the RHEL SE Server

The SE event daemon (storevntd) is a daemon that runs on the RHEL SE server and can be used to import the symaudit log as well as to send out Symmetrix alerts via SNMP, syslog, the OS’s event log, or to a file. The following steps should be followed to configure the storevntd:

1. The following are the main entries should reside in the /var/symapi/config/daemon_options file on the SE server to specify the behaviour of the storevntd (There are some other parameters in the daemon_options file for the storevntd but the following are the basic settings to get storevntd to functioning):

```
storevntd:log_event_targets = file syslog      #Direct the logs and events to a local file and syslog target.

storevntd:log_symmetrix_events = \
    sid=000192601380, audit, tgt=file|syslog;\ #Directs symaudit logs from the specified Symmetrix to the local file and also send the audit event to the syslog server specified below.
    sid=000190301292, audit, tgt=file|syslog;\
    status, tgt=file;\ #Sends status messages to the file target
    events, tgt=file #Sends event messages to the file target

storevntd:log_event_file_name = symevents #Specifies the prefix of the local log file.

storevntd:log_event_syslog_host = 10.40.41.199 #IP address of the SEMS server specified by the GSOC team.
```

2. Add a symauth entry for storevntd. In the example below, a command file called **adduser.cmd** was created with the entry **assign user H:sangs7\storevntd to role Auditor**; to create the symauth entry for storevntd:

```
# symauth -sid 110 commit -f adduser.cmd

Perform an Authorization change on Symmetrix unit 000190103110 (y/[n]) ? y

PREVIEW.....Started.
PREVIEW.....Done.

COMMIT.....Started.
assign user H:sangs7\storevntd to role Auditor.....Done.
COMMIT.....Done.
#
```

3. Start storevntd:

```
# stordaeon start storevntd
Waiting for daemon to start. This may take several seconds.
#
```

4. Run the command to have storevntd start automatically after a system reboot:

```
# stordaeon install storevntd -autostart
```

5. Reload the new options specified in the daemon_options file for the storevntd:

```
# stordaeon action storevntd -cmd reload
storevntd      Action performed
#
```

20.3 Audit Logging

Audit log reconciliation is required for all user-initiated changes on the VMAX3. The storevntd on the RHEL SE management server is used to send the VMAX3 audit events to the ArcSight system via syslog. These events will be parsed and reconciled once received. The symaudit log entries can be viewed by issuing the 'symaudit' command. The Symmetrix user issuing the symaudit command must have the role of SecurityAdmin, Auditor, or Admin.

Citi's Security Event Monitoring System (SEMS) is used to centrally collect security audit events. Custom reports can then be generated from the events collected by SEMS and sent to the log review teams for reconciliation. The following steps need to be followed to integrate the Symmetrix with SEMS for audit log reconciliation:

1. storevntd needs to be configured on the SE management server to send Symmetrix syslog audit events to the SEMS server.
2. The SEMS GSOC team should be contacted to initiate the request for integrating the Symmetrix with SEMS for log reconciliation. The SE management station should be specified as the source device in the request. Use the following distribution: ***CTI Global SOC SDM**
 - a. If PowerBroker is used for user access control and audit logging, the existing log reconciliation process for PowerBroker should be followed for reconciling Symmetrix changes performed by Citi personnel. Only vendor initiated changes should be reconciled using the SEMS system for environments using PowerBroker since PowerBroker cannot track vendor changes which are initiated on the Symmetrix Service Processor. Please note in the request to the SEMS team whether PowerBroker is being used. If Powerbroker is not being used, ensure that symauth is implemented.
 - b. For environments that have implemented symauth for user access control and audit logging, all symaud events should be reported and reconciled with the exception of the events in the exclude list specified in the next step.
3. The GSOC team will then configure Arcsight to receive the Symmetrix syslog audit events and generate daily reports to the regional log review teams.

The GSOC team will apply a filter to the report to only include audit events (fmt=symaud) and an event description containing the following:

- 41) UNMAPPING
- 42) UNMAP
- 43) MAPPING
- 44) MAP
- 45) BASE CONTROL
- 46) CREATE
- 47) ADD
- 48) DELETE
- 49) APPLY
- 50) RENAME
- 51) INITIALIZE
- 52) MODIFY
- 53) REMOVE
- 54) RWENABLE
- 55) RWDISABLE
- 56) CODELOAD
- 57) HBA_PORT_FLAGS_ENABLE
- 58) RENAME_WWN_ALIAS
- 59) CLEAR_WWN_ALIAS
- 60) HBA_PORT_FLAGS_DISABLE
- 61) HBA_PORT_FLAG_OPTION_CLEAR
- 62) REFRESH_DEVMASK_DB
- 63) ADD_DEVICES
- 64) REMOVE_DEVICES
- 65) FAST
- 66) FAST_VP
- 67) FAST_VP
- 68) DELETING
- 69) CREATING
- 70) LICENSE
- 71) DEVMASK
- 72) METADEVICE

-
- 73) ATTRIBUTE
 - 74) BCV
 - 75) CLONE
 - 76) RDF
 - 77) RCOPY
 - 78) QoS
 - 79) CFGCHG
 - 80) RECOVERY

For environments still using PowerBroker, audit events showing up under the 'root' user of the SE management server should also be excluded from the report since PowerBroker has its own log review process. It is important to specify whether PowerBroker is being used in the request to the GSOC team so that they know which filter to apply to the reports for log reconciliation.

20.4 MMCS Anti-Virus

The VMAX3 MMCSs have anti-virus software running on them but virus-definitions do not get updated until the next upgrade of array firmware (i.e. Hypermax OS). 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 MMCSs and configured to use Citi's ePO infrastructure for VirusScan software and virus definition updates. By integrating the VMAX3 MMCSs 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 MMCSs with Citi's ePO infrastructure:

- 1) Create firewall rules for ePO communication.

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

Source	Destination	Protocol/Port Numbers
VMAX3 MMCS	ePO server	TCP/18080 TCP/18443
ePO server	VMAX3 MMCS	TCP/18081 UDP/18082

- 2) 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 VMAX3 MMCS. 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_contact_s.aspx. The following should be done after following the instructions in the document:
 - VMAX3 doesn't require its own policies so the Citi Global Standards policies should be applied.
 - The FramePkg.exe and SiteList.xml files 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 agent has been installed and configured on the VMAX3 MMCS, the MMCS's ePO agent report to the ePO server and should show as an end node.
- 4) 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.

20.5 Data Security

The recycling or reusing of LUNs from one host/application to another is prohibited. The recycling of LUNs compromises the data written by the initial host/application.

It is recommended that all LUNs must be deleted when no longer required by the host to which it was initially presented.

It is also recommended that the host SA should follow guidelines consistent with the data erasure of a physical drive prior to deleting the LUN.

20.5.1 Disk Sanitization

Disk sanitization is the process of physically obliterating data on disks in a manner that precludes recovery of that data by any known recovery methods.

Disk sanitization is a requirement mandated by CITMP IS Policy.

As there is currently no feature available within the VMAX3 that enables overwriting the contents of a disk, it is recommended to follow an approved process in accordance with Global Data Center Secure Media Disposal policies to

- degauss the disks
- use EMC Data Erasure for VMAX3 service to securely erase or dispose of the disks (compliance with Global Data Center Secure Media Disposal policies should be verified and approved by an appropriate Information Security Officer)

NOTE: Current CITMP limitations restricting movement of electronic media warrant that all disk sanitization be completed without the transportation of disks.

20.5.2 Disposal of failed disks

All failed disks within every CTI data center are subject to stringent sanitization or degaussing procedures. As such, all failed media should be processed in accordance with Global Data Center Secure Media Disposal policies.

20.6 CITMP Information Security Policy Convergence

20.6.1 Information Ownership

The head of Distributed Engineering has information ownership of Storage products and all related engineering documents. Storage products are available in the Citi Technology Catalogue. All security and technology best practices for Storage products are the responsibility of CATE Storage and Backup Engineering and can be referenced from the [CSBE website](#).

20.7 Role Based Access Controls (RBAC)

20.7.1 Symauth Role-Based Access Controls

Refer to [Symauth Role-Based User Access Control](#) for instructions on setting up RBAC on VMAX3.

