

# **GNAST DESIGN and CONFIGURATION GUIDE**

**For**

**EMC VMAX3**

**CSBE Storage and Backup Engineering**

**Template Version: 2.0.7**

**Document Version#: 0.5**

**May 15, 2015**



## **Disclaimer**

If printed, content is valid only for the day it was printed. Always refer to the online document for the current version. If the user elects to maintain copies of this procedure, it is their responsibility to verify the currency of the document by checking it against the online version. This document must be promptly removed from use when obsolete.

## Table of Contents

|          |                                                            |           |
|----------|------------------------------------------------------------|-----------|
| <b>1</b> | <b>DOCUMENT PURPOSE .....</b>                              | <b>4</b>  |
| <b>2</b> | <b>OWNER OF STANDARD .....</b>                             | <b>4</b>  |
| <b>3</b> | <b>INTENDED AUDIENCE .....</b>                             | <b>4</b>  |
| <b>4</b> | <b>PRINCIPLES.....</b>                                     | <b>5</b>  |
| 4.1      | Revision History .....                                     | 5         |
| 4.2      | Key Feature Changes.....                                   | 5         |
| <b>5</b> | <b>ASSEMBLY/SOLUTION DESCRIPTION .....</b>                 | <b>5</b>  |
|          | <b>SCOPE.....</b>                                          | <b>6</b>  |
| 5.1      | Technology.....                                            | 7         |
| 5.2      | Applicability .....                                        | 7         |
| <b>6</b> | <b>COMPONENTS .....</b>                                    | <b>7</b>  |
| 6.1      | List of Components .....                                   | 7         |
| 6.2      | Operational Tools.....                                     | 8         |
| <b>7</b> | <b>ASSEMBLY/SOLUTION ARCHITECTURE .....</b>                | <b>8</b>  |
| 7.1      | VMAX3 Architecture.....                                    | 8         |
| 7.1.1    | VMAX Product Components .....                              | 9         |
| <b>8</b> | <b>INSTALLATION &amp; CONFIGURATION INSTRUCTIONS .....</b> | <b>12</b> |
| 8.1      | Installation of Components .....                           | 13        |
| 8.1.1    | Physical VMAX3 Storage Array Installation .....            | 13        |
| 8.1.2    | VMAX3 Hypermax OS Code.....                                | 13        |
| 8.1.3    | VMAX3 MMCS Ethernet Connections.....                       | 13        |
| 8.1.4    | VMAX3 SAN Connections .....                                | 13        |
| 8.1.5    | Solutions Enabler Management Server.....                   | 14        |
| 8.1.6    | Unisphere for VMAX Manager Server .....                    | 15        |
| 8.1.7    | VMAX3 Licensing .....                                      | 15        |
| 8.1.8    | Storage Configuration.....                                 | 15        |
| 8.1.9    | Storage Provisioning.....                                  | 16        |
| 8.1.10   | Event Notification (Tivoli) .....                          | 19        |
| 8.1.11   | Log Reconciliation (SEMs/ArcSight).....                    | 19        |
| 8.1.12   | Automated SBCC (HPNA) .....                                | 19        |



8.1.13 Chargeback and Performance Reporting (EMC ViPR SRM) .....20

8.1.14 MMCS Anti-Virus .....20

**9 CONFIGURATION ASSURANCE ..... 21**

**10 SCALABILITY & CAPACITY ..... 21**

**11 AVAILABILITY MANAGEMENT ..... 22**

**12 REGIONAL OR MERGER PARTNER CONSIDERATION ..... 22**

**13 GLOBAL STANDARD ..... 22**

**14 RELEVANT LINKS ..... 22**

**15 APPENDIX ..... 23**

**Disclaimer**

If printed, content is valid only for the day it was printed. Always refer to the online document for the current version. If the user elects to maintain copies of this procedure, it is their responsibility to verify the currency of the document by checking it against the online version. This document must be promptly removed from use when obsolete.

**Note:** The purpose of this template is to provide a standard document structure for completing designs with detailed instructions on how to install and configure secure technology standards within Citi's global environment. This document follows the official standard GNAST design framework, and applies to all current GNAST technology standards/products inclusive of hardware and software. This template also allows usage of the Secure Build Configuration Checklist (SBCC) into one GNAST Configuration Template.

In some sections, this template can be used as a "pointer" to more widely used guides and SBCC documents (i.e. SBCC). In other cases, it might make sense to use this template and expand as needed.

## 1 Document Purpose

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.

## 2 Owner of Standard

GNAST (Global Network Architecture and Standards Team) is the owner of this standard and the related documents. CATE (Citi Architecture & Technology Engineering) is the custodial organization. CitiNetwork & CitiStorage Platform Engineering technologists will drive and generate content.

