Extreme Network OS 7.0.2d for Extreme VDX devices

Release Notes

Contents

DOCUMENT HISTORY

Version	Summary of Changes	Publication Date
1.0	Initial Release	November 2020

PREFACE

Contacting Extreme Technical Support

As an Extreme customer, you can contact Extreme Technical Support using one of the following methods: 24x7 online or by telephone. OEM customers should contact their OEM/solution provider.

If you require assistance, contact Extreme Networks using one of the following methods:

GTAC (Global Technical Assistance Center) for immediate support

Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.

Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.

GTAC Knowledge - Get on-demand and tested resolutions from the GTAC Knowledgebase or create a help case if you need more guidance.

The Hub - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employee, but is not intended to replace specific guidance from GTAC.

Support Portal - Manage cases, downloads, service contracts, product licensing, and training and certifications.

Before contacting Extreme Networks for technical support, have the following information ready:

Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products

A description of the failure

A description of any action(s) already taken to resolve the problem

A description of your network environment (such as layout, cable type, other relevant environmental information)

Network load at the time of trouble (if known)

The device history (for example, if you have returned the device before, or if this is a recurring problem)

Any related RMA (Return Material Authorization) numbers.

Extreme resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at https://www.extremenetworks.com/support/documentation/.

Document feedback

Quality is our first concern at Extreme, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you.

You can provide feedback in two ways:

- Use our short online feedback form at at https://www.extremenetworks.com/documentation-feedback/
- Email us at documentation@extremenetworks.com

  Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

OVERVIEW

Extreme Network OS 7.0.2b release introduces incremental capabilities to further enable the Layer 3 Extreme® IP Fabrics, Layer 2 VCS ® Data Center Fabrics, Standards (BGP-EPVN) based network Virtualization for Intra DC and DC-Interconnect solutions.

Hardware

The following section lists new hardware introduced with this release as well as hardware that are no longer supported with this release.

New devices

None

New interface modules

None

Deprecated Hardware

None

Software Features

The following section lists new, modified, and deprecated software features for this release. For information about which platforms support these features, refer to the NOS Feature support Matrix.

New Software Features for Network OS v7.0.2d

The following software features are new in this release:

Fix exposure to netkit-telnetd vulnerability CVE-2020-10188

New Software Features for Network OS v7.0.2c

- None

Limitation:

- If the mgmt-vrf default route is configured to be multipath such that paths are resolved via combination of out-of-band and in-band interface, then behavior on removing/addition of these paths is unpredictable.

New Software Features for Network OS v7.0.2b

On October 30, 2017, Extreme Networks, Inc. acquired the data center networking business from Extreme Communications Systems, Inc. The code in this release has been updated to replace technical references to Extreme Communications, Inc. with Extreme Networks, Inc., as appropriate.

For details, refer to the “What’s new in this document” section in the following:

- *Extreme Network OS Command Reference, 7.0.2b*
- *Extreme Network OS Troubleshooting Guide, 7.0.2b*
- *Extreme Network OS YANG Reference Manual,7.0.2b*
- *Extreme Network OS Software Upgrade Guide,7.0.2b*

New Software Features for Network OS v7.0.2

Importing TLS certificates and keys using SCP

This feature allows TLS server certificates (third party CA certificate) and keys to be directly imported without any trust point support.

New Software Features for Network OS v7.0.1

The following software features are new in this release:

SNMP CLI knob to control 3-tuple vs 2-tuple

By default IF MIB objects ifName and ifDescr are retrieved in 2-tuple format. New CLI knob "snmp-server three-tuple-if enable" enables to retrieve in 2-tuple/3-tuple format as per configuration

Two Factor Authentication

Traditional password based authentication is one factor which has security risk as it can be guessed, cracked, or compromised. Two factor authentication uses password as one factor and randomly generated RSA token as second factor. These two factors combined to generate a “passcode”. The passcode is sent to the Radius server for authentication

Duration login

Login Duration enhancement helps to restrict the user’s log-in duration, user can be restricted to login within specific duration, for example, setting login duration as 0900-1100(HHMM format) for a specific user, will restrict that user to login only from 9am to 11 am

IP Fabric enhancements:

Decouple IPv4 and eVPN address-family

To have separate BGP peering session between leaf and spine for underlay and overlay, to decouple overlay neighbor-ship errors from the underlay session and so the routes arenotcompromised.

Layer 3 VNI IMR routes

IMR routes are originated whenever L3VNI is configured under a VRF, even if there are no prefix routes originated, IMR route will establish tunnel-vlan membership on the remote leaf nodes.

Traffic tromboning avoided in IP Fabric by converting ARP into /32 routes.

The MACIP (ARP) routes are converted to /32 prefix routes and installed into the RIB, along with the subnet routes. The /32 host routes help with traffic being delivered directly to the leaf node where host resides, thus avoiding traffic tromboning.

Peer disable AS-check

AS path check can be enforced at the sender side, saving the amount of BGP RIBout memory used to store the routes which are discarded at the receiver and also avoid sending updates and withdraw for those routes, thus improving the convergence time.

Hardware TCAM profiles

Support for increased IPv4 ACL scale has been added as part of hardware TCAM profiles in VDX8770. ACL scale is increased as well as support for number of rules per ACL.

MAPS policy reapply

New command to reapply the MAPS policy globally. If you modify the MAPS policy configuration on a cluster, this command reapplies the policy to the cluster with the updated configuration.

MAC-move-detect feature

The mac-address-table mac-move detect command has been added to support the detection of MAC moves. When this feature is enabled, the default number of MAC-moves that are detected is 20. This limit can be changed by means of the mac-address-table mac-move limit command.

LLDP Enhancement

RASlog support is added in LLDP to capture and report connectivity to peer status change in order to trigger python scripts.

Modified Features

The following software features have been modified in this release:

None

Deprecated Software Features

None

CLI Changes

The following section lists new, modified, and deprecated commands for this release. For details, refer to the Network OS Command Reference.

New Commands for Network OS v7.0.2d

None

New Commands for Network OS v7.0.2c

None

New Commands for Network OS v7.0.2b

None

New Commands for Network OS v7.0.2

The following configuration commands are new in this release:

- show cert-util tlscert
- show cert-util tlsprivkey

New Commands for Network OS v7.0.1c

None

New Commands for Network OS v7.0.1b

None

New Commands for Network OS v7.0.1

The following configuration commands are new in this release:

neighbor <ip address/peer-group> enable-peer-as-check

clear bgp evpn l2route type inclusive-multicast

maps reapply-policy

mac-address-table mac-move detect

area <area_id> nssa default-information-metric metric-type <type> metric <value>

area <area_id> nssa no-redistribution

area <area_id> nssatranslator-always

area <area_id> nssa translator-interval <interval_value>

snmp-server three-tuple-if enable

spanning-tree ieee-bpdu limit-vlan-flood

The following show commands are new in this release:

ARP:

show system internal arp clientlist structures

show system internal arp counter

show system internal arp interface ve 100 show

system internal arp ipv4 vrf all show system

internal arp ipv6 vrf all show system internal arp

l2 clientdb show system internal arp memstats

show system internal arp rib memstats

show system internal arp rib routes

show system internal arp static config interface ve 100

show system internal arp static config vrf all

show system internal arp summary

show system internal arp vrf all

RIB:

show system internal rib clients

show system internal rib ip adj vrf default-vrf

show system internal rib ip route vrf default-vrf

show system internal rib ipv6 adj vrf default-vrf

show system internal rib ipv6 route vrf default-vrf

show system internal ribm memory-stats

show system internal ribm vrf

FIB-ARP:

show system internal fib 0 arp arp-bum-stats

show system internal fib 0 arp counters

show system internal fib 0 arp dai-filter

show system internal fib 0 arp dai-list

show system internal fib 0 arp dai-stats

show system internal fib 0 arp dai-trust

show system internal fib 0 arp dai-usr-acl

show system internal fib 0 arp interface ve100

show system internal fib 0 arp ipv4 vrf all

show system internal fib 0 arp ipv6 vrf all

show system internal fib 0 arp memstats

show system internal fib 0 arp nd-bum-stats

show system internal fib 0 arp summary

show system internal fib 0 arp vrf all

FIB-RIB:

show system internal fib 0 rib clients

show system internal fib 0 rib ip adj vrf default-vrf

show system internal fib 0 rib ip route vrf default-vrf

show system internal fib 0 rib ipv6 adj vrf default-vrf

show system internal fib 0 rib ipv6 route vrf default-vrf

show system internal fib 0 rib memory-stats

Modified Commands for Network OS v7.0.2d

None

Modified Commands for Network OS v7.0.2c

None

Modified Commands for Network OS v7.0.2b

The following command have been modified for this release:

show media

show version

Modified Commands for Network OS v7.0.1b

None

Modified Commands for Network OS v7.0.1

The following commands have been modified for this release:

duplicate-mac-timer <interval> max-count <count>

username <name> access-time HHMM to HHMM

ip mtu

ipv6 mtu

mtu

password-attributes

Deprecated Commands for Network OS v7.0.1b

None

Deprecated Commands for Network OS v7.0.1

no nssa-translator

system tunnel replicator load-balance

API Changes

Network OS follows the YANG model for CLI and NetConf/REST API. Hence relevant changes in above CLI Changes will get mirrored in API Changes as well.

Newly supported standards and RFCs

The following section lists RFCs and other standards newly supported in this release.

This software generally conforms to Ethernet standards in a manner consistent with accepted engineering practices and procedures. In certain cases, Extreme might add proprietary supplemental functions to those specified in the standards, or choose to implement modifications to the standards for performance or behavioral improvements.

None

HARDWARE SUPPORT

Supported devices

Extreme NOS v7.0.2c supports following VDX Switches:

- ExtremeSwitching VDX 6940-144S
- ExtremeSwitching VDX 6940-36Q
- ExtremeSwitching VDX 6740
- ExtremeSwitching VDX 6740T
- ExtremeSwitching VDX 6740T-1G
- ExtremeSwitching VDX8770-4
- ExtremeSwitching VDX8770-8
- ExtremeSwitching VDX2741
- ExtremeSwitching VDX2746

ExtremeSwitching VDX 6940-144S

The ExtremeSwitching VDX 6940-144S is a 2U platform that offers 96 x 10GbE SFP+ downlink ports for server connectivity and also 12 x 40 GbE QSFP+ uplink ports to connect to the aggregation layer. These ports support the following:

- Available in 64, 96 and 144 ports SKU.
- Each 40GbE port can be broken into 4 independent 10GbE ports, providing a total of up to 144 x 10GbE ports in a 2RU form factor.
- 64 port SKU can be upgraded up to 144 ports with Ports On Demand (POD) software license. There are two POD licenses - 16x10GbE for 10GbE server connecting ports and 6x40GbE for the 40GbE uplink ports. The same 6x40GbE POD license can be used to upgrade up to 12x40GbE uplink ports in both 64 and 96 port SKUs.
- Deployable as high-density 10GbE switch for the Top of Rack (TOR) or Middle of Row (MOR) or for End of Row (EOR) configurations.
- Provides optimized on-chip buffer (24MB) and latency (800ns), making it an ideal switch for a wide variety of workloads.
- Interface 97, 98 103 and 104 are dual personality ports. These ports can be configured in 40GbE or 100GbE mode.

ExtremeSwitching VDX 6940-36Q

The ExtremeSwitching VDX 6940-36Q is a 1U platform that offers 36 x 40 GbE QSFP+ ports. Each 40 GbE ports can be further broken out into 4 independent 10 GbE SFP+ ports providing a total of 144 x 10 GbE SFP+ ports. These ports support the following:

- Available in 24 and 36 ports SKU.
- Each 40GbE port can be broken into 4 X 10GbE ports, providing up to 144 x 10GbE ports in a 1RU form factor.
- The 24 port SKU can be upgraded up to 36 ports via 40GbE DPOD license of 12 ports.
- It can be used as a high-density 40GbE spine switch or it can also be used as a leaf switch with dynamic breakout capability.
- It provides optimized on-chip buffer (24MB) and latency (800ns), making it an ideal switch for a wide variety of workloads.

ExtremeSwitching VDX 6740

The ExtremeSwitching VDX 6740 offers 48 10GbE SFP+ ports and 4 ports of 40 Gigabit quad small form-factor pluggable plus (QSFP+), each can be broken out into four independent 10 GbE SFP+ ports, providing an additional 16 x 10 GbE SFP+ ports. No 40 GbE ports are enabled as part of the base license. Four 40 GbE ports can be upgraded via the Ports on Demand (PoD) software license. These ports support the following:

- Available in 24, 48 and 64 port SKU.
- 850-ns microsecond latency for any port to port to assure rapid response for latency-sensitive applications.
- The base SKU is available with 24 ports and can be upgraded up to 48 ports via 10Gbe DPOD license of 8 ports.
- Of the 48 10GbE SFP+ ports, 32 ports can be configured as FlexPorts(FC/Ethernet).
- It has 4 X 40Gbe QSFP ports which can be used for the uplink and VCS fabricformation.
- Each 40GbE port is capable of doing a breakout of 4 X 10GbEports.
- Additional 4X40GbE ports can be added to base version with 2X40GbE POD license increments.
- 100Mb Support – Refer to “Support for 100-Mb interfaces” sections below.

ExtremeSwitching VDX 6740T

The VDX 6740T offers 48 10GbE Base-T ports and 4 ports of 40-gigabit quad small form-factor pluggable plus (QSFP+), each can be broken out into four independent 10GbE SFP+ ports, providing an additional 16 x 10 GbE SFP+ ports. No 40 GbE ports are enabled as part of the base license. Four 40 GbE ports can be upgraded via the Ports on Demand (PoD) software license.

- Available in 24, 48 and 64 port SKU.
- 3 microsecond latency for any port to port to assure rapid response forlatency-sensitive applications.
- The base SKU is available with 24 10GbE Base-T ports and can be upgraded up to48 ports via 10Gbe DPOD license of 8 ports.
- It has 4 X 40 GbE QSFP ports which can be used for uplink and VCS fabricformation.
- Each 40GbE port is capable of doing a breakout of 4 x 10GbEports.

- Each 40GbE port is also capable of doing a FC breakout of 4*8G or 4*16G. These ports can be used to connect to the FOS switches.
- Each 40GbE port is also capable of doing an FC breakout of 4 x8G/16G.
- Additional 4X40GbE ports can be added to base version with 2X40GbE PODlicense increments.
- 100Mb Support – Refer to "Support for 100-Mb interfaces" below.

ExtremeSwitching VDX 6740T-1G

The ExtremeSwitching VDX 6740T-1G offers 48 1000BASE-T ports and two 40 GbE QSFP+ ports in base version. Each 40 GbE port can be broken out into four independent 10 GbE SFP+ ports, providing an additional eight 10 GbE SFP+ ports for uplink. All 48 1000BASE-T ports can be upgraded to 48 10GBASE-T ports via a Capacity on Demand (CoD) software license. Two 40 GbE ports are enabled as part of the base license. The additional two 40 GbE ports can be upgraded via the Ports on Demand (PoD) software license.

- Base version is available with 48 x 1000BASE-T ports and 2 x 40 GbE QSFP+ ports.
- 3-microsecond latency for any port to port to assure rapid response for latency-sensitive applications.
- All 48 x 1000BASE-T ports can be upgraded to 10Gbase-T port with capacity on demand license.
- Additional 2X40Gbe ports can be added to base version with 2X40Gbe PODlicense.
- It has 4 X 40Gbe QSFP ports which can be used for the uplink and VCS fabricformation.
- Each 40GbE port is capable of doing a breakout of 4 X 10GbEports.
- Each 40GbE port is also capable of doing a FC breakout of 4 x8G/16G.
- 100Mb Support – Refer to "Support for 100-Mb interfaces" below.

ExtremeSwitching VDX 8770-4 and VDX 8770-8

The ExtremeSwitching VDX 8770 is available in two form factors; a 4-I/O slot system and an 8 I/O slot system with line-card support for 1-GbE, 10-GbE, 10GbE-T, 40GbE, and 100GbE ports. The ExtremeSwitching VDX 8770 delivers a high-performance switch to support the most demanding data center networking needs, capable of supporting:

- 4 Tbps per slot line-rate design for substantial capacity and headroom.
- ~4-microsecond latency to assure rapid response for latency-sensitiveapplications.
- Up to 384,000 MAC addresses per fabric for extensive virtualizationscalability.
- More than 8000 ports in a single VCS Fabric with Extreme Fabric Multipathing technology, enabling the switch to serve extremely large-scale deployments with the best-possible network utilization.

Supported Blades for VDX 8770

The flexible, modular switch design offers interconnection with other Extreme switches, traditional Ethernet switch infrastructures, and direct server connections. Modular 4-slot and 8-

slot chassis options are available to match the switch to the needs of the organization. These include:

- **ExtremeSwitching VDX 8770-4**: Supports up to 192 1/10 GbE ports, or 108 40GbE ports and 24 100 GbE ports, or a combination.
- **ExtremeSwitching VDX 8770-8**: Supports up to 384 1/10 GbE ports, or 216 40GbE ports and 48 100 GbE ports, or a combination.

The switches support two Management Modules in an active standby configuration. The 4 slot chassis can hold up to 3 Switch Fabric Modules (SFM) and 4 Power supply Units (PSU) while the 8 slot chassis can hold 6 SFMs and 8 PSUs. The switch supports a variety of wire-speed line cards to offer maximum flexibility in terms of port bandwidth as well as cable and connector technology:

- 1 GbE: LC48×1G line card provides up to 48 SFP/SFP-copper ports.
- 10 GbE: LC48×10G line card provides up to 48 SFP+ ports.
- 10 GbE-T: LC48×10GT line card provides up to 48 RJ-45 ports.
- 40 GbE: LC12×40G line card provides up to 12 x 40 GbE QSFP ports.
- 40 GbE: LC27×40G line card provides up to 27 x 40 GbE QSFPports.
- 100 GbE: LC6×100G line card provides up to 6 x 100 GbE CFP2ports.

ExtremeSwitching VDX 2741

• ExtremeSwitching VDX blade switch for Converged Ethernet Fabrics in Sugon TC-6600 Chassis.
• Internal Ports (KR)
  - 28 x 10G Eth (with 1GbE/10GbE auto negotiation)
• External Ports
  - 16 x Flex Ports (10GbE or 16G FC)
  - 4 x 40GbE ports with breakout capability
• Support for low cost of entry base with POD upgrades
• Single ASIC with 850ns low latency
• 128K MAC, 32K ARP and 12K ACL support
• Extreme features as supported in NOS 7.0.0
• Sugon specific features:
  - Feature on Demand (S/W upgrades)

ExtremeSwitching VDX 2746

- ExtremeSwitching VDX blade switch for the Hitachi BladeSymphony 2500chassis
- Internal Ports (KR)
  - 42 x 10G Eth (with 1GbE/10GbE auto negotiation)
- External Ports
  - 14 x Flex Ports (10GbE or 16G FC)
  - 2 x 40GbE ports with breakout capability (also Flex)
- Support for low cost of entry base with POD upgrades:
  - Management integration with Hitachi's chassis management module

Support for 100-Mb interfaces

- Full duplex speed support only for P2P connections
- Limited L2 configuration supported. For example Switchport, LLDP, MTU size, L2 ACL and L3 ACL.
- No support for adding a 100 Mbit/s shared media/hub.
- L3, FCoE, TRILL, PFC configuration are NOT supported on 100 Mbit interfaces.
- Examples for 100 Mbit/s usage are as follows:
  - 100 Mbit/s Host device requirement with IPv4/v6 Connectivity.

Supported power supplies

The following table lists the power supplies that are available for the devices supported in this release:

Part number	Description	Compatible devices
XBR-ACPWR-3000	FRU,3000W AC POWER SUPPLY	VDX 8770-4, VDX 8770-8
XBR-DCPWR-3000	FRU,3000W DC POWER SUPPLY	VDX 8770-4, VDX 8770-8
XBR-250WPSAC-F	FRU,250W,ACPS/FAN,NONPORTSIDE EXHAUST	VDX 6740
XBR-250WPSAC-R	VDX 6740 AC RTF PWR SUPPLY FAN	VDX 6740
XBR-250WPSDC-F	FRU,250W,DCPS/FAN,NONPORTSIDE EXHAUST	VDX 6740
XBR-250WPSDC-R	FRU,250W,DCPS/FAN,PORT SIDE EXHAUST	VDX 6740
XBR-500WPSAC-F	FRU 500W ACPS	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
XBR-500WPSAC-R	FRU 500W ACPS	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
RPS9DC+E	FRU,500W DC PSU PORT SIDE EXHAUST	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
RPS9DC+I	FRU,500W,DCPS/FAN,NONPORTSIDE EXHAUST	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
XBR-1100WPSAC-R	FRU,1100W PSAC,PORTSIDE EXHAUST AF	VDX 6940-144S

Part number	Description	Compatible devices
XBR-1100WPSAC-F	FRU,1100W PSAC,NON-PORT SIDE EXHAUST AF	VDX 6940-144S
XBR-1100WPSDC-01-R	FRU 1100W DCPS,PORTSIDE EXHAUST	VDX 6940-144S
XBR-1100WPSDC-01-F	FRU 1100W DCPS,NON PORTSIDE EXHAUST	VDX 6940-144S

The VDX 8770 switches ship with multiple, field replaceable, load-sharing AC or DC power supplies based on the configuration selected. The PSU SKU is shared by both 4- and 8-slot systems. The VDX 8770-4 ships with a minimum of 2 AC or DC PSU. Additional 2 PSU can be ordered for redundancy. The VDX 8770-8 system ships with a minimum of 3 PSU and additional PSU may be ordered for redundancy:

- XBR-ACPWR-3000 - 3000 W power supply unit AC
- XBR-DCPWR-3000 - 3000 W power supply unit DC

The VDX -6740 switches are both delivered with two internal, redundant, field-replaceable, load-sharing AC or DC power supplies:

- XBR-250WPSAC-F - FRU 250 W AC power supply/fan, non-port-side exhaust airflow
- XBR-250WPSAC-R - FRU 250 W AC power supply/fan, port-side exhaustairflow
- XBR-250WPSDC-F - FRU 250 W DC power supply/fan, non-port-side exhaustairflow
- XBR-250WPSDC-R - FRU 250 W DC power supply/fan, port-side exhaustairflow

The VDX -6740T switches ship with two internal, redundant, field-replaceable, load-sharing AC or DC power supplies:

- XBR-500WPSAC-F -FRU 500 W AC power supply/fan, non-port-side exhaustairflow
- XBR-500WPSAC-R - FRU 500 W AC power supply/fan, port-side exhaustairflow
- XBR-500WPSDC-F -FRU 500 W DC power supply/fan, non-port-side exhaustairflow
- XBR-500WPSDC-R - FRU 500 W DC power supply/fan, port-side exhaustairflow

The VDX -6940-36Q switches ship with two internal, redundant, field-replaceable, load-sharing AC or DC power supplies:

- XBR-500WPSAC-F -FRU 500 W AC power supply/fan, non-port-side exhaustairflow
- XBR-500WPSAC-R - FRU 500 W AC power supply/fan, port-side exhaustairflow
- XBR-500WPSDC-F -FRU 500 W DC power supply/fan, non-port-side exhaustairflow
- XBR-500WPSDC-R - FRU 500 W DC power supply/fan, port-side exhaustairflow

The VDX -6940-144S switches ship with two internal, redundant, field-replaceable, load-sharing AC or DC power supplies:

- XBR-1100WPSAC-F -FRU 500 W AC power supply/fan, non-port-side exhaustairflow
- XBR-1100WPSAC-R - FRU 500 W AC power supply/fan, port-side exhaustairflow
- XBR-500WPSDC-01-F -FRU 500 W DC power supply/fan, non-port-side exhaust airflow
- XBR-500WPSDC-01-R - FRU 500 W DC power supply/fan, port-side exhaust airflow

Supported Optics for Network OS v7.0.1x

For a list of supported fiber-optic transceivers that are available from Extreme, refer to the latest version of the Extreme Optics Family Data Sheet available online at www.extremenetworks.com.

The VDX switches support following optics types listed below. The FC SFP+ optics are supported only on VDX 6740, 2741 and 2746 switches. Breakout optics are supported only for the VDX 8770 (40G line-card), 6740/T, 2741, 2746 and 6940 platforms. The Mellanox (MAM1Q00A) optic is only supported on the VDX 8770, 6740/T and 6940 platforms. The tunable DWDM optics is supported only on VDX 8770, 6740 and 6940-144S platforms 10G ports.

Speed	FRU and Optics SKU	Description	Part Number
1GbE	XBR-000190 (1-pack)	1 GbE copper	57-1000042-01
	E1MG-SX-OM (1-pack)* E1MG-SX-OM-8 (8-pack)*	1000Base-SX	33211-100
	E1MG-LX-OM (1-pack)* E1MG-LX-OM-8 (8-pack)*	1000Base-LX	33210-100
10GbE	10G-SFPP-SR (1-pack) 10G-SFPP-SR-8 (8-pack)	10 Gbps SR	57-0000075-01
	10G-SFPP-LR (1-pack) 10G-SFPP-LR-8 (8-pack)	10 Gbps LR (10km)	57-0000076-01
	10G-SFPP-ER (1-pack) 10G-SFPP-ER-8 (8-pack)	10 Gbps ER (40km)	57-0000085-01
	10G-SFPP-ZR	10 Gbps ZR (80km)	57-1000180-01
	10G-SFPP-ZRD-T	10 Gbps tunable DWDM SFP+ (80km)	57-1000266-01
	10G-SFPP-TWX-0101 (1-pack) 10G-SFPP-TWX-0108 (8-pack)	1m Twinax copper cable	58-1000026-01
	10G-SFPP-TWX-0301 (1-pack) 10G-SFPP-TWX-0308 (8-pack)	3m Twinax copper cable	58-1000027-01
	10G-SFPP-TWX-0501 (1-pack) 10G-SFPP-TWX-0508 (8-pack)	5m Twinax copper cable	58-1000023-01
	10G-SFPP-AOC-0701	10GbE SFP+ Direct Attached Active Optical Cable, 7m, 1-pack	57-1000273-01
	10G-SFPP-AOC-1001	10GbE SFP+ Direct Attached Active Optical Cable, 10m, 1-pack	57-1000274-01
	10G-SFPP-USR	10GE USR SFP+ optic (LC), target range 100m over MMF, 1-pack	57-1000130-01
40GbE	40G-QSFP-QSFP-C-0101	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 1m, 1-pack	58-0000041-01
	40G-QSFP-QSFP-C-0301	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 3m, 1-pack	58-0000042-01
	40G-QSFP-QSFP-C-0501	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 5m, 1-pack	58-0000043-01
	40G-QSFP-4SFP-C-0101	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 1m, 1-pack	58-0000051-01
	40G-QSFP-4SFP-C-0301	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 3m, 1-pack	58-0000052-01
	40G-QSFP-4SFP-C-0501	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 5m, 1-pack	58-0000053-01
	40G-QSFP-SR4	40 GbE SR4 optic	57-1000128-01

Speed	FRU and Optics SKU	Description	Part Number
	40G-QSFP-SR4-INT	40 GbE SR4 (4×10 GbE SFPP break-out capable) Breakout optical cable is not included with this optics	57-1000129-01
	40G-QSFP-SR-BIDI	40 GbE QSFP+ Bi-Directional 100m optics	57-1000339-01
	40G-QSFP-ESR4	40GBase-eSR4 QSFP+ optic (MTP 1x12) 300m over MMF, (10GBASE-SR compatible, breakout), 1-pack	57-1000296-01
	40G-QSFP-ER4	40 GbE 40Km optic	57-1000327-01
	40G-QSFP-LR4	40 GbE 10Km optic	57-1000263-01
	40G-QSFP-LM4	40 GbE 140m multi-mode or 2km single-mode optic	57-1000325-01
	40G-QSFP-QSFP-AOC-1001	40GE Direct Attached QSFP+ to QSFP+ Active Optical Cable, 10m, 1-pack	57-1000306-01
	40G-QSFP-4SFP-AOC-1001	4x10GE Direct Attached QSFP+ to 4 SFP+ Active Optical Breakout Cable, 10m, 1-pack	57-1000307-01
8G FC	XBR-000163 (1-pack) XBR-000164 (8-pack)	8G FC SWL	
	XBR-000153 (1-pack) XBR-000172 (8-pack)	8G FC LWL	
	XBR-000174	8G FC ELWL	
16G FC	XBR-000192 (1-pack) XBR-000193 (8-pack)	16G FC SWL	
	XBR-000198 (1-pack) XBR-000199 (8-pack)	16G FC LWL	
FC QSFP	XBR-000245	4x8G or 4x16G FC QSFP breakout. VDX 6740T, 6740T-1G, 2740 and 2746 only (not applicable for VDX 6740).	
100GbE	100G-CFP2-SR10 (1-pack)	100 GbE CFP2 optic, SR10, for distances up to 100 m over MMF	57-1000284-01
	100G-CFP2-LR4-10KM	100 GbE CFP2 optic, LR4, for distances up to 10 km over SMF	57-1000285-01
	100G-CFP2-ER4-40KM	100 GbE CFP2 optic, ER4, for distances up to 40 km over SMF	57-1000328-01
	100G-QSFP28-SR4	100 GbE SR4 QSFP28 optic for distances up to 100m over MMF. Supported on VDX6940-144S and VDX8770-4/8 platforms.	57-1000326-01
	100G-QSFP28-LR4L-2KM	100 GbE QSFP28 optic for distances up to 2 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms.	57-1000329-01
	100G-QSFP28-LR4-10KM	100 GbE QSFP28 optic for distances up to 10 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms.	57-1000334-01

Note: 100G QSFP28 SR4 optic use core-12 cables, same cables that are used for 40G QSFP optics.

The following 10GbE CWDM optics from Smartoptics are supported on VDX 6740, 6940-144S and 8770. Please note that these are not Extreme parts and is a reference sale. So, the parts needs to be

purchased directly from SmartOptics. **The mark * one is qualified by Extreme.**

Smartoptics 10GbE CWDM SKU	Description
SO-10GE-ZR-C47	10 Gbps CWDM 1470 nm wavelength (70 km)*
SO-10GE-ZR-C49	10 Gbps CWDM 1490 nm wavelength (70 km)
SO-10GE-ZR-C51	10 Gbps CWDM 1510 nm wavelength (70 km)
SO-10GE-ZR-C53	10 Gbps CWDM 1530 nm wavelength (70 km)
SO-10GE-ZR-C55	10 Gbps CWDM 1550 nm wavelength (70 km)*
SO-10GE-ZR-C57	10 Gbps CWDM 1570 nm wavelength (70 km)
SO-10GE-ZR-C59	10 Gbps CWDM 1590 nm wavelength (70 km)
SO-10GE-ZR-C61	10 Gbps CWDM 1610 nm wavelength (70 km)*

Note: The Smartoptics require at least 20km distance or the appropriate attenuation in order for ISL to form.

The VDX 6940x, VDX 8770, and VDX 6740x switches also support the following Quad to Serial Small Form Factor Pluggable Adapters:

Mellanox MAM1Q00A-QSA	Quad to Serial Small Form Factor Pluggable Adapter which can be used with following P/Ns: 10G-SFPP-SR (10G SR) 10G-SFPP-USR (10G USR) 10G-SFPP-LR (10G LR) 10G-SFPP-ER (10G ER) 10G-SFPP-AOC-0701 (10G AOC 7m) 10G-SFPP-AOC-1001 (10G AOC 10m) 10G-SFPP-TWX-0101 (10G 1m Twinax cable) 10G-SFPP-TWX-0301 (10G 3m Twinax cable) 10G-SFPP-TWX-0501 (10G 5m Twinax cable)
CFP2 to QSFP28 conversion module (PN: 80-1008646-01)	CFP2 to QSFP28 conversion module connects the QSFP28 optic (100G optic) in a CFP2 capable port of 2/6x100G line cards in VDX8770-4/8 chassis.

*Note: Legacy Foundry Networks branded optics are not supported.

Note: 100G QSFP28 SR4 optic used in the CFP2 to QSFP28 conversion module uses core-12 cables, same cables that are used for 40G QSFP optics.

SOFTWARE UPGRADE AND DOWNGRADE

Image filenames

Download the following images from www.extremenetworks.com

Image Filename	Description	Supported Device or Module
nos7.0.2d.tar.gz	Network OS v7.0.2d for unix	NA
nos7.0.2d.zip	Network OS v7.0.2d for Windows	NA
nos7.0.2d_all_mibs.tar.gz	Network OS v7.0.2d MIBS	NA
NOS_7.0.2d_Release_Notes	Network OS v7.0.2d Release Note (PDF)	NA
nos7.0.2d.md5	**Network OS v7.0.2d MD5 Checksum**	**NA**

Upgrade/Downgrade considerations

Starting with Network OS v6.0.0, an Extreme 4GB USB drive is required for firmware installation using USB. Extreme 2GB USB drives are not supported.

Migration Path

Recommended upgrade/downgrade migration paths in both fabric cluster and logical chassis cluster modes are summarized in table below.

Note: Firmware download is not available for identical release numbers, such as Network OS 7.0.0 to Network OS 7.0.0.

To From	**5.0.2a**	**6.0.1a**	**6.0.1a2**	**6.0.2**	**7.0.0**	**7.0.2c**	**7.0.2d**
5.0.2a	NA	coldboot	coldboot	coldboot	coldboot default-config	coldboot default-config	default- config
6.0.1a	coldboot	NA	ISSU	ISSU	coldboot	coldboot	coldboot
6.0.1a2	coldboot	coldboot	NA	ISSU	coldboot	coldboot	coldboot
6.0.2	coldboot	coldboot	coldboot	NA	coldboot	Coldboot	coldboot
7.0.0	coldboot default-config	coldboot	coldboot	coldboot	NA	coldboot	coldboot
7.0.2c	coldboot default-config	coldboot	coldboot	coldboot	coldboot**	coldboot	ISSU
7.0.2d	default- config	coldboot	coldboot	coldboot	coldboot	ISSU	NA

NOTES

1. ** CFP2 to QSFP28 conversion module (PN: 80-1008646-01) Version3 downgrade to any release prior to NOS7.0.1 will cause CRC errors on the link.
2. Only Extreme Network Advisor (BNA) v14.0.1 (available separately) supports NOS v7.0.1. It is required to first upgrade to BNA v14.0.1 and then upgrade switches to Network OSv7.0.1.
3. Before downgrading to lower releases, it is recommended to disable all new features that are not supported on lower releases by using the "no" version of the CLIs. Stray configurations left out before downgrade can cause undesired behavior.
4. While upgrading chassis based system, under stress condition (e.g. due to excessive processing load on the processor), some linecards may become faulty during firmware download. To recover, run "power off <linecard>" followed by "power on <linecard>" command.
5. You must remove the IGMP snooping static mrouter configuration from all VLANs before upgrading or downgrading from or to the Network OS 6.0.2x release.
6. **Limitations:
    a) After downgrading from Network OS v7.0.1 to Network OS v5.0.x FCoE devices may not log back in or FCoE configuration may be lost. To recover, reload the switch. Alternate recovery method: re-configure FCoE by removing and adding fcoeport configuration (no fcoeport/fcoeport default) on the affected interfaces.
    b) In rare occurance, 40G links may not come up online after upgrade to 7.0.1, need to do shut/no shut to recover.
    c) In VDX 8770 platforms, After upgrade form 6.0.2 to 7.0.1 with coldboot, SNMP V3 traps are not received for the V3host which is under Rbridge.
    d) Dport test between VDX 6740T and VDX 6940-144S breakout link may fail in upgrade to 7.0.1.

Management IP connectivity

In regards to SNMP, firmware downgrade from Network OS v7.0.x to v6.0.0/v6.0.1x/v5.0.x that do not support "use-vrf" keyword, the host/v3host with use-vrf value as "default-vrf" or "user-defined vrf" is not supported. The host/v3host configuration should set the use-vrf value as "mgmt-vrf" before downgrade.

Also, firmware downgrade from Network OS v7.0.1 to v6.0.0/v6.0.1x/v5.0.x with use-vrf option in host/v3host set to user-defined vrf is not supported. The host/v3host configuration should set the use-vrf value as "mgmt-vrf" or "default-vrf" before downgrade.

Firmware upgrade to v7.0.1 from v6.0.0/v6.0.1x/v5.0.x that do not support "use-vrf" keyword will modify the host/v3host configuration to append "use-vrf" keyword with value of mgmt-vrf and all the existing host/v3host entries will be assigned to mgmt-vrf.

Similarly on downgrade, the "use-vrf" keyword will be automatically removed from the configuration & depending upon the version, it will be put into mgmt-vrf.

The above downgrade/upgrade restrictions holds good for other IP services like Syslog-server, sFlow, NTP, Radius, TACACS and LDAP

For users in 5.x that have configured Inband Management over VE interfaces, may expect to see the configuration fall into Default VRF, however, as noted above, the "use-vrf" keyword pointing to mgmt-vrf will be appended & applied. Thus such customers would need to modify the configuration after upgrade to adapt it according to their needs.

For HTTP services, firmware upgrade to v7.0.1 will add two entries by default under http configuration with "use-vrf" keyword appended with value as "mgmt-vrf" and other entry as "default-vrf".

Firmware downgrade to v6.0.1/6.0.2 with http server on user-defined vrf is not supported. Http server configuration on user-defined vrf should be removed before downgrade.

Firmware downgrade to v6.0.0 or v5.0.x that do not support "use-vrf" keyword, the http server configuration on default-vrf and user-defined vrf are not supported. Http server configuration on default-vrf and user-defined vrf should be removed before downgrade.

Firmware Installation

In fabric cluster mode

- The “firmware download” command is required to be executed by logging on to each individual node.
- Under certain stress conditions firmware download might time out on a node, (e.g. due to excessive processing load on the processor). The firmware download command will recover the system automatically. It is required to wait for completion of recovery before retrying the firmware download command.
- While upgrading firmware on the node, it is recommended not to make any configuration changes before firmware download has been completed successfully.

In logical chassis cluster mode

- The “firmware download logical-chassis” command can be used from the principal node to upgrade one or more nodes in the cluster.
    - Under certain stress conditions firmware download might time out on some nodes, (e.g. due to excessive processing load on the processor) it is recommended to re-run the logical-chassis firmware download command to upgrade these failed nodes and bring their firmware level to be the same as the rest of nodes first before activating any of them.
    - While upgrading the cluster, it is recommended not to make any configuration changes in the cluster until all of the nodes have been upgraded to the same firmware. Otherwise, it may cause cluster segmentation.
    - The firmware download command can also be executed on individual nodes. In such

a case, please follow the procedure from Fabric cluster mode.

General information on installing Extreme Network OS can be found in the Extreme *Network OS Administrator's Guide.* This section includes special considerations and caveats to be aware of

when upgrading to or from this version of Extreme Network OS, as well as recommended migration paths to use to reach this version of Extreme Network OS.