20.8 Local FID Security Guidelines

The following table lists the functional IDs that are configured within the Symmetrix:

Application	Account	Interactive/Non-Interactive	Symauth Role	Password Change Frequency
HPNA	hpnaro	Non-Interactive	Auditor	Not required
ViPR SRM	srm_readonly	Non-Interactive	Auditor	Not required

1) Confirm Basic Attributes of the FID

- a. Read-Only : Yes
- b. Operationally Critical : No
- c. Fixed : No
- d. Privileged : No
- e. Interactive : Yes

2) Provisioning of the FID

- a. CTI Build Team should create the FID on the devices. If GIDA team is currently used for account creation they should be engaged.

3) Inventory of FIDs.

- a. Currently CISAR is the only standard inventory which can be populated by GIDA or automatically through EERS feed (CISAR and EERS have nightly reconciliation process).
- b. Each instance of the FID should be included in inventory containing at least:
 - i. Name of Functional ID
 - ii. Owner of the FID
 - iii. Description / purpose of the FID
 - iv. Privileged / Non-Privileged
 - v. Operationally Critical Component
 - vi. Interactive/Non-Interactive
 - vii. Entitlement review process

4) Entitlement reviews

- a. Entitlement reviews are optional for non-Privileged FIDs.

5) Password Management

- a. Password Management Process is required for privileged and interactive FIDs. PAR is recommended to meet requirement of control over list of users authorized to obtain the password and to store it securely.
- b. Based on CISS 11.5.3.4a password may be set to not expire.

6) Requirements related to SSH keys

- a. Cannot re-use keys across different FIDs.

20.9 Standard Build Configuration Checklist (SBCC) with HPNA

The Hewlett Packard Network Automation (HPNA) solution is used to verify that the VMAX3 configuration is compliant with the [VMAX Standard Build Configuration Checklist](#) (SBCC). Reference [GNAST Design and Standard for SAN Automated Configuration Assurance](#) for instructions for installing the HPNA client on the RHEL SE server. Once installed, a user account called 'hpnaro' is created on the OS. A symauth entry for the 'hpnaro' user will need to be created and assigned the symauth role of 'Auditor' for each Symmetrix array the SE server is managing. The hpnaro user will also need r-x permissions to the /var/symapi/db and /var/symapi/config directories on the SE server.

21 Appendix A – Document Version Control

This section contains a detailed list of the updates made during each major document release. Section 1 is purely an overview.

Version	Changes
0.1	Initial draft
0.3	Minor changes after peer review and released for comments
0.4	Added ePO setup on the MMCS
0.5	Updated Citi standard VMAX3 hardware configurations
0.6	- Updated the VMAX3 document format to match the CSBE Solutions Guide document template - Incorporated all VMAX3 related Solutions Guides (i.e. Solutions Enabler and Unisphere) - Added event ID 1512 (SRP capacity threshold alert) to events to be monitored by Tivoli - Added SRP reserved capacity - Added SRDF section - Added SnapVX section - Added ZSR section - Added bronze hardware configuration
0.8	Added Unisphere for VMAX section

22 Appendix B - List of Symmetrix Event Codes

Event ID	Description
1050	Diagnostic event trace triggered.
1051	Remote (SRDF) diagnostic event trace triggered.
1052	Too many suspend/halt chains encountered switching to Adaptive Copy Write Pending Mode.
1053	Memory bank(s) automatically disabled due to cache error.
1054	Hot Spare invoked against a disk.
1055	M2 mirror resynchronizing with its M1 mirror.
1056	M1 mirror resynchronizing with its M2 mirror.
1057	Disk director not responding.
1058	Data migration completed on all migration devices.
1059	Device resynchronization process started.
1060	Hot Spare invoked against a remote R2 mirror disk.
1061	SIM message initiated to a remote SRDF attached array.
1062	SRDF error occurred.
1063	Fibre Channel front-end has failed or is inoperable.
1064	No SRDF links in an RDF group are operational.
1065	All SRDF links in an RDF group are operational.

1066	Bus Arbiter problem.
1067	Internal temperature too high.
1068	Alarm signal was set, but no alarm found.
1069	Power subsystem alarm or fault has occurred.
1070	Device mirror not ready.
1071	Device mirror write disabled.
1072	SRDF R2 device not ready.
1073	Service processor down or not communicating with array.
1074	Service processor could not complete a call for service.
1075	Abnormal DC voltage (perhaps 12 Volts) situation exists.
1076	Power subsystem environmental sense cable missing.
1077	Power system AC line interruption detected.
1078	Battery system not fully charged.
1079	Latched alarms discovered for the power subsystem.
1080	Single SRDF link in an RDF group is not operational.
1081	Single SRDF link in an RDF group is now operational.
1082	Service processor successfully completed a call for service.
1083	Subsystem unable to set a shared register.
1084	Disabled memory bank error reported to a host.
1085	Validity problem detected during environmental test.
1086	Environmental testing enabled in diagnostic mode.
1087	Communication board data does not match expected value.
1088	Communication board information mismatch.
1089	Failure detected during thermal test.
1090	Power-on-Time inconsistencies detected.
1091	No records found for the last service processor connection time.
1092	Service processor communicating via a serial line.
1093	Remote session to the service processor connected.
1094	Remote session to the service processor denied access.
1095	Remote session to the service processor disconnected.
1096	Service processor detected excessive memory usage.
1097	Battery test detected a failure.
1098	Service processor could not communicate with a director.
1099	Service processor could not query a director.
1100	Service processor is communicating via local director.
1101	Service processor unable to read an environmental sensor.
1102	Service processor detected a failed or unrecognized communications card.
1103	Service processor found environmental readings to be out of limits.
1104	Service processor disk is full.
1105	Service processor detected a smoke detector malfunction.
1106	Service processor detected a smoke detector alert.
1107	Service processor triggered a call home for service.
1108	Database Double Checksum event triggered.

1109	RDF CG trip event triggered.
1110	Service processor successfully rebooted.
1111	SAVE device pool is full.
1112	SRDF/A session inactive.
1113	SRDF/A session active.
1114	Volume (device) not ready.
1115	SAVE device pool is almost full.
1116	Active device in SAVE device pool not ready.
1117	Director not responding.
1118	Timeout writing to an R2 device; writes pending limit reached.
1119	SRDF/A session dropped: write pending limit reached. Host throttling is disabled.
1120	SRDF/A session dropped: write pending limit reached. Host throttling is enabled.
1121	SRDF/A session dropped: device not ready.Tolerance mode is off.
1122	SRDF/A session dropped: device not ready through consistency group.
1123	SRDF/A session dropped: no SRDF links operational.
1124	SRDF/A session dropped: timeout in MSC mode.
1125	SRDF/A session dropped: timeout on Host Adapter.
1126	SRDF/A session dropped: timeout on RA.
1127	Double Checksum (generic) event triggered.
1128	One of the Power Zones is down - shutdown will occur in 20 hours.
1129	One of the Power Zones is down - shutdown will occur in five hours.
1130	SRDF/A session drop requested [host software initiated].
1131	SRDF/A session drop at cycle boundary requested [host software initiated].
1132	SRDF/A session transition out of Asynchronous mode requested [host software initiated].
1133	SRDF/A session transition from Asynchronous to Synchronous mode requested [host software initiated].
1134	SRDF/A session drop requested.
1135	SRDF/A session drop at cycle boundary requested.
1136	SRDF/A session transition out of Asynchronous mode requested.
1137	SRDF/A session transition from Asynchronous to Synchronous mode requested.
1138	SRDF/A session entering transmit idle state.
1139	SRDF/A session recovered from transmit idle state.
1140	SRDF/A session dropped, transmit idle state timeout.
1141	SRDF/A session dropped, no online RAs
1142	Power Supply A multiple fan fault.
1143	Power Supply A single fan fault.
1144	Power Supply A faulted.
1145	Power Supply A shutdown.
1146	Power Supply B multiple fan fault.
1147	Power Supply B single fan fault.