## 3 Intended Audience

The intended audience is CitiNetwork & CitiStorage Engineering, Global and Regional Operations and Engineering personnel responsible for designing, building, supporting, and managing Citi corporate network environments.

## 4 Principles

### 4.1 Revision History

Update the document history each time a draft is issued for review.

Document Version:

A first generation document that hasn't been ratified before is a 0.0 (less than 1.0) during its revision cycle. When adopted it is a 1.0 version. Subsequent generations follow the same method of incrementing the decimal portion of the version number until ratified, and then increment the integer.

| Version | Description                                  | Date       | Author/Organization         |
|---------|----------------------------------------------|------------|-----------------------------|
| 0.1     | Initial Draft                                | 03-27-2015 | Quang Le                    |
| 0.2     | Peer Review                                  | 04-02-2015 | Quang Le,<br>Ludwig Derodel |
| 0.3     | Release for Testing                          | 04-03-2015 | Quang Le                    |
| 0.4     | Added additional McAfee<br>ePO instructions  | 05-07-2015 | Quang Le                    |
| 0.5     | Modified standard hardware<br>configurations | 05-13-2015 | Quang Le                    |

### 4.2 Key Feature Changes

The EMC VMAX3 storage array is the latest generation hardware platform for EMC's Symmetrix product line. The Symmetrix product line is certified within Citi for enterprise tier storage.

## 5 Assembly/Solution Description

The VMAX3 is a monolithic storage array which provides shared storage to host servers via a storage area network (SAN). The main components of the VMAX3 are disk drives, engines, and an Infiniband fabric called the Virtual Matrix Interconnect used for inter-engine transmissions.

The VMAX3 supports 6 Gbps SAS and Flash SAS drives in 2.5" and 3.5" form factor. The drives are housed in either a 120 drive (for the 2.5" drives) or 60 drive (for the 3.5") disk array enclosure (DAE). Up to 4 DAEs can be connected to each engine; the first 2 DAEs will be directly connected to the engine and next 2 will daisy chained.

The engines in the VMAX3 provides cache memory to the storage system, back-end connectivity to the disk drives, front-end connectivity to host servers, and handles all I/O processing. The engines contain two director boards which house the I/O modules for front-end and back-end connectivity, vault flash drives, and the Management Module Control Stations (MMCS). There are 2 MMCSs installed in the 1<sup>st</sup> engine for each VMAX3 storage array. The MMCS is used by the vendor to access and manage the system either by local connection or remotely via [EMC Secure Remote Support \(ESRS\)](#). One MMCS will be the primary and the other will be a standby in case the primary MMCS fails.

If there's more than one engine in the VMAX3, 2 Infiniband switches will be installed for inter-engine transmissions (control, metadata, and user data). Each port on the Infiniband switch operates at 56 Gbps. Infiniband switches are not installed on single-engine systems.

A fibre channel SAN is required for host connections to the VMAX storage array. Reference the [Cisco Fibre Channel Switch Product Page](#) for SAN configuration guidelines and standards. Zoning standards for VMAX3 will be covered in a later section in this document. Host SAN stack configuration guidelines and standards can be found on the [SAN Stacks Product Page](#).

A RHEL SOE management server will need to be setup to manage the VMAX3 storage array. This host will need to be SAN connected to the VMAX3 and will need VMAX3 control devices called gatekeepers provisioned for in-line management communication. The management server will also be used to monitor the storage array and send alerts to the Tivoli Monitoring system. Audit logs will also be captured by the management server and sent to the SEMS audit logging system for audit log review.

A Windows SOE management server will also need to be setup to run the Unisphere for VMAX software. Unisphere is a graphical user interface used to manage and monitor the VMAX3 storage array. Unisphere is required if performance monitoring is being used. To alleviate the need for SAN connectivity for the Unisphere server; the Unisphere server will need the Solutions Enabler software installed and configured to remotely connect to the RHEL SE server through symapi.

## Scope

EMC's VMAX3 Storage array and supporting management products was released by EMC on September 25<sup>th</sup> 2014, however certain key features used in Citi will not be available until a code release sometime in Q4 – as such this certification effort will be split into phases to reflect the products readiness.

### Phase1:

- VMAX3 Storage array
  - o Configured for Non Replicated storage
- Management platforms
  - o Solutions Enabler

**Phase2:**

- VMAX3 Storage array
  - o Configured for Replicated Storage
  - o SNAPVX
- Management platforms
  - o Unisphere for VMAX
  - o SMI-S

**5.1 Technology**

The VMAX3 is a storage platform that will provide shared block-based storage via a fibre channel SAN.

**5.2 Applicability**

Global datacenters requiring block-based enterprise tier storage over a SAN.