Note: Installing Extreme Network OS may be service disruptive and any unsaved running configuration may be lost during the process. In Fabric cluster mode, running-config needs to be saved to startup-config in order to preserve the running-config across reboots. In Logical Chassis mode, running-config is always preserved across reboots. The firmware version migration path determines if the configuration across upgrade/downgrade shall be preserved.

Upgrading to this Release (Best Practices)

In logical chassis cluster mode it is required to upgrade Principal switch at the end if all nodes in the cluster are not upgraded at the same time.

A. Upgrade all nodes in the cluster at same time -- Service Disruptive ClusterWide

- Download the firmware on all the switches running Network OS v7.0.1 using the coldboot option.
- After all switches complete the firmware download, they will be automatically rebooted.
- Since all nodes reboot at the same time, this procedure is service disruptive.

B. Upgrade Odd/Even Nodes (one segment at a time)—Lossless Upgrade:

- This is the most recommended procedure for lossless upgrade. This requires servers to be dual homed.
- Download the firmware in all the odd nodes running Network OS with the coldboot option.
- After these switches complete the firmware download, they will be rebooted automatically. After they boot up, half of the cluster is now on the latest version. Traffic resumes and passes through the other half of the cluster.
- Now download the firmware in all even nodes with the coldbootoption.
- After these switches complete the firmware download, they will be rebooted automatically. After they boot up, the entire cluster is loaded with latest image and up and running

C. Upgrade one node at a time -- Service Disruptive at Node level in theCluster

- Download the firmware in the switch nodes one node at a time in cluster running Extreme NOS 6.0.x using the coldboot option. Principal node in a cluster should be last to be upgraded.
- After a node is upgraded, it will join the existing Network OS v7.0.1 cluster and form Fabric cluster. Eventually, when all the nodes are upgraded, they will form one Network OS 7.0.1 VCS Cluster. The data path will remain intact in Fabric cluster. [Note that no configuration changes are allowed during this time.]

Downgrading to a Previous Release

- In normal circumstances, the SW/0 partition is Active. When an ISSU performed, the SW/1 partition becomes active. In order to ensure config is retained during coldboot downgrade, it is important to have SW/0 partition Active before downgrade. The SW/0 partition can be made Active by reloading the switch before initiating firmware downgrade.
- Alternative: Execute a coldboot downgrade with SW/1 Active.
  - Back-up the config to external server by “copy running file” (for logical chassis cluster) or “copy running start” (for fabric cluster).

- Execute a coldboot downgrade. In FC mode, the startup-config file will be re-applied automatically. In LC mode, copy the 6.0.1x config back by executing “copy filerunning”.

Upgrade/downgrade Considerations for vLAG deployments

There are 2 approaches by which vLAG nodes can be upgraded.

- **Approach 1**: Graceful shutdown of vLAG ports on one node at atime.
- **Approach 2**: Static vLAGs and Dynamic vLAGs without configurationchanges.

vLAG deployment upgrade Illustration

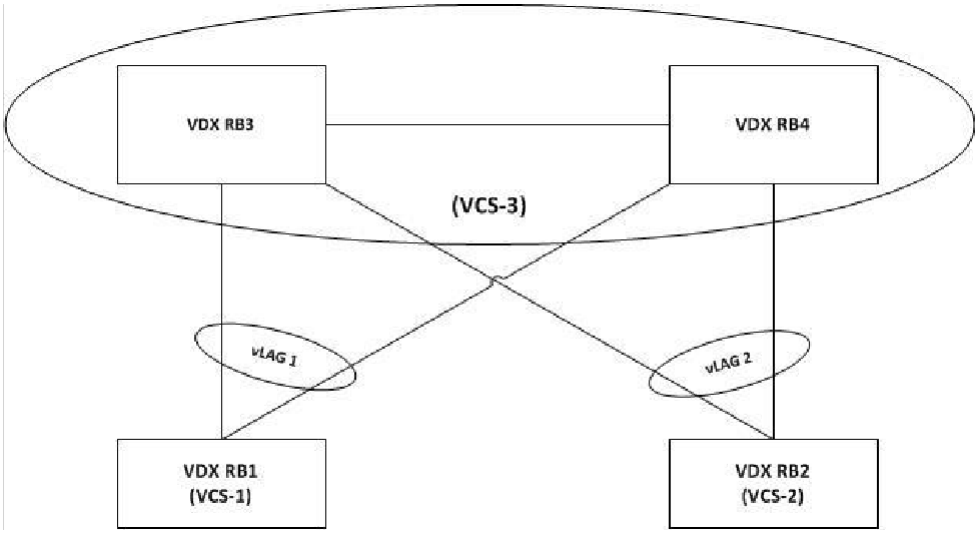

Approach 1: Graceful shutdown of vLAG ports on one node at a time.

Step 1: If in FC mode, shut the port-channel associated with vLAG 1 on RB3. With LC mode, shutting down port-channel takes down entire port-channel including port-channel interfaces on remote RBs. Therefore, if in LC mode, shut all the member ports of the vLAG 1 on RB3.

Step 2: Save running configuration to startup-configuration if VCS is in FC mode. This reduces the vLAG into a single node vLAG/port-channel on RB4. Note: if the vLAG is in static mode, all members of the port-channel should be shutdown. This is due to the static LAG behavior where it may bring up the member links even if the port-channel is admin shut.

Step 3: Upgrade RB3 to the desired Network OS version.

Step 4: After RB3 has rebooted from the Network OS upgrade and is operational, repeat step 1 and 2 on RB4. **Warning:** there will be a complete impact to the data path on vLAG 1 at this time.

Step 5: Promptly perform “no shutdown” on all the interfaces that were shut in step 1 and 2 on RB3. **Note:** if the vLAG is in static mode, it is required to perform “no shutdown” on all the shutdown members of the port-channel.

Step 6: Upgrade RB4 to the desired Network OS version.

Step 7: After RB4 has rebooted after Network OS upgrade and is operational, promptly perform "no shutdown" on all the interfaces that were shut in step 1 and 2 on RB4.

Step 8: Verify RB3 and RB4 were successfully upgraded to the desired Network OS version and the vLAG on RB3 and RB4 was re-established and operational with traffic forwarding.

Step 9: If VCS is in FC mode, perform a "copy running-configuration startup-configuration" on RB3 and RB4 to return the startup-configuration back to the original configuration.

Advantages

- Clean upgrade
- No duplicate primary port issues
- Works well for both static and dynamic vLAGs.

Disadvantages

- Requires manual execution by administrator to perform shutdown/no shutdown on port-channel, allowing for human errors particularly with large numbers of vLAGs.
- Requires precise and efficient execution.
- Impact to the data path for a very small period of time when the vLAG is shut on the second node (RB4).

Approach 2: Static vLAGs and Dynamic vLAGs without configuration changes.

Step 1: Upgrade RB3 to the desired Network OS version and reboot. There are two possible behaviors depending on the *ignore-split* configuration as follows:

Ignore-split on (default): No impact/reconvergence to Static or Dynamic vLAGs. Minimal data path impact observed.

Ignore-split off: For Dynamic vLAGs,

- if RB3 is the primary vLAG node, observe vLAG flap and a few seconds of data path impact.
- if RB3 is not the primary vLAG node, there will be minimal data path impact but no vLAG reconvergence.

Step 2: After RB3 has rebooted from the Network OS upgrade and is operational, RB3 will re-join the vLAG.

Step 3: Upgrade RB4 to the desired Network OS version and reboot. There are two possible behaviors depending on the *ignore-split* configuration as follows:

Ignore-split on (default): No impact/reconvergence to Static or Dynamic vLAGs. Minimal data path impact observed.

Ignore-split off: For Dynamic vLAGs:

- If RB4 is the primary vLAG node, observe vLAG flap and a few seconds of data path impact.
- If RB4 is not the primary vLAG node, there will be minimal data path impact but no vLAG reconvergence.

Step 4: After RB4 has rebooted from the Network OS upgrade and is operational, RB4 will re-join the vLAG with the three possible behaviors as follows:

Advantages:

- No manual administrative configuration required.
- Straightforward upgrade process, no special handling for vLAGs.

Disadvantages:

- Data path impact as detailed above.

Upgrade/downgrade with default configuration

Step 1: Copy and save the running configuration to the RBridge flash or FTP server.
Step 2: If default-config option is available in firmware download command in the active NOS version on the switch, execute firmware download using default-config. If default-config option is not available perform copy default configuration to startup configuration.
Step 3: If the VCS is in FC mode, reboot the RBridge manually. If the VCS is in LC mode, all the RBridge(s) in the VCS will reboot automatically.
Step 4: Downgrade the RBridge(s) to the desired Network OS version and reboot the RBridge(s).
Step 5: Restore the original configuration file by copying the configuration saved in step 1 back to the running-configuration (Individually on each RBridge in FC mode, and from principal RBridge if in LC mode)
Step 6: In FC mode, save the configuration by performing copy running-configuration to the startup-configuration. In LC mode, configuration is automatically saved and is persistent.

Management Gateway IP changes

VDX Fixed-form switches (No L3 license required)

Starting with Network OS v5.x, Management Gateway IP can only be configured under Rbridge-Id context/vrf mgmt-vrf as follows:

```
SW(config)# rbridge-id <RBridge#>
SW(config-rbridge-id-<RBRidge#>)# vrf mgmt-vrf
SW(config-vrf-mgmt-vrf)# address-family ipv4 unicast
SW(vrf-ipv4-unicast)# ip route 0.0.0.0/0 <GW IP Address>
```

Note:

After upgrading to Network OS v5.x or above, remove the old Gateway using “no ip route” command and configure the new route with higher metric to avoid forming ECMP with old and new gateways.

VDX 8770 (with L3 license/without L3 license)

Prior to Network OS v4.0.0, Management Gateway could be configured in two ways based on the availability of L3 license on the node.

- L3 license installed: Configure using command "ip route 0.0.0.0/0 <gateway ip>". Using the command "ip gateway-address" under the management interface will display anerror.
- L3 license not installed: Configure using command "ip gateway-address" under the management interface.

In Network OS v4.0 there is only one option to configure the gateway that is "ip route 0.0.0.0/0 <gateway ip>".

Note:

After upgrading to Network OS v4.0.1 or above, it is required to remove the old Gateway using “no ip route” command and configure the new route with higher metric to avoid forming ECMP with old and new gateways.

Management Services

Telnet, SSH and AAA VRF support

Starting with Network OS 7.0.0, support for TELNET, SSH and AAA (RADIUS, TACACS+ and LDAP) on user defined / default vrf is provided.

CLI Changes for Telnet, SSH, AAA

The following CLI has an additional parameter “use-vrf” to support these features.

[no] ssh server use-vrf <vrf-name> [shutdown]

[no] telnet server use-vrf <vrf-name> [shutdown]

[no] ldap-server host <IPv4|IPv6|hostname> [use-vrf <VRF name>]

[no] tacacs-server host < IPv4|IPv6|hostname > [use-vrf <VRF name>]

[no] radius-server host < IPv4|IPv6|hostname > [use-vrf <VRF name>]

HTTP VRF support

HTTP/HTTPS services are supported on user-defined VRF and default-vrf in addition to mgmt-vrf. CLI option use-vrf is introduced to enable/disable HTTP/HTTPS services on user-defined/default-vrf.

[no] http server use-vrf <vrf-name> shutdown

NTP VRF support

Starting with Network OS 7.0.0, support for NTP on user defined / default vrf and MGMT-VRF in Inband is provided

CLI Changes for NTP

The following CLI has an additional parameter “use-vrf” to support this feature.

[no] ntp server < IPv4|IPv6|hostname > [use-vrf] <mgmt-vrf | default-vrf | non-default-vrf >]

SNMP- Community string maximum length increased to 64:

Maximum length for community string is increased from 16 to 64 characters.

SNMP - Support for traps during hafailover:

Cpstatuschange trap will be triggered during hafailover with cpLastEvent as hafailoverstart and hafailoverdone to notify that hafailover is started and hafailover is completed in the switch.

SNMP-Trap Source IP support:

CLI option source-interface is introduced in host/v3host commands to select the loopback/ve interface IP as source IP in traps.

[no] snmp-server host ip-address <community-string> source-interface {

loopback number|ve vlan_id}]

[no] snmp-server v3host ip-address <username> source-interface {

loopback number|ve vlan_id}]

SNMP context based query:

A single SNMP agent can be supported by multiple instances of the same MIB module by mapping the context name to a virtual routing and forwarding (VRF) instance created within the switch. Each VRF is mapped with a specific key called context name. The context name is used to identify the VRF and fetch the MIB details of the mapped VRF from the underlying modules. In case of snmp v1 and v2c, we need to map the community with the context name.

[no] snmp-server context <context_name> vrf <vrf_name>

[no] snmp-server mib community-map <community-name> context <context-name>

SNMP MIB – VLAN update

During an snmpwalk or snmpgetbulk, all the VLAN interfaces are filtered out from the IF MIB output. Similarly, there is an object “ifNumber” that tells the number of interfaces in the system. The “ifNumber” object is also correspondingly reduced by this number.

SNMP Trap VRF Support

SNMP is able to receive the packets from any VRF including mgmt-vrf/default-vrf and respond to the corresponding VRF from where the SNMP packet is received. The support is also added to send the notification (trap) to the host/v3host configured in the switch through the vrf-name mapped with the host/v3host.

SNMP-Trap CLI

CLI option use-vrf is introduced to get the vrf-id for each client. This option is applicable for both SNMP V1/V2c and V3 versions in host/v3host commands.

[no] snmp-server host ip-address community <comm-string> use-vrf <vrf-name>

[no] snmp-server v3host ip-address <username> [notifytype traps | informs] use-vrf <vrf-name>

To disable per link TRAP under interface

[No] snmp trap link-status

SNMP – IF MIB

To display Interface details when linecard is powered-off

[No] snmp-server offline-if enable

Sflow VRF Support

Sflow can be configured to point to collector in either default-vrf, mgmt-vrf, or non-default vrf..

Sflow-CLI

CLI option use-vrf is introduced to assign the vrf-id for each client.

```
[no] sflow collector <ipv4/ipv6 address>  <port> [use-vrf] <mgmt-vrf | default-vrf | non-default-vrf >
```

Syslog VRF Support

Syslog servers logging can be configured to point to syslog servers in default-vrf, mgmt-vrf, or non-default vrf.

Syslog-CLI

CLI option use-vrf is introduced to get the vrf-id for each client.

```
[no] logging syslog-server <ipv4/ipv6 address> use-vrf <mgmt-vrf | default-vrf | non-default-vrf > [secure [port <xxxx>]]
```

LDAP Support

We support Windows LDAP Active Directory 2008 and 2012

Firmware download, Copy support, Copy config

The use-vrf option is introduced to these commands to specify the name of VRF where the server resides.

Other Management Services

Other management services like REST, Netconf, HTTP, SNMP MIB's would be available in default, user defined and management VRFs.

SCALABILITY AND INTEROPERABILITY

Scalability numbers

All scalability limits are subject to change. The limits noted in this section apply to all the platforms listed unless otherwise specified.

NOS v7.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
Maximum # of dot1Q VLANs (Virtual-Fabric Disabled)	4096	4096	4096	4096
Maximum # of VLANs (dot1Q + Virtual-Fabric)	6000	8192	8192	8192
Maximum # of Service Virtual Fabric VLANs	2000	4096	4096	4096
Maximum # of Transport Virtual Fabric VLANs	1000	1000	1000	1000
Maximum # of MAC addresses per Switch	120000	256000	75000	75000
Maximum # of MAC addresses per Fabric (with CML)	256000	256000	256000	256000
Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for VMware NSX	8000	N/A	8000	8000
Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for Virtual-Fabric Extension	120000	N/A	75000	75000
Maximum # of MAC-based Virtual-Fabric VLAN Classification per switch	256	1024	1000	1000
Maximum # of Classified Virtual Fabric VLANs per Trunk Interface	2000	4096	4096	4096
Maximum # of port profiles (AMPP)	1000	1,000	512	512
Maximum # of VLANS in port profiles	3500	4000	3500	3500
Maximum # of sites (tunnels) in Virtual-Fabric Extension	50	N/A	50	50
Maximum # of dot1q VLANs that can be attached on VxLAN GW for Virtual-Fabric Extension	4000	N/A	4000	4000
Maximum # of Virtual-Fabric (Service + Transport) VLANs that can be extended via Virtual-Fabric Extension	2000	N/A	4000	4000
Maximum # of dot1q VLANs + Virtual-Fabric VLANs enabled on edge-interfaces that can be attached to VxLAN GW and extended via Virtual-Fabric Extension	(2000+1000)	N/A	(2000+1000)	(2000+1000)
Max # of IGMP groups over Tunnels via Virtual-Fabric Extension	6000	N/A	6000	6000

NOS v7.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
Max # of BFD sessions over Virtual-Fabric Extension Tunnels	10	N/A	10	10
Maximum # of dot1q VLANs that can be attached on VxLAN GW for VMware NSX	2000	N/A	2000	2000
Maximum # of VLANs (dot1q VLANs attached to VxLAN GW for NSX + Virtual Fabric VLANs enabled on edge-interfaces)	(2000+1,000)	N/A	(2000+1000)	(2000+1000)
Maximum # of VxLAN tunnels with VMware NSX	250	N/A	250	250
Maximum # of service-nodes with VMware NSX	5	N/A	5	5
Maximum # of MAC Associations for AMPP	8000	4000	8000	8000
Maximum # of per priority pause levels	3	8	3	3
Maximum # of VMware vCenters per Fabric	4	4	4	4
Maximum # of ELD instances in the fabric	2000	2000	2000	2000
Maximum # of IGMP Snooping Interfaces supported	512	512	512	512
Learning rate for IGMP snooping (groups/second)	512	512	512	512
Maximum # of L2 (IGMP Snooping) multicast groups	6000	6000	6000	6000
Maximum # of MLD Interfaces	256	256	256	256
Maximum # of MLD Groups	4000	4000	4000	4000
Learning rate for MLD snooping (groups/second)	512	512	512	512
# of L3 (S,G) forwarding Entries	2000	2000	2000	2000
# of L3 (*,G) forwarding Entries	256	256	256	256
PIM Interfaces Supported	32	32	32	32
IGMP interfaces supported	32	32	32	32
Learning Rate for PIM-SM (flows/second)	32	32	32	32
Maximum # of L2 ACL(ingress/egress) *	3000/120	12000/2000	6128/496	6128/496
Maximum # of L3 ACL ipv4 (ingress/egress) *	1500/120	12000/2000	3064/475	3064/475
Maximum # of class-maps	2048	2048	2048	2048
Maximum # of policy-maps	2048	2048	2048	2048
Maximum # of class-maps per policy map	50	50	50	50
Maximum Total # of L3 ACL ipv6 (ingress/egress) *	500/120	4000/2000	1000/500	1000/500
Maximum # of VF/FCoE interfaces/Logins (Per switch)	1000	1000	1000	1000
Maximum # of Enodes/FCoE Devices per Fabric	2000	2000	2000	2000
Maximum # of NPIV per Port	64	64	64	64
Maximum # of SAN Devices (FC + FCoE) per Fabric	3000	3000	3000	3000

NOS v7.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
Maximum # of MSTP instance	32	32	32	32
Maximum # of VLAN in PVST	128	128	128	128
Maximum # of LAGs (Port Channels)	64	288	144	144
Maximum # of members in a standard LAG	16	16	16	16
Maximum # of members in a Extreme Trunk (10G)	16	8	12	12
Maximum # of members in a Extreme Trunk (40G)	2	NA	3	3
Maximum # of members in a Extreme Trunk (100G)	NA	NA	NA	NA
Maximum # of switches in a Fabric cluster mode **	48	48	48	48
Maximum # of switches in Logical cluster mode **	48	48	48	48
Maximum # of L2 ECMP Paths	16	8	16	16
Maximum # of vLAGs in a fabric	2000	2000	2000	2000
Maximum # of member ports in a vLAG	64	64	64	64
Maximum # of nodes in a vLAG	8	8	8	8
Maximum # of member ports per vLAG per Node	16	16	16	16
Maximum # of Management ACL	256	256	256	256
Maximum # of ARP Entries *	16000	126000	72000	72000
Maximum # of OSPF areas	20	64	20	20
Maximum # of OSPF routers in a single area	64	200	64	64
Maximum # of OSPF adjacencies	100	200	100	100
Maximum # of OSPF routes *	8,000	64,000	10000	10000
# of OSPF Interfaces	100	1,000	100	100
# of OSPF enabled subnets	100	1,000	100	100
# of local subnets in a single area	100	1,000	100	100
Maximum # of OSPFv3 areas	9	9	9	9
Maximum # of OSPFv3 routers in a single area	64	200	64	64
Maximum # of OSPFv3 adjacencies	100	200	100	100
Maximum # of OSPFv3 routes *	1500	64000	1500	1500
# of OSPFv3 Interfaces	100	256	100	100
# of OSPFv3 enabled subnets	100	256	100	100
Maximum # of IPv4 routes in SW *	8000	280000	10000	10000
Maximum # of IPv6 routes in SW *	1500	64000	1500	1500
Maximum # of IPv4 static routes *	2000	40,000	2000	2000
Maximum # of IPv6 static routes *	500	20,000	500	500
Maximum # of VRRP instances per system	255	1024	512	512
Maximum # of VRRP v3 instances per system	255	1024	512	512

NOS v7.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
Maximum # of VRRP instances per interface	32	32	32	32
Maximum # of routers participating in a VRRP-E session	8	8	8	8
Maximum # of virtual IP addresses per VRRP instance	16	16	16	16
Maximum # of FVG instances per system	256	4096	1024	1024
Maximum # of FVG instances per interface	1	1	1	1
Maximum # of routers participating in a FVG session	32	32	32	32
Maximum # of Gateway IP addresses per FVG instance	1	1	1	1
Maximum # of IPv4 routes with ECMP supported *	8000	200000	10000	10000
Maximum # of IPv6 routes with ECMP supported *	1500	64000	1500	1500
Maximum # of L3 ECMP	16	32	32	32
Maximum # of IPv4 interfaces per system *(Ve intf)	2000	4000	2000	2000
Maximum # of IPv6 interfaces per system * (Ve intf)	512	4000	512	512
Maximum # of VRF per node	32	512	512	512
Maximum # of VRFs support protocols per node	32	128	128	128
Maximum # of I-BGP peers	256	512	256	256
Maximum # of E-BGP peers	256	256	256	256
Maximum # of IPv4 BGP routes in HW *	8000	200000	10000	10000
Maximum # of IPv6 BGP routes in HW *	1,500	64000	1500	1500
Maximum # of IPv4 RIB (IN + OUT) Routes *	110000	1300000	110000	110000
Maximum # of IPv6 RIB (IN + OUT) Routes *	110000	1300000	110000	110000
Maximum # BGP IPv4/IPv6 Peer Group	100	250	100	100
Maximum # of BFD sessions per node	100	100	100	100
Maximum # of UDLD enabled interfaces	64	384	144	108
Maximum # of PVLAN domain supported	1000	1000	1000	1000
Maximum # of Secondary vlans per PVLAN supported	24	24	24	24
Maximum # of primary vlans per PVLAN supported in promiscuous mode	24	24	24	24
DHCP IP Helper Addresses per interface	16	16	16	16
DHCP IP Helper VE interfaces	256	1,000	256	256
DHCP IP Helper physical ports	60	384	60	60
DHCP IP Relay Addresses per Node	2000	4000	2000	2000
DHCP IPv6 Relay Address per Node	2000	4000	2000	2000
Max Number of configurable PBR route maps	64	64	64	64

NOS v7.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
Max Number of configurable PBR stanzas	1024	1024	1024	1024
Max Number of HW entries available for PBR	512	8192	512	512
Max Number of configurable next hops within a single PBR stanza	128	128	128	128
Max # of OpenFlow Active Connections	1	1	1	1
Max # of OpenFlow Passive Connections	1	1	1	1
Maximum # of OpenFlow L2 flows	1000	4000	879	879
Maximum # of OpenFlow L3 flows	1000	4000	879	879
Maximum # of Total OpenFlow GROUP	768	768	768	768
Maximun # of OpenFlow GROUP Type ALL	256	256	256	256
Maximun # of OpenFlow GROUP Type SELECT	256	256	256	256
Maximun # of OpenFlow GROUP Type INDIRECT	256	256	256	256
Max # of Buckets per GROUP ALL	16	16	16	16
Max # of Buckets per GROUP SELECT	8	8	8	8
Max # of Buckets per GROUP INDIRECT	1	1	1	1
Max # of ACTIONS per Bucket	3	3	3	3
Max # METERS	1024	4096	1024	1024
Maximum # of MAPS policy	10	10	10	10
Maximum # of MAPS rules	250	250	250	250
Maximum # of MAPS groups	64	64	64	64

* Parameters mentioned are applicable on specific HW profiles. Please check the Network *OS Administrator's Guide* for the specific HW profiles.

**Please consult your Extreme SE for best practices when designing a 48-node VCS Fabric. In Hybrid cluster environment (a cluster involving various VDX platforms), the scalability limit of the cluster is determined by the scalability limit of the lowest denominator. For instance, in such a fabric, if the MAC scalability limit on one VDX platform is lower than the other, then the fabric supports the lower scale value.

IP Fabric Scalability:

IP Fabric Scalability Numbers	VDX-8770	VDX-6940		VDX 6940-144s		VDX 6740, VDX 6740T
	Spine	Spine	Leaf	Spine	Leaf	Leaf
VLANS extended with VxLANs (no. of tunnels * VLANs * ECMP)	NA	NA	16k	NA	16k	16k
Software MAC entries (CML)	NA	200k	200k	200k	200k	200k
Software ARP entries (Conversational ARP)	NA	100k	100k	100k	100k	100k
Software ND entries (Conversational-ND)	NA	50k	50k	50k	50k	50k
BGP eVPN IPv4 routes	200k	200k	200k	200k	200k	200k
BGP eVPN IPv6 routes	64k	2k	2k	2k	2k	2k

BGP eVPN MAC-IP routes BGP eVPN MAC routes	100k 200k	100k 200k	100k 200k	100k 200k	100k 200k	100k 200k
Max # of IP Unnumbered interface	384	36	36	144	144	52
Max # of IP Port channel interface	384	36	36	144	144	52
Max # of members per IP Port-Channel Interface	8	8	8	8	8	8
Max # of Leaf – Spine ECMP	16	16	16	16	16	16
Max # of SAG addresses per interface	32	32	32	32	32	32

HW Profile and Platform Specific Scale Numbers

Route Profile Scale:

VDX 6740, 6740T, 6740T						
NOS v7.0.0 Scalability Numbers	ROUTE PRFILE					
	DEFAULT	IPV4-MAX-ROUTE	IPV4-MAX-ARP	IPV4-MIN-V6	IPV6-MAX-ROUTE	IPV6-MAX-ND
Maximum # of IPv4 routes with ECMP supported *	4000	8000	8000	6000	2000	2000
Maximum # of IPv6 routes with ECMP supported *	1000	0	0	500	1500	1500
Maximum # of OSPF routes *	4000	8000	8000	6000	2000	2000
Maximum # of OSPFv3 routes *	1000	0	0	500	1500	1500
Maximum # of IPv4 BGP routes in HW *	4000	8000	8000	6000	2000	2000
Maximum # of IPv6 BGP routes in HW *	1000	0	0	500	1500	1500
Maximum # of IPv4 routes in SW *	4000	8000	8000	6000	2000	2000
Maximum # of IPv6 routes in SW *	1000	0	0	500	1500	1500
Maximum # of ARP Entries *	16000	16000	16000	16000	16000	16000
Maximum # of IPv6 neighbor cache Entries *	4000	0	0	4000	4000	4000

VDX 6940-36Q, VDX 6940-144S						
NOS v7.0.0 Scalability Numbers	ROUTE PRFILE					
	DEFAULT	IPV4-MAX-ROUTE	IPV4-MAX-ARP	IPV4-MIN-V6	IPV6-MAX-ROUTE	IPV6-MAX-ND
Maximum # of IPv4 routes with ECMP supported *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 routes with ECMP supported *	1000	0	0	500	2000	2000
Maximum # of OSPF routes *	6000	10000	10000	8000	2500	2500
Maximum # of OSPFv3 routes *	1000	0	0	500	2000	2000
Maximum # of IPv4 BGP routes in HW *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 BGP routes in HW *	1000	0	0	500	2000	2000
Maximum # of IPv4 routes in SW *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 routes in SW *	1000	0	0	500	2000	2000
Maximum # of ARP Entries *	43000	49000	73000	49000	6000	6000
Maximum # of IPv6 neighbor cache Entries *	12000	0	0	10000	30000	30000

VDX 8770						
NOS v7.0.0 Scalability Numbers	PROFILE					
	DEFAULT	IPV4-MAX-ROUTE	IPV4-MAX-ARP	IPV4-MIN-V6	IPV6-MAX-ROUTE	IPV6-MAX-ND
Maximum # of IPv4 routes with ECMP supported *	65000	280000	198000	163000	20000	12000
Maximum # of IPv6 routes with ECMP supported *	16000	2000	2000	8000	64000	12000
Maximum # of OSPF routes *	64,000	64,000	64,000	64,000	20000	12,000
Maximum # of OSPFv3 routes *	16000	2000	2000	8000	64000	12000
Maximum # of IPv4 BGP routes in HW *	65000	280000	198000	163000	20000	12000
Maximum # of IPv6 BGP routes in HW *	16000	2000	2000	8000	64000	12000
Maximum # of IPv4 routes in SW *	65000	280000	198000	163000	20000	12000
Maximum # of IPv6 routes in SW *	16000	2000	2000	8000	64000	12000
Maximum # of ARP Entries *	98000	40000	129000	98000	12000	20000
Maximum # of IPv6 neighbor cache Entries *	28000	2000	2000	12000	12000	65000

L2 L3 Multicast Scale :

TCAM PROFILE DEFAULT,DYN-ARP-INS,OPENFLOW				
Maximum # of L2 (IGMP Snooping) multicast groups	1000(openflow)	6000	6000	6000
Maximum # of MLD Groups	0	512	4000	4000
# of L3 (S,G) forwarding Entries	2000	2,000	2000	2000
# of L3 (*,G) forwarding Entries	256	256	256	256

TCAM PROFILE IPV4-IPV6-MCAST				
Maximum # of L2 (IGMP Snooping) multicast groups	1000	6000 (16000)	6000	6000
Maximum # of MLD Groups	500	4000 (8000)	4000	4000
# of L3 (S,G) forwarding Entries	2,000	2,000	2000	2000
# of L3 (*,G) forwarding Entries	256	256	256	256

ACL Scale:

VDX8770-4									
NOS v7.0.0 Scalability Numbers	TCAM PROFILES								
	DEFAULT	DNY-ARP-INSP	IPV4-ACL	IPV4-V6-MCAST	IPV4-V6-PBR	IPV4-V6-QOS	L2-ACL-QOS	L2-IPV4-ACL	OPEN FLOW
Maximum # of L2 ACL(ingress/egress) *	16000/ 2000	12000/ 2000	512/1016	500/ 1000	500/ 1000	500/ 1000	32000/ 2000	16000/ 2000	12000/ 2000
Maximum # of L3 ACL ipv4 (ingress/egress) *	16000/ 2000	16000/ 2000	51000/ 2000	500/ 2000	8000/ 2000	8000/ 2000	5000/ 2000	24500/ 2000	12000/ 2000
Maximum # of L3 ACL ipv6 (ingress/egress) *	500/2000	500/2000	0/2000	500/ 2000	4000/ 2000	4000/ 2000	0/ 1000	0/ 2000	500/ 2000

VDX6940									
NOS v7.0.0 Scalability Numbers	TCAM PROFILES								
	DEFAULT	DNY-ARP-INSP	IPV4-ACL	IPV4-V6-MCAST	IPV4-V6-PBR	IPV4-V6-QOS	L2-ACL-QOS	L2-IPV4-ACL	OPENFLOW
Maximum # of L2 ACL(ingress/egress) *	500/256	500/256	NA	500/256	0/0	0/0	3000/256	1500/256	500/256
Maximum # of L3 ACL ipv4 (ingress/egress) *	1000/256	1000/256	NA	500/256	500/256	500/256	1000/256	1500/256	500/256
Maximum # of L3 ACL ipv6 (ingress/egress) *	500/256	500/256	NA	500/256	500/256	500/256	0/256	500/256	0/256

VDX6740									
NOS v7.0.0 Scalability Numbers	TCAM PROFILES								
	DEFAULT	DNY-ARP-INSP	IPV4-ACL	IPV4-V6-MCAST	IPV4-V6-PBR	IPV4-V6-QOS	L2-ACL-QOS	L2-IPV4-ACL	OPENFLOW
Maximum # of L2 ACL(ingress/egress) *	500/120	500/120	500/120	0/0	0/0	0/0	3000/120	1000/120	500/120
Maximum # of L3 ACL ipv4 (ingress/egress) *	500/120	500/120	500/120	500/120	500/120	500/120	0/120	1500/120	500/120
Maximum # of L3 ACL ipv6 (ingress/egress) *	500/120	500/120	500/120	500/120	500/120	500/120	0/120	0/120	0/120

Compatibility and Interoperability

The following tables list the devices tested for IP storage, FC/FCoE storage and host adapters for VDX as of Network OS v7.0.2. This is a representative list of devices, Network OS v7.0.2 supports all standards-based devices connected to it for these types of storage.

IP Storage

Vendor	Storage Array Model	Protocol	Switch Model	Initiator
EMC	Isilon	NAS	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6
EMC	VG2	NAS	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6
EMC	VNX 5300	iSCSI	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6
EMC	VMAX 40K	iSCSI	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6
HDS	4060	NAS	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6
HDS	4060	iSCSI	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6
NetApp	3170	NAS	6740	Windows 2008 R2, Windows 2012 R2, ESXi 5.5u2, RHEL 6.6

FC/FCoE Storage

Vendor	Storage Array Model	Protocol	Switch Model	Initiator
Netapp	FAS3250-cdot	FC, FCoE	6740, 8770 (FCoE only)	Windows 2012, VMWare
HDS	R800	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	R700	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	HUSVM	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	DF850	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	DF800	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
EMC	CX4-120	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2
	VMAX 40K	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2, Windows 2012
	VNX-5300	FC. FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2
	VNX-5500	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2, Windows 2012
	VSP	FC, FCoE	6740	RHEL 6.5, Windows 2012
IBM	DS8100	FC	6740/T	Windows 2012 R2
	Flash 840	FC	6740/T	Windows 2012 R2
	XIV	FC	6740/T	Windows 2012 R2
HP	MSA2040	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P10000	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P6500	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P6300	FC, FCoE	6740, 8770 (FCoE only)	RHEL 7.0, Windows 2012, Windows 2012 R2
	P4330	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P9500	FC, FCoE	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2

Network Adapters

Vendor	Model	Protocol	Switch Model	OS
HP	526FLR	FCoE	6740	Windows 2012. SuSE 12
	554FLR	FCoE	6740	RHEL 7.0, Windows 2008 R2 SP1, RHEL 6.5
	CN1000E	FCoE	6740, 8770	RHEL 7.0, SuSE 12
	CN1000R	FCoE	6740, 8770	Windows 2012 R2, VMWare ESXi 5.5
	CN1000Q	FCoE	6740, 8770	Windows 2012, RHEL 6.6
	CN1100R	FCoE	6740	Windows 2012 R2,
	CN1000Q	FCoE	6740	Windows 2012, RHEL 5.1
	CN1000E	FCoE		RHEL 6.5
Emulex	OCe10102	FCoE	6740	RHEL 6.5
	LPe16002	FC	6740	RHEL 6.5, Windows 2008, Windows 2012
	LPe16202	FCoE	6740	RHEL 6.5
	90Y3556 (IBM)	FCoE	2740	Windows 2008 R2, Windows 2012 R2
	OCe14102	FCoE	6740	Windows 2012 R2, RHEL 6.5
	OCe11002-FM	FCoE	6740	Windows 2008 R2, RHEL 6.4
	90Y3556	FCoE	6740	Windows 2012 R2, Windows 2008 R2
Qlogic	1020	FCoE	6740	Windows 2012
	1860	FCoE	6740	RHEL 6.5, 6.3, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1, Solaris 10
	2672	FC	6740	RHEL 6.5, Windows 2008
	8152	FCoE	6740	ESX 5.1
	8142	FCoE	6740	Windows 2012, RHEL 6.5
	2672	FC	6740	RHEL 6.5
	2762	FC	6740	RHEL 5.1, Windows 2012

ADDITIONAL CONSIDERATIONS

Limitations and Restrictions

Command Line Interface

- Break command is not supported. ctrl-c can be used as an alternative.
- Few commands may not display paginated output.
- For few clear and show commands "?" will not show all options for VRF. Tab completion will give all possible values.
- For certain commands (including "no" form with some commands), "?" shows unsupported additional options.
- Some CLI commands will generate an "Error:Access denied" message upon failure. This means the operation failed on the switch and may not be related to permissions.
- Tab completion and <ctrl>-c (cancel) does not work for some commands.
- Incorrect range might be displayed in the help text for some of the show commands.
- Range support is available for all the interfaces in Network OS v7.0.1. Following limitations are applicable:
    - Interface range command is supported on breakout ports of same connector. Range is not supported involving breakout ports of multiple connectors.
    - Interface range command does not support mix of regular ports and breakout ports.
    - Range command is not supported across multiple slots of the chassis.
    - Range command for rbridge-id is not supported.
    - In some instances, there could be a delay in starting of operation specified in the range command after being issued.
    - When range issued for very large subset (e.g 4k VLAN, 2k port-channels, etc.), timeout can occur or user may temporarily see switch being unresponsive or with high CPU utilization. Extreme recommends using range in smaller chunks. Especially, while configuring VLANs/VEs and Port-channels, Extreme recommends range to be less than 500.
    - Range prompt doesn't get updated when few or all of interface in that range are deleted. Therefore, user should exit from Range submode if few or all interfaces are deleted that are part of that range. New configuration performed on same range submode may give unpredictable results.
    - On a large VCS cluster, configurations performed on Range of physical interfaces and port-channels may spike high memory usage.
- System does not warn user on deleting the IP config when VRF is configured.
- If "switchport trunk allowed vlan all" is already configured on any interface, then VLAN creation using range command will be slow as each VLAN will get provisioned individually.
- Some unsupported debug commands may be seen in Network OS v7.0.1. Extreme recommends not to run them on switches:
    - Show confd-state –, for debugging purpose only.
    - Show parser dump –, for debugging purpose only.
    - Show notification stream –, for debugging purpose only.

