



UPTIME BULLETIN

A Newsletter from EMSD, a Division of EMC

For VNX/VNXe

ACHIEVING BEST PERFORMANCE WITH VAAI

Block OE version 05.32.000.5.209 offers EMC's latest fixes and enhancements for VAAI performance. Many previously discovered performance issues have been addressed in this code, however a small number of customers have still observed less than optimal performance with either their xcopy operations, or in some cases with the performance of non-xcopy IO during xcopy operations to the same pool. Cross functional EMC tiger teams continue to work to resolve these issues on an urgent basis. Multiple potential avenues are being pursued.

Any findings learned as a result of the above efforts will be incorporated into future code releases. In the meantime, EMC's extensive lab testing has uncovered an ESX host setting which, if applied to all connected ESX hosts, may offer remediation to many customers suffering sub-optimal xcopy and non-xcopy performance issues during VAAI enabled xcopy operations. This solution requires that all attached storage supports a MaxHWTransferSize of 16MB. Upwardly adjusting the xcopy transfer size on ESX hosts from the default of 4MB to 16MB improves performance significantly in some cases, especially instances where fewer than 6 concurrent xcopy operations are being performed. A larger xcopy transfer size results in fewer concurrent xcopy I/Os using up the host queue depth. As a result, more of the queue depth is available for non-xcopy host I/O. The result is a fairer balance between xcopy and non-xcopy I/O on the ESX host side. To upwardly adjust the xcopy transfer size, the following command must be issued on all attached ESX hosts:

```
# esxcfg -advcfg -s 16384 /DataMover/MaxHWTransferSize
```

To verify the setting:

```
# esxcfg -advcfg -g /DataMover/MaxHWTransferSize
```

Validation of support for the 16MB transfer size for all storage systems in the node should be confirmed before implementation.

VOLUME 9

June 2014

VAAI performance	1
Mission Critical Center	1
Performance or bugcheck issues with Win2012 trim/unmap or RP	2
Celerra LUNs in ~physical storage group	2
Target code revisions and key fixes	3
Upgrading R33 storage arrays to R33.038 or later	4
VNX online documentation	4



Follow us on the web!

Search the VNX Series page for *Uptime Bulletin* here:

<https://support.emc.com/products/12781>

EMC MISSION CRITICAL CENTER

Mission Critical Center (MCC) is a collaborative initiative driven by EMC's Enterprise and Mid-Range Systems Division (EMSD) and the Data Protection & Availability Division (DPAD). Their goal is to ensure customers can build trusted IT infrastructure while also accelerating their organization's ability to capitalize on cloud and big data. MCC ensures continuous availability and continuous data protection by deploying a complex customer configuration from applications, through virtual layers, down to the storage arrays. Furthermore, MCC test runs pre-GA beta cycles (releases and service packs) through rigorous failure injection, failover and recovery operations through all system components while monitoring and reporting on workloads and identifying engineering improvements needed in the product.

Next quarter, the Uptime Bulletin will go into more detail about what MCC does, how it benefits EMC's customers, how MCC simulates years of production run time, and what role VNX plays within MCC's quarterly projects.

COMMENTS?
IDEAS?

WE'D LIKE YOUR FEEDBACK ABOUT THE UPTIME BULLETIN. SEND US YOUR IDEAS FOR FUTURE TOPICS AT :
UptimeforVNX@emc.com

CUSTOMER DOCUMENTATION

- <https://mydocuments.emc.com/VNX>
- <http://emc.com/vnxsupport>





VNX STORAGE PROCESSOR RESPONSE TIMES CAN INCREASE OR MAY INITIATE A BUGCHECK WHEN WINDOWS 2012 ISSUES TRIM/UNMAP COMMANDS TO THIN LUNS

As part of disk operations to reclaim free space from thin LUNs, Windows 2012 Server can issue large numbers of the SCSI command 0x9E/0x12 (Service Action/Get LBA Status). This SCSI command results in what is called a “DESCRIBE_EXTENTS” I/O on the VNX Storage Processor (SP.) These commands are used as part of the Trim/Unmap process to see if each logical block address (LBA) that has been freed up on the host’s file system is allocated on the VNX thin LUN. The host would then issue Unmap SCSI commands to shrink the allocated space in the thin LUN, thus freeing up blocks that were no longer in use in the file system. RecoverPoint also issues these same SCSI commands when the Thin LUN Extender mechanism is enabled, which can cause similar performance issues. See knowledgebase article [KB174052](#) for more information about the RecoverPoint variation of this issue and how to prevent it.

Windows 2012 appears to remain in the initial range of LBA for its “Get LBA Status” requests and does not accomplish the objective of the Trim/Unmap operation. This has been corrected in Windows 2012 R2. The repeated requests generate a high load on the SP CPU increasing proportionately with the number of volumes/LUNs that have this Trim/Unmap operation in progress at the same time. In some cases, the performance degradation can be so severe that it can cause one or both storage processors to reboot with a bugcheck. One tell-tale symptom for this particular bugcheck is that it most often occurs at or very near 3:00 AM, which is a default time when Windows 2012 server may perform Trim/Unmap operations.

