



NexentaStor

Release Notes 4.0.4

Date: September 2015

Subject: NexentaStor Release Notes

Software: NexentaStor

Software Version: 4.0.4

Part Number: 7000-nxs-4.0.4-000017-B

Copyright © 2015 Nexenta Systems™, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose, without the express written permission of Nexenta Systems (hereinafter referred to as “Nexenta”).

Nexenta reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. Nexenta products and services only can be ordered under the terms and conditions of Nexenta Systems’ applicable agreements. All of the features described in this document may not be available currently. Refer to the latest product announcement or contact your local Nexenta Systems sales office for information on feature and product availability. This document includes the latest information available at the time of publication.

Nexenta, NexentaStor, NexentaEdge, and NexentaConnect are registered trademarks of Nexenta Systems in the United States and other countries. All other trademarks, service marks, and company names in this document are properties of their respective owners.

Contents

What is New in this Release?	1
Recommendations	2
VAAI Block Changes	3
Warning about upgrading without preparation	4
Planning for upgrade	5
Identifying VMFS datastore extents and checking ATS status	5
Disabling ATS on Existing Datastores	6
Enabling ATS on Existing Datastores	7
Enabling ATS on NexentaStor	8
Reverting to pre-FP3 checkpoints	8
IPMI Support	9
Determining the Version of the Appliance	9
System Requirements	9
List of SMB Supported Client Operating Systems	9
Upgrading	11
Upgrading Minor Versions of NexentaStor 4.0.x	11
Change in NMV Port	11
Upgrading from Version 3.1.x to 4.0.4	11
Upgrading from Latest Version of 3.1.6 to 4.0.4 with an Internet Connection	12
Upgrading from Latest Version of 3.1.6 to 4.0.4 without an Internet Connection	14
Upgrading to Version 4.0.4 After Rolling Back to Latest Version of 3.1.6	14
Enhancements	16
Resolved Common Vulnerabilities and Exposures (CVE)	17
Resolved Issues	18
Known Issues	21

What is New in this Release?

NexentaStor 4.0.4 is a maintenance update that delivers significant improvements in the areas of general stability, scalability and performance. NexentaStor 4.0.4 builds on all the fixes previously released as of version 4.0.3-FP4 and addresses customer reported issues, as well as issues found internally by Nexenta Engineering since 4.0.3. The following changes and enhancements are worth highlighting:

- **Fibre Channel and iSCSI Block services hardening**

Nexenta engineering took a holistic view of known Comstar and block services issues to refactor the block stack in NexentaStor 4.0.4. The result addresses a large number of known issues and delivers dramatic improvements in reliability of both Fibre Channel and iSCSI block services.

- **VMware certification for VAAI Block services**

With the Fibre Channel and iSCSI changes in NexentaStor 4.0.4, NexentaStor is now fully certified for VMware VAAI Block for ESXi 5.5 and ESXi 6.0. See the VMware Storage Compatibility Guide for more details. For more details in terms of NexentaStor support for ESXi 6.0, review NEX-3648 in the [Known Issues](#) section

- **ZFS file delete scalability improvements**

Previous releases of NexentaStor have exhibited performance challenges when having to delete large amounts of data at once. In VMware environments, that issue can be triggered by VM deletions from mounted datastores and exhibits itself as APD (All Path Down) events affecting virtual machines.

NexentaStor 4.0.4 contains a number of core ZFS enhancements improving the ability to handle deletes of large files and minimizing the occurrence of APD events in typical large scale VMware environments. In addition to changes in 4.0.4, testing has shown that ZFS record sizes of 32KB and above help in mitigating instances of APD events when doing parallel deletes of large files.

Nexenta engineering is actively working on this area of ZFS and we expect to have additional scalability improvements in upcoming Fix Packs. See NEX-3890 under [Known Issues](#) for more details.

- **Auto-Sync reliability and functionality extensions**

NexentaStor 4.0.4 improves Auto-Sync usability and reliability. In addition to addressing a number of customer reported issues, Auto-Sync in 4.0.4 contains a number of functionality enhancements exposed in NMV. Specifically, NexentaStor 4.0.4 adds the ability to change snapshot ownership, formally allowing the creation of clones from auto-sync snapshots. It also enhances failback capabilities, removing the need to do a full resync before failing back to the primary site. See NMV and the Auto-Sync user guide for more details.

- **Support for Emulex LPe16002 16Gbps FC HBAs**

NexentaStor 4.0.4 has been certified with 16Gbps Fibre Channel using Emulex LPe16002 HBAs. Please see the latest Hardware Compatibility List for configurations supporting this HBA.

- **New SNMP Manager available to display SNMP trap information**

With the release of NexentaStor 4.0.4 a new SNMP manager is available in NexentaStor. The SNMP manager can now listen and act on SNMP traps received from external systems. This can be used to automate a NexentaStor appliance and have it react to traps generated by other systems in a data center

- **RSF monitoring of Fibre Channel ports**

NexentaStor 4.0.4 now monitors front-end Fibre Channel ports and will automatically trigger controller fail-over in the event of a concurrent failure of all front-end ports on a controller.

- **Chassis Management updates**

NexentaStor 4.0.4 adds chassis management for the following storage enclosures:

- Dell/Compellent SC280, 84 bay enclosure
- Seagate OneStor SP2584, 84 bay enclosure
- Supermicro SMC-847E2C, 44 bay enclosure
- Supermicro SC216BE2C, 24 bay enclosure

- **New provider for Microsoft Volume Shadow Copy Services (VSS)**

A plug-in provider for integrating NexentaStor 4.0.4 with Microsoft VSS is now available. Contact [Nexenta Support](#) for more details.

- **Support for Campus/Metropolitan clusters**

With the newly released MetroHA support for NexentaStor 4.0.4, the capability to create highly available clusters that stretch over a metropolitan distance has been added to NexentaStor 4.0.

