NexentaStor 4.0.5
Release Notes

NexentaStor 4.0.5 delivers new features (see What's New in 4.0.5?), along with product enhancements and fixes that address and improve stability. Each NexentaStor 4.0.x release builds on previous fixes and enhancements, while addressing customer-reported problems and issues found internally by Nexenta engineering.

NexentaStor 4.0.5 Release History

- NexentaStor 4.0.5: 12/13/16

Document Revision History

Date	Description
June, 2017	Added a note in the release notes about disabling USB3 in the BIOS settings before installing NexentaStor 5.0.3

Table of Contents

New Features, Enhancements, Resolved and Known Issues

NexentaStor 4.0.5 encompasses several features, improvements, and fixes in multiple areas.

What's New in 4.0.5?

This section summarizes the notable enhancements and changes in 4.0.5.

Self-Encrypting Drive Support

With NexentaStor 4.0.5, Nexenta expands its Software Defined Storage portfolio with support for self-encrypting drives (SED). Self-encrypting drives use an internal AES-256 algorithm in the firmware to encrypt and decrypt drive data. The encryption/decryption process is combined with an Authentication Key (AK) to lock or unlock a drive.

Nexenta 4.0.5 uses the Gemalto SafeNet external key manager to store authentication keys in accordance with the Trusted Computing Group KMIP specifications. No keys are stored on the Nexenta appliance. Users will now be able to create individual pools taking advantage of the securities offered by SEDs, providing a solid foundation for addressing data-at-rest security requirements.

Nexenta SED supports Seagate ST4000NM0043 CONSTELLATION 4TB and HGST HDD-A4TBHUS726040ALS211 4TB drives for creation of pools in clustered or standalone configurations. Additional drives that meet the TSG Enterprise specification and are supported by the specific DELL or SMC JBODs may also be considered. For questions, please contact your sales representative.

Support for New Chassis and Devices

NexentaStor 4.0.5 adds chassis management for the following storage enclosures:

- AMAX STorMax NX224
- DELL MD1280 JBOD
- Ericsson HDS8000 Chassis Management
- HGST Storage Enclosure 4U60
- SanDisk InfiniFlash IF-100 and IF-150
- Toyou NCS3700
- Zstor Q12NexentaStor 4.0.5 also adds auto-pool builder support for the following:

Toshiba and SanDisk devices for SLOG or Cache support

- LTO400MO – SanDisk 400Gb (L2ARC)
- PX04SHB040 – Toshiba 400Gg WI (SLOG)
- PX04SMB040 - Toshiba 400GB RI (L2ARC)
- PX04SRB048 – Toshiba 480Gb RI (L2ARC)

HGST device SLOG additions

- HUSMH8010BSS200
- HUSMH8020BSS200
- HUSMH8040BSS200
- HUSMH8080BSS200

Ericsson devices

- CACHE: 3 DWPD – L2ARC (PX04SVB096) (a.k.a. cache)
- SLOG: 10 DWPD – SLOG (PX04SMB080) (a.k.a log / ZIL)

SMB Enhancements

NexentaStor 4.0.5 provides several critical SMB fixes and enhancements.

Security Enhancements

NexentaStor 4.0.5 continues to focus on ensuring the highest level of security. As such, we have included fixes for several CVEs and other security issues, as follows:

- Upgraded OpenSSL to version 1.0.2
- Prevention for Click Jacking
- Support for newer Apache – v2.4.23
- Removed conflicting TLS/SSL directives; other general TLS/SSL improvements
- Prevention for HTTP request smuggling attack against chunked request parser
- Prevention for Test-cgi Script Information Disclosure Vulnerability

Driver Additions and Enhancements

NexentaStor 4.0.5 packages and provides support for the following new drivers:

- Intel 40GbE network adapter – XL710
- Change to make the LSI mr-sas driver as default
- Changes and improvements to the mpt_sas driver
- QLogic 16G FC
- ATTO Celerity FC-162E Gen 5 and Celerity FC-162P Gen 6 16Gb FC cards
- bnxe driver for newer QLogic 10GE NICs

In addition to the individual drivers listed above, NexentaStor now supports the necessary driver set for HP Gen9 servers.

Resolved Issues in 4.0.5

Table 1 lists the issues that have been resolved in NexentaStor 4.0.5.

