



CATE Storage and Backup Engineering
Citigroup Architecture & Technology Engineering

EMC XtremIO

CSBE-14008-SG

Solutions Guide

Version 1.1

©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	26-Feb-2014	Jeff Hernandez	Ludwig Derodel Bob Small Tim Armstrong	CATE Storage Review
0.2	28-Feb-2014	Jeff Hernandez	Global Stakeholders	Release for Comments
0.3	23-Jul-2014	Jeff Hernandez	Global Stakeholders	Release for Comments
1.0	4-Nov-2014	Jeff Hernandez	Global Stakeholders	Release for Production
1.1	15-Jan-2015	Jeff Hernandez	Global Stakeholders	Release for Production – 3.0.1 Bash Vulnerability Fix

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	6
8	Assumptions and References	7
8.1	Assumptions	7
8.2	How to use this document	7
8.3	References	7
9	General Design	8
9.1	Overview	8
9.2	Citi Support and Management	8
9.2.1	Array Management	8
9.2.2	Enterprise Monitoring and Alerting	8
9.2.3	Email Alerting	9
9.2.4	Audit Logs	9
9.2.5	Performance	9
9.3	Remote Vendor Support	9
9.3.1	Email Alerting	9
9.3.2	EMC Secure Remote Support (ESRS)	9
9.4	Host Connectivity	10
9.5	Software and Features	10
9.6	XtremIO Hardware	11
9.6.1	System Overview	11
9.6.2	XtremIO Cluster Configurations	13
9.6.3	XtremIO Disk Array Enclosure	13
9.6.4	XtremIO Power & Weight Specifications	13
9.6.5	XtremIO Connectivity	13
10	Prerequisites for a new XtremIO Environment	14
10.1	Required Licensing for Array	14
10.2	Required Installation Steps	15
11	XtremIO System Features	15
11.1	Thin Provisioning	15
11.2	SAS Back-end	15
11.3	High Availability	15
11.4	XtremIO Data Protection	16
12	Storage Provisioning Best Practices	17
12.1	Thin Provisioning	17
12.1.1	Volume (LUN) Creation	17
12.2	Initiator Groups	18
12.2.1	Adding Initiators	18
12.3	LUN Mapping	20
12.3.1	LUN Expansion and Shrinking	20
12.3.2	LUN Sizing Guidelines	21
12.3.3	Monitoring XtremIO Capacity with XMS Dashboard	21
12.3.4	Storage Capacity Alerting	22
12.3.5	VAAI Thin Provision Limit	23
12.3.6	Capacity Monitoring with SRM	24
12.3.7	LUN Migration	24
13	Data Protection Best Practices	25
13.1	Local Protection	25

13.2	Remote Protection.....	25
13.3	Backup Standards.....	25
14	Security Best Practices.....	26
14.1	Security Standards.....	26
14.2	Audit Logging.....	26
14.3	Data Security.....	26
14.3.1	Disk Sanitization.....	27
14.3.2	Disposal of failed disks.....	27
14.4	CITMP Information Security Policy Convergence.....	27
14.4.1	Information Ownership.....	27
14.5	Role Based Access Controls (RBAC).....	27
14.5.1	Roles Defined in XMS.....	27
14.5.2	Roles defined for use in Citi.....	28
14.5.3	Adding a Domain User.....	28
14.5.4	Functional IDs.....	31
14.5.5	Root ID Access.....	32
14.5.6	Access Control List (ACL) with IPTables.....	32
15	Alerting and Monitoring.....	35
15.1	SNMP and Email Alerting.....	35
15.1.1	Adding an Array to the Monitoring System.....	35
15.1.2	SNMP alerting with Tivoli NetCool.....	35
15.1.3	Adding an SNMP Recipient and SNMP Properties.....	35
15.1.4	SNMP MIB.....	37
15.2	Call Home and Email Alerting.....	38
15.3	Remote Support (EMC ESRS).....	38
15.4	Dossier Collection.....	38
15.5	Event Monitor.....	39
15.6	Performance Monitoring using XMS.....	40
15.6.1	Exporting Performance History.....	40
15.6.2	Advanced Real Time Monitors.....	41
16	Host Connectivity and Configuration.....	43
16.1	IP Fabric and iSCSI Guidelines.....	44
16.2	iSCSI Portal Tables.....	45
16.3	iSCSI Ports Configuration.....	45
17	Local Protection.....	46
17.1	Snapshots.....	46
18	XMS Management Server.....	47
18.1	XMS Installation Procedure.....	47
19	Appendix A: Storage Controller Failover.....	49
20	Appendix B: Upgrade Procedures.....	54
21	Document Version Control.....	55

7 Introduction

At this time the XtremIO array is being certified as restricted only for use with VDI deployments.

This document should serve as a guide when configuring XtremIO storage arrays in Citi's storage environment for block-based storage utilizing IP Fabrics. The information compiled in this document has been tested and certified by CATE Storage and Backup Engineering.

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 XtremIO storage arrays and their usage at Citi. However, the document does not cover the internals of the XtremIO.

This document is intended for use by all offices within Citi who intend to purchase and implement an XtremIO 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
CATE Storage and Backup Engineering site	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/default.aspx
CATE Storage ESRS standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/esrs.aspx
Citi DNS and NTP configuration standards	http://ipns.ny.ssmb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc

9 General Design

This section explains the general design for the implementation of an EMC XtremIO for block-based storage, utilizing IP fabric for host and storage connectivity.

9.1 Overview

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



9.2 Citi Support and Management

9.2.1 Array Management

The XtremIO storage array is managed via a stand-alone Linux-based virtual appliance, called the XtremIO Management Server (XMS), which can be accessed by typing the XMS management IP address in a browser (recommended). Each XtremIO cluster requires its own XMS host which can be a physical or virtual server; a virtual server is required at Citi. Please refer to the [XMS Management Server](#) section of this guide for further details.

9.2.2 Enterprise Monitoring and Alerting

SNMP alerting should be configured by Citi Storage team to send appropriate SNMP traps to the Enterprise Monitoring System. The monitoring infrastructure can then filter the alerts to generate notifications or problem tickets as appropriate. Please refer to the [Alerting and Monitoring](#) section of this guide for further details.

9.2.3 Email Alerting

Email alerting should be configured by Citi Storage Integration team to send notifications to the appropriate Citi Storage Support team. Please refer to the [Alerting and Monitoring](#) section of this guide for further details.

9.2.4 Audit Logs

Audit logs should be configured by Citi Storage Integration personnel in order to forward XMS syslog notifications to the appropriate Citi SEMS environment to meet compliance requirements. Please refer to the [Audit Logging](#) section of this guide for further details

9.2.5 Performance

The XMS dashboard provides storage system performance statistics in real time as well as a 7 day history. These may be leveraged to improve performance and assist in determining where issues may be occurring. Please refer to [Performance Monitoring using XMS](#) for configuration information.

9.3 Remote Vendor Support

9.3.1 Email Alerting

Email home should be configured by EMC certified engineer to use Citi's SMTP server to send the alerts to EMC support. Please refer to the [Call Home and Email Alerting](#) section for details.

9.3.2 EMC Secure Remote Support (ESRS)

XtremIO utilizes the ESRS Gateway for remote support activities. EMC support will connect to any XtremIO cluster via the ESRS Gateway for remote troubleshooting and problem resolution. Please refer to the [Remote Support](#) section for details.

9.4 Host Connectivity

XtremIO has initially been certified for the VDI use case only. It is recommended that hosts be configured with redundant 10GbE network adapters connected via redundant IP fabric switches, with paths to both storage controllers of the X-brick. Dual non-routed vlans should be configured for path redundancy. The X-brick supports fiber channel connectivity to redundant SANs; however, this configuration is currently restricted within Citi. Please refer to the Host Connectivity and Configuration section for details.

9.5 Software and Features

The table below lists the XtremIO Software components and their current status in Citi.

Feature	Item	Function	Citi Status
XtremIO RTU OE	XtremIO Right to use	Licensed per X-brick cluster	Supported
XtremIO Mgmt	XtremIO Management license	Licensed per X-brick cluster	Supported
XIOS	XtremIO Operating System	XIOS runs on storage controllers and manages system functions	Supported
XMS	XtremIO Management Server	Virtual management server with web-based interface and CLI used to manage X-brick cluster. (Physical XMS server is not supported)	Supported
Thin Provisioning	Native thin provisioning	All volumes have thin provisioning enabled	Supported
Inline Data Reduction	In-line de-duplication	Always on deduplication	Supported
Local Protection	Snapshots	Block storage snapshots	Restricted
Remote Protection	None	Array to array block storage replication	Not Supported
Reporting	SRM	Centralized reporting – SRM (formerly Watch4Net)	Supported
Monitoring	SNMP Alerting	Centralized Enterprise Monitoring	Supported
Audit Log	Syslog Forwarding	Ability to forward audit log for SEMS reconciliation	Supported
Vendor remote support	ESRS for XtremIO	Monitors XtremIO and notifies vendor of error events	Supported
Path Management	NMP, MPIO, MPxIO, DMP, Powerpath	Host Multi-Pathing – ESX NMP is only certified use case at this time	Supported
Host/Array Connectivity	iSCSI 10GbE	Supported connectivity protocol	Supported
Host/Array Connectivity	FC 8Gb	Restricted connectivity protocol	Restricted

Note: With the latest service pack release, support for the following is not available: remote replication and centralized management.