Recommendations

- **NFS settings**

With the release of 4.0.4, the defaults for the maximum number of nfsd and lockd threads have been increased from 16 to 256. This increase will allow NexentaStor to handle larger number of concurrent NFS requests.

As the default nfsd and lockd threads in earlier versions of NexentaStor 4.0x was 16, when you upgrade NexentaStor to 4.0.4 from earlier versions, the older settings of NFS threads will be retained. To improve NFS performance and scalability ensure that you increase the nfsd and lockd thread counts. No changes are required in the case of fresh installation of 4.0.4.

❖ *To change the NFS settings, using NMC:*

1. Log in to NMC.

2. Type:

```
nmc:/$ setup network service nfs-server configure
nmc:/$ Option ? configure
```

```
nmc:/$ Server Version      : 4
nmc:/$ Client Version     : 4
nmc:/$ Concurrent NFSD Servers : 32
```

Increase the number of concurrent NFSD servers.

❖ *To change the NFS settings, using NMV:*

1. Go to **Data Management > Shares > NFS Server > Configure**

2. In the Concurrent NFSD Servers field, change the value of the maximum number of concurrent NFS requests.

- **ESXi settings**

To prevent NFS connection issues with any NexentaStor deployed in VMware ESXi environments, you should review the settings proposed in the VMware KnowledgeBase article '[NFS connectivity issues on NetApp NFS filers on ESXi 5.x/6.0 \(2016122\)](#)' and make appropriate changes on the ESXi server.

VAAI Block Changes

Fix Pack 3 and above for NexentaStor 4.0.3 have all VAAI support disabled by default. This was done to protect customer data while still allowing ESXi environments to run with VAAI enabled at the ESXi level to benefit from other backend storage systems that do provide that functionality.

While NexentaStor 4.0.4 is fully VAAI Block certified, it also ships with VAAI support disabled by default. In order to use VAAI Block on a fresh NexentaStor 4.0.4 system, the user will have to first enable the functionality on NexentaStor.

Upgrading from a previous version of NexentaStor requires some planning when it comes to VAAI support. Removing Atomic Test and Set (ATS) functionality from a storage system disrupts VMFS locking mechanism for datastores that have it enabled and causes loss of access to the datastore. ESXi enables ATS by default on creation with VMFS5.

When planning an upgrade, consider that:

- NexentaStor 4.0.3 versions up to FP2 included shipped with VAAI enabled by default.
- NexentaStor 4.0.3 version with FP3 and later ship with VAAI disabled by default.

Regardless of what version of 4.0.x you are upgrading from, or what its defaults were, the customer's settings will be retained during the upgrade. Meaning, regardless if settings were default or changed by active manipulation of settings by customer; if you had VAAI disabled prior to upgrade it will remain disabled after upgrade; vice versa if setting was enabled.

The following table provides some high level guidance, based on starting NexentaStor version to get to a VMware environment with VAAI fully enabled.

Table 2-1: Steps to enable VAAI

Upgrading from	Procedure
4.0.3 FP2 or earlier version	<ol style="list-style-type: none"> 1. Enable ATS on NexentaStor ¹ 2. Upgrade NexentaStor
4.0.3 FP3 or later version	<ol style="list-style-type: none"> 1. Enable ATS on NexentaStor ¹ 2. Upgrade NexentaStor 3. Re-enable ATS on existing ESXi servers

¹ To change ATS settings on NexentaStor requires a reboot. You can minimize reboots by applying configuration changes before upgrading your system, as they will take effect if the settings are on the system at upgrade. A `vaaictl` script, provided by Nexenta, may be used to enable VAAI features. Download the `vaaictl` script from <https://nexenta.com/products/downloads/nexentastor> in the NexentaStor Utilities section. See section [Enabling ATS on NexentaStor](#) for more details.

The procedures and background information provided here summarize information provided in VMware knowledge base articles. You are recommended to consult the original documents. The following KB articles were reviewed by Nexenta in consultation with VMware:

[1033665](#)

[2037144](#)

[2006858](#)

[2030416](#)

[2094604](#)

In case of any issues with these procedures, you should first open support cases with VMware and then, as needed, with Nexenta as joint support.

Warning about upgrading without preparation

VMware and Nexenta both recommend against an upgrade that changes defaults without first configuring ATS locking for affected datastores to be consistent with the new defaults.

As per VMware KB [2037144](#), datastores configured to use ATS-only locking fail to mount after an upgrade that changes defaults to disable and do not show up in the vSphere client datastore view. In this situation Nexenta recommends that you revert to the previously running snapshot, thereby reverting the change in defaults, performing the preparatory steps outlined below, then returning to the upgrade checkpoint.

If you are running ESXi 6.0 with multi-extent datastores mounted by multiple hosts, you should consult KB [2094604](#) and open a support case with VMware as necessary, given the following notice in that KB:

The combination of one host using ATS-only and another host using SCSI Reserve/Release might result in file system corruption.

This can result from other procedures to disable ATS documented by VMware but not

recommended by Nexenta for this situation, including others from the listed KBs.

Planning for upgrade

You should schedule a maintenance window to prepare for upgrades on ESXi hosts and complete NexentaStor upgrades. As per KB [2030416](#), the datastore must be inactive (guests must either be migrated off the datastore or powered off) before disabling ATS:

- All virtual machines must be migrated off the affected datastore, or powered off, prior to running the below steps.

Identifying VMFS datastore extents and checking ATS status

VMFS datastores use one or more extents. All procedures assume that a datastore uses extents backed exclusively by NexentaStor LUNs. In case of datastores using extents backed by storage from more than one vendor, you should open a support case with VMware to confirm appropriate procedures and identify possible further sources of risk, opening joint support cases with Nexenta and other storage vendors as appropriate.

