



NexentaStor FAQ

Here you will find a collection of questions and answers related to NexentaStor installation and deployment

1 - FAQ: Machine signature has changed on my appliance. Please advise.

Older versions of the appliance software sometimes produced situations when the appliance's machine signature would change. Fresh **install** of the version **3.0.4** or later fixes the problem. Notice the emphasis on "install". Yes, 3.0.4 has a fixed machine signature mechanism that produces globally unique machine signatures that will NOT change with BIOS upgrades and hardware reconfigurations. However, to make use of this new mechanism you need to install from scratch..

In very rare cases, you may **also** need to upgrade the system's BIOS. Here's a specific recommendation for Supermicro mobos that have an older BIOS version:

1. Do a fresh install of NexentaStor v3.0.4 or later
2. Update a new BIOS version for the X8DTN+-F motherboard available on the Supermicro web site.
3. Once the new BIOS version is installed, the machine ID is stays fixed despite several test reboots. The issue of changing machine ID with Supermicro chassis after each reboot is now resolved.

2 - How do you calculate storage space?

NexentaStor Enterprise Edition imposes a simple, unified, capacity-based limitation - a limit on the **total physical** ("raw") storage space, excluding hot spares, log devices and cache devices.

For instance, let's say you have 8 (eight) 1TB drives. The following table illustrates a few possible examples of deploying these drives, and the resulting licensed capacity:

Configuration	Storage space counted towards license limit
1) single volume (ZFS pool) with no spares, cache or log devices	8TB
2) multiple volumes with no spares, cache and log devices	8TB
3) single volume with 6 drives in a RAIDZ configuration and 2 hot spares	6TB
4) two volumes with 4 drives total in MIRROR configurations, 2 hot spares and 2 SSDs slogs	4TB
5) single volume with 4 drives in a RAIDZ2 configuration, 2 spares, 1 log and 1 cache device	4TB

Therefore:

ZIL slogs (a.k.a. log) devices, L2ARC (a.k.a. cache) devices, AND spare drives are NOT counted.

Your appliance's licensed storage capacity is simply a sum of all ("raw") drive sizes, excepting logs, caches and spares.

And finally, unlimited **site license** does not have any capacity limitation.

3 - FAQ: Does the appliance support SNMP?

Simple Network Management Protocol (SNMP) is a legacy network management protocol used to monitor and administer devices on the IP network. SNMP is standardized by the Internet Engineering Task Force (IETF). An SNMP-managed network consists of three key components:

Managed device (for instance NexentaStor appliance)

SNMP agent (for instance, NexentaStor built-in agent)

SNMP manager that executes SNMP Get, GetNext, GetBulk, Set operations, as well as SNMPv3 Inform, to monitor and control the device. SNMP Traps are also supported.

NexentaStor provides a built-in SNMP agent that implements MIB-II a.k.a RFC-1213. More exactly, the agent supports **SNMPv2-MIB (MIB-II), HOST-RESOURCES-MIB, DISMAN-EVENT-MIB, NOTIFICATION-LOG-MIB**, and provides basic information to monitoring system using SNMP protocol. This information includes following sections:

Basic system (hostname, uptime, contacts, etc)

Network interfaces (IP/MAC addresses, inbound/outbound traffic, errors, other)

System load (CPUs, memory, folder usage)

Processes

NexentaStor supports SNMP traps. NexentaStor supports all three versions of SNMP protocol, including SNMPv3.

Setting up SNMP agent

Basic setup will be done using the following NMC command (example):

```
nmc$ setup network service snmp-agent configure
Read-only community string : public
Read-write community string : private
User name for SNMPv3 agent : alpha
Password for SNMPv3 agent. Must be at least 8 characters : pass1234
Destination of SNMP-traps : 192.168.1.55
```

The example above configures SNMP agent with a (sample) user account. Note that SNMPv3 uses MD5 for authentication and DES for privacy. Once SNMP agent on the appliance is configured, you can execute a simple MIB walk from a remote appliance, for instance:

```
# snmpwalk -v 3 -a md5 -A pass1234 -u alpha -l authPriv -x DES -X pass1234 SystemUpTime
```

The example above assumes that the remote host has snmpwalk utility (NexentaStor does include it), and that the user/password is specified as 'alpha' and 'pass1234', respectively. A sample output follows below:

```
SNMPv2-MIB::sysDescr.0 = STRING: abc
SNMPv2-MIB::sysObjectID.0 = OID: NET-SNMP-MIB::netSnmpAgentOIDs.3
DISMAN-EVENT-MIB:: SystemUpTimeInstance = Timeticks: (25964852) 3 days, 0:07:28.52
SNMPv2-MIB::sysContact.0 = STRING: johnsmith
SNMPv2-MIB::sysName.0 = STRING: devel119
SNMPv2-MIB::sysLocation.0 = STRING: corpoffice
SNMPv2-MIB::sysORLastChange.0 = Timeticks: (6) 0:00:00.06
SNMPv2-MIB::sysORID.1 = OID: SNMP-FRAMEWORK-
MIB::snmpFrameworkMIBCompliance
```

...

```
IF-MIB::ifNumber.0 = INTEGER: 2
IF-MIB::ifIndex.1 = INTEGER: 1
IF-MIB::ifIndex.2 = INTEGER: 2
IF-MIB::ifDescr.1 = STRING: lo0
IF-MIB::ifDescr.2 = STRING: ae0
IF-MIB::ifType.1 = INTEGER: softwareLoopback(24)
IF-MIB::ifType.2 = INTEGER: other(1)
```

IF-MIB::ifMtu.1 = INTEGER: 8232
IF-MIB::ifMtu.2 = INTEGER: 1500
IF-MIB::ifSpeed.1 = Gauge32: 127000000
IF-MIB::ifSpeed.2 = Gauge32: 10000000
IF-MIB::ifPhysAddress.1 = STRING:

...

RFC1213-MIB::atIfIndex.2.1.192.168.1.111 = INTEGER: 2
RFC1213-MIB::atIfIndex.2.1.192.168.1.119 = INTEGER: 2
RFC1213-MIB::atIfIndex.2.1.192.168.1.200 = INTEGER: 2

...

