



Release Notes | NOS 4.0.1

Last modified: July 23, 2014 (2014-07-23-11:14 GMT-07:00)

Features

In addition to the NOS 4.0 features, NOS 4.0.1 includes the following new features.

- Configurable remote syslog forwarding enables you to send logs to a remote server using the TCP/UDP protocols. Each log in `/home/nutanix/data/logs/` is prefixed with the name of the module (for example, `cassandra`) generating the information. [ENG-12848, ENG-17953]
- If you are using the multi-cluster management feature (also known as Prism Central), you can access convenient and automated cluster NOS upgrades through a web console upgrade dialog. Automatic software alerts notify you of available upgrades, which you can install manually or automatically. [FEAT-382]
- Volume Shadow Copy Service (VSS) support for Hyper-V hosts [FEAT-632]

NOS 4.0 includes the following new features.

- General performance and stability improvements. [FEAT-400, FEAT-377, FEAT-376, FEAT-379]
- NOS introduces feature-based licensing for nodes and clusters to enable you to administer your environment according to your current and future needs. All Nutanix models ship with the default Starter license, with upgrades to Pro and Ultimate licenses available. Licenses are administered through the web console and Nutanix Support Portal. See the *License Management* topics in the *Web Console Guide*. [FEAT-405]
- Disaster recovery support for Windows Server 2012 R2 Hyper-V. NOS disaster recovery options now include support for Hyper-V hypervisors. Support is seamless through the web console and nCLI. [FEAT-365]
- Ability to monitor and manage multiple clusters through a single web console. Known as Prism Central, this user-installed VM enables centralized management of user-specified clusters. [FEAT-381]
- Convenient and automated cluster NOS upgrades through a web console **Upgrade Software** dialog . Automatic software alerts notify you of available upgrades, which you can install manually or automatically. [FEAT-382]
- Monitor and manage cluster health through the The Prism Central Health dashboard. This health dashboard displays dynamically updated health information about hosts, disks, and VMs across the registered clusters. [FEAT-385]
- Nutanix clusters are now "block aware", which means that redundant copies of any data required to serve I/O are placed on nodes that are not in the same block. With block awareness, where multiple nodes in a block fail, the cluster can continue to run because cluster configuration data exists on other blocks. Guest VMs can continue to run because redundant copies of guest VM data and metadata exist on other blocks [FEAT-391]
- Enhanced data protection on a protection domain basis, allowing creation of scheduled backup of copies of VMs and NFS files with specific remote retention policies for remote sites through the Prism UI. Through calendaring, scheduling can be as granular as one hour or one month. [FEAT-397]
- Disk space savings through capacity tier duplication, also known as on-disk deduplication. This addition to existing deduplication features can help increase greater VM density per node. [FEAT-392]

- The Nutanix hosted backup service on Amazon enables Nutanix customers to back up and restore copies of virtual machines and files to and from an on-premise cluster and an Amazon Web Services (AWS) Elastic Cloud 2 (EC2) server. As this feature is available only as a tech preview and is not recommended in a production environment, contact your Nutanix Support Reliability Engineer for details. [FEAT-395]
- PowerShell cmdlets for controlling Nutanix clusters. You can download and install the cmdlets from the web console to your Windows desktop, then administer a cluster through PowerShell. The Nutanix PowerShell cmdlets are described in the NOS 4.0 documentation. [FEAT-387]
- Redundancy factor 3, a configurable option that allows a Nutanix cluster to withstand the failure of 2 nodes or drives in different blocks. By default, Nutanix clusters have redundancy factor 2, which means they can tolerate the failure of a single node or drive. [FEAT-464]

Installation

You can upgrade to NOS 4.0.1 from the following versions:

- 4.0
- 3.5.x
- 3.1.x

If you are upgrading from version 3.1.x or 3.5.x, use the Nutanix *Upgrade Guide* for complete instructions. If you are upgrading from version 4.0, also see the *Web Console Guide* for information about NOS software upgrades.

After upgrading from a version of NOS lower than 4.0, apply licenses to the cluster through the web console.

Product Mixing Restrictions

While a Nutanix cluster can include different products, there are some restrictions. The following compatibility matrix lists the supported combinations.



Caution: Do not configure a cluster that violates any of the following rules.