9.6 XtremIO Hardware

9.6.1 System Overview

The XtremIO Storage Array is an all flash system, based on scale-out architecture. The system uses building blocks called X-Bricks which can be clustered together to grow performance and capacity as required. Currently single X-Brick systems are being certified and NDU expansion is not yet supported by EMC.

Each X-brick is comprised of:

- One 2U Disk Array Enclosure (DAE), containing:
 - 25 eMLC SSDs (standard X-Brick 10TB) or 13 eMLC SSDs (5TB Starter X-Brick)
 - Two redundant power supply units (PSUs)
 - Two redundant SAS interconnect modules
- Two 1U Battery Backup Units (BBUs)
- Two 1U Storage Controllers (redundant storage processors)
 - Each Storage Controller includes:
 - Two redundant PSUs
 - Two 8GB/s Fiber Channel (FC) ports
 - Two 10Gbe iSCSI ports
 - Two 40GB/s Infiniband ports
 - One 1Gb/s management port

Front View



Rear View



9.6.2 XtremIO Cluster Configurations

Description	One X-Brick
Max eMLC SSDs per Cluster	25
X-Brick Storage Controllers	2
Processors Per Cluster	4 x 8-core 2.1 GHz Xeon
Total Memory Per Cluster	512 GB
Max 6Gbps SAS Back-End Buses per Cluster	4 x 4-Lane
Max 8Gbps FC Front-End Ports per Cluster	4
Max 10Gbps Ethernet iSCSI - Front-End Ports per Cluster	4
# Infiniband Switches per Cluster	0
X-Brick Cluster Rack Space ¹	6U
Battery Backup Unit (BBU) 2 x PS	2
Usable Capacity	7.58 TB
Max # defined initiators per cluster	512
Max # defined initiators per initiator groups	64
Max # defined volumes per cluster	8192
Max volume size per cluster	256TB**
Max # defined users per cluster	50

** Thin provisioning of more than 130TB can be approved via RPQ

9.6.3 XtremIO Disk Array Enclosure

Description	25 x 2.5" Drive DAE
Drive Types Supported ²	2.5" SSD / 400Gb eMLC Flash
Enclosure Height (Rack Units)	2U
Controller Interface	6 Gb SAS

9.6.4 XtremIO Power & Weight Specifications

Description	Power Consumption	Heat Dissipation	IEC C14 Connections	30A AC Circuits Single Phase	Weight
Single X-brick	750VA	2,559 BTU	4	2	562.5 lb (255.1kg)

9.6.5 XtremIO Connectivity


Connection per Storage Controller	Ports	IP Address	Speed
Management	1	1	1Gb
IPMI	1	1	1Gb
iSCSI	2	2	10GbE

Each X-brick has two storage controllers with each one having 2 x 10GbE LC iSCSI connections and a single 1 Gb Cat6 management connection. One iSCSI port per controller should be connected to a switch port on VLAN A and the other connected to a switch port on VLAN B for


¹ This includes a secondary BBU (1U) as part of the standard build for single X-bricks at Citi.

² A 20TB X-Brick with 800Gb eMLC SSDs is now available but is not certified at this time.

redundancy. The 1Gb MGMT1 connections on each controller require one IP address for management.



Note: Due to IPMI being identified as a security vulnerability the IP addresses for IPMI must be changed to dedicated internal access by EMC on initial install. This will require a change in cabling where the IPMI ports have a cross connection between storage controllers within the X-brick as in the example below.



10 Prerequisites for a new XtremIO Environment

The following items are required in order to implement an XtremIO storage array according to CSBE standards:

- A virtual XtremIO management server (XMS) which EMC provides as a virtual appliance for ESX. This server will run the GUI and XMCLI interface for XtremIO storage management application.
- Management IPs on the GRN for the XtremIO storage controllers and XMS management connections.
- Citi standard monitoring infrastructure environment.
- Citi standard SEMS environment.
- Appropriate physical floor space, cooling, power, and redundant storage IP fabric connectivity.
- 10GbE network connectivity is required between the XtremIO storage and host ports for iSCSI environments.

10.1 Required Licensing for Array

A new X-brick installation will include the following licenses

- XtremIO Right to Use

- XtremIO Management License

10.2 Required Installation Steps

- Physical installation of array into enterprise class data center
- Request cabling for XtremIO connectivity
- Assign relevant naming standards to XMS and storage controller IP addresses
- Request management IP's for storage controllers and XMS connectivity as well as IP's for 10GbE iSCSI connections.
- Request installation of XMS virtual appliance on production ESX cluster within same DC
- Coordinate with EMC to upgrade XtremIO cluster to latest Citi approved Service Pack and disable IPMI access.
- Configure XtremIO for e-mail home to EMC and Citi Storage Support
- Work with EMC to setup XMS for remote support via ESRS Gateway
- Enable IPtables on XMS and Storage Controller for limiting access.
- Configure and enable Citi enterprise monitoring
- Configure and enable Citi audit log forwarding.

11 XtremIO System Features

11.1 Thin Provisioning

XtremIO storage is natively thin provisioned using a small internal 8KB block size. All volumes in the system are thin provisioned, meaning that the system consumes capacity only when it is actually needed. XtremIO determines where to place the unique 8KB blocks physically inside the cluster after it calculates their fingerprint IDs. This means it never pre-allocates or thick provisions storage space before writing.

11.2 SAS Back-end

The XtremIO array utilizes 6 Gb/s Serial Attached SCSI (SAS) backend, with 4 lanes each. The SAS backend connection can address 4 storage devices at the same time on a port.

The XtremIO includes 2 SAS backend ports per storage controller. Currently the XtremIO utilizes only one back-end SAS port for connectivity to the DAE per storage controller. Enabling all backend ports is scheduled for a future release.

11.3 High Availability

The Storage Controllers and DAE in the cluster are equipped with dual power supplies. The cluster also has dual Battery Backup Units.

Any information not committed to the SSD is kept in multiple locations, called Journals. Each software module has its own journal, which is not kept on the same Storage Controller (SC) and can be used to restore data in case of unexpected failure. These journals are kept on the storage controllers with battery backed up power supplies. In case of a problem with the Battery Backup Unit (BBU), the journal fails over to another SC. In case of global failure, the BBU ensure that all Journals are written to vault drives in the SC and the cluster is then turned off.

11.4 XtremIO Data Protection

An X-brick is preconfigured as a single redundancy group, which eliminates the need to select, configure, and tune RAID groups. XDP uses double parity RAID6-like data protection. Dedicated hot spares are not required for rebuild; instead it leverages the “hot” space concept, where any free space available in the array can be utilized for failed drive reconstructions. The cluster always reserves sufficient distributed capacity for performing a single rebuild even at 100% physical capacity utilization. When an SSD fails the cluster rebuilds it to restore redundancy, this rebuild reconstructs data from the failed SSD onto the remaining SSDs in the protection group. Capacity utilization on a given protection group should never exceed 95%, it is recommended to reserve the capacity of at least one SSD in case of a dual SSD failure scenario. The X-brick can recover from multiple SSD failures that do not occur simultaneously. A dual drive failure that occurs concurrently is not yet supported.

11.6 Inline Data Reduction

The always on inline deduplication feature enables the removal of redundancies from data before it is written to the flash media. As data enters the cluster, data reduction is performed in real-time and not as a post-processing operation. Data blocks are stored according to their content and not according to their user level address within the volumes. If host data is written over and over again, it will be de-duplicated and no additional writes to the SSDs will occur.

12 Storage Provisioning Best Practices

12.1 Thin Provisioning

XtremIO provides array-based thin provisioning. Volumes are a collection of logical devices (LUNs) which are created for allocation to host initiator groups. The concept of managing pools or the need to configure RAID groups is not applicable.

Note: Thin provisioning of more than 130TB must be approved via EMC RPQ.

12.1.1 Volume (LUN) Creation

When creating a volume, the following inputs are required

- Volume Name – The volume naming convention below is recommended:
 - o Data Volume: *vm_esxclustname_dsXX(XX will be numbers...01, 02, 03, etc.)*
 - Size
 - Volume Type (Normal 512 LBs should be used)
1. Go to Menu Bar > Select Configuration > Under Volumes Tab > Click Add
Provide Volume Name, Size, Type, and leave Small IO, Unaligned IO, and VAAI TP Alerts Disabled > Next

Name	Size	Type	Small IO...	Unalign...	VAAI TP Alerts

0 volumes


Current Logical Storage: 3.58TB free

Current Physical Storage: 3.36TB free

< Back Next > Finish Cancel

With the VDI use case, alerting at the volume level is not recommended at this time due to increased number of warnings that will be received on the VMware virtual center console for each ESX host.