To enumerate mounted datastore and identify which extents they use, log into the ESXi console and type the following from the ESXi console (ssh into the ESXi host(s), using what is also termed "tech support mode"):

```
~ # esxcli storage vmfs extent list
Volume Name  VMFS UUID          Extent Number Device Name
Partition
-----
-----
ham01-zv01   546fcc0f-d40379dd-5ae5-002590daef96      0
naa.600144f0c140cf6e0000546fca5d0002      1
lrtsex01-ds-01 53b43e1d-d4ab8871-1a8d-002590daef96      0
t10.ATA_____ST1000NM00332D9ZM173_____Z1W11CAL
3
```

To confirm that an extent is backed by a NexentaStor block device, use "esxcli storage core device list -d <device>", as in our example:

```
~ # esxcli storage core device list -d naa.600144f0c140cf6e0000546fca5d0002
naa.600144f0c140cf6e0000546fca5d0002
  Display Name: NEXENTA Fibre Channel Disk
(naa.600144f0c140cf6e0000546fca5d0002)
  Has Settable Display Name: true
  Size: 2097152
  Device Type: Direct-Access
  Multipath Plugin: NMP
  Devfs Path: /vmfs/devices/disks/naa.600144f0c140cf6e0000546fca5d0002
  Vendor: NEXENTA
  Model: COMSTAR
  Revision: 1.0
  SCSI Level: 5
  Is Pseudo: false
  Status: on
```



```

Is RDM Capable: true
Is Local: false
Is Removable: false
Is SSD: false
Is Offline: false
Is Perennially Reserved: false
Queue Full Sample Size: 0
Queue Full Threshold: 0
Thin Provisioning Status: yes
Attached Filters:
VAAI Status: unknown
Other UIDs:
vml.0200010000600144f0c140cf6e0000546fca5d0002434f4d535441
Is Local SAS Device: false
Is USB: false
Is Boot USB Device: false
No of outstanding IOs with competing worlds: 32

```

Devices exported from NexentaStor are evident because the Vendor field is set to NEXENTA. For each mounted datastore using NexentaStor-exported extents, use "vmkfstools -Phv1 /vmfs/volumes/<datastore>" to confirm that ATS is enabled, as in our example:

```

~ # vmkfstools -Phv1 /vmfs/volumes/ham01-zv01
VMFS-5.60 file system spanning 1 partitions.
File system label (if any): ham01-zv01
Mode: public ATS-only
Capacity 2 TB, 725.6 GB available, file block size 1 MB, max file size
62.9 TB
Volume Creation Time: Fri Nov 21 23:34:39 2014
Files (max/free): 130000/129619
Ptr Blocks (max/free): 64512/63162
Sub Blocks (max/free): 32000/31911
Secondary Ptr Blocks (max/free): 256/256
File Blocks (overcommit/used/overcommit %): 0/1353841/0
Ptr Blocks (overcommit/used/overcommit %): 0/1350/0
Sub Blocks (overcommit/used/overcommit %): 0/89/0
Volume Metadata size: 814383104
UUID: 546fcc0f-d40379dd-5ae5-002590daef96
Partitions spanned (on "lvm"):
    naa.600144f0c140cf6e0000546fca5d0002:1
Is Native Snapshot Capable: YES
OBJLIB-LIB: ObjLib cleanup done.

```

The "ATS-only" output in the mode line indicates that the datastore is configured to use ATS. If that element is not present in the mode line, the datastore is not configured to use ATS.

Disabling ATS on Existing Datastores

If you need to disable VAAI, Nexenta recommends disabling ATS on a per-device basis, consistent with VMware's recommendation in KB [2006858](#):

Disabling VAAI entirely on the ESXi host may introduce issues in the environment. Instead of disabling VAAI for all devices, you can be disable it only for the affected LUN without impacting other LUNs.

To identify VMFS datastores using NexentaStor storage and check their ATS status, see the section [“Identifying VMFS datastore extents and checking ATS status”](#).

To disable ATS, use `"vmkfstools --configATSONly 0 /vmfs/devices/disks/<extent>"`, as in our example:

```
~ # vmkfstools --configATSONly 0 /vmfs/devices/disks/
naa.600144f0c140cf6e0000546fca5d002:1
```

The command will produce the following output, including a prompt to confirm the change of settings:

```
VMware ESX Question:
VMFS on device naa.600144f0c140cf6e0000546fca5d0002:1 will be upgraded to or
downgraded from ATS capability. Please ensure that the VMFS-5 volume is not in
active use by any local or remote ESX 4.x servers.

Continue with configuration of ATS capability?

0) _Yes
1) _No

Select a number from 0-1: 0

Checking if remote hosts are using this device as a valid file system. This may
take a few seconds...
Downgrading VMFS-5 on 'naa.600144f0c140cf6e0000546fca5d0002:1' from ATS
capability...done
```

In case of any other output, you are recommended to open a support case with VMware, requesting joint support from Nexenta as appropriate.

Once ATS-only mode has been disabled for the datastore, you may proceed with the upgrade, checking guest I/O afterwards. VMware KB [2006858](#) provides a list of symptoms to check in case resulting problems with storage availability are suspected or apparent.

Enabling ATS on Existing Datastores

To identify VMFS datastores using NexentaStor storage and check their ATS status, see the section [“Identifying VMFS datastore extents and checking ATS status”](#).

To enable ATS, use `"vmkfstools --configATSONly 1 /vmfs/devices/disks/<extent>"`, as in our example:

```
~ # vmkfstools --configATSONly 1 /vmfs/devices/disks/
naa.600144f0c140cf6e0000546fca5d002:1
```

The command will produce the following output, including a prompt to confirm the change of settings:

```
VMware ESX Question:
VMFS on device naa.600144f0c140cf6e0000546fca5d0002:1 will be upgraded to or
downgraded from ATS capability. Please ensure that the VMFS-5 volume is not in
active use by any local or remote ESX 4.x servers.

Continue with configuration of ATS capability?
```

- 0) _Yes
- 1) _No

Select a number from 0-1: 0

```
Checking if remote hosts are using this device as a valid file system.
This may take a few seconds...
Downgrading VMFS-5 on 'naa.600144f0c140cf6e0000546fca5d0002:1' from
ATS capability...done
```

In case of any other output, you are recommended to open a support case with VMware, requesting joint support from Nexenta as appropriate.