    - Auto upgrade command in config mode
- During “copy running-config startup-config” or “copy support” user mightsee occasional and temporary CPU spikes (up to ~30-40%).
- show mac-address-table command on console with include option can not be aborted witha break/ctrl-C. Use a telnet session for the same.
- Short form of MAC-Address is not supported as filter in “showrunning-config”.
- For IP access lists, display filtering based on sequence number alone does not work as expected.
- Certain oscmd commands may not work or give a different output under admin login
- If an alias exactly matches a partial keyword anywhere in the command line, pressing the TAB key for CLI command completion will claim that the input is invalid, and pressing the ENTER key will first replace the partial keyword with the alias expansion string. To avoid this, make sure that any partial keywords are not an exact match for an alias name.
- The authentication mode with primary & secondary sources of authentication cannot be updated to a configuration containing only the primary source. For example, the authentication mode cannot be changed from “radius local or radius local-auth-fallback” to ‘radius’. The workaround is to remove the existing configuration and then configure it to the required configuration.
- The “logging syslog server” command returns an error on the “secure” keyword. Use “secure port” to assign a nondefault port number.
- OSPFv3 on default VRF can be created without mentioning VRF name but while removing default VRF needs to be provided like "no ipv6 router ospf vrfdefault-vrf".

Platform

- After “chassis disable” it is recommended to wait for 60 seconds for VDX fixed-form switches and 300 seconds for VDX 87xx before performing the next “chassisenable”.
- Chassis-name is limited to 15 characters.
- 1G copper SFPs do not support exchanging flow-control settings during the auto-negotiation process. It is recommended to configure static mode of configuration of flow-control on both the ends of the desired link.
- 1G Optical ports should use the same speed config (speed auto or speed 1000) on both sides of the link for a proper link up.
- The VDX6940-36Q and VDX6940-144S requires 40 seconds between the removaland insertion of the 100G QSFP28 optics in order to establish a stablelink.
- System verification/ offline diagnostics tests need “chassis disable” before the testand “chassis enable” followed by immediate reboot.
- After “power-off line-card <x>” please wait for 120 seconds before doing the next “power-on line-card <x>” to avoid hitting a known defect where some interfaces might remain in administratively shut state.
- The speed on the management interface for VDX 8770 can be hardset to desired speed after configuring speed as auto. The speed on VDX 6740x and 6940x is supported only in auto mode.

- Multiple OIR (Online insertion and removal) of 40G LR optics when connected to ICX/FCX may cause link to remain down. Performing “shutdown” followed by “no shutdown” of the interface will recover the link.
- VDX 6740/6740T/6740T-1G/6940 platforms do not support IP fragmentation. MTU errors are reported in “show interface” as “Errors” under the “Transmit Statistics”.
- When a switch fan or PSU is removed or is faulty, switch status LED will blink green on VDX6940-144S and amber-green on VDX6940-36Q and VDX6740.
- For 6940 platform family, if all ports in a given trunk-group are used as ISLs, it is recommended to configure only 1 lossless priority on the switch.
- Logical SAN is not supported in fabric cluster mode.

Line cards

- The VDX 8770 supports following line-cards only on Network OS v4.1.2 and above:
  - LC48×10G
  - LC12×40G
  - LC48×10GT
  - LC27×40G
  - LC6x100G
- It is required to upgrade the chassis to the line-card’s supported Network OS version before plugging the line-card into the chassis.
- If there exists a configuration for a line-card on the slot of VDX8770, before inserting a new line-card of other type in the same slot, it is required to remove the configuration of the old line-card from that slot. The “no line-card” command should be used to remove the old line-card configuration from the slot where the new line-card is to be inserted. The new line card may be faulted with appropriate code if the new line-card is plugged into the slot which has configuration of a line card of other type.

USB

- Starting with Network OS v6.0.0, Extreme 4GB USB drive support is added. But, Extreme 2GB USB drives should still work as before.

Licensing

- On VDX platforms that have Flexport FC capable interfaces, enabling FibreChannel ports requires only the FCoE license to be installed and does not require any Port Upgrade license. The Port Upgrade license only controls Ethernet ports (number of ports or speed supported).
- An Integrated Routing license is NOT required on FOS-based SAN platforms running FOS 7.0.1 or above for FCR interoperability connectivity with VCS fabrics and the VDX6740x. Please refer to the FOS v7.0.1 Admin Guide documentation on configuring FOS platforms for connectivity to VDX 674x switches and VCS fabrics.
- The Layer 3 license is required on VDX8770 switches to enable Layer 3 feature set including OSPF, VRRP, BGP, VRF etc. A separate Layer 3 license is not required on VDX fixed-form factor switches as Layer 3 features are included in the default license.
- The Advanced Services License provides a single upgrade option to enable Fibre Channel over Ethernet (FCoE) and Layer 3 features on VDX8770 switches.

VCS

- Loopback connection is not supported in VCS mode. If a loopback connection is done (either using loopback plugs or port to port connections on the same switch), those interfaces become ISL interfaces.
- A node with default configuration will not join a cluster if the intermediate nodes between the node being defaulted and rest of the cluster are also undergoing reload. If the node boots up earlier than the intermediate nodes, it will form its own VCS and not join the parent cluster. In such situations, reload the node that is required to join the cluster.
- Logical Chassis Cluster Mode:
    - When a new switch is added to an existing VCS Fabric and if the new switch takes the role of principal node, the other switches in the fabric will receive the configuration of the distributed features such as Virtual IP and VM-Aware Network Automation from the newly added switch. This will cause the existing distributed configuration to be overwritten by the newly added switch in the principal role. This can be avoided by following the new switch addition procedures in *the Network OS Administrator's Guide*.
    - After a cluster reboot, Extreme recommends to do both "show fabric all" and "show vcs" to ensure that cluster is entirely formed without any issue. User might see that 'show vcs' takes an additional 2-3 minutes to show all participating switches. This is an existing behavior and doesn't affect data path functionality in mostcases.
- "show fabric isl" & "show fabric trunk" may show the interfaces in random order without sorting.
- The default-configuration behavior may be different depending on the default-configuration triggers.
- The snapshot restore feature in VCS should be used to restore the local configuration and not the global configurations.
- Usage of Rbridge-range option to configure Rbridge context specific configurations is not recommended.
- Fastboot option is not recommended as a preferred method of reloading theswitch.
- VCS for NOSv7.0.1:
  Note the following results for the given actions.

Default-config trigger	Global Config (i.e. virtual-fabric)	Local Config (i.e. SFP breakout)
copy default-config startup-config	Preserved	Preserved
VCS-ID and/or Rbridge-ID change	Preserved	Removed
firmware download default-config	Removed	Removed
write-erase	Removed	Removed

VCS Node Replacement

When a switch is added into the cluster using the **VCS replace** command, the default configuration is updated for DPOD and ISL interfaces. For example, when the **no fabric isl enable** command and the **no fabric trunk enable** command is configured, the command configuration reverts back to **fabric isl enable** command and **fabric trunk enable** command**.**

When **reserve** is configured under the dpod configuration, the following initial configuration is displayed.

```
device(config)# dpod 1/0/2 reserve
dpod 1/0/2
    reserve
  !

 interface TenGigabitEthernet 1/0/2
    description khnfa0203_TE3/0/12
    channel-group 5 mode active type standard
    no fabric isl enable
    no fabric trunk enable
    fabric neighbor-discovery disable
    lacp timeout short
    no shutdown
  !
```

When the **vcs replace** command is executed, the **reserve** configuration, and the **no fabric isl enable** and **no fabric trunk enable** commands are removed as displayed in the example below.

```
dpod 1/0/2
  !
 interface TenGigabitEthernet 1/0/2
    description khnfa0203_TE3/0/12
    channel-group 5 mode active type standard
    fabric isl enable      <<< missing "no"
    fabric trunk enable    <<< missing "no"
    fabric neighbor-discovery disable
    lacp timeout short
    no shutdown !
```

Logical Chassis

- Configurations are not auto preserved on mode transitions (between Fabric Cluster and Logical Chassis mode). Please follow the mode transition procedure as outlined in the Network OS Admin Guide. Non-default User Id/password will be lost when migratingfrom FC to LC.
- User should not make configuration change during Logical Chassis firmware upgrade or while ISL toggling to prevent the switch segmenting from the cluster due to configuration mis-match.
- Upon Node segmentation from the cluster, user should run “copy default start” orexercise the default-config boot feature on the segmented switch to bring it back to thecluster.
- For Netconf and SNMP, user has to poll using individual node ManagementIP.
- Creating a snapshot with “\” in snapshot-id creates the snapshot file with incorrect name.
- Config snapshot cannot be restored on pizza box platform when SW1 isactive.
- There will not be any raslog to the user when replacement of a nodefails.
- With large configs, while a switch is rejoining a fabric with default config, “%Error:Couldnot find Interface" may be printed temporarily. The switch will recover and join thefabric.
- Config changes during principal switch-overs are not supported and may segmentthe cluster.
- Disabling virtual-fabric may take up to 10 minutes depending on the number of ISLs and VLAN interfaces configured in the VCS.

Extreme Trunks

- The VDX 6740, VDX 6740T, VDX 2741 Extreme trunk (BTRUNK) can support up to 16 member links with a maximum throughput of 160G using 16x10G ports in the same trunk group. On these platforms traffic may not be distributed evenly across all member of a trunk at lower traffic rates.
- The VDX 6740, VDX 6740T, VDX 2741, VDX 2746 and VDX 6740T-1G Extreme trunk (BTRUNK) can support up to 2x40G member links in the same trunk group for a maximum throughput of 80G. The VDX 2746 can support 4x40G.
- The VDX 8770 Extreme trunk (BTRUNK) can support up to 8 member links with a maximum throughput of 80G using 8x10G ports in the same trunk group. Full link utilization of 8 ports in a trunk group is achievable with larger packet size (>128 Bytes).
- In the VDX 6940-36Q and VDX 6940-144s, only 63 port-channels are supported including LACP and Extreme PO.
- The VDX 6940-36Q Extreme trunk (BTRUNK) can support up to a maximum throughputof 120G using 3x40G or 120G using 12x10G breakout ports in the same trunkgroup.
- The VDX 6940-144S Extreme trunk (BTRUNK) can support a maximum throughput of120G using 3x40G or 12x10G links in the same trunk group.
- In order for two 40G ports on VDX 8770 to form Extreme trunk, it is required that the ports be in breakout mode and in same trunk group. Breakout optics with a single QSFP optical cable must be used.

- Prior to Network OS v4.1.0, “fabric trunk enable” configuration on the 40G interfaces on VDX 8770 is allowed, however it does not provide non-breakout mode trunk capability to the ISLs.
- Upgrading from any version before Network OS v4.1.x will change the default configuration on 40G interfaces on VDX 8770 from “fabric trunk enable” to “no fabric trunk enable” to accurately indicate the capability. Configuring “fabric trunk enable” directly on the 40G interfaces is accordingly blocked.

Breakout Interfaces

- VDX 8770 supports only static breakout of 40G ports. It is required to power OFF and ON linecard for the 40G ports on it to be converted into 10G breakout ports andviceversa.
- VDX 6940-36 and 6940-144S supports only static breakout of 40G ports. It is required to reboot the switch for the 40G ports on it to be converted into 10G breakoutports
- For VDX 6740, 6740T, 2741 and 6740T-1G platforms, the LED state for a breakout interface is deterministic. For all other supported platforms, the LED state for a breakout interface is non-deterministic.
- In breakout mode, the ‘show media’ CLI will display the same media information for all breakout interfaces, except for temperature, Tx voltage, Tx bias current and Rx power. These parameters would be displayed on per line basis. The TX Power Field in the show media command is not supported by the 40G optics.
- On 40G native mode - Breakout configuration is not blocked. If configured on one side, other side of link won’t able be identify peer port config is breakout and link won’t be stable.
- On VDX 6740T/6740T-1G/2746, the breakout ports are FlexPort capable, and may be configured to connect to FC switches with 4x16G breakout supported cables and optics.
- On VDX6940-144S, breakout connection using non-breakout cable is notsupported.

Dual-personality Ports

- Interface can be brought up in 100GbE or 40GbE mode. This feature is supported on VDX 6940-144S.
- Only static configuration is supported, the switch need to be rebooted for the dual personality mode change to take effect.
- Configuring 40GbE dual personality interface in 100GbE mode would result in the other two 40GbE interface in the port-group being disabled.

1G Mode

- RMON stats are calculated incorrectly for packet sizes 64-127 bytes.
- 1G ports cannot form ISL links. Only 10G ports can be used to form ISLlinks.
- Extreme Trunks cannot be formed with 1G. Extreme Trunks are only supported on10G.
- A LAG cannot be created between 1G and 10G ports.
- FCoE configuration is NOT supported on 1G ports.
- DCBX configuration for FCoE is not supported on 1G ports.
- For 1G optics used in VDX6740 and VDX6940-144S, port speed should be set to Auto on both sides. If one side is speed 1000 and other side is Auto, link may not comeonline.

vLAG

- LAGs are created with default speed of 10G. Therefore Extreme recommends end user to set required speed manually based on member speed using “speed” command.
- When configuring LACP LAG between VDX and non-Extreme switches it is highly recommended to enable the vLAG ignore-split on the VDX . Ignore split option isenabled by default.
- The port-channel interface "load-balance" is not the same as "fabric port-channel <#> load-balance"
  - The port-channel interface “***load-balance***” command configures load-balancing on the actual vLAG member links (effective on Rbridges directly participating in the vLAG).
  - The "***fabric port-channel <#> load-balance***" configures load-balancing on Rbridges NOT participating in the vLAG, but connecting to neighboring vLAG participating Rbridges.

Virtual IP Address Support

- A separate gateway cannot be configured for Virtual IP address. Default gateway will bethe same as the gateway address for the management port of thisswitch.
- For VCS Virtual IP address to work correctly, the management port’s IPv4 or IPv6 address should be assigned, functional and both address should be in samesubnet.
- There is no Virtual MAC address associated with the Virtual IP address, physical MAC will be used.
- Chassis Virtual-IP is only supported on the VDX8770.

Security, Management ACLs, Authentication, Authorization

- Login authentication service (aaa authentication login cli):
  - With “local” option specified as secondary authentication service, local authentication will be tried only when the primary authentication service (TACACS+/RADIUS/LDAP) is either unreachable or notavailable.
  - Behavior of “local” option in pre-4.1.0 releases is changed to the “local-auth-fallback” option.
  - When login authentication configuration is modified, the user sessions are not logged out. All connected user sessions can be explicitly logged out using “clear sessions” CLI.
- ACLs are not supported for egress traffic flows on managementinterfaces.
- Configuring TACACS+ or RADIUS without a key is not supported. If no key is configured, the switch uses a default key of “sharedsecret”. The use-vrf option should be used to enter any additional parameters such as retries, timeout or key.
- Same NTP server configuration with different vrf not supported.

- There is a possibility that locked user accounts will get unlocked after a reboot if the running-config (before reboot) is different from startup-config of user accounts.
- Encrypted text (taken from running-config of any user account password with encryption turned on) should not be used as input for clear-text password for the same user. This may result in login failure of the user subsequently.
- When the ACL is applied to a management interface, only the top 256 rules will be applied if the ACL contains more than 256 rules.
- It is advised to not to apply ACL with 12k rules to management interface.
- When more than 250 rules ACL's are configured (over supported scale), they may be partially installed & effective.
- Access to ONLY the following Active Directory (AD) servers is supported by Extreme LDAP client:
    - Windows 2000
    - Windows 2003
    - Windows 2008 AD
- IPv6 RA Guard feature is not supported on VDX 8770 although the CLIs are visible.

SPAN & RSPAN

- CPU-originated packets cannot be output spanned.
- If SPAN has to be supported to multiple locations, please use RSPAN on VLAN.
- On VDX 8770 and SPAN in VCS feature, ISL can be source port, but the destination has to be on the same RBridge.
- Spanning of LAG port is not supported. To span a LAG, user should individually enable spanning on all the member ports of the LAG. However flow based SPAN is supported on LAG port.
- A profiled port cannot be a SPAN destination.
- After ISSU upgrade on VDX 8770, Port Based SPAN may not work.
- SPAN destination port statistics will keep incrementing even when port is operational or admin down.

MAC Learning Considerations in VCS

- Under rare circumstances, end user might see mac address sync up issues on few nodes of a cluster (where 1 or more MAC addresses might be missing in some nodes). Extreme recommends to do "clear mac-address-table dynamic" in such cases.
- Static mac addresses will be displayed even when interfaces are down. This may cause blackholing of the traffic.
- Under certain conditions, MAC addresses may not be learnt even though ARP's may be learnt for those same MAC addresses.
- Under certain conditions, multicast traffic destined for static multicast address will flood on to other VLANS.

PVLAN

- Following PVLAN features are not supported:
  - IGMP on PVLANs but there is no error message displayed if operator configures IGMP snooping on PVLAN
  - ARP & Routing in PVLAN domain
  - Enabling Routing in Primary and Secondary Vlans.
  - CLI to enable Local Proxy ARP on primary VLAN.
  - IP Configuration on PVLANs
  - Ve Configuration on both Primary and Secondary Vlans
  - AMPP on PVLANs
  - In case of MSTP if a primary VLAN is added to the instance automatically secondary VLAN also added to the instance.
  - When the operator wants to delete the host association on a host port recommended to use "no switchport" rather than "no switchport private-VLAN host-association". This is applicable only when the host port is untagged. When the host port is tagged both the commands can be used.
  - Primary VLAN ID needs to be lower than the secondary VLAN IDs. If primary VLANID is greater than secondary there is an issue with config replay.
  - In Logical Chassis mode source macs may not learn on PVLAN configured ports, after deleting some of the secondary vlans for which the traffic is not flowing.

UDLD

- The UDLD protocol is not supported on the members of a Extreme trunk.
- The UDLD protocol is not compatible with Cisco's proprietary UDLD protocol.
- UDLD needs to use the higher timeout in Scale and Stress environment. UDLD may flap during HA failover and ISSU.

STP/DiST

- VDX does not support tunneling non-standard BPDUs and thus IEEE BPDUs (0180:C200:0000) generated as tagged packets in STP/RSTP/MSTP modes may not be tunneled successfully across VCS Fabric. However, VDX supports tunneling standards' based BPDUs such as untagged IEEE BPDUs and tagged or untagged PVST BPDUs (0100:0CCC:CCCD). Post 3.0.1, the tagged IEEE BPDU can be tunneled across VCS fabric using command: "tunnel tagged-ieee-bpdu" under interface configuration.
- In Fabric Cluster mode, global spanning-tree configurations (STP enable, STP Vlan configurations, STP over vLAG configurations) have to be performed in all the switches in VCS at the same time. For example, to run spanning-tree, it has to be enabled on all the switches including switches that don't have any edge ports.
- By default global spanning-tree and interface level spanning-tree will be disabled, user has to explicitly enable on the desired ports. VLAN spanning-tree state is default enabled.
- BPDU tunnel configurations are permitted only when spanning-tree is disabled in VCS.

- For cisco proprietary Per Vlan Spanning Tree protocols (PVST and RPVST) user needs to configure the Extreme switch to send BPDU on Cisco multicast destination mac address “0100.0ccc.cccd” for non-native VLANs. By default, NOS 6.0.1 software uses Extreme "0304.0800.0700" multicast mac to send BPDU's on non-native VLANs.

  Since Cisco boxes use Cisco multicast mac address to send spanning tree BPDU on non-native VLANs, this configuration is needed in VDX switches to interoperate. This is an interface specific configuration.

  Below is the example to configure Cisco BPDU mac for PVST and RPVST under interface mode:

```
VDX 6740-VCS1# conf t
VDX 6740-VCS1(config)# protocol spanning-tree rpvst
VDX 6740-VCS1(config-rpvst)# exit
VDX 6740-VCS1(config)# interface Port-channel 100
VDX 6740-VCS1(config-Port-channel-100)# spanning-tree bpdu-mac ?
Possible completions:
  0100.0ccc.cccd   Cisco Control Mac
  0304.0800.0700   Extreme Control Mac
VDX 6740-VCS1(config-Port-channel-100)# spanning-tree bpdu-mac 0100.0ccc.cccd
VDX 6740-VCS1(config-Port-channel-100)# exit
VDX 6740-VCS1(config)#
```

Edge Loop Detection (ELD)

- ELD is supported on the edge ports that are connected either by end-hosts OR another switch OR another VCS.
- ELD is also supported for edge interfaces connected to hosts.
- ELD may not be enabled after line-card powercycle.
- The edge-loop-detection port-priority with the higher number takes priority for shutting down the loop interface. If the port-priority is the same, the highest interface ID followed by the highest Rbridge-ID are used as the deciding metric.

Long Distance ISL Ports

- Long distance ISL configuration (“long-distance isl” command) is not allowed if CEE Map/fcoeport is configured on any edge ports in the same port group.
- CEE Map modification is not allowed when long distance ISL is configured.
- A maximum of three PFCs can be supported on a long distance ISL configuredplatform.
- When long distance ISL is configured on the switches, all ports in the port group will be bounced.
- Both side of long distance link should have long-distance-isl configuration. Otherwise end to end PFC might not work appropriately.
- For 10Km/Extended Range long distance configuration all other ISLs in the port group will be disabled.
- For 2Km/5 Km long distance configuration, one other ISL will be allowed to come online in the port group.
- For 2 km, 5 km and 10 km long-distance, use Extreme-supported Long Range (LR) optics for direct connectivity.
- For 30 km long-distance, use Extreme-supported Extended Range (ER) optics for direct connectivity.

- The “long-distance isl” command based extended fabrics are supported only on 10G interfaces.
- The 40G and 100G interfaces do not support “long-distance isl” command, however can extend distances for non-lossless traffic up to 40Km using standard ISLs.
- On standard ISLs, the 10G, 40G and 100G interfaces support lossless traffic up to 1Km.
- The “long-distance-isl” command will not be supported on the SO-10GE-ZR-CX, 10G-SFPP-ZR, and 10G-SFPP-ZRD-T 80km optics.
- The SO-10GE-ZR-CX, 10G-SFPP-ZR, and 10G-SFPP-ZRD-T 80km optics requires a minimum distance of 20km in order to successfully form a standard ISL connection

AMPP and Port-Profiles

- Port-profile status does not reflect the remote interface info in VCS fabric mode.
- Native VLAN support inside AMPP does not honor the global enable/disable flag.
- SPAN destination port cannot be a profiled port.

  Extreme recommends deleting all manually created port-profiles when migrating from a legacy AMPP environment to VM Aware Network Automation.
- Vmkernel related port-profiles removed/reapplied during HA operations may result in vmotion failures.
- MAC-based classification allowed only on access port-profile and C-tag classification allowed only on trunk port-profile.
- When a port becomes a profiled-port, all SERVICE VFs in that domain are provisioned on this port.
- “Switch trunk allow VLAN all” can only be present in one domain, it cannot co-exist with other c-tag based classifications in that domain.
- User is not allowed to edit/delete the default-profile-domain when Service VF is disabled.
- New port-profile is not auto added to the default domain when Service VF is enabled. It can only be explicitly added to or removed from the default profile-domain.
- On disabling Service VF UpgradedVlanProfile should be re-configured with “switchport trunk allowed VLAN all” in Default-profile-domain if it is removed/modified.
- Newly created port-profiles which is not part of any domain should be added to the default-profile-domain explicitly while disabling the Service VF.
- SERVICE VF classification cannot conflict across port-profiles in the same port-profile domain, but it can conflict across PP in different domains. i.e. a port-profile-domain cannot contain conflicting SERVICE VF classifications.

vCenter

- Receiving more than five vCenter events within a span of 30 seconds, results in asset discovery getting initiated. Post discovery cluster configuration will be in sync with vCenter.
- vCenter auto-profile is automatically added/deleted to the default port-profile-domain in Service VF enabled/disabled mode.
- Modifying/editing the auto port-profiles in the default-domain is not recommended, which may cause auto-pp application failure during vCenter operation and end up in traffic failure.
- vCenter auto-profile does not support SERVICE VF classification.

- Output of show vnetwork vss displays the vmnic against the vSwitch even after the removal of the vmnics from the vSwitch through vCenter. Recovery happens in the next auto-discovery cycle.

QoS

- LC27x40G and LC12x40G linecards do not respond to incoming Ethernet pause (802.3x) and PFC frames in native 40G interface mode. This does not impact throttling of incoming traffic in response to PFC or Pause frames transmitted by the linecards. In order for LC27x40G and LC12x40G linecards to respond to pause frames, it is required to configure the 40G interfaces in breakout mode and use 40G breakout optics with regular native 40G cables.
- It is recommended to use the same CoS tail-drop threshold on all members of a port-channel to avoid unpredictable behavior.
- In a hybrid logical-chassis, if a user configures a platform specific feature, it will be configured only on the rbridges which support that feature
- Asymmetric pause is supported on 1G port interfaces.
- It is recommended to enable maximum 2 PFC s on edge interfaces on VDX 6740/6740T and 6940-36Q platforms. Flow control is disabled by default on all interfaces.
- Priority 7 is reserved for control traffic on VDX switches. User data traffic should use priorities 0 through 6. Priority 3 is used for the FCoE lossless traffic by default.
- ExtremeSwitching VDX architecture prioritizes Unicast traffic over Broadcast or Multicast traffic under port congestion.
- The interface queues operate in Strict Priority mode when there are no ISLs online on the switch. This could result in potential starvation of the CPU queue if line-rate traffic flows through an interface.
- Byte count is not supported for match ACL statistics on the VDX 6740/6740-T and 6940-36Q.
- Byte count is not supported for RED statistics on either the VDX 8770 or the VDX 6740/6940-T and 6940-36Q.
- For 6940-36Q its not recommended to configure “log” option in ACL for Flow based QoS and System based QoS as it may lead to throughput issues with larger packet size.
- The “count log” option in ACL is not supported for Flow based QoS and SysFBQ.
- The CLI “qos trust cos” is not applicable in VCS mode. However, “show qos int” will show as cos is trusted on ports on which “cos-mutation” or “cee default” config is applied.
- Configuring an interface with a nondefault DSCP-to-traffic class-map is allowed. However, configuring an interface with a nondefault CoS-to-traffic class-map is not supported.

FCoE

- On switches shipped with NOS7.0.1, the default mode of operation is Access Gateway for the VDX 6740, 6740T, 6740T-1G, VDX 2741, VDX 2746. The VDX 2741 was supported in 5.0.x (not supported in 6.0.x) and is upgradable to 7.0.1.
    - Logical SANs have been supported since Network OS v6.0.0. If user needs to enable Fibre Channel Forwarder (FCF) mode, switch needs to be configured in FCF mode. If the switch is upgraded from a lower NOS version (pre 4.1.2 builds) to v7.0.1, it will be in FCFmode.
    - FCoE logical SAN is not supported in an FC cluster.
    - Extreme recommends not having FCoE ports and Long Distance ISL ports in the same port-

groups. This configuration will NOT be prevented by the CLI; however it can result in unpredictable behavior for FCoE traffic.

- If the FCoE FCMAP is changed from default to non-default, user might see that FCoE login may not happen. Please toggle the interface using “shutdown” followed by “no shutdown” to work this around.
- When an RBridge is removed from VCS cluster, it does not need to be manually removed from fcoe fabric-map.
- MAC binding for remote SANs is not honored during config replay.
- In case the FIF is multiple hops away from the FCF, it is recommended that the interfaces of the intermediate switch shouldn’t be configured with the same remote san as that of the FIF/FCF
- VLAN's which are reserved for FCoE may not be used for any other purpose. This is true for both Fabric Cluster and Logical Chassis modes.
- Extreme recommends that for all LAGs with FSB, the fcoeport config must be applied on the LAG itself. And for all LAGs with directly attached CNAs, the fcoeport config must be applied on the member ports.
- Binding an enode mac to FCoE interface is not allowed in range context, as only one enode mac can be bound to one FCoE interfac.e
- While providing range for FCoE interfaces, it’s recommended to provide the range only in ascending order. For ex: interface fcoe 1/48/11-38 is recommended, interface fcoe 1/48/38-11 is not recommended.
- FCoE traffic may not be mirrored using RSPAN. Workaround is to use SPAN.
- In use cases with FSB, it is noticed that after converting dynamic port-channel to static, hosts and targets don’t see each other.
- In NETWORK OS V6.0.1 and later, up to four FCoE Vlans are supported in VDX . But on a single VDX , All member ports in a LAG have to be configured with the same FCoE Vlan. Different LAG can be configured with different FCoE Vlan.
- In NETWORK OS V6.0.1 and later, it is recommended user define different fabric-map for Remote Logical SAN and Local Logical SAN configuration. If user configures a fabric-map to work on Remote Logical SAN first and then later change the same fabric-map to become Local Logical SAN, it may cause FCoE port continuously flapping.
- In NETWORK OS V6.0.1 and later, when FCoE CNA connect through VDX 6940-36Q/VDX 6940-144S to a Remote Logical SAN, if user performs certain operation in AG switch (e.g. N-port failover, VF-port remapping, fcmap change etc), FCoE CNA may fail to login. The workaround is to do shut and no shut on the FCoE port on which FCoE CNA is connected.
- LC27x40G and LC12x40G linecards do not respond to incoming Ethernet pause (802.3x) and PFC frames in native 40G interface mode. This does not impact throttling of incoming traffic in response to PFC or Pause frames transmitted by the linecards. In order for LC27x40G and LC12x40G linecards to respond to pause frames, it is required to configure the 40G interfaces in breakout mode and use 40G breakout optics with regular native 40G cables. This is required for FCoE traffic.

FlexPorts

- The port transceiver optic type must match the configured FlexPort type. If a port is configured as Fibre Channel then an appropriate FC SFP+ transceiver must be used; likewise when the port is configured as an Ethernet port then an appropriate Ethernet SFP+ transceiver must be used. The same applies to QSFP+ transceivers – the transceiver type must match the configured Flexport type for the QSFP links.
- Only Extreme-branded FC SFPs are supported.
- Setting the connector-group speed to HighMixed allows only one FC port speed (16G) but the port speed configuration can still be set to auto.
- Changing the connector-group speed always disrupts any other active ports in the connector-group, independent of configured Flexport type.
- The FCoE Base license is required to enable any FibreChannel interface.

Fibre Channel

- F_Port can support only 63 NPIV devices.
- Loop devices are not supported.
- Long distance is not supported on Fibre Channel ports.
- Proprietary features such as QoS, D-Port, FAPWWN are not supported on Fibre Channel ports.
- Credit Recovery is supported on Fibre Channel ports.
- FEC is supported on Fibre Channel E/Ex ports only (no support on F/Nports).
- Trunking is not supported on Fibre Channel ports running at 2G or 4Gspeeds.
- On the VDX 6740, VDX 6740T, VDX 6740T-1G, VDX 2741 and VDX 2746 platforms Fibre Channel trunks are limited to 2 FC trunks per portgroup.
- To configure a connector-group as Fibre channel need to have all the interfaces in the connector group as type Fibre channel.
- Fibre Channel trunks only form with ports running at the samespeed.

Access Gateway

- The switch can be operated as Fibre Channel Forwarder (FCF) by disabling Access Gateway mode.
- AG does not bridge the VCS and SAN fabrics because hosts connected to the AG switch are registered in the SAN name-server only. Therefore, all zoning operations for AG are done on the SAN fabric.
- At least one N-port must be online in order for FCoE devices to login.
- After enabling Remote Logical SAN on AG switch, FCoE devices connected to AG switch will not login with “fcoeport default” provisioning and needs to be configured as“fcoeport <logical-san>”.
- Cannot configure the default Logical SAN to the interfaces in the FCF-groupswitch.

IP Fabric

BGP eVPN :

- RD should be unique across the VLANs/VRFs and across the leaf nodes.
- If the leaf nodes are in different BGP AS, then ignore-as option should be specified to the route-target configuration under eVPN instance.
- BGP MAC route dampening is applicable only for frequent MAC moves across leaf nodes not part of vLAG pair.
- BGP neighbor next-hop-unchanged should be configured under eVPN address-family on all of the nodes in the IP fabric.
- On a vLAG pair, eVPN instance configuration should be symmetric.
- If the leaf nodes are in the same BGP AS, "allowas-in 1" should beconfigured.
- On VDX6740, part of a 2 node VCS, remote VTEP destination should not be reachable via another node in the VCS.
- For VRF extended using L3VNI over eVPN, at least one prefix should be advertised by both of the leaf nodes extending the VRF.
- It is recommended to configure different BGP AS numbers on each set of spine nodes when connecting 2 PoDs.
- Traffic tromboning is not supported for IPV6 in IP Fabric with /128 routes.
- In the scale environment with a large number of /32 routes, traffic disruption may be seen upon reload or HA failover.
- Tunnel creation is triggered by BGP NH installation resulting in creating more tunnels than configured which might be seen at the Border Leaf.

ARP/ND Suppression:

- ARP/ND suppression should be configured on the VLAN if IPv4/IPv6 Static Anycast Gateway is being configured.
- Upto 512 VLANs are supported with DAI hardware profile. Default hardware profile supports upto 32 VLANs.
- ARP/ND suppression feature is supported only on VDX 6740, 6940, 6940-144s platforms.

Conversational ARP:

- It is recommended to enable both Conversational-ARP and Conversational-MACtogether.

Static Anycast Gateway:

- ARP/ND suppression should be configured on the VLAN if IPv4/IPv6 Static Anycast Gateway is being configured.
- Static Anycast Gateway address/static Anycast MAC configuration should be identical for a given VLAN across leaf nodes in IP Fabric.
- IP services/protocols cannot be enabled on an interface where only Static Anycast Gateway address is configured.
- VRRP/VRRP-E configuration should be identical for a given VLAN across leaf nodes in IP Fabric. But it is recommended to use Static Anycast Gateway.

- All VLANs having Static Anycast Gateway configuration should be extended into eVPN on a vLAG pair.

ND/RA

- Proxy ND is not supported.Static

IPv4

- IP Directed Broadcast is not supported under non-default VRF context. It is supported only in Default-VRf context.

BFD

- Static Route BFD, BGP BFD and OSPFv2/v3 BFD
    - For Single HOP BFD sessions configured with source IP as secondary IP is not supported, since significance of Source IP in BFD configuration is only to determine on which interface BFD session should be started and hence interfaces' Secondary IP is not used as source in BFD PDU.
    - BFD is not supported on leaked routes.
    - BFD for multi-HOP BFD neighbor reachable via multiple paths with equal cost(ECMP) will not be supported since BFD requires BFD session to be created for the neighbor for each ECMP path.
    - BFD is not supported for OSPFv2 & OSPFv3 virtual links.
    - For single hop BFD sessions, BFD will consider the interval values that are configured on interface, and not the non-default values that are configured within the global command.
- BFD for VxLAN tunnels
    - BFD session may not come online or may flap if VCS cluster is in transient state during reload, vLAG failover, fabric split, chassis disable/enable and such scenarios. It is required to have a stable VCS cluster in order for BFD sessions on VxLAN tunnels to work as expected.
    - BFD parameters are not configurable on VCS VxLAN Gateway for Type NSX. The parameters are derived from NSX controller.

VRRP

- VRRP and VRRP-E cannot be enabled together on VDX 6740 and VDX 6740Tplatforms.
- IPv6 and IPv4 VRRP sessions cannot be configured with the same VRRP group-ID on the same Layer3 interface.
- If an IPv6 VRRP session is configured with only global VIP address without Link-Local VIP, VIP configuration will fail for that session during download of configuration fromfile.
- VRRP v4 or v6 can be enabled with VRRP-E v4 and v6 on the VDX 6940family.
- VRRP v4 and v6 cannot be enabled together on an interface on the VDX 6940 family.
- "show vrrp summary" and "show ipv6 vrrp summary" will display all sessions in default vrf. In earlier NOS versions, these commands displayed sessions across all vrf.

Fabric Virtual Gateway (FVG)

- FVG co-existence with VRRP/VRRP-E in VDX 6740 and VDX6740T
    - FVG ipv4 or FVG ipv6 with default global mac cannot be enabled with VRRP but can be enabled with VRRPE-E.
    - FVG ipv4 or FVG ipv6 with non-default global mac cannot be enabled either with VRRP or VRRPE-E.
- FVG co-existence with VRRP/VRRP-E in VDX 6940
    - FVG ipvx with non-default global mac: when the global gateway-mac-address is changed using the "gateway-mac-address" command to something other than the default mac. for eg. 0000.1111.2222.
    - There are two groups of protocols
        - Group 1:
            - VRRP ipv4
            - VRRP ipv6
            - FVG ipv4 with non-default global mac
            - FVG ipv6 with non-default global mac
        - Group 2:
            - VRRPE ipv4
            - VRRPE ipv6
            - FVG ipv4 with default global mac
            - FVG ipv6 with default global mac
        - A maximum of only two protocols from group 1 can be enabled at a time.
        - All protocols of group 2 can be enabled at a time.
        - If 2 protocols from group 1 are enabled, no protocol from group 2 can be enabled. While if only 1 of the group 1 protocols is enabled, all the group 2 protocols can be enable at the same time.
- Fabric Virtual Gateway (FVG) is not applicable in IP Fabric environment, Static Anycast Gateway to be used to achieve similar functionality.

OSPFv3

- OSPFv3 HA with Graceful restart is not supported but GR-helper mode functionality is supported. VRF-Lite-Capability CLI and support for Down bit handling is not available in OSPFv3 as in OSPFv2. When the BGP4+ route from the MPLS cloud is redistributed into OSPFv3 domain the redistributed route is always installed in the OSPFv3 routing table.

BGP

- Conditional advertisement of default-route using route-map match prefix not supported.
- Over a link-local eBGP session, updates are not carrying the new nexthop that is set usinga route-map.

Layer 2/Layer 3 Multicast

- The following PIM features are not supported in this release:
  - IP version 6
  - VRF
  - Configuring the switch as the BSR (Bootstrap Router) candidate.
  - Configuring the switch as the Rendezvous Point or Rendezvous Point candidate. The RP must be configured outside the VCS cluster.
- In Fabric Cluster mode, IGMP Snooping must be enabled in all the switches in VCS Fabric Cluster mode
- Statistics for MLDv1 is done on a VLAN basis across VCS.
- Multiple IP subnetting support: PIM FHR and LHR operation are not supported on secondary subnets.

VRF

- Under VRF submode there is a syntax change for the address-family ipv4command.
  Old format: address-family ipv4 [max-route <value>]
  New format:
  address-family ipv4 unicast
  max-route <value>
  Note: “max-route” command is now moved to address-family submode.
- There is no provision to configure “max-routes” for default-vrf.
- There is no use case for “rd” configuration in VRF and this command will be deprecated in next release.
- On configuring VRF on an interface, all previous IP config on that interface will be deleted.
- Removing VRF address family on a non-default VRF will delete all relevant address-family configurations including the interface and protocol configuration for that VRF.