Compatibility Matrix

	NX-1020 ¹	NX-1050	NX-2000	NX-2050	NX-3000	NX-3050	NX-6000	NX-7000
NX-1020	•							
NX-1050 ²		•	•	•	•	•	•	•
NX-2000		•	•	•	•	•		•
NX-2050		•	•	•	•	•	•	•
NX-3000		•	•	•	•	•	•	•
NX-3050		•	•	•	•	•	•	•
NX-6000 ³		•		•	•	• ⁴	•	•
NX-7000		•	•	•	•	•	•	•

1. NX-1020 nodes cannot be mixed with other products. The maximum cluster size is 8 nodes.
2. NX-1050 nodes can be mixed with other products in the same cluster only when they are running 10 GbE networking; they cannot be mixed when running 1 GbE networking. If NX-1000 nodes are using the 1 GbE interface, the maximum cluster size is 8 nodes. If the nodes are using the 10 GbE interface, the cluster has no limits other than the maximum supported cluster size that applies to all products.
3. NX-6000 nodes cannot be mixed with NX-2000 nodes in the same cluster.
4. Because it has a larger Flash tier, NX-3050 is recommended over other products to be mixed with NX-6000.

-
- All nodes in a cluster must be the same hypervisor type (ESXi, KVM, or Hyper-V).
 - All Controller VMs in a cluster must have the same NOS version.
 - Mixed Nutanix clusters comprising NX-2000 nodes and other products are supported as specified above. However, because the NX-2000 processor architecture differs from other models, vSphere does not support enhanced/live vMotion of VMs from one type of node to another unless Enhanced vMotion Capability (EVC) is enabled. For more information about EVC, see the vSphere 5 documentation and the following VMware knowledge base articles:

- *Enhanced vMotion Compatibility (EVC) processor support* [[1003212](#)]
- *EVC and CPU Compatibility FAQ* [[1005764](#)]

Enabling EVC on an existing vCenter cluster requires shutting down all VMs. Because EVC is a cluster setting, it applies to all hosts at the same time. For these reasons, adding a node with a different processor type to a vCenter cluster requires shutting down the Nutanix cluster.

Resolved Issues

The following critical issues were resolved in NOS 4.0.1:

- nCLI commands erroneously issued in a multicluster configuration might display incorrect results. [[ENG-14134](#)]
- Disk usage statistics incorrectly reported by ESXi disk tools. [[ENG-14461](#)]
- If cassandra is compacting an unusually large number of files, the stargate process might fail. [[ENG-15519](#)]
- Restore operation might fail if the remote reverse SSH tunnel port on an AWS cloud node is in use. [[ENG-16225](#)]
- Stats aggregator process can generate a large amount of login error messages and events. [[ENG-16535](#)]
- Attempting to restore a container from backup after the container was deleted might cause the cerebro process to fail. [[ENG-16588](#)]
- Clusters running the KVM hypervisor might rarely experience an intermittent hung operation state during an upgrade. [[ENG-17195](#)]
- Using the ncli domain-join command does not work when attempting to join a NOS Hyper-V cluster to a Windows Server 2008 Active Directory domain. [[ENG-17601](#)]
- Attempts to create virtual machine might fail due to incorrect path in environment. [[ENG-17960](#)]

The following critical issues were resolved in NOS 4.0:

- More detail required for disk status. [[ENG-8453](#)]
- A large number of timeouts resulting from continuous ping commands during an upgrade might cause VMs to disconnect. [[ENG-9996](#)]
- VM migration can stall if the datastore is intermittently unavailable. [[ENG-10779](#)]