Once ATS-only mode has been enabled for the datastore, you may proceed with the upgrade, checking guest I/O afterwards. VMware KB [2006858](#) provides a list of symptoms to check in case resulting problems with storage availability are suspected or apparent.

Enabling ATS on NexentaStor

To identify VMFS datastores using NexentaStor storage and check their ATS status, see the section [Identifying VMFS datastore extents and checking ATS status](#).

To change ATS settings on NexentaStor requires a reboot. You can minimize reboots by applying configuration changes before upgrading your system, as they will take effect if the settings are on the system at upgrade.

The `vaaictl` script may be used to enable VAAI features. Download the `vaaictl` script from <https://nexenta.com/products/downloads/nexentastor> in the NexentaStor Utilities section. Copy the script to admin's home directory, and make sure that it is executable by setting the execution mode of the file:

```
# -bash-4.2$ chmod 555 ./vaaictl
```

To apply changes, ssh into the device as admin and run the following script:

```
# -bash-4.2$ sudo ./vaaictl --enable
```

If you need to disable VAAI thereafter, you can follow the previous steps, changing the `vaaictl` invocation either to restore defaults, allowing the system to follow product defaults again:

```
# -bash-4.2$ sudo ./vaaictl --default
```

or to disable explicitly:

```
# -bash-4.2$ sudo ./vaaictl --disable
```

All settings changes require reboot. You can check current settings using this invocation, which will warn you if the config file has been changed without a system reboot, which means that settings on the running system are uncertain:

```
# -bash-4.2$ ./vaaictl --status
```

Reverting to pre-FP3 checkpoints

Reverting checkpoints will revert defaults. As long as the above preparatory steps have been

successfully completed, datastores with ATS-only locking disabled will continue to be accessible.

IPMI Support

Some storage vendors no longer support traditional SES-based JBOD monitoring. In such cases, you can use IPMI monitoring as a substitute for traditional SES-based monitoring.

Determining the Version of the Appliance

NexentaStor 4.0.4 is the current release of NexentaStor. You can determine the version of the appliance using NMC.

❖ *To determine the version of the appliance using NMC:*

◆ **Type:**

```
nmc:/$ show appliance version
```

System response:

```
NMS version: 40-0-47
NMC version: 40-0-39
NMV version: 40-0-45
Release Date: June 20 2015
Operating System: Nexenta/illumos (version 4.0.4)
```

System Requirements

For system requirements for each environment, refer to the “System Requirements” section in the *NexentaStor Installation Guide*.

List of SMB Supported Client Operating Systems

Network clients can access files on NexentaStor using the Server Message Block (SMB) protocol if NexentaStor can properly authenticate the domain users according to the permissions specified in the domain to which NexentaStor is joined and has an active machine account.

The following table describes the versions of Domain Controllers and client Operating Systems that have been successfully tested to work with NexentaStor.

Table 2-2: SMB Compatibility Matrix

	Windows Server R2 2012	Windows Server 2012	Windows Server R2 2008	Windows Server 2008	Windows Server R2 2003	Windows Server 2003	Workgroup Mode
Windows 2012 R2	X	X	X	X	X	X	X

	Windows Server R2 2012	Windows Server 2012	Windows Server R2 2008	Windows Server 2008	Windows Server R2 2003	Windows Server 2003	Workgroup Mode
Windows 8	X	X	X	X	X	X	X
Windows 2012	X	X	X	X	X	X	X
Windows 2008 R2	X	X	X	X	X	X	X
Windows 7	X	X	X	X	X	X	X
Windows 2008	X	X	X	X	X	X	X
Windows 2003 R2	X	X	X	X	X	X	X
Windows XP	X	X	X	X	X	X	X
Windows 2003	X	X	X	X	X	X	X
Red Hat/CentOS 6.5	X	X	X	X	X	X	X
Ubuntu 12.04 LTS	X	X	X	X	X	X	X
Mac OS X 10.9.2	X	N/S	N/S	N/S	N/S	X	X

N/S represents those OSs that are not tested with the indicated Domain Controller.

Upgrading

To upgrade between minor versions of NexentaStor 4.0.x, see: [Upgrading Minor Versions of NexentaStor 4.0.x](#)

To upgrade from NexentaStor 3.1.x to 4.0.x, see [Upgrading from Latest Version of 3.1.6 to 4.0.4 without an Internet Connection](#), [Upgrading from Latest Version of 3.1.6 to 4.0.4 with an Internet Connection](#), [Upgrading from Latest Version of 3.1.6 to 4.0.4 without an Internet Connection](#)

Upgrading Minor Versions of NexentaStor 4.0.x

You can upgrade from Version 4.0.x release or from 4.0.3 FP releases with a single command.

❖ *To upgrade the appliance to 4.0.4 from 4.0.x or 4.0.3 FPx releases:*

1. Type:

```
nmc:/$ setup appliance upgrade
```

Change in NMV Port

In NexentaStor 4.0.x, the default NMV port has been changed to 8457.

In order to avoid potential conflicts with other networking vendors, Nexenta registered the port 8457 with IANA. For Nexenta releases 4.x and above, web access no longer uses the default port 2000 and uses port 8457. Note when upgrading from release 3.x to 4.x, NMV is now accessed by port 8457.

Upgrading from Version 3.1.x to 4.0.4

Upgrading NexentaStor from version 3.1.x to 4.0.4 is a three-step process. First you need to upgrade to the latest version of 3.1.6 and then upgrade to latest 4.0.3 and then to 4.0.4.

