

Lenovo ThinkSystem NE2572 RackSwitch

Release Notes

For Lenovo Cloud Network Operating System 10.10

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.

First Edition (April 2019)

© Copyright Lenovo 2019
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 10.10 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 10.10:

- *Lenovo Network Application Guide for Lenovo Cloud Network Operating System 10.10*
- *Lenovo Network Command Reference for Lenovo Cloud Network Operating System 10.10*
- *Lenovo Network Python Programming Guide for Lenovo Cloud Network Operating System 10.10*
- *Lenovo Network REST API Programming Guide for Lenovo Cloud Network Operating System 10.10*
- *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/entie-1.7.1/licenses/>. For details on how to obtain open source code, please contact Lenovo Support.

Hardware Support

CNOS 10.10 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



Updating vLAG Switches

Following are the steps for updating the software and boot images for switches configured with vLAG:

1. Having an earlier than CNOS 10.10 version on both vLAG switches, load the CNOS 10.10 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, ISL automatically enters in the **Active** state.

5. A syslog notification appears alerting the user that the vLAG OS version is mismatched, but vLAG becomes operational with all its instances in the **Formed** state, after startup-delay expires.

6. On the current 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 current 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, ISL automatically enters in the **Active** state, and vLAG becomes operational with all its instances in the **Formed** state, after startup-delay expires.

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. This menu permits the following actions:

```
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              | ><: Select Screen  |
| System Language     [English]                     | ^v: Select Item   |
|
| System Date         [Thu 05/11/2017]              | Enter: Select     |
| System Time         [11:13:37]                    | +/-: Change Opt. |
| Access Level        Administrator                  | F1: General Help  |
|
|                                     | F2: Previous Values |
|                                     | F3: Optimized Defaults |
|                                     | F4: Save & Exit     |
|                                     | ESC: Exit          |
-----
Version 2.17.1245. Copyright (C) 2017 American Megatrends, Inc.
```

- 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 NE2572 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.

Controller-less VXLAN (TNO + VXLAN) on VMWare and Nutanix

This feature integrates controller-less VXLAN (VXLAN bridging without requiring an SDN controller) for VMWare and Nutanix applications with the automated Layer 2 port provisioning provided by ThinkAgile Network Orchestrator.

ARP Suppression

This feature reduces traffic across controller-less VXLAN DC by keeping learning packets local to the VTEP. This is done by avoiding ARP packet flooding, which occupies core network bandwidth.

Open Linux Shell

This feature provides the capability for a CNOS administrator to load and install system applications that provide enhanced functionality that is not part of the standard CNOS distribution.

Linux shell access provides CNOS users with support for installing different applications.

You must install secure applications. For any user-installed applications, scripts or programs, the user is wholly responsible for the sanity and integrity of the system.

Standalone Unified Fabric Port (UFP)

This feature uses vPorts to create isolation between vNICs within the compute node and maintain that isolation within the switch. Unified Fabric Port (UFP) defines an architecture that allows the logical subdivision of a single high speed physical networking link connecting a server Network Interface Controller (NIC) or Network Adapter port to an edge switch port such that each subdivision appears to the server operating system or hypervisor as an independent physical NIC, and each subdivision can be uniquely identified at the edge switch port, which can then be configured to provide the same isolation and functionality as if each logical subdivision were a separate physical connection.

Port Mode Change Without Reboot

This feature provides flexibility of configuration by allowing port speeds to be changed without resetting the switch.

802.1X Port Authentication

802.1X is a standardized framework that provides port-based access control. This protocol authenticates the network clients by using a unique information like personal credentials (username/password) or digital certificates. In 802.1X architecture, the switch is a Network Access Server (NAS), and relies on a central identity/attributes repository to authenticate and authorize the client. The 802.1X standard specifies the Extensible Authentication Protocol (EAP) as a method to provide authentication and authorization using a centrally administered Authentication Server and defines the EAP over LAN (EAPoL) message format.

Boot Profiles

Traditional switches provide a fixed number of forwarding table entries, which are referred to as the size of that respective table. The forwarding tables refer to the Layer 3 Route table, ARP cache table, and MAC address table. The numbers are fixed according to capability of underlying ASIC. The latest ASICs from Broadcom allow these forwarding tables to share a common pool of resources, which in turn allows customizing the share for each of those tables. The Boot Profiles feature leverages capabilities in the Broadcom ASIC to allow forwarding tables to share a common pool of resources. A new profile was added, allowing you to choose between two predefined table sizes.

Telemetry Data Set Extension

New features allow CNOS telemetry to store data using an embedded Redis database. When used, the database can help streamline the REST/Python API processing for pull requests by allowing direct access from the Python API layer to the database.