- Alert manager might crash if the configured SNMP server fails to respond to transmitted traps. [ENG-10903]
- Cluster upgrade stalls if Service VM has insufficient free space on its home directory. [ENG-11446]
- Incorrect Cassandra ring calculation based on misconfiguration can affect performance. [ENG-11522]
- A large number of oplog records might affect disk performance, [ENG-11662, ENG-12448]
- Imprecise Cassandra ring calculation can affect performance. [ENG-11862]
- Throughput is slow when the I/O size is small and the outstanding I/O operations is low. [ENG-12046]
- Broken Cassandra ring can disrupt VM network communication. [ENG-12161]
- Privileges assigned to an Active Directory Organizational Unit (OU) by using the authconfig command are inherited by and apply to groups and users in that OU. [ENG-12376]
- Multiple disks can go offline when there are many extent group delete operations, which can cause more high priority operations to wait in the queue for too long (triggering disks to go offline). [ENG-12486]
- The ipmitool can hang indefinitely on some NX-2000 systems. [ENG-12725]
- Smart Support emails can be blocked by customer SMTP servers. [ENG-12731]
- The default Frequency setting in alert policies should be modified for critical alert notifications [ENG-12740]
- Multiple disks can go offline when there are many extent group delete operations, which can cause more high priority operations to wait in the queue for too long (triggering disks to go offline). [12486]
- Web console displays deduplication alerts when deduplication is disabled. [ENG-13062]
- A hung stargate process kernel thread might affect performance. [ENG-13148]
- Cassandra stops and restarts when encountering a corrupt sstable file. [ENG-13433]
- ipmitool might show imprecise NX-2000 power supply status and generate false alerts. [ENG-13805]
- A replication operation can use maximum network bandwidth if Max Bandwidth setting is not specified when creating a remote site. [ENG-13881]
- Hyper-V did not support double-byte languages. [ENG-13981]
- Service VM does not shut down as part of upgrade. [ENG-13998]
- The Advanced Setting Enable Compression selection is cleared and disabled when updating a compression enabled container.[ENG-14011]
- Cassandra health monitoring processes can fail during periods of high communications traffic timeouts. [ENG-14083]
- Cluster might exhibit stability issues if it is manually created with an in-use IP address. [ENG-14316]
- Guest VM data becomes unavailable when the cluster is under high load of small I/O requests. [ENG-14413]
- Nutanix configurations do not accept equal-signs in passwords stored by zeus. [ENG-14447]
- Guest VM data becomes unavailable while running IOmeter tests that use atypically small I/O size [ENG-14460]
- boot_disk_replace script fails to copy configuration information. [ENG-14469]
- Hyper-V setup script setup_hyperv.py fails if a cluster returns multiple IP addresses. [ENG-14496]
- The diagnostics.py script unmounted a datastore under test before all user VM destroy tasks were completed, in some cases resulting in an unresponsive host. [ENG-14737]
- The Entity field might not be populated when creating an Entity or Metric Chart on the Prism Analysis dashboard. [ENG-14806]
- The disk_replace script fails when no disk controller is found. [ENG-14918]
- NOS needs to further verify jumbo frames configuration when using Hyper-V. [ENG-14954]
- NFS snapshots may fail if a large number of snapshot delta files cannot be deleted before snapshot operations complete. [ENG-15209]
- Guest VM data becomes unavailable when the cluster is under high load of small I/O requests. [ENG-15399]
- Cluster unavailability occurs during upgrade when node UUID is incorrectly configured. [ENG-15536]
- Management Server IP address incorrectly reported during add node or upgrade operations. [ENG-15541]

- Suboptimal guest VM performance results from Cassandra consuming excess memory on one node in a cluster. [ENG-15776]
- Safari browser users might be unable to browse to a cluster or perform cluster configuration (cluster_init) tasks with IPv6 addressing enabled. [ENG-15795]
- Nutanix HA might fail if an ESXi operation deletes the autopath route. [ENG-15880]
- Pre-allocated file space can increase after a snapshot operation. [ENG-16035]
- Excessive number of warning messages might affect stargate performance as logs are continuously flushed. [ENG-16056]
- Upgrade to latest version of OpenSSL. [ENG-16177]
- Incorrect deletion of a protection domain can affect data availability. [ENG-16550]
- A cluster might create a large number of log entries in the vCenter database after an upgrade. [ENG-16967]

Notes and Cautions

Nutanix Core and Management Interfaces

The following items apply to all Nutanix clusters.

- Disaster recovery (DR) support has been verified for up to 50 VMs per protection domain and the following replication topologies: 1-to-1, 1-to-2, 2-to-1, and bidirectional 1-to-1. The following topologies have not been tested: 1-to-many (more than 2), many-to-1, and many-to-many. Check with Nutanix technical support before implementing one of the untested topologies or more than 50 VMs in a protection domain.
- When using a remote site for replication, DR requires that the source site run the same or higher Nutanix OS (NOS) release as the remote (target) site. For example, if the remote site is upgraded to release 3.5, the source site must also be upgraded to 3.5 for DR to function properly. (DR can function properly when the remote site is at a lower release than the source site.) [ENG-11059]
- Do not take a snapshot of the Controller VM. Doing so can cause the Controller VM to stop functioning and compromise data integrity. The cluster is designed to handle failures without the use of VM snapshots. [ENG-12442]
- Statistics are not gathered during upgrade to a new version of NOS. Graphs on the home page of the web console have gaps that correspond to the time of upgrade. [ENG-16576]
- When users try to log on to the web console and are implementing valid Common Access Card (CAC) certificates, Microsoft Internet Explorer versions 10 and 11 might require additional steps for access. Nutanix Knowledge Base article 1549 describes this browser issue. [ENG-18624]
- syslog messages from Cassandra are truncated by the remote syslog server if the messages are longer than 32 characters. [ENG-18758]
- After performing an upgrade to NOS 4.0, also install Nutanix Cluster Check (NCC) version 1.1 on any controller VM. [ENG-18622]

vSphere Hypervisor

The following items apply only if the Nutanix cluster is running vSphere.