BGP-VRF

- Local-as <num> can be configured for particular VRF under "address-family ipv4 unicastvrf <vrfname>" and is not supported under “address-family ipv6 unicast vrf <vrf-name>”.
- “maxas-limit in” can be configured for particular VRF under "address-family ipv4 unicastvrf <vrfname>" and is not supported under “address-family ipv6 unicast vrf <vrf-name>”.
- When route-map is applied to BGP, and route-map has multiple ‘set ip next-hop’ statements in a single instance, BGP will consider the last ‘set ip next-hop’ in theroute-map.

ACL

- L2 User ACL deny rule can prevent trapping of L3 controlframes.
- IPv6 ACLs at ingress are not applicable for packets with Link local sourceaddress.
- ACL Logging at egress can impact forwarding traffic at highrates.
- Counters for hard-drop ACLs may not countaccurately.
- Statistics are not supported for hard-drops at Egress.
- For VDX 8770, IPV6 Egress ACLs, Match on DSCP value compares only 4 LSBs instead of all 6 DSCP Bits.
- ACL with “Routed” keyword functions only for VE/Router Port MACs. It does not work for VRRP Routed.
    - Work-around: Apply default mode ACLs (No “routed” keyword).
- For Private VLANs, Egress ACLs on Primary VLAN is applied only for all traffic which ingresses primary VLAN i.e.
    - If the traffic ingresses from Primary VLAN but gets translated to Secondary VLAN at egress, ACL on primary VLAN at egress is still applicable toit.
    - If the traffic ingresses from Secondary VLAN but gets translated to Primary VLAN at egress, ACL on primary VLAN at egress is still not applicable toit.

Policy-based Routing (PBR)

- If a PBR route-map is applied to an interface that is actively participating in a control protocol and the ACL specified in the route-map also matches the control protocol traffic the control protocol traffic will be trapped to the local processor and not be forwarded according to the route-map.

Inter-VRF Leaking (Static)

- S+ symbol routes indicates leaked routes.
- VRF route leak cascading is not supported– only one level ofindirection.
- User should avoid making Static, dynamic and connected route conflict with routes in target VRF when configuring route leak.
- For bidirectional traffic with router leak, user needs to configure route leak in both direction separately.
- Route leak configuration to next hop IP on the same box on different VRF is not a valid configuration, but CLI will be accepted.
- Precaution needs to be taken when leaking default routes - this can result in routingloops.

- Switch management from non-management VRF by leaking route from non-managementto management VRF is not supported.

DHCP IP Helper

- There is no HA support for DHCP relay statistics. When a switchover happens, the statistics will not be replicated to the new active MM.
- Clients may not converge in some IP Fabric environment. Care should be taken to not configure DHCP IP helper and Static Anycast Gateway on the sameinterface.
- Two DHCP OFFER per one DHCP DISCOVER and two DHCP ACK forsingle DHCP request seen IP fabric setup

Dynamic ARP Inspection (DAI)

- The ARPs learnt on trusted ports would be deleted when DAI is enabled or DAI filter changed.
- Static ARPs not permitted by DAI filter would be promoted to active state. Administrator is responsible for configuring static ARPs in sync with DAIACLs.
- ARP packets more than 190 bytes on a DAI enabled VLAN will bedropped.
- ARP access-list with longer names is not effective (greater than 20characters)

DHCP-based Firmware Download (DAD – DHCP Automatic Deployment)

- In order for successful version upgrade using DAD method, switch should undergo 2 reloads. For switch in factory default, there is additional reboot to cancel bare metalmode.
- If firmware download is skipped only config download is allowed.
- For dual MM chassis, dual MM must be in sync for DAD to function.
- DAD is a disruptive. ISSU is not supported.
- In FIPS mode, DAD is not supported.
- Cluster principal node failover is not supported.
- DAD over in-band is not supported. Virtual Fabrics is not supported with DAD. You must disable Virtual Fabrics before starting the DAD process in the global configuration file or in the script.
- DAD must complete and succeed on Principal node alone before turn on power for all secondary nodes.
- When the switch is in Factory default, DAD is enabled upon power up theswitch
- DAD executes only if the switch configuration is the default configuration. If the configuration on the switch is not the default configuration, DADexits.
- If the switch is in the default configuration before DAD is triggered, DHCP will remain enabled after the deployment completes. However, this setting can be overwritten by the switch-specific configuration file or the dad.py script.
- You must enable DHCP in the DCMD default configuration to ensure that the switch receives its IP address from the preconfigured DHCP server.
- The factory default DAD runs only once in a DHCP-enabled environment. Irrespective of whether this process is a success or failure, DAD will not be triggered again after a reboot or power off. You can run DAD manually using the dhcp auto-deployment enable command if required.
- Must set ztp=0 in dad configuration file since ZTP (Zero Touch Provisioning) is enabled by default.

- The "vcsmode" value in dad.conf MUST be set to "LC" regardless of whether the existing cluster is in LC or FC mode. If "vcsmode" set to "FC" value in dad.conf, the DAD request can fail.
- DAD is enabled automatically upon switch reboot when you use NOSCLI “write erase” command.

Zero Touch Provisioning (ZTP) consideration
DAD supports up to two nodes for IP fabric in logical chassis mode
All nodes can either be powered up at the same time or enabled from the CLI. This is the key difference vs regular DAD process.

Link State Tracking

- The "track enable/disable" command can only be used to enable or disable the tracking. In order to remove tracking configuration from internal database for a given interface "track remove all" command must be used.

- When there are no uplink interfaces configured, the track disable command will remove tracking configuration from internal database and this behavior is applicable only in 6.0.1a patch and not in prior releases.

- If track min-link number is greater than the number of uplinks, then the downlink will be shutdown with a warning message.

- After toggling the line card using “power-off / on”, LC related interfaces that are configured as uplink interfaces are not seen in “show track summary” clioutput.

OpenFlow

- Interoperability support only with Extreme Controller aka. BVC/BSC.
- Once an interface becomes OpenFlow enabled, very limited set of conventional commands are allowed which includes some of the QoS related configuration commands. For complete list of allowed commands please refer to “*NETWORK OS V6.0.1 SDN Configuration Guide*”
- Priority-tagged frames are not supported.
- L3 Generic flows (incoming port as “*wildcard*”) are not supported.
- PUSH/POP operations can only be associated with action type OFPAT_OUTPUT inside a flow-mod.
- Type of an existing GROUP cannot be changed.
- Existing “clear counter all” command applies to OpenFlow ports as well.
- As part of ISSU, all controller driven configurations will be lost. Controller is expected to re-program after re-connection.
- Uncontrolled Line-Card failover would need power-cycle to recover hardware resources which were in use for the feature to continue to work.
- Uncontrolled failover on 6740 and 6940 would need power-cycle to recover hardware resources for the feature to continue to work.
- Pre-provisioned flow-mods will not be replayed to a new slot coming online. GROUP and METER configurations will be always replayed.
- On the ExtremeSwitching VDX 8770, queue statistics should be interpreted as wire-vlan (COS) priority statistics.

- Actual number of supported flow-mods (L2/L3) may be less since MAX scale values include per port default miss entries, and single LLDP entry is needed for topology discovery. This applies to all supported platforms.
- For layer 3 rules, switch can't differentiate between tagged and untagged flows when matching against rules. This applies to all supportedplatforms.
- Filtering options are not supported for show openflow CLIs. Show openflow commands with filter option show the complete output.
- For the port based flow mod, if the ivid reference is active, egress tagging is not cleared. The new flow mod will not be installed If the previous flow mod has created the egress tagging behavior. This case has to be handled by work-around flow mods or take the port off from openflow and bring it back.
- With default rcv-queue and after coldboot, group select traffic may not be correct, need to do shut/no shut on the interface. This issue is not there with non-defaultrcv-queue.
- With large number of flows, "show openflow flow <>" may take 20 seconds to display packet counts.
- "Module Unknown" is shown for CLI "show open resources" in VDX6940-144S.
- Openflow is not supported on Lag/vlag or port-channel interface.

Layer 2 and Layer 3 ISSU on VDX 6740x

The ISSU functionality on the VDX 6740x (and derivatives) has been added in Network OS 5.0.1. This functionality leverages the HA model that has been delivered on the VDX 8770. It involves running dual-Network OS images on the multi-core control processor. This allows for non-disruptive (to Layer 2, Layer 3, and FCoE traffic) upgrade/downgrade of Network OS 5.0.1 and subsequent minor releases/patches.

ISSU functionality on the VDX 6740x (and derivatives) covers forwarding of Layer 2, Layer 3, and FCoE traffic through the VDX device. Protocols that involve the sending and receiving of Layer 2 and Layer 3 control packets on the VDX device itself are not covered by ISSU. For example, ISSU covers the forwarding of control packets for protocols such as VRRP and OSPF sent by hosts other than the VDX. ISSU allows for non-disruptive upgrades when the VDX is forwarding control packets for other hosts. ISSU does not currently allow for non-disruptive upgrades when the VDX itself is configured for protocols such as VRRP and OSPF and is sending and receiving control packets.

The implementation is based on a type-1 hypervisor.

Vyatta Controller

- Controller does not update the config database based on the flow rejected notification/group rejected notification/meter rejected notification/deletenotification/hard timeout aging notification from switch. Workaround : User needs to delete the flow from the config database and program the correct flow.
- In rare scenario, Controller sends the asynchronous messages leading to flow rejections e.g. flow-mods (associated with group/meter) are rejected after reconnection due to flow-mods being programmed before group/meter config. Work around is that the user needs to delete the group/meter/flow from the config database and program themagain.
- In scale scenario, few flow-mods are not programmed after reconnection. Work around is that the user needs to delete the missing flow-mods and program themagain.

- Topology/Change of interface states are not reflected correctly on BVC.
- Topology with multiple links are not reflected on BVC. BVC shows only single link between the switches.
- Refer to BVC 1.3.0 release noted for all the knownissues/workaround.
- Limitations while configuring flows using BVC:

  1.) MAC addresses- Mac addresses needs to be in uppercase. - Public bug (https://bugs.opendaylight.org/show_bug.cgi?id=2949)

  2.) Ip addresses should have mask – if it is just host say 10.19.18.20 it should be like 10.19.18.20/32 - Public bug(https://bugs.opendaylight.org/show_bug.cgi?id=2950)

  3.) 0s in Ipv6 addresses are rounded ,eg 0000:0000:0000 is rounded to :: . But this is failing in comparison logic and both are treated differently. So use 0000:0000 where :: is there.

  4.) There are some default values ,eg: max_length=0 . They should be set , eventhough they are 0.

  5.) "vlanid-present" in vlan based flows is one field . If you put true, config vs operational will be out of sync (that means flows will have different ids). If you put false or remove the field, flow will not be configured.

Auto QoS for NAS

- From Network OS v5.0.1 onwards, 'nas auto-qos' configuration appears below 'cee-map' configurations in running-config. In earlier versions, it was the other wayround.

As a result of this, if file replay is done using the Network OS v6.0.0 config (with auto-nas configuration) on any previous version (say, Network OS v4.1.0), 'nas auto-qos' configuration will be lost.
User will have to reconfigure 'nas auto-qos' configuration manually.

REST API

- REST configuration for startup-config datastore is not supported.
- Only one command can be configured with one REST request. Configuring multiple commands in a single request is not supported.
- Versioning in API is not supported
- Pagination and Range is not supported.
- Higher level of resource can be obtained with the header "-H "Resource-Depth:x".
- Action related operational commands are not supported.
- Maximum 30 sessions are supported.
- An FCoE Base license is required for the FCoE device to log in. Each FCoE device must have a VF port to log in.

NetConf

- Netconf configuration for startup-config datastore is not supported
- Configuring multiple commands in a single request is supported for configuration/deletion of vlan, switch port, trunk port, VE and rules under IP ACLonly.
- Range is not supported.
- On a large cluster (of 32 nodes or more) and with scaled up configuration, it is recommended to query configuration using rBridge ID filter. In extreme scenario, querying cluster wide configuration without specifying rbridge ID filter might cause switch to run out of memory.
- Maximum 16 sessions supported.

VXLAN Gateway for VMware NSX

- VCS VXLAN Gateway for NSX-MH/NSX-V, is supported only on VDX 6740, VDX6740T,VDX 6740T-1G, VDX 6940-36Q and VDX 6940-144S
- VCS VXLAN Gateway for NSX-MH/NSX-V, is supported only in the VCS Logical Chassismode.
- A maximum of 4 RBridges are supported in a VXLAN enabled VCS Cluster. VXLAN Gateway should be enabled on all the RBbridges of the VCS Cluster.
- Only 1 VTEP Gateway is supported in a VXLAN enabled VCS Cluster.
- VxLAN GW for VMware NSX-MH/NSX-V, and VF Extension cannot be enabled in the same VCS fabric.
- VMware NSX-MH vSwitch with vSphere version 5.5 (ESXi 5.5), and KVM on Ubuntu 12.04 are supported as hypervisors.
- Only one-to-one VLAN to VNI mapping is supported.
- Service and Transport VF cannot be attached to VxLAN GW.
- Tunnel interfaces cannot be used as SPAN (Switch port Analyzer) destination.
- Only Ingress ACL can be applied on tunnels.
- Ingress/Egress QoS policies cannot be applied to tunnels.
- Unicast/Multicast routing between VXLAN and VLAN/VXLAN is not supported.
- BFD should be enabled for all Service node tunnels.
- ALL the VE interfaces should run VRRP-E with the same VRID and same virtual-macto terminate the incoming packets on other VLANs.
- Tunnels egressing/ingressing through an ISL port is supported only on VDX 6940-36Q and VDX 6940-144S.
- Fabric-Virtual-Gateway (FVG) based VTEP is not supported. CLIs for configuring FVG as VTEP are available under overlay-gateway, however these CLIs should not be used as the functionality is not available in this release.
- We certified with NSX 6.3.0 for NSX-V cert using defect628238.

VF Extension using VxLAN

- VF Extension overlay-gateway (VTEP) is supported only on the VDX 6740, VDX 6740T, VDX 6740T-1G, VDX 6940-36Q, and VDX 6940-144S.
- VF Extension overlay-gateway is supported only in the VCS Logical Chassismode.
- VDX 8770 can be in the same VCS fabric where VF-Extension functionality is enabled.
- VxLAN Tunnels are supported over ISL links.
- VF Extension overlay-gateway can be enabled on maximum 4 Rbridges in a VCSCluster.
- VxLAN GW for VMware NSX and VF Extension cannot be enabled in the same VCS fabric.
- Only 1 VF Extension overlay-gateway is supported in a VCS Cluster.
- Only one-to-one VLAN to VNI mapping is supported.
- Tunnel interfaces cannot be used as SPAN (Switch Port ANalyzer)destination.
- Only Ingress ACLs can be applied to tunnels.
- Ingress/Egress QoS policies cannot be applied to tunnels.
- Multicast routing between VXLAN and VLAN/VXLAN is notsupported.
- L3 routing protocols and static routes over tunnels are notsupported.
- Connected subnet L3 forwarding is supported over tunnels.
- Tunnels egressing/ingressing through an ISL port is supported only on VDX 6940 as a VTEP beginning with NOS v6.0.1. Such topologies and configuration must be removed before downgrading to any version below NOS 6.0.1.
- Fabric-Virtual-Gateway (FVG) based VTEP is not supported. CLIs for configuring FVG as VTEP are available under overlay-gateway, however these CLIs should not be used as the functionality is not available in this release.

TCAM Profiles

- The TCAM profiles the user can create may not match the max scale numbers dueto reserved routes/entries which are created for internal use.
- Use count field is added to show the number of entries currently in use.

Management VRF

Layer 3 protocols such as OSPF/BGP/PIM/VRRP/VRRPe are not supported on Management VRF. The following are not supported on in-band ports when they are part of Management VRF:

- DHCP Client functionality
- Auto-config address
- Out-of-band management ports can only be part of Management VRF.
- Switch cannot be managed from leaked routes pointing to Management-VRF.
- Address family on Management VRF cannot be removed.

Conversational MAC Learning

- Conversational MAC Learning and 'Disable Source MAC Learning' cannot beenabled simultaneously on VDX 674x platform.

System level Flowbased QoS

- System Flow based QOS is not supported on the Egress direction.
- QoS can operate on either of three modes – MLS, CEE and MQC. Henceonce service-policy is configured, the interface will be assumed to be in MQC mode and existing MLS and CEE commands will not be supported on the same interface. Un-configuring the policy will put the interface back to default mode which is MLS.
- For Policer, aggregation is possible only within a single chip. Hence when policer is applied on port- channel interface, multi-chip aggregation is not expected.
- SFLOW as action is not supported on Port-Channel interface.
- Any ACL that is used in Flowbased QoS class-map as a match criterion is considered as "QoS ACL" and is special in nature. Hence behavior in some aspects may differ from that of regular "User ACL".
- System based QoS is not supported in egress direction.

Port level Flowbased QoS

- Policer action or SPAN action or both can be applied in egress direction for Port Level Flowbased QoS.
- No other QoS actions are supported in egress direction for port level flowbased QoS.

Non-trivial Merge

- Non-trivial merge is not supported for global configuration. There are a few exceptions in local configuration as well which are not supported for non-trivial merge. This is because these configurations modify global configuration indirectly.
- Modifying the local configurations listed below will result in both a local and global configuration change thereby causing configuration mismatch when ISLs are brought up during fabric formation resulting in node segmentation.

Command (Local Configuration)	Description
/hardware/flexport <interface tuple>/type fibre-channel	Converting an Ethernet interface to Fibre-Channel causes global configuration changes because the Ethernet interface can have configurations in these global configs L2Sys, SPAN, IGMPs, MLDs.
/rbridge-id <#>/vrf <name>	The creation of a VRF on an RBridge will internally create a global partition object which is not visible to the user and used to track the same VRFs created across rbridges in thecluster.

- The Element Manager GUI is intended for use with the VDX 2741 and VDX 2746 platforms only, and may not be used with any other VDX platform.

HA on TOR switches

- HA failover is supported when a user-space daemon is terminated. However, HA failover is not supported on kernel panic. When kernel panic happens, the entire switch will be rebooted for recovery.

Logical Chassis HA

- HA failover and unplanned failover is supported on VDX 8770 only.
- When the principal switch in the VCS cluster undergoing MM failover, it will remain as the principal switch after the MM failover. All the secondary nodes will first disconnect from it when the MM failover starts and then rejoin as the VCS cluster is reformed. At the fabric level, the cluster will remain intact and there will be no traffic disruption.
- When the secondary switch undergoing MM failover, the switch will disconnect and rejoin the VCS cluster after reestablishing connection with the principal switch and the rest of the cluster will stay intact. At the fabric level, the cluster will remain intact and there will be no traffic disruption.
- RMON HA is not supported.
- vMotion during HA failover is not supported.
- If UDLD is enabled, HA is supported with a higher range for the UDLD hello time (> ow1 sec)
- HA is not supported for OpenFlow feature, however, system level ISSU is supported. For ISSU, it is recommended that the controller is disconnected first, all flows are cleared using "clear OpenFlow all" command and then perform the upgrade.

Interoperability

- In a VPC environment where the ExtremeSwitching VDX side has the active LACP settings and the Cisco side has the passive settings on the vLAG, the port-channel takes over 30 seconds to come up.
  Workaround: Reverse the settings and have the ExtremeSwitching VDX LACP settings passive and the Cisco side set as active. The port channel will then restore after about 10 seconds.
- VDX interop with Cisco Nexus switch with 'peer-switch' enabled on VPC is not supported.
- When interoperating with a Extreme 8000, it is recommended to set the ***mac-aging*** time to 0 on the VDX switch to prevent any adverse impact caused by certain errors generated by the Extreme 8000.
- ADX HA Sync packets use UDLD PDU's which may be dropped by VDX . To enable forwarding, we recommend configuring dot1q tagging to treat UDLD packets as data-packets to be forwarded across VCS.Virtual Fabric.
- PIM-SM is not supported on Virtual Fabric on VDX8770.
- For frames forwarded on a transport fabric, ingress CTAG tagging is preserved at the egress port regardless of the egress tagging classification.
- Default-VLAN can only be configured using TRANSPORT VF IDs.
- The "no vcs virtual-fabric enable" command execution time is dependent on the number of ISLs and VLANs in the VCS.
- The virtual-fabric resource allocation are platform dependent as follows:
  - VDX 8770 – no limitation
  - VDX 6740/6740T/6740T-1G – uses TCAM table
  - VDX 6940-36Q – virtual-fabric transport and service VLANs use TCAM and EXM table respectively.

MAPS

- BNA version 14.0.1 is supported with NOS7.0.1
- MAPS is supported on VDX 2741, 2746, 6740, 6940 and 8770platforms.
- MAPS port level alerting in NOS V7.0.1 is not available for Flex Ports configured inFiber Channel mode.
- RX_SYM_ERR MAPS messages are displayed when breakout cable is connected on a 40G interface that is not configured for breakout.
- When line card on the remote end of the link is powered off, MAPS generates Insertion/Removal notification for the SFPs on the local side. These can beignored.
- 100G SFP threshold monitoring is not supported onVDX6940-144s.

Maintenance Mode

- Port-channel configuration changes while a node is in maintenance-mode is notsupported.
- Configuration replay of a saved configuration file or snapshot containing both maintenance-mode and port-channels is not supported.

BNA

Recommendations to customer when the cluster size is 32 or more nodes

- Make sure the lazy polling period is 60 minutes.
- Disable event based polling in such large clusters. Essentially this means there will not be any update from the cluster for BNA till the lazy period iselapsed.

Miscellaneous

- ExtremeSwitching VDX switches load balance internal and external traffic based on hash functions using standard network headers as keys. Due to this implementation, users may experience traffic imbalance depending upon application flow definition.
- Packet drops will be seen for a short duration due to routing changes with link flaps and/or node failovers.
- On both ISL and Edge ports, sFlow sampling is supported only in inbounddirection.
- Sflow collectors are not queried in snmp v1, v2 & v3 versions.
- L2 packets may not be sampled on line-card power OFF & ON.
- If multiple VLANs are configured on a switch, then in order to enable certain features such as IGMP or PVST it is recommended that specific features be enabled on a per-VLAN basis instead of enabling them globally.
- The VLANs 4087-4095 and 1002 are reserved and used for internal clusteroperations.
- “Clear ip route all” need to be issued once the maximum number ofroutes supported by a router is exceeded.
- SNMP supports 2k OCTET-STRING size for MIB objects.
- Snmpwalk operation on TCP MIB (RFC 4022) may become very slow and timeouts may happen on all VDX platforms. The snmpwalk timeout should be set to at least 3 seconds while walking the TCP MIB.

- Under rare conditions, the switch may bootup with the default configuration upon power-cycling the switch.
- Firmware downgrade is not blocked if the scale configured would not be supported in the downgraded release.
- Under rare conditions, after disabling keepalive timeout followed by shut & no shut of the port-channel link may prevent FCoE logins through that port-channel.
- On rare instances of HA failover, SFM may turn faulty. Workaround is to manually reseat the card.
- On rare instances of ISSU, HA failover, line-card may turn faulty. Workaround is to reset the line-card.
- PCAP utility is not supported on standby MM on VDX 8770.
- Please make sure to not have large no of unreachable tacacs+ accounting server configured, else it might cause unit to reboot. This issue is hit only with large config (4K VLAN etc and 20K lines or config).
- Configuration of more than one In-band management port on a single switch is not recommended.
- Under certain stress conditions the 'copy support' command might time out for some modules. In such cases, it is recommended to retry 'copy support' with a higher timeout multiplier value.
- It is highly recommended to copy the configuration file to running-config and then save the running-config to startup-config, instead of directly copying the external configuration file to startup-config, especially when using fabric distributed features such as Zoning, VM Aware Network Automation and Virtual IP.

Defects

TSBs - Critical Issues to Consider Prior to Installing This NOS Release

Technical Support Bulletins (TSBs) are produced to provide detailed information about high priority defects or issues present in NOS releases. The following sections specify all current TSBs that have been identified as being a risk to or resolved with this specific version of Extreme Network OS. Please review carefully and refer to the complete TSB for relevant issues prior to migrating to this version of code. TSBs can be found at https://extremeportal.force.com/ (note that TSBs are generated for all Extreme platforms and products, so not all TSBs apply to this release).

TSB Issues Resolved in Network OS v7.0.1

TSB	Summary
TSB 2016-236-A	A component within the 10G Base-T power circuitry can overheat causing the switch to power off and be unusable. Only 6740-T switches configured with "Port Side Exhaust" (-R fan direction) are at risk to this specific component overheating condition. VDX 6740-T switches with "Port Side Intake" provide sufficient airflow over the specific component to preventoverheating. The number of failures observed in the field have been very low. The number of total failures of VDX 6740-T switches, counting all reasons including this specific failure, are well below the predicted failure rate for this HW platform. Switch will power down and becomeunusable. The switch may power down due to a detected thermal event or it may power down due to loss of an internal voltage. There may not be any log recorded ahead of the power down. A SW solution to increase the nominal fan speed to ensure sufficient airflow over the circuitry to prevent any overheating of the specific component has been developed. Upgrading to one of the specified firmware versions or any higher version of the Extreme Network OS will provide an increased fan speed and airflow to prevent any overheating. The nominal power consumption of the switch will only be increased by 20W with this change, and the maximum power consumption is not affected by this change.

Network OS v7.0.2b Caveats

None

Network OS v7.0.2b Caveats

None

Network OS v7.0.2 Caveats

None

Network OS v7.0.1c Caveats

None

Network OS v7.0.1b Caveats

SSH

- Shutting down SSH server does not close all existing SSH login sessions. Shutdown SSH server.

IP Fabric

- ISSU upgrade to NOS7.0.1a can cause 19 seconds of traffic loss if BFD is configured.Please disable BFD.

VLAN - Virtual LAN

- Packets are flooded to all the vlan member interfaces of the remote node even though Static-mac blinding is configured. When Multicast mac address is used and Vlan member interfaces are present in the remote nodes of the cluster

DHCP - Dynamic Host Configuration Protocol

- PV4 DHCP relay statistics does not display the count for DHCP offer and DHCP Ack packets relayed by switch from the DHCP server to DHCP client. Execution of "show vlan brief" CLI after reload.

Monitoring

- Multiple FFDC and Core files seen on firmware downgrade. When 8770-8 chassis runningsix new version line-cards 6x100G, 27x40G, 48x10G-T with three or more of type 6x100G is downgraded from 7.0.1b version to lower version.

OSPF

- When overlapping routes (such as routes with same network address, but different mask/prefix-length) are redistributed into NSSA area, withdrawal of more specific route inadvertently results into the withdrawal of less specific route.

- After doing HA failover on M4, HA is not in SYNC and observed ONM crash. HA failover and OSPFv3 - IPv6 Open Shortest Path First
  timing issue can trigger the issue, very rare to occur.

CLI

- After setting the default RA interval using "ipv6 nd ra-interval 600" the line does not disappear from config anymore. There is no work-around.

Network OS v7.0.1a Caveats

None

Network OS v7.0.1 Caveats

BFD

Although the BFD timer values are allowed to be configured below default values of 200 ms (VDX8770) and below 500 ms (VDX6740, 6940), only default values and above are recommended.

VxLAN

- For VXLAN tunnel packets, the IP MTU check on egress is bypassed to allow larger size packets. Any fragmentation occurring on the underlay transit nodes will result in failure of VxLAN termination at the destination VTEP. So, if a packet of size greater than configured L3 MTU of 9018 Bytes is forwarded through the tunnel, the packet will pass through and the transit node shall fragment or discard the packet based on the fragmentation support on the node and the DF bit set on the packet.

  Note:

  DF bit is set on VDX6940 and not set on packets originating from VDX6740

  Packet Fragmentation is supported on VDX8770 and not supported on VDX6740 and 6940 platforms.

- On occurrence of events that may bring down the tunnel on an R-Bridge, there could be few seconds of traffic interruption due to a default de-bounce-timer which is set to 2 secs, this could delay the fail-over of the traffic to redundant path. A debug command “show system internal tnlmgr de-bounce-timer 0 0” can be utilized to reduce the traffic impact, however, the command settings are not persistent across reloads.

- On sending IGMP queries over VF_Extension Tunnel with VLAG as underlay, packets might loop over the tunnel .Queries come back from the same tunnel interface from which its egressedout.

- “show ip igmp groups interface tunnel <tunnel_id>" cli shows all IGMP interfaces instead of just the tunnel interface.
- Adding and Removing RBridges under overlay-gatway may take longer than expected time if large number of VLANs are configured in the fabric.

Long Distance ISL

- The "long-distance-isl" functionality on an interface will not be preserved although "long-distance-isl" configuration is displayed in running-config when the following actions are performed:
    1. Configuring "long-distance-isl" on an "administratively down" ISL interface.
    2. VCS or switch reload/Chassis disable-enable/interface shut-no shut/Firmware download with "coldboot" option
- It is recommended the user configure any "long-distance-isl" configuration while the ISL interface is in the "administratively up" state.
- If the "long-distance-isl" persistent issue is encountered, the user can recover by manually removing the "long-distance-isl" configuration and reconfigure.

Loopback interfaces

- On topologies where same IP address is configured on loopback interfaces on multiple nodes in a cluster, performing admin down of loopback interfaces may result in pingissues.

Route distribution

- When redistribute bgp metric command is unconfigured, the configuration is not completely removed. It is required to configure redistribution without metric and then unconfigure again to unconfigure it completely.

FCoE/Access Gateway

- If a node with FCoE interfaces configured with local logical SAN is reloaded, the FCoE logins may fail to come online. In order to recover, remove and configure the respective local logical SAN fabric-map.

BNA/NetConf/REST

- Special character ‘$’ under the custom RPC “bna-config-cmd” cannot be used for Netconf and REST API for performing copy operation.
- REST API deletion on the main resource will remove all the sub-resources under it. For Example, REST API delete Operation without specifying ACL name will remove all the ACLs in the system. Specify the ACL name in the request in order to delete particular ACL from theconfig.
- For large scale VCS fabrics with more than 4000 ports, querying the cluster with BNA/REST APIs may result in switch software exception. For this purpose it is not recommended to enable BNA monitoring or querying with REST APIs for large VCS fabrics.

AAA Configuration

- The number of user accounts is limited to 60. Adding any additional accounts and performing add/remove user operations may result in a Switch Software Exception.

Sync Failure Error

- If an error “CRITICAL, VDX8770-4, FSS Error on service component [ethsw1:eswc]: sync-failure: -994” is observed when DHCP IP helper functionality is enabled between 2 different VRFs please contact Extreme Support for defect confirmation and recovery steps.

Mac Loop Detect Feature:

- “Loop detection may not take action of shutting down the interfaces in a high scale environment with greater than 20K macs flapping at a time”.
- “MAC-move detect feature may shutdown the Server port under certainconditions”.

Port Channel Scalability:

- Under certain circumstances, port-channel configured with Extreme protocol, may limit the maximum scale number to a lower value.
- Port-channel vLAG/LAG may not re-establish after issuing “no vlag-commit-mode disable”. User may require to delete and re-configure the port-channel interface and memberlinks.

AMPP/vCenter:

- Event notification is not received for the second host move, when more than one host is moved from one data-center to another in vCenter 6.0.0. The hosts would still be part of old data-center and workaround is to initiate a manual discovery
- Event notification is not received when the VLAN of two identical port-groups are modified and the running config doesn't change. Workaround is to initiate a manualdiscovery.
- Output of show vnetwork vmpolicy command is not displaying the VM name and datacenter-id for a cloned VM. Workaround is to initiate a manualdiscovery.

OpenFlow:

- With default rcv-queue and after coldboot group select traffic may not be correct, need to do shut/no shut on the interface. This is not observed with non-defaultrcv-queue.
- With large number of flows, “show openflow flow <>” may take 20 seconds to displaypacket counts
- Filtering options (e.g. show | include) will not work for show openflow commands. show commands will display the complete output.
- "Module Unknown" is shown for CLI "show open resources" in VDX6940-144S.

• Hardware Profile:

- When modifying the route-table profile type and maximum-path using the hardware-profile command, the user should only change one parameter at a time. Otherwise the maximum-path setting will be incorrect. If the issue already occurred, the user can re-run the command to set the maximum-path with the correct value.

• Copy Config command:

- In VDX6940-144S, 100G mode configuration replay can fail when executing "copy <file> running-config" if DPOD license is not reserved. To work around this issue, the user can manually reserve the license and then run “copy <file> running-config”.

Syslog:

- Syslog server configured with same IP across the VRFs in inband will not receive the messages.

Closed with code changes for NOS v7.0.2d

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as November 2020 in Network OS v7.0.2d.

Parent Defect ID:	NOS-67866	**Issue ID:**	NOS-67866
Severity:	S4 - Low		
Product:	Network OS	**Reported in Release:**	NOS7.0.2c
Technology Group:	Security	**Technology:**	Security Vulnerability
Symptom:	VDX is vulnerable if telnet connection enabled.		
Condition:	VDX is vulnerable if telnet connection enabled.		
Workaround:	Disable the telnet feature and use SSH for secure login to switch		

Parent Defect ID:	NOS-67674	**Issue ID:**	NOS-67908
Severity:	S3 - Medium		
Product:	Network OS	**Reported in Release:**	NOS7.0.2b
Technology Group:	Monitoring	**Technology:**	MAPS - Monitoring and Alerting Policy Suite
Symptom:	Continuous RASLOG, SNMP, EMAIL notification for Rx_sym_err.		
Condition:	Whenever any symbol errors are logged		
Workaround:	None		

Closed with code changes for NOS v7.0.2c

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as July 2019 in Network OS v7.0.2c.

Parent Defect ID:	NOS-48024	**Issue ID:**	NOS-67136
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Management
Reported in Release:	NOS6.0.2h	**Technology:**	SNMP - Simple Network Management Protocol
Symptom:	Unexpected system reload		
Condition:	Reload trigged when polling ipRouteTable (1.3.6.1.2.1.4.21) SNMP		
Workaround:	None		

Parent Defect ID:	NOS-66864	**Issue ID:**	NOS-67137
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Other
Reported in Release:	NOS7.4.0	**Technology:**	Other
Symptom:	NOS version 7.0.2b, 7.2.0b and 7.4.0 will not allow the LC48X10GT to boot up due to a missing file		
Condition:	LC48X10GT will not boot up online		
Workaround:	None		

Parent Defect ID:	NOS-67061	**Issue ID:**	NOS-67138
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Management

Reported in Release:	NOS7.0.2a	**Technology:**	Configuration Fundamentals
Symptom:	Linecard reset by the detection of Multibit Parity error condition show in RAS Log. "[HWK2-5651], 14604207/0, L4/0 | Active, ERROR, VDX8770-8, CHIP0: Interrupt: DMC dmc mbit err."		
Condition:	Linecard reset.		
Workaround:	None		

Parent Defect ID:	NOS-66900	**Issue ID:**	NOS-67139
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Layer 2 Switching
Reported in Release:	NOS6.0.2h	**Technology:**	LAG - Link Aggregation Group
Symptom:	Some of the mac addresses will not be learnt after firmware upgrade.		
Condition:	On firmware upgrade.		
Workaround:	None		

Parent Defect ID:	NOS-66967	**Issue ID:**	NOS-67140
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Layer 2 Switching
Reported in Release:	NOS6.0.2h	**Technology:**	LAG - Link Aggregation Group
Symptom:	MAC address remains inactive even after the PVLAG link is down		
Condition:	When we shut and no shut the PVLAG interface multiple times		
Workaround:	None		

Parent Defect ID:	NOS-66866	**Issue ID:**	NOS-67141
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Security
Reported in Release:	NOS6.0.2f	**Technology:**	AAA - Authentication, Authorization, and Accounting
Symptom:	Unexpected Reload		
Condition:	When TACACS authorization fails on re-try		
Workaround:	None		

Parent Defect ID:	NOS-66958	**Issue ID:**	NOS-67142
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Management
Reported in Release:	NOS6.0.2g	**Technology:**	Configuration Fundamentals
Symptom:	Dcmd process died and switch reloaded while doing "copy running-config startup-config"		
Condition:	Running the CLI "copy running-config startup-config" when standby in rebooting.		

Parent Defect ID:	NOS-47973	**Issue ID:**	NOS-67147

Severity:	S1 - Critical		
Product:	Network OS	**Technology Group:**	Data Center Fabric
Reported in Release:	NOS6.0.2e	**Technology:**	VCS Fabric
Symptom:	Unexpected reload.		
Condition:	There are certain conditions that may cause the IPv6 traffic to be sent to the CPU.		
Workaround:	None		

Parent Defect ID:	NOS-54704	**Issue ID:**	NOS-67148
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Management
Reported in Release:	NOS7.1.0b	**Technology:**	SNMP - Simple Network Management Protocol
Symptom:	Unexpected reload		
Condition:	when two tlvs are received but tlv type of one is either zero or unsupported.		
Workaround:	None		

Parent Defect ID:	NOS-66761	**Issue ID:**	NOS-67149
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Data Center Fabric
Reported in Release:	NOS6.0.2h	**Technology:**	Logical Chassis
Symptom:	Unexpected reload		
Condition:	Switch receives the more Fiber Channel traffic and it leads to corruption.		

Parent Defect ID:	NOS-36884	**Issue ID:**	NOS-67150
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Layer 2 Switching
Reported in Release:	NOS5.0.1	**Technology:**	VLAN - Virtual LAN
Symptom:	"show interface description" and "show interface trunk" do not display all port-channels.		
Condition:	Multi-node cluster with Port channels configured and most of the port-channels have ports from non-principal nodes.		
Workaround:	None		

Parent Defect ID:	NOS-67017	**Issue ID:**	NOS-67151
Severity:	S2 - High		
Product:	Network OS	**Technology Group:**	Layer 2 Switching
Reported in Release:	NOS6.0.2ha	**Technology:**	LAG - Link Aggregation Group
Symptom:	Momentary traffic black-hole is observed		
Condition:	When we shut the peer LAG port or reload the peer device		
Workaround:	None		

Parent Defect ID:	NOS-66998	**Issue ID:**	NOS-67152

Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Data Center Fabric
Reported in Release:	NOS7.0.0	**Technology:**	Logical Chassis
Symptom:	No reason at RASLOG when interface administratively down with out user config.		
Condition:	When exception created on port due to ASIC fault.		
Workaround:	None		

Parent Defect ID:	NOS-66927	**Issue ID:**	NOS-67153
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	IP Multicast
Reported in Release:	NOS7.1.0a	**Technology:**	IPv4 Multicast Routing
Symptom:	switch reloaded after running "show ip pim mcache"		
Condition:	Running the show ip pim mcache		

Parent Defect ID:	NOS-66997	**Issue ID:**	NOS-67156
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	**Technology:**	ARP - Address Resolution Protocol
Symptom:	Repeated IPAD-1001-log entries even though there is no change to mgmt interface status		
Condition:	mgmt-vrf default route is resolved via inband interface		
Workaround:	do not configure default route with nh pointing to inband interface		

Parent Defect ID:	NOS-66898	**Issue ID:**	NOS-67170
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	**Technology:**	BGP4+ - IPv6 Border Gateway Protocol
Symptom:	Unexpected reload.		
Condition:	When we run the "show ip bgp neighbors" after giving the "no neighbor activate" followed by "clear ip bgp neighbour"		
Workaround:	None		

Parent Defect ID:	NOS-67109	**Issue ID:**	NOS-67179
Severity:	S3 - Medium		
Product:	Network OS	**Technology Group:**	Management
Reported in Release:	NOS7.0.2b	**Technology:**	CLI - Command Line Interface
Symptom:	Un expected reload.		
Condition:	switch reloaded after issuing show startup-config in multiple sessions.		
Workaround:	Avoid parallel request of show startup-config		

Closed with code changes for NOS v7.0.2b

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as November 12, 2018 in Network OS v7.0.2b.
The rebranding UBoot change is available beginning with release NOS v7.0.2b.