1. Upgrade to the latest version of 3.1.6.
2. Then, upgrade to latest version of 4.0.3.
3. Then, upgrade to 4.0.4

See [Upgrading from Latest Version of 3.1.6 to 4.0.4 with an Internet Connection](#).

During the upgrade, NexentaStor services and volumes remain available to network clients. The upgrade operation requires system restart. Therefore, it is recommended that the upgrade process be scheduled during a system maintenance window. All NexentaStor services and volumes are not available during the restart.

Upgrading from Latest Version of 3.1.6 to 4.0.4 with an Internet Connection

When upgrading if your machine ID changes, visit the respective [Customer Portal](#) or [Partner Portal](#) and provide the following to obtain the new license key.

- Old license key
- Sales order that applies to the old license key
- New machine ID

Also verify that your environment meets the following prerequisites:

- The network interface card is included in the hardware compatibility list for NexentaStor 4.0.x.
- No third-party applications or packages are present on your NexentaStor appliance. You may have third-party packages installed if you changed repository sources on your NexentaStor appliance. The upgrade will result in the loss of components that are not included with the NexentaStor build.

❖ *To upgrade from version 3.1.6.x to 4.0.4, using NMC:*

1. If you have not upgraded to latest version of NexentaStor 3.1.6, upgrade to it by typing:

```
nmc :/$ setup appliance upgrade
```

System response:

```
Cleanup upgrade caches? (y/n)
```

If you choose Yes, prepare to wait for software upgrade to complete a bit longer. Follow the prompts on the screen and answer based on your requirements. This upgrade creates Rollback Checkpoint. You may view the list of all available checkpoints by using 'show appliance checkpoint' command or proceed with 4.0.4 upgrade.

To upgrade to latest 4.0.3, you may run the NMC command `setup nexentastor upgrade -r <release>`. Running this command will automatically disable the multi-NMS and restart NMS.

2. Upgrade to NexentaStor latest 4.0.3 version, by typing:

```
nmc :/$ setup nexentastor upgrade -r 4.0.3
```

This upgrades the system to 4.0.3 FP4.

3. Then upgrade to NexentaStor 4.0.4 by typing:

```
nmc :/$ setup appliance upgrade
```

System response:

```
Proceed to automatically disable multi-NMS and restart NMS?
```

4. Type y.

Multi-NMS is disabled and NMS is restarted.

System response:

The upgrade process may take some time up to 30 seconds to complete.

Do you know if your hardware has been certified for 4.0.x? (y/n)

5. Type y if your hardware is listed in the Hardware Certification List (HCL).

Upgrade NexentaStor Appliance from version 3 to version 4.

This process include upgrade kernel, drivers, base system and appliance.

WARNING: We can't guarantee third-party software will continue to work properly after upgrade.

WARNING: The system should be restarted at the end of the process.

Proceed? (y/n)

6. Type y.

System response:

NexentaStor is upgrading.

During the upgrade, do not switch off or restart the NexentaStor appliance.

7. NexentaStor notifies you about the upgrade process.

The first phase of upgrade has completed successfully

Reboot now to finish upgrade to 4.0?

8. Continue to use NexentaStor 3.1.6 or reboot to activate NexentaStor 4.0.

Nexenta does not recommend continuing to work using NexentaStor 3.1.6 after the first stage of the upgrade is completed. You may postpone the restart if you have incomplete archiving tasks. Otherwise, proceed with the reboot. When rebooting, all NexentaStor services and datasets are unavailable for network clients.

9. Verify that syspool is mounted:

1. In NMV, click **Settings > Appliance**.

2. In the **Upgrade Checkpoints** pane, click **View**.

You should see the list of upgrade checkpoints.

Warning: After you upgrade the volume version, back up your system. Backups created for mirrored pools with earlier volume versions may not be available after the upgrade.

10. Optionally, upgrade NexentaStor volumes to version 28 by typing:

```
nmc:/$ setup volume <volname> version-upgrade
```

11. Repeat [Step 10](#) for all NexentaStor volumes

12. After seamless upgrade from 3.x to 4.x, `nfsmapid_domain` setting is not maintained, and must be reset manually.

To reset manually, SSH to the system

Set the `nfsmapid_domain`

To set the `nfsmapid_domain`, log in to bash:

```
nmc:/$ option expert_mode =1
```



```
nmc:/$ !bash
```

Type:

```
# sharectl set -p nfsmapid_domain=<domain> nfs
```

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

Upgrading from Latest Version of 3.1.6 to 4.0.4 without an Internet Connection

Before you upgrade your appliance without an Internet connection, review [Upgrading from Latest Version of 3.1.6 to 4.0.4 with an Internet Connection](#). Then verify that your environment meets all prerequisites described in this section.

❖ *To upgrade from latest version of 3.1.6 to 4.0.4 with no Internet connection:*

1. If you are unable to connect to the Internet to upgrade your system, contact support@nexenta.com for the ISO image.
2. Mount or burn the ISO image.

Complete [Step 2](#) to [Step 11](#) in [Upgrading from Latest Version of 3.1.6 to 4.0.4 with an Internet Connection](#).

Upgrading to Version 4.0.4 After Rolling Back to Latest Version of 3.1.6

Generally, Nexenta does not recommend that you roll back a NexentaStor appliance to version 3.1.6 after the upgrade to 4.0.4 on a production system. If you upgrade the volume version during the upgrade to version 4.0.4, the data and system volumes will be unavailable in version 3.1.6, since volume version 28 is not supported in version 3.1.6.

Rollback and upgrade is somewhat acceptable on a testing system.

During the upgrade, NexentaStor creates a flag file `/volumes/.config/.3_to_4_upgrade`. If you try to run the upgrade after rolling back to version 3.1.6, the upgrade fails.