TABLE 1: NEXENTASTOR 4.0.5 ENHANCEMENTS

Component	Key	Description
Command, Daemons	SUP-549	Resolved an issue where users could not remove the Domain Admin from the local administrators group.
HA, COMSTAR, Fibre Channel	NEX-3648	Resolved, the underlying issue was addressed by VMware in ESXi 6.0 Update 2.
HA, Plugin	NEX-5147	Resolved an issue where users creating mappings from NMV might intermittently encounter an STMF error.
Installation	NEX-6485	Resolved an issue where systems with large amounts of memory were not getting allocated a large enough swap space at the time of installation.
Kernel	NEX-6018	Resolved a condition where, under certain heavy loads, iSCSI LUNs in an off-lining state could result in hung threads.
Kernel	NEX-6064	Resolved a condition where SSD's (solid state drives) could fail to respond because of an internal failure, causing the SAS adapter to be triggered to re-enumerate the disk topology, leading to system hangs.

Kernel	NEX-6135	Resolved an issue where exports of pools with datasets shared via SMB could be delayed while the SMB datasets with quotas enabled were being unshared.
Kernel	NEX-7273	Resolved an issue where RENAME operations could become inoperable when enabling nbmand on an NFSv4 filesystem.
Kernel	NEX-7362	Resolved an issue where stuck threads in door calls to idmapd could cause CIFS to hang.
Kernel	NEX-8441	Resolved an issue where a failing SMB2 lock request could prevent certain files from opening.
NMC	NEX-8672	Added the “-e” option for the NMC “lunsync” command.
Protocols	NEX-2522	Addressed a potential SMB process hang while waiting for taskq threads to exit.
Protocols	NEX-5956	Resolved an issue applying ACLs when copying or duplicating files in the macOS Finder.
Protocols	NEX-8077	Enhancements and improvements to help increase SMB performance and stability.
Protocols	NEX-8495	Incorporated fix for Illumos 7483 bug to eliminate a panic upon an SMB flush request (SMB1 or SMB2) with an open named pipe.

Resolved Security Issues in 4.0.5

Table 2 describes security issues resolved in NexentaStor 4.0.5.

TABLE 2: NEXENTASTOR 4.0.4-FP5 RESOLVED ISSUES

Component	Key	Resolved Issue Description
Appliance Mgmt	NEX-2408	Resolved an issue with clickjacking, also known as a UI redress attack. Clickjacking is a method in which an attacker uses multiple transparent or opaque layers to trick a user into clicking a button or link on a page other than the one they believe they are clicking. The attacker is in effect "hijacking" clicks meant for one page and routing the user to an illegitimate page.
Appliance Mgmt, Packaging	NEX-5857	Resolved an issue with an OpenSSL X509_ATTRIBUTE memory leak (CVE-2015-3195) (http-openssl-cve-2015-3195) that allowed remote attackers to obtain sensitive information from process memory by triggering a decoding failure.
Packaging	NEX-4658	Resolved an issue where an HTTP request smuggling attack was possible due to a bug in parsing of chunked requests. A malicious client could force the server to misinterpret the request length, allowing cache poisoning or credential hijacking if an intermediary proxy was in use.

Packaging	NEX-5867	Resolved an issue where TLS server used a Diffie-Hellman group with a prime modulus of less than 2048 bits in length. It is estimated that an academic team can break a 768-bit prime and that a state-level actor can break a 1024-bit prime.
Packaging	NEX-6043	Resolved an issue where TLS/SSL Server supported 3DES Cipher Suite (ssl-3des-ciphers). Since 3DES only provides an effective security of 112 bits, it is considered close to end of life by some agencies.
Packaging	NEX-6187	Resolved an issue with TLS protocol 1.2 and earlier, when a DHE_EXPORT cipher suite is enabled on a server but not on a client, the protocol does not properly convey a DHE_EXPORT choice. This enabled man-in-the-middle attackers to conduct cipher-downgrade attacks.
Packaging	NEX-6188	Resolved an issue where TLS/SSL server did not support modern, secure ciphers. TLS/SSL server only supported ciphers known to be vulnerable to attack.
Packaging	NEX-6189	Resolved an issue where the Transport Layer Security (TLS) versions 1.0 (RFC 2246) and 1.1 (RFC 4346) include cipher suites based on the DES (Data Encryption Standard) and IDEA (International Data Encryption Algorithm) algorithms. DES and IDEA algorithms are no longer recommended for general use in TLS and have been removed from TLS version 1.2.
Packaging	NEX-6190	Resolved an issue where the TLS/SSL server supported export cipher suites, intentionally crippled to conform to US export laws. Symmetric ciphers used in export cipher suites typically do not exceed 56 bits.
Packaging	NEX-6191	Resolved an issue where TLS/SSL server supported RSA-based cipher suites intentionally weakened due to export control regulations. This enabled an attacker to launch man-in-the-middle attacks and monitor or tamper with sensitive data against clients susceptible to the vulnerability.
Packaging	NEX-6193	Resolved an issue where the TLS server used a Diffie-Hellman group with a prime modulus of less than 1024 bits in length. It is estimated that an academic team can break a 768-bit prime and that a state-level actor can break a 1024-bit prime.
Packaging	NEX-8719	Resolved an issue where the web server made a test script available that revealed details of the web server configuration to anyone who can connect to the machine.
Packaging	NEX-8720	Resolved an issue with the HTTP TRACE method that is normally used to return a full HTTP request back to the requesting client for proxy-debugging purposes. An attacker can create a webpage using XMLHTTP, ActiveX, or XMLDOM to cause a client to issue a TRACE request and capture the client's cookies. This effectively results in a Cross-Site Scripting attack.

