

Lenovo ThinkSystem NE1032/NE1032T/NE1072T RackSwitch

Release Notes

For Lenovo Cloud Network Operating System 10.9

LenovoTM

Note: Before using this information and the product it supports, read the general information in the *Safety information and Environmental Notices* and *User Guide* documents on the *Lenovo Documentation CD* and the *Warranty Information* document that comes with the product.

Second Edition (December 2018)

© Copyright Lenovo 2018
Portions © Copyright IBM Corporation 2014

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Lenovo and the Lenovo logo are trademarks of Lenovo in the United States, other countries, or both.

Release Notes

This release supplement provides the latest information regarding Lenovo Cloud Network Operating System 10.9 for the Lenovo ThinkSystem NE1032/NE1032T/NE1072T (referred to as NE1032/NE1032T/NE1072T throughout this document).

This supplement modifies and extends the following Cloud NOS documentation for use with CNOS 10.9:

- *Lenovo Network Application Guide for Lenovo Cloud Network Operating System 10.9*
- *Lenovo Network Command Reference for Lenovo Cloud Network Operating System 10.9*
- *Lenovo Network Python Programming Guide for Lenovo Cloud Network Operating System 10.9*
- *Lenovo Network REST API Programming Guide for Lenovo Cloud Network Operating System 10.9*
- *Lenovo ThinkSystem NE1032/NE1032T/NE1072T Installation Guide for Lenovo Cloud Network Operating System*

These publications are available from the following website:

http://systemx.lenovofiles.com/help/topic/com.lenovo.systemx.common.nav.doc/overview_rack_switches.html

Please keep these release notes with your product manuals.

Note: The Lenovo Cloud Network OS is based on the Embedded Linux Integration Environment (ELIE). To obtain open source code licenses, go to <https://github.com/lenovo/ELIE/tree/master/elic-1.7.1/licenses/>. For details on how to obtain open source code, please contact Lenovo Support.

Hardware Support

CNOS 10.9 software is supported on the NE1032/NE1032T/NE1072T high performance Layer 2-3 network switches.


The NE1032/NE1032T/NE1072T is a 1U GbE switch and it can be mounted horizontally or vertically, depending on your application. Mounting options are available for a variety of rack systems.

For superior reliability, the NE1032/NE1032T/NE1072T uses redundant, hot-swap power supply modules and three (NE1032 and NE1032T) or five (NE1072T) hot-swap fan modules. Module options are available for either front-to-rear airflow, or rear-to-front airflow.

The NE1032 contains the following ethernet ports:

- Thirty-two 10 Gigabit Ethernet (GbE) Small Form Pluggable Plus (SFP+) ports which also support legacy 1 GbE connections


Figure 1. NE1032 front panel



The NE1032T contains the following ethernet ports:

- Twenty-four 1G/10G BASE-T RJ45 ports
- Eight 10G SFP+ Uplinks
- Ten Gigabit Ethernet (GbE) Small Form Pluggable Plus (SFP+) ports which also support legacy 1 GbE connections


Figure 2. NE1032T front panel



The NE1072T contains the following ethernet ports:

- Forty-eight 1G/10G BASE-T RJ45 ports
- Six 40 GbE Quad Small Form Pluggable Plus (QSFP+) ports, all of which ports can optionally be used as four 10 GbE ports

Figure 3. NE1072T front panel



Supplemental Information

This section provides additional information about configuring and operating the NE1032/NE1032T/NE1072T and CNOS.

The Boot Management Menu

The Boot Management menu allows you to switch the software image or to recover from a failed software download.