Warning: The Nutanix *Upgrade Guide* contains instructions to upgrade to vSphere 5.1. It is critical to follow these instructions as unique steps are required to ensure proper functioning of the cluster after upgrade.

- To prevent the ESXi host from running out of memory, in the **Resources > Memory** configuration of the Controller VM ensure that **Reserve all guest memory** is selected. [ENG-7469, ENG-9031]
- (vSphere 5.1 only) APD must be disabled for NFS on every ESXi host in the cluster by setting **Advanced settings > Misc > APDHandlingEnable = 0** in vCenter. [ENG-7056]

- (vSphere 5.1 NX-2000 only) If the multiextent module is not loaded, the Controller VM will not start. [ENG-7475]
- If a 1 GbE interface is connected, ESXi may prefer it over a 10 GbE interface leading to poor performance. To avoid this situation, delete the 1 GbE interface from vSwitch properties. Setting the 1 GbE interface to **Unused** does not always prevent it from being used.

For more information about this known issue, see VMware KB articles [2030006](#) and [2008144](#). Nutanix has observed this condition in both vSphere 5.0 and 5.1. [ENG-8625, ENG-11495]

- If an ESXi datastore is inactive, the Nutanix user interface (web console or nCLI) will not be able to report the host ID or any hypervisor vdisk information. [ENG-10900]
- Do not upgrade VMware Tools on Nutanix Controller VMs. [ENG-7770]
- Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM) support has been verified for 50 VMs per vStore and two vStores for a VMware protection group. Larger configurations have not been tested, and only a 1-to-1 mapping topology has been tested.

Hyper-V Hypervisor

The following items apply only if the Nutanix cluster is running Hyper-V.

- Hyper-V hosts must be connected to a 10 GbE switch that has jumbo frames enabled. Cluster creation will fail if the switch does not support jumbo frames or does not have jumbo frames enabled.
- If you are upgrading from NOS version 3.5.2, as a final step to complete the upgrade, log on any controller VM in your cluster and run the `setup_hyperv.py setup_hosts` script. [ENG-18486]
- The Hyper-V extensible switch is not supported on current versions of NOS. [ENG-18486]
- Data protection items
 - The HA settings of a VM are not retained when it is migrated to a remote site. These settings must be manually configured in SCVMM after migration. [ENG-15545]
 - After restoring a snapshot for a protection domain, perform a **Refresh Virtual Machines** on the hosts in SCVMM to ensure that the correct VMs are reported. [ENG-15194]
 - Only one restore snapshot operation can be performed at a time. Wait for the operation to complete before starting another. [ENG-16071]
 - VMs in **Paused** state do not retain this state during migration or failover. On the recovery site paused VMs will be in **Powered Off** state. [ENG-16024]
 - VMs in **Save** state cannot be migrated. [ENG-16024]
 - If multiple disks in different paths are attached to a VM, restoring the VM from snapshot fails. [ENG-17141]
 - When configuring and enabling remote site replication, all local and remote Hyper-V hosts must be connected to a switch that has jumbo frames enabled. [ENG-18179]
 - NOS data protection does not support protection of Hyper-V check-pointed virtual machines. [ENG-18324]



Warning: Nutanix native snapshots and replication is not supported for VMs configured with Hyper-V host snapshots. If such VMs are configured with Nutanix native snapshot or replication, unexpected snapshot behavior could result.

KVM Hypervisor

The following items apply only if the Nutanix cluster is running KVM.

- Guest VMs may hang or experience slow performance with emulated devices. Specify the `--paravirt` option when creating guest VMs with `virt_install` to use paravirtualized devices. The `virt_install` script installs emulated devices by default. Paravirtualized device drivers for Windows are available at http://www.linux-kvm.org/page/WindowsGuestDrivers/Download_Drivers. [ENG-12944]
- To use the VM management scripts provided by Nutanix, the cluster must have a storage pool and a container named `default`.