**6 Components****6.1 List of Components**

The following tables list the required components and versions for the VMAX3 solution.

| Component Name                           | Capability Description                                                                                                                |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Base Components                          |                                                                                                                                       |
| Hypermax OS                              | 5977 is the certified version of VMAX3 firmware                                                                                       |
| Solutions Enabler                        | Solutions Enabler 8.x CLI management software                                                                                         |
| RHEL Solutions Enabler Management server | Fiber attached SOE RHEL server running Solutions Enabler management software. Used for management, logging, monitoring, and alerting. |
| Windows Unisphere server                 | SOE Windows server running Unisphere and Solutions Enabler. Used for management, monitoring, and reporting.                           |
| ESRS                                     | ESRS is used to provide remote management of the VMAX3 storage array via the IP network.                                              |
|                                          |                                                                                                                                       |

GNAST Bill of Material (BoM) NOTE: Restricted access:

<https://catecollaboration.citigroup.net/domains/networks/bom/default.aspx>

## 6.2 Operational Tools


The following tools listed in the table will have each VMAX storage array integrated into their infrastructure to help with monitoring the solution.

| Name                               | Capability Description               | Result |
|------------------------------------|--------------------------------------|--------|
| Tivoli                             | Event Notification                   |        |
| Hewlett Packard Network Automation | SBCC Automation                      |        |
| Arcsight                           | SEMS Audit Log Reconciliation        |        |
| EMC SRM                            | Chargeback and Performance reporting |        |
| McAfee Virus Scan                  | Anti-Virus for the MMCS              |        |

## 7 Assembly/Solution Architecture

### 7.1 VMAX3 Architecture

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.



### 7.1.1 VMAX Product Components

**VMAX3 Storage array** – configured with 2.7GHz, 48 Core Ivy-Bridge engines with 2TB of memory per engine (1TB usable since cache is mirrored).

VMAX3 storage array supports performance, and capacity disks to maximize array performance across all drive types utilizing Fully Automated Storage Tiering (FAST) to fine tune performance according to application workload requirements.




Each system bay can house 1-2 engines. Each 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 connectivity<sup>1</sup>. 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

---

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

engine. The picture below depicts the components of a VMAX3 engine along with the logical port numbers for the front-end connections.



(EMC, 2015)

Single-engine system bays can contain 2-6 disk array enclosures (DAE). Dual-engine system bays can contain 2-4 DAEs. The DAEs come in two sizes, 120 drives (for 2.5" drives) and 60 drives (for 3.5" drives). The Citi standard configuration 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.



System bay also contains power supplies with LI-ON battery backup, Virtual matrix infiniband interconnect, and component cabling.

## 8 Installation & Configuration Instructions

The following sections of this document discuss the installation and configuration steps of the VMAX3 storage array. The following are the high-level steps to installing and configuring the VMAX3 storage array:

1. Physical installation of the VMAX3 storage array
2. VMAX3 Hypermax OS Code
3. Connect VMAX3 Management connections
4. Setup Vendor Remote Management (ESRS)
5. Connect VMAX3 to SAN
6. Setup Management Infrastructure
7. Licensing
8. Provision storage to host servers