Defect ID:	DEFECT000629607		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Hardware Monitoring
Reported In Release:	NOS 4.1.3	**Technology:**	Monitoring
Symptom:	The switch may power down due to a detected thermal event or it may power down due to loss of an internal voltage. There may not be any log recorded ahead of the power down		
Condition:	VDX8770-4 and VDX8770-8 with any line card part number that has a "-xx" PRIOR to these listed: VDX8770-6X100G-CFP2: 60-1002821-17 (i.e. -16 or lower is affected) VDX8770-48X10G-T: 60-1002909-16 (i.e. -15 or lower is affected) VDX8770-27X40G-QSFP: 60-1002792-12 (i.e. -11 or lower is affected)		
Workaround:	None. If any line cards in the system are below the revisions indicated above, the software must be upgraded to one of the versions above		

Defect ID:	DEFECT000662379		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Data Center Fabric
Reported In Release:	NOS 7.3.0	**Technology:**	Logical Chassis
Symptom:	VDX device will startup with the default configuration and reports the same using below RAS log. [DCM-3053], SW/0 | Active, INFO, VDX-VCS, Dcmd database corruption was detected. The system will startup with the default configuration for this database.		
Condition:	Sudden Power Cycle the device can cause the issue.		

Workaround:	We can use below workaround for planned outrage or power cycle. **NOS7.1.0 and Later Releases**: 1. Execute chassis power-cycle-db-shutdown command through NOSCLI. 2. Reload the switch after the below RASLOG: [DCM-1015], SW/0 | Active, INFO, VDX6740T, Switch is prepared for power-cycle. No CLIs will work henceforth. Reload or power cycle to make switch fully functional. **For releases prior to NOS7.1.0:** Root level command. root> shutdowndcmdb 2018/09/20-20:58:29 : shutdowndcmdb : Shutting Down Database ..... (New)

Closed with code changes for NOS v7.0.2a

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of July 31, 2018 in Network OS v7.0.2a.

Defect ID:	DEFECT000609232		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS7.1.0	**Technology:**	VCS
Symptom:	1. Compact Flash grows and disk full issue can occur. 2. Unexpected DCMd daemon termination.		
Condition:	No special condition or configuration required to hit this issue.		
Workaround:	For syslog.log file growing beyond limit below is preventive workaround: Please comment below two lines from syslog config & template files and reread config file from root using below command [killall]. It should stop all those IO messages. Please also verify that syslog utility is working fine after the workaround applied to make sure all is fine before we try on customer production environment. /etc/syslog-ng/syslog-ng.conf /etc/syslog-ng/syslog-ng.template #destination df_kern { pipe("/var/log/kmsg"); }; #log { source(s_all); filter(f_kern); destination(df_kern); }; root > /usr/bin/killall -HUP syslog-ng		
Recovery:	Empty /var/log/syslog.log file if it is growing beyond 1 Mb.		

Defect ID:	DEFECT000633219		
Technical Severity:	Medium	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	VRRPv3 - Virtual Router Redundancy Protocol Version 3
Reported In Release:	NOS7.1.0	**Technology:**	Layer 3 Routing/Network Layer
Symptom:	VDX experience unexpected reload due to DCMd daemon termination.		
Condition:	Execution of CLI: sw0(config-vrrp-extended-group-1)#track network 0.0.0.0/0 priority 50.		

Defect ID:	DEFECT000645906		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	FCoE - Fibre Channel over Ethernet
Reported In Release:	NOS5.0.2	**Technology:**	Layer 2 Switching
Symptom:	FCOE flapping on some FCOE devices until reloaded server after adding new VDX into VCS		
Condition:	Cluster disturbance		
Recovery:	Recovery-------- Apply "shut/noshut" on problematic physical interfaces		

Defect ID:	DEFECT000646316		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	ARP - Address Resolution Protocol
Reported In Release:	NOS6.0.2	**Technology:**	Layer 3 Routing/Network Layer
Symptom:	Unexpected reload of switch.		
Condition:	Removing L3 configs (in specific IPv4 addresses) and defaulting the config for VDX.		

Defect ID:	DEFECT000646540		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	OSPFv3 - IPv6 Open Shortest Path First
Reported In Release:	NOS7.1.0	**Technology:**	Layer 3 Routing/Network Layer
Symptom:	Message generic error at CLI console		
Condition:	While removing OSPFv3 configuration		

Defect ID:	DEFECT000647282		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Hardware Monitoring
Reported In Release:	NOS7.0.1	**Technology:**	Monitoring
Symptom:	1G port link flapped in VDX6740-T.		
Condition:	On VDX6740-T if the peer end is connected to Intel NIC, auto negotiation will fail, resulting in flapping of 1G port.		

Defect ID:	DEFECT000647840		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS6.0.2	**Technology:**	VCS
Symptom:	System may undergo unexpected reload		
Condition:	Media removal while media data is reading		
Workaround:	shut/ no shut media removed interface		

Defect ID:	DEFECT000647847		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	LAG - Link Aggregation Group
Reported In Release:	NOS6.0.2	**Technology:**	Layer 2 Switching
Symptom:	Unexpected reload		
Condition:	In rare a case, DB corruption happens at the time of port-channel deletion.		

Defect ID:	DEFECT000648291		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	CLI - Command Line Interface
Reported In Release:	NOS7.0.1	**Technology:**	Management
Symptom:	Help string update for SSH related CLIs. Keyword "etc..." got removed.		
Condition:	For the below CLIs sw0(config-rbridge-id-1)# ssh server key-exchange ? ssh server cipher ssh server mac ssh client key-exchange ssh client cipher ssh client mac		

Defect ID:	DEFECT000651850		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	SNMP - Simple Network Management Protocol
Reported In Release:	NOS7.2.0	**Technology:**	Management
Symptom:	SNMP sysName query returns hostname instead of FQDN.		
Condition:	When SNMP sysName OID is queried.		

Defect ID:	DEFECT000652192		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	OSPF - IPv4 Open Shortest Path First
Reported In Release:	NOS7.0.1	**Technology:**	Layer 3 Routing/Network Layer
Symptom:	"OSPF-1003 - Received Invalid LS packet" RASLOGs get flooded.		
Condition:	Unexpected reload of the switch.		

Defect ID:	DEFECT000654900		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	OpenStack Integration
Reported In Release:	NOS7.1.0	**Technology:**	Network Automation and Orchestration
Symptom:	1G Port won't come online.		
Condition:	Connected 1G with 10G at other end.		

Defect ID:	DEFECT000655619		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Management GUI
Reported In Release:	NOS7.0.1	**Technology:**	Management
Symptom:	BNA hangs on VDX logical chassis firmware downgrade from 7.0.1c to 7.0.1b or 6.0.1f to 6.0.1e		
Condition:	VDX logical chassis firmware downgrade using BNA.		
Workaround:	Use NOS CLI for firmware downgrade rather than BNA.		

Defect ID:	DEFECT000656869		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS7.1.0	**Technology:**	Data Center Fabric
Symptom:	Port does not came online on VDX 6740-T platform		
Condition:	Port didn't came online when the peer server is CentOS was rebooted multiple times.		

Defect ID:	DEFECT000657045		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	HTTP/HTTPS
Reported In Release:	NOS7.0.2	**Technology:**	Security
Symptom:	HTTPS will be enabled if expired TLS certificate and key is imported to device using scpuser credentials. HTTPs should not be enabled if the certificate is expired.		
Condition:	When expired TLS certificate is imported to device using scpuser credentials, HTTPS can be enabled even with expired TLS certificate.		
Workaround:	Expired TLS certificate should not be imported to device.		
Recovery:	Import valid TLS certificate.		

Defect ID:	DEFECT000657950		
Technical Severity:	Medium	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Syslog
Reported In Release:	NOS7.0.1	**Technology:**	Monitoring
Symptom:	After adding a VDX to an existing VCS using "vcs replace", the newly added VDX is unable to send messages to a remote syslog server.		
Condition:	The newly added or reconnected VDX will be able to see "logging syslog-server" settings in "show run", but it will not be able to send syslog messages to that remote server		
Workaround:	After this issue has occurred on a newly added non-principal, it is possible to recover by removing and re-applying the "logging syslog-server" setting on the VCS principal rbridge.		

Defect ID:	DEFECT000658011		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	TACACS & TACACS+
Reported In Release:	NOS7.1.0	**Technology:**	Security
Symptom:	Tacacs accounting functionality does not work properly.		
Condition:	In VCS cluster node rejoin operation can cause this issue.		

Defect ID:	DEFECT000658079		
Technical Severity:	Low	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	OSPF - IPv4 Open Shortest Path First
Reported In Release:	NOS7.0.1	**Technology:**	Layer 3 Routing/Network Layer
Symptom:	Route summarization does not happen even after configuring it on the device.		
Condition:	This issue is seen when configured route summarization prefix triggers OSPF Appendix E calculation with the existing Type 3 LSAs.		

Defect ID:	DEFECT000658974		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Configuration Fundamentals
Reported In Release:	NOS7.2.0	**Technology:**	Management
Symptom:	Default-config operations [copy default-config startup-config, FWDL with default-config, netInstall] does not preserve the DHCP configuration on management interface		
Condition:	Performing default-config operations [copy default-config startup-config, FWDL with default-config, netInstall].		

Defect ID:	DEFECT000659778		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS6.0.2	**Technology:**	VCS
Symptom:	For VDX 6740 and VDX 6740-T, during the firmware upgrade, customer experienced Compact Flash card issue with the following scenario: - Rolling reboot - Console log message of ?SCSI_REQ_SENSE failed cmd 0x03 returned 0x70 0x06 0x28 0x00? and/or ?Hypervisor Reset Flush?		
Condition:	Un-correctable internal errors occurred on the Compact Flash card that used to store programs and data.		
Workaround:	Recovery using netinstall is possible, but some units fail again after some time even after a netinstall procedure has recovered the system.		

Defect ID:	DEFECT000659781		
Technical Severity:	Medium	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS6.0.2	**Technology:**	VCS
Symptom:	For VDX 6740 and VDX 6740-T, during the firmware upgrade, customer experienced Compact Flash card issue with the following scenario: - Rolling reboot - Console log message of ?SCSI_REQ_SENSE failed cmd 0x03 returned 0x70 0x06 0x28 0x00? and/or ?Hypervisor Reset Flush?		
Condition:	Un-correctable internal errors occurred on the Compact Flash card used to store programs and data.		
Workaround:	Recovery using netinstall is possible, but some units fail again after some time even after a netinstall has recovered the system.		

Defect ID:	DEFECT000660553		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS7.0.1	**Technology:**	Data Center Fabric
Symptom:	Cannot configure IPv6 (/126) address on VIP for VRRP-E.		
Condition:	Configuring IPv6 address (/126) for VRRP-E		

Defect ID:	DEFECT000660697		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Logical Chassis
Reported In Release:	NOS7.0.1	**Technology:**	Data Center Fabric
Symptom:	unexpected core files fills up disk.		
Condition:	"show logging raslog rbridge-id" CLI execution for multiple rbridge at the same time.		

Defect ID:	DEFECT000661476		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	CLI - Command Line Interface
Reported In Release:	NOS7.2.0	**Technology:**	Management
Symptom:	"Please check the valid CLI format, host IP address, and the permission and space left on the remote directory." Error message comes on terminal.		
Condition:	Change in RSA host key of the management server.		

Defect ID:	DEFECT000661527		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Software Installation & Upgrade
Reported In Release:	NOS7.0.2	**Technology:**	Management
Symptom:	Lost configuration during upgrade tsd terminated with core dump		
Condition:	Upgrade from 6.0.2e to 7.0.2		

Defect ID:	DEFECT000661579		
Technical Severity:	Critical	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Management GUI
Reported In Release:	NOS6.0.2	**Technology:**	Management
Symptom:	"write erase" removes /var/spool/cron/root crontab config file and as a result all crontab functionality is impacted. Ex: /var/log/syslog.log file can grow beyond 100k as log rotation doesn't work.		
Condition:	execute "write erase" .		

Defect ID:	DEFECT000661695		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	LAG - Link Aggregation Group
Reported In Release:	NOS6.0.2	**Technology:**	Layer 2 Switching
Symptom:	Traffic disruption in the cluster due to unresponsive rbridge		
Condition:	In rare conditions, the ISLs stay up on an unresponsive rbridge.		

Defect ID:	DEFECT000661782		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Hardware Monitoring
Reported In Release:	NOS6.0.2	**Technology:**	Monitoring
Symptom:	Slow kernel memory leak due to 'aapl_malloc+0x38/0x8c [dce_blade_module]'. Leak is 4MB per day.		
Condition:	Memory leak of 4MB per day due to 'aapl_malloc+0x38/0x8c [dce_blade_module]'		

Defect ID:	DEFECT000662522		
Technical Severity:	Critical	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	VLAN - Virtual LAN
Reported In Release:	NOS6.0.2	**Technology:**	Layer 2 Switching
Symptom:	Traffic impact or packet loss between directly connected hosts.		
Condition:	Traffic impact or packet loss between directly connected hosts.		

Defect ID:	DEFECT000663071		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Hardware Monitoring
Reported In Release:	NOS6.0.2	**Technology:**	Monitoring
Symptom:	No asic parity error messages in RASlog.		
Condition:	Switch did not go to faulty state even though there were parity errors.		

Closed with code changes for NOS v7.0.2

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as November 22, 2017 in Network OS v7.0.2.

Defect ID: DEFECT000550982	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Switch management port does not generate a ColdStart trap if a Management port is configured to acquire the IP address via DHCP.	
Condition when switch is configured to acquire IP address via DHCP, then we will observe this issue.	
Workaround: If IP is configured statically, the issue will not happen.	

Defect ID: DEFECT000579904	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.2	**Technology:** AAA - Authentication, Authorization, and Accounting
Symptom: Command set field on the Windows based TACACS server is empty	
Condition 1. When TACACS server is windows based 2. Accounting is enabled	

Defect ID: DEFECT000600591	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: Logs are dumped on the screen, when there is a read failure on SFPs connected to the port.	
Condition Accessing information about the SFPs inserted in the ports.	
Recovery: Disable the port and re-enable it.	

Defect ID: DEFECT000619425	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.1.0	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: Traffic loss on Port-channel interface	
Condition If Global MTU is smaller than Port-channel MTU or Global MTU is configured and un-configured, user may see traffic loss on port-channel interface.	
Workaround: Configure MTU same as port-channel on Port-channel member interfaces	

Defect ID: DEFECT000623805	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.2	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom: l2traceroute request failure via REST API	
Condition After 1000 successful l2traceroute request, any subsequent request for l2traceroute will fail.	
Recovery: Reload of the switch recovers from this state	

Defect ID: DEFECT000628230	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** VLAN - Virtual LAN
Symptom: member interface of a port-channel fails to learn source MAC addresses resulting in flooding.	
Condition port-channel member interface is configured as fcoeport after port-channel is configured as switch port.	

Defect ID: DEFECT000631176	
Technical Severity: Low	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.1.0	**Technology:** CLI - Command Line Interface
Symptom: Ambiguity in IP MTU field of "show interface" output. Cosmetic issue, no functional impact.	
Condition L3 Interface is configured back to L2.	

Defect ID: DEFECT000634769	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** Configuration Fundamentals
Symptom: SCP file transfer fails	
Condition Using double quotes ("") for file names with spaces causes SCP to fail on certain servers.	
Workaround: Use file names without spaces	

Defect ID: DEFECT000636143	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** LAG - Link Aggregation Group

Symptom:	Cosmetic issue. some of fields (actor system id, Receive link count, Transmit link count, Individual and ready) won't display properly at "show port-channel detail | nomore" output.
Condition:	Rare scenario. Execution of "show port-channel detail | nomore".

Defect ID:	DEFECT000637104		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 2 Switching
Reported In Release:	NOS7.1.0	**Technology:**	VLAN - Virtual LAN
Symptom:	Switch does not allow user to configure L2 configuration on physical interface after config and unconfig of "ip directed broadcast" cli.		
Condition:	After config and unconfig of "ip directed broadcast" cli		

Defect ID:	DEFECT000637684		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Data Center Fabric
Reported In Release:	NOS5.0.2	**Technology:**	Logical Chassis
Symptom:	Unexpected reload		
Condition:	Deleting the zone member entry in upper case which is configured as upper case.		
Workaround:	Use lower case letters only to delete the zone member entry.		

Defect ID:	DEFECT000637797		
Technical Severity:	Medium	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Network Automation and Orchestration
Reported In Release:	NOS6.0.2	**Technology:**	NETCONF - Network Configuration Protocol
Symptom:	DHCP Auto Deployment upgrade from NOS 6.0.2x to NOS 7.0.1x fails.		
Condition:	DHCP Auto Deployment upgrade from NOS 6.0.2x to NOS 7.0.1x.		

Defect ID:	DEFECT000638197		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.1	**Technology:**	BGP4 - IPv4 Border Gateway Protocol
Symptom:	peer-group configuration may not exist after the firmware upgrade		
Condition:	This happens when the peer-group has only the BFD configuration		
Workaround:	Reconfigure the peer-group		

Defect ID:	DEFECT000639403

Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.2	**Technology:** Port Mirroring
Symptom: Monitor traffic may not appear on the monitor ports	
Condition Two monitor sessions are created with same destination but different source and first monitor session is removed.	

Defect ID: DEFECT000640057	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS7.2.0	**Technology:** OpenStack Integration
Symptom: VDX6940-36Q and -144S may cause FFDC (First Failure Data Capture) on 4x10g breakout ports 1:1, 17:1, and 18:1.	
Condition When VDX reloads unexpectedly, it might fail over to new active GOS (e.g., SW1) then VDX is vulnerable to this issue.	
Recovery: Use NOSCLI command "HA Failover" to manually failover (e.g., back to SW0).	

Defect ID: DEFECT000640567	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: Displaying routes with detail option asterisk are shown on all routes returned. This should only be on the selected route.	
Condition Execution of "show ip route detail" CLI	

Defect ID: DEFECT000641485	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Management cluster/VCS goes offline when ISL between two nodes goes down even though the connectivity could have been established through the other nodes' ISL.	
Condition It happens rarely when the new link/connectivity happens slowly.	

Defect ID: DEFECT000642278	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Snmpwalk of in-band Ve interfaces fails with timeout error.	

Condition	SNMP packets are discarded when the ingress interface differs from the egress interface.

Defect ID:	DEFECT000642884	
Technical Severity:	Medium	**Probability:** Low
Product:	Extreme Network OS	**Technology Group:** Monitoring
Reported In Release:	NOS7.0.0	**Technology:** Hardware Monitoring
Symptom:	The following warning will be logged on some interfaces which are installed with 'SR' SFP+ The mentioned threshold in the logs looks like a 10G LR threshold even though the installed SFP+ is 'SR 'Sfp Current for port x/0/y, is below low boundary(High=85, Low=15). Current value is Z mA' on 10G SR SFP+'	
Condition	This will occur only on interfaces where already inserted 10G 'LR' SFP+. are replaced with a 10G 'SR' SFP+ and the link is up	

Defect ID:	DEFECT000643124	
Technical Severity:	Medium	**Probability:** Medium
Product:	Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release:	NOS7.2.0	**Technology:** OpenStack Integration
Symptom:	VDX 6740 or 6940 40g port connected to SLX 9xxx 40g port, and then rebooting or power-cycling the SLX switch. Upon reboot, the 40g link may not come online	
Condition	VDX 6740 or 6940 40g port connected to SLX 9xxx 40g port, and then rebooting or power-cycling the SLX switch. Upon reboot, the 40g link may not come online	
Workaround:	Try NOSCLI shut / no shut on VDX and or SLXswitch.	
Recovery:	Try NOSCLI shut / no shut on VDX and or SLX switch.	

Defect ID:	DEFECT000643696	
Technical Severity:	Medium	**Probability:** Medium
Product:	Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.1	**Technology:** OSPFv3 - IPv6 Open Shortest Path First
Symptom:	Occasionally in a VCS consisting of two VDX running as ASBR., a few type7 LSAs are not generated on one of the RBridge after reloading VDXs at times.	
Condition	A VCS cluster with 2 VDXs and distributing 127 routes their own VE interfaces into OSPF Area 21 (NSSA).	

Defect ID:	DEFECT000644087	
Technical Severity:	Low	**Probability:** Low
Product:	Extreme Network OS	**Technology Group:** Management
Reported In Release:	NOS6.0.0	**Technology:** Software Installation & Upgrade
Symptom:	VDX6940-2U with new DRAM may encounter machine-checks errors during or after FWDL and can cause unexpected switch reload.	
Condition	VDX6940-2U with new HW component [DRAM]	

Workaround: When a newer HW component has been detected, firmware download is being blocked by pre-install script to avoid degrading the system performance of BR-VDX6940-144S platform.	
Recovery: Perform FWDL to a newer NOS version where new uboot change does exist.	

Defect ID: DEFECT000644227	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** ARP - Address Resolution Protocol
Symptom: mac learning stops after ARP limit is exceeded and then ARP entries are cleared with "clear arp"	
Condition Scaling ARP to limit	

Defect ID: DEFECT000645046	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.2.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: ifInErrors counter does not increment when CRC errors are seen on interface.	
Condition CRC error occurs on interfaces and ifInErrors counter is polled.	

Defect ID: DEFECT000645982	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** ICMP - Internet Control Message Protocol
Symptom: Packet Loss in IP Fabric topology.	
Condition ARP/IP moves from one mac to another.	

Defect ID: DEFECT000646528	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Unexpected reload	
Condition In rare scenarios, MAC address age out results in corrupt data.	

Defect ID: DEFECT000647389	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** IP Addressing

Symptom:	CLI prohibits user from adding multiple /31 subnets under L3 interfaces.
Condition	Configuring multiple /31 subnets under L3 interfaces.

Defect ID:	DEFECT000647398
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS4.1.3	**Technology:** VCS Fabric
Symptom:	Unexpected reload.
Condition	Rare scenario. During the cluster formation.

Defect ID:	DEFECT000647433
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** IP Fabric
Symptom:	L2 VNI and tunnel IP value in the BGP route update is set to "zero" .
Condition	In IP fabric topology, when a route-map with set condition is applied to evpn peer.

Defect ID:	DEFECT000648098
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VPN
Reported In Release: NOS7.1.0	**Technology:** EVPN - Ethernet VPN
Symptom:	GARP Doesn't flood to hosts to updated their ARP cache irrespective of whether ARP suppression is enabled/disabled.
Condition	Ipfabric environment where L2VPN is enabled.

Defect ID:	DEFECT000648655
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** CLI - Command Line Interface
Symptom:	Displaying generic error message.
Condition	When scp fails displaying common error message.

Defect ID:	DEFECT000648729
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.2.0	**Technology:** OSPF - IPv4 Open Shortest Path First

Symptom: OSPF vulnerabilities CVE-2017-3224, CVE-2017-3752, CVE-2017-6770	
Condition Existing code has above vulnerabilities in OSPF.	

Defect ID: DEFECT000649012	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.2.0	**Technology:** IP Fabric
Symptom: Unexpected reload	
Condition Dampening configuration under BGP	

Defect ID: DEFECT000651945	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.2	**Technology:** Hardware Monitoring
Symptom: Unexpected reload.	
Condition Rare scenario. Internal polling of memory statistics.	

Defect ID: DEFECT000651956	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.1.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: SNMP traps will not be seen.	
Condition Chassis IP and VCS IP not configured.	

Defect ID: DEFECT000652746	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.2	**Technology:** OpenStack Integration
Symptom: Mac learning won't happen for some of the ports on VDX 6740T-1G platform.	
Condition Interface configured with 100MB speed. Seen when connected to certain power-tower units via 100mb interface, or to Avaya CLAN 100mb. May occur on other non-VDX 100mb link partners as well.	
Workaround: No workaround for 100mb. May try 1gb if link partner supports it.	
Recovery: May try 1gb if link partner supports it. Recommend upgrade VDX firmware for fix.	

Defect ID: DEFECT000652894

Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Unexpected reload.	
Condition Execution of CLI(vcs replace rbridge-id) during the cluster re-join.	
Workaround: Avoid the CLI during cluster re-join	

Defect ID: DEFECT000655163	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.2	**Technology:** RADIUS
Symptom: Unable to import TLS server certificate and keys without trust point association and use these two to establish TLS connection.	
Condition This is a feature enhancement. So customer will hit this scenario every time when they try to import TLS certificate and key without trust point.	
Workaround: Use crypto functionality (which uses trust point association) to import TLS certificate and key and establish TLS connection.	
Recovery: N/A	

Closed with code changes for NOS v7.0.1c

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS v7.0.1c.

Defect ID: DEFECT000564498	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS4.1.3	**Technology:** Port Mirroring
Symptom: "show interface status" command shows incorrect status for internal VDX ports in the switch.	
Condition If we have internal embedded ports in the VDX switch, then we will observe this issue.	

Defect ID: DEFECT000592902	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: False SNMP traps observed for link status of management interface.	
Condition Very rare scenario to hit this false statustraps	
Recovery: Restart SNMP process.	

Defect ID: DEFECT000601293	
Technical Severity: Medium	**Probability:** High

Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: COS Priority tag frames egressed as Untagged frames	
Condition Over VxLAN tunnel COS Priority tag frames are egressed as untagged frames.	

Defect ID: DEFECT000603775	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom: VDX picks up global MTU value instead of local MTU after VDX reboot, when global MTU is configured and some interfaces has default/local MTU configured,	
Condition VDX reload can hit the issue when global MTU is configured and some interfaces has default/local MTU configured ,	

Defect ID: DEFECT000611018	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.1.0	**Technology:** Logical Chassis
Symptom: VDX is not able to join cluster and stuck at "Awaiting Rejoin" state after reload.	
Condition VDX reload can hit the issue in VCS cluster.	

Defect ID: DEFECT000611303	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.1	**Technology:** AMPP - Automatic Migration of Port Profiles
Symptom: Unexpected reload.	
Condition After configuring vCenter and enabling CDP on ESXi vSwitch, due to very mild memory leak.	

Defect ID: DEFECT000612699	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** LAG - Link Aggregation Group
Symptom: Unexpected reload of VDX	
Condition In rare cases, deleting PO or reload of LC may cause the VDX to for an abrupt reload due to software daemon termination.	

Defect ID: DEFECT000612967	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.1.0	**Technology:** Security Vulnerability
Symptom: Shutting down SSH server does not close all existing SSH login sessions	
Condition Shutdown SSH server	
Recovery: Close all existing login sessions using "clear sessions" command, please note this command will close telnet sessions as well.	

Defect ID: DEFECT000613368	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** CLI - Command Line Interface
Symptom: No Single CLI to clear the system wide counters.Added the same as enhancement.	
Condition Single CLI to clear system wide counter.	

Defect ID: DEFECT000614007	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** IP Fabric
Symptom: VE ip address is advertised as /32 route through EVPN to remote leaf.	
Condition When VE interface is in shutdown state, VE ip address is advertised as /32 route through EVPN to remote leaf.	

Defect ID: DEFECT000617887	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: On upgrade from 6.x to 7.x, one of the python CLI libraries may not be carried forward with all the changes & might impact some of the python scripts.	
Condition Under certain unknown condition, when upgrading from 6.x to 7.x.	
Recovery: Copy CLI.py file manually to restore the script function. After restoration, test using below CLI & it should appear as below with "splitlines()" instead of "split()": sw0:FID128:root> grep -A 2 get_output /etc/fabos/Dcmd/python/CLI.py def get_output(self): return (self.output.splitlines())	

Defect ID: DEFECT000618373	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.1.0	**Technology:** CLI - Command Line Interface
Symptom: MTU values get changed to global MTU value.	
Condition User configured default MTU values are not shown in running config. Due to this copy running to any file is not storing MTU values of such interfaces. After downloading the save configuration back to VDX, VDX MTU values get changed to global MTU value.	

Defect ID: DEFECT000625751	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Rolling reboot of VDX after firmware download.	
Condition Firmware download in cluster. Odd nodes has upgraded from NOS 4.1.3x to NOS 5.x and then to NOS 6.x while Even nodes were still in NOS 4.1.3x.	
Workaround: Perform firmware download from NOS 4.1.3x to NOS 5.x for Odd nodes and then Even nodes. Once first phase is done please upgrade from NOS 5.x to NOS 6.x for Odd nodes and then Even nodes.	

Defect ID: DEFECT000626037	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.1.0	**Technology:** VLAN - Virtual LAN
Symptom: VLAN creation fails along with "Error: VLAN creation failed due to lack of sufficient resources" error.	
Condition When GVLAN is configured along with "switchport trunk allowed vlan all" on more that 64 interface.	

Workaround:	Disabling GVLAN or removing "switchport trunk allowed vlan all" from some interface to make it less than 64 interface.

Defect ID:	DEFECT000629138
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** LLDP - Link Layer Discovery Protocol
Symptom:	VDX experience unexpected reload due to onmd process termination.
Condition:	VDX can experience unexpected reload due to onmd process termination when it is connected with ESX servers / VMware VMs and any lldp operation occur.

Defect ID:	DEFECT000630676
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** Software Installation & Upgrade
Symptom:	NOS BNA 14.0.1 and 14.0.2 throws the error "Firmware image download reboot operation has timed out", even the FW downgrade was successful.
Condition:	Firmware Download through BNA on VDX running in FC Cluster mode.

Defect ID:	DEFECT000631440
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** ARP - Address Resolution Protocol
Symptom:	ARP is not learnt from the incoming packet on the source interface when /31 addressing is used.
Condition:	ARP learning when /31 addressing is used

Defect ID:	DEFECT000631591
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.1	**Technology:** AAA - Authentication, Authorization, and Accounting
Symptom:	Unable to log into VDX after password change
Condition:	Some alpha numeric passwords having more than 24 character length will not allow the user to log on with the new password.
Workaround:	Create passwords with length less than 24 character.

Defect ID:	DEFECT000632115
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.2.0	**Technology:** Logical Chassis
Symptom:	Node cluster rejoin operation fails with VCS-1006 error. VCS-1006, ERROR, , Event: VCS node rejoin, Coordinator IP: <IP>, VCS ID: <id>, Status: rBridge ID <ID> failed to rejoin VCS cluster, Reason: Remote Location is not available.
Condition:	Cluster rejoin operation.

Defect ID:	DEFECT000632419
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** xSTP - Spanning Tree Protocols
Symptom:	MSTP 'FORWARD-DELAY' lost configured non-default value and become default value after VDX reload.
Condition:	VDX reload.

Defect ID: DEFECT000633384	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: Unexpected reload due to OSPF daemon termination.	
Condition When same external LSA is received from multiple ASBRs	

Defect ID: DEFECT000633831	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: When a VCS cluster reformation happens, existing FCOE hosts gets logged out.	
Condition Adding/removing an ISL or adding/removing a switch from cluster that results in the fabric reformation.	
Recovery: Execute the CLIs "shutdown" and "no shutdown" on the respective interfaces from which FCOE hosts logged out.	

Defect ID: DEFECT000634094	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: short traffic outage while shut/no shut best BGP router.	
Condition Having two adjacent BGP evpn router and shut/no shut one of the best BGP router.	

Defect ID: DEFECT000634129	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: Route filtering using distribution list will not happen after HA failover.	
Condition If distance for inter area route is not configured to non-default value and HA failover occurs.	

Defect ID: DEFECT000634192	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: Tunnel terminated packets that generate the ARP's may keep looping across VCS nodes in the same VCS	
Condition VDX puts incorrect source information cause the issue	

Defect ID: DEFECT000634366	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.1.0	**Technology:** CLI - Command Line Interface
Symptom: There is no option to view the ACL log buffers for VE interfaces.	
Condition 'show access-list-log buffer interface' does not have option for VE interfaces.	
Workaround: Can check the ACL logging packet in phy interface as well	

Defect ID: DEFECT000634370	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.1.0	**Technology:** Configuration Fundamentals

Symptom:	show running-config rbr route-map output is being sorted based on the action (permit or deny) and then the sequence number instead of just sequence number.
Condition:	Execution of "show running-config rbr route-map" CLI.

Defect ID:	DEFECT000634372		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS7.1.0	**Technology:**	CLI - Command Line Interface
Symptom:	route map is allowing to change the existing action without any warning/error.		
Condition:	Modifying route map action.		

Defect ID:	DEFECT000634673		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	VPN
Reported In Release:	NOS7.0.1	**Technology:**	EVPN - Ethernet VPN
Symptom:	Frames on udp port 4789 are wrongly treated as vxlan packets and not subjected to tunnel source suppression in transit nodes		
Condition:	Using of UDP port 4789		

Defect ID:	DEFECT000634766		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS7.0.0	**Technology:**	Software Installation & Upgrade
Symptom:	"no system-mode maintenance" CLI fails with operation failed.		
Condition:	"no system-mode maintenance" CLI execution.		

Defect ID:	DEFECT000635078		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS7.0.1	**Technology:**	CLI - Command Line Interface
Symptom:	Rbridge-range command is not functioning properly.		
Condition:	Execution of Rbridge-range command		

Defect ID:	DEFECT000635101		
Technical Severity:	Low	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	IP Multicast
Reported In Release:	NOS6.0.1	**Technology:**	IGMP - Internet Group Management Protocol
Symptom:	Memory leak in igmpd.		
Condition:	When debug igmpd command is enabled and leads to error condition "Illegal multicast group address".		

Defect ID:	DEFECT000635328		
Technical Severity:	Low	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS7.1.0	**Technology:**	SNMP - Simple Network Management Protocol
Symptom:	Cannot find module message appears during installing and compiling FOUNDRY-SN-NOTIFICATION.mib file.		
Condition:	Installing and compiling FOUNDRY-SN-NOTIFICATION.mibfile		

Defect ID:	DEFECT000635411

Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom: Routing Administrative distance not honored on 6940 border leaf using EVPN and BGP	
Condition Routing Protocol running on 6940 border leaf using EVPN and BGP.	

Defect ID: DEFECT000635440	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** IP Addressing
Symptom: "show ip route <prefix>/<mask> longer" is not showing the expected results.	
Condition When we have more than 500 entries	

Defect ID: DEFECT000635844	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS4.1.3	**Technology:** VLAN - Virtual LAN
Symptom: Newly added VLAN is showing up in RPVST disabled/discarding state on vLAG members.	
Condition When STP is enabled on a PO interfaces, it is enabled only for vports and not for main physical/po interfaces for PVST/RPVST. In this scenario whenever a RBridge leaves cluster, STP module runs State Machine to re compute the topology and it can hit the issue.	

Defect ID: DEFECT000636084	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: "unqualified SFP transceiver" error appears on VDX for Extreme optic part number 33210-100.	
Condition Inserting Extreme optic with part number 33210-100 on VDX running can hit the error.	

Defect ID: DEFECT000636497	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: VDX experience unexpected reload due to DCMd daemon termination.	
Condition No system-mode maintenance activity can cause the issue.	

Defect ID: DEFECT000636649	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** LAG - Link Aggregation Group
Symptom: Unable to reach to end host when VDX is routing packets.	
Condition VDX put Incorrect Router Source MAC addresses.	
Workaround: Reload the switch	

Defect ID: DEFECT000637538	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** IP Fabric
Symptom: In a IP fabric topology, the leaf (or border leaf) does not forward L3 packets (e.g. ICMP) destined to hosts on to peer leaf nodes. This causes the end to end L3 traffic between leaf nodes to fail.	

Condition	The prefix route learnt over evpn-bgp was incorrectly programed in Linux kernel, thereby forwarding traffic to wrong destination.
Recovery:	Issuing following command at NOS should help to recover from the problem state. clear arp no-refresh vrf <vrf-name>

Defect ID:	DEFECT000637753		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS7.1.0	**Technology:**	Configuration Fundamentals
Symptom:	VDX returns "<html><body> 401 authentication needed </body></html>" response.		
Condition	Proxied REST configuration request using persistent connection Authentication-Token.		

Defect ID:	DEFECT000637857		
Technical Severity:	High	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Layer 2 Switching
Reported In Release:	NOS6.0.2	**Technology:**	VLAN - Virtual LAN
Symptom:	VDX do not learn MAC addresses.		
Condition	VDX stops learning MAC addresses when specific configuration exist and it has reached to certain number of lines in configuration.		
Workaround:	1. configure "mac-address-table consistency-check suppress" 2. reload the affected switches or 1. configure dummy vlan 2. Assign it to interface/Po 3. reload		
Recovery:	The same above steps works we may expect unexpected reloads during this		

Defect ID:	DEFECT000638990		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	VCS
Reported In Release:	NOS7.0.1	**Technology:**	Logical Chassis
Symptom:	VDX is not able to join cluster and stuck at "Awaiting Rejoin" state after reload.		
Condition	VDX has "switchport port-security" configuration and reload occur in VCS cluster.		

Defect ID:	DEFECT000639081		
Technical Severity:	Medium	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS6.0.2	**Technology:**	Software Installation & Upgrade
Symptom:	vLAG/Po interface flaps for several seconds during firmware upgrade.		
Condition	vLag/PO has an inactive link while upgrade.		

Defect ID:	DEFECT000640567		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Data Center Fabric
Reported In Release:	NOS7.0.1	**Technology:**	IP Fabric
Symptom:	Displaying routes with detail option asterisk are shown on all routes returned. This should only be on the selected route.		
Condition	Execution of "show ip route detail" CLI		

Defect ID:	DEFECT000642711		
Technical Severity:	Critical	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer

Reported In Release: NOS5.0.2	**Technology:** BFD - BiDirectional Forwarding Detection
Symptom: BGP & BFD is not working as expected.	
Condition: Bringing down the loopback interface and VDX has "bfd interval 300 min-rx 300 multiplier 3" configuration under interface.	

Defect ID: DEFECT000642983	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** IP Fabric
Symptom: BFD goes down and BGP does not goes down during unnumbered interface shutdown.	
Condition: Unnumbered interface shutdown, when ECMP paths exist.	