RFC1213-MIB::atIfIndex.2.1.224.0.0.0 = INTEGER: 2
RFC1213-MIB::atPhysAddress.2.1.192.168.1.55 = Hex-STRING: 00 1F D0 9E D2 25
RFC1213-MIB::atPhysAddress.2.1.192.168.1.77 = Hex-STRING: 00 22 15 D4 F5 BC
RFC1213-MIB::atPhysAddress.2.1.192.168.1.111 = ""
RFC1213-MIB::atPhysAddress.2.1.192.168.1.119 = Hex-STRING: 00 0C 29 8A 55 2A
RFC1213-MIB::atPhysAddress.2.1.192.168.1.200 = Hex-STRING: 00 0C 29 17 16 6D
RFC1213-MIB::atPhysAddress.2.1.224.0.0.0 = Hex-STRING: 01 00 5E 00 00 00
RFC1213-MIB::atNetAddress.2.1.192.168.1.55 = Network Address: C0:A8:01:37
RFC1213-MIB::atNetAddress.2.1.192.168.1.77 = Network Address: C0:A8:01:4D
RFC1213-MIB::atNetAddress.2.1.192.168.1.111 = Network Address: C0:A8:01:6F
RFC1213-MIB::atNetAddress.2.1.192.168.1.119 = Network Address: C0:A8:01:77
RFC1213-MIB::atNetAddress.2.1.192.168.1.200 = Network Address: C0:A8:01:C8
RFC1213-MIB::atNetAddress.2.1.224.0.0.0 = Network Address: E0:00:00:00
IP-MIB::ipForwarding.0 = INTEGER: notForwarding(2)
IP-MIB::ipDefaultTTL.0 = INTEGER: 255
IP-MIB::ipInReceives.0 = Counter32: 77366
IP-MIB::ipInHdrErrors.0 = Counter32: 0
IP-MIB::ipInAddrErrors.0 = Counter32: 0
IP-MIB::ipForwDatagrams.0 = Counter32: 0
IP-MIB::ipInUnknownProtos.0 = Counter32: 0
IP-MIB::ipInDiscards.0 = Counter32: 0
IP-MIB::ipInDelivers.0 = Counter32: 80748
IP-MIB::ipOutRequests.0 = Counter32: 62375
IP-MIB::ipOutDiscards.0 = Counter32: 18
IP-MIB::ipOutNoRoutes.0 = Counter32: 0
IP-MIB::ipReasmTimeout.0 = INTEGER: 15 seconds
IP-MIB::ipReasmReqds.0 = Counter32: 0
IP-MIB::ipReasmOKs.0 = Counter32: 0

```
IP-MIB::ipReasmFails.0 = Counter32: 0
```

```
IP-MIB::ipFragOKs.0 = Counter32: 0
```

...

Alternatively, use any compliant MIB browser to walk the appliance's MIB. Advanced configuration of the NexentaStor SNMP agent can be done via:

```
nmc$ setup network service snmp-agent edit-settings
```

For details on SNMP configuration, please see 'help snmpd.conf'.

```
nmc$ help snmpd.conf
```

4 - FAQ: How can I evaluate NexentaStor commercial extensions (plug-ins)?

NexentaStor extension, a pluggable module or a plugin, can be easily added and removed. A plug-in implements a certain well-defined extended functionality and uses the same Storage Appliance API (SA-API) as all the rest of the software components including 2nd tier storage services, fault management and statistic collection runners, and the management console and management web GUI. At installation time, the plug-in integrates itself with the appliance's core software.

All available NexentaStor licensed features (plug-ins) are listed here. Plug-ins extend the capabilities of NexentaStor in the areas of Data Replication, Data Security, High Availability, Virtualization, Performance Benchmarks, and Storage Management.

Certain NexentaStor plug-ins are free and can be used with all available NexentaStor editions. Other plug-ins are commercial. Please note:

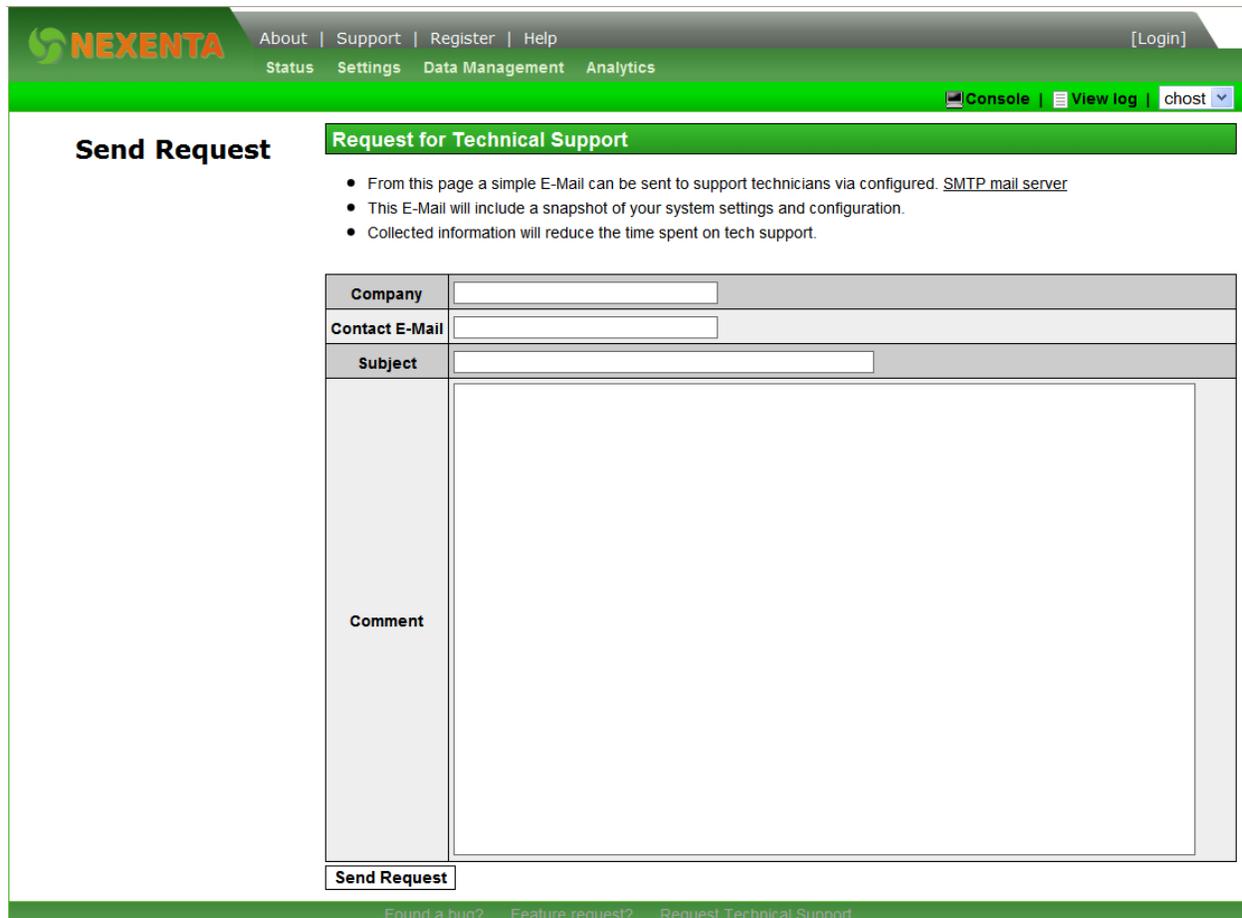
Commercial plug-ins cannot be used with NexentaStor Community Edition.