1148	Power Supply B faulted.
1149	Power Supply B shutdown.
1150	Link Card Controller A temperature high.
1151	Link Card Controller B temperature high.
1152	Supplemental Power Supply internal fault.
1153	Supplemental Power Supply battery end of line.
1154	Supplemental Power Supply low input AC Voltage.
1155	Entries are being written to the audit log at an unusually high rate.
1156	Audit log has lost its redundancy due to SFS mirror being offline.
1157	Audit log entries have been overwritten in an unusually short period of time.
1158	SRDF/A session dropped, write pending limit reached on a cache partition.
1159	There has been no SRDF/A cycle switching within the past hour.
1160	The job flow control setting has changed on RDF group.
1200	Device state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1201	Array state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1202	Director state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1203	Port state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1204	Disk state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1205	Device configuration has changed.
1206	[Snap SAVEdev SRDF/A Delta Set Extension] pool state has changed to [Not Present Unknown Online Write Disabled Offline Failed].
1207	[Snap SAVEdev SRDF/A Delta Set Extension] pool configuration has changed.
1208	<NAME> Pool utilization is now <NN> percent.
1209	Symmetrix external lock has been [released acquired].
1210	Disk hot spare [is no longer invoked against a failed disk has been invoked against a failed disk].
1211	Number of available disk hot spares is N.
1212	Virtual device is now N percent allocated.
1213	Virtual device is now N percent used.
1214	Director configuration has changed.
1215	Port configuration has changed.
1221	Disk state is now [Offline Offline Spare Online Online Spare Online Degraded] (was: [Offline Offline Spare Online Online Spare Online Degraded])
1230	Array configuration has changed.
1231	Device Masking database has changed.
1232	Access Control definitions have changed.

1233	Dynamic RDF operation performed on device.
1234	Snap session created, activated or deleted.
1235	BCV device pairing has changed.
1236	HPUX device identifier has changed.
1237	Device name has changed.
1238	Device nice name has changed.
1239	OpenVMS device identifier has changed.
1240	Device reservations data has changed.
1241	Time since last SRDF/A cycle switch exceeds minimum cycle time by n seconds.
1242	SRDF/A cycles now using nn percent of the cache available for it.
1243	Write pending data is now using nn percent of the cache.
1244	Component state has changed to {Not Present Unknown Online Write Disabled Offline Failed}.
1245	Component state has changed to Online Online [Degraded] Offline Unknown
1246	Disk Group has changed. or Disk Group has been deleted. or Disk Group has been created.
1247	Disk has spare coverage. or Disk no longer has spare coverage.
1248	The spare capacity percentage for diskgroup on engine %d has changed
1249	Disk has spare coverage. or Disk no longer has spare coverage.
1280	Cache partitioning configuration has changed.
1281	Dynamic mapping configuration for a device has changed.
1282	Meta configuration for a device has changed.
1283	Initiator group has changed.
1284	Storage group has changed.
1285	Director Port group has changed.
1286	Masking view has changed.
1287	Feature registration database has changed.
1288	Application registration database has changed.
1289	optimizer,FAST tiers have changed
1290	optimizer,FAST policies have changed
1291	optimizer,FAST associations have changed
1292	optimizer,Optimizer/FAST time windows have changed
1293	optimizer,Optimizer/FAST control parameters have changed
1294	Storage group configuration has changed.

1295	Initiator group configuration has changed.
1296	Port group configuration has changed.
1297	Masking view configuration has changed.
1298	Storage group scope has changed.
1299	Initiator group scope has changed.
1300	Masking view scope has changed
1301	Indicates that a change has been made to the FAST (Fully Automated Storage Tiering) control parameters on the Symmetrix array.
1400	User authorization rules have changed.
1401	Audit log is at nn percent of capacity (before wrapping).
1402	Security Alert from Audit Log.
1403	Security Failure Alert from Audit Log.
1404	Message from Audit Log.
1500	Optimizer Swap activity (from Audit Log).
1501	Optimizer Move activity (from Audit Log).
1502	Optimizer configuration change (from Audit Log).
1503	FAST Controller Swap activity (from Audit Log).
1504	FAST Controller Move activity (from Audit Log).
1505	FAST Controller configuration change (from Audit Log).
1506	Optimizer Rollback activity (from Audit Log).
1507	User approval is required for a Config Change plan generated by the Optimizer/FAST Controller.
1508	The FAST (DP or VP) controller has switched to a different state.
1509	The Optimizer has switched to a different mode.
1510	%s FAST policy allocated capacity is now %u percent.
1511	FAST Tier (%s) performance changed
1512	The allocation for SRP has changed.
1513	The Snapshot capacity for SRP has changed to [0..100] percent."
1601	GNS is now using [0..100] percent of the available group space
1602	GNS is now using [0..100] percent of the available group directory entries.
1650	The polling cycle is complete.

23 Appendix C – 2-Site Multi-Session Consistency

This section deals with issues that are specific to an implementation requiring multi-session consistency, and which are over and above those described in the Generic SRDF Configuration Section.

23.1 Overview of MSC

MSC offers the ability to provide SRDF/A replication of multiple RDF groups with I/O consistency across all the groups, whilst still retaining the ability to failover individual RDF groups to deal with, for example, application server failures, without impacting the replication of other applications.

SRDF/A MSC provides consistency protection for an asynchronous-mode consistency group by performing cycle switching and cache recovery operations across all SRDF/A sessions in that group. Both the RDF daemon and SYMAPI contain SRDF/A MSC logic to determine whether to commit or discard data in the cache when SRDF/A sessions have been deactivated.

Consistency here means that the devices act in unison to maintain the integrity of a database distributed across multiple RDF groups within a single Symmetrix, or indeed across multiple Symmetrix. If a source R1 device in the consistency group cannot propagate data to its corresponding target R2 device, SRDF/A MSC suspends data propagation from all R1 devices in the consistency group, halting all data flow to the R2 targets.

MSC cycle-switching and cache recovery operations across all SRDF/A sessions in the consistency group ensure a consistent R2 copy of the database during the time that the interruption occurred. When an interruption occurs, SRDF/A MSC examines the status of SRDF/A sessions and data. Based on the results, MSC either commits the last copy cycle to the R2 side or discards it, converting the track to a remote invalid on the R2 device.

Achieving data consistency across multiple RDF groups requires that the cycle switch process be coordinated, and that the switch occurs during a very brief period (usually a few ms) when no host writes are being serviced. MSC provides a single coordination point to drive the cycle switch process.

Implementing consistency protection involves building a composite group (a grouping of device groups) and enabling the composite group for consistency protection, at which time the group becomes a *consistency group* and is managed by the new RDF daemon.

An interruption that causes deactivation of the SRDF/A sessions and the suspension of the consistency protection can occur in several ways. Deactivation occurs when one or more R1 source devices in a consistency group cannot propagate data to their corresponding R2 target devices. For example:

- RDF links between the R1 and R2 might go down for an extended period of time
- The RDF directors on the R1 side or R2 side might fail

If I/O cannot be propagated from an R1 device in the consistency group to its R2 device, MSC suspends remote mirroring on ALL RDF links in the consistency group, in a manner that honours dependent write I/Os.

23.2 Choice of consistency models

For datacentres with many applications and storage frames, consistency comes in two extreme flavors (intermediate variations are also possible e.g. consistency at the array level):

A) Enterprise Consistency: All the devices for all applications are in one consistency group. Either SRDF/CG (for sync mode) or MSC (for async) is used to manage consistency among frames. In case of STAR, one STAR instance is managing the system. The main benefit is that in a DR scenario, all the applications on the remote side(s) have the same point in time, and the dependencies between them are preserved. However, this solution is much less flexible, as application or server level failover is very difficult, or impossible, to provide as applications are not independent of each other from a replication perspective. This is an issue as the business will commonly want to do application level failover for the following reasons:

- **DR testing.** For example, before a new application goes live, there is often an extensive end-to-end testing of the application failover/failback process. This of course needs application level failover to avoid disrupting anyone else. Once live, there are also scheduled DR invocation tests too. For this purpose alone, application level failover is essential.
- **Dealing with component failures in servers.** For example, in London, the Storage Management team execute failovers about 3 times a month for this purpose alone as not all applications have local HA via clustering, but instead rely on SRDF and their business continuity server at the remote site.