Known Issues in 4.0.5

Table 3 lists the issues known in NexentaStor 4.0.5 as of November 2016.

TABLE 3: NEXENTASTOR 4.0.5 KNOWN ISSUES

Components	Key	Known Issue Description	Workaround
Appliance Mgmt	NEX-3226	Modifying a syslog configuration using NMV may truncate the syslog.conf file.	Do not use the GUI to modify syslog server settings. Use this NMC command: setup network service syslog-daemon edit-settings syslog.conf
Appliance Mgmt	NEX-5041	Users are currently unable to monitor and report issues with SMCX10 server power supplies and fans.	If iPMI is being used, check the documentation to determine if the IPMI vendor supports accessing power supply and fan data.
Autosync	NEX-5228	Users may encounter an error when destroying and recreating auto-sync jobs.	If this issue is encountered, contact Nexenta Support for assistance with setting up a new auto-sync job.
Autosync	NEX-5830	There is an intermittent, rare condition where an NFS outage may lead to the rrdaemon utilization to jump to 100%.	There is no known workaround.
Autosync	NEX-5835	The unmap_zvols general-flag does not work in flip direction, failing with a "dataset is busy" error.	Use a local before_replication Action Script to a) save existing COMSTAR mappings on the new destination, then b) delete COMSTAR mappings on the new destination. **Note:** Contact Nexenta Support for assistance with implementing this workaround.
Autosync, Kernel	NEX-5239	Recursive auto-sync may stop functioning correctly if new child datasets are added between auto-sync runs.	In the event of this issue, do the following: 1. Disable the job. 2. Determine which datasets need to be removed at the destination. 3. Re-enable the job. **Note:** Contact Nexenta Support for assistance with this procedure as necessary.
COMSTAR	NEX-4246	The default setting of stmf_sbd:stmf_standby_fail_reads=0 can result in impacted performance.	In Microsoft-only server environments, setting stmf_sbd:stmf_standby_fail_reads = 1, followed by a reboot, resolves this issue. However, in any mixed-server environment, this value must be left at zero to support proper LUN discovery. The following commands will create a system checkpoint, then add the necessary setting to /etc/system. Please issue these commands in NMC and then reboot, note that both cluster nodes should be updated: nmc@:> setup appliance checkpoint create nmc@:> options expert_mode=1 nmc@:> !echo "set stmf_sbd:stmf_standby_fail_reads=1" >> /etc/system nmc@:> setup appliance reboot
HA	NEX-3191	There may be an export failure on failover in clusters with a large number of nfs mounts	If an automatic failover times out, manually initiate the failover.

