

Lenovo ThinkSystem NE2572 RackSwitch

Release Notes

For Lenovo Cloud Network Operating System 10.9

Lenovo[™]

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 NE2572 (referred to as NE2572 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 NE2572 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 NE2572 high performance Layer 2-3 network switches.


The NE2572 is 1U in height and can be mounted horizontally or vertically, depending on your application. Mounting options are available for a variety of rack systems.

For superior reliability, the NE2572 uses redundant, hot-swap power supply modules and hot-swap fan modules. Module options are available for either front-to-rear airflow or rear-to-front airflow.

The NE2572 contains the following ethernet ports:

- Forty-eight 25 GbE Small Form Pluggable 28 (SFP28) ports
- Six 100 GbE Quad Small Form Pluggable 28 (QSFP28) ports - each QSFP28 port can optionally be used as four 25 GbE SFP+ ports

Figure 1. NE2572 front panel



Supplemental Information

This section provides additional information about configuring and operating the NE2572 and CNOS.

The BIOS Menu

The Basic Input/Output System (BIOS) menu allows you to have complete system control at boot.

You can interrupt the startup process of the switch and enter the BIOS menu from the serial console port. When the system displays the following message, press **Delete** or **Esc**.

```
Press <DEL> (Terminal Not applicable) or <ESC> to enter setup...
```

The BIOS menu appears.

```
Aptio Setup Utility - Copyright (C) 2017 American Megatrends, Inc.
Main Advanced Event Logs Security Boot Save & Exit
-----
| BIOS Information                                     |Choose the system
| BIOS Vendor          American Megatrends          |default language
| Core Version         5.009
| Compliancy           UEFI 2.3; PI 1.2
| Project Version      0ACBZ 0.33 x64
| Build Date and Time  05/10/2017 11:09:14
|
| Memory Information
| Total Memory        8192 MB (DDR3)
|
| OEM Version         ALPHA.05.33.0206B
|
| System Language     [English]
|
| System Date         [Thu 05/11/2017]
| System Time         [11:13:37]
|
| Access Level        Administrator
|
|-----|
|>: Select Screen
|^v: Select Item
|Enter: Select
|+/-: Change Opt.
|F1: General Help
|F2: Previous Values
|F3: Optimized Defaults
|F4: Save & Exit
|ESC: Exit
|-----|
Version 2.17.1245. Copyright (C) 2017 American Megatrends, Inc.
```

This menu permits the following actions:

- Monitoring system configuration
- Setting user passwords
- Switching to Secure Boot Mode
- Performing key provisioning

The Grub Menu

The Grub menu allows you to switch the software image. The menu appears on the screen automatically during the switch startup process.

```
Welcome to GRUB!

                                GNU GRUB  version 2.00

+-----+
|CNOS slot 1
|CNOS slot 2
|Recovery Mode
|ONIE
|
|
+-----+

Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, `e' to edit the commands
before booting or `c' for a command-line. ESC to return
previous menu.
```

Note: For more information on ONIE, please see the *Lenovo ThinkSystem NE10032 RackSwitch ONIE User Guide*.

Rescue Mode

The Rescue Mode option allows you to recover from a failed firmware or boot image upgrade using TFTP or a USB drive.

To enter Rescue Mode, select **Recovery Mode** in the GRUB menu. The following menu appears.

```
Entering Rescue Mode.
Please select one of the following options:
  T) Configure networking and tftp download an image
  U) Install image from USB stick
  F) Run filesystem check
  I) Select which image to boot
  C) Reset configuration to factory default
  Z) Reset the Network Administrator (admin) password
  B) Reset the password required to enter privileged exec mode
  R) Reboot
  E) Exit

Option? :
```

The Rescue Mode menu allows you to perform the following actions:

- To recover from a failed software or boot image upgrade using TFTP, press **T** and follow the screen prompts.
- To recover using an image from a USB stick, press **U**.

- To check if the switch is ready to run Cloud NOS, press **F**. It performs a check to see if the filesystem is optimally partitioned and updates it accordingly.
- To select which firmware image to boot, press **I**.
- To reset the switch configuration to factory defaults, press **C**.
- To reset the Network Administrator account (admin) password, press **Z**.
- To reset the password required to enter Privileged EXEC mode, press **B**.
- To restart the reload process from the beginning, press **R**.
- To exit the Rescue Mode menu, press **E**.

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 MACs are attached.

VRF Enhancement

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

Support 1Gb network connectivity on the 25G (SFP28) ports

This feature adds support for 1G (SX, LX, Cu) SFP device over the native 25G ports; requires ports to be configured in 10G mode.

Flexible speed types support

This feature removes previous restriction and allows flexible support of all port speeds, including 10G, 25G, 40G, 50G, and 100G.

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 NE2572 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:

```
NE2572(config)# interface port-channel <LAG number (1-4096)>  
NE2572(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)

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)
- To ensure no traffic is lost while upgrading a vLAG topology to CNOS 10.9 from an earlier version, you must enable backwards compatibility (ID: 146411). To do so, follow these steps:

Note: The **show vlag information** command displays the configured and operational role for each VLAG peer. The following steps assume that the Primary/Secondary terms refer to the operational role.

1. Having an earlier than CNOS 10.9 version on both vLAG switches, load the CNOS 10.9 image on both of them.

Note: Do not reboot the switch.

2. On the Primary switch, shutdown all the port-channels associated with the vLAG instances. All traffic goes through the Secondary switch.

Note: Do not save the configuration with the disabled port-channels.

3. Reload the Primary switch so that the current Secondary switch becomes the new Primary.

Note: The operational VLAG roles are swapped after this step.

4. After reboot, the auto-recovery timer starts on the switch but the ISL remains in the `Inactive` state. Use the **ecp compatibility-mode** command on this switch so that the ISL becomes `Active` and the vLAG instances are brought up (after the startup delay timer expires).

5. A syslog notification appears alerting the user that the vLAG OS version is mismatched, but vLAG is operational with all its instances in the `Formed` state.

6. On the Primary switch, shutdown all the port-channels associated with the vLAG instances. All traffic goes through the Secondary switch.

Note: Do not save the configuration with the disabled port-channels.

7. Reload the Primary switch so that the current Secondary switch becomes the new Primary.

Note: The operational VLAG roles are swapped after this step (the initial roles from the beginning of this procedure are restored).

8. After reboot, the auto-recovery timer starts on the switch, but the ISL remains in the `Inactive` state. Use the **no ecp compatibility-mode** command on the other switch so that the ISL becomes `Active` and the vLAG instances are brought up (after the startup delay timer expires).

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)

1G CuSFP Transceivers

Due to hardware limitations of the switch ASIC, when using a 1 GbE CuSFP transceiver, the link state change can be detected with a delay of 2-3 seconds. Also, during reload, you may see a temporary link up state even though the link is down configuration wise. This may have an impact to link failovers with this type of transceiver. The switch stabilizes and resumes under normal operation. (ID: 147546, 147507, 145840, 149018)

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:

```
NE2572(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)