Commercial plug-ins require Free Trial Edition.

Note that plug-ins are not downloadable from the website. Pluggable modules are distributed exactly in the same way as NexentaStor software upgrades and updates: via built-in reliable transactional upgrade mechanism (see the Support section of our website for more information).

5 - FAQ: How to ask for technical support

NexentaStor web GUI provides a simple and effective mechanism to generate support requests. In Nexenta Management View (NMV), simply click on Support top level menu.



The screenshot shows the 'Request for Technical Support' form in the Nexenta Management View (NMV) web GUI. The page has a green header with the Nexenta logo and navigation links: About, Support, Register, Help, Status, Settings, Data Management, Analytics, and a [Login] button. A secondary navigation bar includes Console, View log, and a user dropdown menu (chost). The main content area is titled 'Send Request' and 'Request for Technical Support'. It contains a bulleted list of instructions: 'From this page a simple E-Mail can be sent to support technicians via configured. SMTP mail server', 'This E-Mail will include a snapshot of your system settings and configuration.', and 'Collected information will reduce the time spent on tech support.' Below the list is a form with fields for 'Company', 'Contact E-Mail', and 'Subject', followed by a large 'Comment' text area. A 'Send Request' button is located at the bottom of the form. At the very bottom of the page, there are links for 'Found a bug?', 'Feature request?', and 'Request Technical Support'.

This will result in an email to Technical Support, with your text (a problem description, feature request or any feedback) carrying a number of automatically attached log and configuration files, and possibly certain performance data.

The equivalent NMC command is called 'support':

```
nmc$ support
```

This is further illustrated in Quick Start Guide document, Section "Documentation and Technical Support".

You can always send a regular email to Technical Support. Using Tech Support form, however, is the preferred and recommended way as it allows technicians to start working on the problem immediately.

6 - FAQ: How do I replace a failed device in a mirror VDEV?

The intention of this document is to describe steps for replacing a faulty device in a ZFS pool without any available hot spares. This document focuses on a replacement of a device in a mirrored configuration. Replacement of a device in a mirror calls for detaching, and later re-attaching the device, which is different from replacing a device in a RaidZ-1, Raid-Z-2, or a RaidZ-3 VDEV configuration.

Download the document [here](#).

7 - FAQ: Multi NMS

<code>nmc\$ show appliance nms property grep srvpool</code>	0	0
<code>srvpool_cnt_initial</code>	2	Initial size of the pool of Management Servers
<code>srvpool_port_range_min</code>	2011	TCP port for the first Management Server in the pool to
<code>srvpool_interface_serialization_timeout</code>	60	Time interval to serialize API calls through a given
<code>srvpool_cnt_max</code>	0	Maximum size of the pool of Management Servers
<code>srvpool_affinity_timeout</code>	60	Time interval to retain client => NMS-from-the-pool

Those are multi-nms related knobs. The prefix “srvpool” stands for “pool of nexenta management servers” as the name implies. Full NMC help is available as well. To switch the feature off, set `srvpool_cnt_max = 0`

Some plug-ins disable the feature at install time, e.g., HA Cluster:

```
$Plugin::RsfCluster::IPC_SRVPOOL = 'disable';
```

8 - FAQ: How to Add Virtual IP Addresses to an HA Cluster Configuration

There are times when having a single Virtual Internet Protocol Address (VIP) tied to a Service Group of your HA Cluster is not sufficient, and multiple VIPs are required. A good example of this is where you might have dedicated networks for different security levels, or slow and fast networks, etc. Please keep in mind that this is an advanced configuration option, and failure to configure correctly may lead to loss of HA services. As such, please follow with caution, and if unsure, engage Support prior to making changes.

It is possible to configure multiple Virtual IPs for any Service Group, but there are some things to keep in mind.

Do not assign VIPs from the same subnet to physical interfaces with different bandwidth capabilities, such as 10GbE and 1GbE. This may result in inbound packets coming from the faster 10GbE interface and leaving via the slower 1GbE.

It is acceptable to create an aggregate out of available interfaces and assign a VIP to the aggregate interface.

Please, complete the following steps via NMC:

1. Create a configuration checkpoint to have a known good configuration snapshot, prior to making any changes in case further changes result in a malfunctioning configuration.

```
nmc@clus-a-node-1:/$ setup appliance checkpoint create
```

2. Validate current network configuration to determine which interface(s) are available to provision new VIP. Pay close attention to the netmask setting on the interface(s). If a network interface to which the VIP will be assigned is already configured with an IP address, make sure to validate and use the same network mask when configuring the VIP. It is not desirable to have multiple interfaces on the same subnet configured with different network masks.

```
nmc@clus-a-node-1:/$ show network interface
```

```
==== Interfaces ====
```

```
lo0: flags=2001000849 mtu 8232 index 1
    inet 127.0.0.1 netmask ffffffff
e1000g0: flags=1001000843 mtu 1500 index 2
    inet 10.10.100.100 netmask fffffc00 broadcast 10.10.103.255
```

```

ether 0:c:29:ed:4a:a8
e1000g0:1: flags=1001000843 mtu 1500 index 2
    inet 10.10.100.101 netmask ff000000 broadcast 10.255.255.255
e1000g1: flags=1000842 mtu 1500 index 3
    inet 0.0.0.0 netmask 0
    ether 0:c:29:ed:4a:b2
...snipped...

