

# XtremIO Bulletin

Q4 2016

## Target Code

Dell EMC has established target code versions for each product to ensure stable and reliable environments. As a best practice, Dell EMC recommends that you operate at target code levels. There may be features, enhancements or fixes in newer code that may benefit or be required for a specific environment. Your Dell EMC account team will provide you with guidance. Unless otherwise directed by Dell EMC personnel, continue to use target code.

Code Families	Target Code
4.x, 3.x, 2.x	4.0.4-41

Note: There are cases that non-target code may be required. Consult your Dell EMC XtremIO account representative for details.

For the latest adoption rate for all XtremIO code families, refer to this link:

[https://support.emc.com/docu39695\\_Target-Revisions-and-Adoption-Rates.pdf](https://support.emc.com/docu39695_Target-Revisions-and-Adoption-Rates.pdf)

## XtremIO Code Recommendations

The official target code release is not the only generally available version of XIOS and XMS codes. Dell EMC KnowledgeBase has an article that specifies the XtremIO cluster configurations with recommended target code versions. Refer to this link:

<https://support.emc.com/kb/463324>

## KnowledgeBase Articles

Several recent KnowledgeBase articles have been published that address known issues for specific XtremIO environments and provide requirements for XIOS and XMS codes.

KB489171: XtremIO recommendation for customers using RecoverPoint with an XtremIO cluster. Refer to this link:

<https://support.emc.com/kb/489171>

KB491532: XtremIO requires performing Online Cluster Expansion using 4.0.15. Refer to this link:

<https://support.emc.com/kb/491532>

KB491612: XtremIO recommendation regarding NDU of 4 X-Brick (or greater) 40TB clusters to XIOS 4.0.10-33 or XIOS 4.0.15-15. Refer to this link:

<https://support.emc.com/kb/491612>

## Code Features

### XIOS and XMS 4.0.4

- Target code.
- Security Enhancement – Allowing the modification of the default SSH key to a unique SSH key per XMS.

- Enhanced Online Cluster Expansion of Single to Dual X-Brick - The Technician Advisor (internal Dell EMC Global Services tool) has been modified to support a more interactive user experience in this release performed by Dell EMC Global Services personnel.
- RecoverPoint - Optimizations and enhancements intended to improve replication operations.
- Consult the XtremIO 4.0.4-41 Release Notes for a complete list of fixes.

## XMS 4.2.0 and 4.2.1

- XMS and XIOS are now identified as different versions. However, they are included in the same upgrade bundle.
- Supports XtremIO clusters running version 4.0.2 and higher.
- Microsoft platform integration: PowerShell, SMI-S Provider, and VSS support. Enhancements create a seamless management experience for Microsoft Hyper-V, SCVMM and Azure platforms.
- Simplified, streamlined management interface. New HTML5 WebUI Technical Preview to obtain customer feedback to improve the WebUI usability.
- Efficient XtremIO configuration management. XMS cluster configuration import, export, and back-up.
- Displaying device reservations. New XMCLI command displays SCSI-2 and SCSI-3 persistent reservations on a device.
- New alert framework to reduce excessive alerts by consolidating repeating alerts.
- Reduced daily log bundle size.
- REST API enhancements.
- Consult the XtremIO 4.2.0 and 4.2.1 Release Notes for a complete list of fixes.

## XIOS 4.0.10 and 4.0.15

- Improved diagnostics adds new intelligent hardware control
  - Automatic error detection and deactivation of faulty SAS Connections, DIMMs, and Infiniband components.
  - Exceeding the configurable “High Ambient Temperature” threshold will trigger a graceful cluster shutdown to protect array data.
- Automated Health check during Non-Disruptive Upgrade.
- Scalability improvements in the number of Volume mappings per initiator group (now 4,096 Vol/IG).
- Fix for the “glibc getaddress stack overflow” vulnerability (CVE-2015-7547).
- ESX Datastore Accessibility. Administrators can choose between All Paths Down and Permanent Device Loss behaviors when host to device access is interrupted during a planned outage.

- Fix for higher latency in RecoverPoint environments with very low or continuous RPO (supported in 4.0.15).
- Consult XtremIO 4.0.10 and 4.0.15 Release Notes for a complete list of fixes.

## Simple Support Matrix

Simple Support Matrices are produced through extensive testing by Dell EMC E-Lab. They provide supported interoperability configurations that include storage arrays, operating systems, multipathing software and more. For XtremIO, RecoverPoint and other Dell EMC Simple Support Matrices, refer to this link:

<https://elabnavigator.emc.com/elab/modernHomeSSM>

## The OS Corner

### ESX PDL and APD Behavior

Permanent Device Loss (PDL) is a condition where all paths to a device are marked as “Dead.” Because the storage adapter cannot communicate to the device, its state is “Lost Communication.” Similarly, All Paths Down (APD) is a condition where all paths to a device are also marked as “Dead.” However, in this case, the storage adapter displays the state of the device as “Dead” or “Error.”