2. Volume folders can be created to logically group objects for simpler management > Create new folder or click Finish
 - Folder name - Providing a relevant name is recommended (Ex. VDI Volumes)




12.2 Initiator Groups


XtremIO uses the term “Initiators” to refer to ports that have access to volumes. These can be an iSCSI NIC or FC HBA. Initiators can be assigned to an Initiator Group (IG) for ease of management. The initiators within an IG share access to one or more of the X-brick volumes using LUN mapping. IQN’s for iSCSI initiators should be obtained from ESX server administrators.

12.2.1 Adding Initiators

1. Go to Menu Bar > Select Configuration > Under Initiator Groups Tab > Click Add




2. Provide Initiator Group Name > Add > Type relevant Initiator Name and Port Address (Provided by ESX SA) > OK



The 'Add Initiator' dialog box contains the following fields and controls:

- Initiator Name:** Text input field with 'HOST1_iSCSI' entered.
- Initiator Port Address:** Dropdown menu with 'iqn.1998-01.com.vmware:xx' selected.
- Discovery Section:**
 - User Name:** Text input field.
 - Password:** Password input field.
 - Confirm Password:** Password input field.
 - Clear Discovery Credentials:** Button.
- Authentication Section:**
 - User Name:** Text input field.
 - Password:** Password input field.
 - Confirm Password:** Password input field.
 - Clear Authentication Credentials:** Button.
- Buttons:** 'Refresh', 'OK', and 'Cancel' at the bottom.

3. Initiator Group folders can also be created to logically group objects for simpler management > Create New Folder with relevant name or Click Finish




The 'Add New Initiator Group' dialog box features the following elements:

- Title Bar:** 'Add New Initiator Group' with standard window controls.
- Section Header:** 'Move Initiator Group to Folder' with a folder icon.
- Instruction:** 'Please select a folder in which to create the new initiator group or use the new folder button to create the desired folder'.
- Select parent folder:** A list box containing 'CitiUser Servers' and 'GDSE Servers'. A 'New Folder' button is located to the right.
- Selected parent folder:** A text field showing '/'.
- Navigation Buttons:** '< Back', 'Next >', 'Finish', and 'Cancel' at the bottom.


12.3 LUN Mapping

In order to allow initiators in an initiator group to access a volume's disk space, you must map the volume to the initiator group.

1. Select Volume you want to map from Volumes tab and Initiators that needs to access this volume from Initiator Groups tab



2. Select Map All > Blue line appears between Volume and IG > Click Apply to enable




12.3.1 LUN Expansion and Shrinking

With the current XtremIO release, the ability to expand LUNs is available without having to unmap the volume. Shrinking a LUN is not supported at this time. LUN expansion is restricted at this time due to standard sized LUNs being utilized with the VDI use case.

12.3.2 LUN Sizing Guidelines

While it is possible for the LUN size to be as large as 256TB per X-brick, thin provisioning of more than 130TB is not recommended and must be approved via EMC RPQ. Host OS and application limits need to be taken into consideration when determining LUN size.

12.3.3 Monitoring XtremIO Capacity with XMS Dashboard




Term	Definition
Overall Efficiency	The disk space saved by the XtremIO array – Calculated as volume capacity / physical space used
Deduplication Ratio	Real-time inline data reduction ratio – Calculated as Volume Capacity used / Physical Space Used
Thin Provisioning Savings	% of volume capacity not in use – Free Volume Capacity / Allocated Volume Capacity
Volume Capacity	Total disk space defined by the volumes – Displays physical space used compared to logical space used
Physical Capacity	Total physical capacity available to all LUNs in cluster - Displays Used versus Free


12.3.4 Storage Capacity Alerting

The cluster is monitored using the Sys Ud Ssd Space parameter which represents Physical Capacity used on the array. This event activates when the physical capacity free reaches <10%. The system begins triggering a minor alert which must be confirmed manually before it clears. When the free space drops to 1% the alert changes to major and triggers another alert for forwarding. It is required to update the alert definitions for Sys Ud Ssd Space Limited so that the initial 90% alert will have a Major severity versus the default of Minor. Alerts are visible using the **Alerts & Events** option in the XMS Menu Bar under the Alerts or Events tab.


1. Go to Alerts & Events>Select small triangle icon on right side for Display Alert Definitions



2. Sort by Alert Name column>Double click Sys Ud Ssd Space Limited alert




3. Edit the alert definition by changing severity from Minor to Major and confirm Activity Mode: Enabled> Select OK




12.3.5 VAAI Thin Provision Limit

The additional ability to alert on total array capacity is available using the VAAI TP Limit. This setting provides a percentage to be set on the cluster configuration that will trigger once that percentage of physically used capacity is reached. Configuring this option can be done using the following steps.


1. Go to Administration > Cluster Limits



2. Confirm Don't Limit is unchecked > Enter 85% in the Limit to section > Apply



3. A minor alert similar to below will generate once the threshold is met.



Alerts will generate an SNMP trap to be sent to the Enterprise Monitoring System. Below is the information on the thresholds and their settings.

Threshold Type	Threshold Range	Default	Alert Severity	Alert Acknowledge
TP Soft Limit	Custom	100%/85% is recommended	Minor	Manual acknowledgment required
Built-in	90%	90%	Major(Manual update required)	Manual acknowledgment required
Built-in	99%	99%	Critical	Manual acknowledgment required

Allowing total allocation to exceed 90% of total capacity puts you at risk of running out of space and impacting all applications using LUNs on the XtremIO storage array.

12.3.6 Capacity Monitoring with SRM

With the current XtremIO release and SRM, multiple XtremIO systems can be discovered using the XtremIO SolutionPack. This feature allows for a variety of capacity and performance metrics to be viewed in one dashboard. At the time of this publishing SRM is going through certification at Citi and more details on capacity reporting will be available in a future update to this document.

12.3.7 LUN Migration

With the current XtremIO release, LUN migration is not available. If LUN migration is required, host based methods will need to be utilized.

13 Data Protection Best Practices

13.1 Local Protection

The XtremIO Storage Controllers are clustered so as to protect from the failure of a single storage controller (SC). In the case of an SC failure, the LUNs will continue to be available via the surviving SC. Each storage controller also has built in redundancy with dual power supplies, mirrored write cache as well as multiple paths connecting to the disk subsystem.

13.2 Remote Protection

At the time of writing, replication at the storage array level is not available for EMC XtremIO. With the current VDI use case, array based replication is not required.

System/data owners should ensure that any requirement for replication for COB purposes is addressed at the host or application level. For details on certified replication approaches please consult with the application and platform owners.

Further information is also available in the [Open Systems Data Replication Playbook](#)

13.3 Backup Standards

All data is accessed via the hosts to which LUNs are allocated. With the current VDI use case for XtremIO, backups are not taken of this environment.

14 Security Best Practices

14.1 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.

14.2 Audit Logging

Syslog forwarding for the EMC XtremIO environment should be integrated with Citi’s Secure Event Monitoring System (SEMS) for Audit Log Reconciliation. SEMS is used to centrally collect all security events. Custom reports can then be generated from the events collected by SEMS and sent to the log review teams for reconciliation. For details on configuring syslog forwarding go to the [Event Monitor](#) section.

14.3 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.

14.3.1 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 XMS/XIOS 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 XtremIO 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.

14.3.2 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.

14.4 CITMP Information Security Policy Convergence

14.4.1 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](#).

14.5 Role Based Access Controls (RBAC)

14.5.1 Roles Defined in XMS

The following roles are defined for users and groups in XMS.

XtremIO Role	Privileges
Administrator	Full privileges for storage and security tasks.
Configuration	Full privileges for storage tasks. No privileges for security tasks.
Read-Only	Read-only privileges for storage and monitor tasks. No privileges for security tasks.

14.5.2 Roles defined for use in Citi

With the current XtremIO release LDAPS/Active Directory integration is required for user authentication.

For everyday administration, Citi administrators should perform one of the following roles:

- Configuration
- Read-Only

This section breaks down the responsibilities and corresponding XMS role.

Citi Role	Description	XMS Role
Storage Engineer	The role of the administrator is to • Setup and make configuration changes to storage array • Create / manage user accounts (including role association)	Administrator
Storage Administrator	The role of the Storage Administrator is to • Make configuration changes to the storage system.	Configuration
Storage Monitor	The storage monitor will only be given read-only access to the storage system for troubleshooting and reporting purposes	Read-Only
System Administrator	A system account should be maintained in the unexpected event of being unable to manage the system using the above roles. The system account should be used to create the initial Administrator accounts and then should be immediately uploaded to a certified password storage system to ensure that these accounts are not used for system administration. Should there be a need to use the system account; the requester should use an auditable and traceable Change Management system to request the use of the account.	Administrator (admin acct)

14.5.3 Adding a Domain User

14.5.3.1 User Accounts


XMS should be setup to authenticate users with Citi's Active Directory (AD) for the following reasons:

- Central management of user accounts
- Enforces user account and password policies as defined in AD
- Assign XMS roles to AD groups instead of having to map roles to each individual user.
- Provides segregation of duties for Storage Operations to assign user management functions including entitlement reviews to ISA.