Static NAT Support

Static Network Address Translation (NAT) provides the capability to configure one-to-one, many-to-few, and many-to-one mappings of the inside local addresses to the inside global addresses. This feature can provide either only IP addresses translations or both IP addresses and UDP/TCP L4 port number translations from inside to outside traffic, and outside to inside traffic.

Static NAT is useful when a host from private network must be accessible by a fixed address from external network.

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) is an effort to converge two of the different physical networks in today's data centers. It allows Fibre Channel traffic (such as that commonly used in Storage Area Networks, or SANs) to be transported without loss over 10Gb Ethernet links (typically used for high-speed Local Area Networks, or LANs). This provides an evolutionary approach toward network consolidation, allowing Fibre Channel equipment and tools to be retained, while leveraging cheap, ubiquitous Ethernet networks for growth.

FCoE Initialization Protocol (FIP) Snooping

FCoE Initialization Protocol (FIP) snooping is a feature of FCoE. In order to enforce point-to-point links for FCoE traffic outside the regular Fibre Channel topology, Ethernet ports used in FCoE can be automatically and dynamically configured with access control lists (ACLs).

Using FIP snooping, the switch examines the FIP frames normally exchanged between the FCF and ENodes to determine information about connected FCoE devices. This information is used to create narrowly tailored ACLs that permit expected FCoE traffic to and from confirmed Fibre Channel nodes, and deny all other undesirable FCoE or FIP traffic.

MAC-move Notification

This feature enables MAC-move notification, which generates syslog messages when the switch detects that a MAC address has moved between switch ports.

MAC-move Loop Detection

This feature enables MAC-move loop detection, which detects MAC-move loops and error-disable ports that appear in the loops.

BIOS Firmware Update

There is a new release of BIOS firmware to version 020BB. This version addresses issues with running the `halt` and `poweroff` commands under a third party NOS. Lenovo CNOS is not impacted. To update your BIOS firmware please follow the instructions in the BIOS update section of the Application Guide.

100G BiDi Transceiver

Added support for 100G BiDi transceiver.

OSPFv2 with VRF

Using VRF together with OSPF enables multiple OSPF instances to be run on the switch.

NTP Server Support

The switch can be configured to act as an NTP server.

Link Flap Dampening

Link-Flap Dampening (LFD) allows the switch to separately monitor each interface for flapping events. If flapping events on an interface exceed the maximum allowed number of flaps within a time interval, then the switch automatically shuts down that interface.

PKI CRL, CDP Check and Credential Expiry

Added support for Public Key Infrastructure (PKI) CRL, CDP check, and credential expiry alert.

vLAG Orphan Ports and SVI Shutdown

Added the vLAG Orphan port option to set a port that is not part of a vLAG aggregation to also go to suspend state on secondary vLAG switch when the ISL fails with the peer vLAG switch alive.

Similarly, added the option to shutdown on a secondary vLAG switch the Layer 3 SVIs that the ISL LAG belongs to in order to disable Layer 3 forwarding while being in dual-active state.

Known Issues

This section describes known issues for CNOS 10.10 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 10.10 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)s. After the firmware upgrade, the LACP configuration remains the same, but LACP enabled ports now transition to suspended state when no LACPDU)s 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)s, 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 10.10 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)
- When upgrading the firmware to CNOS 10.10 from 10.9.3, an incorrect syslog appears, alerting the user that the VLAG peer is using an older image than 10.9.x.

```
%VLAG-5-ECP_VERSION_MISMATCH: vLAG peer is running an image older than 10.9.x. Use the upgrade procedure described in the Release Notes to ensure there is no traffic loss!
```

This message does not affect the upgrade process and should be ignored. (ID: 141257)

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)

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.

IGMP

- The number of IGMPv3 source entries (for a multicast group) is limited to 64 when VLAG is enabled. (ID: 171858)

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:

```
%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: 158487)

DCI MP-BGP L2VPN

- BGP-EVPN is not backwards compatible with earlier CNOS releases. (ID: 157557)
- 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)

Firmware Update Error

- When upgrading the firmware to CNOS 10.10 from 10.7.1, 10.8.1 or 10.9.1, an incorrect error message appears:

```
WARNING: Downgrade from 10.7.x or later to 10.6.x or earlier version
detected.Prior to proceeding with the downgrade, please save/backup
the 10.6.x or earlier version startup-config file by using the 'copy
file config-10-6' command.
After the downgrade, once the switch comes up on 10.6.x or earlier
version, the saved 10.6.x or earlier version startup-config file needs
to be copied over to the running and startup configuration manually.
```

This message does not affect the upgrade process and should be ignored. (ID: 157431)

- When upgrading from an earlier version to 10.10, make sure both the software image and the boot image are updated. (ID: 161315)

Port Statistics

At times, the switch displays a slightly incorrect interface packet rate when running port statistics commands. To avoid reading an incorrect packet rate, we recommend that you run such statistics commands several times. (ID: 166285)