Defect ID: DEFECT000643646	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** IP Fabric
Symptom: BFD (Bi-directional forwarding detection) state may remain in INIT state on the leaf switch after shut/no shut of L3 port channel to the spine switch.	
Condition: BFD session remains in INIT state which makes longer time to converge for protocols registered with BFD.	

Defect ID: DEFECT000643924	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: Reachability from border leaf switches to multi homed hosts behind leaf switches will be affected under the following condition. - leaf switches have a static route configured to multi homed hosts network and next hop as one of its connected networks.	
Condition: Connectivity between border leaf switches to multi homed hosts behind leaf switches will be affected.	
Workaround: Deleting and adding static route programmed in leaf switches should be done per rbridge basis.	

Closed with code changes for NOS v7.0.1b

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as January 30, 2017 in Network OS v7.0.1b.

Defect ID: DEFECT000550177	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.1	**Technology:** Metro VCS
Symptom: A host will experience lost connectivity when a separate host in a separate VLAN uses the same IP address and sends a gratuitous ARP from that IP address. The VDX will learn from that gratuitous ARP an ARP entry for the IP address on a ve interface that doesn't have an address in the same subnet.	
Condition This issue seems to happen when the wrong host in the wrong VLAN is trying to use an IP address already in use by another host in the correct VLAN for the IP address of the subnet.	

Defect ID: DEFECT000555460	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** IP Addressing
Symptom: 'ICMP unreachables are always sent' displayed in the configuration even when disabled in the configuration	
Condition: Default ICMP unreachable is not set	

Defect ID: DEFECT000562722	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS5.0.2	**Technology:** Rate Limiting and Shaping
Symptom: ipv6 icmpv6 rate-limiting does not work per interface	
Condition: The above situation occurs under two conditions 1. More than one ipv6 interface 2. Different rate-limiting value configured When both of the above conditions met, then the recently configured rate-limiting value applied to all interfaces	

Defect ID: DEFECT000583274	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.1	**Technology:** Logical Chassis
Symptom: Port-group configuration fails	
Condition: Inserting the new Linecard OR power off/on with out removing the existing port-group configuration.	

Defect ID: DEFECT000587135	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.2	**Technology:** TRILL - Transparent Interconnection of Lots of Links
Symptom: Unexpected reload on standby MM in rare scenario.	
Condition: While changing VCS.	

Defect ID: DEFECT000593537	
Technical Severity: High	**Probability:** Low

Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** IP Addressing
Symptom: Host ARP is learnt even when host IP subnet does not match to VE IP subnet.	
Condition: Host is connected to a VLAN where the Ve IP subnet is different than the host IP subnet.	
Workaround: Disable proxy ARP on VE	

Defect ID: DEFECT000596720	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** CLI - Command Line Interface
Symptom: When IPv6 nd prefix is configured with a prefix flag(no-autoconfig/no-onlink/offlink) enabled and if the same prefix is updated later with different lifetime values, then the already configured prefix flag will not be present in the running configuration of that prefix.	
Condition: This issue happens when an IPv6 prefix configuration is updated with lifetime values provided a prefix flag(no-autoconfig/no-onlink/offlink) was already configured.	
Workaround: NA	

Defect ID: DEFECT000597104	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: Under rare scenarios of leaking routes between VRF's, the switch may get reloaded due to "termination of process ribmgr"	
Condition: When leaking routes from one VRF to another & presence of those same routes in target VRF as connected routes.	
Workaround: Reconfigure to avoid leaking routes between VRF's OR ensure that the leaked routes are not present in target VRF as local routes.	

Defect ID: DEFECT000598878	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Configuration Fundamentals
Symptom: A stale default-route gets applied in the running configuration of the secondary nodes in cluster environment during configuration replay.	
Condition: The issue arises when secondary nodes disconnect and re-join the cluster provided DHCP is enabled.	

Defect ID: DEFECT000599306	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: Vrf information is missing for some interfaces while displaying output of "show ip interface brief" command.	
Condition: This issue is seen, then "show ip interface brief" is executed repeatedly in multiple terminals.	
Workaround: If "show ip interface brief" executed from multiple terminals, then it not should be executed too quickly. Let the command output display completed on one terminal before starting on other terminal.	

Defect ID: DEFECT000602062	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** Access Gateway
Symptom: Console logs appear when snmpwalk is performed.	
Condition: When snmpwalk is performed for community/user associated with IPv6 ACL.	

Defect ID: DEFECT000602861	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.1	**Technology:** Logical Chassis
Symptom: High disk usage that ended up out of space.	
Condition Postgres log file(Dcmd.Linux.powerpc.pg_ctl.log) unconditionally growing	
Recovery: Delete Dcmd.Linux.powerpc.pg_ctl.log file.	

Defect ID: DEFECT000604049	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** VCS Fabric
Symptom: Potential for Name Server fail-over and recovery to the Standby Control Processor if the overall scale of the VCS cluster exceeds the limit described within the "Conditions for Publication" section.	
Condition The maximum number of elements within a cluster cannot exceed 32,767 prior to having this modification to increase scale to 80,000. Entities that contribute to this count are: - RBridges - Ports (physical and virtual) - Devices that appear in the Name Server The maximum assignable port indexes are listed here by platformtype: Chassis-based systems (Director class) 1800 VDX 6740/VDX 6740T/VDX 6740T-1G : 1200 VDX 6940 : 1312 For example, one Director-class RBridge accounts for 1 (for the RBridge itself) + 1800 (maximum assignable port indexes) + <FC/FCoE device count>. Thus, if we have 500 devices, this would translate to 1 + 1800 + 500 = 2301 (of the total allowable 32767). Here are some sample combinations in terms of RBridge composition within a cluster, where a cluster-wide FC/FCoE device count is presumed to be 3000: • 16 Directors • 14 Directors + 3 VDX 6940/ 3 VDX 6740 • 12 Directors + 5 VDX 6940 / 6 VDX 6740 • 8 Directors + 11 VDX 6940	
Workaround: Limit cluster composition in a manner compliant with the maximum values described within the "Conditions for Publication" section.	
Recovery: Limit cluster composition in a manner compliant with the maximum values described within the "Conditions for Publication" section.	

Defect ID: DEFECT000604131	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** Multi-VRF
Symptom: If local route exists from a route source and a leaked route is added from the same route source for the prefix, the routing table is updated with the new leaked route.	
Condition Issue is seen if dynamic route leak is configured with prefixes matching the local prefixes.	
Workaround: There should not be overlap between local and leaked prefixes	

Defect ID: DEFECT000604338

Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** xSTP - Spanning Tree Protocols
Symptom: After the VDX switch was reloaded, the configured MSTP hello time was ignored and uses default value of 2 seconds. The value configured in the system does not change, NOS simply ignores it and uses default value.	
Condition The configured the MSTP hello time is not persistent after VDX switch was reloaded.	
Recovery: The configured MSTP hello-time is persistent after the VDX switch was reloaded. That is if MSTP hello-time was configured for example set to 5. The value was applied after the VDX was reloaded.	

Defect ID: DEFECT000605042	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: 'snmp-server' command doesn't update the values of 3 input parameters.	
Condition Inputting all the 3 parameters contact, location and sys-descr on a single line of execution.	
Workaround: Configure each of the input parameter separately.	

Defect ID: DEFECT000605923	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: FCoE VLAN creation and subsequent fabric map may fail. Also customer can experience "%% Error: VLAN creation failed due to lack of sufficient resources" error while creating VLAN.	
Condition When more than 64 ports are configured with 'switchport trunk allowed vlan all' configuration and tried to create VLAN or FCoE VLAN.	
Workaround: Do not configure more than 64 ports with 'swtchport trunk allowed vlan all' configuration.	
Recovery: Remove 'switchport trunk allowed vlan all' configuration if it is configured on more than 64 interfaces and try creating VLAN or FCoE VLAN and fabric map.	

Defect ID: DEFECT000605998	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: just once “clear ip ospf all” is done in RB01, a tremendous “LSA flush rcvd Type:5” message lasted to pop up forever and network got unstable.	
Condition In huge scale OSPF setups when there are more than 10 neighbors and OSPF peer has to retransmitan LSA to all these neighbors this issue is seen as each neighbor is added to retransmit queue multiple times	
Workaround: Decrease the number of LSAs and neighbors	

Defect ID: DEFECT000610510	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: OSPF routes are uninstalled from one or more VRF's, causing traffic disruption. Router LSA's do not refresh.	
Condition Occurs when there are many OSPF session across many VRF's, with total OSPF routes exceeding 1500.	
Recovery: Flap OSPF neighbor sessions.	

Defect ID: DEFECT000610816

Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** LAG - Link Aggregation Group
Symptom: VDX throws FVCS-1005 RASLOG message followed by an unexpected reboot.	
Condition The user may experience this issue when attempting to change or undo the active Port Channel in a Redundancy Group using the 'no port-channel <PortChannel ID> active' command.	
Workaround: When changing the active Port Channel in a Redundancy Group, it is best to avoid using the 'no port-channel <PortChannel ID> active' command. It is advisable to delete the Redundancy Group and recreate it when wanting to change the Active Port Channel in a Redundancy Group.	

Defect ID: DEFECT000610873	
Technical Severity: Low	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** Static Routing (IPv4)
Symptom: The secondary node running configuration has default route applied to the default-vrf when joining the cluster.	
Condition This issue occurs whenever the user tries to add a node to the cluster.	

Defect ID: DEFECT000611149	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VPN
Reported In Release: NOS7.1.0	**Technology:** EVPN - Ethernet VPN
Symptom: Tunnel extension is not getting removed.	
Condition L2VNI IMR route removed	

Defect ID: DEFECT000611400	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS7.1.0	**Technology:** Rate Limiting and Shaping
Symptom: Switch can go for a reboot when the slot values are provided well outside the permissible range in the 'bp-rate-limit command'. Permissible range for slot is '0-16'	
Condition The issue is seen only when the command is executed by providing the slot values well outside the permissible range.	
Workaround: Ensure that slot values are provided only in the valid range '0-16'	
Recovery: Remove any of the slot values provided outside the permissible range of '0-16'	

Defect ID: DEFECT000611688	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.1.0	**Technology:** Hardware Monitoring
Symptom: VDX 6940 and 6940-144S may show CRC errors on ports cabled with QSFP 40 GbE active copper transceiver cables.	
Condition VDX 6940 and 6940-144S may show CRC errors on ports cabled with QSFP 40 GbE active copper transceiver cables.	
Workaround: Replace QSFP 40 GbE active copper transceiver cables with optical QSFP transceivers. Then reboot the switch.	
Recovery: Replace QSFP 40 GbE active copper transceiver cables with optical QSFP transceivers. Then reboot the switch.	

Defect ID: DEFECT000612821	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer

Reported In Release: NOS6.0.2	**Technology:** VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: VRRP-1002 raslog message is not displayed.	
Condition When Master to backup change happens .	

Defect ID: DEFECT000613594	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.2	**Technology:** Logical Chassis
Symptom: show commands couldn't be accepted due to "application communication failure".	
Condition Deletion of snmp-community config after ISSU upgrade from NOS 502a to NOS 502b1, can cause the issue of show command.	
Workaround: Remove the snmp community config before the upgrades and apply it back	

Defect ID: DEFECT000613895	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS7.1.0	**Technology:** Rate Limiting and Shaping
Symptom: Pizza box should accept only module id "0" in the "bp-rate-limit heavy module add ..." config. But the pizza box is accepting non-existent module id.	
Condition The configuration command "bp-rate-limit heavy module add ..." can accept module id other than "0" on the pizza box.	
Recovery: The configuration command "bp-rate-limit heavy module add ..." can accept module id "0" on the pizza box.	

Defect ID: DEFECT000614000	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS4.1.3	**Technology:** Logical Chassis
Symptom: After invoke "copy support" CLI, the CLI will block, not return to user prompt for more than 2 hours, then the switch will go reboot.	
Condition The failure shows on early NOS release, and is very rare to happen. So far it never failure the same in the field.r	

Defect ID: DEFECT000614390	
Technical Severity: Critical	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** ICMP - Internet Control Message Protocol
Symptom: Very rarely we could see 5% of the ICMP replies are dropped in software and random interval.	
Condition The issue can be happened when we have ARP requests from 1000 different hosts at the rate of 25 ARP's/sec, and at the same time pinging VE or VRRP IP on the same SVI at 1 ICMP/sec	

Defect ID: DEFECT000614988	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** IPv6 Addressing
Symptom: "ipv6 nd prefix" CLI command displays incorrect default value for lifetime and preferred lifetime parameter.	
Condition Execution of "ipv6 nd prefix" CLI.	

Defect ID: DEFECT000615165	
Technical Severity: High	**Probability:** Medium

Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** IPv6 Addressing
Symptom: "ipv6 nd prefix <IPv6> no-autoconfig" config can get lost.	
Condition Config-replay from backup configuration file when "ipv6 nd prefix <IPv6> no-autoconfig" is configured with valid and preferred life time default values.	

Defect ID: DEFECT000615176	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** CLI - Command Line Interface
Symptom: CLI command "show support" will not show few core files from system daemon crash on usual place which is /core_files	
Condition When there is crash by one of management daemon then core file doesn't get saved on regular system path /core_files	

Defect ID: DEFECT000615242	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** AMPP - Automatic Migration of Port Profiles
Symptom: MACs on Linux Virtual Machines with VMWare Tools installed may not get programmed on VDX.	
Condition When VMWare Tools are installed on Virtual Machines, Both IPV4 and IPV6 address gets reported from Vmware to VDX. VDX is unable to handle very long IP Strings and ignores such vnics(MACs)	
Workaround: Either disable IPV6 on the Virtual Machines or don't install VMware tools on the Virtual Machines	
Recovery: Disable IPV6 on Virtual Machines or remove VMware tools and re-run the discovery cycle	

Defect ID: DEFECT000615380	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP packets will be dropped in the box where DHCP Relay is configured.	
Condition DHCP Relay listens on standard well-known BOOTPS and BOOTPC ports (i.e. 67 and 68). If any other ports are used for communication between DHCP Client and DHCP Server can cause the issue.	
Workaround: As a workaround, use standard BOOTPS and BOOTPC (i.e. 67 and 68) UDP ports for communication between DHCP Relay and DHCP Server.	
Recovery: Use of standard BOOTPS and BOOTPC (i.e. 67 and 68) UDP ports for communication between DHCP Relay and DHCP Server will recover the system.	

Defect ID: DEFECT000615564	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** LAG - Link Aggregation Group
Symptom: If a port channel interface is configured as tracking interface for an interface which exists before this port channel interface in output of "show running-config" then during replay of this configuration file will cause the issue. It throws the error that it can not find particular port channel interface.	
Condition This issue can occur during configuration file replay in which a port channel can be configured as tracking interface.	

Defect ID: DEFECT000615646	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer

Reported In Release: NOS7.0.1	**Technology:** IPv6 Addressing
Symptom: Prefix is advertised in the IPv6 RA messages even though it is configured with "no-advertise" option.	
Condition Prefix is configured using "ipv6 nd prefix" with "no-advertise"	
Workaround: Do not configure prefix if it should not be present in IPv6 RA messages.	

Defect ID: DEFECT000615651	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** IPv6 Addressing
Symptom: ipv6 nd prefix <prefix> with "off-link" option does not work.	
Condition execution of ipv6 nd prefix <prefix> CLI with "off-link" option	
Workaround: NA	

Defect ID: DEFECT000616334	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** IP Fabric
Symptom: L3 traffics are not forwarded correctly.	
Condition The environment have lots of flows which generate more than 3K hash results and some hash values are shared by 2 or more flows.	
Workaround: Reduce the total flows or consider re-arrange the private subnet prefix if there are private subnet.	
Recovery: Clear the host table.	

Defect ID: DEFECT000616998	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** BGP4+ - IPv6 Border Gateway Protocol
Symptom: ipv6 routes are not installed in RIB when next hop is link local address.	
Condition ipv6 bgp route	

Defect ID: DEFECT000617049	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** VLAN - Virtual LAN
Symptom: Static-MACfor multicast-mac address floods the packet after removing static-ARP and static-MAC entry and re-configuring.	
Condition Static multicast MAC configured as static ARP.	

Defect ID: DEFECT000617399	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** BGP4+ - IPv6 Border Gateway Protocol
Symptom: When VDX receives a BGP update message with duplicate path attribute, It does not send an error message back to neighbor about malformed packet.	
Condition Handling of malformed BGP packets received by VDX.	

Defect ID: DEFECT000617886	
Technical Severity: Critical	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** VLAN - Virtual LAN
Symptom: VDX experience unexpected reload due to Out-Of-Memory condition. Also some of the ports are unable to transmit.	

Condition	Known to happen with 10G ports that have copper-pigtail connector. And the link-partner is not a Extreme device.

Defect ID:	DEFECT000618254
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.1	**Technology:** Logical Chassis
Symptom:	Unable to use REST API to configure prefix-list out for router bgp.
Condition	REST API to configure prefix-list out for router bgp.

Defect ID:	DEFECT000618268
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** High Availability
Symptom:	HA Sync failure after ISSU upgrade
Condition	With 6X100G LC in chassis during ISSU upgrade

Defect ID:	DEFECT000618317
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.1.0	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom:	Termination of raslogd process after upgrading from 7.0.1 to 7.1.0
Condition	In cluster environment after updating firmware.
Recovery:	Raslogd will restart automatically.

Defect ID:	DEFECT000618691
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom:	Customer will notice that the direct unicast DHCP packets between Client and Server are also getting trapped.
Condition	When number of DHCP packets getting exchanged between Server and Client are huge (say 1000 pps rate), other protocols like OSPF will have impact.

Defect ID:	DEFECT000618713
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS7.1.0	**Technology:** OpenStack Integration
Symptom:	The VDX6940-144S 10G passive cable (1m and 3m) interfaces do not display the interface "link down" RASLOG message when the corresponding 10G interface on the remote end is shut down
Condition	Shutting down 10G interfaces when remote switch is a VDX6940-144S connected with 10G passive cables (1m and 3m)
Workaround:	Shut the 10G interface on the local interface
Recovery:	Shut the 10G interface on the local interface

Defect ID:	DEFECT000619405
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration

Reported In Release: NOS6.0.2	**Technology:** OpenStack Integration
Symptom: CRC errors when using 40g DAC (direct attach copper) cable with VDX6940	
Condition 40g DAC (direct attach copper) cable with VDX6940	

Defect ID: DEFECT000619719	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.1	**Technology:** SSH - Secure Shell
Symptom: Telnet/ssh for default-vrf enables though user configured as disabled.	
Condition If node disconnected and re-joined to the fabric after "no telnet server use-vrf default-vrf" OR "no ssh server use-vrf default-vrf"	
Workaround: Disable Telnet/ssh using "telnet server use-vrf default-vrf shutdown" or "ssh server use-vrf default-vrf shutdown".	
Recovery: After node rejoins the fabric, to disable the telnet/ssh, execute the CLIs "telnet server use-vrf default-vrf shutdown" or "no telnet server use-vrf default-vrf" for telnet and "ssh server use-vrf default-vrf shutdown" or "no ssh server use-vrf default-vrf".	

Defect ID: DEFECT000619807	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VPN
Reported In Release: NOS7.0.1	**Technology:** EVPN - Ethernet VPN
Symptom: Tunnel membership vlans will get deleted.	
Condition HA failover.	
Workaround: Avoid HA failover.	

Defect ID: DEFECT000620197	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: Configuration of OSPF authentication key is not applied when done using config-replay.	
Condition The issue is observed for below sequence of steps: 1. Configure OSPF authentication key on interface using CLI. 2. Save running configuration using command: copy running-config flash://<file-name> 3. Remove configured OSPF authentication key using CLI. 4. Replay saved configuration by using command: copy flash://<file-name> running-config It is observed that OSPF authentication key is not applied after step-4 though it was expected to be applied on the interface.	
Workaround: After config-replay fails to configure OSPF authentication key on the interface, it is possible to configure authentication key using CLI.	

Defect ID: DEFECT000620617	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.1	**Technology:** xSTP - Spanning Tree Protocols
Symptom: VDX6940 device may see traffic loss if HA failover or ISSU operation is performed from nos7.0.1 to nos7.0.1a release.	
Condition 1) RSTP is configured 2) HA failover or ISSU is performed	
Recovery: Disable/enable spanning-tree protocol on the interface	

Defect ID: DEFECT000621402	
Technical Severity: High	**Probability:** High

Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Inband Management
Symptom: Telnet access to VDX is blocked via default-VRF and user defined VRF.	
Condition Firmware install with "no-activate" option.	
Workaround: Activate the firmware which was installed with "no-activate" option.	

Defect ID: DEFECT000621408	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.1	**Technology:** Security Vulnerability
Symptom: Though telnet service on MGMT-VRF is shutdown, telnet access to VDX is allowed via MGMT-VRF.	
Condition 1. Shutdown telnet service on MGMT-VRF 2. Firmware install with "no-activate" option 3. Recover the firmware using "firmware recover"	
Recovery: Activate the partially installed firmware.	

Defect ID: DEFECT000622750	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** IPv6 Addressing
Symptom: When the user updates an IPv6 prefix with preferred lifetime alone, valid lifetime changes to default value.	
Condition The issue happens only when the user updates the preferred lifetime value to an already configured IPv6 prefix with valid and preferred lifetime.	

Defect ID: DEFECT000623309	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.1	**Technology:** OpenStack Integration
Symptom: CRCs occur on VDX 6940-36Q when DAC (direct attached copper) cable is used with DELL NIC server, and DELL NIC server is running traffic towards VDX 6940-36Q while the VDX switch is booting up.	
Condition CRCs occur on VDX 6940-36Q when DAC (direct attached copper) cable is used with DELL NIC server, and DELL NIC server is running traffic towards VDX 6940-36Q while the VDX switch is booting up.	
Workaround: Reboot VDX switch, with DAC cabled ports administratively down state (save port configuration as "no shut", then reboot); or stop all traffic coming into DAC cabled ports from DELL NIC servers.	
Recovery: Reboot VDX switch, with DAC cabled ports administratively down state (save port configuration as "no shut", then reboot); or stop all traffic coming into DAC cabled ports from DELL NIC servers.	

Defect ID: DEFECT000623618	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.0	**Technology:** IP Fabric
Symptom: Host ARP is learnt even when host IP subnet does not match to VE IP subnet.	
Condition Host is connected to a VLAN where the Ve IP subnet is different than the host IP subnet.	
Workaround: Disable proxy ARP on VE	

Defect ID: DEFECT000623711	
Technical Severity: High	**Probability:** High

Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.1.0	**Technology:** VCS Fabric
Symptom: Any packet transmitted from CPU gets dropped on FCport.	
Condition Can happen only on FC port.	

Defect ID: DEFECT000624388	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.1	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom: This is the debug enhancement to capture the VxLAN packet at RTE tool	
Condition Debugging tool enhancement	

Defect ID: DEFECT000624394	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.1.0	**Technology:** VLAN - Virtual LAN
Symptom: Continuous ASIC errors causes chip fault.	
Condition Heavy ASIC activity can cause the issue.	

Defect ID: DEFECT000624621	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.0.1	**Technology:** PIM - Protocol-Independent Multicast
Symptom: Multicast functionality daemon "PIMd" goes down with memory leak.	
Condition On enabling PIM, for every 60 seconds there is a memory leak of 5K bytes.The leak will be observed even if PIM is not enabled on router, but standby becomes as active node.	
Workaround: Do not enable PIM on router. Do not trigger HA failover	
Recovery: Disable PIM on router and reboot the router. Do not enable PIM after reboot. Do not trigger HA failover after reboot.	

Defect ID: DEFECT000624701	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.0	**Technology:** Security Vulnerability
Symptom: NOS/SLX kernel (NOS/Host/TPVM) are all vulnerable. User can overwrite the etc/password with root access.	
Condition CVE-2016-5195 - kernel > 2.6.22 can hit this Dirty COW issue.	

Defect ID: DEFECT000624805	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** IP Addressing
Symptom: Show command is not showing "ip icmp unreachable" under physical interface.	
Condition After configuring the "ip icmp unreachable" under physical interface.	
Workaround: This is a cosmetic issue and can be ignored.	

Defect ID: DEFECT000625243	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	**Technology:** IP Addressing

Symptom:	"show ip in ve <>" does not show "ip icmp address mask" enabled/disabled status.
Condition	Execution of "show ip in ve <>" CLI.

Defect ID:	DEFECT000625527
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.1.0	**Technology:** PIM - Protocol-Independent Multicast
Symptom:	Multicast functionality daemon "PIMd" goes down with memory leak.
Condition	On enabling PIM, for every 60 seconds there is a memory leak of 5K bytes.
Workaround:	Do not enable PIM on router.
Recovery:	Disable PIM on router and reboot the router. Do not enable PIM after reboot.

Defect ID:	DEFECT000625982
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.0.1	**Technology:** IPv4 Multicast Routing
Symptom:	Multicast functionality daemon "PIMd" goes down with memory leak.
Condition	On enabling PIM, for every 60 seconds there is a memory leak of 5K bytes.The leak will be observed even if PIM is not enabled on router, but standby becomes as active node.
Workaround:	Do not enable PIM on router. Do not trigger HA failover
Recovery:	Disable PIM on router and reboot the router. Do not enable PIM after reboot. Do not trigger HA failover after reboot.

Defect ID:	DEFECT000626555
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.1.0	**Technology:** PIM - Protocol-Independent Multicast
Symptom:	Multicast functionality daemon "PIMd" goes down with memory leak
Condition	PIM enable configuration.

Defect ID:	DEFECT000626886
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom:	When a VCS cluster reformation happens, existing FCOE hosts gets logged out.
Condition	Adding/removing an ISL or adding/removing a switch from cluster that results in the fabric reformation.
Recovery:	Execute the CLIs "shutdown" and "no shutdown" on the respective interfaces from which FCOE hosts logged out.

Defect ID:	DEFECT000628198
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.1.0	**Technology:** Configuration Fundamentals
Symptom:	Firmware install ftp using prompted path comes up in Fabric Cluster.
Condition	Firmware install execution

Defect ID:	DEFECT000628238
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.2.0	**Technology:** VXLAN - Virtual Extensible LAN

Symptom: VDX does not connect to NSX6.3.0 EA controllers.
Condition VDX to NSX6.3.0 EA controllers connection.

Defect ID: DEFECT000628474	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS7.0.1	**Technology:** NETCONF - Network Configuration Protocol
Symptom: "RD auto" configuration failing under evpn instance for primary node in vcs when configured through netconf.	
Condition "RD auto" configuration through netconf.	

Defect ID: DEFECT000629513	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.2.0	**Technology:** Logical Chassis
Symptom: Very rare scenario, MAC is not learned on port-channel	
Condition Create and add vlans to port-channel after HA failover	
Recovery: Enable and disable STP on interface.	

Defect ID: DEFECT000630071	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: VDX comes up with default config.	
Condition execution of "write erase" in past.	

Defect ID: DEFECT000630310	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: After unconfiguring and configuring dhcp relay address, DHCP offer is not forwarded to client from VDX.	
Condition DHCP relay	

Defect ID: DEFECT000630819	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: Execution of "show ip interface brief" CLI has missing default vrf status for port-channel	
Condition Execution of "show ip interface brief" CLI	

Defect ID: DEFECT000630999	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.1.0	**Technology:** LAG - Link Aggregation Group
Symptom: 1G SFP [SN 57-1000042-01] does not come online with fix speed	
Condition SFP SN- 57-1000042-01: Remove cable, Remove SFP, Insert SFP,	

Defect ID: DEFECT000631332

Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS4.0.0	**Technology:** Syslog
Symptom: Some Internal RAS log [Ex: BL-5282] are important and good to monitor those, but we don;t display internal RAS log on Console and we also don;t redirect them to syslog server.	
Condition RAS log monitoring through Console or syslog.	

Closed with code changes for NOS v7.0.1a

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as September 9th, 2016 in Network OS v7.0.1a.

Defect ID: DEFECT000543303	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS4.1.3	**Technology:** OpenStack Integration
Symptom: VDX can experience unexpected reload due to a daemon termination.	
Condition When any telnet session is in middle of some CLI through pipe option and stays there for 7 days, VDX daemon will terminate all existing socket sessions and try to reconnect. Distributed module is not able to clear the stuck session and reconnect fails as a result after 1 hrs it will get terminated.	

Defect ID: DEFECT000559340	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Syslog
Symptom: if we configure two syslog server (one with mgmt-vrf and another one with default-vrf) then first log message fails to send to syslog server.	
Condition Recently we have introduced VRF support for syslog and enhancing all corner cases.	

Defect ID: DEFECT000562737	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS4.0.1	**Technology:** OpenStack Integration
Symptom: SNMP trap of topology change will be sent from the switch, when switchport configuration is done on an interface where spanning-tree is shutdown.	
Condition Topology change trap will be observed, when switchport configuration is done on an interface in spanning-tree shutdown state.	

Defect ID: DEFECT000579138	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.1	**Technology:** Logical Chassis
Symptom: Very rare case chassis name set CLI fails.	
Condition After upgrade to 6.0.1	
Recovery: Reload and re-apply the CLI	

Defect ID: DEFECT000581797	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS

Reported In Release: NOS4.1.3	**Technology:** Logical Chassis
Symptom: "waiting for pending actions to exit" warning message appears on console session and eventually VDX experience unexpected reload.	
Condition: When any session is in middle of some CLI through pipe option and stays there for 7 days, VDX throws "waiting for pending actions to exist" error message and DCM gets terminated eventually.	

Defect ID: DEFECT000584668	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** IP Addressing
Symptom: The access-list configured for the interface with "Routed" keyword may not act upon the traffic flows on ingress destined towards the VRRP-E VIP.	
Condition: Applicable for the traffic destined to VRRP-E Virtual-mac coming on ingress & access-list using "Routed" keyword.	
Workaround: The traffic flows destined towards VE MAC will not be impacted & thus hosts can be configured to point to VE IP address as default gateway. Alternatively, remove the "routed" keyword in the access-list.	

Defect ID: DEFECT000584733	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Management GUI
Symptom: Firmware downgrade from NOS6.x to NOS5.x using BNA display below error in BNA even though firmware download process is successful on VDX. Download Failed: (Unknow Error Code: 1) Other Errors: Firmware Image download reboot operation has timed out	
Condition: Firmware downgrade from NOS6.x to NOS5.x using BNA cause the issue.	

Defect ID: DEFECT000586577	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** Multi-VRF
Symptom: VDX switch can go for unexpected reload after configuring no vrf.	
Condition: If static route leaks exist in system and we configure "no vrf" command.	

Defect ID: DEFECT000589286	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** UDLD - Uni-Directional Link Detection
Symptom: Link of 1G Copper SFP comes up too early during the power-cycle on VDX 6940-144S	
Condition: Power-cycle on VDX 6940-144S with 1G Copper SFP	

Defect ID: DEFECT000590465	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: channel-group configurations for port-channel member interfaces are lost upon reload.	
Condition: VDX replays configuration through file [startup-config] when configuration has been defaulted and it causes channel-group configuration lost.	

Defect ID: DEFECT000592597	
Technical Severity: Medium	**Probability:** Medium

Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** Software Installation & Upgrade
Symptom: Allowing for N+2 version upgrade with default config.	
Condition It's a RFE to allow upgrade of N+2 version.	

Defect ID: DEFECT000592647	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** NTP - Network Time Protocol
Symptom: Timezone set might fail	
Condition Particular timezone related files got corrupted. It is very rare scenario to hit.	
Recovery: Delete the failed timezone file under /usr/share/zoneinfo/ . Configure the timezone .	

Defect ID: DEFECT000592874	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Hardware Monitoring
Symptom: In very rare scenario they can observe interface flap	
Condition Due to excessive symbol errors	

Defect ID: DEFECT000595754	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** IPv6 Addressing
Symptom: Disabling autoconfig (autonomous address-configuration flag) for an IPv6 prefix in NOS 6.0.2 has no impact on router-advertisement.	
Condition Disabling autoconfig	

Defect ID: DEFECT000596280	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** IP Addressing
Symptom: Unable to delete an ACL.	
Condition When ACL is associated to the management interface of one or more switches in the VCS and the switch gets removed from VCS.	

Defect ID: DEFECT000596781	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** IPv6 Addressing
Symptom: Lifetime configuration value of VE interface IPv6 nd prefix is reset to infinite.	
Condition Doing "shutdown" and "no shutdown" configuration on the VE interface	

Defect ID: DEFECT000596932	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: Interfaces may not join into Dynamic LAG.	
Condition Static lag creation before dynamic LAG.	
Workaround: Configuring dynamic LAG first and then static	
Recovery: Delete the static LAGs and re-add the same.	

Defect ID: DEFECT000597053	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.1	**Technology:** VCS Fabric
Symptom: In rare scenario, VDX can send packets with TTL=0. Which can cause the connectivity issues.	
Condition VxLAN packets terminated on VDX6940 & BUM forwarder on other ISL partner.	
Recovery: Configure static MAC address for the specific IP address.	

Defect ID: DEFECT000598345	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS5.0.0	**Technology:** Rate Limiting and Shaping
Symptom: slow learning of hosts ARP entries in 6740 platform	
Condition In rare scenarios when there is a sudden burst of routed traffic.	

Defect ID: DEFECT000598524	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS6.0.2	**Technology:** Rate Limiting and Shaping
Symptom: rte_cap_acl debug tool won't work for 6940 platforms	
Condition Enable the rte_cap_acl tool support for 6940 platforms	

Defect ID: DEFECT000598657	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.1	**Technology:** SSH - Secure Shell
Symptom: Unexpected reload.	
Condition Rare scenario where remote host IP becomes NULL.	

Defect ID: DEFECT000598663	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS5.0.2	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom: DCMd daemon terminated and sudden reload occurred.	
Condition If customer has big cluster and actively executing CLI commands through script or monitoring tools [BNA] then Principal node receives too many message to handle and it hit this issue.	
Workaround: Please reduce any command execution frequency.	

Defect ID: DEFECT000600169	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom: IP MTU configuration is not working for VE interface when IP address or L3 VNI association is not present.	
Condition When IP MTU is configured, it is not applied on the VE interface.	
Workaround: Configure IP MTU followed by the configuration of the IP address.	

Defect ID: DEFECT000600482	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer

Reported In Release: NOS6.0.2	**Technology:** ARP - Address Resolution Protocol
Symptom: Inter-VLAN traffic that is routed on VDX is failing for specific end hosts.	
Condition Frequent ARP addition/deletion, topology changes, high CPU bound ARP/L3 traffic	
Recovery: 'clear arp no-refresh' command	

Defect ID: DEFECT000601715	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** CLI - Command Line Interface
Symptom: When copying files to/from VDX switch to TFTP server we are seeing errors when 'use-vrf' option is specified.	
Condition Copying files to/from VDX switch to TFTP server.	

Defect ID: DEFECT000601917	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** Multi-VRF
Symptom: change of MAC address of a host connected to VCS is not updated across user defined vrfs in ARP table.	
Condition Incorrect MAC address will be replied for an ARP request.	

Defect ID: DEFECT000601985	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** xSTP - Spanning Tree Protocols
Symptom: VDX switches running in a VCS cluster may encounter CIST Spanning-tree interoperability problem with certain Juniper switches where BPDU's sourced by the VDX may be dropped by the partner.	
Condition When VDX running in VCS cluster running distributed CIST spanning-tree & VDX switches are configured as spanning-tree root.	
Workaround: Change the spanning-tree root to partner switch.	

Defect ID: DEFECT000602227	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: SNMP OID 1.3.6.1.2.1.17.1.3 displays ‘No such instance' in output	
Condition snmpwalk for SNMP OID 1.3.6.1.2.1.17.1.3	

Defect ID: DEFECT000602239	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.2	**Technology:** ACLs - Access Control Lists
Symptom: VDX experience unexpected reload after configuring permit statement on standard ACL applied to management interface.	
Condition Configuration of permit statement on standard ACL applied to management interface.	
Workaround: NA	

Defect ID: DEFECT000602751	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** xSTP - Spanning Tree Protocols

Symptom:	User tries firmware downgrade and will hit error message as, User need to clean the config and then only downgrade can be done.
Condition	When "system-oui" configuration is done under "protocol spanning-tree" configuration mode and subsequently, a downgrade is done.
Workaround:	User needs to remove the config with "no system-oui" command under "protocol spanning-tree" mode.

Defect ID:	DEFECT000602764		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 2 Switching
Reported In Release:	NOS7.0.0	**Technology:**	xSTP - Spanning Tree Protocols
Symptom:	After spanning tree system OUI feature enabled and then disabled, the firmware download is failed.		
Condition	Doing spanning tree system OUI enable and disable. Then performing the firmware download.		

Defect ID:	DEFECT000603443		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Layer 2 Switching
Reported In Release:	NOS6.0.1	**Technology:**	LAG - Link Aggregation Group
Symptom:	Changing LACP timeout option in VDX can cause LACP PDUs to be sent at short intervals when neighboring device is cisco Nexus 7k. Changing LACP timeout option from long to short and again to long in both the devices can cause this behavior.		
Condition	LACP timeout option in VDX internally remains as short though configuration is shown as long.		

Defect ID:	DEFECT000603778		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS6.0.2	**Technology:**	IPv6 Addressing
Symptom:	When both "IPv6 vrrp-suppress-interface-ra" and "IPv6 VRRP VIP" are configured the IPv6 RA response to the IPv6 RS contains link-local address instead of the VIP address.		
Condition	Configure both "IPv6 vrrp-suppress-interface-ra" and "IPv6 VRRP VIP"		

Defect ID:	DEFECT000604054		
Technical Severity:	Medium	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS7.1.0	**Technology:**	BGP4 - IPv4 Border Gateway Protocol
Symptom:	Loopback interfaces are showing bogus IP MTU value, when global MTU is configured.		
Condition	Execution of "show ip interface lo <ID>" when global MTU is configured.		

Defect ID:	DEFECT000605476		
Technical Severity:	Low	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.1	**Technology:**	BGP4 - IPv4 Border Gateway Protocol
Symptom:	The route advertised by eBGP peer is not installed in the routing table.		
Condition	This issue occurs only in the self-referencing scenario i.e. when route prefix overlaps with the prefix of next hop from where the route is received.		
Workaround:	Isolate the bgp peering in a different subnet so that their prefix does not overlap with the routes being advertised between them		

Defect ID:	DEFECT000605776		
Technical Severity:	Low	**Probability:**	Low
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS6.0.2	**Technology:**	CLI - Command Line Interface

Symptom: New script will help to clear all the counters with single command	
Condition It is an enhancement	
Workaround: Use the individual commands to clear the counters	

Defect ID: DEFECT000608446	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.1	**Technology:** Logical Chassis
Symptom: VDX generates FFDC core file and throws Software 'verify' error on console.	
Condition Execution of copy default-config startup-config from VCS primary node.	
Workaround: NA	
Recovery: LC gets automatically recovered.	

