

Lenovo ThinkSystem NE2572 RackSwitch

Release Notes

For Lenovo Cloud Network Operating System 10.8

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 (July 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 10.8 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 10.8:

- *Lenovo Network Application Guide for Lenovo Cloud Network Operating System 10.8*
- *Lenovo Network Command Reference for Lenovo Cloud Network Operating System 10.8*
- *Lenovo Network Python Programming Guide for Lenovo Cloud Network Operating System 10.8*
- *Lenovo Network REST API Programming Guide for Lenovo Cloud Network Operating System 10.8*
- *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 10.8 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 (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                   |>=: Select Screen|
| System Language [English]                           |^v: Select Item  |
|                 |                                   |Enter: Select    |
|                 |                                   |+/-: Change Opt. |
| System Date     [Thu 05/11/2017]                   |F1: General Help |
| System Time     [11:13:37]                          |F2: Previous Values|
|                 |                                   |F3: Optimized Defaults|
| Access Level    Administrator                       |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 Quick Start 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 fixes, enhancements, and other changes.

802.1Q encapsulation

802.1Q encapsulation inserts a 802.1Q tag into Ethernet frames transmitted through a Layer 3 routed port that identifies to which VLAN the frame belongs.

802.1Q encapsulation adds a 802.1Q tag after the source and destination MAC address fields of the outgoing routed packet. The tag contains the VLAN identifier which can be configured for each Layer 3 routed interface.

BGP Unnumbered and BFD

BGP unnumbered can be used in conjunction with Bidirectional Forwarding Detection (BFD). It extends the capabilities of BGP unnumbered to make setting up large configurations easier and faster by enabling BFD on all BGP unnumbered neighbors.

Changing the Default Network Administrator Password

After logging onto the switch for the first time or after updating the switch firmware to 10.8 from an earlier version, a screen prompt appears requesting you to change the default network administrator password.

```
You are required to change your default 'admin' user password immediately
Changing password for admin.

Current Administrator password:

Choose a strong password (Min 8 chars, at least 1 uppercase, 1 lowercase,
1 number etc.)
New Administrator password:
Retype new Administrator password:
```

You must enter the current network administrator password, then set up a new password and then enter it a second time for confirmation.

CNOS requires the use of strong passwords for users to access the switch. Strong passwords enhance security because they make password guessing more difficult.

Data Center Interconnect

Data Center Interconnection (DCI) is a wide area architecture between data centers. It ensures that data is consistent, it allows virtualized workloads to be moved quickly, and it makes for an easier deployment of high availability solutions between different geographical locations.

DCI enables the bridging of Layer 2 services between multiple locations across existing Layer 3 IP networks. Virtual Extensible LAN (VXLAN) provides a high availability DCI framework for connecting two or more data centers. It does not require specialized hardware and it is deployable using any network devices that support VXLAN.

Lenovo offers a DCI solution that can be configured statically or by using Multi-protocol Border Gateway Protocol (MP-BGP) Ethernet Virtual Private Network (EVPN). High Availability (HA) topologies using Virtual Link Aggregation Group (vLAG) are possible for both configuration options.

Dynamic Port Mode

Switch ports configured in 25Gbps hardware profile port mode can have their port speed changed to 10Gbps or back to 25Gbps without needing to reload the switch.

You need to configure the same hardware profile port mode and the same port speed on both end-ports of the link. You also need install transceivers that support both 25Gbps and 10Gbps speeds, else you need to change the transceiver each time the port speed is modified.

Explicit Congestion Notification

Explicit Congestion Notification (ECN) extends the functionality of Weighted Random Early Detection (WRED) by marking packets instead of discarding them when the average queue length exceeds the configured threshold. Network devices configured with WRED and ECN use the marking of packets as a signal that the network is congested and packet transmission is slowed down.

IPv4 Next-hop Health Check

IPv4 next-hop health checking is used to verify if the IPv4 next-hop is reachable or not. It uses Address Resolution Protocol (ARP) to periodically send ARP request to the next-hops that have associated routes present in the switch's routing table.