```

3. We can obtain further details for interface which will be used for the new VIP by adding to command from prior step. Interface must be physically connected to correct network. In our example, we select interface e1000g1, and validate its configuration.

```
nmc@clus-a-node-1:/$ show network interface e1000g1
```

```
==== Interface: e1000g1 =====
```

```
e1000g1: flags=1000842 mtu 1500 index 3
```

```
    inet 0.0.0.0 netmask 0
```

```
    ether 0:c:29:ed:4a:b2
```

```
==== Interface e1000g1: statistics =====
```

Name	Mtu	Net/Dest	Address	Ipkts	Ierrs	Opkts	Oerrs	Collis	Queue
e1000g1	1500	default	0.0.0.0	89235	0	12	0	0	0

4. In addition, it is a good idea to validate the routing table on each node. It is prudent to confirm that ALL nodes in the cluster share the same entries in the routing table. Any differences should be investigated.

```
nmc@clus-a-node-1:/$ netstat -nr
```

```
Routing Table: IPv4
```

Destination	Gateway	Flags	Ref	Use	Interface
default	10.10.100.1	UG	5	2770	
10.0.0.0	10.10.100.101	U	2	0	e1000g0
10.10.100.0	10.10.100.100	U	6	14291	e1000g0
127.0.0.1	127.0.0.1	UH	23	5714	lo0

...snipped...

5. Assuming that we do not have entries in **/etc/hosts**, we need to update entries on ALL nodes in the cluster. We need to add an entry for the new VIP. Consider being descriptive and using the Service Group name as part of the VIP hostname in the **/etc/hosts** file. Command **setup appliance hosts** will launch the vi-editor, and allow for modification of the hosts file.

We are adding the following entry to **/etc/hosts**:

```
10.10.100.102 clus-a-pool-a-01
```

```
nmc@clus-a-node-1:/$ setup appliance hosts
```

6. We will next validate **/etc/netmasks**, and for this VIP we are expecting to use a /22 netmask. Changes to **/etc/netmasks** should again be made on ALL nodes. Command **setup appliance netmasks** will launch the vi-editor, and allow for modification of the netmasks file.

We are adding the following entry to **/etc/netmasks**:

```
10.10.100.0 255.255.252.0
```

```
nmc@clus-a-node-1:/$ setup appliance netmasks
```

7. We should now have everything in place to configure the new Virtual IP. We can make this change from any node in the cluster. We will be presented with a a list of cluster groups, and need to select a group to be modified from the list. In this instance we are selecting **clus-1_pool_a**.

```
nmc@clus-a-node-1:/$ setup group rsf-cluster clus-alpha-01 vips add
Service      :
clus-1_pool_a
```

```
-----
Select HA Cluster service. Navigate with arrow keys (or hjkl), or
Ctrl-C to exit.
```

8. Next, we are presented with a configuration VIPs currently in place, and a prompt to add our new VIP. In this case we are going to use hostname **clus-a-pool-a-01**, which we added to **/etc/hosts** earlier.

Your current network configuration for service 'clus-1_pool_a':

```
VIP1: clus-a-pool-a
```

```
clus-a-node-1.homer.lab:e1000g0
```

```
clus-a-node-2:e1000g0
```

```
VIP2 Shared logical hostname:
```

IPv4 address or hostname that maps to the failover IP interface on all the appliances in the HA Cluster. This IP address or hostname is expected to be used by services such as NFS/CIFS/iSCSI to reliably access shared storage at all times.. Press Ctrl-C to exit.

9. At this point we are prompted with a list of interfaces available to us on the system, and as we planned earlier, we will select e1000g1 from this list. Names and numbers of interfaces will vary from system to system. We are going to see the prompt to select interface for this VIP repeat for each node in the cluster. We are expecting that all nodes will use the same interface.

VIP2 Shared logical hostname: clus-a-pool-a-01
VIP2 Network interface at clus-a-node-1.homer.lab:
e1000g2 e1000g3 e1000g0 e1000g1

10. Now we have to specify the mask used for this VIP, and while we already have /etc/masks updated, we are still going to enter the mask here, assuming that if there any changes in the future to the logic of HA configuration, we are already covered and do not have to reconfigure. We are using the same mask as in /etc/netmasks.

Verifying logical (failover) IP address '10.10.100.102' ...Success.
VIP2 Failover Netmask:

Optionally, specify IP network mask for the failover interface '10.10.100.102'. The mask includes the network part address and the subnet part. Press Enter to use default netmask: 255.0.0.0.. Press Ctrl-C to exit.

11. We are prompted with changes that are about to be committed, and we accept these changes. If we are done adding VIPs, we are going to say Yes to 'Stop adding VIPs?'.

The IP configuration to be used with 'clus-a-pool-a-01' is: 10.10.100.102/255.255.252.0.
Please confirm that this configuration is correct ? Yes
Stop adding VIPs? Yes
Success.

12. At this point we need to validate that the VIP was added, and on the system where this service group is active, we can validate that the VIP is now in service. We can simply ping the VIP to validate its state. We should make sure that we can ping the VIP from ALL nodes in the cluster.

```
nmc@clus-a-node-1:/$ ping clus-a-pool-a-01
clus-a-pool-a-01 is alive
```

9 - FAQ: Supermicro SBB onboard SAS expander issues

The Supermicro SUPERSERVER 6036ST-6LR Storage Bridge Bay (SBB) design includes the X8DTS-F serverboards. For best optimized solution of Nexenta and SBB together we require that you use an HBA connected into the available PCIe slots for additional storage. Please note that the 4-lane SAS connector at the rear for external storage connections cannot be used in a Nexenta environment due to SAS device timeouts, invalid DWORD counts, and running disparity error counts.

If needed, additional external storage can use HBAs connected into the available PCI-express slots.

Also note that direct attachment of SATA drives is not supported by Supermicro in the SBB. If a SATA drive is used, then a Supermicro AOC-LSISS9252 SATA/SAS interposer is required.

This information is dated 15-Jun-2011. Check the Supermicro documentation for pertinent changes.

10 - FAQ: SuperMicro Storage Bridge Bay SATADOM Power Cable Issue SuperMicro SBB SATADOM Power Cable Issue

Nexenta has become aware of an issue regarding the SuperMicro Storage Bridge Bay and the onboard SATA slot, in particular when used with a SATA DOM. The SuperMicro motherboard in the SBB does not provide power via that SATA slot, and as such, you will need an appropriate power cable to connect the SATA DOM to the small power slot provided directly next to the SATA slot. This cable needs to be very specific - able to plug in to the SATA DOM being used and into the slot on the motherboard, and is REQUIRED (as again, the SATA port is not providing power as some other motherboard do and some SATA DOM's are built to accept).

If you run into issues connecting a SATA DOM to the power slot with provided cabling, please contact SuperMicro or Nexenta Systems.

11 - FAQ: When is a volume upgrade needed?