B) Per Application Consistency: Each application maintains its consistency independently, with no enterprise consistency. In case of disaster, applications will belong to different point-in-times and even though each of them will be consistent, transactions that involve multiple applications might be inconsistent (registered in one database and not another). Manual work may be required to bring the inter-dependent applications into a consistent state. In SRDF terms, this configuration means a separate consistency group (for the case of SRDF/S), a separate SRDF group (SRDF/A app), or a separate STAR instance (for SRDF/STAR). The major advantage here is the ability to failover individual applications. This solution is available today with limits on the number of apps resulting from limitations on SRDF groups, STAR instances, etc.

Very similar to application-level failover/consistency is server level failover/consistency, but there are some important differences:

- Some applications have a 1:1 mapping with a data server. In such cases, server level and application level consistency are one and the same.
- Some applications have multiple data servers. Here, data consistency between the data servers may be necessary and the volumes both servers form the consistency domain.
- Some servers have multiple applications. There are two variants here:
 - Each application on the server has its own set of LUNs. Each application can be failed over independently.
 - Applications on the server share LUNs. Hence, all applications that share LUNs need to be failed over together as there is a hard dependency between the applications. Where this is a common model, server level failover may be the norm.

In each case, the consistency can be provided either by putting all LUNs in a consistency domain into the same RDF group. All must then be failed over together. If some level of independent failover is needed, then MSC can be used to bind multiple RDF groups together in a consistent manner. Single RDF groups can then be failed over, although of course your inter-application consistency is then broken, so you may have to fail over all dependent applications anyway to maintain application consistency.

In summary, the required consistency domains must be identified, and the most appropriate mechanism (single RDF group vs MSC) identified to provide this with the required level of failover granularity.

The two solutions A & B described above are inherently contradictory. Should one sacrifice enterprise consistency for the sake of flexibility i.e. being able to manage each application independently? This depends on the requirements and recovery plans of the business. Generally

what is needed is “business consistency”, which is defined here as consistency between a set of inter-dependent applications. That is a much smaller domain than the enterprise or datacentre. Different business units believe they can address this in different ways:

- **Most do not expect there to be automated or storage-enforced consistency between applications.** They state that they already have alternative mechanisms to deal with recovery post disaster; many of these apps were built in days before SRDF. We in engineering do not have visibility of the recovery plans, or indeed necessarily the relevant expertise, to validate these, and have to take business statements at face value.
- **Some wish to rely on Storage Array consistency technology.** This is an approach that has been adopted in Singapore. CSBE also support this model on a smaller scale (e.g. using MSC for interdependent applications within an array).

The focus of this Solutions Guide is on the application level consistency and failover model. Limited inter-application data consistency is also at the storage level via MSC. Enterprise level consistency is not supported.

23.3 Running MSC between frames

CSBE strongly discourage spreading a single application across multiple frames, regardless of SRDF mode. The remote copy of application data sets spread over multiple storage arrays are at risk of being inconsistent in the event of a rolling disaster or other failures. Consistency groups would need to be used to protect against such a condition when an application is spread over multiple frames. This is undesirable for a number of reasons:

- MSC is less robust than single RDF group SRDF/A replication on a single frame. MSC depends on the availability of an external daemon running on a server to execute the cycle switches.
- Use of MSC requires additional Consistency Group licenses on all source and target VMAX. This is over and above SRDF/A licensing. Hence, if you have an application that spans, say, 3 frames, you need 6 CG licenses. These are licensed per Symm, and based on total capacity and hence are expensive to implement as the cost of the licenses is high.
- CSBE has insufficient Symmetrix (a minimum of 4 would be needed) to test and support multi-frame consistency configurations and hence would need to rely exclusively on EMC testing.

In very rare occasions, there will be a requirement to support MSC between arrays where a single application is too large to fit within a single array. Consistency Groups technology should not be used to deal with applications that span multiple arrays due to poor data management due to the cost and increased operational complexity associated with this technology.

Where consistency is needed between multiple applications, this is supported by CSBE in the case where applications are all on the same frame. However, the above bulleted caveats still apply.

MSC between applications located on separate frames cannot be tested by CSBE and at the time of writing is not supported. If necessary, this could be supported, but we would need to rely exclusively on EMC for testing and on-going support. Again, the above bulleted caveats apply.

23.4 High-availability of MSC replication

23.4.1 Overview

It is important to understand that SRDF/A cycle switching for individual RDF groups is controlled by the VMAX itself, whereas for consistency-enabled Composite groups (a.k.a. Consistency

Groups), switching is controlled by daemons (storrdfd) running on fibre-attached hosts. This is inherently less robust than VMAX-controlled cycle switching.

EMC have implemented a live-live control system to improve availability. MSC daemons can be running on multiple servers and the active master for each MSC session can change with each cycle. As a consequence, the MSC service is highly available by design, and cycle switching (i.e. replication) will occur as long as there is at least one storrdfd running on a connected SAN management server. No clustering is required to provide high availability.

However, it is important to note that:

- If no storrdfd daemons are available:
 - Cycling switching of the MSC-controlled consistency group will cease
 - The “time R2 is behind the R1” time will increase
 - IOs will be stored in the production VMAX SRDF/A cache (at least until this becomes exhausted)

Once a daemon comes on-line again, it takes over replication control of the MSC group again. Hence, it is vital to ensure that multiple servers are configured for MSC control to ensure a continual replication service.

- The cycle number shown by symrdf query and symstat will vary when an RDF group is under MSC control.

```
# symrdf -cg <cg-name> query -rdfa
```

shows MSC cycle number which is reset to 0 when MSC control was enabled

```
# symrdf -g <dg-name> query -rdfa
```

shows MSC cycle number which is reset to 0 when MSC control was enabled

```
# symstat -type cycle
```

shows the original cycle number for the RDF group prior to MSC being enabled. Hence, different groups can be on different cycle numbers, even if they are under MSC control and switching in step with each other. Slightly confusing, especially as symstat gives different output than `symrdf -g <dg-name> query -rdfa` for a given RDF group.

23.5 Configuring MSC support on management servers

23.5.1 Gatekeeper requirements

One gatekeeper is required for every 3 Composite Groups to be put under MSC control. In addition, several (typically 3) gatekeepers are required for status querying.

Note that it is possible to dedicate up to 20 Gatekeepers for exclusive use by the storrdfd daemon to ensure that there are plenty available for MSC control. This is achieved using the parameter

```
storrdfd:rdfd_num_dedicated_gks=7
```

in the `/var/symapi/config/daemon_options` file.

23.5.2 High-availability of RDF daemon

MSC replication control is via the RDF daemon storrdfd. High-availability of the MSC replication function is provided by ensuring that there is a minimum of two RDF daemons active at any one time at the source site.

The RDF daemon is enabled with the following entry in /var/symapi/config/options

```
SYMAPI_USE_RDFD = ENABLED
```

The daemon can be set to start automatically every time the local host is booted. This is known as installing the daemon. If the Daemon is stopped for some reason (other than deliberately via the storrdfd shutdown command), it can be automatically restarted by in internal Solutions Enabler watchdog mechanism. A combination of the watchdog mechanism and auto-start option above can be used to ensure that the daemon is always running.

The RDF daemon (storrdfd) can be “installed” with the following command.

```
# stordaeon install storrdfd -autostart
```

Installing the daemon ensures that the watchdog daemon restarts the storrdfd if it fails. The –autostart ensures that the storrdfd is started at system boot time. A daemon that is shutdown (rather than dies or is killed) is not restarted by the watchdog daemon.

Note that any symrdf command (e.g. symrdf list) will cause a shut down daemon to start up again.

The status of the storrdfd can be checked with the command

```
# stordaeon show storrdfd
```

```
Sarbbunker1# stordaeon show storrdfd
```

```

Daemon State           : Running
Daemon Start Time      : Thu Apr 13 09:18:45 2005
Version                : T6.0-623 (0.43)
Auto-Restart by Watchdog : Enabled

Total Number of Connections : 1
Number of Active Connections : 1
Total Number of Requests    : 0
```

23.6 Use of Composite groups

MSC is controlled via Composite groups. To implement MSC for a set of device groups, it is necessary to:

1. Create the Composite Group enabled for RDF daemon controlled consistency, and populate the Composite Group with the devices from the device groups that are to be bound together from a consistency perspective.
2. Replicate composite group to other SRDF/A management servers
3. Enable the Composite group for consistency i.e. enabling Multi-Session Consistency replication.