The purpose of differentiating PDL and APD in ESX 5.x and higher is to inform the operating system whether the paths are permanently or temporarily down. This affects whether or not ESX attempts to re-establish connectivity to the device.

A new feature in XtremIO 4.0.15 allows the storage administrator to configure the array to not send “PDL” as a response to a planned device removal. By default, XtremIO configures all ESX initiators with the PDL setting.

In planned device removals where the cluster has stopped its services and there are no responses received by the ESX host for I/O requests, ESX Datastores will respond with PDL or APD behavior depending on the XtremIO setting.

An XMCLI command is used to enable all ESX initiators as “APD” or revert back to “PDL.” A new option is available for the `modify-clusters-parameters` admin-level command which modifies various cluster parameters: `device-connectivity-mode=<apd, pdl>`

Refer to the XtremIO Storage Array User Guide ([https://support.emc.com/products/31111\\_XtremIO/Documentation/](https://support.emc.com/products/31111_XtremIO/Documentation/)) or `help` in XMCLI for more information on its usage.



## Single VMware ESX Host. Heterogeneous Dell EMC Storage.

In a typical SAN architecture, it is possible that hosts will have connectivity to multiple array types. Within the Dell EMC family of Fibre Channel arrays, an administrator can leverage the flexibility of the SAN to connect one host to more than one storage type in order to achieve the best performance for an application. Dedicated HBAs are not necessary when a host is presented with storage from Dell EMC XtremIO, Dell EMC VMAX, Dell EMC VNX, and Dell EMC VPLEX. It is important to verify that host settings are properly configured per documented recommendations. Customers currently wishing to connect their ESX hosts to two or more heterogeneous arrays, where Dell EMC XtremIO is one of them, should use this document and the included guidelines from Dell EMC to address these conflicts.

For more information, refer to this link:

<https://support.emc.com/kb/303782>

### AIX Queuing

In order to avoid potential storage performance bottleneck, queue depth settings should be appropriately configured. First, it is important to understand the number of queues, the types of queues and their functions, and the settings.

AIX disk and adapter drives have two queues; in-service queue and a wait queue. In-service queue IOs are sent to the storage and the queue slot is available when the IO is complete. IOs in the wait queue are sent to the in-service queue when a queue slot becomes available. The hdisk queue sits on top of the adapter queue in the host IO stack. From the hdisk device driver point of view, IOs in queue at the adapter driver level are considered in-flight.

The following table includes the performance tunables for AIX queuing:

Driver Name	Queue Name	Attribute	Recommended Value
hdisk driver	in-service	queue_depth	256 (lower if throttling is required)
	wait	None	N/A
fcs adapter driver	in-service	num_cmd_elems	2048 256 for VFCs
	wait	None	N/A
	DMA memory	max_xfer_size	Adapter and bandwidth dependent. See note A.

Note A: IBM documentation reference link: [IBM AIX/VIOS Disk and Adapter IO Queue Tuning](#) paper

The total IOs in-flight are based the number of adapter ports and the queue\_depth and num\_cmd\_elems settings for each of those ports.

## Optimization and Support Parameters

The XtremIO Host Configuration Guide is regularly updated with the latest best practices to optimize performance and maximize availability for hosts accessing XtremIO storage systems.

In this quarterly bulletin, the AIX ODM (Object Data Manager) is highlighted. Dell EMC recommends minimum versions of the Dell EMC ODM support package depending on the version of the AIX and the multi-pathing software (PowerPath or native MPIO).

AIX	PowerPath
AIX 7.1 and higher	EMC.AIX.6.0.0.4.tar.Z
AIX 6.1 and higher	EMC.AIX.5.3.0.9.tar.Z

AIX	MPIO
AIX 7.1 TL3 SP5 and higher	EMC.AIX.6.0.0.5.tar.Z
AIX 6.1 TL9 SP5 and higher	EMC.AIX.5.3.1.0.tar.Z

The Dell EMC ODM fileset EMC.AIX.6.0.0.5 also includes additional LUN information in the `lscfg -v1` output that is not included in version 6.0.0.4. The NAA identifier in the output of the command can be useful for troubleshooting whether PowerPath or native MPIO is in use.

```
>lscfg -v1 hdiskpower8
hdiskpower8          U78C9.001.WZS0E1X-P1-C3-T1-L5
PowerPath Device

Manufacturer.....XtremIO
Machine Type and Model.....XtremApp
ROS Level and ID.....4030
Device Specific.(Z0).....FNM00153900356
Device Specific.(Z1).....514F0C5C26E00015
NAA Identifier for the LUN  ————
```