14.5.3.2 XMS LDAP User Authentication

XMS has the ability to integrate with AD to authenticate users. To setup XMS AD user authentication, the ISA will need to do the following:

1. XMS will need a functional ID and an associated password to be created in AD which it will use to bind to the directory. For further information on AD LDAP, please refer to the [CATE Microsoft Active Directory site](#).
2. AD Groups should be created for each of the XMS roles needed to manage the cluster: Read-Only, Configuration, and Administrator.
3. Once the groups have been created in AD, users should be assigned membership to the AD group with their corresponding role.
4. This is a one-time configuration setting. Once this setting is made, all user management tasks will be done on AD, not XMS. To configure AD user authentication in XMS, do the following:
 - a. Log into XMS using the default admin account.
 - b. Go to Menu Bar > Administration > Security Tab > LDAPs Configuration




- c. Select add from right side and enter the AD settings. All of this information should be provided by CATE Directory Engineering through the request process. Cache expiration should be lowered to 1 hour.


- d. Add server URLs,




- e. Add active directory group information



- f. Load CA Certificate – A Citigroup LDAP certificate should be pasted in PEM format in this field. This certificate can be found at the following URL
<https://digitalcertificate7.citigroup.com/cda-cgi/clientcgi?action=start>




The output displayed after this option is selected is the PEM Encoding. Cut from -
 ----BEGIN CERTIFICATE----- statement to the ----END CERTIFICATE-----
 statement with no extra spaces or lines then paste in the certificate field.



- g. Select Ok to enable settings

Once the AD settings have been configured and successful connection to the replicated SSO LDAP infrastructure has been established, users can login with their AD credentials. For cross-domain access an additional LDAP configuration must be entered following the same steps as documented above. An example login screen is provided below:



14.5.4 Functional IDs

With the current XtremIO release there are local accounts that provide specific privileges for EMC personnel to manage the device. These functional IDs must be managed by the storage operations teams and need to have passwords changed when appropriate. They should not be used by operational teams as a means of accessing the system, unless AD is inaccessible.

Once AD integration is configured, only the default Administrator account should be available for limited use in specific support cases.

User	Purpose	Role	Interactive	Privileged Access	Password Expiration
root	Administrative user created for Storage Controller & XMS (SSH access only)	Admin	Yes	Yes	Yes (PAR)
admin	XMCLI & GUI built-in administrative user for emergency access	Admin (cannot manage tech user account)	Yes	Yes	None (dual key control)
tech	XMCLI & GUI built-in administrative user for EMC support	Admin	Yes	Yes	None (dual key control)
xinstall	User for initial configuration/ installation on XMS used by EMC support	Admin	Yes	Yes	None
xmsupload	User created for uploading SW images to XMS for install and upgrade	User	Yes	Yes	None
xmsadmin	Built in administrative user for direct access to XMCLI used by EMC support	User	Yes	Yes	None

14.5.5 Root ID Access

With the current Service Pack release, root access will be required in certain troubleshooting scenarios on both XMS and the storage controllers. The following are some examples of cases where root access will be required:

1. Collecting traces on live cluster for problem investigation.
2. Troubleshooting cluster connectivity issues
3. Some FRU replacement procedures
4. Manually collect cluster logs in case dossier can't collect the logs directly or there is a need for additional debug logs which are not collected by dossier.
5. SSH from XMS to Storage Controllers
6. Managing ACL rules with IPTables

14.5.6 Access Control List (ACL) with IPTables

With EMC support requiring access to XMS for ESRS capability there must be an ACL implemented for limiting access from the XMS and storage controllers to any other devices on the Citi GRN. IPTables is the internal firewall software on XMS and the storage controllers that will create the ACL to limit SSH access. This is an important compensating control for restricting outbound connections from XtremIO cluster components.

Warning: EMC has recently enabled specific ACL rules on the XMS and storage controllers without any manual entries required. Any additional changes to the IPTables ACL must be reviewed and approved by CATE and EMC support before implementing.

14.5.6.1 Environment Setup

To validate the environment is setup after Hotfix is applied use the following command:

```
xmcli (admin)> show-clusters-info
Cluster-Name Index State Conn-State Activation-Time Start-Time SW-Version PSNT
b2-xtremio1 1 active connected Mon Dec 9 17:19:47 2013 Tue Aug 5 12:36:00 2014 2.4.1-12_hotfix_1.4
Encryption-Mode Encryption-Supported Mode-Switch-St
atus Mode-Switch-New-Mode SSH-Firewall-Mode
APM00134910382 disabled No none
disabled locked
```

Under SSH-Firewall-Mode you should see locked instead of unlocked.

14.5.6.2 Required entries in ACL

SSH is used to communicate with XMS and the storage controllers. Based on this requirement SSH should only be allowed as follows:

Source	Target	Purpose	Ports
Storage management host	XMS	SSH access to XMS for management	SSH - 22
ESRS gateway	XMS	SSH access to XMS for EMC troubleshooting	SSH - 22
XMS	Storage Controller	SSH access from XMS to storage controllers for troubleshooting. No SSH access to any other targets.	SSH - 22
Storage controllers	None	No SSH access to any other targets	SSH - 22

14.5.6.3 Sample IPTables Entries

The following is a sample of the IPTables ACL that is enabled when the Hotfix is applied to the XtremIO cluster. These rules limit SSH access on XMS and storage controllers.

XMS rules

```
[root@csbe-vxms1 ~]# iptables -L
```

```
Chain INPUT (policy DROP)
```

target	prot	opt	source	destination	
ACCEPT	all	--	anywhere	anywhere	
ACCEPT	all	--	anywhere	anywhere	state RELATED,ESTABLISHED
ACCEPT	tcp	--	anywhere	anywhere	tcp dpt:ssh
ACCEPT	udp	--	anywhere	anywhere	udp dpt:ntp
ACCEPT	tcp	--	anywhere	anywhere	tcp dpt:http
ACCEPT	tcp	--	anywhere	anywhere	tcp dpt:https
ACCEPT	udp	--	anywhere	anywhere	udp dpt:bootpc
ACCEPT	icmp	--	anywhere	anywhere	icmp echo-request

```
Chain FORWARD (policy ACCEPT)
```

target	prot	opt	source	destination
--------	------	-----	--------	-------------

```
Chain OUTPUT (policy ACCEPT)
```

target	prot	opt	source	destination	
ACCEPT	tcp	--	anywhere	anywhere	tcp dpt:ssh state RELATED,ESTABLISHED
ACCEPT	tcp	--	anywhere	xtremio-1.nam.nsroot.net	tcp dpt:ssh tcp
ACCEPT	tcp	--	anywhere	xtremio-3.nam.nsroot.net	tcp dpt:ssh tcp
ACCEPT	tcp	--	anywhere	xtremio-5.nam.nsroot.net	tcp dpt:ssh tcp
ACCEPT	tcp	--	anywhere	xtremio-7.nam.nsroot.net	tcp dpt:ssh tcp
DROP	tcp	--	anywhere	anywhere	tcp dpt:ssh state NEW

SC1 rules

```
[root@b2-xtremio1-x1-n1 ~]# iptables -L
```

```
Chain INPUT (policy DROP)
```

target	prot	opt	source	destination	
ACCEPT	all	--	anywhere	anywhere	
ACCEPT	all	--	anywhere	anywhere	
ACCEPT	tcp	--	anywhere	anywhere	tcp dpt:vce
ACCEPT	all	--	anywhere	anywhere	state RELATED,ESTABLISHED
ACCEPT	tcp	--	anywhere	anywhere	tcp dpt:ssh
ACCEPT	icmp	--	anywhere	anywhere	icmp echo-request
ACCEPT	udp	--	anywhere	anywhere	udp dpt:bootpc
iscsi	all	--	anywhere	anywhere	
iscsi	all	--	anywhere	anywhere	

```
Chain FORWARD (policy ACCEPT)
target    prot opt source      destination

Chain OUTPUT (policy ACCEPT)
target    prot opt source      destination
ACCEPT    tcp  -- anywhere   anywhere     state ESTABLISHED /* XIO_LOCK_OUT_CON */
DROP      tcp  -- anywhere   anywhere     state NEW /* XIO_LOCK_OUT_CON */

Chain iscsi (2 references)
target    prot opt source      destination
ACCEPT    tcp  -- anywhere   anywhere     tcp dpt:iscsi-target

SC2 rules
[root@b2-xtremio1-x1-n2 ~]# iptables -L
Chain INPUT (policy DROP)
target    prot opt source      destination
ACCEPT    all  -- anywhere   anywhere
ACCEPT    all  -- anywhere   anywhere
ACCEPT    tcp  -- anywhere   anywhere     tcp dpt:vce
ACCEPT    all  -- anywhere   anywhere     state RELATED,ESTABLISHED
ACCEPT    tcp  -- anywhere   anywhere     tcp dpt:ssh
ACCEPT    icmp -- anywhere   anywhere     icmp echo-request
ACCEPT    udp  -- anywhere   anywhere     udp dpt:bootpc
iscsi     all  -- anywhere   anywhere
iscsi     all  -- anywhere   anywhere

Chain FORWARD (policy ACCEPT)
target    prot opt source      destination

Chain OUTPUT (policy ACCEPT)
target    prot opt source      destination
ACCEPT    tcp  -- anywhere   anywhere     state ESTABLISHED /* XIO_LOCK_OUT_CON */
DROP      tcp  -- anywhere   anywhere     state NEW /* XIO_LOCK_OUT_CON */

Chain iscsi (2 references)
target    prot opt source      destination
ACCEPT    tcp  -- anywhere   anywhere     tcp dpt:iscsi-target
```

Important: After IPTables configuration changes are applied on XMS and storage controllers it is critical to validate access to and from each device.