23.6.1 Scalability restrictions

EMC has tested up to 250 RDF groups in an MSC CG. Citi has only tested up to 30 due to lab facility constraints. Citi found in testing that with large numbers of RDF groups, the transfer of the delta sets does not occur at the same rate for all groups. Hence, some RDF groups will complete their individual cycles later than others. This can lead to cycle extension as all delta sets must be transferred or applied before the group cycle switching can take place.

23.6.2 Creating the Composite Group

The following is an example of how to setup a composite group called multi that binds together devices in RDF groups 3, 4 and 5. The source VMAX has a SID of 1380.

Put all the devices in RDF groups into a composite group called multi. The `-rdf_consistency` flag places the composite group under the control of the rdf daemon.

```
# symcg create -type rdf1 -rdf_consistency multi
# symcg -cg multi addall dev -rdfig 3 -sid 1380
# symcg -cg multi addall dev -rdfig 4 -sid 1380
# symcg -cg multi addall dev -rdfig 5 -sid 1380
```

Note that it is important to clean up composite groups when no longer needed (using `symcg delete`) as these are stored in the VMAX (not the local host as with a device group) and can cause issues if another host attempts to use `symcg` commands. If one does not do this, the `storrdm` daemon keeps seeing groups in the VMAX that are no longer being used.

The opposite of `addall` is `rmall`. The following commands will remove the Composite Group multi.

```
# symcg -cg multi rmall -rdfig 3
# symcg -cg multi rmall -rdfig 4
# symcg -cg multi rmall -rdfig 5
# symcg delete multi
```

Note that creating the composite group has no immediate effect on the replication of the individual RDF group cycles.

```
# symcg show multi
```

“RDF Consistency Protection Allowed” entry should be Yes.

At this point “RDF Consistency Enabled” should be “No”.

23.6.3 Synchronizing Composite Group definitions across management server

The definition of each composite group defined for MSC replication must be replicated to all management servers responsible for providing an RDF daemon for MSC control. This is best achieved by using GNS in a 2-site implementation.

GNS provides the ability to store Composite Group definitions in a shared repository located on each Symmetrix array, which then becomes visible to all hosts that are configured to run Solutions Enabler. This allows all GNS-enabled hosts to see the same group definitions across the Symmetrix environment, while sharing real-time updates to group definitions and configurations made by other hosts.

Instructions for enabling GNS can be found in the document “Solutions Enabler 7.x” Solutions Guide.

23.6.4 Enabling MSC replication

At this point, RDF consistency will not be enabled, but individual RDF group sessions will be active (i.e. cycling under VMAX control)

```
# symrdf -cg multi query -rdfa
```

In output from the above command, check “RDFA MSC Info” section. MSC Session status should be “N/A”. Consistency State will also be “N/A”

Then check that under the RDFA Info section for each RDF group, that the Session Status should be “Active”. This indicates that the cycling of that group is under VMAX control, rather than MSC control.

Also, at this point at this point “RDF Consistency Enabled” should be “No”.

Now enable consistency between the individual device group devices (i.e. across the entire composite group) and confirm that multi-session consistency is active

```
# symcg -cg multi enable
```

```
# symrdf -cg multi query -rdfa
```

In output from the above command, check “RDFA MSC Info” section. MSC Session status should now be “Active”, although Consistency State will remain “INCONSISTENT” until 2 cycles have completed.

Also check that under the RDFA Info section for each RDF group, that the Session Status is “Active - MSC”. This indicates that the cycling of that group is under MSC control, rather than VMAX control (where “Session Status” would be “Active”).

Also, at this point at this point “RDF Consistency Enabled” should be “Yes”.

When MSC is enabled, it will:

- wait until all current cycles have completed. This will result in some individual most session cycles being extended beyond their usual durations
- All groups will then begin cycle switching in unison

Individual session cycle numbers (as shown by symstat) will not be reset and may vary. However, all delta set switches across all the sessions will now begin to occur simultaneously. Symrdf query -rdfa will show the cycle # as initially reset to 0 when MSC is enabled.

During this process, the RDFA MSC Info (observable with # symrdf -cg multi query -rdfa) will change undergo the following changes:

MSC Session Status	N/A →	Inactive →	Active →	Active
Consistency State	N/A →	Inconsistent →	Inconsistent →	Consistent

MSC replication will stop if:

- No daemons are available

- One (or more) of the RDF groups cannot replicate (link down). Replication of the other groups will then be held to ensure consistency.

23.7 2-Site Multi-Session Consistency Operations

23.7.1 Independent RDF group failover/failback

One of the benefits of Multi-Session Consistency in an SRDF/A environment is that it allows one to guarantee IO consistency at the R2-end between RDF groups (and hence individual applications), whilst still permitting failover independence for each application.

The process for failing over a single RDF group within a Composite Group replicating under MSC control is as follows.

1. Temporarily disable consistency on the composite group
2. Temporarily remove the devices in device group to be failed over from the Composite Group. Composite group definition on all local servers running RDF daemons for MSC control must be updated by GNS or manually.
3. Re-enable consistency on the composite group
4. Failover the device group
5. Confirm other groups in composite group are still replicating
6. Confirm CG definitions on remote CG servers are also up to date.

Scripts called **2site_MSC_failover** and **2site_MSC_failback** are available to automate the process. See Appendix B for further details.

In a real DR scenario, once failover to the SRDF/A target site has been executed, the PIT BCV should be split off to provide a golden copy of the last image of data from the production site.

For example, to failover RDF group 3 (device group DG03) on VMAX 1380, which is part of a Composite Group multi replicating under MSC control:

```
# symcgc disable -cg multi
# symcgc -cg multi rmall dev -rdfg 3
```

Ensure that local Composite Group definitions on local servers controlling MSC replication are updated, either automatically via GNS (strongly recommended) or manually, before proceeding with re-enabling MSC replication.

```
# symcgc enable -cg multi
# syrmdf -g DG03 failover
# symstat -type cache -sid 1380 -reptype rdfa -i 10 -g all
```

GNS remote mirroring should be used to ensure updated CG definitions are visible to all management servers at the remote site too.

The process for failing back the RDF group and adding it back under MSC control is:

1. Failback RDF group
2. Temporarily disable consistency on the composite group
3. Add devices in RDF group back into Composite Group
4. Composite group definition on all local servers running RDF daemons for MSC control must be updated by GNS or manually.
5. Re-enable consistency on the composite group
6. Confirm all groups in composite group are replicating together

```
# syrmdf -g DG03 failback
# symcgc disable -cg multi
# symcgc -cg multi addall dev -rdfg 03 -sid 1380
```

Ensure that local Composite Group definitions on local servers controlling MSC replication are updated, either automatically via GNS (strongly recommended) or manually, before proceeding with re-enabling MSC replication.

```
# symcg enable -cg multi
# symstat -type cache -sid 1380 -retype rdfa -i 10 -g all
```

GNS remote mirroring should be used to automatically replicate the updated CG definition to the remote site. If GNS is not in use, the remote CG definition will need to be updated manually.

23.7.2 Composite group failover/failback

An entire Composite Group can also be failed over. This would be required in the case where there has been an issue at the source site and DR is being invoked. In such cases, replication usually already have ceased due to the MSC Consistency Group tripping. Failover via the Composite group (rather than via the underlying RDF groups) ensures consistency at the R2 end.

When MSC replication is halted due to the MSC Composite Group tripping, a clean up process is required to determine if the receive Delta Sets (on the Target VMAX) for each RDF group should be applied to the R2 or discarded. A description of this process is beyond the scope of the present document, but can be found in Appendix H of the SRDF/A Evaluation document.

Scripts called **2site_CG_failover** and **2site_CG_failback** are available to automate the process. See Appendix B for further details.

The command to failover a consistency group is

```
# symrdf -cg <CGname> failover -force -symforce -nop
```

The **-force** option allows you to perform control operations on SRDF devices when they are not in the expected state. In this case, it is needed as consistency is enabled.