To re-run the upgrade to version 4.0.4, delete the `/volumes/.config/.3_to_4_upgrade` file and run the `setup nexentastor upgrade` command again.

❖ *To rerun the upgrade from version 3.1.6 to 4.0.4, using NMC:*

1. Log in to bash:

```
nmc:/$ option expert_mode =1
```

```
nmc:/$ !bash
```

2. Type:

```
# rm /.config/.3_to_4_upgrade
```

3. Exit bash by typing:

```
# exit
```

4. Run:

```
nmc:/$ setup nexentastor upgrade -r <release>
```

Example:

```
nmc:/$ setup nexentastor upgrade -r 4.0.3
```

5. Then upgrade to NexentaStor 4.0.4 by typing:

```
nmc:/$ setup appliance upgrade
```

Enhancements

This section lists enhancements in NexentaStor 4.0.4.

Key	Description
NEX-1099	Changed behavior for clearing the faulted runners list to always require administrative intervention.
NEX-1160	Provided appliance management of NIS client for use with NFS.
NEX-1955	Changed the default value of the ses-check runner ival_anti_flapping setting to address possible false values being reported by JBOD sensors.
NEX-2510	Enhanced NMC Kerberos client support for principal creation and keytab population interoperating with MIT KDCs.
NEX-2694	Updated pool creation logic to prevent using two disks from the same tray.
NEX-2788	Restored 3.x functionality for changing snapshot ownership for Auto-services.
NEX-2849	Incorporated Emulex driver updates.
NEX-2865	Added functionality to Auto-sync to accommodate multiple Auto-snap schedules at once.
NEX-2951	Restored 3.x functionality for ACL replication by auto-tier for Microsoft Active Directory.
NEX-3082	Added FC port status monitoring to RSF.
NEX-3134	Enhanced functionality of the SNMP traps.
NEX-3145	Added a user confirmation response prior to deleting Auto-sync snapshot sources.
NEX-3663	Ported Illumos kmem_reap fixes #5376, #5498, and #5514.

Resolved Common Vulnerabilities and Exposures (CVE)

The following CVEs have been incorporated into NexentaStor 4.0.4.

CVE	Description
(CVE-2015-1798) (CVE-2015-1799) (CVE-2015-2781)	The symmetric-key feature in the receive function in ntp_proto.c in ntpd in NTP 4.x before 4.2.8p2 requires a correct MAC only if the MAC field has a nonzero length, which makes it easier for man-in-the-middle attackers to spoof packets by omitting the MAC.
(CVE-2003-1418)	The mod_cgid module in the Apache HTTP Server before 2.4.10 does not have a timeout mechanism, which allows remote attackers to cause a denial of service (process hang) via a request to a CGI script that does not read from its stdin file descriptor.
(CVE-2014-6273)	Buffer overflow in the HTTP transport code in apt-get in APT 1.0.1 and earlier allows man-in-the-middle attackers to cause a denial of service (crash) or possibly execute arbitrary code via a crafted URL.
(CVE-2013-5704)	The mod_headers module in the Apache HTTP Server 2.2.22 allows remote attackers to bypass "RequestHeader unset" directives by placing a header in the trailer portion of data sent with chunked transfer coding. NOTE: the vendor states "this is not a security issue in httpd as such."
(CVE-2014-0231)	The mod_cgid module in the Apache HTTP Server before 2.4.10 does not have a timeout mechanism, which allows remote attackers to cause a denial of service (process hang) via a request to a CGI script that does not read from its stdin file descriptor.
(CVE-2014-0118)	The deflate_in_filter function in mod_deflate.c in the mod_deflate module in the Apache HTTP Server before 2.4.10, when request body decompression is enabled, allows remote attackers to cause a denial of service (resource consumption) via crafted request data that decompresses to a much larger size.
(CVE-2013-2566)	The RC4 algorithm, as used in the TLS protocol and SSL protocol, has many single-byte biases, which makes it easier for remote attackers to conduct plaintext-recovery attacks via statistical analysis of ciphertext in a large number of sessions that use the same plaintext.

Resolved Issues

This section describes the issues fixed in NexentaStor 4.0.4.

Functional Area	Key	Description
Appliance Management	NEX-2617	Resolved issue where configuring LDAP could lead to a DNS resolution hang and a possible fault.
Appliance Management	NEX-3596	Restored 3.x functionality to allow setting the nfsmapid_domain via NMS.
Autosync	NEX-2385	Resolved issue where the Auto-sync service could timeout under certain conditions.
AutoSync	NEX-3167	Resolved issue with changing back auto-sync direction after running reverse-service.
AutoSync	NEX-3420	Resolved condition where "Use of uninitialized value \$fmri" error is logged each time nmwatchdog starts.
AutoSync	NEX-3512	Addressed case where rotated Auto-services logs would not refresh in NMV.
AutoSync	NEX-3531	Resolved condition where "Destroy Any Snapshots" did not properly comply with the retention policy settings.
AutoSync	NEX-3573	Fixes for 3.x, autosync snapshot does not expire after seamless upgrade.
AutoSync	NEX-3656	Corrected the functionality of Auto-sync flip-direction mode. See Auto-sync User Guide or What is New in this Release? to learn about the options available for flip-directions.
Autosync	NEX-3009	Resolved issue where nmwatchdog could incorrectly report that snapshots were not replicated.
CIFS	NEX-3430	Resolved issue where backups via SMB mount could fail when using Symantec Backup Exec.
CIFS	NEX-3576	Resolved RPC error 1726 when displaying open files via Microsoft Management Console Computer Management snap-in.
Commands	NEX-3203	Resolved inconsistency in NTP configuration between a fresh install and a seamless upgrade from 3.1.x.
COMSTAR	NEX-2723	Enhanced handling of SCSI write commands when used with VAAI extensions.
COMSTAR	NEX-2787	Resolved SCSI target issue resulting in unresponsiveness to ESXi clients.
COMSTAR	NEX-2862	Resolved SCSI Target issue where issuing many UNMAP operations could cause system unresponsiveness.
Daemons	NEX-2893	Addressed issue where a stale volume configuration cache could potentially lead to degraded status.

