NexentaStor

Release Notes 3.1.6 FP4

Date: August 4, 2015

Subject: NexentaStor Release Notes 3.1.6 FP4

Software: NexentaStor

Software Version: 3.1.6 FP4

Part Number: 7000-nxs-3.1.6-FP4-000010-D

www.nexenta.com

This page intentionally left blank

Contents

What is New in this Release?

NexentaStor 3.1.6 FP4

3.1.6 FP4 provides better control over the upgrade process from 3.1.x to 4.0.x. You must upgrade first to the latest version of 3.1.6 as the first step in any 4.0.x upgrade from 3.1.x. The following describes the process to upgrade to 4.0.4:

1. Upgrade to 3.1.6 FP4
2. Then, upgrade to latest version of 4.0.3
3. Then, upgrade to 4.0.4

In some rare instances, your machine ID may change when upgrading. If this happens 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. Refer to NexentaStor 4.0.4 release notes.

NexentaStor 3.1.6 FP3

The following changes are included in NexentaStor 3.1.6 FP3.

VAAI Block Changes

To avoid kernel panics that may interrupt services, VAAI block feature is removed in NexentaStor 3.1.6 FP3. NexentaStor 3.1.6 FP3 no longer supports the following SCSI commands.

- compare_and _write (0x89)
- copy (0x83 and 0x84)

- write_same (0x41/0x93)

The procedures and background information provided here summarize information provided in VMware knowledge base articles. Customers 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, customers 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 without first disabling VMware VMFS ATS-only locking.

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

Customers running ESXi 6.0 with multi-extent datastores mounted by multiple hosts 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 TS documented by VMware but not recommended by Nexenta for this situation, including others from the listed KBs.

Planning for upgrade

Customers should schedule a maintenance window to prepare for upgrades on ESXi hosts and complete NexentaStor upgrades. VMware assesses the impact thus, as per KB 2006858 (italics added):

- Disabling ATS locking takes effect immediately and does not require a VMware host reboot. *But if there is a host that has the volume mounted successfully, you may need to restart all the hosts in order to clear the ATS-Only.* This is because the host accessing with ATS-Only is preventing a Non-ATS-Only host accessing it.
- This change must be implemented on each host individually. All hosts able to connect to this storage must be configured consistently in regards to ATS.

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.

Preparing ESXi hosts for upgrade

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 disable it only for the affected LUN without impacting other LUNs.

VMFS datastores use one or more extents. ATS settings must be modified on the devices underlying these extents.