For more host-specific details on AIX and other operating systems, refer to the complete XtremIO Host Configuration Guide:

[https://support.emc.com/docu56210\\_XtremIO-Host-Configuration-Guide.pdf?language=en\\_US](https://support.emc.com/docu56210_XtremIO-Host-Configuration-Guide.pdf?language=en_US)

### Customer Success Story

Hollard Insurance Group is Africa's largest privately owned insurance group. XtremIO has completely eliminated the data copy process for this customer by leveraging the iCDM capabilities of the array. Long term retention backups now run in less than 15% of the original time with no additional investment in hardware or software. The IT department now plans to consolidate all BI, data and analytics workloads onto XtremIO, thereby eliminating separate physical envi-

ronments for BI and analytics applications. XtremIO has also accelerated their virtualization project by allowing the teams to do more activities during production hours.

Hollard's decision to migrate to XtremIO has enhanced its business performance, improved productivity and saved costs. Hear Mark Nicholson talk about how XtremIO is a big part of their 'modern datacenter' strategy in the testimonial found [here](#).

## Environmental Considerations

### Online Cluster Expansions

An XtremIO cluster can be expanded non-disruptively to increase capacity and performance as more workload is added to the array. Customers starting with a smaller cluster can expand single-to-dual online cluster expansions as well as a multi-to-multi online cluster expansions.

With any installation, the storage administrator should assume that the cluster will be expanded even if there are no immediate plans to do so. With proper planning, expansion is simplified by localizing hardware and racks. Avoiding dispersed clusters, especially when expanding from single to dual X-Bricks, increases serviceability and availability. Plan for it in advance if possible.

### Prevent vXMS from Hanging

The VMware vCenter snapshot of a virtual machine preserves the state and data at a point in time. When the XtremIO vXMS is deployed on a VMware infrastructure, it is often included in the group of VMs that are snapped. There are processes running on the vXMS that cannot be properly quiesced when snapshots are performed. Snapshots of the vXMS instance are known to cause performance related issues, such as the vXMS becoming unresponsive.

Refer to this link: <https://support.emc.com/kb/476656>

## Ecosystem Updates

We know that XtremIO is a very powerful platform on its own, but we are also investing a lot of our efforts into building a complete ecosystem around the storage array. Here are some updates on that front from 2016:

- CommVault IntelliSnap now supports XtremIO to provide additional options for backup and copy creation.
- SolarWinds can be used to manage a customer's environment consisting of XtremIO.
- AppSync 3.0 release adds support for VPLEX and XtremIO deployments. This extends the iCDM benefits to those customers who are using XtremIO behind a VPLEX array.

- AppSync 3.1 release adds support for general file system repurpose, faster refresh operations using native XtremIO snap-and-reassign and staggered snapshot deletion.
- VPLEX 5.4.2 release allows for passing space reclamation commands from the host to the backend storage. It also enabled automated provisioning through VPLEX Integrated Array Services (VIAS).
- ProtectPoint enables direct backups from XtremIO to Data Domain without having any backup servers in the data path. The newest release adds support for SQL Server and Exchange, in addition to other applications that have already been supported since Q4, 2015.
- 15 licenses of RecoverPoint for Virtual Machines and vROPs management plugin is now included with every XtremIO order.
- vRO plugin for XtremIO is now available which facilitates the automation and orchestration of storage tasks and augments the capabilities of vRealize Orchestrator.

## Technical Advisory

ETA 485417: XtremIO, RecoverPoint: Trim, unmap, and zeroing operations performed against XtremIO volumes larger than 128GB replicated by RecoverPoint during an XtremIO High Availability event may result in data integrity issues on the replicated volumes.

For more information on this ETA, go to this link:

<https://support.emc.com/kb/485417>

ETA 484910: XtremIO: Storage Controller may not start on a Cluster with Encryption enabled following a power-cycle or cluster shutdown.

For more information on this ETA, go to this link:

<https://support.emc.com/kb/484910>

ETA 208184: XtremIO: SCSI3 reservations may result in service disruption on XtremIO 3.0.x clusters.

For more information on this ETA, go to this link:

<https://support.emc.com/kb/208184>

### Important Links and Updates

Documentation/Knowledgebase/Downloads

[https://support.emc.com/products/31111\\_XtremIO/Documentation/](https://support.emc.com/products/31111_XtremIO/Documentation/)

Additional resources (Videos, Testimonials, Datasheets)

<https://www.emc.com/en-us/storage/xtremio/resources.htm>

Security Update for Multiple Components

<https://support.emc.com/kb/486035>