The **-symforce** option allows you to force an operation that overrides instances where they are normally rejected, usually as they will result in data loss. Note: To enable the **-symforce** option for RDF use, a behavior parameter called **SYMAPI_ALLOW_RDF_SYMFORCE** in the options file must be set to **TRUE**.

In this case **-symforce** is needed as the post-trip MSC clean up process can leave the invalid tracks tables in a state that requires the use of this flag.


In a real DR scenario, once failover to the SRDF/A target site has been executed, the PIT BCV should be split off to provide a golden copy of the last image of data from the production site.

Failback process requires the use of just the **-force** flag

```
# symrdf -cg <CGname> failback -force -nop
```


23.8 Cycle time under MSC

When RDF groups are under Multi-Session Consistency (MSC) control, cycle time increases as a function of latency. This is NOT due to lack of bandwidth but due to the overhead of the latency in controlling the SRDF/A delta set switching process. The following graph shows an example of how cycle time extends for MSC-controlled replication as a function of latency. The minimum cycle time was set to 30 seconds.



Note that the cycle time is also a function on the number of RDF groups under MSC control, and not simply a singular function of latency. The following graph demonstrates how cycle time changes as a function of the number of RDF groups under MSC control for a fixed latency (in this case 250ms).

Hence, MSC extend times beyond preset



use of will cycle the

minimum in most implementations with any significant latency.

24 Appendix D – Template Operation Scripts Summary

The following template scripts have been tested in the SRDF/A proof-of-concept test environment. This environment had protected R2s. They are not guaranteed to work in all locations due to variations in local practice relating to:

- Device group management (both creation and “replication” of device group definitions to local and remote SAN management servers).
- Formation of target R2s. These scripts are designed to work with protected R2s. If R2s are comprised of unmirrored R2 + BCV, then additional steps may be required to manage these additional BCVs.
- Local naming conventions for logical devices and device groups.

Nevertheless, they do provide fully working scripts which can be used “as is” where the various assumptions are fulfilled. Alternatively, they can be taken and have minor modifications made to them to work within your environment. Operations groups that choose to make use of these templates must take on on-going support of their implementation in their region.

If your location is using PowerBroker for running SYMCLI, then the scripts will need to be modified to accommodate this.

Unless otherwise stated, all scripts should be run from a SAN management server located at the production site.

A package containing these template scripts can be obtained from the GDS distribution server at:

SE platform	SE versions	Script package PRsrdfa version	Location of GDS distribution server
Solaris 10	7.0 → 7.1	1.2	/net/stealth/export1/home1.localhost/sw/PKG/PRsrdfa_1.2.dstream
Solaris 10	7.1 → 7.5	1.3	/net/stealth/export1/home1.localhost/sw/PKG/PRsrdfa1.3.dstream
RHEL 5.x	7.0 → 7.1	1.2	/net/stealth/export1/home1.localhost/sw/Linux/PRsrdfa-1.2-0_E0.noarch.rpm (The RHEL package may also be found on the yum repositories).
RHEL 5.x	7.1 → 7.5	1.3	/net/stealth/export1/home1.localhost/sw/Linux/PRsrdfa-1.3-0_E0.noarch.rpm (The RHEL package may also be found on the yum repositories).
RHEL 5.x RHEL 6.x	7.6.X → 8.x	1.4	//net/stealth/export1/home1.localhost/sw/Linux/soe5products-x86_64/PRsrdfa-1.4-0_E0.noarch.rpm (The RHEL package may also be found on the yum repositories).

Note from the author. The VMAX only script **2site_add_vmax_device** is not included in this package (to avoid it being run accidentally on DMX) and can instead be found in Appendix E of this document.

24.1 Summary of scripts

Monitoring scripts	
mon_srdfa_inactive	Logs errors in /var/adm/messages if inactive SRDF/A sessions found
mon_srdfa_cycle	Logs errors in /var/adm/messages if current cycle time for any replication group is greater than a preset threshold
mon_srdfa_cache	Logs errors in /var/adm/messages if overall SRDF/A cache utilization is greater than a preset threshold
srdfa_health	Single script combining all three of the above functions
mon_srdfa_dse_pool	Logs errors in /var/adm/messages if DSE pool utilization threshold is exceeded.

Generic Automation scripts	
seq_srdfa	Sequential automated execution for multiple DGs of SRDF/A control scripts that take a DG as a single argument.

Adding SRDF/A devices	
2site_add_vmax_device	Add new devices to existing RDF group and device group. Configures PIT BCVs for SRDF/A recovery. VMAX only script.

Recovery scripts	
resume_srdfa	Script to resume SRDF/A replication that take a point-in-time copy before the resynchronisation starts

Failover/failback scripts	
2site_async_failover	Controls failover of async RDF groups
2site_async_failback	Control failback of async RDF groups
2site_async_swap	Controls personality swap of async RDF groups
2site_MSC_failover	Control independent failover of an RDF group participating in multi-session consistency replication
2site_MSC_failback	Control independent failback of an RDF group formerly participating in multi-session consistency replication
2site_CG_failover	Control failover of an <i>entire</i> Composite Group participating in multi-session consistency replication
2site_CG_failback	Control failback of an <i>entire</i> Composite Group participating in multi-session consistency replication

24.2 Script mon_srdfa_inactive

24.2.1 Overview

This script will monitor for inactive SRDF/A sessions using symcfg. It accomplishes this by looking for all RDF groups in asynchronous mode that are not active. The script will work for SRDF/A sessions under both MSC and native VMAX replication control.

24.2.2 Invocation and required arguments

```
# Invocation:  mon_srdfa_inactive  <SID> <support_team>
#
# Arguments:   <SID> - Symmetrix ID
#              <support_team> - String with support team name
#                               (e.g. "SAN Support team")
#
# Example:     mon_srdfa_inactive 1380 "SAN support team"
```

24.2.3 Usage notes

Script **mon_srdfa_inactive** should be run from cron at regular intervals on a SE server being used to manage SRDF/A. Once every 5 minutes may be a suitable time interval.

The script will write a message to /var/adm/messages using the logger command if one of the criteria for raising an alert is met.

All messages will include the string SRDFA_ALERT. You will need to instruct your local Enterprise Monitoring (EMS) Team to monitor the /var/adm/messages file on the server running the SRDF/A monitoring script(s) and set their system to raise an alert to the appropriate support team should a message with the string SRDFA_ALERT be detected. The contents of the message should also be forwarded by EMS to the support team.

24.3 Script mon_srdfa_cache

24.3.1 Overview

This script will monitor the symmetrix wide cache that has been assigned for SRDF/A usage. If the usage exceeds a preset percentage, it will log an error.

24.3.2 Invocation and required arguments

```
# Invocation:  mon_srdfa_cache <SID> <cache_threshold>
                                     <support_team>
#
# Arguments:  <SID> - Symmetrix ID
#              <cache_threshold> - % cache threshold
#              <support_team>    - String with support team name
#                               (e.g. "SAN Support team")
#
# Example:    mon_srdfa_cache 1380 80 "SAN support team"
```

24.3.3 Usage notes

Script **mon_srdfa_cache** should be run from cron at regular intervals on a SE server being used to manage SRDF/A. Once every 5 minutes may be a suitable time interval.

The script will write a message to /var/adm/messages using the logger command if one of the criteria for raising an alert is met.

All messages will include the string SRDFA_ALERT. You will need to instruct your local Enterprise Monitoring (EMS) Team to monitor the /var/adm/messages file on the server running the SRDF/A monitoring script(s) and set their system to raise an alert to the appropriate support team should a message with the string SRDFA_ALERT be detected. The contents of the message should also be forwarded by EMS to the support team.

24.4 Script mon_srdfa_cycle

24.4.1 Overview

This script will monitor the cycle time of all SRDF/A session usage. If the time since the last cycle switch exceeds a preset value for any SRDF/A sessions, it will log an error for each such session.

24.4.2 Invocation and required arguments

```
# Invocation:  mon_srdfa_cycle <SID> <cycle_threshold> <support_team>
#
# Arguments:   <SID> - Symmetrix ID
#              <cycle_threshold> - maximum cycle time threshold
#              <support_team>    - String with support team name
#                                (e.g. "SAN Support team")
#
# Example:     mon_srdfa_cycle 1380 60 "SAN support team"
```

24.4.3 Usage notes

Script **mon_srdfa_cycle** should be run from cron at regular intervals on a SE server being used to manage SRDF/A. Once every 5 minutes may be a suitable time interval.