Lenovo Intelligent Network Controller

Lenovo Intelligent Network Controller (LINC) is a Software-Defined Networking (SDN) solution from Lenovo which provides network virtualization with service chaining capability to offer flexibility, agility, and scalability needed for modern cloud data centers.

Lenovo Intelligent Network Controller is based on the OpenContrail R3.2 release.

Multiple VRF Instances

CNOS now supports multiple Virtual Routing and Forwarding (VRF) instances. VRF allows multiple instances of a routing table to coexist and to work simultaneously on a switch. You can segment network paths without using multiple devices while running separate routing protocol instances and forwarding tables.

VRF now supports up to nine instances: eight instances for data traffic and one instance for management traffic.

Private VLANs

Private VLANs provide Layer 2 isolation between the ports within the same broadcast domain. Private VLANs can control traffic within a VLAN domain and provide port-based security for host servers.

REST/Python API Enhancement

REST and Python APIs have been introduced or updated for the following:

- OSPF
- BGP DiffServ Code Point (DSCP) marking
- Routes Maps (GET only)
- 802.1Q Encapsulation
- IPv4 Next-hop Health Check
- Network Time Protocol (NTP)
- Network Virtualization (NWV)
- Private VLANs
- Virtual Extensible LAN (VXLAN)
- Data Center Interconnect (DCI)
- Telemetry
- Secure Shell (SSH)
- SSH Public Key Authentication
- Telnet
- Network Policy Agent (NPA)
- REST Server
- Hardware Switch Controller (HSC)

Secure Hash Algorithm

Network Time Protocol (NTP) now supports Secure Hash Algorithm 1 (SHA-1) encrypted authentication.

Open Shortest Path First (OSPF) and Bidirectional Forwarding Detection (BFD) now support SHA-256 encrypted authentication.

SSH Public Key Authentication

Public Key Authentication (PKA) is a cryptographic system that offers an alternative for local SSH authentication, by providing the switch with a pair of keys: a public key and a private key.

When connecting to switch via SSH, the SSH server encrypts messages exchanged during the session using the public key. Only the SSH client that has the matching private key can decrypt the messages and thus successfully authenticate with the SSH server.

System Logging Enhancement

When configuring remote syslog servers on the switch, you can now also specify the transport protocol (TCP or UDP) and the port used for communication.

User Action Logging Enhancement

User actions executed on the switch are logged and saved to a file without displaying them on the console, or on Telnet and SSH terminals. User action logging has a severity level of informational (6).

You can now enable user action logging when using the CLI, REpresentational State Transfer (REST), or Simple Network Management Protocol (SNMP).

vSphere Distributed Switch Support for ThinkAgile Network Orchestrator

The ThinkAgile Network Orchestrator (TNO) now supports VMware's vSphere Distributed Switch (vDS) through integration with vCenter. TNO can be used with both VMware and Nutanix environments. TNO's dynamic port and security (QoS) configuration is now integrated with VMware's distributed vSwitch.

Known Issues

This section describes known issues for CNOS 10.8 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)

Firmware Downgrade

When performing a firmware downgrade from CNOS 10.7 or later to CNOS 10.6 or an earlier version, the following warning message appears:

```
WARNING: Downgrade from 10.7.x or later to 10.6.x or earlier version detected. This downgrade would require a manual configuration change when the switch comes up on 10.6.x or earlier version. Prior to proceeding with the downgrade, please keep 10.6.x or earlier version startup-config file handy by using 'copy file config-10.6' command.
```

The warning message is incorrect. The actual command is:

```
NE2572# copy file config-10-6
```

The warning message will be corrected in the next firmware release. (ID: 141384)

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 10.8 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)

REST API

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

System Logging

The switch does not send syslog messages to syslog servers defined using DNS that are accessible through the default Virtual Routing and Forwarding (VRF) instance. DNS defined syslog servers using the management VRF instance are able to receive syslog messages from the switch. (ID: 135438)

vLAG

When upgrading the firmware of vLAG peer switches one by one to CNOS 10.8 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)