Functional Area	Key	Description
Fibre Channel	NEX-3295	Resolved condition where cloning vmdk files larger than 450Gb via SVMotion could occasionally fail.
HA	NEX-2508	Addressed RSF memory fault caused by invalid cache.
HA	NEX-3049	Resolved inconsistency between cluster virtual IP addressing.
HA	NEX-3085	Improved error handling when exporting service under cluster control.
HA	NEX-3161	Corrected cluster service handling of uppercase letters in node names.
HA	NEX-3391	Resolved condition where stmfha backup-globals was not synchronizing between nodes.
HA	NEX-3405	Resolved condition during a failover where delays could be caused by RSF performing DNS lookups for non-existent hostnames.
HA	SUP-949	Resolved condition where adding a pool to cluster control would restart services on existing pools, leading to a brief outage.
Kernel	NEX-1823	Incorporated Illumos #5911 to minimize the impact of I/Os during large deletes. See NEX-3890 under Known Issues for further details.
Kernel	NEX-1825	Resolved SCSI target condition preventing LUN discovery for Windows initiators using standby ALUA path.
Kernel	NEX-941	Resolved issue where ZFS was failing to replace an unavailable drive with a hot spare.
Kernel	SUP-930	Resolved a panic caused by freed and reused buffers.
Kernel, chassis management	NEX-2591	Resolved issue where FRU field for drives was not being appropriately updated.
Kernel, Protocols	NEX-2936	Incorporated fix for Illumos #4950 to address situation where directories and/or files could not be deleted from the FS once refquota set on the FS was exceeded.
NFS	NEX-2529	Resolved condition where the NFS service may enter maintenance mode after changing the hostname or during a volume import if the hostname had been previously changed.
NFS	NEX-3019	Address NFSv3 specification issue by forbidding writes across exported folder filesystem boundaries.
NFS	NEX-3095	Fixed leaks and other minor issues in the NFS server code.
NFS	NEX-3505	Resolved NFS Authentication errors with certain anonymous access settings.
NFS	NEX-3758	Enhanced detection and logging of NFSv3 stale locks.
NMS	NEX-3018	Added functionality for configuring an IP alias on top of a VLAN interface.

Functional Area	Key	Description
NMS	NEX-607	Resolved security issue where permissions for user home directory were set improperly by default.
NMV	NEX-1670	Enhanced generation of NMV key pair.
NMV	NEX-2713	Corrected error message for NMV password change.
Protocols	NEX-2464	Resolved defect in NMS which caused an error while joining Microsoft Active Directory on the inactive node of a cluster.
Seamless Upgrade	NEX-2609	Resolved situation where connectivity issues could be caused after an upgrade from 3.x to 4.x on clients without anonymous read/write enabled.
Seamless Upgrade	NEX-915	Changed defaults to disable SMART.
Security	NEX-3084	Resolved condition when inability to communicate with NTP server could trigger nms-check to go to maintenance mode.
SMB	NEX-2525	Resolved Microsoft Management Console Share and Storage Management snap-in error when setting share-level security using Windows Server 2008 R2.
ZFS	NEX-3209	Corrected situation where folders were not properly recognizing upper-case letters with certain normalization settings.
ZFS	NEX-3485	Addressed condition where deferred deletes could lead to loss of NFS access on cluster failover.

Known Issues

This section lists the known issues in NexentaStor 4.0.4 as of July 2015.

Table 2-3: Known Issues in 4.0.4

Functional Area	Key	Description	Workaround
CIFS	NEX-1999	Potential, but highly infrequent panic from SMB session management. This panic will be fixed in upcoming release	No known workaround
Commands + Daemons	NEX-2941	Editing share level ACL through MMC intermittently fails after a cluster fail-over.	Execute the following command via bash: svcadm restart network/smb/server
COMSTAR	NEX-2971	Failure to clean up I/Os can cause intermittent FC link resets.	No known workaround
NFS	NEX-3648	Manual failovers hang and cause loss of communication with FC LUNs configured in an ESXi 6.0 cluster.	This issue is not unique to NexentaStor and has been posted to VMware. While we work this issue with VMware the Customer is recommended to use ESXi 5.5 Server for clustered environments using FC LUNs.
HA	NEX-3191	Export failure on failover in clusters with large number of nfs mounts and auto-sync jobs.	If an automatic failover times out, manually initiate the failover from NMV or NMC.
HA	NEX-3394	Issues with cluster failover after upgrading a system using PGR3 Reservations to a later release using SCSI-2 Reservations.	Configurations using STEC SAS SSD's 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 SSD's as cache or log devices are not affected by this restriction.
HA	NEX-3504	Intermittent issue where systems are unable to import full volumes due to certain sub-shares being created prior to their mount.	Retry the operation if this condition is encountered.