15 Alerting and Monitoring

15.1 SNMP and Email Alerting

The XtremIO SNMP alerts should be configured to send to an Enterprise Monitoring trap host. The received traps are then filtered and (as appropriate) will create a problem ticket or generate an SMS / Paging alert.

15.1.1 Adding an Array to the Monitoring System

Tivoli monitoring and alerting service is provided by Citi's ESM team. The ESM team can be contacted at ***ESM Distributed Event Delivery** for instructions on requesting the systems to be monitored. Once the request has been submitted, the ESM team should provide the Tivoli server that will be used to send the XtremIO SNMP traps to.

It should also be requested to the monitoring team to have the X-bricks setup in the ping scanning list i.e. if a ping of the array does not respond then this should be configured to generate an alert.

15.1.2 SNMP alerting with Tivoli NetCool

When there is an alert event, the array will send an SNMP alert to the trap host configured in the system. In order for the trap host to create a problem ticket or send out SMS/ Paging alert, it is important to check that a valid entry exists for each of the storage controllers in [Aperture](#).

Notes


- Resiliency at the Tivoli NetCool (trap host) level is not available at this time. i.e. if the configured SNMP trap host server is down, all SNMP alerts sent by the storage system will not be received by the Tivoli NetCool environment³

15.1.3 Adding an SNMP Recipient and SNMP Properties


The following steps should be taken to configure SNMP alerting on the XtremIO:

1. Add recipient and properties for SNMP alerting.
 - a. From the menu bar select Administration icon, then in the left pane select SNMP Configuration.
 - b. Check the Send SNMP Notification box
 - c. Select Add to open Recipient Dialog Box, type in the recipient IP address then click OK to confirm and close dialog box.
 - d. Select the SNMP version (V1 or V2C) V1 is recommended within Citi.
 - e. Type in Community field (Default is public read only)
 - f. Type in Trap port field (Default is 162)


³ When configuring a new array, it is suggested to check with ESM team if a solution with failover/load-balancing is available for the trap host (not available at time of documentation)



- g. Click Apply to Save Changes
- h. From menu bar, select Alerts & Events>Events Tab> Display Event Handlers icon on upper right corner.



- i. Select Add



- j. Select the Category, Severity, and Entity from the drop down then select Send e-mail and Send SNMP trap select OK. All severity levels should be selected in order to receive Major and Critical events.

Note: There is currently not a Send Test Event option within XMS, enabling all properties and then creating a test user account can trigger a test alert to be sent to Tivoli NetCool console.

15.1.4 SNMP MIB

The below table describes the SNMP trap variables that are generated by the XtremIO and examples of relevant information from a triggered event.

OID/ Trap#	Trap Variable Name	Sample Alert Data
1	xioEventType	Create
2	xioEventThreshold	alert_def_sys_stopping
3	xioEventObject	X2-SC1 [3]
4	xioEventSeverity	Major
5	xioEventState	Outstanding
6	xioEventDesc	Raised alert of type alert_def_disconnected_from_node_mgr: object: X2-SC1 [3] severity: major
7	xioEventXmsTime	Fri Feb 21 20:01:48 2014

XtremIO Severity	Netcool Severity	ESM Action
Information	1	No Ticket, Event will be cleared in 24 hrs in Netcool
Minor	3	No Ticket, possible e-mail alert if configured in SMRDB
Major	5	Ticket or email or both based on user config in SMRDB
Critical	5	Ticket or email or both based on user config in SMRDB

The various events are being reviewed within Citi to determine where there should be changes to the default severity. This is then used by the ESM team to create a rule file that maps an eventID to the Citi severity level. If a trap is received that is not defined in the rule file, the severity of the trap will be set to informational and will not be processed by any automation within Netcool.

In order for the ESM team to complete the SNMP alerting setup, they also require the MIB file for the array. The MIB file is available on the [CSBE XtremIO Product page](#).

15.2 Call Home and Email Alerting

Email home should be configured on the storage array for remote alerting. The template should be setup to also send an email to Citi's support team responsible for monitoring the XtremIO environment each time the device emails home to EMC. Citi's SMTP server should be specified to send out the email home alerts under the SMTP Information.

Notes

- It is only possible to configure a single SMTP Server IP Address (DNS names are not supported). When selecting a mail relay host, ensure that a Virtual IP (VIP) address is used for the SMTP Server to provide resiliency.
- Do not use an SMTP server that is using storage on the array you are configuring.

15.3 Remote Support (EMC ESRS)

ESRS supports access to the XMS both outbound and inbound. Since the XMS is a virtual appliance on the Citi GRN the access to as well as from XMS and the storage controllers must be restricted. In order to facilitate this restriction [ACL with IPTables](#) must be used to enable access controls on XMS and storage controllers. Once alerts are received by EMC support from XMS they will have the ability to initiate access remotely to XMS for proper troubleshooting..

Please reference the [ESRS Solutions Guide](#) for more details.


15.4 Dossier Collection


In many troubleshooting scenarios the need to collect a dossier will be required. The dossier is a collection of logs from the XMS and storage controllers that EMC support will review. The way to obtain a dossier is via the CLI with the **create-debug-info** command. An example of using this command is below:

```
xmcli (jh53740)> create-debug-info
The process may take a while. Please do not interrupt.
.....
.....
.....
Successfully collected dossier on X2-SC2
Successfully collected dossier on X1-SC2
Successfully collected dossier on X1-SC1
Successfully collected dossier on X2-SC1
Successfully collected dossier on xms
Created /var/www/xtremapp/DebugInfo/b2-xtremio1_2014_08_07_1025.tar.gz

Debug      info      collected      and      may      be      accessed      via
https://10.40.43.62/xtremapp/DebugInfo/b2-xtremio1_2014_08_07_1025.tar.gz
```

Once the collection is complete you can copy the URL into a browser and download the tar file.







15.5 Event Monitor

For Audit Log Reconciliation, it is required that central event monitoring be configured for all arrays. Within XMS the ability to forward syslog events to a remote server can be configured as follows. .


1. Go to Menu Bar > Select Administration > Notifications tab > Syslog Configuration



2. Select Add > Type in hostname of Syslog server and click OK




3. Select check box for Send Syslog notifications then click Apply



15.6 Performance Monitoring using XMS

XMS gathers real time performance statistics at the Block Size, Initiator Groups, and Volumes level in the Dashboard view. Bandwidth, IOPS, and Latency metrics for Reads and Writes are updated on a 5 second interval in the Performance pane. The ability to select statistics from the following time frames is available: Last Hour, Last 6 Hours, Last 24 Hours, Last 3 days, Last Week.



15.6.1 Exporting Performance History

In the current release, performance history is only kept for 7 days. The performance history can be exported via the CLI with the **export-performance-history** command. An example of the using this export is below.


```
xmcli (admin)> export-performance-history
Exported performance history can be accessed via
http://x.x.x.x/xtremapp/performance_history.csv
xmcli (admin)>
```

Once the export is complete you can copy the URL into a browser and download the CSV file.


15.6.2 Advanced Real Time Monitors

In order to view specific components in real time there is the ability to create a custom monitor for a particular monitor type.


1. Go to Menu Bar > Select Monitor > New to add new monitor




2. Type in a Monitor name and select a monitor type from the list



3. Select the elements to include and properties to monitor then Next



4. Review monitor parameters and hit Finish




5. The monitor will be available in the main viewing pane on the right.

Monitor				
Volumes Monitor1				
Name	Read-BW(MB/s)	Write-BW(MB/s)	U-Read-BW(MB/s)	S-Read-
VDI datastore25	0	82	0	0
VDI datastore24	0	1,227	0	0
VDI datastore23	0	165	0	0
VDI datastore22	460	1,357	0	0
VDI datastore21	0	459	0	0
VDI datastore20	0	377	0	0
VDI datastore19	0	1,843	0	0
VDI datastore18	0	47	0	0
VDI datastore17	14	239	10	10
VDI datastore16	0	374	0	0
VDI datastore15	0	53	0	0
VDI datastore14	0	234	0	0
VDI datastore13	32	109	0	0
VDI datastore12	8	683	4	4
VDI datastore11	0	302	0	0
VDI datastore10	0	93	0	0
VDI datastore9	0	161	0	0
VDI datastore8	0	2,078	0	0
VDI datastore7	0	1,923	0	0
VDI datastore6	0	159	0	0
VDI datastore5	0	119	0	0
VDI datastore4	0	72	0	0
VDI datastore3	0	34	0	0
VDI datastore2	0	1,124	0	0
VDI datastore1	0	501	0	0
Infra2	0	0	0	0
Infra1	0	0	0	0

Total: 27 Volumes. Last updated: Feb 26, 2014 12:38:15 PM

16 Host Connectivity and Configuration

Currently, CSBE has only certified host connectivity to the EMC XtremIO storage arrays via iSCSI connected to top of rack IP fabric switches. Below is a diagram of the logical iSCSI layout.