If you suspect you may be seeing these symptoms, you should open up a ticket with VNX support to investigate fully. See knowledgebase article [KB182688](#) for extensive information about this known issue. The knowledgebase solution offers a workaround that involves disabling the feature that is triggering the Trim/Unmap operation, and that needs to be applied to all Windows 2012 server hosts attached to the storage system. You can also avoid this issue by exclusively using Thick LUNs as opposed to Thin LUNs. A VNX Block OE remediation for this issue is planned for the next major release within the R33 family of VNX Block OE. A timetable for a fix within the R32 code family has not yet been committed. For further questions, please open a ticket with support.

A UNIFIED VNX SYSTEM ONCE RUNNING 05.31.000.5.509 OR EARLIER MAY EXPERIENCE ADMINISTRATIVE PROBLEMS CAUSED BY PREVIOUSLY CREATED CELERRA PRIVATE LUNS IN THE ~PHYSICAL STORAGE GROUP.

A Unified VNX storage system once running the 05.31.x family of code may have its Celerra private LUNs present in the ~physical storage group. If this issue occurs, then both storage processors may at some point go unmanaged after they are both rebooted for any reason. Alternatively, if the storage system is upgraded to 05.32.000.5.209, some administrative actions may not be possible, such as binding LUNs, expanding storage pools, or creating FAST Cache. If you see any of these symptoms, please escalate to VNX support immediately for resolution. Refer to knowledgebase solution [KB91236](#) for more information about how to identify this issue.

The latest Block OE 05.32.000.5.215 and File OE 7.1.76.4 have mitigation for this issue.

RELEASE 31 FAMILY OF VNX CODE GOES END OF SERVICE LIFE (EOSL) IN JUNE, 2014

The Release 31 family of VNX code will go End of Service Life in June of 2014. As such, we will no longer be listing the recommended “target” or

“latest” code versions for this code branch on page 3 of the quarterly VNX *Uptime Bulletin*. Previously released VNX Block OE 05.31.x patches, hot fixes, and workarounds will remain available for an additional 6 months beyond the EOSL date.

EMC will no longer provide root cause analysis for any system running VNX Block OE 05.31.x

after June 30, 2014. In cases where EMC has not already released a solution or a workaround for a newly discovered critical bug, the only recommendation may be to upgrade to VNX for Block OE 05.32.x.



VNX/VNXE TARGET VERSIONS



EMC has established target revisions for each product to ensure stable and reliable environments. As a best practice, EMC recommends that you operate at target code levels or above to benefit from the latest enhancements and fixes available. Search using the term “adoption rates” in <http://support.emc.com> for current VNX/VNXe target code adoption rates.

VNXe OS VERSION	RELEASE DATE	STATUS
2.4.1.21519	12/16/13	Target
2.4.2.21519	12/16/13	Latest Release
UNIFIED VNX CODE VERSIONS (7.1 & R32)	RELEASE DATE	STATUS
7.1.74.5 (VNX for File)	12/13/13	Target
7.1.76.4 (VNX for File)	5/15/14	Latest Release
05.32.000.5.209 (VNX for Block)	12/13/13	Target
05.32.000.5.215 (VNX for Block)	5/15/14	Latest Release
UNIFIED VNX CODE VERSIONS (8.1 & R33)	RELEASE DATE	STATUS
8.1.2.51 (VNX for File)	2/28/14	Target
8.1.2.51 (VNX for File)	2/28/14	Latest Release
05.33.000.5.052 (VNX for Block)	5/2/14	Target
05.33.000.5.052 (VNX for Block)	5/2/14	Latest Release

See Product Release Notes for a full list of enhancements per new code release.

VNXe OE CODE ENHANCEMENTS IN RELEASE 2.4.1.21519

- Fixes an issue preventing recovery of virtual disks larger than 2 TB.
- Fixes a performance issue that could occur when an interface is down.
- Fixes an issue that could prevent users from extending CIFS shares.
- Fixes multiple issues that could prevent users from managing the storage system.

FILE CODE ENHANCEMENTS IN BOTH 7.1.74.5 & 8.1.1.33

- Fixes an issue that prevented a file system with dedup enabled and replicated to a pre-7.1 Celerra/VNX from being unmounted. Can hang replication operations.
- Fixes an issue which could hang a data mover when reading a 4GB or larger file.

FLARE CODE ENHANCEMENTS IN RELEASE 05.32.000.5.209

- Fixes a VAAI performance issue.
- Unisphere “One button shutdown” feature.

FLARE CODE ENHANCEMENTS IN RELEASE 05.32.000.5.215

- Fixes multiple SP bugchecks.
- Resolves an issue with Celerra control LUNs incorrectly in the ~physical storage group, resulting in Celerra LUNs offline.
- Has a fix that attempts to prevent a network storm on the management port or the service LAN port from resulting in an SP panic.
- ESRS/UDoctor now resides on the SP.
- Improved drive error handling for some drive types.
- Resolves an issue that prevents disks from coming back online after they have been shut down due to an over temperature condition in the enclosure.