Defect ID: DEFECT000608838	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: SFP Interface goes into administratively down state. Ex: [NSM-1028], 5673/2457, SW/0 | Active | DCE, ERROR, <hostname>, Incompatible SFP transceiver for interface TenGigabitEthernet 1/0/45 is detected	
Condition Execution of "[no] snmp trap link-status" command on an un-tunable SFP interface.	
Workaround: Please do not disable "snmp trap link-status" which is enabled by default on all interface.	
Recovery: Enable tunable-optics configuration and then disable it on impacted interface as below: tunable-optics sfpp channel 1 no tunable-optics sfpp channel Make interface up again: no shutdown	

Defect ID: DEFECT000608995	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** ARP - Address Resolution Protocol
Symptom: Traffic to/from dhcp host is not routed when the dhcp IP is assigned to a new host. The ARP for such host does not age out when age out timer expires.	
Condition DHCP Server is sending ACK packets to relay agent even when the client address is known. Mostly seen with Windows dhcp server.	
Workaround: Modify DHCP server settings so that it will send reply directly to dhcp client when client IP is present in the received dhcp message.	

Defect ID: DEFECT000610937	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: Gateway for default route obtained through DHCP remains in running configuration under mgmt-vrf even after deleting DHCP config and reloading the switch.	
Condition Invalid gateway for default route may appear after reloading the switch.	

Defect ID: DEFECT000611059

Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.1.0	**Technology:** Logical Chassis
Symptom: VDX experience unexpected reload due to DCMd daemon termination.	
Condition When Principal fail-over occurs, secondary nodes DB transaction cleanup fails on standby partition due to timing condition.	

Defect ID: DEFECT000611576	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.2	**Technology:** Logical Chassis
Symptom: Getting "% Error: VLAN string length(1139) is more than maximum length 1023" on reboot.	
Condition VDX with allowed vlan configuration string length more than 1023 can hit the issue at boot up & configuration replay time.	

Defect ID: DEFECT000613777	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP request packets are dropped on VDX, and are not relayed to DCHP server(s).	
Condition This affects only DHCP request packets with option-82. For example, an intermediate layer 2 node may have inserted option 82 in the packet and then forwarded to the VDX.	
Workaround: A workaround script is available to disable option-82 check on VDX	
Recovery: A workaround script can be used to recover from this issue	

Defect ID: DEFECT000614353	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.1	**Technology:** Hardware Monitoring
Symptom: After inserting a media 'SFP transceiver for interface XYZ is inserted' RASLOG is missing.	
Condition a media/SFP insertion	

Defect ID: DEFECT000615075	
Technical Severity: Low	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** Licensing
Symptom: LED on unlicensed and shutdown VDX 40G ports are slow blinking amber after boot. Expected behavior is off since it is unlicensed.	
Condition After reload, the single QSFP amber LED should only blink slow amber when all the 4 internal links/ports are offline and the port has a 40G Port Upgrade license reservation; otherwise it should be turned off (ie, no color/black).	

Defect ID: DEFECT000615168	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: User may not be able to login to VDX switch after ISSU upgrade from NOS7.0.0 to NOS7.0.1 with no activation option. Existing telnet sessions will not be impacted. Cold boot upgrades are not impacted	
Condition ISSU upgrade from NOS7.0.0 to NOS7.0.1 with no activation option	
Workaround: Perform ISSU firmware Install with auto activation. (Don't install firmware on a single node VDX running NOS7.0.0 to NOS7.0.1 with "noactivate" option. In case "logical-chassis" keyword used,	

	don't install firmware on single or multiple nodes in cluster running NOS7.0.0 to NOS7.0.1 without "auto-activate" option.) Existing open telnet sessions will not be impacted. Thus if you really want to perform ISSU firmware install without activation, then keep principal node telnet/ssh/console session open with infinite terminal time out using "terminal timeout 0". This will allow to perform "firmware activate" to recover from the impacted state
Recovery:	The user may login using another node in the VCS that is not yet upgraded and carry out principal switchover to make that VDX a Principal switch. Once done, execute "firmware activate". Alternatively, add a new switch to VCS cluster, make it principal and run command "firmware activate" which would recover all switches in VCS cluster

Defect ID: DEFECT000616987	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** BFD - BiDirectional Forwarding Detection
Symptom:	BFD session is not switched over to other available links if existing BFD session is deleted and added back.
Condition	BFD session remains in INIT state thus causing registered protocols with BFD to converge in longer duration.

Defect ID: DEFECT000617919	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** IPv6 Addressing
Symptom:	Unable to configure update-source for ipv6 interface, it throws syntax error: "xx/x/101" is an invalid value.
Condition	Configure update-source for ipv6 interface which is greater than 99.

Defect ID: DEFECT000619467	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Inband Management
Symptom:	ZR optics are undetected and shows data access errors when connected to edge ports other than xx/x/1
Condition	ZR optic connected to edge ports other than xx/x/1
Workaround:	Connect ZR optic on first interface.
Recovery:	Connect ZR optic on first interface and reseat the other interface ZR optic.

Closed with code changes for NOS v7.0.1

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as May 25, 2016 in Network OS v7.0.1.

Defect ID: DEFECT000556411	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: The RASLOG FCPH-1003 generated on console wrongly shows interface type as "Fi" instead of "Fcoe" with wrong tuple information. Functionality is not broken, only port type in raslog is printed wrongly.	
Condition When logins with Duplicate WWN are attempted on multiple ports at same time with Ethernet port being the port on which second login is attempted.	

Defect ID: DEFECT000562543	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: IP ACL for SNMP community and v3 user config lost after loading the config back to running-config from back-up config	
Condition When we do config upload of running configuration with SNMP IP ACL's applied on SNMP community/ v3 users.	

Defect ID: DEFECT000567339	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** ARP - Address Resolution Protocol
Symptom: debug arp packet output shows destination mac address of ARP request as ffff:ffff:ffff, instead of 0000:0000:0000	
Condition debug arp packet command is executed	

Defect ID: DEFECT000573107	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.1	**Technology:** ACLs - Access Control Lists
Symptom: When we applied IP ACL on SNMP community/user configuration, then wildcard subnet mask on IP ACL is not working on SNMP. But subnet mask on IP ACL is working fine on SNMP.	
Condition When we have wildcard subnet mask on IP ACL applied for SNMP configuration, then we will observe this issue.	

Defect ID: DEFECT000576391	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol

Symptom:	The HA failover start trap may not be seen for the HA fail over event, for user defined VRF in VDX-6740 platforms.
Condition:	The HA failover start trap may not be seen for user defined VRF, for the the HA fail over event.

Defect ID:	DEFECT000577171		
Technical Severity:	Medium	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Network Automation and Orchestration
Reported In Release:	NOS7.0.0	**Technology:**	OpenStack Integration
Symptom:	The NOSCLI command "show openflow interface" does not reflect the actual operating speed of the OpenFlow interface		
Condition:	If the interface speed has been manually changed to something else which is not same as suggested via interface name		

Defect ID:	DEFECT000577822		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Network Automation and Orchestration
Reported In Release:	NOS5.0.1	**Technology:**	OpenStack Integration
Symptom:	Errors [crc, encoding...] on 8G links.		
Condition:	The issue is only seen on 8G links to 3Par storage devices using 16G SFPs		
Workaround:	Changing the SFP to 8G SFP and running at 8G speed the issue was not seen.		

Defect ID:	DEFECT000577928		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Monitoring
Reported In Release:	NOS7.0.0	**Technology:**	MAPS - Monitoring and Alerting Policy Suite
Symptom:	Slot number need to be verified while creating groups onVDX8770 .		
Condition:	Currently the API which converts slot/port to port index doesn't throw error while creating group on VDX8770.		

Defect ID:	DEFECT000578258		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Data Center Fabric
Reported In Release:	NOS7.0.0	**Technology:**	IP Fabric
Symptom:	Traffic loss may be observed for destination subnets under non-default VRF advertised over BGP-EVPN using L3VNI.		
Condition:	Leaf nodes extending the VRF over BGP-EVPN are not advertising any prefix route.		
Workaround:	One of following options may be chosen: (1) Redistribute connected routes under VRF into BGP VRF. (2) Configure static IP route and redistribute into BGP VRF. (3) Configure network or static-network under BGP VRF instance.		

Defect ID:	DEFECT000579234		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 2 Switching
Reported In Release:	NOS7.0.0	**Technology:**	VLAN - Virtual LAN
Symptom:	Multicast traffic destined for static multicast address, will flood, if the mac is configured on remote node of VCS.		
Condition:	Static multicast MAC is configured in a remote node within a VCS, with no local interface part of the group.		

Defect ID: DEFECT000580478	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VLAN - Virtual LAN
Symptom: Sometimes, SFP removal messages are displayed incorrectly even though the media is present, when a chassis disable is executed after failover or ISSU.	
Condition Media presence check is incorrect on the new active partition after failover or ISSU	
Recovery: 'no shut' on the interface would make the correct Media presence state consistent.	

Defect ID: DEFECT000581205	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: In rare case, snmpv3 traps will not be received when any host is configured as SNMPv3 trap recipient under rbridge mode.	
Condition Configure snmpv3 host under rbridge mode.	
Recovery: reconfigure the specific v3host config under rbridge mode.	

Defect ID: DEFECT000581259	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.0	**Technology:** IP Fabric
Symptom: Even though overlay-gateway configuration is deactivated, BGP discovered dynamic tunnels are still present. Traffic loss will be observed if remote Leaf nodes send traffic over dynamic tunnels.	
Condition Overlay-gateway configuration is deactivated using "no activate" command.	
Workaround: Avoid deactivating the overlay-gateway using "no activate" command. Instead detach the RBridge from overlay gateway configuration.	

Defect ID: DEFECT000582010	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** ARP - Address Resolution Protocol
Symptom: Under rare conditions, some of the hosts may lost IP connectivity with the VDX switch acting as a layer-3 gateway.	
Condition This would occur if the MAC to the IP association of a VDX learnt ARP changes. ie. For the same IP address, the MAC changes from say MAc1 to Mac2.	
Recovery: "clear arp no-refresh" would clean the ARP table and recover from the problem state.	

Defect ID: DEFECT000582119	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.0.0	**Technology:** IGMP - Internet Group Management Protocol
Symptom: The tunnel terminated IGMP frames sent to other nodes can loop back to the source node. The CPU generated IGMP frames are not getting source suppressed in active-active gateway.	
Condition This happens in specific tunnel topology with multicast root rbridge and BUM forwarder. The tunnel terminated IGMP frames sent to other nodes are trapped and flood back on the vlan by control path. These packets can loop back to source node.	
Recovery: Shut down the tunnel	

Defect ID: DEFECT000583123	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.2	**Technology:** Logical Chassis
Symptom: There is a time delay(debounce-timer delay) of approx 1sec between underlay network down and tunnel down because of which traffic impact may occur for this debounce-timer duration .	
Condition The above mentioned time delay happens whenever tunnel goes down. Now customer is provided with the following knob to suppress the debounce-timer delay. [no] system tunnel suppress-debounce	

Defect ID: DEFECT000584215	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** xSTP - Spanning Tree Protocols
Symptom: IEEE BPDU packets are flooded from one VF to another, in the absence of "spanning-tree ieee-bpdu limit-vlan-flood" configuration.	
Condition IEEE BPDU packet are received at the ingress port of a switch configured with VFs.	

Defect ID: DEFECT000584364	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.0	**Technology:** Syslog
Symptom: User can configure the user defined vrf in cluster, though the user defined vrf is not configured on all the rbridge's.	
Condition In cluster though the vrf is not configured on all the rbridge's, it is allowing to configure syslog-server on user defined vrf.	

Defect ID: DEFECT000584709	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VLAN - Virtual LAN
Symptom: Physical or port-channel is not added back to normal VLAN in a particular sequence.	
Condition Physical or port-channel is not added back to normal VLAN after changing a private VLAN to a normal VLAN on a primary VLAN	
Workaround: Delete private VLAN and create the same again instead of changing the type on a private VLAN.	
Recovery:	

Defect ID: DEFECT000585043	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** BGP4+ - IPv6 Border Gateway Protocol
Symptom: when multi-hop BFD session is created, default BFD interval will be shown for loopback interface in show bfd output	
Condition Default BFD interval will be shown for loopback interface in show bfd output	

Defect ID: DEFECT000585392	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** CLI - Command Line Interface
Symptom: BNA unable to discover NOS switch, when NOS switch is connected to a FCR.	

Condition	1. Have a setup with VCS connected to FCR 2. Install BNA 14.0.1 3. Shift to IP tab. 4. In Discovery dialog, add IP of one of the VCS switches and click ok. 5. Observe the device is not discovered and shows "Discovery Failed" mesage.

Defect ID: DEFECT000585445	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.0	**Technology:** Logical Chassis
Symptom: Some 40 GbE ports on VDX 6940-144S may not come online after cold boot.	
Condition Some 40 GbE ports on VDX 6940-144S may not come online after cold boot.	
Workaround: Execute noscli command shut / no shut on the 40 GbE port to bring it online.	
Recovery: Execute noscli command shut / no shut on the 40 GbE port to bring it online.	

Defect ID: DEFECT000585723	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: MIB walk for IP Forwarding MIB will return with an error with unnumbered interfaces.	
Condition MIB walk of IP Forwarding MIB and has ECMP routes with unnumbered/L3 VNI interfaces will lead to error.	

Defect ID: DEFECT000585903	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.0	**Technology:** MAPS - Monitoring and Alerting Policy Suite
Symptom: IPMAPS Custom policy modifications are not reflected.	
Condition IPMAPS Custom policy modifications are not dynamically reflected.	
Workaround: Revert to default policy, and then reapply custom policy. Run CLI "enable policy <policy_name> actions <actions_list>" then we can re-enable the same policy to reflect the changes made. Here actions_list can be same as what was already configured.	

Defect ID: DEFECT000585927	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Port Mirroring
Symptom: Mirrored VXLAN packets outer header was getting removed while going out on destination mirror port. This gives misleading information when validating the VXLAN mirroring. The data path traffic goes out properly but mirrored copy has the outer header stripped only in VXLAN frames.	
Condition The VXLAN packets outer header is not handled properly and causing the stripped packet to go out on destination mirror port	
Recovery: This is not functional data path issue, but mirrored information shows wrong details.	

Defect ID: DEFECT000585960	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.0	**Technology:** VCS Fabric

Symptom:	40G Interface is administratively (or) protocol down with FFDC raslogs
Condition	Admin operations on 40G Interface.
Recovery:	Shut/no-shut both interfaces on either side of the link

Defect ID:	DEFECT000585970		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.0	**Technology:**	VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom:	On VDX 8770 switch, maximum VRRPv2 and VRRPv3 sessions supported on an interface are 16 in NOS6.0.x. This limit got increased to 32 in NOS7.0.0. Firmware downgrade from NOS7.0.0 to NOS6.0.x need to be blocked in case if more that 16 sessions are present on an interface.		
Condition	Issue can be seen if more that 16 VRRPv2 and VRRPv3 sessions are configured on an interface and firmware is downgraded from NOS7.0.0 to NOS6.0.x. In this case only 16 sessions will get enabled and rest will be disabled.		
Workaround:	As a workaround user should delete/unconfigure more than 16 VRRPv2/VRRPv3 sessions present on an interface in NOS7.0.0 before downgrading it to NOS6.0.x.		

Defect ID:	DEFECT000586001		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.0	**Technology:**	ARP - Address Resolution Protocol
Symptom:	IPv4 DHCP relay packets forwarded through a VxLAN tunnel is trapped but not forwarded unless ARP is forcefully resolved.		
Condition	Running DHCP Relay in IP Fabric EVPN.		
Workaround:	Resolve ARP forcefully.		

Defect ID:	DEFECT000586178		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	VCS
Reported In Release:	NOS7.0.0	**Technology:**	Logical Chassis
Symptom:	Non-existent port-channel shows up in "show fcoe interface ethernet"		
Condition	1. Create a port-channel 2. Add members to it and make it fcoe-provisioned 3. Delete the port-channel		
Workaround:	Remove FCOE provisioning from port-channel before deleting it		

Defect ID:	DEFECT000586252		
Technical Severity:	Medium	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Management
Reported In Release:	NOS7.0.0	**Technology:**	Configuration Fundamentals
Symptom:	Physical or port-channel is not added back to normal VLAN in a particular sequence.		
Condition	Physical or port-channel is not added back to normal VLAN after changing a private VLAN to a normal VLAN on a primary VLAN.		
Workaround:	Delete private VLAN and create the same again instead of changing the type on a private VLAN.		

Defect ID:	DEFECT000586338		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Data Center Fabric
Reported In Release:	NOS7.0.0	**Technology:**	Logical Chassis

Symptom:	IN VDX 6940-144S, link flap occurs on 40 G ISL ports if breakout configuration mis-matched with any adjacent ports.
Condition	In VDX 6940-144S, a 40 G Port with breakout QSFP is not configured as breakout may cause adjacent 40 G port to flap, whereas its peer port is configured as 40G breakout.
Workaround:	In VDX 6940-144S, configure 40 G port as breakout if the peer port is configured as 40G breakout. After that, the link flap on the port will stop.

Defect ID:	DEFECT000586856	
Technical Severity:	High	**Probability:** Low
Product:	Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.0	**Technology:** BGP4+ - IPv6 Border Gateway Protocol
Symptom:	BGP add path is not showing up all the 4 available paths.	
Condition	Sometimes when the RR is reloaded the BGP add path is not showing up all the 4 available paths.	

Defect ID:	DEFECT000586973	
Technical Severity:	High	**Probability:** High
Product:	Extreme Network OS	**Technology Group:** Security
Reported In Release:	NOS7.0.0	**Technology:** LDAP - Lightweight Directory Access Protocol
Symptom:	LDAP authentication is not working	
Condition	LDAP authenticaion is not working via inband deafult and non-deafult-vrf	

Defect ID:	DEFECT000587170	
Technical Severity:	Medium	**Probability:** Low
Product:	Extreme Network OS	**Technology Group:** VCS
Reported In Release:	NOS6.0.1	**Technology:** Logical Chassis
Symptom:	Continuous occurrence of ECC correctable errors	
Condition	This is very rare scenario to occur.	

Defect ID:	DEFECT000587276	
Technical Severity:	High	**Probability:** Medium
Product:	Extreme Network OS	**Technology Group:** Monitoring
Reported In Release:	NOS6.0.2	**Technology:** Hardware Monitoring
Symptom:	Blocked or stopped fan units may not show up as faulty.	
Condition	This was a defect in the original release of this product.	

Defect ID:	DEFECT000587419	
Technical Severity:	High	**Probability:** High
Product:	Extreme Network OS	**Technology Group:** Monitoring
Reported In Release:	NOS7.0.0	**Technology:** Syslog
Symptom:	Ipv6 syslog-server not working via inband def-vrf and user-vrf. When multiple server configured as default-vrf or user-defined vrf.	
Condition	Was deferred from 7.0.0 but fixed in 7.0.1.	

Defect ID:	DEFECT000587615	
Technical Severity:	High	**Probability:** Medium
Product:	Extreme Network OS	**Technology Group:** Management
Reported In Release:	NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom:	The SNMP V3 traps may not be received for the SNMP v3 host recipients configured under rbridge mode.	

Condition	The trap may not be received after upgrade from NOS6.0.1a to NOS7.0.0 with cold boot option

Defect ID:	DEFECT000587617
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom:	Static leaked VRF route can not be imported into BGP RIB-IN and can't advertise via eBGP/iBGP using network/static command.
Condition	Advertise static leaked VRF route via BGP.
Workaround:	Use "redistribute static" command to leak the static VRF route into BGP RIB-IN and then can advertise it via eBGP/iBGP.

Defect ID:	DEFECT000587637
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS7.0.0	**Technology:** NETCONF - Network Configuration Protocol
Symptom:	Netconf RPC "get-interface-detail" does not provide physical interfaces details. It provides only port-channel details.
Condition	This issue will happen only when number of port-channels configured are equal to or more than 70. If number of port-channels are less than 70, this issue will not be encountered.
Workaround:	Total number of port-channels configured should be less than 70.
Recovery:	If total number of port-channels configured are exceeding 70, delete few port-channels to reduce the total count to be less than 70.

Defect ID:	DEFECT000587654
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom:	The ECMP configuration in the hardware profile shows incorrect values.
Condition	This will only happen when a user changes both route-table profile type and maximum-path at the same time using the hardware-profile command
Workaround:	The user can change the route-table profile type and maximum-path one at a time.
Recovery:	The user can re-run the hardware-profile command to set the maximum-path with the correct value.

Defect ID:	DEFECT000587767
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.0	**Technology:** Logical Chassis
Symptom:	Possible for Edge port interfaces to stay inactive after chassis enable command.
Condition	This issue can occur in releases prior to NOS 7.0. If multiple attempts to issue the chassis enable command is failed and the command is retried, it is possible that the configuration replay will be blocked after the chassis enable succeeds.
Recovery:	Issue chassis disable then chassis enable.

Defect ID:	DEFECT000587804
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.0	**Technology:** IP Fabric

Symptom:	Even though there are no matching EVPN import route-targets configured under VRF, imported EVPN routes are present in BGP VRF table.
Condition	EVPN import route-target(s) is/are removed while matching routes are present in BGP-EVPN and imported into BGP VRF table.
Workaround:	Issuing "clear bgp evpn neighbors all soft in" command should cleanup the routes which are still imported in BGP VRF instance after matching EVPN import route-targets are removed.

Defect ID:	DEFECT000587925		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 2 Switching
Reported In Release:	NOS5.0.2	**Technology:**	FCoE - Fibre Channel over Ethernet
Symptom:	Syslog daemon generates a silent core file as it is restarted to reload configuration. There is no crash or loss of traffic in this case.		
Condition	Defect exists in previous releases of NOS. Core file is generated due to SIGTERM signal received by syslog instead of SIGHUP.		

Defect ID:	DEFECT000588001		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	VCS
Reported In Release:	NOS7.0.0	**Technology:**	AMPP - Automatic Migration of Port Profiles
Symptom:	Traffic may flood though the source mac was seen behind profiled port		
Condition	Port-profile is configured on a vlag and 'clear-mac-address table' command is executed more than 10 times in short interval.		

Defect ID:	DEFECT000588178		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Network Automation and Orchestration
Reported In Release:	NOS7.0.0	**Technology:**	OpenStack Integration
Symptom:	Interface remains protocol down after speed change		
Condition	speed change config performed on an interface which is not in protocol up state.		
Recovery:	shut/no-shut the interface		

Defect ID:	DEFECT000588190		
Technical Severity:	High	**Probability:**	Medium
Product:	Extreme Network OS	**Technology Group:**	Data Center Fabric
Reported In Release:	NOS7.0.0	**Technology:**	IP Fabric
Symptom:	Aggregate route(s) configured under BGP VRF instance are not exported into BGP-EVPN.		
Condition	BGP VRF address-family is removed and added back.		
Workaround:	Remove the aggregate route configuration under BGP VRF instance and configure it again.		

Defect ID:	DEFECT000588238		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.0	**Technology:**	BGP4 - IPv4 Border Gateway Protocol
Symptom:	'Invalid value' error is thrown for 'vni add' command under 'rbridge-id evpn-instance' mode.		
Condition	Configure 'vni add' command under 'rbridge-id evpn-instance' mode. If the value falls in 10000000-15999999 range.		
Workaround:	1. Use a VNI range in 'vni add' command that is less than 10000000-15999999. 2. Use 'vni <vni-number>' CLI under 'rbridge-id evpn-instance' mode.		

Defect ID: DEFECT000588451	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom: IPv6 traffic may not forward when it received on tunnel	
Condition When bigger VNI like 10000000 configured as l3vni	

Defect ID: DEFECT000588519	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: When the RBridge responsible for Multicast distribution over VXLAN Tunnels is powered off, there is a multi-second delay before the multicast stream changes to the standby RBridge.	
Condition Issue when the RBridge responsible for multicast distribution is powered off or the ISL cables are physically disconnected.	

Defect ID: DEFECT000588730	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.2	**Technology:** NETCONF - Network Configuration Protocol
Symptom: When querying the VDX netconf server an invalid yang model “ietf-netconf-notifications-ann” is advertised.	
Condition This issue will show up when trying to view the mounted netconf capabilities for a VDX mounted with Extreme SDN Controller (BSC).	

Defect ID: DEFECT000588822	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.1	**Technology:** TRILL - Transparent Interconnection of Lots of Links
Symptom: An ISL (Inter Switch Link) flap is seen on VDX6940.	
Condition This can be seen due to un-handled internal memory parity error interrupts.	

Defect ID: DEFECT000588918	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.1	**Technology:** VCS Fabric
Symptom: Customer encountered an unexpected VDX6740 reload:RAS logs & stack trace for the reset as below:2016/02/16-01:02:57, [SEC-1203], 795596, SW/0 | Active, INFO, NINMUM03-DC1-R107-NSPL-RTR-049, Login information: Login successful via TELNET/SSH/RSH. IP Addr: A.B.C.D.2016/02/16-01:03:13, [SEC-3022], 795597, SW/0 | Active, INFO, NINMUM03-DC1-R107-NSPL-RTR-049, Event: logout, Status: success, Info: Successful logout by user [admin].2016/02/16-06:38:11, [HSL-1012], 795598, INFO, VDX6740T-1G, Interface lo is link up2016/02/16-06:38:11, [HSL-1012], 795599, INFO, VDX6740T-1G, Interface eth0 is link up2016/02/16-06:38:11, [HSL-1012], 795600, INFO, VDX6740T-1G, Interface eth1 is link up2016/02/16-06:38:11, [HASM-1004], 795601, INFO, VDX6740T-1G, Processor reloaded - Software Fault:Kernel Panic.2016/02/16-06:38:11, [HASM-1026], 795602, WARNING, VDX6740T-1G, The last reboot is due to Kernel Panic in kernel .NOSCLI show support:Tue Feb 16 09:25:17 IST 2016**********	

Condition: When high rate of TFTP ip_directed broadcast packets are sent destined to known subnets.	

Defect ID: DEFECT000589893	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Hardware Monitoring
Symptom: Request for Enhancement to optimize the fan speed to achieve better temperature distribution for the VDX 6740T & VDX6740T-1G switches	
Condition: Applies only to the VDX6740-T-R & VDX6740T-1G-R switches running port-side exhaust fans	

Defect ID: DEFECT000589911	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.2	**Technology:** VCS Fabric
Symptom: Data loss is seen when an ISL port is flapped in a VCS that is employing VXLAN to connect to the remote data center VCS fabric.	
Condition: Flapping ISL link in a VCS fabric connecting to remote data center network using VXLAN/VTEP technology, would incur 1 to 2 seconds of data loss.	

Defect ID: DEFECT000589967	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.0.0	**Technology:** IGMP - Internet Group Management Protocol
Symptom: error message seen on console while trying to configure Query-Interval on L3 physical interface Or PO. The queries will be generated at default time interval = 125 sec.	
Condition: This issue is seen when user try to configure Query-Interval on PO/Physical interface being in shut state. The config will not be applied as long as interface is in "protocol-down state"	
Workaround: Customer should bring the interface in "protocol up" state before applying Query-Interval config. Once the interface is up, Config will succeed.	

Defect ID: DEFECT000590478	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS5.0.2	**Technology:** IPv4 Multicast Routing
Symptom: mcasgt process termination	
Condition: The issue is seen when multicast routes are added and deleted from the system, which leaves some amount of memory leak, which grows over time and causes a system crash.	
Workaround: Yes	

Defect ID: DEFECT000590808	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: Hidden commands under debug and foscmd hide group were not shown as part of show running config even after unhiding and configuring them. Even the copy running to file was not having the configuration after copy command was executed after unhiding.	
Condition: Config commands under hide group "debug" and "foscmd" have to be executed after unhiding respective hide group. Post this, executing "show running config" will not show these unhiden configurations.	

Defect ID: DEFECT000591223

Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: This is an enhancement that introduces a new CLI under rbridge-id sub-mode to configure the behaviour of some IF-MIB attributes: ifName and ifDescr. If this knob is configured to 3-tuple, then the above 2 objects will be of 3-tuple format. Else, they will be of 2-tuple format. These 2 attributes will also be in the same format during Link Up/Down Trap generation.	
Condition This is applicable only for ifName and ifDescr attributes of IF MIB and the linkUp/Down traps.	

Defect ID: DEFECT000591225	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom: SNMP IP ACL config mismatch between the Frontend & Backend database.	
Condition Reload with default config will retain the IP ACL data for SNMP community string.	

Defect ID: DEFECT000591700	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS6.0.1	**Technology:** QoS - Quality of Service
Symptom: BUM traffic has higher latency compare to data traffic.	
Condition BUM traffic use store and forward method and data traffic use cut through method.	

Defect ID: DEFECT000592128	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** IP Addressing
Symptom: Software Fault: A rare memory corruption issue in the tty driver caused Kernel Panic and rebooting of the switch.	
Condition The issue was introduced in the 2.6.34 kernel and the same was addressed by a open source fix in the tty driver.	

Defect ID: DEFECT000592398	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.0	**Technology:** Logical Chassis
Symptom: During multi-cast tree formation, a rbridge with a configured root priority level may not take effect for the tree's formation. Instead, the configured rbridge behaves as though it has a default or lowest priority configuration. However, when displaying the running configuration, it shows theexpected tree root priority configuration.	
Condition Following an operation where a rbridge boots up with a default configuration, and then downloads it's configuration from the active cluster, a non-default setting for the rbridge's multi-cast root priority may not take affect. This may happen such as after a 'vcs replace' operation.	
Recovery: Rebooting the affected node forces it to refresh the effective priority value for the multi-cast tree root priority. Alternatively, explicitly changing the priority to a different value and then setting it back to the original desired value causes the priority to be updated. However, setting the root priority to a different value may affect the multi-cast tree formation depending on the temporary priority specified.	

Defect ID: DEFECT000592617	
Technical Severity: Medium	**Probability:** Medium

Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** xSTP - Spanning Tree Protocols
Symptom: IEEE BPDU Local VLAN tunnel CLI allowed to be configured when protocol spanning tree is already configured or vice versa.	
Condition: When both STP protocol and IEEE BPDU Local VLAN tunnel CLI are enabled at the same time.	

Defect ID: DEFECT000593092	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.1	**Technology:** Security Vulnerability
Symptom: Security vulnerabilities.	
Condition: Unix open source code [openssh & openssl] is vulnerable. Please take a look into Extreme CVE [Common Vulnerabilities and Exposures] list to get detail of which CVE is fixed in which NOS release.	

Defect ID: DEFECT000593245	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** Multi-VRF
Symptom: Ping Round-Trip-Times fluctuate between 4 and 16 ms.	
Condition: Happens in 6.0.2a and later releases.	

Defect ID: DEFECT000593960	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: With 3-tuple format configured for ifDescr and ifName, the linkUp/Down traps generated still contain ifDescr var-bind in 2-tuple format.	
Condition: This is related to ifDescr var-bind in the linkUp/Down trap only.	

Defect ID: DEFECT000594223	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Software Installation & Upgrade
Symptom: TFTP server/service was enabled by default.	
Condition: Any device from outside can try to connect VDX using TFTP and VDX burn its resources unnecessary.	

Defect ID: DEFECT000594815	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** VLAN - Virtual LAN
Symptom: The execution of command "show vlan brief" will cause the box to reboot.	
Condition: This issue may be seen when all the following conditions are met. 1. There are more than 40 nodes in a Logical Chassis. 2. VFAB is enabled on the cluster. 3. There are 10 vlans configured. 4. There are more than 1000 ports configured on each vlan. 5. show-vlan-brief was executed.	
Workaround: Instead of "show vlan brief", the user can execute "show interface trunk" to check the vlan-port configurations.	

Defect ID: DEFECT000595226	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.1	**Technology:** Syslog
Symptom: IPv4 and IPv6 syslog servers were not working when configured together as default/non-default VRF.	
Condition Defect exists in 7.0.0 also.	

Defect ID: DEFECT000595395	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: IP DHCP Relay is not working properly when enabled on VRRP-E master interface	
Condition Operating IP DHCP Relay together with VRRP-E	
Workaround: toggle the VE interface	

Defect ID: DEFECT000595980	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: When tunnel tagged-ieee-bpdu is enabled on any of the interface, protocol spanning-tree is allowed to be configured.	
Condition Tunnel tagged-ieee-bpdu configured before configuring protocol spanning tree.	

Defect ID: DEFECT000596257	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** Software Installation & Upgrade
Symptom: After reload, though the uplink interface is down, the downlink tracking interface is still up.	
Condition All the downlinks interface are brought up , irrespective of the uplink interface state after reboot.	

Defect ID: DEFECT000596496	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Protocol spanning-tree configuration will not be allowed even after removing the "spanning-tree ieee-bpdu limit-vlan-flood" and "tunnel tagged-ieee-bpdu" configuration.	
Condition When all the switches in the VCS are configured with "spanning-tree ieee-bpdu limit-vlan-flood" and one or more switches are removed from VCS.	
Recovery: Copy running configuration to remote. Reload the switch with default configuration and copy back the running configuration.	

Defect ID: DEFECT000597782	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** VLAN - Virtual LAN
Symptom: The management MAC and one of the VE MACs may conflict.	
Condition This is a software defect that has affected the VDX6940-36Q and VDX6940-144S since their release.	

Defect ID: DEFECT000600386

Closed with code changes for NOS 7.0.1

Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom: When EVPN related configuration is done in RBridge range mode, the nodes might reboot	
Condition: When EVPN related configuration is done in RBridge range mode	
Workaround: Instead of using RBridge range mode, use RBridge mode to configure the individual rbridges	
Recovery: If issue is encountered, reboot the RBridges to recover	

Closed without code changes for Network OS v7.0.2

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of November 20, 2017 in Network OS v7.0.2.

None

Closed without code changes for Network OS v7.0.1c

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change in Network OS v7.0.1c.

None

Closed without code changes for Network OS v7.0.1b

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of January 30, 2017 in Network OS v7.0.1b.

None

Closed without code changes for Network OS v7.0.1a

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of September 9th, 2016 in Network OS v7.0.1a

Defect ID: DEFECT000564498	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS4.1.3	**Technology:** Port Monitoring
Symptom: VDX Ports are showing up as spfabsent as status in "show interface status"	
Condition We have verified the code changes done for displaying "show interface status" between 4.1.2a and 4.1.3. In 4.1.2a, we have the changes as per defect xxxxxx. Hence it works fine in 4.1.2x release. But in 4.1.3, we have modified the code to address a defect which corrects display status for copper ports, in which we have missed to handle the code to display correct status for these internal ports in the switch.	
Workaround: None	

Closed without code changes for Network OS v7.0.1

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of May 25, 2016 in Network OS v7.0.1.

Defect ID: DEFECT000552520	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** Static Routing (IPv4)
Symptom: Memory leak observed with repeated addition/deletion of VRFs using an automated script	
Condition: Adding and deleting VRFs repetitively	
Workaround: Remove routes before deleting the VRF	

Defect ID: DEFECT000556025	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: On port channel, Fabric Watch provides incorrect message SFP is absent when link is shut and SFP is not removed.	
Condition: On Chassis VDX switches, SFP info from Fabric Watch may mislead when port is shut.	

Defect ID: DEFECT000558216	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: Attaching VE interface to another rbridge is taking more time.	
Condition: With more than 2K VE interfaces created, attaching a VE interface to another rbridge takes more 1 sec for each VE interface.	

Defect ID: DEFECT000579176	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: BFD may not work over Layer 3 Port Channels when the gateway address and nexthop pointing to port channel overlap	
Condition: Running BFD with Layer 3 Port Channels	

Defect ID: DEFECT000581124	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS7.0.0	**Technology:** OpenStack Integration
Symptom: 40G Interface is either protocol down (or) administratively down after "no shut" operation.	
Condition: Breakout Config operation performed on 40G Interface connected to a 40G Interface	

Defect ID: DEFECT000583324	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Medium

Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: ISL fails to come up due to Trunking Error	
Condition When port is enabled between VDX8770 LC48x10G and VDX 6940 4x10G breakout interfaces	
Recovery: Issue shut followed by no shut on the port.	

Defect ID: DEFECT000585015	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.0	**Technology:** AMPP - Automatic Migration of Port Profiles
Symptom: traffic may flood for the non-profiled macs even if the global knob is disabled.	
Condition 'no allow non-profiled-macs' is configured.	
Workaround: configure and remove 'allow non-profiled-macs' again.	

Defect ID: DEFECT000586205	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.0	**Technology:** Syslog
Symptom: Syslog server not working via inband under def-vrf, mgmt.-vrf and user-vrf, all having different ip address	
Condition Was deferred from 7.0.0 and fixed in 7.0.1.	

Defect ID: DEFECT000587880	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** IPv6 Addressing
Symptom: IPv6 DHCP relay SOLICIT packets are not getting intercepted after coldbooot upgrade.	
Condition Running DHCP relay in a IP Fabric EVPN environment sometimes.	
Workaround: Delete and re-configure the same L3 interface where relay config is present.	

Defect ID: DEFECT000591172	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** IP Addressing
Symptom: It is seen that configuration of global VE interface is missing in output of" show running-cofig". Configuration of same global VE was already present in protocol daemon, hence configuration of global VE again is not allowed.	
Condition In a rare scenario during global VE configuration.	

Defect ID: DEFECT000592812	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** VCS Fabric
Symptom: The Dcmd process may terminate and cause an MM to failover.	
Condition Given a large node cluster with 4,000 tengigabit ethernet interfaces or more, a NETCONF get-config request can cause the Dcmd process to hit an out of memory condition and cause it to terminate.	
Workaround: Use the vcs-rbridge-context NETCONF RPC to set an rbridge filter prior to issuing the NETCONF get-config request. This will limit the get-config results to the rbridge context set and reduce memory usage of the Dcmd process.	

Closed without code changes for Network OS v7.0.0

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of May 3, 2016 in Network OS v7.0.0

Defect ID: DEFECT000393266	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS3.0.0	**Technology:** Configuration Fundamentals
Symptom: "fcsp auth-secret dhchap" command with incorrect node field fails with in appropriate error message.	
Condition "fcsp auth-secret dhchap" command with incorrect node field	

Defect ID: DEFECT000396994	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS3.0.0	**Technology:** SSH - Secure Shell
Symptom: Garbled characters may be seen on SSH session during login	
Condition SSH login to the management interface, mostly seen after changing the user password.	

Defect ID: DEFECT000409067	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS3.0.0	**Technology:** ACLs - Access Control Lists
Symptom: Command "show access-list" is displaying rules attached on the management interface with thewrong protocol information. In the case where protocol is “0”, it is showing protocol as “ip”.	
Condition Adding ACL rule with "0" as protocol number.	
Workaround: Explicitly use “4” or "ip" while applying IP protocol to ACL rule rather than using the protocol number "0".	