For some NexentaStor releases, new versions of the ZFS software are available for data volumes (or ZFS pools). This happens for upgrades to:

NexentaStor 3.x from NexentaStor 2.x
NexentaStor 3.1 from NexentaStor 3.0.x

It is not required that a volume version be upgraded in order to continue to function using its versioned features. However, new functionality can be useful for some use cases.

Caveat Emptor: Once a volume is upgraded, it can no longer be accessed by older versions of NexentaStor that do not support the same version. There is no method for downgrading volume versions.

The list of volume versions and their features is available from NMC:

```
nmc@myhost:/$ show volume version -h
```

The following versions are supported:

VER DESCRIPTION

- 1 Initial ZFS version
- 2 Ditto blocks (replicated metadata)
- 3 Hot spares and double parity RAID-Z
- 4 zpool history
- 5 Compression using the gzip algorithm
- 6 bootfs pool property
- 7 Separate intent log devices
- 8 Delegated administration
- 9 refquota and reservation properties
- 10 Cache devices
- 11 Improved scrub performance
- 12 Snapshot properties
- 13 snapused property
- 14 passthrough-x aclinherit
- 15 user/group space accounting
- 16 stmf property support
- 17 Triple-parity RAID-Z
- 18 Snapshot user holds
- 19 Log device removal
- 20 Compression using zle (zero-length encoding)
- 21 Deduplication
- 22 Received properties
- 23 Slim ZIL

- 24 System attributes
- 25 Improved scrub stats
- 26 Improved snapshot deletion performance
- 27 Improved snapshot creation performance
- 28 Multiple vdev replacements

For more information on a particular version, including supported releases, see the ZFS Administration Guide.

NexentaStor releases have volume version support as follows

NexentaStor Release	Volume Version
2.2	14
3.0	22
3.0.3	26
3.1	28

The current volume versions are visible from NMC:

```
nmc@myhost:/$ show volume mydata version
System volumes' ZFS version: 28.
VOLUME      VER COMMENT
mydata      28 Version is up to date
```

In some cases, the volume-check runner will report that a volume can be upgraded. The notification of upgrade is also available from the volume status (or zpool status) commands in NMC.

The volume (pool) status will show a report similar to:

```
nmc@myhost:/$ show volume status mydata
volume: mydata
state: ONLINE
status: The pool is formatted using an older on-disk format. The pool can
still be used, but some features are unavailable.
action: Upgrade the pool using 'zpool upgrade'. Once this is done, the
pool will no longer be accessible on older software versions.
scan: none requested
config:
NAME      STATE  READ WRITE CKSUM
mydata    ONLINE  0  0  0
```

```
c11t0d0 ONLINE 0 0 0
errors: No known data errors
```

To upgrade a volume, use the NMC setup command:

Caveat Emptor: Once a volume is upgraded, it can no longer be accessed by older versions of NexentaStor that do not support the same version. There is no method for downgrading volume versions.

```
nmc@myhost:/$ show volume mydata version
System volumes' ZFS version: 28.
VOLUME      VER COMMENT
mydata      26  The volume is out of date, and can be upgraded. After
          upgrade, the volume will no longer be accessible by
          older software versions. See usage (-h) for details
```

```
nmc@myhost:/$ setup volume mydata upgrade-version
Unexpected option 'upgrade-version'
```

```
nmc@myhost:/$ setup volume mydata version-upgrade
Upgrade volume 'mydata' from version '26' to system version '28'? Yes
Volume 'mydata' upgraded to version 28.
```

See also the NMC commands:

```
setup volume version version-upgrade
setup folder version-upgrade
show volume version
show appliance syspool
```

12 - FAQ: Troubleshooting AD join CIFS issues step by step

CIFS troubleshooting

1. If they are running Win2k8, make sure they are on SP2 or R2.
2. Make sure the user joining the domain is a Domain Admin.
3. Find out how many DCs they have.
4. Set the preferred DC with `sharectl set -p pdc=IPADDRESSOFSERVER`
5. Check to see if `LMAUTH_LEVEL` is set to 2. if not, set it using `sharectl set -p lmauth_level=2`

6. Make sure Netbios is turned on the DC. Client for Microsoft Networks is on and File/Printer Sharing is On in network control panels.
7. From !bash, does #dig @NAMESERVER _ldap._tcp.dc._msdcs.domainname SRV +short return SRV records?
8. From !bash can you do a kinit?
9. restart smb/server service. At this point, see if you can join the domain. If not, as a last resort: Customer may not want this security policy changed.
10. disable SMB packet signing

Win2K8

Start the Group Policy Management tool on the DC and set the following:

Computer Configuration\Policies\Administrative Templates\System\NetLogon\Allow Cryptography Algorithms Compatible with Windows NT 4.0 -> Enabled

Then run: gpupdate /force

Win2K3

To disable SMB packet and secure channel signing enforcement on Windows Server 2003–based domain controllers:

1. Open **Active Directory Users and Computers**, right-click the **Domain Controllers** container, and then click **Properties**.
2. Click the **Group Policy** tab, then click **Edit**.
3. Under **Computer Configuration**, go to the Windows Settings\Security Settings\Local Policies\Security Options folder.
4. In the details pane, double-click **Microsoft network server: Digitally sign communications (always)**, and then click **Disabled** to prevent SMB packet signing from being required.
5. Click **OK**.
6. In the details pane, double-click **Domain member: Digitally encrypt or sign secure channel data (always)**, and then click **Disabled** to prevent secure channel signing from being required.
7. Click **OK**.
8. To apply the Group Policy change immediately, either restart the domain controller, or type **gpupdate /force** at a command line, and then press ENTER.
9. Try rejoining the domain.

On Win2k8 if that doesn't solve it, also make these group policy updates as well:

1. Default Domain Policy -> Computer Configuration -> Windows Settings -> Security Settings -> Local Policies -> Security Options -> Network security: Minimum session security for NTLM SSP based (including secure RPC) clients
2. Default Domain Policy -> Computer Configuration -> Windows Settings -> Security Settings -> Local Policies -> Security Options -> Network security: Minimum session security for NTLM SSP based (including secure RPC) servers
3. Changed the Default Domain Policy from Not Configured to Configured with the "Require 128-bit encryption" unchecked.
4. Rejoin the domain...
 1. Verify you have joined the domain using smbadmin

13 - FAQ: Appliance OS (syspool) size recommendation

When deciding how large to make your system pool, a good rule of thumb is that the minimum syspool size should be the larger choice of these two options:

1. 32 GB
2. System RAM total size + 16 GB

We do not recommend syspools less than 32 GB in size for production deployments. Also it is worth noting that larger is never bad, just potentially un-utilized. Erring on the side of caution is suggested.