Table 2-3: Known Issues in 4.0.4

Functional Area	Key	Description	Workaround
HA	NEX-3769	Systems with 2 or more pools with similar names can be susceptible to unintended LU deletions upon failover.	There are two possible workarounds. 1. Edit /opt/HAC/RSF-1/bin/purge-stmf.sh replace this line: volume_pattern=`echo "/rdsk/\$1" sed 's#/#\\V/#g` with this: volume_pattern=`echo "/rdsk/\$1/" sed 's#/#\\V/#g` 2. Make sure that service names don't allow for undesired pattern matching. Having services named mypool and mypool2 will cause problems, but combination mypool1 and mypool2 is fine.
Installation	NEX-1881	Under certain circumstances, NexentaStor clusters can have mismatched controller numbers between the nodes.	Contact system installer or support provider 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.
Kernel	NEX-2940	Disk pools with a failed sTEC drive as a single ZIL can cause a system panic when users attempt to remove the failed ZIL.	Use redundantly configured (mirrored) ZILs.
Kernel	NEX-2966	On busy NFS servers NLM does not always release deleted files causing what looks like a space leak.	Execute the following command via bash: svcadm restart nlockmgr
Kernel	NEX-3043	Alternating I/Os to datasets of different record sizes can cause long zio_cache reaps.	No known workaround
Kernel	NEX-3585	Intermittent issue where VM slack in non-ARC ZFS kmem caches can degrade ARC performance.	No known workaround
Kernel	NEX-3717	Toshiba THNSNJ96 SATA SSD is not recognized upon hotplug.	Avoid hotplugging when replacing Toshiba THNSNJ96 SATA SSDs.

Table 2-3: Known Issues in 4.0.4

Functional Area	Key	Description	Workaround
Kernel	NEX-3734	ZFS allows the user to set a duplicate mountpoint path on two different ZFS filesystems, leading to broken volume services.	Check pool for duplicate mountpoints before failover, perform manual remediation.
Kernel	NEX-3890	Users performing several parallel deletions of very large files built with small recordsize can see performance degradation during the delete.	1) Serialize delete operations 2) Use a larger recordsize (for example, 128k)
Kernel	NEX-928	When using ZEUS IOPS drives in a JBOD, a 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, ZFS	NEX-1760	ZFS exhibits long kmem reap times in certain situations	No known workaround
NMC	NEX-3969	Upon upgrade, systems with a time/date set incorrectly can boot to the incorrect checkpoint.	Before starting an install or upgrade, ensure that the system time/date are set correctly. If this issue is encountered, reboot the system to the correct checkpoint.
NMS	SUP-737	NMV may over time grow heap memory while failing to reclaim allocations.	Restart NMS if large amounts of memory are being used.
NMV	NEX-2782	On very large configurations, NMV "Create New Volume" may not show all profiles and available drives.	On systems exhibiting this behavior, manually create volumes without using metis.
NMV	NEX-4198	After upgrading to NexentaStor 4.0.4 there are instances when NMV pages do not display information related to HAC, iSCSI mapping.	Clear any cached files for the NexentaStor RSF GUI. Details for clearing the cached files for each browser can be found in Wikipedia knowledge base article. Once you clear the cache file force reload the NMV page. https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache

Table 2-3: Known Issues in 4.0.4

Functional Area	Key	Description	Workaround
NMV	NEX-3511	Users may encounter an ACLCollector.pm error when editing CIFS settings on folders.	<p>You can use MMC command console from Windows to add/modify/delete ACL settings on CIFS shares.</p> <ol style="list-style-type: none"> 1) Make sure to start MMC as Administrator on the windows client. 2) Ensure that you have an ACL in place on the root of the share, that has sufficient rights to perform ACL modifications. 3) Launch MMC from windows (%windir%\system32\compmgmt.msc /s) 4) Right click on "Computer Management" at top of window and select "Another computer". Enter IP address or FQDN in box. Alternatively, you can click on browse and select from AD server, which system to manage. 5) Launch application and click on "System Tools" and then click on "Shared Folders" and then "Shares". 6) Select the share to modify and right click and select properties and select top tab "Security". 7) click on "edit" to add/remove/modify users/groups <p>If you get any errors, then check that you have correct permissions to view/modify security settings.</p>

Table 2-3: Known Issues in 4.0.4

Functional Area	Key	Description	Workaround
Protocols	NEX-3941	Users performing a query with a PowerShell script, running through MS Task Scheduler, can encounter a script failure with STATUS_BUFFER_OVERFLOW error.	<p>Set NStor SMB server to SMB1:</p> <pre>sharemgr set -p smb2_enable=false smb svcadm restart smb/server</pre> <p>or</p> <pre>svccfg -s smb/server setprop smbd/smb2_enable=false svcadm refresh smb/server svcadm restart smb/server</pre> <p>or on MS client disable SMBv2</p> <p>Disables the SMBv2 and SMBv3 on the MS client:</p> <pre>Set-SmbServerConfiguration -EnableSMB2Protocol \$false reboot client</pre> <p>to re-enable SMBv2 and SMBv3 on the client:</p> <pre>Set-SmbServerConfiguration -EnableSMB2Protocol \$true reboot client</pre>

Table 2-3: Known Issues in 4.0.4

Functional Area	Key	Description	Workaround
Protocols	NEX-4053	Customers issuing a dir filename command on a Windows client may encounter a hang when smb2 is enabled on the Nexenta Server.	<p>Use the following steps to disable SMB2:</p> <p>Set NStor SMB server to SMB1:</p> <pre>sharemgr set -p smb2_enable=false smb svcadm restart smb/server</pre> <p>or</p> <pre>svccfg -s smb/server setprop smbd/smb2_enable=false svcadm refresh smb/server svcadm restart smb/server</pre> <p>or on MS client disable SMBv2</p> <p>Disables the SMBv2 and SMBv3 on the MS client:</p> <pre>Set-SmbServerConfiguration -EnableSMB2Protocol \$false reboot client</pre> <p>to re-enable SMBv2 and SMBv3 on the client:</p> <pre>Set-SmbServerConfiguration -EnableSMB2Protocol \$true reboot client</pre>
Seamless Upgrade	NEX-3606	After seamless upgrade from 3.x to 4.x, nfsmapid_domain setting is not maintained, and must be reset manually.	<p>SSH to the system and run the following command to set the nfsmapid_domain:</p> <pre>sharectl set -p nfsmapid_domain=<domain> nfs</pre>

Global Headquarters

451 El Camino Real, Suite 201
Santa Clara, California 95050
USA

7000-nxs-4.0.4-000017-B