You can interrupt the boot process and enter the Boot Management menu from the serial console port. When the system displays Memory Test, press **Shift + B**. The Boot Management menu appears.

```
The system is going down for reboot NOW!
INIT: reboot: Restarting system
...

Press shift-B for startup menu or shift-R for recovery mode: ..
Running Startup Menu
...

Boot Management Menu
  I - Change booting image
  C - Change configuration to factory default
  R - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
  P - Reset the Network Administrator (admin) password
  B - Reset the password required to enter privileged exec mode
  Q - Reboot
  E - Exit
Please choose your menu option:
```

The Boot Management menu allows you to perform the following actions:

- To change the booting image, press **I** and follow the screen prompts.
- To reset the switch configuration to factory defaults, press **C** and follow the screen prompts.
- To boot in recovery mode, press **R**.
- To reset the Network Administrator account (admin) password, press **P** and follow the screen prompts.
- To reset the password required to enter Privileged EXEC configuration mode, press **B** and follow the screen prompts.
- To reload the switch, press **Q**. The reloading process will start again.
- To exit the Boot Management menu, press **E**. The reloading process continues.

Recovering from a Failed Software Upgrade

Use the following procedure to recover from a failed software upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports TFTP (preferred) or Xmodem download (for example, HyperTerminal, SecureCRT, PuTTY). If using Xmodem, make sure the following settings are in effect:
 - Speed: 9,600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None
3. To access the Rescue Mode menu, you must interrupt the boot process from the Console port. Boot the NE1032/NE1032T/NE1072T, and when the system begins displaying Memory Test progress (a series of dots), press **Shift + R**. The Rescue Mode menu will be displayed:

```
Entering Rescue Mode.
Please select one of the following options:
  T) Configure networking and tftp download an image
  U) Install image from USB stick
  P) Physical presence test (low security mode)
  F) Run filesystem check
  I) Select which image to boot
  C) Reset configuration to factory default
  R) Reboot
  E) Exit

Option?:
```

- If you choose option **T** (TFTP download), go to step [4](#).

4. **TFTP download:** The switch prompts you to enter the following information:

```
Performing TFTP rescue. Please answer the following questions (enter 'q'
to quit):
IP addr   :
Server addr:
Netmask   :
Gateway   :
Image Filename:
```

- a. Enter the required information and press **Enter**.
- b. You will see a display similar to the following:

```
Host IP    : 10.10.98.110
Server IP  : 10.10.98.100
Netmask    : 255.255.255.0
Broadcast  : 10.10.98.255
Gateway    : 10.10.98.254
Installing image NE1032/NE1032T/NE1072T-CNOS-10.9.2.0_OS.imgs from
TFTP server 10.10.98.100
```

- c. When you see the following prompt, enter the image number where you want to install the new software and press **Enter**.

```
Install image as image 1 or 2 (hit return to just boot image): 1
```

- d. The following message is displayed when the image download is complete. Continue to step 7.

```
Entering Rescue Mode.
Please select one of the following options:
  T) Configure networking and tftp download an image
  U) Install image from USB stick
  P) Physical presence test (low security mode)
  F) Run filesystem check
  I) Select which image to boot
  C) Reset configuration to factory default
  R) Reboot
  E) Exit

Option?:
```

5. Image recovery is complete. Perform one of the following steps:

- Press **R** to reboot the switch
- Press **E** to exit the Boot Management menu
- Press **Esc** to re-display the Boot Management menu

New Features in This Release

This release of Lenovo Cloud Network OS contains the following significant enhancements.

Management ACLs

Management ACLs allow you to add an access list to incoming VTY (Virtual Teletype) lines. This way, you can control who can access the switch. The Management ACL attached to VTY lines applies to all the packet destined to the switch CPU, coming from interfaces that are part of VRFs where MACLS are attached.

VRF Enhancement

This feature adds support to increase non-management VRF (VPN Routing/Forwarding instance) from 8 to 64.

Converged Enhanced Ethernet (CEE)

Converged Enhanced Ethernet (CEE) refers to a set of IEEE standards developed primarily to enable FCoE, requiring enhancing the existing Ethernet standards to make them lossless on a per-priority traffic basis, and providing a mechanism to carry converged (LAN/SAN/IPC) traffic on a single physical link. CEE features can also be used in traditional LAN (non-FCoE) networks to provide lossless guarantees on a per-priority basis, and to provide efficient bandwidth allocation.

The CEE capabilities include:

- Priority-based Flow Control (PFC) as per IEEE 802.1Qbb
- Enhanced Transmission Selection (ETS) as per IEEE 802.1Qaz
- Data Center Bridging Capability Exchange (DCBX) - CEE DCBX(1.01) & IEEE DCBX (802.1Qaz).