14 - FAQ: Joining AD without Domain Admin

A domain admin needs to give these rights to a user who doesn't have domain admin privileges to join a computer to the domain.

1. Create a new computer Object.
2. Start ADSI Edit
3. Find and select the Computer Object you just created.
4. Right Click on Properties and select the Security Tab.

5. Add the User that is going to be used to join the domain on the Nexenta Host.
6. Click Advanced.

7. Select the user that is going to join the domain from the Nexenta host, and click Edit.
8. Click on the properties tab.
9. Enable Write userPrincipalName and Write userAccountControl

Now you can go into NMV and join the domain using a non domain admin account.

15 - FAQ: Installation of NexentaStor on a Virtual Machine

Process to install NexentaStor on a Virtual Machine

1. In the VMware Menu, Go to File -> New -> Virtual Machine.
2. In the New Virtual Machine Wizard, pick Typical/Custom configuration for the VM you are creating.
To keep things simple to begin with, pick Typical.
3. In the next screen which pops out, specify the location to "Install From" :
 - a. If you prefer to boot from Nexenta install CD, specify the location of the CD.
 - b. If you prefer to install from disk image file, specify the location of the "iso"/disk image file.
ex : C:\Nexenta\NexentaStor-COM-2.2.0.iso
4. In the next menu, specify OS (Sun Solaris) and version (Solaris10/Solaris 10 64 bit).
5. Next, specify the name of the VM and location for the VM file.
Ex : Name of VM: NexentaVM
Location of VM File : C:\Documents and Settings\Mukund\My Documents\My Virtual Machines\Solaris 10
6. Provide disk capacity on the next screen.
7. once you are ready to Create Virtual Machine, click finish.
8. New VM is created with Guest OS as Nexenta.
9. After installation, configure the NexentaStor appliance as per the guidelines in the Nexenta User Guide.

16 - FAQ: NMS and DBUS restart steps

If you attempt to login to your appliance via NMV or NMC and receive errors about timeouts, or any time when you can't login to the NMV or NMC but can login to bash (either directly as root and you get kicked to bash instead of NMC with an error, or you login as admin because logging in as root is timing out), it is often correctable with the following steps. Please note that if you are not comfortable following these instructions, it is perfectly OK to submit a support case requesting assistance, instead. Also, if these steps do not correct the issue, please submit a support case.

Note that these instructions expect a certain level of familiarity with the underlying bash shell and OpenSolaris command environment. You should be able to copy and paste the commands shown below directly to your prompt if you desire, but any mistakes in copying and pasting or any unexpected results may cause additional issues, so again, if you are uncomfortable running commands in the underlying shell, please submit a support case. Note that we only support the commands as shown below - running any other commands or deviations of the below commands is unsupported.

First Attempt Steps

Run the following command (as root in a bash prompt [not NMC prompt]): `svcs nm{s,v,cd} nbs nmdtrace dbus`

Example of the command and output (bolded sections here for emphasis and won't appear bold in your prompt):

```
root@homesan:/volumes# svcs nm{s,v,cd} nbs nmdtrace dbus
STATE      STIME  FMRI
online     Jul_10  svc:/system/dbus:default
online     Jul_10  svc:/application/nbs:default
online     Jul_10  svc:/application/nmdtrace:default
online     Jul_10  svc:/application/nms:default
online     Jul_10  svc:/application/nmcd:default
offline    5:00:39  svc:/application/nmv:default
```

Note that in the above example, the command showed that nmv (`svc:/application/nmv:default`) was offline. In the event you see one or more services offline, for each service you see in an "offline" or "maintenance" state, type (replacing [service] with the name of the service, in this above case, nmv, without the []): `svcadm clear [service]`

Example of this: `svcadm clear nmv`

After you have done this, run: `svcadm -v enable -rs nmv`

Then wait 60 seconds and check if NMV and NMC access has been restored. If not, try Second Attempt Steps below.

Second Attempt Steps

If you follow the First Attempt Steps and services are still not responding, especially if you see errors like this:

```
remote host 'localhost', port '2001': org.freedesktop.DBus.Error.NoReply: The reply timeout expired. Server is taking a long time to respond
```

Or any similar error mentioning DBus and timeouts or retries, you should run the command:
`svcadm -v disable nm{s,v,cd} nbs nmdtrace dbus`

Example output of that command below:

```
root@amber:/volumes# svcadm -v disable nm{s,v,cd} nbs nmdtrace dbus
svc:/application/nms:default disabled.
svc:/application/nmv:default disabled.
svc:/application/nmcd:default disabled.
svc:/system/dbus:default disabled.
svc:/application/nmdtrace:default disabled.
svc:/application/nbs:default disabled.
```

After running this, wait a full 60 seconds, and then type: `svcs nm{s,v,cd} nbs nmdtrace dbus`

Example output of what you should see below:

```
root@amber:/volumes# svcs nm{s,v,cd} nbs nmdtrace dbus
STATE      STIME    FMRI
disabled   14:23:41 svc:/application/nmcd:default
disabled   14:23:41 svc:/application/nmv:default
disabled   14:23:48 svc:/application/nmdtrace:default
disabled   14:24:12 svc:/application/nms:default
```

```
disabled    14:24:12 svc:/application/nbs:default
disabled    14:24:13 svc:/system/dbus:default
```

If you run this command and see something like this:

```
root@amber:/volumes# svcs nm{s,v,cd} nbs nmdtrace dbus
STATE      STIME    FMRI
disabled   14:23:41 svc:/application/nmcd:default
disabled   14:23:41 svc:/application/nmv:default
disabled   14:23:48 svc:/application/nmdtrace:default
online     Jul_08   svc:/system/dbus:default
online     Jul_08   svc:/application/nbs:default
online*    14:23:41 svc:/application/nms:default
```

Then you need to wait another 60 seconds and check again. If you still see anything but the first printout (with all services showing as disabled), wait a third 60 seconds and check one last time. If you are still seeing services not in disabled state, please submit a support case. Assuming within this timeframe you eventually see an output like the first output shown with all services disabled, the next thing to do is check a few common trouble spots. First, run this command:
cat /etc/hosts | grep -E 'localhost|loghost'

Example output below:

```
root@homesan:/volumes# cat /etc/hosts | grep -E 'localhost|loghost'
::1          localhost
127.0.0.1    localhost
192.168.2.15 homesan homesan.home.local loghost
```

You should see output like the above, but yours will of course vary depending on your environment. What you are looking for, though, is that there is a line of "127.0.0.1" with "localhost" at some point in the list of names after it. If you do not have a "localhost" definition pointing at 127.0.0.1, you have a problem. If you know how to edit /etc/hosts and add a line of "127.0.0.1 localhost", or modify an existing "127.0.0.1" line to include localhost, you can do so; if not, please contact support for assistance (be sure to include the output of all the commands you've run up to this point in your case submission).