The script will write a message to `/var/adm/messages` using the `logger` command if one of the criteria for raising an alert is met.

All messages will include the string `SRDFA_ALERT`. You will need to instruct your local Enterprise Monitoring (EMS) Team to monitor the `/var/adm/messages` file on the server running the SRDF/A monitoring script(s) and set their system to raise an alert to the appropriate support team should a message with the string `SRDFA_ALERT` be detected. The contents of the message should also be forwarded by EMS to the support team.

24.5 Script `srdfa_health`

24.5.1 Overview

This is the primary SRDF/A monitoring script. It embodies all the functionality of the previous 3 scripts and provides:

- Monitoring for inactive SRDF/A sessions
- Monitoring for cache utilization above a preset threshold
- Monitoring for cycle times above a preset threshold

If script `srdfa_health` is used, (recommended) it is not necessary to run the previous three `mon_srdfa_*` scripts.

24.5.2 Invocation and required arguments

```
# Invocation:   srdfa_health  <SID> <cache_threshold>
                                     <cycle_threshold> <support_team>

#
# Arguments:   <SID> - Symmetrix ID
#               <cache_threshold> - cache threshold (percentage)
#               <cycle_threshold> - cycle threshold (seconds)
#               <support_team>    - String with support team name
#                                   (e.g. "SAN Support team")
#
# Example:     srdfa_health 1380 50 90 "SAN support team"
```

24.5.3 Usage notes

Script **`srdfa_health`** should be run from cron at regular intervals on a SE server being used to manage SRDF/A. Once every 5 minutes may be a suitable time interval.

The script will write a message to `/var/adm/messages` using the `logger` command if one of their criteria for raising an alert is met.

All messages will include the string `SRDFA_ALERT`. You will need to instruct your local Enterprise Monitoring (EMS) Team to monitor the `/var/adm/messages` file on the server running the SRDF/A monitoring script(s) and set their system to raise an alert to the appropriate support team should a message with the string `SRDFA_ALERT` be detected. The contents of the message should also be forwarded by EMS to the support team.

24.6 Script mon_srdfa_dse_pool

24.6.1 Overview

This script will monitor the pool usage of a nominated DSE pool. If the usage exceeds a preset percentage, it will log an error.

24.6.2 Invocation and required arguments

```
# Invocation:  mon_srdfa_dse_pool <SID> <pool_threshold> <dse_pool>
<support_team>
#
# Arguments:  <SID>                - Symmetrix ID
#              <pool_threshold>    - % pool threshold
#              <dse_pool>          - dse_pool name
#              <support_team>      - String with support team name
#                                   (e.g. "SAN Support")
#
# Example:  mon_srdfa_dse_pool 0847 60 DSE "SAN Support team"
#
```

24.6.3 Usage notes

Script **mon_srdfa_dse_pool** should be run from cron at regular intervals on a SE server being used to manage SRDF/A. Once every 5 minutes may be a suitable time interval.

The script will write a message to /var/adm/messages using the logger command if one of the criteria for raising an alert is met.

All messages will include the string SRDFA_ALERT. You will need to instruct your local Enterprise Monitoring (EMS) Team to monitor the /var/adm/messages file on the server running the SRDF/A monitoring script(s) and set their system to raise an alert to the appropriate support team should a message with the string SRDFA_ALERT be detected. The contents of the message should also be forwarded by EMS to the support team.

24.7 Script seq_srdfa

24.7.1 Overview

This script can be used to automate for multiple Device Groups (DGs) the running of other SRDF/A control scripts that take a DG as their single argument. The commands are run sequentially in the order specified in the file containing the list of DGs.

24.7.2 Invocation and required arguments

```
# Invocation:  seq_srdfa <script> <dglist>
#
# Arguments:  <script>    - Script to be run
#              <dglist>   - File with list of DGs in order they
#                           should run
#
# Example:    seq_srdfa 2site_async_failover dglist
```

24.7.3 Usage notes

This script has been provided to help automate the running of the same control script for multiple DGs. The most obvious usage is to automate the failover of multiple DGs in the event of a disaster (e.g. using script 2site_async_failover as in the above example), but it can also be used to speed up house keeping tasks such as resuming SRDF/a sessions after they have dropped.

24.8 Script 2site_add_vmax_device

24.8.1 Overview

This script requires Enginuity 5874 (or later) and SE 7.0 (or later).

This script can be used to add new devices to an existing RDF group with RDF devices in a 2-site SRDF/A configuration without suspending replication. See Section 6.2.3 for further details. The listing for this VMAX only script can be found in Appendix E.

The script also sets up the remote SRDF/A PIT BCV at the async target site to the local device group as a remote BCV.

The script executes the following steps:

- Place new devices in a temporary holding RDF group and synchronize, if not already in that state
- Suspend new devices in holding RDF group with consistency exempt
- Disable consistency for target device group, if enabled
- Move the new device(s) from the holding RDF group to the target RDF group and resume RDF pair
- Add the new device(s) into the local device group (GNS master)
- Wait until RDF is fully synced up and enable consistency
- Add the remote SRDF/A PIT BCV into the device group and establish
- Remind user to update device group at remote site

24.8.2 Invocation and required arguments

```
# Invocation: 2site_add_vmax_device <DGName> <RDF group> <RDF pairsfile>
#              <BCV pairing file> <local SID>
#              <Holding RDF group>
#
# Arguments:  <DGName>           - Device Group Name
#              <RDF group>       - RDF group number
#              <RDF pairsfile>   - RDF pairsfile
#              <BCV pairing file> - BCV pairing file
#              <local SID>       - local Symmetrix ID
#              <Holding RDF group> - Holding RDF group
#
# Example:
#
# 2site_add_vmax_device AppA 12 ./pairsfile_AppA_newdisks./bcv_pairs_newdisks
1380 250
```

24.8.3 Usage notes

The devices defined in <RDF pairsfile> will be synchronized in the “holding” RDF group <Holding RDF group>, with the R1 being on the Symmetrix referred to by the –sid <SID> qualifier. The “holding” RDF group must already exist.

Some sites may pre-synchronize their RDF group pairs in the holding RDF group in advance. Hence, this script tests if the devices to be added to the target RDF group are already in the holding RDF group **and also already synchronized**. If they are, then the above createpair step is skipped. If the RDF pairs to be allocated to the application are already in the holding RDF group, but NOT

synchronized (e.g. in suspended state), then they will need to be manually put into synchronized first before the script is run, else the script will fail and exit when the createpair step fails.

The script also requires the usual RDF pairs file <RDF pairsfile> that defines which local and remote devices are to be formed into an R1-R2 pair. One row is needed for each R1-R2 pairing. The first entry in a row is the Symmetrix Volume ID of the device in the local Symmetrix to be used as the R1. The second entry is the Symmetrix Volume ID of the device in the remote box to be used as the R2. Only metaheads should be included in the file.

Note that the volume ID for the R1 and R2 will usually be the same. For example

```
0001 0001
0005 0005
```

The script also requires an input file <BCV pairing file> that defines which BCV is to be paired with which async target R2. One row is needed for each STD-BCV pairing. The first entry in a row is the Symmetrix Volume ID of the R2 and the second entry is the Symmetrix Volume ID of the BCV to be associated with the R2. Only metaheads should be included in the file.

For example the following BCV pairing file would pair remote BCV 0521 with R2 00D1 and remote BCV 0525 with R2 00D5.

```
00D1 0521
00D5 0525
```

24.9 Script resume_srdfa

24.9.1 Overview

When using Solutions Enabler 6.3.1 or later, the program symrecover can be used instead of resume_async (symrecover was not available when resume_async was first written). See earlier section for further details.

Prior to starting resynchronisation after SRDF/A link recovery, a BCV copy of the R2 volumes should be split off (consistently) to ensure that there is a consistent (albeit aging) point-in-time copy at all times in the remote site during the re-synchronization process when the R2 will be inconsistent.

