Lenovo ThinkSystem NE2572

Release Notes

For Lenovo Cloud Network Operating System 10.6

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 (January 2018)

Release Notes

This release supplement provides the latest information regarding Lenovo Cloud Network Operating System 10.6 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* 10.6:

- *Lenovo Network Application Guide for Lenovo Cloud Network Operating System 10.6*
- *Lenovo Network Command Reference for Lenovo Cloud Network Operating System 10.6*
- *Lenovo Network Python Programming Guide for Lenovo Cloud Network Operating System 10.6*
- *Lenovo Network REST API Programming Guide for Lenovo Cloud Network Operating System 10.6*
- *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/elie-1.7.1/licenses/. For details on how to obtain open source code, please contact Lenovo Support.

Hardware Support

CNOS 10.6 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

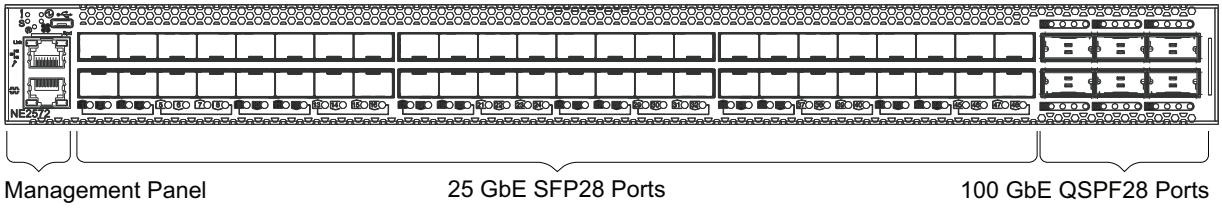

Supplemental Information

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.

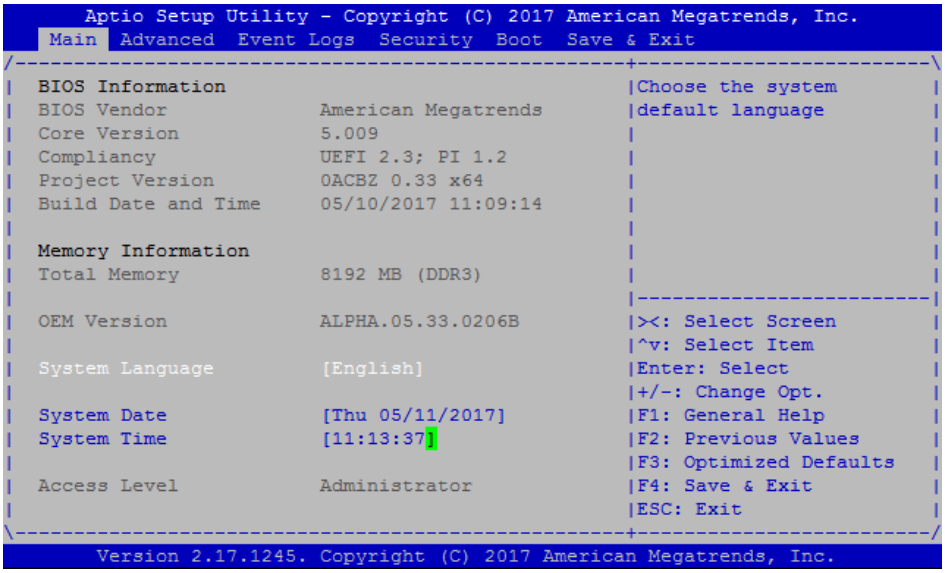

This menu permits the following actions:

- 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.

Access Control List Logging

ACL logging provides insight into traffic as it passes the network or is dropped by the switch. ACL logging can be CPU intensive and can slow traffic going through the network device. There are two main factors that contribute to the CPU load increase from ACL logging:

- process switching of packets that match log-enabled access control entries (ACEs)
- generating and transmitting log messages

Access Control Lists Remarks

A remark is comment text that can be added to an Access Control List (ACL) to make the ACL configuration easier to understand.

Address Resolution Protocol Refresh

Each Address Resolution Protocol (ARP) entry is checked periodically to determine its state. Based on the entry's state, ARP undertakes certain actions, like refreshing the entry or removing it from the ARP table.

Border Gateway Protocol and Differentiated Services

Border Gateway Protocol (BGP) works with the Differentiated Services (DS) computer networking architecture. You can use differentiated services with BGP to provide low latency to critical network traffic, such as VoIP, while providing best-effort service to noncritical services.