FLARE CODE ENHANCEMENTS IN RELEASE 05.33.000.5.038

- Fixes an issue where the BBU self test could execute reiteratively and this could cause various FRUs to go offline/reboot.
- Fixes an issue whereby storage processors could reboot after as few as 80 days.
- Fixes a serious deduplication issue which could lead to data inconsistency under certain circumstances.

FLARE CODE ENHANCEMENTS IN RELEASE 05.33.000.5.052

- Contains all of the previous enhancements from R33.051 and a fix to prevent a possible LUNs offline or SP bugcheck due to interaction between the dedup feature and ESX host VAAI hardware acceleration functionality. This is the best available code for customers using the dedup feature. Further fixes/enhancements to this feature are due out in the Q3 and Q4 releases of R33.
- Fixes an issue where pool LUNs can go offline due to frequent trespassing. (Fix first appears in R33.051).
- USM will now auto-notify when new OE software is available (as of R33.051).

FILE CODE ENHANCEMENTS IN BOTH 8.1.2.51

- Fixes an issue where the in-memory file system usage grows faster than the actual usage, leading to a discrepancy in accounting or in some cases the inability to write.
- Fixes an issue where NFS access may become blocked when NFS clients attempt to access mount paths which do not exist.
- Fixes a potential watchdog panic when a datamover runs multiple concurrent backup sessions for a long time.



THE IMPORTANCE OF UPGRADING ALL RELEASE 33 VNX STORAGE SYSTEMS TO 05.33.000.5.038 OR LATER

At the time of this publication, EMC's official target version of code within the R33 family is 05.33.000.5.052. EMC strongly recommends that all customers running code within the R33 family upgrade to 05.33.000.5.038 or later (preferably 05.33.000.5.052) to take advantage of some major fixes detailed below.

Block OE version 05.33.000.5.038 (also referred to as 33.038) was released on 01/13/2014 as a critical fix service pack intended to fix, among other things, the following 3 issues:

1. ETA [175619: VNX Systems may experience a reboot after 80 days of runtime unless a code upgrade or a staggered reboot is implemented as described in this advisory](#). VNX storage systems running code in the R33 family that is older than 05.33.000.5.038 may experience unscheduled reboots of both of their storage processors after as little as 80 days of uptime. The exact number of days varies by model. The SP reboots may come in close enough proximity to cause the loss of SP cache. See the linked ETA for specific details and steps that may be taken to mitigate risk prior to your code upgrade.
2. An issue that may cause LUNs in block deduplication pools to go offline under certain circumstances. This issue does not affect smaller block deduplication pools. To be safe however, all customers running R33 code and using block deduplication are urged to upgrade to 05.33.000.5.038 or later. If the block Deduplication LUNs go offline due to inconsistency, the array must be upgraded before the LUNs can be brought back online.
3. An issue where the BBU self-test could run on a daily basis instead of weekly. This issue was partially fixed in 05.33.000.5.035; however the full fix is in 05.33.000.5.038. In addition to running every day, the BBU self-test can also run reiteratively. During the time that it is running, it is possible for various FRUs to go offline. FRUs most commonly affected are Management Modules, certain IO Modules, or even the storage processor itself. This issue has resulted in many unnecessary hardware replacements. Replacement hardware will not remedy the problem. The issue is resolved in 05.33.000.5.038 and later code.

Additional customers who should prioritize a code upgrade

For customers who are running 05.33.000.5.051 and using both hardware acceleration features (VAAI or ODX) and the block deduplication feature of the VNX storage system, EMC recommends upgrading to 05.33.000.5.052 as soon as possible. Customers using the block deduplication feature in VNX2 storage systems with VMware ESX server or Microsoft Windows 2012 server may lose access to a set of deduplication enabled LUNs. The deduplication LUNs may go offline and require recovery when attached to hosts that utilize xcopy or ODX hardware acceleration. If this issue occurs, data may not be completely recoverable. See [ETA186733](#) for more information about this issue.

In addition to the above fixes, there are also a number of resolved bugchecks and other critical fixes in releases 05.33.000.5.038 and later. For a complete list of fixes, please see the *Release Notes*. The current EMC recommended Target and Latest code release for the R33 family is 05.33.000.5.052.

AN ADDITIONAL SOURCE FOR VNX DOCUMENTATION

By visiting <https://mydocuments.emc.com> and selecting **VNX Series**, you can learn more about your VNX product, read the most up-to-date list of known issues, and create customized installation and maintenance procedures tailored to your environment. Visit the Related Documentation links at the bottom of the VNX Series page to download a zip file of commonly accessed content. Please visit the site and send us your feedback!

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license. EMC², EMC, E-Lab, Powerlink, VNX, VNXe, Unisphere, RecoverPoint, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. Copyright © 2014 EMC Corporation. All rights reserved. Published in the USA, June, 2014.