## 8.1 Installation of Components

### 8.1.1 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
```

### 8.1.2 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.

### 8.1.3 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.

### 8.1.4 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:



### **Zoning**

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).

### **8.1.5 Solutions Enabler Management Server**

The RHEL 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 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. Reference the Solutions Enabler 8.x Solutions Guide for instructions on Solutions Enabler configuration and setup. The following high-level steps will need to be completed in order for the SE server to manage the VMAX3:

- 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 managements servers and the server 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.

**Example:** # symconfigure -sid 973 commit -cmd "set port 1d:4 show\_aclx\_device=disable;"

**Note:** The 'Show\_ACLX\_Device' flag is different from the 'Access\_Logix(ACLX)' flag. The 'Access\_Logix(ACLX)' should be enabled by default on all FA ports which allows devices to be mapped to FA ports using symaccess. The 'map' and 'unmap' commands will no longer work on ports with the 'Access\_Logix(ACLX)' flag enabled.

### 8.1.6 Unisphere for VMAX Manager Server

Unisphere for VMAX is EMC's graphical user interface (GUI) used to manage Symmetrix storage arrays. Unisphere can be installed on a Citi SOE Windows server and requires the Solutions Enabler software to discover and manage the storage arrays. Unisphere is needed for VMAX3 performance monitoring and ViPR SRM reporting. Reference the CSBE Unisphere for VMAX Solutions Guide for installation and setup instructions.

### 8.1.7 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:

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

The following command to list the licenses on the VMAX3:

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

### 8.1.8 Storage 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.

### 8.1.9 Storage Provisioning

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 namgdse14\_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 namgdse14\_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 namgdse14\_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 namgdse14\_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 will 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 namgdse14\_mv -sg namgdse14\_sg -pg 1d7-1d7-3d7-4d7\_pg -ig namgdse14\_ig

## **8.1.10 Event Notification (Tivoli)**

The Tivoli Monitoring system will be used to alert the Citi storage teams for component failures or threshold alerts. Tivoli alerting is setup on the RHEL SE management server. Refer to the CSBE Solutions Enabler Solutions Guide for setup instructions.

## **8.1.11 Log Reconciliation (SEMs/ArcSight)**

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. Refer to the CSBE Solutions Enabler Solutions Guide for setup instructions.

## **8.1.12 Automated SBCC (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 the Solutions Enabler Solutions Guide for instructions on setting up HPNA on the Solutions Enabler management server to verify the VMAX3 configuration.

### 8.1.13 Chargeback and Performance Reporting (EMC 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 with need to be created for ViPR SRM and assigned the symauth role of Monitor for SRM to gather the data.

### 8.1.14 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<br>TCP/18443 |
| ePO server | VMAX3 MMCS  | TCP/18081<br>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\\_contacts.aspx](http://catecollaboration.citigroup.net/domains/security/isp_standards/hsestds/Antivirus%20%20McAfee%20ePolicy%20Orchestrator/AV%20Documentation/ePOusefulinfo_contacts.aspx). The following should be done after following the instructions in the document:
  - 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.

## 9 Configuration Assurance

See section [8.1.12](#) Automated SBCC (HPNA).

## 10 Scalability & Capacity

The following VMAX3 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) |
|--------|-------|--------------|--------------------|------------------------------------------------|----------------------|
| Silver | 100K  | 2            | 1TB                | 32 x 800GB/EFD/R5(3+1),<br>432 x 600GB/SAS/R1  | 130                  |
| Silver | 200K  | 2            | 2TB                | 64 x 800GB/EFD/R5(3+1),<br>670 x 600GB/SAS/R1  | 210                  |
| Silver | 200K  | 2            | 2TB                | 64 x 800GB/EFD/R5(3+1),<br>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.

## 11 Availability Management

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.

## 12 Regional or Merger Partner Consideration

None.

## 13 Global Standard

The VMAX3 uses both host-based access control and user-based access control to ensure that only authorized hosts and users can access the system for management. Host-based access control is handled via symacl which limits management connectivity to specified hosts with granular control over what actions that host can make on storage system. User-based access is handled by symauth which limits authenticated users to actions for the symauth role they are assigned. Users are authenticated to the management systems using LDAP and Active Directory on the SOE management servers for VMAX3. Reference the Solutions Enabler Solutions Guide for more details on symacl and symauth configuration.

## 14 Relevant Links

| Document/Web Site                                                               | Location                                                    |
|---------------------------------------------------------------------------------|-------------------------------------------------------------|
| EMC VMAX Family (100K, 200K, 400K) Documentation Set                            | <a href="http://support.emc.com">http://support.emc.com</a> |
| EMC VMAX3 Family with HYPERMAX OS 5977 Release Notes                            | <a href="http://support.emc.com">http://support.emc.com</a> |
| Solutions Enabler Symmetrix Array Management CLI Product Guide                  | <a href="http://support.emc.com">http://support.emc.com</a> |
| EMC VMAX3 Family New Features – A Detailed Review for Open Systems Environments | <a href="http://support.emc.com">http://support.emc.com</a> |

|                                                                                           |                                                             |
|-------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| EMC Simple Suoport Matrix (EMC VMAX 400K, VMAX 200K, and VMAX 100k) Director Bit Settings | <a href="http://support.emc.com">http://support.emc.com</a> |
| EMC VMAX Family (VMAX 100K, 200K, 400K) Planning Guide                                    | <a href="http://support.emc.com">http://support.emc.com</a> |
| EMC VMAX Family Security Configuration Guide                                              | <a href="http://support.emc.com">http://support.emc.com</a> |
| EMC VMAX Reliability, Availability, and Serviceability Technical Notes                    | <a href="http://support.emc.com">http://support.emc.com</a> |
| Drive Sparing in EMC Symmetrix VMAX Family Systems                                        | <a href="http://support.emc.com">http://support.emc.com</a> |
| EMC VMAX3 Service Level Provisioning with Fully Automated Storage Tiering (FAST)          | <a href="http://support.emc.com">http://support.emc.com</a> |
| VMAX Standard Build Configuration Checklist (SBCC)                                        | <a href="#">CSBE Symmetrix Product Page</a>                 |
| CSBE Solutions Enabler 8.0.2 Solution Guide                                               | <a href="#">CSBE Symmetrix Product Page</a>                 |
|                                                                                           |                                                             |

## 15 Appendix