To enumerate mounted datastores 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
lrtsesx01-ds-01  53b43e1d-d4ab8871-1a8d-002590daef96              0
t10.ATA_____ST1000NM00332D9ZM173_____________________________Z1W
11CAL           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 below:

```
~ # 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 NextaStor-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.

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, customers 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.

SMB Changes

To see only the files and directories to which you have access, select Access Based Enumeration (ABE) in the CIFS share option. You may enable ABE to filter large directories or to hide the inaccessible files/folders and to share the SMB property.

NTP Vulnerabilities

Updated the NTP version to 4.2.8 which fixes multiple NTP vulnerabilities including remote attackers from defeating cryptographic protection mechanisms.

NexentaStor 3.1.6 FP2

NexentaStor 3.1.6 FP2 fixed issues clustered around Seamless Upgrade.

NexentaStor 3.1.6 FP1

NexentaStor 3.1.6 FP1 fixed the Shellshock Bash vulnerabilities that affect all software that uses the Bash shell and parses values of environment variables.

NexentaStor 3.1.6

The following new features and enhancements are included in the NexentaStor release 3.1.6:

- Improved handling of intermittently faulty devices
- Support for Microsoft Server 2012 Cluster™ and Cluster Shared Volumes (CSV)™
- Reduced HA failover times
- Changed the default settings for SCSI reservations

SCSI Reservation Changes

To ensure reliability of NexentaStor deployments, the default SCSI reservation setting for HA Cluster has been changed. In NexentaStor 3.1.6 the default setting is SCSI-2.

In NexentaStor 3.1.5 the default setting for SCSI reservation is PGR-3. After the upgrade to version 3.1.6, the setting will automatically change to SCSI-2 reservations.

For new deployments, SCSI-2 reservation is assigned automatically.

System Requirements

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

Upgrading

Warning:	If you use Cisco UCS C240 M3 server with LSI-9271 RAID controller, we do not recommend that you upgrade to 3.1.6. The upgrade will result in storage devices going offline after the upgrade.

Upgrading with an Internet Connection

Applicable Versions: 3.1.0, 3.1.1, 3.1.2, 3.1.3, 3.1.3.5, 3.1.4, 3.1.4.1, 3.1.4.2, 3.1.5, 3.1.6, 3.1.6 FP1, 3.1.6 FP2 or 3.1.6 FP3

To upgrade from 3.1.0, 3.1.1, 3.1.2, 3.1.3, 3.1.3.5, 3.1.4, 3.1.4.1, 3.1.4.2, 3.1.5, 3.1.6, 3.1.6 FP1, 3.1.6 FP2 or 3.1.6 FP3, use the NMC command `setup appliance upgrade`.

Note:	Upgrading from 3.1.3.5 will not affect the IDMU capabilities.

Upgrading without an Internet connection

Applicable Versions: 3.0.5, 3.1.1, 3.1.2, 3.1.3 or 3.1.3.5, 3.1.4, 3.1.4.1, 3.1.4.2, 3.1.5, 3.1.6, 3.1.6 FP1, 3.1.6 FP2, 3.1.6 FP3

If you are unable to connect to the Internet to upgrade your system, contact support@nexenta.com.

Upgrading from 3.0.5 and earlier with an Internet Connection

❖ *To upgrade the appliance to 3.1.6 FP4 and ensure all updates are installed:*

1. Type:

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

Note:	For HA Cluster configurations, it is vital to install the mapmgr patch prior to upgrading to NexentaStor 3.1.6.

For more information, refer to the 060711 Technical Bulletin and Field Information Notice 2011-02 available in the Self-Service Portal.

Resolved Issues

NexentaStor 3.1.6 FP3

Table 3-1: Resolved issues in NexentaStor 3.1.6 FP3

Key	Description	Component(s)
NEX-3077	Added Access-Based Enumeration (ABE) to CIFS share options.	Appliance Mgmt
NEX-3013	Removed VAAI in the appliance.	Comstar
NEX-3074	Fixed seamless upgrade failures when upgrading from 3.1.6-FP3 to later releases.	NMC, Seamless Upgrade
NEX-3094	Fixed upgrade failures when upgrading from 3.1.3.5 to 3.1.6 FP3.	Upgrade
NEX-3006	Fixed multiple NTP vulnerabilities. Backported NTP fixes for CVE-2014-9293 in NexentaStor 3.1.x.	Packaging
NEX-3087	Improved retention of diagnostic information by creating individual cores based on unique program binary names. Prior to 3.1.6 FP3, If multiple cores occur on the same program, binary names for older core will be overwritten.	Packaging
SUP-942	Resolved issues when importing a zpool on a clustered system that is sharing an filesystem using CIFS that is a member of an AD domain.	Protocols

NexentaStor 3.1.6 FP2

Table 3-2: Resolved issues in NexentaStor 3.1.6 FP2

Key	Description	Component(s)
NEX-2202	Fixed the issue with the password in NMV in the SNMP AGENT: CONFIGURATION page.	Appliance Management

Key	Description	Component(s)
NEX-2764	Fixed NMS to not use `cfgadm unconfigure` on any controllers of type fc-fabric while rescanning HBAs, which causes delay in setups with more than 1000 targets.	NMS
NEX-2203	Fixed the issue with the password in the NDMP SERVER CONFIGURATION page in NMV.	NMV
NEX-1492	Resolved a condition that may occur during seamless upgrade where default user passwords may revert back to the system defaults.	Seamless Upgrade
NEX-1493	Resolved a condition where custom user information may not persist after seamless upgrade completes.	Seamless Upgrade
NEX-2638	Resolved conditions where checkpoints may not have persisted after seamless upgrade.	Seamless Upgrade
NEX-2805	Multi-NMS will be disabled automatically for seamless upgrade while running the nmc command `setup nexentastor upgrade`.	Seamless Upgrade
NEX-2822	Resolved issue where custom users created in NMV were unable to login after seamless upgrade.	Seamless Upgrade
NEX-2405	Fixed the NMS properties to 4x defaults before saving appliance configuration during seamless upgrade to prevent restoring 3x defaults in upgraded 4x version.	Seamless Upgrade

NexentaStor 3.1.6 FP1

Table 3-3: Resolved Issues in NexentaStor 3.1.6 FP1

Key	Description	Functional Area
NEX-2658 NEX-2635	Security Update to address vulnerability CVE-2014-6271. This vulnerability CVE-2014-6271 could allow for arbitrary code execution.	Packaging
NEX-2657 NEX-2642	Security Update to address vulnerability CVE-2014-7169. This vulnerability CVE-CVE-2014-7169 involved bash allowing code execution via specially-crafted environment.	Packaging

NexentaStor 3.1.6 Known Issues as of February 2015

Changes to Auto-Sync using NMV may not take effect

Description: 13278

In NMV, when attempting to change Auto-Sync options, the changes do not take effect.

In some conditions, HA systems failover does not occur when running Auto-Sync

Description: 13366, NEX-315

In some conditions, HA systems when running Auto-Sync, may not failover as expected.

Workaround:

Remove Auto-Sync from each node, and mark the pool as repaired.

In some conditions, the Confguard plugin does not send Email notifications for changes

Description: 13282

When using the NexentaStor Confguard plugin, which continuously monitors system configuration, in some conditions, it does not send email alerts.

Workaround:

In NMC, type:

```
nmc:/$ setup confguard report send
```

mptsas kstat namespace collision during Installer boot

Description: NEX-2276

Due to certain timing and specific sequence of operation in the NexentaStor Installer, you may see kernel statistics warning messages when booting from the NexentaStor ISO.

Example:

```
WARNING: kstat_create('unix', 1, 'mpt_sas_nexus_enum_tq'): namespace
collision

WARNING: kstat_create('mpt_sas', 1, 'fm'): namespace collision
```

Workaround:

The kernel statistics error messages do not affect the installation and system operations. You can safely ignore these messages.

Upgrades from 3.1.3 to 3.1.6 displays error "mandb: cannot open /usr/share/man/man8/pwconv.8"

Description: 11556

When upgrading from 3.1.3 to 3.1.6, errors occurs during the upgrade sequence referencing ‘`mandb: can't open /usr/share/man/man8/pwconv.8`’

The errors are benign and may be safely ignored. These errors are resolved later in the upgrade sequence.

Upgrade path: Unable to remove mapping entry from zvol 'tank1/zvol_1'

Description: NEX-2208, NEX-2270

This bug affects HA-Cluster deployments only.

You may have issues with `mapmgr` during the upgrade to 3.1.6:

- During the upgrade process symlink `/usr/bin/mapmgr -> /opt/HAC/RSF-1/bin/mapmgr` is replaced with the binary file `/usr/bin/mapmgr`. This results in `nms-comstar` and `rsfmon` using a different version of mapmgr, which causes RPC issues with block target operations
- During the upgrade of the NexentaStor HA Cluster nodes from version 3.1.5 to 3.1.6, the shared volume gets exported. After the upgrade you may be unable to add the shared volume back under cluster control.

Workaround:

To fix this issue you must move the mapmgr binary file to the old location and replace the binary file that the upgrade adds to `/opt/HAC/RSF-1/bin/mapmgr` with a symlink.

❖ *To fix the mapmgr issue:*

2. Log in to bash:

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

3. Move the mapmgr file:

```
# mv /usr/bin/mapmgr /usr/bin/mapmgr.old
```

4. Replace the binary file with symlink:

```
# ln -s /opt/HAC/RSF-1/bin/mapmgr /usr/bin/mapmgr
```

sTec © ZeusRAM™ SAS SSD recommended upgrade of the Firmware to version C023

Description:

sTec ZeusRAM™ SAS SSD firmware version C023 has been certified by Nexenta on 3.1.6 and it is strongly recommended to upgrade to this firmware version.

Upgrade:

Contact sTec for details on how to obtain and upgrade this firmware.

Active Directory member IDs may persist when leaving a domain

Description: 13310

After leaving an Active Directory Domain NexentaStor may retain the ID's of members of the domain.

Workaround:

- *To remove the IDs:*

1. Open NMC. At the NMC prompt, type:

```
# option expert_mode=1
# !bash
```

2. List the local group members using the following command in bash

```
# smbadm show -m
```

3. After identifying any group members associated with the removed domain, use the following command to remove:

```
# smbadm remove-member -m MEMBER GROUP
```

The "setup configuration save" and "setup configuration restore" commands do not save and restore iSCSI settings

Description: NEX-1956

The `setup appliance configuration save` and `setup appliance configuration restore` commands do not save/restore iSCSI settings, such as targets, mappings, etc. Only iSCSI Target Portal Group settings are restored.

Workaround:

- *To restore iSCSI configuration, using NMC:*

1. Disable multi-NMS:

```
nmc:/$ setup appliance nms property srvpool_cnt_initial -p 0
```

2. Restart NMS:

```
nmc:/$ setup appliance nms restart
```

3. Restore appliance configuration:

```
nmc:/$ setup appliance configuration restore all
```

4. Restart NMS:

```
nmc:/$ setup appliance nms restart
```

NMV issue when trying to list iSCSI initiator with IQN containing non-ASCII characters

Description: 12866

NMV issue when trying to list iSCSI initiator groups immediately after successfully creating IQN with non-ASCII characters.

Workaround:

Use ASCII characters for the iSCSI initiator IQN.

Auto-tier job stopped due to ACL callback error

Description: SUP-812

Auto-tier job fails with "ACL callback error" when "Copy ACL" is set to "Yes" and the "rsync+ssh" protocol is selected. The auto-tier job completes successfully when "Copy ACL" is set to "No"

Microsoft Identity Management for UNIX (IDMU) authentication not supported

Description: 13535, SFR-56

If you have previously configured the idmap service to use "idmu" mapping strategy through NMC by typing:

```
nmc:/$ option expert_mode=1
nmc:/$ !bash
# svccfg -s svc:/system/idmap setprop \ config/
  directory_based_mapping = astring: idmu
```

You need to either:

(A) Remove this configuration change (reverting to the default "name" based mapping strategy)

(B) Delay your upgrade until enhancement #13535 (support idmap config/ directory_based_mapping = idmu) is implemented

❖ *To determine whether you have this idmap configuration setting, look for these results in NMC:*

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

```
nmc:/$ !bash
# svccfg -s svc:/system/idmap listprop config
```

If you see the following response IDMU is enabled:

```
config/directory_based_mapping astring idmu
```

While we believe this impacts a very small number of existing customers, we are currently working on a solution.

Multi-NMS prohibits use of proxies when upgrading

Description: SUP-542, SUP-561

When upgrading using the NMC command `setup appliance nms property upgrade_proxy` the proxy settings are not honored if Multi-NMS is enabled. Multi-NMS is enabled by default.

Workaround:

- *Disable Multi-NMS, using NMC:*

1. Decrease the size of NMS-pool:

```
nmc:/$ setup appliance nms property srvpool_cnt_max -p 0 -y
```

2. Restart NMS:

```
nmc:/$ setup appliance nms restart
```

Issues with SuperMicro® physical view in NMV

Description: 13297

When using SuperMicro hardware where the shared JBOD(s) have been manually setup using NMC, and in an HA configuration, it is possible the physical view of the JBOD within NMV is not correctly displayed on one of the nodes.

Workaround:

Manually configure the shared JBOD(s) using the NMC `setup jbod model` command on both nodes prior to failover. If the service has already been failed over, issue the NMC `setup jbod model` on the node missing the JBOD physical view.

Global Headquarters
451 El Camino Real, Suite 201
Santa Clara, CA 95050
USA

7000-nxs-3.1.6-FP4-000010-D