Secondly, it is preferable if you get at least one hit on "loghost". It is perfectly acceptable to see "loghost" on multiple entries, that is fine and nothing to be concerned about, but you need to see "loghost" on at least one line, and at least one of those lines should point at some IP of local appliance (unless you've been told otherwise by a certified Nexenta engineer). Any IP will do, as long as it is one this appliance you are on has bound to it. Usually it is on the management IP, but not always.

Assuming you've verified you have 127.0.0.1 referenced to localhost and have at least one local IP with a loghost entry, the next thing to do is run this command and look at the output: `ps -Aef | grep -E 'hddisco|zfs\ |zpool\ ' | grep -v 'grep -E hddisco'`

Example output (what you should see) below:

```
root@homesan:/volumes# ps -Aef | grep -E 'hddisco|zfs\ |zpool\ ' | grep -v 'grep -E hddisco'
root@homesan:/volumes#
```

As you can see from the above, nothing was returned. This is what you want. If it returns any other lines of output, you need to file a support case (be sure to include all the output of the commands you've run up to this point in your case submission), as it is quite likely that something is going on with the system that will disallow NMS from being restarted until it is corrected or finishes, and a support technician at Nexenta can help diagnose this and a proper next step. Assuming the above command returned no output as shown above, next type: `svcadm -v enable nm{s,v,cd} nbs nmdtrace dbus`

Example output below:

```
root@amber:/volumes# svcadm -v enable nm{s,v,cd} nbs nmdtrace dbus
svc:/application/nms:default enabled.
svc:/application/nmv:default enabled.
svc:/application/nmcd:default enabled.
svc:/application/nbs:default enabled.
svc:/application/nmdtrace:default enabled.
svc:/system/dbus:default enabled.
```

Then wait a full 60 seconds. Once 60 seconds has elapsed, type this command: `svcs nm{s,v,cd} nbs nmdtrace dbus`

Example output below:

```
root@amber:/volumes# svcs nm{s,v,cd} nbs nmdtrace dbus
STATE      STIME    FMRI
online     14:36:19 svc:/system/dbus:default
online     14:36:19 svc:/application/nbs:default
online     14:36:20 svc:/application/nmdtrace:default
online     14:36:26 svc:/application/nms:default
online     14:36:26 svc:/application/nmv:default
online     14:36:29 svc:/application/nmcd:default
```

If you see an output like the above, you can now attempt to login to the NMV or NMC and see if they respond. If they do, you should be in good shape. If they do not, you need to submit a support case (be sure to paste all the output from the commands you ran above into the support case). If you do not see an output like the above, and instead see an output where services have a status of offline and one or more have a status of online*, wait another 60 seconds and check again. Repeat up to 5 times (it can take up to 5 minutes in some extreme cases). If after 5 minutes your output still doesn't look like the above, please submit a support case (be sure to include all the output from the above commands in your case submission).

These steps can help clear up transient NMS and DBUS issues, commonly stemming from temporary outages. If you find your appliance loses connectivity to NMV or NMC and the above instructions help, but have to be repeated multiple times in a month, there is a strong possibility that something is unhealthy about your NexentaStor appliance, and you should submit a full debug-level log case to Nexenta Support for analysis. As always, if you have any questions or problems with any of the above, or need Nexenta's assistance with correcting problems, please submit a support case and we will be glad to assist you.

17 - FAQ: Generate and send config and diag from NMC without SMTP

From the NMC prompt:

```
nmc@nexstor:/$run diagnostics
```

This creates a file called general-diag.txt.

```
nmc@nexstor:/$show all > config.txt
```

```
nmc@nexstor:/$show appliance syslog dmesg >dmesg.txt
```

```
nmc@nexstor:/$tar -czvf diag.tgz general-diag.txt config.txt dmesg.txt
```

```
nmc@nexstor:scp diag.tgz user@hostname.com:/tmp
```

18 - FAQ: Nexenta Core Platform | Nexenta CP

Where Did Nexenta Core Platform Go?

Nexenta Core Platform (AKA Nexenta CP, AKA NexentaOS), has been moved to the Illumos community at <http://www.illumian.org>, to fit with the new name change and the tighter alliance with the illumos community.

We recently moved to an illumos kernel, but continue the debian packaging, illumos + debian = illumian.

The mirror of the ncp repo is still alive at apt.zpool.org.

19 - FAQ: PXE booting NexentaStor

PXE-Boot Installation of NexentaStor

Introduction

The Preboot eXecution Environment (PXE, pronounced pixie) is used to boot computers using a network interface, with a bootable image located at a remote specified location. PXE employs a combination of DHCP and TFTP, to locate the bootable image, and boot from it.

Prerequisites

1. NexentaStor ISO image.
2. NFSv3 server.
3. DHCP server with the ability to modify the scope options.
4. TFTP server

Procedures

NOTE: For this tutorial our TFTP and NFS services are running on the same host with IP <host-ip>. The NFS export defined in /etc/exports is referred to as <nfs-root> and the root TFTP directory is referred to as <tftpboot>.

Copy the NexentaStor ISO image to your server and mount the ISO using the loop device.

```
mount -o loop <path-to-iso> <mnt-point>
```

Copy the contents of the ISO to a folder that is exported via NFS.

```
cp -r <mnt-point>/* <nfs-root>/nexenta
```

Create the following directories in <tftpboot>.

```
mkdir -p <tftpboot>/boot/grub  
mkdir -p <tftpboot>/platform/i86pc  
mkdir -p <tftpboot>/platform/i86pc/kernel
```

Copy the kernel, miniroot and loader from the ISO image to the tftpboot directory.

```
cp -r <mnt-point>/platform/i86pc/kernel/* <tftpboot>/platform/i86pc/kernel/  
cp <mnt-point>/platform/i86pc/miniroot <tftpboot>/platform/i86pc/  
cp <mnt-point>/boot/grub/pxegrub <tftpboot>/boot/grub/
```

Create the boot loader menu, menu.lst, in the <tftpboot>/boot/grub/ directory. The menu.lst file defines the grub boot options.

```
vim <tftpboot>/boot/grub/menu.lst
```

Add the following text to the boot loader menu, menu.lst.

```
default 0  
timeout 10  
  
title Nexenta PXE Installation  
kernel$ /platform/i86pc/kernel/$ISADIR/unix -B iso_nfs_path=<host-ip>:<nfs-root>/nexenta  
module$ /platform/i86pc/miniroot
```

Configure the DHCP server scope options to point to your TFTP server and the appropriate boot file.