16.1 IP Fabric and iSCSI Guidelines


The following guidelines should be followed when connecting a host to the XtremIO array.

- The XtremIO storage system and hosts should be connected to two iSCSI VLANs for redundancy. These will be Layer 2 non-routed VLANs.
- iSCSI connections will be 10GbE SFP from host and storage. Jumbo frames should be enabled on the VLAN as well as the iSCSI array ports. The recommended MTU setting is 8800 across compute NICs, network switch, and iSCSI storage ports.
- Each iSCSI LUN will be presented out of each XtremIO iSCSI portal resulting in each LUN being seen 4 times by the ESX servers, twice per iSCSI VLAN.
- To easily differentiate IP connections it is recommended to configure
 - o Odd IP addresses on controller 1
 - o Even IP addresses on controller 2
- ESX native multi-pathing must be configured on the host side for proper round robin access to the target iSCSI ports on the X-bricks. For further details see [VMware VDI 2.0 Storage Design Document](#).


16.2 iSCSI Portal Tables

As part of the implementation of the XtremIO, the iSCSI ports on the storage controllers must be configured for communication with the host iSCSI initiators. This requires setting up iSCSI portal tables via XMS.

1. Go to Menu Bar > Select Administration > iSCSI Network Configuration



2. Click Add > Select target port > Enter Non-routed VLAN information/Subnet bits. Port defaults to 3260 and VLAN to 0.




Note: For a given storage controller, one iSCSI port must be configured for VLAN A and the second iSCSI port for VLAN B. VLAN tagging on the XtremIO is not currently supported therefore verifying the port configuration and cabling is recommended.

16.3 iSCSI Ports Configuration


The XtremIO iSCSI ports on the storage controllers should be configured with Jumbo Frames enabled. The same MTU of 8800 must be set on compute NICs, iSCSI storage ports, and network switch.

1. Go to Menu Bar > Select Administration > iSCSI Ports Configuration>



1. Go to Menu Bar > Select Administration > iSCSI Ports Configuration>
2. Select Enable Jumbo Frames check box and change MTU to 8800 > Apply

Note: The warning message below will appear and host connectivity could be impacted. It is recommended to do this during initial installation and before device is in production or during green zone.



17 Local Protection

17.1 Snapshots

Snapshots are now supported in the current release. This functionality is restricted at this time and will be released in a future Solutions Guide.

18 XMS Management Server

The XMS will be installed from an EMC provided OVA package for each X-brick cluster. It should be deployed on a production vSphere cluster within the same data center that is not running on XtremIO storage. The specifications for XMS are as follows:

Operating System	Memory	CPU	Disk
Linux CentOS 6.2	8 GB vRAM	2 vCPUs	900 GB (Thin Provisioned)

It is recommended to manage and monitor the XtremIO storage array through the XtremIO storage management application installed on a Windows management server or desktop. XMS is preinstalled with the CLI and GUI. It must have network access to management and IPMI addresses on all the X-brick storage controllers and be located within the same data center.

The XMS GUI is initially installed via internet browser (e.g. Microsoft Internet Explorer) and a subsequent Java client will be the primary method for access. This client software communicates with the XMS using standard TCP/IP protocols and can be used in any location that allows the client to access the XMS. The full java client must be installed on the Windows management server or desktop as Java Web Start is required.

18.1 XMS Installation Procedure


Installation Steps

1. Open up a browser and type in XMS IP address:



Note: Adding this URL to restricted sites may be required depending on your security settings

2. Select launch and check the accept risk box to run application



3. Login screen will appear and allow you to enter the credentials



4. Once login is verified the dashboard will appear



Note: A shortcut on the desktop will be created for future access

19 Appendix A: Storage Controller Failover

XtremIO Health Check & SC Failover Simulation

1. Upload the "Health Check" script to XMS

"Health Check" script should be uploaded to /images/scripts as user xmsupload/xmsupload

Requires "tech" account to execute

Running health Check script on an XIO with IPMI set to Dedicated.

```
system_health-v1.9-s3.0.0.py.gpg
[root@lab12vdi-extremio01 scripts]# xmcli -u tech
Password:
Connect XMS on 10.40.43.151:443: version 3.0.0 build 42
xmcli (tech)>
xmcli (tech)> run-script script="system_health-v1.9-s3.0.0.py"
Script exited with status: 0
Script output:
Assuming xms is local lab12vdi-extremio01
Verifying logs for critical errors: OK
DIMM alerts: OK
guid2lid: OK
IB cables: OK
IB errors: OK
IB fencing consistency: OK
Initiator: 12c7-b1460g8-01 doesn't have enough connections: 0 (at least 2 needed)
Initiators connectivity: FAIL
IPMI active port: OK
Fan and power redundancy: SKIPPED: run in full mode only
LCC connectivity: OK
Local disk: OK
System Modules health: OK
Nodes physical memory: OK
PSU health: OK
The current ipmi configuration is: dedicated, however the ipmi state is: unknown (expected "disconnected")
IPMI conn state isn't as expected
The current ipmi configuration is: dedicated, however the ipmi state is: unknown (expected "disconnected")
IPMI conn state isn't as expected
Storage controllers health: FAIL
SSD health: OK
IB errors: OK
Balanced targets performance: OK
Fiber channel firmware status verification: OK
UPS health: OK
Not enough free space on xms: 2450748
XMS free space: FAIL
Xenw health: OK
XMS physical memory: OK
XMS recovery for live systems: OK
Errors found while executing the system health scripts, please review the relevant items and re-run.
xmcli (tech)>
```

→ this failure is code bug in the script, will be fix in the next release

↑

needs to space cleanup for the old dossier and old code

Note: XMS free space: 3GB space is required

2. Before Testing: Disable SYR Notification to stop alerting:

#modify-syr-notifier disable

3. Run the health check script

xmcli (tech)> run-script script="system_health-v1.9-s3.0.0.py"

IPMI is set to management:

Initiator Fail error is a code bug in the health check script will be fixed in the next release.

```

Storage-Controller-Name Index Active-IPMI-Port
X1-SC1 1 management
X1-SC2 2 management
xmcli (tech)> show-storage-controllers
Storage-Controller-Name Index Mgr-Addr IB-Addr-1 IB-Addr-2 IPMI-Addr Brick-Name Index Cluster-Name Index State Health-State
Conn-State IPMI-State Powered-State
X1-SC1 1 10.40.87.45 169.254.0.1 169.254.0.2 10.40.87.47 X1 1 lab12vdi-xtremio01 1 healthy healthy
connected connected on
X1-SC2 2 10.40.87.46 169.254.0.17 169.254.0.18 10.40.87.48 X1 1 lab12vdi-xtremio01 1 healthy healthy
connected connected on
xmcli (tech)> run-script script="system_health-v1.9-s3.0.0.py"
Script exited with status: 0
Script output:
Assuming xms is local lab12vdi-xtremio01
Verifying logs for critical errors: OK
DIMM alerts: OK
guid2lid: OK
IB cables: OK
IB errors: OK
IB fencing consistency: OK
Initiator: 12c7-b1460g8-01 doesn't have enough connections: 0 (at least 2 needed)
Initiators connectivity: FAIL
IPMI active port: OK
Fan and power redundancy: SKIPPED: run in full mode only
LCC connectivity: OK
Local disk: OK
System Modules health: OK
Nodes physical memory: OK
PSU health: OK
Storage controllers health: OK
SSD health: OK
IB errors: OK
Balanced targets performance: OK
Fiber channel firmware status verification: OK
UPS health: OK 3Gb space needed
XMS free space: OK
Xenv health: OK
XMS physical memory: OK
XMS recovery for live systems: OK
Errors found while executing the system health scripts, please review the relevant items and re-run.
xmcli (tech)>

```

4. Check ESX path

Note: All 4 iSCSI ports should be UP before starting the failover. The ESX hosts should have 4 Active paths before and after test.

2 active paths with a single SC after deactivated.

Check ESX paths before and after each deactivate/activate action as well as check for pending alerts.