Components	Key	Known Issue Description	Workaround
		and replication jobs.	
HA	NEX-5593	Mappings added via the idmap command can potentially be lost upon an HA failover.	Instead of the idmap command, use NMV GUI 'Identity Mapping' to perform the removal, then recreate the mappings. If a high number of mappings make this workaround prohibitive, please contact Nexenta Support.
HA	NEX-9032	You can configure pools in a cluster with mismatched controller IDs between nodes.	Verify that the controller IDs for both nodes match before deploying the cluster.
HA, Appliance Mgmt, COMSTAR	NEX-6040	There is currently no indication when a COMSTAR configuration becomes out of sync between two HA nodes.	Go to bash and manually check for the existence of the .comstar/config-out-of-sync file that is created in the root of the clustered volumes.
HA, COMSTAR	NEX-5315	ALUA can be unintentionally disabled after a hard reset of the passive node.	Re-enable ALUA if needed via NexentaStor GUI (NMV) or command-line interface (NMC).
Installation	NEX-1881	Under certain circumstances, NexentaStor clusters can have mismatched controller numbers between the nodes.	Contact the system installation engineer or Nexenta Support to manually reconcile controller numbers.
Installation	NEX-3488	Unable to boot NexentaStor from a drive with 4k native sector size.	Use 512 native or 512 emulated drives for NexentaStor installations.
Installation	NEX-5548	GRUB boot loader fails checksum verification and prevents booting of checkpoints after an upgrade.	If you experience this issue, boot to a recovery CD and reinstall GRUB in the following way: Boot recovery console from the install media, then perform the following procedure (assuming c1t0d0s0 and c1t1d0s0 are the mirrored boot disks): # zpool import -f syspool # zfs list -r syspool rootfs-nmu-003 # mkdir /tmp/syspool # mount -F zfs syspool/rootfs-nmu-003 /tmp/syspool # rm -f /tmp/syspool/etc/zfs/zpool.cache # bootadm update-archive -R /tmp/syspool # cd /tmp/syspool/boot/grub # installgrub -f -m stage[12] /dev/rdsk/c1t0d0s0 # installgrub -f -m stage[12] /dev/rdsk/c1t1d0s0 # umount /tmp/syspool # sync # reboot
Kernel	NEX-928	When using ZEUS IOPS drives in a JBOD, an mptsas deadlock may occur due to a poor connection with the backplane.	Ensure that required components are installed and properly configured when using ZEUS IOPS drives in a JBOD.
Kernel	NEX-1760	ZFS exhibits long kmem reap times in certain situations.	There is no known workaround.

Components	Key	Known Issue Description	Workaround
Kernel	NEX-2899	The "zfs send -I" command performed on a snapshot will close any files opened within that snapshot, leading to potential IO errors.	It is recommended that users clone the snapshot first, then access files from the clone, if your use case permits.
Kernel	NEX-2940	Disk pools with a failed sTEC drive as a single ZIL can cause a system panic when users try to remove the failed ZIL.	Use redundantly configured (mirrored) ZILs.
Kernel	NEX-3043	Alternating I/Os to datasets of different record sizes can cause long zio_cache reaps.	There is no known workaround.
Kernel	NEX-3585	An intermittent issue where VM slack in non-ARC ZFS kmem caches can degrade ARC performance.	There is no known workaround.
Kernel	NEX-3734	ZFS allows setting of a duplicate mount point path on two different ZFS filesystems, creating broken volume services.	Check the pool for duplicate mount points before failover, then perform manual remediation.
Kernel	NEX-4393	In some situations, the slow I/O diagnosis engine may identify disks experiencing high latency. Slow I/O may produce a message that indicates an attempt to retire a disk.	Unless slow I/O disk retirement has been explicitly enabled, disregard this message. By default, slow I/O will not attempt to retire a device.
Kernel	NEX-5308	GRUB menu mistakenly reports 32-bit in a 64-bit environment, possibly leading to issues when upgrading via undocumented methods.	Ignore the 32-bit entry in the GRUB menu. The environment is 64-bit. **Note:** Upgrading NexentaStor should always be performed using the NMC 'setup appliance upgrade' command.
NMS	SUP-737	Over time, NMV may grow heap memory while failing to reclaim allocations.	Restart NMS if large amounts of memory are being used.
NMS	NEX-4237	Unexpected behavior after failover may result from restoring an old system configuration after making nameservice changes.	Contact Nexenta Support for assistance resolving if this issue is encountered.
NMS, Plugin	NEX-2097	A failover that occurs when the COMSTAR configurations between two cluster nodes cluster are not synchronized can cause the configuration to not be restored.	Contact Nexenta Support for assistance synchronizing the COMSTAR configurations.

Recommendations:

- NexentaStor 4.0.5 does not support USB3. So you must disable USB3 in the BIOS before installing NexentaStor 4.0.5.
- All customers using VMware must follow VMware recommendations for maintaining VMware vSphere.
- For more details on the issue related to simultaneous file modifications in a mixed SMB/NFS environment, refer to the KB article number 1359 titled "Considerations when using multi-protocol file locking" in the Customer and Partner Portal.
- For issues related to Linux clients failing to see recovered paths after a clustered node reboot, refer to the KB article number 1361 titled "Linux clients do not see stand-by path" in the Customer and Partner Portal.

Upgrading to 4.0.5

This section covers how to upgrade NexentaStor to version 4.0.5. First read the following information, then follow the appropriate upgrade instructions.