Note: This feature is supported on NE1032T and NE1072T.

SLP over non-default VRF support

This enhancement added support for SLP over non-default VRFs.

Policy-Based Routing (PBR)

PBR allows you to route traffic based on defined policies rather than entries in the routing table. Such policies are defined based on the protocol, source IP, or other information present in a packet header. PBR provides a mechanism for applying the defined policies based on Access Control Lists (ACLs). The route-map must be assigned on a Layer 3 interface (routed port or SVI).

BGP EVPN

DCI MP-BGP L2VPN interoperability with other vendors is now supported.

MP-BGP EVPN now supports Type 3 – Inclusive Multicast Route. Type 3 routes are required for Broadcast, Unknown Unicast and Multicast (BUM) traffic delivery across EVPN networks. Type 3 advertisements provide information about P-tunnels used to send BUM traffic.

Known Issues

This section describes known issues for CNOS 10.9 on the Lenovo ThinkSystem NE1032/NE1032T/NE1072T RackSwitch.

Note: Please check the Change History documentation posted with the Switch Firmware to check if any of these issues have been fixed in the latest release.

BGP

When the switch detects that it cannot reach the next-hop in a route, the BGP instance does not deactivate the route from its routing table. The route is deactivated only when the BGP session with the next-hop times out. (ID: 135910)

Copying Configuration Files

When copying a new configuration file over the switch's current running configuration, CNOS does not overwrite the old configuration. Instead, it appends the new configuration over the old one, leaving settings that are not present in the new configuration intact. It only overwrites settings present in both configurations. (ID: 133055)

LACP

The default behavior of Link Aggregation Control Protocol (LACP) Individual has been changed. Now, by default, LACP enabled ports transition to suspend.

When upgrading the firmware to CNOS 10.9 from 10.7 or an earlier version, the following LACP behavior occurs:

- For Link Aggregation Groups (LAGs) with default configuration, before the upgrade, LACP enabled ports transition to individual state when not receiving Link Aggregation Control Protocol Data Units (LACPDU). After the firmware upgrade, the LACP configuration remains the same, but LACP enabled ports now transition to suspended state when no LACPDU are received.

To configure the LACP enabled ports as individual, use the following command on a LAG:

```
NE1032(config)# interface port-channel <LAG number (1-4096)>
NE1032(config-if)# no lACP suspend-individual
```

- For LAGs with non-default configuration, if before the upgrade LACP enabled ports transition to suspended state when not receiving LACPDU, then after the firmware upgrade this behavior remains the same. However, the LACP configuration is changed to reflect the new default behavior: the LACP individual setting for LAGs is removed from the switch's running configuration. (ID: 132111)

Network Virtualization Gateway

The following limitations exist:

- For optimal performance, we recommend that the number of VLAN-VXLAN Network Identifier (VNI) mappings does not exceed 1,000 entries. Going above this limit leads to longer convergence times when attaching or detaching the Lenovo hardware Layer 2 gateway to or from NSX logical switches. (ID: 99467)
- For optimal Equal Cost Multiple Paths (ECMP) load balancing, we recommend that only Layer 3 routed ports are used for connecting to spine switches. (ID: 123627)
- When broadcast, unknown unicast, and multicast (BUM) traffic is received on the switch, it is replicated on all member ports of the same VXLAN Network, except the port that is the source port of the BUM traffic. This is displayed in the source port statistics as dropped packets. (ID: 95658)
- Throughput statistics of the southbound interface of the VXLAN gateway do not display on the NSX GUI. (ID: 113832)
- After Bidirectional Forwarding Detection (BFD) failover, when using the default value for the probe interval of 300 ms, the active service node election takes approximately three seconds to occur. (ID: 116882)
- The VLAN used in VLAN-VXLAN Network Identifier (VNI) mapping must be used exclusively for switching within the associated VNI domain. Different access vPorts belonging to the same VLAN must be mapped to the same VNI. Hence, only one-to-one VLAN-VNI bindings are supported. (ID: 100606, 123143)
- In case all network ports go down on one of the vLAG switches, all ingress traffic received on the local access ports is flooded to all other local access ports from the same VXLAN Network Identifier (VNI). Flooding stops when the network ports are back up. (ID: 110732)
- In High Availability (HA) mode, the recommended maximum number of local unicast MAC addresses is 32,000. If this limit is exceeded, MAC address synchronization between the vLAG switches might not work properly. More than 32,000 unicast MAC addresses can be used, but the synchronization process fails to function normally. (ID: 113145)
- When using NSX manager 6.4.1, the remote FDB (forwarding database) table is not correctly updated on the remote gateways, when changing the PVID or the access VLAN on the local access port. This issue has been fixed starting with NSX Manager 6.4.2. (ID: 159591)