If the link between sites is down for an extended period, and there is a large build up of invalid tracks on the R1 side, resumption of the link with SRDF/A (the default behaviour in Open Systems) can cause a heavy surge of link traffic, created by the large backlog of invalid tracks being added to the cycles generated by the production traffic. This can lead to an unacceptably high demand being placed on Symmetrix cache, causing SRDF/A to drop again. To avoid this risk, it is best to briefly disable SRDF/A when restarting after a large build up of invalid tracks.

This script handles the creation of the PIT copy and also the switch into and out of adaptive copy mode before and after resynchronisation takes place.

The script executes the following steps:

- Verify remote BCV is fully established
- Split remote BCV
- Wait until finished splitting in the background
- Set mode to adaptive copy
- Establish RDF
- Wait until RDF is full synced up
- Set mode to async mode
- Enable consistency
- Establish remote BCV
- Wait until remote BCV fully established before exiting

24.9.2 Invocation and required arguments

```
# Invocation:  resume_async <DGName> <RDF group>
#
# Arguments:   <DGName>          - Device Group Name
#              <RDF group>       - RDF group number of async link
#
# Example:     resume_async App1 4
```

24.9.3 Usage notes

The script requires that the PIT BCV is established. Note this script will gracefully exit if BCVs are not established as this may indicate that the SRDF/A outage occurred during a previously incomplete resync operation and hence the BCV should NOT be established against a potentially inconsistent R2. Alternatively it may indicate that they are being used for COB test purposes.

24.10 Script 2site_async_failover

24.10.1 Overview

This script executes an SRDF failover down the async leg in a 2-site SRDF/A configuration.

The script executes the following steps:

- Temporarily disable consistency on the device group
- Confirming no outstanding tracks not committed to async R2 side
- Failover the device group

24.10.2 Invocation and required arguments

```
# Invocation:      2site_async_failover <DGname>
#
# Arguments:      <DGName>      - Device groups to be failed over
#
# Example:        2site_async_failover DG21
```

24.10.3 Usage notes

Trivial script included for completeness.

24.11 Script 2site_async_failback

24.11.1 Overview

This script executes an SRDF failback down the async leg in a 2-site SRDF/A configuration.

The script executes the following steps:

- Failback RDF group
- Wait until there are no invalid tracks on the async leg
- Renable consistency

24.11.2 Invocation and required arguments

```
# Invocation:      2site_async_failback <DGname>
#
# Arguments:      <DGName>      - Device groups to be failed back
#
# Example: 2site_async_failback DG21
```

24.11.3 Usage notes

The script will loop whilst the R1 is updating from the R2 and not exit until this process has completed.

24.12 Script 2site_async_swap

24.12.1 Overview

With dynamic RDF, you can swap the R1/R2 personality of RDF pairs. Source R1 devices become target R2 devices, and vice versa. Swapping SRDF devices allows the original R2 side to take over operations while retaining a remote mirror on the original R1 side i.e. the original R1 becomes an R2 and vice versa (see figure below).

This is used in some regions to ensure that when a failover to COB has occurred, that off-site COB is still available (i.e. back to the original production site). For example, in the diagram below, if an application that usually runs in Site A has failed over to Site B due to server hardware failure, then personality swapping ensures that synchronous off-site replication still continues.

This script conducts a personality swap of an RDF group running SRDF/A

The script executes the following steps:

- Temporarily disable consistency on the device group
- Confirming no outstanding tracks not committed to async R2 side
- Execute failover
- Execute swap
- Execute establish
- Enable consistency on the device group
- Print out current RDF group type

24.12.2 Invocation and required arguments

```
# Invocation:      2site_async_swap <DGname>
#
# Arguments:      <DGName>      - Device groups to be swapped
#
# Example:  2site_async_swap DG21
```

24.12.3 Usage notes

Note that “fast-swap” (failover –establish) is not supported on SRDF/A enabled devices and hence is not used in this script.

24.13 Script 2site_MSC_failover

24.13.1 Overview

One of the benefits of Multi-Session Consistency (MSC) in an SRDF/A environment is that it allows one to guarantee IO consistency at the R2-end between RDF groups (and hence individual applications), whilst still permitting failover independence for each application.

This script executes an SRDF failover **of a single RDF group** participating in Multi-Session Consistency (MSC) replication down the async leg in a 2-site SRDF/A configuration.

This is an unusual process, as it requires that the RDF group to be failed over be removed from the Composite Group being used for MSC replication before the failover can take place. MSC replication of the remaining RDF groups in the re-defined Composite group is then restarted.

The script executes the following steps:

- Temporarily disable consistency on the composite group (CG)
- Remove devices to be failed over from CG
- Re-enable consistency on the composite group (CG)
- Confirm no outstanding tracks not committed to async R2 side
- Failover the RDF group (DG)

24.13.2 Invocation and required arguments

```
# Invocation: 2site_MSC_failover <CGName> <DGname>
               <RDF group to failover>
#
# Arguments:  <CGName>      - Composite Group Name
#             <DGname>      - Device group to be failed over
#             <RDF group to failover>
#                               - RDF group number to be failed over
#
# Example:    2site_MSC_failover App1CG DG22 22
```

24.13.3 Usage notes

No special notes

24.14 Script 2site_MSC_failback

24.14.1 Overview

One of the benefits of Multi-Session Consistency (MSC) in an SRDF/A environment is that it allows one to guarantee IO consistency at the R2-end between RDF groups (and hence individual applications), whilst still permitting failover independence for each application.

This script executes an SRDF failback **of a single RDF group** formerly participating in Multi-Session Consistency (MSC) replication down the async leg in a 2-site SRDF/A configuration.

This is an unusual process, as it requires that the RDF group to be failed over be replaced into the Composite Group being used for MSC replication after the failover has taken place. MSC replication of the all RDF groups in the Composite Group is then restarted.

The script executes the following steps:

- Failback the device group (DG)
- Wait until there are no invalid tracks on the async leg
- Temporarily disable consistency on the composite group (CG)
- Add devices from RDF group just failed back into CG
- Re-enable consistency on the composite group (CG)

24.14.2 Invocation and required arguments

```
# Invocation: 2site_MSC_failback <CGName> <DGname>
               <RDF group to failback>
#
# Arguments:  <CGName>      - Composite Group Name
#              <DGname>      - Device groups to be failed back
#              <RDF group to failback>
#                               - RDF group number to be failed back
#
# Example: 2site_MSC_failback App1CG 22
```

24.14.3 Usage notes

No special notes

24.15 Script 2site_CG_failover

24.15.1 Overview

An entire Composite Group can also be failed over. This would be required in the case where there has been an issue at the source site and DR is being invoked. In such cases, replication usually already have ceased due to the MSC Consistency Group tripping. Failover via the Composite group (rather than via the underlying RDF groups) ensures consistency at the R2 end.

This script executes a Composite Group failover down the async leg in a 2-site SRDF/A configuration.

The script executes the following steps:

- Failover the composite group

24.15.2 Invocation and required arguments

```
# Invocation:      2site_CG_failover <CGname>
#
# Arguments:      <CGName>                - Composite group to be failed
over
#
# Example:      2site_CG_failover App1CG
```

24.15.3 Usage notes

When MSC replication is halted due to the MSC Composite Group tripping, a clean up process is required to determine if the receive Delta Sets (on the Target VMAX) for each RDF group should be applied to the R2 or discarded. A description of this process is beyond the scope of the present document, but can be found in Appendix H of the SRDF/A Evaluation document.

The failover command uses the `-force` and `-symforce` options. These are sometime required if the failover is occurring after a rolling disaster when the MSC has tripped and replication has stopped. The SRDF/A MSC cleanup process can result in Delta Sets being discarded on the R2 side leaving the Symmetrix in a state with invalids on both sides. For further details see Appendix H of the SRDF/A Evaluation document.

24.16 Script 2site_CG_failback

24.16.1 Overview

This script executes a Composite Group failback down the async leg in a 2-site SRDF/A configuration.

The script executes the following steps:

- Failback the composite group

24.16.2 Invocation and required arguments

```
# Invocation:      2site_CG_failback <CGname>
#
# Arguments:      <CGName> - Composite group to be failed back
#
# Example:        2site_CG_failback App1CG
```

24.16.3 Usage notes

The `-force` flag is needed as consistency is still necessarily enabled on the Composite Group.