- Upgrading from 4.0.x
- Upgrading from 3.1.x

Note: To upgrade the HA Cluster Plugin, see *HA Cluster User Guide*.

Before You Upgrade

- **Review system requirements**, SMB-supported client operating systems, the NMV port number (8457 for all 4.x releases), and other installation changes that occurred with previous 4.0.x releases. This information is available in the Upgrade sections of the previous NexentaStor 4.0.x Release Notes.
- **Review the *NexentaStor 4.0.5 Installation Guide*** for additional details on upgrading to NexentaStor 4.0.5.
- **Review the Hardware Components List** (HCL) to ensure that your current hardware is compatible with upgrading to NexentaStor 4.0.5.
- **Ensure that you don't have any 3rd-party packages running** on NexentaStor. Upgrading NexentaStor will cause those packages to be deleted.
- **Allow Auto-Sync and Auto-Snap jobs to finish processing** before upgrading NexentaStor. Rebooting into NexentaStor is required to complete the upgrade process.

Frequently Asked Questions and Guidelines for Upgrading

Version	**Question**	*How do I know which NexentaStor version I currently have installed?*
	Answer	To determine the NexentaStor version you currently have installed, use the following command at the `nmc` prompt: `nmc:/$ show appliance version`
Availability	**Question**	*Do services and volumes remain available to clients during the upgrade and required restart?*
	Answer	During the upgrade, NexentaStor services and volumes remain available to network clients. During the required system restart after upgrading, however, NexentaStor services are not available; therefore, we recommend that you schedule the upgrade and restart during a scheduled system maintenance window.
New ID?	**Question**	*What if I'm upgrading NexentaStor onto a new machine or a system with significantly different hardware?*
	Answer	If your machine ID has changed, visit the Customer Portal or Partner Portal and get a new License Key. To do so, you'll need to provide: • The old license key • The sales order that applies to the old license key • The new machine ID
3rd party apps	**Question**	*Can I upgrade with third-party applications on my NexentaStor appliance?*
	Answer	You might have third-party packages installed if you changed repository sources on your NexentaStor appliance. *Upgrading with third-party packages installed on NexentaStor will result in the loss of components that are not included with the NexentaStor build.*
Upgrade after rollback?	**Question**	*Can I roll back NexentaStor to a previous version after upgrading to the latest FixPack?*
	Answer	You can roll back to a previous 4.0.x release; however, we do not recommend rolling back to 3.1.6-FP4 after upgrading to the latest FixPack on a production system. In particular, if you upgrade the volume version, rollback to 3.1.6-FP4 will not be possible.
No Internet	**Question**	What if I don't have an Internet connection?
	Answer	If you do not have an Internet connection and want to upgrade NexentaStor, contact support@nexenta.com.

Incorrectly named checkpoint	**Question**	What steps should I take if, when upgrading specifically from 4.0.4-FP3 to 4.0.4-FP5, I discover an incorrectly named checkpoint?
	Answer	If you experience this issue, do the following: 1. Find the current checkpoint (marked 'Yes' in the CURRENT column) using NMC: nmc@my_liler:/$ show appliance checkpoint rootfs-nmu-215 Mar 2 9:27 2016 upgrade Yes Yes 4.0.4-FP3 2. Enable expert mode: nmc@my_filer:/$ option expert_mode = 1 3. Enter into bash: nmc@my_filer:/$!bash You are about to enter the Unix ("raw") shell and execute low-level Unix command(s). Warning: using low-level Unix commands is not recommended! Execute? (y/n) <press y> Select Yes 4. Using nano or vim to edit /etc/default/versions file: nmc@my_filer:/$ nano /etc/default/versions 5. Locate the line that contains the current checkpoint name. In that line, change 4.0.4-FP3 to 4.0.4-FP5: syspool/rootfs-nmu-215 4.0.4-FP5 6. Enter into the NMC by pressing crtrl+d. 7. Issue nmc command: setup appliance nms restart 8. Verify that name is correct: nmc@my_filer:/$ show appliance checkpoint
Upgrade timeout	**Question**	What steps should I take if an upgrade fails with a 'Failed to gain exclusive access, operation timed out' message while auto-sync replication is in progress?
	Answer	Wait until auto-sync replication completes, and then retry the upgrade.
Reset HA reservation drives	**Question**	What steps should I take if an upgrade process resets the number of HA reservation drives to the default number of 2?
	Answer	The correct number of reservation drives should be one more than the parity level of the pool. For example, a RAIDZ2 pool should have 3 reservation drives. To query the current number of reservation drives, issue the following command from a root bash shell: root@nexenta35:/# /opt/HAC/RSF-1/bin/rsfcdb list_props |grep prop_scsi2_drive_count To change the number of reservation drives, use the following command from the root bash shell. Not that this command automatically replicates the setting to the other cluster node: root@nexenta35:/# /opt/HAC/RSF-1/bin/rsfcdb update prop_scsi2_drive_count 3 Please contact Nexenta Support if you require assistance setting the number of reservation drives.
Cluster failover	**Note**	There can be issues with cluster failover after upgrading a system using PGR3 Reservations to a later release using SCSI-2 Reservations.
	Guideline	Configurations using STEC SAS SSDs as data drives with firmware revision E50x, or earlier, should not be upgraded until the device manufacturer issues a firmware update to resolve this issue. Configurations using STEC SAS SSDs as cache or log devices are not affected by this restriction.
Incorrect checkpoint	**Note**	Upon upgrade, systems with incorrectly set time and date can boot to an incorrect checkpoint.
	Guideline	Before starting an installation or upgrade, ensure that the system time and date are set correctly. If this issue is encountered, reboot the system to the correct checkpoint.
Manual reset of nfsmapid_domain	**Note**	After a seamless upgrade from 3.x to 4.x, the nfsmapid_domain setting is not maintained and must be reset manually.
	Guideline	SSH to the system and run the following command to reset the nfsmapid_domain: sharectl set –p nfsmapid_domain=<domain> nfs