- Shared vdisks (`vdisk create shared` command) are not supported for systems running KVM. [ENG-8950]
- Local backup and remote replication are not supported for systems running KVM. [ENG-8906]
- A DMA read error might occur when rebooting the Controller VM on a system running KVM. If this occurs, reboot the host. [ENG-9065]
- VM memory usage reported in the web console might be higher than what is reported from the `top` command. This can occur because the web console displays the highest memory usage it has seen while the `top` command displays the current memory usage. [ENG-10811]

Hardware

- (NX-1000/2000/3050/6000) With certain switches, connection instability of the dedicated IPMI network interface causes the hypervisor host to crash and fail to start until IPMI is disconnected. [ENG-10732]
- (NX-1000/2000/3050/6000) If a Nutanix cluster does not include node 1 of a block, the cluster will not generate power supply alerts. [ENG-10194]
- (NX-2000 only) The CPU temperature alert is not supported. [ENG-10215]
- (NX-3000 only) After turning on a node, the hard drive LEDs flash red, amber, and green until the Controller VM starts (up to 5 minutes). [ENG-7372]
- (NX-2000 only) Redundancy factor 3 is not supported on nodes that have less than 300 GB of SSD space. NX-2000 systems with Fusion-io ioDrive PCIe-SSDs do not have adequate space to support redundancy factor 3. NX-2000 systems with Fusion-io ioDrive2 or Intel PCIe-SSDs do have adequate space to support redundancy factor 3. [ENG-13418]

Supported Versions

Third-Party Software

Component	Version
VMware ESXi	5.0 5.1 5.5
VMware VMFS	5.0
Windows Server	2012 R2 Datacenter Edition
KVM	CentOS 6.4 2.6.32-358.6.2.el6.x86_64
Controller VM	CentOS 6.3 2.6.32-279.9.1.el6.nutanix.x86_64
Diagnostic VM	CentOS 6.4 2.6.32-358.6.1.el6.x86_64

Hypervisor Versions

Hypervisor support varies by Nutanix product.

Product	ESXi 5.0U3	ESXi 5.1U2	ESXi 5.5	KVM	Hyper-V ¹
NX-1000	•	•	•	•	•
NX-2000	•	•	•		
NX-3000	•	•	•	•	

Product	ESXi 5.0U3	ESXi 5.1U2	ESXi 5.5	KVM	Hyper-V ¹
NX-3050	•	•	•	•	•
NX-6000	•	•	•	•	•
NX-7000		•	•		

1. Hyper-V is restricted to nodes that have 64 GB hypervisor boot drives.

Nutanix Software

Component	Version
Installer	danube-4.0.1-stable
Core	danube-4.0.1-stable
Serviceability	danube-4.0.1-stable
Toolchain	danube-4.0.1-stable
Diagnostics	danube-4.0.1-stable

Diagnostics Expected Results

The following test results are typical for a four-node cluster running ESX 5.1 Update 2, except where noted as ESX 5.5 Update 1 (indicated as 5.5). Smaller clusters will have lower results, and larger clusters will have higher results. These numbers are for release 4.0.1.

The results from diagnostics are expected to be within 7% of the number given in this table.



Note: Results for Model NX-3550 were achieved with replication factor of 3. All other model results were achieved with replication factor of 2.

Metric	NX-6000	NX-3050	NX-2000	NX-1000
Sequential write bandwidth (MBps)	845 (6220) 767 (6220, 5.5) 1325 (6260), 1297 (6260, 5.5)	1318 (3450), 1443 (3460), 1127 (3550)	890 (2400), 955 (2450)	368 (1420), 372 (1420, 5.5) 373 (1450 1GbE), 760 (1450 10GbE), 745 (1450 10GbE, 5.5)
Sequential read bandwidth (MBps)	1905 (6220) 1902 (6220, 5.5), 3612 (6260), 3630 (6260, 5.5)	3655 (3450), 3634 (3460), 4548 (3550)	3002 (2400), 3530 (2450)	1903 (1420), 1909 (1420, 5.5) 1891 (1450 1GbE), 1899 (1450 10GbE), 1900 (1450 10GbE, 5.5)

Metric	NX-6000	NX-3050	NX-2000	NX-1000
Random read IOPS	58523 (6220), 59005 (6220, 5.5), 106040 (6260), 107030 (6260, 5.5)	104027 (3450), 108973 (3460), 130384 (3550)	50211 (2400), 59278 (2450)	55267 (1420), 55491 (1420, 5.5) 45002 (1450 1GbE), 44998 (1450 10GbE), 45516 (1450 10GbE, 5.5)
Random write IOPS	55061 (6220), 59794 (6220, 5.5), 75556 (6260), 74016 (6260, 5.5)	71457 (3450), 78373 (3460), 74543 (3550)	35778 (2400), 50557 (2450)	43061 (1420), 42674 (1420, 5.5) 46167 (1450 1GbE), 49542 (1450 10GbE), 50689 (1450 10GbE, 5.5)