Border Gateway Protocol Unnumbered

Border Gateway Protocol (BGP) unnumbered is useful for quickly setting up large configurations for CLOS based network design. In a multi-chassis system you have a set of lower layer or line card switches that all interconnect through a different set of upper layer or fabric card switches to the fabric chassis.

BIOS Menu Rescue Mode

The Rescue Mode option allows you to recover from a failed firmware or boot image upgrade using TFTP or a USB drive, reset the switch configuration to factory defaults, reset the network administrator account password, and reset the password required to enter Privileged EXEC configuration mode.

This option is available from the BIOS Menu. For more details, see "The BIOS Menu" on page 5.

Dynamic ACL/QoS Provisioning for Lenovo HCI Solution with Nutanix

Guest virtual machines (VMs) can now have Access Control Lists (ACLs) attached to them and these ACLs follow the guest from one host to another within the same Nutanix cluster. IP and MAC security filters, QoS and Queue filters for prioritization can be deployed. Filters must be associated with the VM in the configurations of all switches attached to the Nutanix hosts - this can't be done from PRISM.

Dynamic Peer BFD Support

BFD now provides sub-second failure detection between two dynamic BGP peers.

Hybrid Bridge Port Mode

Hybrid bridge port mode is a trunk bridge port mode that lets you have more than one egress untagged VLAN. Like a trunk port, a hybrid port can carry multiple VLANs to receive and pass traffic for them. The hybrid bridge port mode lets you control which VLANs receive tagged egress traffic and which VLANs receive untagged egress traffic. Unlike trunk bridge port mode, the native VLAN is not considered the only VLAN that can send untagged traffic.

IP Subnet VLAN Assignment

IP subnet VLAN assignment lets you configure the switch to assign a VLAN based on the IP subnet for incoming untagged or priority-tagged packets. This feature can also assign a priority to untagged traffic.

IPv6 Neighbor Table Threshold

The IPv6 Neighbor Table threshold has been increased from 8,192 to 16,000 entries.

Layer 2 Failover

The main purpose of Layer 2 Failover is to support Network Adapter Teaming. With Network Adapter Teaming, all the NICs on each server share the same IP address and are configured into a team. One NIC is the primary link and the other is a standby link.

The Manual Monitor (MMON) enables you to configure a set of ports or LAGs to monitor for link failures (a monitor list), and another set of ports or LAGs to disable when the number of forwarding monitor links is less than the trigger limit (a control list). When the switch detects a link failure on the monitor list, it automatically disables the items in the control list. When server ports are disabled, the corresponding server's network adapter can detect the disabled link and trigger a network adapter failover to another port or LAG on the switch or another switch.

Layer 2 Failover works together with static LAGs, Link Aggregation Control Protocol (LACP), and with Spanning Tree Protocol (STP).

Lenovo Ganglia Plug-in

Lenovo Ganglia plug-in helps the telemetry agent to integrate with Ganglia. It provides pull and push mode support to collect telemetry BST statistics, translates data to Ganglia metrics, and provides a visualization tool.

Lightweight Directory Access Protocol

Lightweight Directory Access Protocol (LDAP) is a protocol for accessing distributed directory information services over a network. LDAP is used for authentication and authorization. With an LDAP client enabled, the switch will authenticate a user and determine the user's privilege level by checking with one or more directory servers instead of a local database of users. This prevents customers from having to configure local user accounts on multiple switches; they can maintain a centralized directory instead.

NSX Gateway

NSX is a VMware solution for network virtualization. Lenovo's NSX gateway provides a Hardware Switch Controller (HSC) VxLAN Tunnel Endpoint (VTEP) that enables legacy network devices to communicate with Virtual Machines (VMs) on NSX hypervisors.

Note: Lenovo is actively working with VMware to complete the certification for NSX Gateway. NSX Gateway support is available in this release while VMware certification is pending. Lenovo will fully support the NSX Gateway implementation throughout the certification process. Please work with the controller vendor on issues related to the controller.

Proxy Address Resolution Protocol

Proxy Address Resolution Protocol (ARP) is a technique in which a device on a given network answers ARP requests intended for another device. The proxy ARP is aware of the location of the traffic's destination and offers its own MAC address as the destination. The received traffic is then routed by the proxy device to the intended destination via another interface or a tunnel. Proxy ARP is primarily used when hosts in the connected subnet are separated by features such as a private VLAN.