If you are running a Windows DHCP server navigate to Start->Administrative Tools->DHCP and expand your DHCP scope. Right click on the Scope Options and select 'Configure Options...'. Add the following options and values.

```
066 Boot Server Host Name -> <host-ip>  
067 Bootfile Name -> /boot/grub/pxegrub
```

If you are running a Linux DHCP server using ISC DHCP then you must modify the /etc/dhpc3/dhcpd.conf file and add the following options to the scope.

```
filename "/boot/grub/pxegrub";  
option grubmenu "/boot/grub/menu.lst";  
next-server <host-ip>;
```

After modifying dhcpd.conf you must restart the dhcpd service.

Enable PXE on your client NIC and power on the system. The client should boot to a grub menu with the 'Nexenta PXE Installation' option.

Appendix 1 - Modifying miniroot

The kernel path in the pxegrub menu is read during multiple stages of the bootstrap processes. The path relative to the root of the tftp directory and the miniroot image must be identical for the boot process to proceed successfully. The user may want to define their own directory structure on the TFTP server so the miniroot image and kernel from multiple NexentaStor releases can be hosted on the same system. The following instructions describe how to modify the miniroot image with a new kernel path.

The following procedures must be executed on a Solaris based system.

Uncompress the miniroot image.

```
#> mv miniroot miniroot.gz  
#> gunzip miniroot.gz
```

Mount the uncompressed miniroot image.

```
#> lofiadm -a miniroot  
#> mount /dev/lofi/<x> /mnt/miniroot
```

Create a symbolic link from the root of the filesystem that matches the path from the root of your TFTP server to the kernel directory.

```
#> mkdir -p /<new>/<path>  
#> ln -s /platform/i86pc/kernel/ /<new>/<path>/.
```

Unmount the miniroot image.

```
#> umount /mnt/miniroot  
#> lofiadm -d /dev/lofi/<x>
```

Compress the miniroot image.

```
#> gzip miniroot  
#> gzip -t miniroot.gz  
#> mv miniroot.gz miniroot
```

You can now replace the default miniroot image included with the NexentaStor release with the updated image. With these procedures you can store the kernel files in any directory on your TFTP server.

Appendix 2 - Loading pxegrub From pxelinux

Many users already have an existing pxelinux based PXE server configured in their environment. Unfortunately pxelinux cannot load the Solaris based installer; however, you can configure pxelinux to chain load pxegrub. The following instructions detail how to add a menu option to pxelinux to chain load pxegrub.

Copy the pxegrub loader from the ISO image to tftpboot root directory.

```
cp <mnt-point>/boot/grub/pxegrub <tftpboot>/pxegrub.0
```

NOTE: You must append '.0' to the end of the file name otherwise pxelinux will not load the file.

Add the following clause to your pxelinux menu.

```
label pxegrub  
kernel pxegrub.0
```

20 - FAQ: Difference Between Versions?

What are the differences between versions of NexentaStor?

The differences between the Community and Trial and Enterprise versions are:

Feature	Community Edition	Free Trial Edition	Enterprise
Time limitations	None	45 days	None
Storage capacity	18 TB	Unlimited	Unlimited
Commercial plugins			
HA functionality	-	Yes	Yes
Virtualization	-	Yes	Yes
Data Replication	-	Yes	Yes
Data Security	-	Yes	Yes
Support and upgrades			
Technical Support	-	Yes	Yes
Maintenance Agreement	-	-	Yes
Storage Upgrades	-	Yes	Yes

21 - FAQ: Doesn't flash wear out quickly when used for writes?

Yes, flash does wear out, but today's enterprise flash vendors have done a lot to improve wear-leveling algorithms, reserve spare flash cells, and otherwise extend the longevity of flash in enterprise devices. Don't buy a USB thumb drive for your ZIL, and know that not all SSDs are created equal. Consider something on Nexenta's Hardware Supported List (HSL), and contact a sales representative or sales engineer for specific guidance.

22 - FAQ: What happens if a flash drive fails?

First of all, data isn't lost; losing a ZIL/SLOG or L2ARC device is simply an impact to performance. Any data in the L2ARC also is in the spinning HDDs, so losing an L2ARC device simply means that any requests for the data now must be serviced by reads from the spinning HDDs. If multiple flash drives are used for L2ARC, they are used in a striped fashion, and the remaining devices will continue to be used if one fails. Any data in the ZIL/SLOG also is in the ARC until it is flushed to the spinning HDDs, which NexentaStor does every five seconds by default. Thus, only if the SLOG device fails and controller power is lost within five seconds would data be lost. This is extremely rare, but many customers prefer the comfort of mirrored SLOG devices nonetheless, so the options to stripe or mirror data across SLOG devices both exist.

23 - FAQ: Can I use PCIe-based flash devices?

Yes, but the tradeoff here is much like that described in the question asking, **“Can I put L2ARC in the head nodes in HA Clusters?”** PCIe-based flash devices obviously must reside in the head nodes, so if a head fails or is powered off, the other node cannot utilize that flash. (There is no cross-node cache mirroring in ZFS today.) While that is the downside, the upside is that PCIe-based flash devices offer lower I/O latencies than SSD flash devices. So, where performance is most critical, some customers prefer to have PCIe-based flash devices in the head nodes and then plan to restore their HA Clusters to the optimal state as quickly as possible if something happens.

24 - FAQ: What about an all-SSD pool?

Many Nexenta customers use all-SSD pools. Where the budget allows, this offers a very low-latency solution and negates the benefit of the L2ARC and SLOG (unless SSDs with differing performance characteristics are used).

25 - FAQ: Does Nexenta have its own SMB development team or does it simply take and port Samba's effort?

We have our own strong SMB team led by a senior CIFS expert recruited from Sun. Our current in-kernel CIFS implementation is feature-rich and stable, and we are currently developing a next generation CIFS stack for our NexentaStor platform to keep up with the latest technology trends.

Nexenta is a registered trademark of Nexenta Systems Inc., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Nexenta Systems Inc.

© Nexenta Systems Inc. 2014 All Rights Reserved.