REST API

VMware vSphere Distributed Switch (vDS) information retrieval is not supported using REST API. (ID: 138329)

SolarWinds

After setting up a SFTP/SCP server using SolarWinds, the switch fails most times to establish a connection to the server. We recommend using other SFTP/SCP servers. (ID: 135539)

Telemetry

Sometimes, when the egress CPU queue is full, Buffer Statistics Tracking (BST) gets disabled and telemetry messages are no longer sent to remote monitoring servers. Use REST to re-enable BST. (ID: 139654)

vLAG

- When upgrading the firmware of vLAG peer switches one by one to CNOS 10.9 from 10.7 or an earlier version, vLAG configuration consistency checks fail on both vLAG peer switches until the upgrade process is complete for both. (ID: 131917)

vNIC Statistics

When rebooting the switch, vNIC statistics can be quickly retrieved by toggling vnic-stats state. (ID: 131081)

QLogic Firmware

QLogic QL41262 does not establish 25G link with passive SFP28 DAC cable when QL41262 is set to auto negotiation. (ID: 148199)

Management Port Traffic

When upgrading from 10.8 to 10.9, make sure both the software image and the boot image are updated. Otherwise, all management port traffic will be affected. (ID: 143713)

Lenovo HCI Solution with Nutanix

ThinkAgile Network Orchestrator is broken on Nutanix Prism v5.9.1 release. You must have Nutanix AOS version 5.0.2 – 5.8 or 5.9.2 – 5.10 installed for the ThinkAgile Network Orchestrator to work properly.

Logging

- When establishing a second SSH or Telnet session with the switch, terminal logging may not properly initialize for the session. (ID: 158589) The following message appears:

```
NE1032(cnos:default) %IMI-6-TTY_LOGGING_INIT_IGNORE: Terminal logging
initialization ignored due to log service busy (cannot obtain mutex
lock), no log messages will forward to the corresponding terminal
(/dev/pts/0)
```

- If the logging service is busy, syslog messages are cached and logged once the service is free. In this situation, the log message timestamp is not sequenced in the log file or on the terminal. (ID: 158588)

DCI MP-BGP L2VPN

- BGP-EVPN is not backwards compatible with earlier CNOS releases. (ID: 157557)
- Juniper VXLAN tunnels are not removed if Layer 3 connectivity is lost. However, they restore correctly when Layer 3 connectivity is recovered. In some rare scaling scenarios, this can cause the maximum number of tunnels to be exhausted. (ID: 158207)
- Tunnels to other vendors will remain in the up state if BGP protocol/neighbors are shut down manually. (ID: 158106) Each time BGP/BGP neighbor is shut down the following message appears:

Losing BGP evpn connectivity while DCI BGP evpn is enabled will result in traffic loss.

Note: In these situations, disabling/enabling DCI mode BGP EVPN deletes the inactive tunnels.

- HW clean-up messages are not displayed when the NWV mode is disabled on the newly elected vLAG primary. (ID: 156759)
- Some internal messages are duplicated in certain scenarios. This is not causing any functional issues. (ID: 157434)
- If the same MAC/IP is used on different VNIs, the output of the **show ip bgp l2vpn evpn mac/ip** command only displays one BGP route. Using the full command, for example: **show ip bgp l2vpn evpn ip 10.100.0.1 virtual-network 101 rd 100.100.1.5:101** displays the other routes. (ID: 158222)