Verify no Alerts in the array by running: # show-alerts

Command verification before De-Activating Storage Controllers

xmcli (tech)> show-alerts

xmcli (tech)> show-initiators-performance frequency=5

Initiator-Name Index Write-BW(MB/s) Write-IOPS Read-BW(MB/s) Read-IOPS BW(MB/s) IOPS Total-Write-IOs Total-Read-IOs

```

Initiator-Name Index Write-BW(MB/s) Write-IOPS Read-BW(MB/s) Read-IOPS BW(MB/s) IOPS Total-Write-IOPS Total-Read-IOPS
12c7-b1460g8-01 1 0.000 0 0.000 0 0.000 0 0 4760
12c7-b1460g8-04 2 0.000 0 0.000 0 0.000 0 6120 11064
xmcli (tech)> show-initiators
Initiator-Name Index Port-Type Port-Address IG-Name Index Parent-Folder-Name Index Certainty-State Chap-Authentication-Initiator-User-Name
Chap-Discovery-Initiator-User-Name Chap-Authentication-Cluster-User-Name Chap-Discovery-Cluster-User-Name
12c7-b1460g8-01 1 iscsi iqn.1998-01.com.vmware:12c7-b1460g8-01-78948f2f VDI ESXservers 1 / 1 ok
12c7-b1460g8-04 2 iscsi iqn.1998-01.com.vmware:12c7-b1460g8-04-2043d62f VDI ESXservers 1 / 1 ok

xmcli (tech)> show-initiators-connectivity
Existing Initiator List:
Name Index Port-Type Port-Address Num-Of-Conn-Targets
12c7-b1460g8-01 1 iscsi iqn.1998-01.com.vmware:12c7-b1460g8-01-78948f2f 2
12c7-b1460g8-04 2 iscsi iqn.1998-01.com.vmware:12c7-b1460g8-04-2043d62f 2

xmcli (tech)>
xmcli (tech)> show-target
Unknown command: 'show-target'. Try 'help'
xmcli (tech)> show-targets
Name Index Port-Type Port-Address Port-State Health-Level Storage-Controller-Name Index TG-Name Index MTU Jumbo-Frames Certainty-State
X1-SC1-fc1 1 fc 21:00:00:24:ff:58:0a:48 down level_1_clear X1-SC1 1 Default 1 1500 False ok
X1-SC1-fc2 2 fc 21:00:00:24:ff:58:0a:49 down level_1_clear X1-SC1 1 Default 1 1500 False ok
X1-SC1-iscsi1 5 iscsi iqn.2008-05.com.xtremio:001e679f591e down level_1_clear X1-SC1 1 Default 1 8800 True ok
X1-SC1-iscsi2 6 iscsi iqn.2008-05.com.xtremio:001e679f591f up level_1_clear X1-SC1 1 Default 1 8800 True ok
X1-SC2-fc1 11 fc 21:00:00:24:ff:58:06:3c down level_1_clear X1-SC2 2 Default 1 1500 False ok
X1-SC2-fc2 12 fc 21:00:00:24:ff:58:06:3d down level_1_clear X1-SC2 2 Default 1 1500 False ok
X1-SC2-iscsi1 15 iscsi iqn.2008-05.com.xtremio:001e679f15aa down level_1_clear X1-SC2 2 Default 1 8800 True ok
X1-SC2-iscsi2 16 iscsi iqn.2008-05.com.xtremio:001e679f15ab up level_1_clear X1-SC2 2 Default 1 8800 True ok
xmcli (tech)> run-script script="system_health-v1.9-s3.0.0.py"

```

5. De-activating Storage Controller:

```

xmcli (tech)> show-alerts
Index Description Severity Raise-Time Entity Name Index Alert-Type State Alert-Code
xmcli (tech)> show-storage-controllers
Storage-Controller-Name Index Mgr-Addr IB-Addr-1 IB-Addr-2 IPMI-Addr Brick-Name Index Cluster-Name Index State Health-State Enabled-State Unorderly-Stop-Reason
Conn-State IPMI-State Powered-State
X1-SC1 1 10.40.87.45 169.254.0.1 169.254.0.2 10.40.87.47 X1 1 lab12vdi-xtremio01 1 healthy healthy enabled none
connected connected on
X1-SC2 2 10.40.87.46 169.254.0.17 169.254.0.18 10.40.87.48 X1 1 lab12vdi-xtremio01 1 healthy healthy enabled none
connected connected on
xmcli (tech)> deactivate-storage-controller sc-id=2
Are you sure you want to deactivate Storage Controller X1-SC2 [2]? (Yes/No): Yes
Storage Controller X1-SC2 [2] deactivation initiated
xmcli (tech)> show-storage-controllers
Storage-Controller-Name Index Mgr-Addr IB-Addr-1 IB-Addr-2 IPMI-Addr Brick-Name Index Cluster-Name Index State Health-State Enabled-State Unorderly-Stop-Reason
Conn-State IPMI-State Powered-State
X1-SC1 1 10.40.87.45 169.254.0.1 169.254.0.2 10.40.87.47 X1 1 lab12vdi-xtremio01 1 healthy healthy enabled none
connected connected on
X1-SC2 2 10.40.87.46 169.254.0.17 169.254.0.18 10.40.87.48 X1 1 lab12vdi-xtremio01 1 disconnected healthy user_disabled none

```

Verify XENVS are stabilized meta data load before Activating Storage Controllers

Alerts Events					
Search Alerts: ☐ Hide acknowledged alerts					
Severity	Code	Date and Time	Entity	Entity Details	Description
Major	1500202	Oct 16, 2014 11:47:13 AM	Battery Bac...	X1-BBU	BBU is disconnected. No communication with the BBU is available.
Major	0404702	Oct 16, 2014 11:47:13 AM	Storage Co...	X1-SC1	The cluster has detected a failure of the journal.
Major	1501103	Oct 16, 2014 11:47:07 AM	Battery Bac...	X1-BBU	Battery Backup Unit bypass mode is activated (there is no protection against power dips or outages). Invalid error.
Major	1500803	Oct 16, 2014 11:47:07 AM	Battery Bac...	X1-BBU	Battery Backup Unit is overloaded. Invalid value.
Major	1501303	Oct 16, 2014 11:47:07 AM	Battery Bac...	X1-BBU	The cluster detected that the Battery Backup Unit needs to be replaced. Invalid error.
Major	0402602	Oct 16, 2014 11:47:07 AM	Storage Co...	X1-SC1	The Storage Controller was disabled by user request.
Major	1501504	Oct 16, 2014 11:47:06 AM	Battery Bac...	X1-BBU	BBU is disconnected from the cluster.
Major	1600504	Oct 16, 2014 11:47:06 AM	DAE Contro...	X1-DAE-LCC-A	The bottom DAE LCC is disconnected.

```

xmc11 (tech1)> show-alerts
Index Description
Raise-Time Entity Name Index Alert-Type State Alert-Code Severity
17 The Storage Controller is disconnected from the XMS.
Thu Oct 16 11:47:36 2014 Node XI-SC1 1 disconnected_from_node_mgr outstanding 0400503 major
16 Storage Controller InfiniBand port 2 is down.
Thu Oct 16 11:47:16 2014 Node XI-SC2 2 node_ib2_port_down outstanding 0402302 major
15 Storage Controller InfiniBand port 1 is down.
Thu Oct 16 11:47:16 2014 Node XI-SC2 2 node_ib1_port_down outstanding 0401902 major
14 BBU is disconnected. No communication with the BBU is available.
Thu Oct 16 11:47:13 2014 UPS XI-BBU 1 ups_node_disconnected outstanding 1500202 major
13 The cluster has detected a failure of the journal.
Thu Oct 16 11:47:13 2014 Node XI-SC1 1 journal_failed outstanding 0404702 major
12 Battery Backup Unit bypass mode is activated (there is no protection against power dips or outages). Invalid error.
Thu Oct 16 11:47:07 2014 UPS XI-BBU 1 ups_bypass_active_failure_unexpected_value outstanding 1501103 major
11 Battery Backup Unit is overloaded. Invalid value.
Thu Oct 16 11:47:07 2014 UPS XI-BBU 1 ups_overrid_failure_unexpected_value outstanding 1500803 major
10 The cluster detected that the Battery Backup Unit needs to be replaced. Invalid error.
Thu Oct 16 11:47:07 2014 UPS XI-BBU 1 ups_failure_unexpected_value outstanding 1501303 major
9 The cluster has detected a potential risk for the journal health or persistency. Journaling was disabled by the cluster to keep data consistency and persistency.
Thu Oct 16 11:47:07 2014 Node XI-SC1 1 journal_fault outstanding 0400703 major
2 The Storage Controller was disabled by user request.
Thu Oct 16 11:47:07 2014 Node XI-SC1 1 node_user_disabled outstanding 0402602 major
8 Storage Controller has stopped (immediate).
Thu Oct 16 11:47:07 2014 Node XI-SC1 1 node_unorderly_stop outstanding 0402503 major
7 Storage Controller is disconnected from the cluster.
Thu Oct 16 11:47:06 2014 Node XI-SC1 1 node_fru_disconnected outstanding 0402704 major
6 BBU is disconnected from the cluster.
Thu Oct 16 11:47:06 2014 UPS XI-BBU 1 ups_fru_disconnected outstanding 1501504 major
5 The bottom DAE LCC is disconnected.
Thu Oct 16 11:47:06 2014 JBODController XI-DAE-LCC-A 1 jbodcontroller_fru_disconnected outstanding 1600504 major
xmc11 (tech1)> show-xenvs
XEnv-Name Index CPU(%) CSID State Storage-Controller-Name Index Brick-Name Index
XI-SC1-E1 1 0 10 inactive XI-SC1 1 XI 1
XI-SC1-E2 2 0 11 inactive XI-SC1 1 XI 1
XI-SC2-E1 3 0 12 active XI-SC2 2 XI 1
XI-SC2-E2 4 0 13 active XI-SC2 2 XI 1
xmc11 (tech1)>