Reserved VLANs

Some features, such as OpenFlow and Layer 3 ports (routed ports), require internal VLANs for their operations. You can expand the range of internal VLANs by reserving a contiguous block of VLANs to guarantee the delivery and operation of features with such requirements. These reserved VLANs cannot be created, deleted, modified, or manipulated, unless the reserved VLAN range is reset to the default range (4000-4094).

Sampled Flow

Sampled Flow (sFlow) a technology for monitoring traffic in networks containing switches and routers. The sFlow monitoring system consists of an sFlow Agent and a central sFlow Collector. The sFlow architecture and sampling techniques were designed for providing continuous site-wide and enterprise-wide traffic monitoring of high speed switched and routed networks.

Scheduled Switch Reloads

This feature lets you schedule a switch reload for a later time, enabling you to perform switch upgrades during off-peak hours.

Secure Mode

Secure mode allows you to determine which protocols can be enabled. In secure mode, only secured traffic and secured authentication management are allowed.

Syslog User Action Logging

You can configure whether user actions are logged to the console or to terminals.

VLAN Access Control Lists over Switch Virtual Interfaces

A VLAN ACL is an Access Control List (ACL) that can be assigned to a VLAN rather than to a switch port as with IPv4 ACLs. This is particularly useful in a virtualized environment where traffic filtering and metering policies must follow virtual machines (VMs) as they migrate between hypervisors.

ACLs can now be also configured over Layer 3 Switch Virtual Interfaces (SVIs).

vNIC Statistics for Lenovo HCI Solution with Nutanix

Statistical data is now generated for traffic in and out of a vNIC as defined in PRISM. These statistics are retrieved from the switch(es) attached to the Nutanix hosts.

Warning Message Displayed when Changing the Default Network Administrator Password

The following warning message is displayed reminding the users to change the default Network Administrator password as soon as possible:

```
Warning: Please change the default Network Administrator password as soon
as possible. Note that in the next CNOS release (10.7.x or later), user
will be forced to change the default password upon first successful
login.
```

Weighted ECMP Routes

In traditional ECMP with next hops, each hop is added to the ECMP multipath once so traffic is distributed equally among the next hops. However, in some scenarios, traffic may use one path more than others, causing congestion. A lack of balance can occur because the ECMP hashing algorithm only considers the source, destination, or both when selecting the path, but not the use of the link.

Weighted ECMP lets you configure the multipath based on the use of each link, thus avoiding congestion and obtaining a better balance of traffic. This is achieved by adding the next hops in the multipath from one to four times.

Known Issues

This section describes known issues for CNOS 10.6 on the Lenovo ThinkSystem NE2572 RackSwitch.

NSX Gateway

The following limitations exist:

- For optimal performance, it is recommended 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, it is recommended 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)
- In the switch Application-specific Integrated Circuit (ASIC), the lookup to VxLAN-translation takes place first, before the dot1q tunnel VLAN. The dot1q tunnel VLAN is not used during VxLAN processing. Therefore, dot1q tunnel feature in conjunction with VxLAN is not operational. (ID: 101708)
- 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 the vLAG instance goes down on one of the vLAG switches, the MAC addresses learned on that instance are not be moved to the ISL. This means that the traffic will be flooded on both ISL and other potential access ports. The flooding on ISL means that the traffic eventually gets to the vLAG instance on the peer, so no traffic is lost. The hosts on the other access ports drop the traffic as it is not addressed to them. (ID: 108729)
- 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)

Privileged EXEC Mode Password Persistence

When upgrading the switch firmware image from CNOS version 10.3 or 10.4 to CNOS version 10.6, if there exists a previously configured encrypted password used to enter Privileged EXEC mode, it persists across the upgrade process. It is overwritten only when configuring a new clear text password and the switch running configuration is saved.

If a previously configured encrypted password is still used for entering Privileged EXEC configuration mode after the upgrade process, then only the first eight characters are checked when entering the password.

When downgrading the switch firmware image from CNOS version 10.6 to a previous version, if a previously configured Privileged EXEC encrypted password is present in the switch startup configuration file, then the password persist across the downgrade process. The password is required to enter Privileged EXEC configuration mode after the downgrade process is done. (ID: 119771)