Upgrading From 4.0.x

1. `nmc:/$ setup appliance upgrade -s`
2. Complete the upgrade.
3. Reboot your system.

Upgrading From 3.1.x

Upgrading NexentaStor from version 3.1.*x* to the latest 4.0.5 requires these interim upgrade steps:

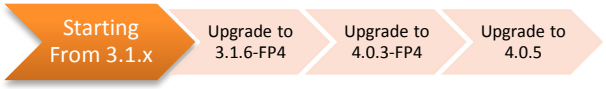

Step 1. Upgrade to 3.1.6-FP4

Upgrade to NexentaStor 3.1.6-FP4:

1. `nmc:/$ setup appliance upgrade`
2. `Complete the upgrade.`
3. Reboot your system.

Step 2. Upgrade to 4.0.3-FP4

Now upgrade to 4.0.3-FP4:

1. `nmc:/$ setup nexentastor upgrade -r 4.0.3`
2. Complete the upgrade.
3. Reboot your system.

Step 3. Upgrade to 4.0.5

And finally, upgrade to 4.0.5:

1. `nmc:/$ setup appliance upgrade -s`
2. Complete the upgrade.
3. Reboot your system.

Note: NMV for 4.x is now accessed at port 8457.

What to Expect During Upgrade

The instructions in this section show the basic upgrade processes from the most common starting points; however, depending on your previous installation and configuration, you may encounter additional configuration questions. We've addressed the most common ones here:

Cleanup the upgrade caches? Selecting Yes creates a rollback checkpoint, which is useful if you need to roll back to the installation prior to upgrading. You can view a list of available check points by using the `show appliance checkpoint` command.

Disabling and restarting multi-NMS? Upgrading NexentaStor requires that multi-NMS is disabled and restarted.

Is your hardware certified? Use the Hardware Certification List (HCL) to ensure that your hardware is compatible with NexentaStor. Using incompatible hardware may cause unexpected results and may also void your license. If your existing hardware is not included in the HCL, contact Nexenta Support.

Reboot the system? Yes, to complete the upgrade process, you'll need to reboot the system into NexentaStor. You can continue to work in a previous version—for example, if you have a process running that hasn't completed; however, rebooting into NexentaStor is required to complete the upgrade process.

For additional information about upgrading NexentaStor, see the *NexentaStor 4.0.5 Installation Guide*.

Completing Additional Configurations

Upgrading Data Volumes if You Will Not Be Booting in to a 3.1.x Checkpoint

1. Upgrade NexentaStor volumes to use ZFS feature flags:
   `setup volume <volname> version-upgrade`
2. Repeat to upgrade all NexentaStor volumes.

Resetting nfsmapid

After upgrading from 3.x, you will need to manually reset the `nfsmapid_domain` setting:

1. SSH to the system.
2. Log in to bash and type:
   `option expert_mode =1`
   `!bash`
3. Type:
   `sharectl set -p nfsmapid_domain=<domain> nfs`