```

6. Activating Storage Controllers after verifying Host path and Cluster/storage controllers had failover completely and meta data is stabilized.

```

xmcli (tech)> activate-storage-controller sc-id=1
Storage Controller X1-SC1 [1] activation initiated
xmcli (tech)>
xmcli (tech)>

xmcli (tech)> show-storage-controllers
Storage-Controller-Name Index Mgr-Addr IB-Addr-1 IB-Addr-2 IPMI-Addr Brick-Name Index Cluster-Name Index State Health-State Enabled-State Unorderly-Stop-Reason
Comm-State IPMI-State Powered-State
X1-SC1 1 10.40.87.45 169.254.0.1 169.254.0.2 10.40.87.47 X1 1 lbl2vdi-xtremio01 1 healthy degraded enabled none
connected connected on
X1-SC2 2 10.40.87.46 169.254.0.17 169.254.0.18 10.40.87.48 X1 1 lbl2vdi-xtremio01 1 healthy healthy enabled none
connected connected on
xmcli (tech)> show-xenvs frequency=10
XEnv-Name Index CPU(%) CSID State Storage-Controller-Name Index Brick-Name Index
X1-SC1-E1 1 0 10 active X1-SC1 1 X1 1
X1-SC1-E2 2 0 11 inactive X1-SC1 1 X1 1
X1-SC2-E1 3 0 12 active X1-SC2 2 X1 1
X1-SC2-E2 4 0 13 active X1-SC2 2 X1 1
XEnv-Name Index CPU(%) CSID State Storage-Controller-Name Index Brick-Name Index
X1-SC1-E1 1 10 10 active X1-SC1 1 X1 1
X1-SC1-E2 2 0 11 inactive X1-SC1 1 X1 1
X1-SC2-E1 3 6 12 active X1-SC2 2 X1 1
X1-SC2-E2 4 0 13 active X1-SC2 2 X1 1
XEnv-Name Index CPU(%) CSID State Storage-Controller-Name Index Brick-Name Index
X1-SC1-E1 1 99 10 active X1-SC1 1 X1 1
X1-SC1-E2 2 0 11 inactive X1-SC1 1 X1 1
X1-SC2-E1 3 99 12 active X1-SC2 2 X1 1
X1-SC2-E2 4 99 13 active X1-SC2 2 X1 1

```

Verification Command to check connectivity

```

xmcli (tech)> show-initiators-connectivity
Existing Initiator List:
Name Index Port-Type Port-Address Num-Of-Conn-Targets
12c7-b1460g8-01 1 iscsi iqn.1998-01.com.vmware:12c7-b1460g8-01-78948f2f 1

xmcli (tech)> show-initiators
Initiator-Name Index Port-Type Port-Address IG-Name Index Parent-Folder-Name Index Certainty-State Chap-Authentication-Initiator-User-Name
Chap-Discovery-Initiator-User-Name Chap-Authentication-Cluster-User-Name Chap-Discovery-Cluster-User-Name
12c7-b1460g8-01 1 iscsi iqn.1998-01.com.vmware:12c7-b1460g8-01-78948f2f VDI ESXservers 1 / 1 ok

xmcli (tech)> show-initiator-group
Usage: show-initiator-group property=value list
PROPERTY MANDATORY DESCRIPTION VALUE
-----
duration No Monitor Duration seconds
frequency No Monitor Intervals seconds
ig-id Yes IG ID name or index

** Error: CLI Command Syntax Error: Missing some of mandatory properties list: ['ig-id']
xmcli (tech)> show-initiator-groups
IG-Name Index Parent-Folder-Name Index Certainty-State
VDI ESXservers 1 / 1 ok
xmcli (tech)> show-initiator-groups-performance
IG-Name Index Write-BW(MB/s) Write-IOPS Read-BW(MB/s) Read-IOPS BW(MB/s) IOPS Total-Write-IOPS Total-Read-IOPS
VDI ESXservers 1 0.000 0 0.000 0 0.000 0 0 2930
xmcli (tech)> show-targets
Name Index Port-Type Port-Address Port-State Health-Level Storage-Controller-Name Index TG-Name Index MTU Jumbo-Frames Certainty-State
X1-SC1-fc1 1 fc 21:00:00:24:ff:58:0a:48 down level_1_clear X1-SC1 1 Default 1 1500 False ok
X1-SC1-fc2 2 fc 21:00:00:24:ff:58:0a:49 down level_1_clear X1-SC1 1 Default 1 1500 False ok
X1-SC1-iscsi1 5 iscsi iqn.2008-05.com.xtremio:001e679f591e down level_1_clear X1-SC1 1 Default 1 8800 True ok
X1-SC1-iscsi2 6 iscsi iqn.2008-05.com.xtremio:001e679f591f up level_1_clear X1-SC1 1 Default 1 1500 False ok
X1-SC2-fc1 11 fc 21:00:00:24:ff:58:06:3c down level_1_clear X1-SC2 2 Default 1 1500 False ok
X1-SC2-fc2 12 fc 21:00:00:24:ff:58:06:3d down level_1_clear X1-SC2 2 Default 1 1500 False ok
X1-SC2-iscsi1 15 iscsi iqn.2008-05.com.xtremio:001e679f15aa down level_1_clear X1-SC2 2 Default 1 8800 True ok
X1-SC2-iscsi2 16 iscsi iqn.2008-05.com.xtremio:001e679f15ab up level_1_clear X1-SC2 2 Default 1 8800 True ok
xmcli (tech)> show-targets

```

7. Verify Storage Controllers are healthy

```

xmcli (tech)> show-storage-controllers
Storage-Controller-Name Index Mgr-Addr IB-Addr-1 IB-Addr-2 IPMI-Addr Brick-Name Index Cluster-Name Index State Health-State Enabled-State Unorderly-Stop-Reason
Comm-State IPMI-State Powered-State
X1-SC1 1 10.40.87.45 169.254.0.1 169.254.0.2 10.40.87.47 X1 1 lbl2vdi-xtremio01 1 healthy healthy enabled none
connected connected on
X1-SC2 2 10.40.87.46 169.254.0.17 169.254.0.18 10.40.87.48 X1 1 lbl2vdi-xtremio01 1 healthy healthy enabled none

```

Run the health check script (as tech user)

```
xmcli (tech)> run-script script="system_health-v1.9-s3.0.0.py"
```

Check XtremIO alerts

```
xmcli (tech)> show-alerts
```

8. Check ESX paths

9. Run through steps 5-8 for second controller failover simulation

10. Enable SYR Notification after the testing

```
xmcli (admin)> modify-syr-notifier enable
```

20 Appendix B: Upgrade Procedures

As part of the upgrade from 3.0.0-42 to 3.0.1-11 a larger XMS server is required. This is the primary change with the newer software release. The following are the high level steps for upgrading the XtremIO cluster.

1. Verify all host paths are active with ESX team
2. EMC will run health check script and collect dossier
3. Download new OVA file for deployment.
4. Copy any configuration data from existing XMS server (Ex. IP, /etc/shadow, etc...)
5. Shutdown existing XMS server
6. VM team deploys the new XMS from OVA file. (2 vCPU, 8GB vRAM, 900GB storage)
7. Upload 3.0.0-42 and 3.0.1-11 software files to new XMS
8. EMC support will run XMS install for 3.0.0-42
9. EMC support runs XMS recovery (XMS restore)
10. EMC support reconfigures Timezone and NTP servers
11. EMC will disable the SYR, SNMP, Email notifications
12. EMC will run test-cluster-upgradability and the health check script to verify no issues exist
13. EMC will run upgrade-cluster command to initiate the upgrade to 3.0.1 (Single X-brick should be 45 mins to an hour in duration)
14. EMC will run various commands and health check script to verify no issues exist
15. Create-debug-info will be run and dossier should be uploaded

21 Document Version Control

This section contains a detailed list of the updates made during each document release. Section 1 is purely an overview.

Version	Changes
0.1	Initial draft – general CATE peer review
0.2	Release For Comment
0.3	- Updates following global storage review 14.2 & 15.4 Audit log forwarding capability 14.5.3 LDAPS functionality 14.5.4-5 FID and root requirements - Removed reference to specific service pack version 14.5.6 ACL rules with IPTables 12.3.4 Capacity Monitoring with SRM
1.0	- Release for Production 9.6.2 XtremIO Cluster Configurations 9.6.4 XtremIO Power & Weight Specifications 9.6.5 XtremIO Connectivity - IPMI 12.3.4 Storage Capacity Alerting - Update Alert Definition 12.3.6 VAAI Thin Provision Limit 14.5.3.2 XMS LDAP Cross Domain Authentication 15.5 Event Monitor - Hostname 16.3 iSCSI Ports Configuration - Appendix A: Storage Controller Failover
1.1	- Release for Production - GNU BASH Vulnerability Fix 9.6.2 XtremIO Cluster Configuration (Usable TB / Thin Provision 130TB) 11.1 & 12.1 Thin Provisioning (130TB) 12.3.2 LUN Size Guidelines (130TB) 14.5.6 ACL w/ IPTables – Removed HotFix reference 18 XMS Management Server – Virtual XMS specifications vRAM & Storage increased 20 Appendix B: Upgrade Procedures