Defect ID: DEFECT000482263	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS4.1.0	**Technology:** VLAN - Virtual LAN
Symptom: Spanning tree does not converge properly and can lead to traffic loop.	
Condition Issue is seen when PVST or RPVST is configured and spanning tree is enabled on few switchports with VLAN-MAC classification.	
Workaround: Do not enable spanning tree on switchports having VLAN-MAC classification.	
Recovery: Disable spanning tree on all switchports having VLAN-MAC classification.	

Defect ID: DEFECT000489529	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS3.0.1	**Technology:** AAA - Authentication, Authorization, and Accounting
Symptom: Admin cannot create a user-role which would dis-allow `show running-config' command	
Condition When configuring AAA service for authorization using RBAC model.	

Defect ID: DEFECT000491044	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS4.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: When port-channel is either admin down or operationally down, the command "show interface status" shows the interface status to be 'not connected'	
Condition When 'shutdown' command is issued on the port-channel or when the member ports are brought operationally down.	

Defect ID: DEFECT000491465	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.0.1	**Technology:** VMWare
Symptom: On upgrade from 3.x with vCenter configured we may see some Vcenter configuration changes not getting updated on one of the nodes in the cluster.	
Condition On upgrade from 3.x to higher releases with vCenter configuration.	
Workaround: Remove the vCenter configuration before upgrade. And on successful upgrade re-configure the vCenter. This work-around may result in traffic loss.	
Recovery: In case of facing the issue, vCenter configuration should be removed and re-applied.	

Defect ID: DEFECT000492196	**Technical Severity:** High
Reason Code: Feature/Function Not Supported	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.1.0	**Technology:** Configuration Fundamentals
Symptom: The show media output may report varying power levels for various optics in the same switch. This difference may be ignored.	
Condition This may be seen on 10G optics.	

Defect ID: DEFECT000492427	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.1.0	**Technology:** Configuration Fundamentals
Symptom: NETCONF "get-config" response will be slower based on the configuration.	
Condition Retrieving running configuration through NETCONF will be slower and the response time will increase substantially if the configuration is large.	
Workaround: Use sub-tree filtering in the get-config which will be faster.	

Defect ID: DEFECT000503858	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.0	**Technology:** VCS Fabric
Symptom: In case of errors during config file download, there is no reference to the line number which caused the error	
Condition Observed when the user downloads a config file onto the switch	

Defect ID: DEFECT000515693	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.0	**Technology:** Fabric Authentication
Symptom: Ethernet LED does not glow green when a port is enabled. This is applicable to Lenovo VDX embedded platforms.	
Condition No loss of system functionality. Useful data for Lenovo embedded VDX customers.	

Defect ID: DEFECT000517443	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.0	**Technology:** Configuration Fundamentals
Symptom: Browser level login box will be displayed. Element manager will not login If credentials entered in the box.	
Condition The defects will be seen in following conditions: 1. When invalid credentials are entered in NOS Element Manager login page 2. When element manager gets timed out	
Recovery: Click 'Cancel' on the browser level login box and login using the Element Manager login page	

Defect ID: DEFECT000521284	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.0	**Technology:** Configuration Fundamentals
Symptom: A bad sfm might cause a panic while diags test is running on it.	
Condition Detection of a Bad SFM Card after a chassis reboot.	

Defect ID: DEFECT000527393	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Configuration Fundamentals
Symptom: Web Element Manager hardware view doesn't match with the Physical hardware view	
Condition This condition occurs when port and link is up in the Switch module	

Defect ID: DEFECT000527401	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Configuration Fundamentals
Symptom: Web Element manager doesn't report session expiry and doesn't report current status on removing the management module from the switch.	
Condition This condition occurs when the management module is removed from the switch	
Workaround: Management module should be removed after exiting the active web element manager session	
Recovery: Exiting or logging out of the Web element manager session manually	

Defect ID: DEFECT000528475	**Technical Severity:** Medium
Reason Code: Feature/Function Not Supported	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.1	**Technology:** Fabric Authentication
Symptom: On configuring OUI based port security on the port and sending traffic with a different OUI could lead to traffic loss.	
Condition If customer is planning to use OUI based port security on ports.	

Defect ID: DEFECT000529660	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** Configuration Fundamentals
Symptom: Web Element Manager session gets logged out before 30 minutes.	
Condition: Web Element Manager session will be logged out before 30 minutes in the following conditions: 1) If there is an unauthorized request sent to switch (or) 2) Internal bad server error	

Defect ID: DEFECT000531718	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Configuration Fundamentals
Symptom: The command may take upto 1.5 minutes to throw an error when trying to configure non-existing interface.	
Condition: Occurs only in large cluster with high scale implemented.	

Defect ID: DEFECT000538035	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** Configuration Fundamentals
Symptom: Browser level message will be displayed for user notification in Network OS Element manager.For e.g. If there are any error condition, the message will be append with unwanted characters and lines	
Condition: User performing invalid configuration in Network OS Element manager. For e.g. Trying to set a role which doesn't exists in the device	

Defect ID: DEFECT000546734	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.1	**Technology:** RA Guard - Router Advertisement Guard
Symptom: Router Advertisement packets will be accepted even though RA Guard is configured on LAG and vLAG interfaces.	
Condition: 1) Configure Port Channel 2) Enable RA Guard 3) Router Advertisement packets will still be accepted	
Workaround: Shut/no shut Port Channel	

Defect ID: DEFECT000546936	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS5.0.1	**Technology:** OpenStack Integration
Symptom: show arp is not supported on interfaces instantiated in management VRF	
Condition: executing show arp on interfaces instantiated in management VRF	

Defect ID: DEFECT000549648	**Technical Severity:** High
Reason Code: Already Fixed in Release	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** Software Installation & Upgrade
Symptom: When doing a firmware downgrade from 5.x to 4.x, the messages appearing on the screen are requested to be re-worded for easier understanding.	
Condition When doing a firmware downgrade from 5.x to 4.x.	

Defect ID: DEFECT000552701	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: In a Remote Logical SAN configuration, after a FCoE device logout, later when the same FCOE device try to login, occasionally the login may fail.	
Condition Shutting the FCOE interface for a Remote Logical SAN login can result in a FIP Clear Virtual Link (CVL) to be sent from FCF through FIF to the FCoE CNA.. When FCoE CNA receives CVL, FCoE CNA normally will initiate the FIP Discovery and then follow with FLOGI. But sometime FCoE CNA does not initiate FIP Discovery but sends FLOGI directly which then causes login failure.	
Recovery: Toggle the interface port connected to the FCoE Converged Network Adapter(CNA) to re-initiate the FIP Discovery.	

Defect ID: DEFECT000553496	**Technical Severity:** Medium
Reason Code: Already Fixed in Release	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** ICMP - Internet Control Message Protocol
Symptom: When numeric option is not specified in the ping command, the output does not display the host name in a string format.	
Condition When numeric option is not specified in the pingcommand.	

Defect ID: DEFECT000553915	**Technical Severity:** High
Reason Code: Feature/Function Not Supported	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.0.0	**Technology:** Configuration Fundamentals
Symptom: Command "supportsave" does not support TFTP protocol.	
Condition Specifying the TFTP as transfer protocol isn't allowed.	
Workaround: To invoke command "supportsave", valid transfer protocol values are: File transfer protocol (FTP), Secure copy (SCP), or Secure FTP (SFTP).	

Defect ID: DEFECT000554573	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: When two hosts with same WWN login, the switch can undergo a series on unexpected reboots. The Duplicate WWN feature is currently not supported on NOS AG.	
Condition Two hosts with the same WWN try to login	
Workaround: Please make sure there are no two hosts with same WWN.	

Defect ID: DEFECT000555171	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Hardware Monitoring
Symptom: VDX 40Gb port may go administratively down after reboot, or reboot of partner VDX switch.	
Condition When VDX switch is rebooted, or partner VDX switch is rebooted.	
Workaround: Shut/no shut the port	
Recovery: Shut/no shut the port	

Defect ID: DEFECT000556146	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.1	**Technology:** AMPP - Automatic Migration of Port Profiles
Symptom: Shut on one of the members of the VLAG makes the source MAC addresses learnt on other members of the VLAG disappear.	
Condition Conflicting type of traffic is received simultaneously on multiple links of VLAG	
Recovery: No port-profile-port followed by port-profile-port.	

Defect ID: DEFECT000556823	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** BFD - BiDirectional Forwarding Detection
Symptom: static route BFD session is not coming UP.	
Condition When User configures Static route single/Multi BFD with neighbor who has this neighbor IP configured on remote interface. Topology where BFD packets will be sent on interface which does not have reachability to destination.	
Workaround: symmetric source/destination pair with static route BFD.	

Defect ID: DEFECT000557278	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Port Mirroring
Symptom: MAPS RASLOGs for RX_SYM - RX Symbol Errors - are seen when PO is "shut"	
Condition MAPS needs to be enabled and PO needs to be shut.	

Defect ID: DEFECT000557518	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Virtual Fabrics
Symptom: In virtual-fabric enabled mode, port-profile-port configuration on an interface can take upto5 minutes only the first time.	
Condition If the system is in virtual-fabric enabled mode, first time configuration of "port-profile-port" or "port-profile-port domain <domainName>" , provided the port-profile-domain has "switchport trunk allowed vlan all" config	
Workaround: No workaround	

Defect ID: DEFECT000558057	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Port Mirroring
Symptom: When chassis disable command is executed on VDX 6940-144S the following Fabric watch RASLOG's may be observed: [FW-1038], 3196, SW/0 | Active, WARNING, sw0, Sfp RX power for port x/x/x, is below low boundary	
Condition There is no functionality loss. User can ignore the additional RASLOGs.	

Defect ID: DEFECT000558616	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS5.0.2	**Technology:** Hardware Monitoring
Symptom: When VDX switches uses DC power supply, the output of "show environment power" OR "show chassis" does not display the serial number.	
Condition When DC power supply is used with VDX switches.	

Defect ID: DEFECT000558687	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.1	**Technology:** VCS Fabric
Symptom: The following error is thrown: "%%Error: Platform hardware limitation or resource limitreached."	
Condition In Fabric Cluster mode, when 4K vlans are configured and user tries to configure vlans beyond 4K, the above error is thrown. However, there are no issues till 4K vlans are configured.	
Recovery: Reload the system when this issue isobserved .	

Defect ID: DEFECT000558898	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS4.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: Primary port of a VLAG may send back the BUM traffic to same VLAG member.	
Condition When the “no channel-group” command is issued to the VLAG member that is the last port in a port-group	
Workaround: Re-configuring the port-channel on all member ports of a VLAG	

Defect ID: DEFECT000559516	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: Extended ACL rules with protocol IPv6 does not take effect on traffic on management interface	
Condition Extended ACL rule with protocol as IPv6	
Workaround: Create two separate rules for protocols TCP and UDP instead of a single IPv6 rule	

Defect ID: DEFECT000559589	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Hardware Monitoring
Symptom: 10Gb port may be stuck in offline state.	
Condition After power cycle or reboot of VDX switch, 10Gb port may be stuck in offline state.	
Recovery: Shut/no shut the port.	

Defect ID: DEFECT000559631	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: FCoE login does not work	
Condition On changing the LAG config from local SAN to remote SAN on LAG.	

Defect ID: DEFECT000559741	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Access Gateway
Symptom: FCoE logins won't come through LAG.	
Condition Changing LAG from Local SAN to Remote SAN.	

Defect ID: DEFECT000559806	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.1	**Technology:** OpenStack Integration
Symptom: "show openflow flow" does not display Flowids in sorted order.	
Condition User tries to view flow-mods installed in the switch.	

Defect ID: DEFECT000559907	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: Copying configuration from file takes long time.	
Condition The configuration file contains overlay-gateway with more than 1000 "map vlan" and 500 "site" configurations.	

Defect ID: DEFECT000560127	**Technical Severity:** High
Reason Code: Feature/Function Not Supported	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS5.0.1	**Technology:** OpenStack Integration
Symptom: Traffic loss can be observed between linecards going from one LC to another.	
Condition Traffic loss can be observed between linecards [Ex: LC # 2] with another linecard [LC #1] being installed.	

Defect ID: DEFECT000560280	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** Multi-VRF
Symptom: Under conditions of VRF scale and operation, sometimes Address Family may not be instantiated on creating a VRF	
Condition Configuring a VRF in complex network topologies	

Defect ID: DEFECT000560607	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP client binding does not happen on one of the VDX 6740s in the cluster, when clients are configured on a tagged vlan. DHCP Relay statistics show that discover packets are not received at the relay.	
Condition DHCP Clients configured on tagged vlan.	

Defect ID: DEFECT000560620	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** IP Addressing
Symptom: Ping fail between VDX 6940 and VDX 8770 .	
Condition Executing ping command	
Workaround: Atleast one L3 interface needs to be enabled in a VRF for a route to be added in the kernel or stack	

Defect ID: DEFECT000560681	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: FCoE Host in a Local Logical SAN may not be able to see the FC Target when VDX 6740 where the FC Target is connected to, is dynamically changed from the Remote Logical SAN to the Local Logical SAN configuration.	
Condition When a VDX6740 is reconfigured from a Remote Logical SAN to a Local Logical SAN, FC Target connected through the VDX6740 will become un-accessible in VCS.	
Workaround: Reboot the VDX6740 that has been moving from Remote Logical SAN to Local Logical SAN.	
Recovery: To prevent hitting this issue, user can do delete fcf-group, change VDX6740 from AG mode to non-AG mode, and reboot VDX6740. Then proceed on modifying the fabric-map configuration from Remote Logical SAN to Local Logical SAN.	

Defect ID: DEFECT000560802	**Technical Severity:** Medium
Reason Code: Already Fixed in Release	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** FCoE - Fibre Channel over Ethernet
Symptom: The switch may accept the new fibrechannel or FCoE logins even if there is Duplicate WWN detection. This fabric login policy for Duplicate WWN is not supported in NOS6.0.0.	
Condition When the firmware downgrade from NOS6.0.1 to NOS6.0.0 happens with the non-default fabric login policy i.e new-login for DuplicateWWN is configured on the switch.	
Workaround: move the fabric login policy to old-login for Duplicate WWN (default mode) and then downgrade the switch to NOS6.0.0	

Defect ID: DEFECT000560931	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** OSPF - IPv4 Open Shortest Path First
Symptom: Delay in OSPF session establishment in VDX6940 144S on IP address delete and reconfigure.	
Condition Running OSPF after IP address reconfiguration	

Defect ID: DEFECT000561037	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.1	**Technology:** OpenStack Integration
Symptom: Openflow related hardware resources may not get de-allocated in a clean manner resulting in inconsistent behavior in data-path forwarding. New configurations may also fail to get programmed and failures will not be reported back to controller.	
Condition If ISSU was performed with controller driven flows installed in the system.	
Workaround: Controller needs to be disconnected first, clear all installed flows using "clear openflow all" command and then trigger ISSU.	

Defect ID: DEFECT000561046	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: Internal Server Error is returned as the status even though the intended functionality works fine when the rbridge-id is changed through REST request.	
Condition When the rbridge-id is modified through REST request.	
Workaround: Modify the rbridge-id through CLI.	

Defect ID: DEFECT000561179	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS6.0.1	**Technology:** QoS - Quality of Service
Symptom: Modular QoS CLI (MQC) policy config will be lost on a port in downgrade from NOS6.0.1 to NOS5.0.1 if MQC has shaper config on lossless priority.	
Condition Per port scheduler configuration (i.e MQC policy having scheduler and shaper configuration applied in egress direction on a port).	
Workaround: Before downgrading to NOS5.0.1, make sure that active MQC policy should not have shaper config on lossless priority.	
Recovery: Remove shaper config lossless priority in MQC policy.	

Defect ID: DEFECT000561210	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: Unsupported features are getting configured on 100 MB interface if the link comes up with 100 MB using autonegotiation. Please check the NOS 6.0.1 Release Notes for features that are supported with 100MB	
Condition When speed is configured as an 'auto' on an interface which is connected to an 100 MB peer link.	

Defect ID: DEFECT000561274	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** Static Routing (IPv4)
Symptom: Under conditions of scale (512 VRFs), few routes may go missing in VDX 6940-144S after an ISSU upgrade.	
Condition ISSU in VDX 6940 144S	

Defect ID: DEFECT000561304	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.1	**Technology:** VCS Fabric
Symptom: VDX 6940 platforms in FC mode may undergo an unexpected reload during firmware upgrade or downgrade operations if there is security related configuration - tacacs server,radius sever ip, aaa	
Condition Switch is in FC mode, Security Configuration (tacacs server,radius sever ip,aaa) is present and Switch undergoes a firmware upgrade/downgrade.	
Workaround: Remove " tacacs server,radius sever ip,aaa" on all the vcs nodes.	
Recovery: Powercycle will recover the Switch.	

Defect ID: DEFECT000561506	**Technical Severity:** High
Reason Code: Already Fixed in Release	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: VDX6940 and VDX6740 may take longer to boot up in 6.0.1 than in 6.0.0. (the delay is about 1 minute with default configuration).	
Condition When the reload command is executed.	

Defect ID: DEFECT000561605	**Technical Severity:** High
Reason Code: Already Fixed in Release	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.1.3	**Technology:** NTP - Network Time Protocol
Symptom: tsd module termination and VDX abrupt reload	
Condition Fail to reach NTP server due to network reachability issue	
Recovery: Fix NTP server reachability issue.	

Defect ID: DEFECT000561706	**Technical Severity:** High
Reason Code: Already Fixed in Release	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.1	**Technology:** VLAN - Virtual LAN
Symptom: If static MAC is configured for mulicast MAC address, the packets destined to the MAC address will be flooded to all member ports in VLAN	
Condition Configuring a static muticast MAC address will result in this issue.	

Defect ID: DEFECT000561713	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: When snmpwalk operation is done on TCP MIB (RFC 4022), the operation may become very slow and may experience timeouts.	
Condition This issue is seen when snmpwalk is done only on TCP MIB.	
Workaround: The snmpwalk should be done with a timeout set to at least 3 seconds. This will help in avoiding the timeout during the snmpwalk operation on the TCP MIB.	

Defect ID: DEFECT000562896	**Technical Severity:** Medium
Reason Code: Feature/Function Not Supported	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.1	**Technology:** Logical Chassis
Symptom: "no fabric isl enable" and " no fabric trunk enable" configurations on ports 0/33 to 0/48 got reverted back to "fabric isl enable" and " fabric trunk enable"	
Condition Upgrade from no4.1.2a1 to nos5.0.1d	
Workaround: 1. Before upgrade, save the running config to a file 2. Upgrade to 5.0.1d 3. Copy default to startup 4. Replay saved config from the file	
Recovery: After upgrade, manually configure "fabric isl enable" and " fabric trunk enable" on ports 0/33 to 0/48	

Defect ID: DEFECT000562941	**Technical Severity:** Medium
Reason Code: Feature/Function Not Supported	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.2	**Technology:** AAA - Authentication, Authorization, and Accounting
Symptom: when upgrading from 4.1.2a to 5.0.2 with tacacs-server config	
Condition This happens only upgrades from 4.x to 5.x with tacacs-server configuration	
Workaround: copy the config before upgrade, upgrade with default-config and replay the config	

Defect ID: DEFECT000563273	**Technical Severity:** High
Reason Code: Feature/Function Not Supported	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.1.3	**Technology:** Configuration Fundamentals
Symptom: High CPU usage is seen on 6730 platform.	
Condition High memory usage due to memory leak which eventually result in CPU stuck at 100%.	
Recovery: 1. Login into Switch as root 2. Capture the “top” output 3. Run: echo 10240 > /proc/sys/vm/min_free_kbytes 4. Issue command: sysctl -q -p /etc/sysctl.conf	

Defect ID: DEFECT000563673	**Technical Severity:** High
Reason Code: Feature/Function Not Supported	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS3.0.1	**Technology:** Configuration Fundamentals
Symptom: We will observe unexpected reload of VDX 6720 platform switch.	
Condition: In 6720 platform switch, when packet size greater than 1536 is received on the management port, then we will observe this issue.	
Workaround: Setting MTU size of eth0 as 9512 will avoid this issue. But it should be set after every reload of switch, as MTU change is not persisted after reload.	

Defect ID: DEFECT000564101	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.0	**Technology:** Security Vulnerability
Symptom: Using restrict_ssh script on VDX 6740x platforms will not do a GOS sync for the SSHd config file.	
Condition: Using restrict_ssh on VDX 6740s platforms will have a GOS sync issue.	

Defect ID: DEFECT000564347	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Traffic Management
Reported In Release: NOS5.0.2	**Technology:** Rate Limiting and Shaping
Symptom: After upgrade to 5.0.2a-bld-02, we may observe some traffic forwarding issue with 1Gigabit connections on 8770 platform	
Condition: When we have 8770 platform with 1Gigabit connections, we will observe this issue.	

Defect ID: DEFECT000564498	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS4.1.3	**Technology:** Port Mirroring
Symptom: "show interface status" command shows incorrect status for internal VDX ports in the switch.	
Condition: If we have internal embedded ports in the VDX switch, then we will observe this issue.	

Defect ID: DEFECT000565415	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Added new enums in Fruclass to indicate the slotnames correctly 87 sfm(12), 88 lineCard(13), 89 managementModule(14) 90 }	
Condition: The slotnames are not insync with CLI slotnames	

Defect ID: DEFECT000565954	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.1	**Technology:** Logical Chassis
Symptom: Principal node of a logical cluster will not reach 'ready' state to accept any configuration commands. When any configuration command is issued, it displays error as 'Cluster formation is in progress'.	
Condition Cluster reformation is triggered due to node joining or leaving the fabric or due to firmware upgrade. Cluster formation might not complete due to underlying communication issues.	
Recovery: Remove the node from the cluster either by issuing reload command or by issuing 'chassis disable' command. This will trigger some other node in the cluster to become principal. Once the cluster is stable, bring back this switch to the cluster by issuing 'chassis enable'.	

Defect ID: DEFECT000569750	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** Software Installation & Upgrade
Symptom: When upgrading VDX switches from 5.0.2x to 6.0.1x, certain rules configured for the RBAC feature, would change to a different set of commands.	
Condition When upgrading switches from 5.0.x to 6.0.x & RBAC rules are configured.	
Workaround: It is recommended to take a backup of the rules before the upgrade to 6.x from 5.x so that it can be restored after upgrade.	

Defect ID: DEFECT000577381	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS4.1.3	**Technology:** VRRPv3 - Virtual Router Redundancy Protocol Version 3
Symptom: Under rare conditions, certain hosts lose IP connectivity with selected few devices when hosts connected on VDX6710/20/30 configured as Layer-3 gateways for the host.	
Condition Usually when the set of MAC addresses are toggled between one interface to another on the switch.	
Recovery: Clearing the ARP on the VDX switches restores the connectivity.	

Defect ID: DEFECT000578967	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS6.0.2	**Technology:** PIM - Protocol-Independent Multicast
Symptom: On rare occassion, there might be stale S,G entries seen on starting the multicast traffic	
Condition Not known	
Workaround: Not known	
Recovery: Clear ip pim mcache cleanup the stale entries	

Defect ID: DEFECT000585841	**Technical Severity:** High
Reason Code: Already Fixed in Release	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: Customer uses pVLAG between VDX 6740 and third party vendor switch for underlay of L2 extension tunnel. After pVLAG member PO's go through fail over, the L2 extension traffic drops for upto 10 seconds.	
Condition The issue occurs when the primary member port-channel of a port-channel group goes through a fail over. Under such circumstances, a 10 second tunnel traffic is lost.	

Known Issues for Network OS v7.0.2c

This section lists open software defects with Critical, High, and Medium Technical Severity as of July, 2019 in Network OS v7.0.2c.

Known Issues for Network OS v7.0.2

This section lists open software defects with Critical, High, and Medium Technical Severity as of November 20, 2017 in Network OS v7.0.2.

Defect ID: DEFECT000554319	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Switch does not generate a ColdStart trap on the VE interface configured in mgmt-vrf.	
Condition: When switch is configured with VE interface in mgmt-vrf, then we will observe this issue.	

Defect ID: DEFECT000619146	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: ISSU upgrade from 7.0.1 to NOS7.0.1a can cause some traffic loss if BFD is configured.	
Condition: ISSU upgrade to NOS7.0.1a when BFD is configured	
Workaround: BFD can be disabled during upgrade.	

Defect ID: DEFECT000646180	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	**Technology:** IPv6 Addressing
Symptom: Unexpected reload	
Condition: Bulk (L3anycast) configuration through NETCONF	
Workaround: Single configuration in one query should be done.	

Defect ID: DEFECT000646908	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.1.0	**Technology:** Hardware Monitoring
Symptom: The source IP for SNMP traps is not deterministic.	
Condition: When VCS virtual IP address is configured and SNMP traps are enabled.	

Defect ID: DEFECT000647847	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** LAG - Link Aggregation Group
Symptom: Unexpected reload	
Condition: In rare a case, DB corruption happens at the time of port-channel deletion.	

Defect ID: DEFECT000648164	
Technical Severity: Medium	**Probability:** Medium

Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: SNMP responding on VCS IPv6 instead of management IPv6 address.	
Condition: Both MM/Chassis IPv6 and virtual vcs IPv6 addresses are configured.	
Workaround: Have the management IPv6 configured latest	

Known Issues for Network OS v7.0.1c

This section lists open software defects with Critical, High, and Medium Technical Severity in Network OS v7.0.1c.

Defect ID: DEFECT000619425	
Technical Severity: High	Probabilitty: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported in Release: NOS7.1.0	Technology: VXLAN - Virtual Extensible LAN
Symptoms: Traffic loss on Port-channel interface.	
Condition: If Global MTU is smaller than Port-channel MTU or Global MTU is configured and un-configured, user may see traffic loss on port-channel interface.	
Workaround: Configure MTU same as port-channel on Port-channel member interfaces.	

Defect ID: DEFECT000641475	
Technical Severity: High	Probabilitty: Low
Product: Extreme Network OS	Technology Group: Security
Reported in Release: NOS7.0.1	Technology: User Accounts & Passwords
Symptoms: Configuration of invalid encrypted password for existing user with encryption level as 7 it is getting accepted without throwing error.	
Condition: VDX switch allows to change password as invalid encrypted password for existing user.	

Known Issues for Network OS v7.0.1b

This section lists open software defects with Critical, High, and Medium Technical Severity as of February 17th, 2017 in Network OS v7.0.1b.

Defect ID: DEFECT000612967	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported in Release: NOS7.1.0	Technology: Security Vulnerability
Symptoms: Shutting down SSH server does not close all existing SSH login sessions	
Condition: Shutdown SSH server	

Defect ID: DEFECT000619146	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported in Release: NOS7.0.1	Technology: IP Fabric
Symptoms: ISSU upgrade to NOS7.0.1a can cause 19 seconds of traffic loss if BFD is configured.	
Condition: ISSU upgrade to NOS7.0.1a when BFD is configured	
Workaround: Please disable BFD.	

Defect ID: DEFECT000628176	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported in Release: NOS7.0.1	Technology: VLAN - Virtual LAN
Symptoms: Packets are flooded to all the vlan member interfaces of the remote node even though Static-mac blinding is configured.	
Condition: When Multicast mac address is used and Vlan member interfaces are present in the remote nodes of the cluster	

Defect ID: DEFECT000629678	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported in Release: NOS7.0.1	Technology: DHCP - Dynamic Host Configuration Protocol
Symptoms: IPV4 DHCP relay statistics does not display the count for DHCP offer and DHCP Ack packets relayed by switch from the DHCP server to DHCP client	
Condition: when DHCP client and DHCP server are present in different VRF	

Defect ID: DEFECT000626331	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported in Release: NOS7.0.1	Technology: VLAN - Virtual LAN
Symptoms: User configured Vlan names are not displayed after reload of cluster in "show vlan br". It changes to default Vlan name	
Condition: Execution of "show vlan brief" CLI after reload.	

Defect ID: DEFECT000627922	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Hardware Monitoring
Reported in Release: NOS7.0.1	Technology: Monitoring
Symptoms: Multiple FFDC and Core files seen on firmware downgrade.	

Condition: When 8770-8 chassis running six new version line-cards 6x100G, 27x40G, 48x10G-T with three or more of type 6x100G is downgraded from 7.0.1b version to lower version.

Defect ID: DEFECT000630802	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported in Release: NOS7.0.1	Technology: OSPFv3 - IPv6 Open Shortest Path First
Symptoms: After doing HA failover on M4, HA is not in SYNC and observed ONM carsh.	
Condition: HA failover and timing issue can trigger the issue, very rare to occur.	

Defect ID: DEFECT000596775	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported in Release: NOS7.0.1	Technology: OSPFv3 - IPv6 Open Shortest Path First
Symptoms: When the user configures IPv6 RA interval with the default value 600, the running-config shows the default RA value without suppressing it.	
Condition: The issue is seen by the user every time the RA interval is configured with the default value.	

Known Issues for Network OS v7.0.1a

This section lists open software defects with Critical, High, and Medium Technical Severity as of September 9th, 2016 in Network OS v7.0.1a.

None

Known Issues for Network OS v7.0.1

This section lists open software defects with Critical, High, and Medium Technical Severity as of May 25, 2016 in Network OS v7.0.1.

Defect ID: DEFECT000590114	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.1	**Technology:** AMPP - Automatic Migration of Port Profiles
Symptom: If user configures two AMPP port-profiles, one of them configured with access VLAN x and other configured with trunk VLAN x, Then ,In that case , It will not be shown as conflicting in "show port-profile name <pp1-name> name <pp2-name> validate" command output.	
Condition When user creates 2 port-profiles, one port-profile with access VLAN x and other port-profile with trunk VLAN x and executes "show port-profile name <pp1-name> name <pp2-name> validate" CLI.	

Defect ID: DEFECT000591616	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: Switch goes for an unexpected reload with the REST request.	
Condition When the switch is pounded with the REST requests from multiple concurrent sessions simultaneously and continuously over a long period of time.	
Workaround: As far as possible, send REST requests to configure the switch from one session only. Multiple sessions can be used for retrieving information from the switch with GET requests.	

Defect ID: DEFECT000592879	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Configuration Fundamentals
Symptom: After LC power on/off in VDX8770, uplink interfaces from the LC are missing on show track summary output.	
Condition When Link State Tracking (LST) configuration is present on a linecard, after slot power off/on the uplink configuration will be lost.	
Workaround: Uplinks need to be reconfigured again after slot power on.	

Defect ID: DEFECT000594276	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.1	**Technology:** Logical Chassis
Symptom: Under a high scale of VCS nodes, the configuration applied for a range of interfaces across the VDX nodes may cause principal node to encounter an unexpected reload.	
Condition When issuing a configuration command under an interface range in a large cluster (32+ nodes)	
Workaround: Avoid using interface range option in large clusters (32+ nodes) & instead configure the interfaces individually.	

Defect ID: DEFECT000594793	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Software Installation & Upgrade
Symptom: System may display:	

	"qman_recovery_exit_local: DEBUG: the FQID 516 has dest_wq as chaqman_recovery_exit_local: DEBUG: the WQ lengths for pool channel of portal 1 on cpu1 are: 0:0:0:0:0:0:0:0"
Condition	This bug appears when partitions are switched with heavy traffic.
Recovery:	Reboot the system.

Defect ID:	DEFECT000595709
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.1	**Technology:** Logical Chassis
Symptom:	System reloads on VDX8770.
Condition	This occurs with 512 or more VRRP sessions enabled and "debug vrrp packets" is turned on.
Workaround:	"debug vrrp packets" should not be turned on in a scaled environment.

Defect ID:	DEFECT000596280
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** IP Addressing
Symptom:	Unable to delete an ACL.
Condition	When ACL is associated to the management interface of one or more switches in the VCS and the switch gets removed from VCS.

Defect ID:	DEFECT000596480
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** Configuration Fundamentals
Symptom:	On execution of CLI "track remove all" complete Link State Tracking (LST) configuration should get removed from a port. In case of port-channel interface protocol daemon is not clearing the LST configuration hence it is displayed in output of show command.
Condition	Execution of "track remove all" CLI for a port-channel interface for which Link State Tracking (LST) configuration is present.
Workaround:	As a workaround user can remove the configuration one by one by executing respective 'no' CLIs.

Defect ID:	DEFECT000596868
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom:	The global MTU value cannot be deleted through REST API.
Condition	Issue happens when the user tries to delete the global mtu using the DELETE request through the REST interface.
Workaround:	Using the PATCH request with the default value as a work around. The effect of this is same as deleting the config.

Defect ID:	DEFECT000598965
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom:	Local configuration related to global configuration may not restore on "config snapshot restore".
Condition	Customer using snapshot feature may see issues when running "attached rbridge-id add 1" missing from running-config after "vcs config snapshot restore rbridge-id <rb-id> snapshot-id <snapshot-id> .
Workaround:	Customer should configure the missing configurations again.

Defect ID: DEFECT000598972	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: Switch might go for an unexpected reload when any configuration update is performed on a range of interfaces.	
Condition: On a large cluster with scaled up configurations, performing any configuration on a range of interfaces by entering into interface range sub-mode might cause switch to run out of memory and thereby causing it to reload.	
Workaround: Required configuration update can be made on individual interfaces one at a time instead of performing it on a range of interfaces. Configuration update on multiple interfaces can stillbe performed by using comma (,) as separators instead of hyphen (-) when specifying the range. For ex, to shutdown interfaces 1 to 5, use "interface te 1,2,3,4,5" instead of "interface te1-5".	

Defect ID: DEFECT000599203	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: The SNMP IPV4 traps may not be received through in-band interface.	
Condition: The SNMP traps may not be received through in-band interface after upgrade from 6.0.1 to 7.0.1.	
Workaround: Configure source-interface in the SNMP host / v3host recipients.	

Defect ID: DEFECT000599289	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.1	**Technology:** ACLs - Access Control Lists
Symptom: Applying Access Control List (ACL's) with 12K rules on management interface takes more than 3 minutes to enforce it.	
Condition: When Access Control List (ACL's) is configured with 12K rules.	

Defect ID: DEFECT000599778	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.1	**Technology:** TACACS & TACACS+
Symptom: LDAP/RADIUS/TACACS+ server configurations are not displayed in the same order in which they were added.	
Condition: 1. Configure multiple TACACS+/RADIUS/LDAP servers(max 5) 2. Remove few server entries 3. Add those servers entries back	
Workaround: Remove all Server entries and configure those servers back in the desired order.	

Defect ID: DEFECT000600022	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.1	**Technology:** Metro VCS
Symptom: When VDX 8770 is in chassis-disabled state, the far-end 100 GbE link partners using QSFP28 optics may see intermittent link flaps. After VDX 8770 is chassis-enabled, there is a low probability that the 100 GbE port may not come online.	
Condition: When VDX 8770 is in chassis-disabled state, the far-end 100 GbE link partners using QSFP28 optics may see intermittent link flaps. After VDX 8770 is chassis-enabled, there is a low probability that the 100 GbE port may not come online.	

Recovery:	Execute "shut" on 100 GbE link partner port connected to VDX 8770 to stop the port from flapping intermittently. After the VDX 8770 is chassis-enabled, execute "no shut" on the 100 GbE link partner to re-enable the port.

Defect ID:	DEFECT000600057	
Technical Severity:	High	**Probability:** Low
Product:	Extreme Network OS	**Technology Group:** VCS
Reported In Release:	NOS7.0.1	**Technology:** Logical Chassis
Symptom:	Switch might not rejoin the cluster when reloaded using 'fastboot' command.	
Condition	Reloading switch using 'fastboot' command on VDX6940 and VDX6740 platforms when SW1 partition is active might lead to this issue.	
Workaround:	Reload the switch using 'reload' command which is more graceful way of reloading.	
Recovery:	Bring the switch which failed to join the cluster to default configuration using command 'copy default-config startup-config'. On reload, switch rejoins the cluster and regains older configuration.	

Defect ID:	DEFECT000600066	
Technical Severity:	High	**Probability:** High
Product:	Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.1	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom:	DHCP IPv4 Relay forwarded DISCOVER packet is not getting forwarded through remote leaf node in BGP-EVPN IP Fabric.	
Condition	While deploying DHCP Relay in BGP-EVPN IP Fabric.	
Recovery:	Disable “conversational-arp".	

Defect ID:	DEFECT000600169	
Technical Severity:	High	**Probability:** High
Product:	Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release:	NOS7.0.1	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom:	IP MTU configuration is not working for VE interface when IP address or L3 VNI association is not present.	
Condition	When IP MTU is configured, it is not applied on the VE interface.	
Workaround:	Configure IP MTU followed by the configuration of the IP address.	

Defect ID:	DEFECT000600185	
Technical Severity:	Medium	**Probability:** High
Product:	Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release:	NOS7.0.1	**Technology:** OpenStack Integration
Symptom:	When VDX-8770 is in chassis-disable state, the "show media" command will not show 100 GbE ports.	
Condition	When VDX-8770 is in chassis-disable state, the "show media" command will not show 100 GbE ports.	
Workaround:	After the chassis is enabled using "chassis enable" command, "show media" will show the 100 GbE ports.	
Recovery:	After the chassis is enabled using "chassis enable" command, "show media" will show the 100 GbE ports.	

Defect ID:	DEFECT000600197	
Technical Severity:	Medium	**Probability:** High
Product:	Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer

Reported In Release: NOS7.0.1	**Technology:** BGP4 - IPv4 Border Gateway Protocol
Symptom: "show running-config overlay-gateway <name> vlan <vlan number>" throws error "% No entries found " even when VLAN is present.	
Condition This happens only when a filter is specified after "overlay-gateway <name>". Otherwise command works fine when no filter is specified.	
Workaround: Instead of using the filter, use the " | include <string>" for filters like following: "show running-config overlay-gateway <name> | include "vlan <vlan-number>"	

Defect ID: DEFECT000600230	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS7.0.1	**Technology:** IP Fabric
Symptom: "show running-config rbridge-id evpn-instance <vni -name> vni add <vni-range>" throws an error message.	
Condition Customer doing show running configuration with VNI range in EVPN instance.	
Workaround: Use the following command: "show running-config rbridge-id evpn-instance vni add" .	

Defect ID: DEFECT000600377	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: The SNMP walk may fail and SNMPV3 trap may not be received for the user configured under RBridge.	
Condition The SNMP walk may fail and SNMPV3 trap may not be received only for the SNMPV3 user configured under rbridge after upgrade from 7.0.0 to 7.0.1.	
Recovery: Reconfigure the user under RBridge after the successful upgrade from 7.0.0 to 7.0.1.	

Defect ID: DEFECT000600385	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VLAN - Virtual LAN
Symptom: Duplicate ARP entries are observed.	
Condition This can happen after an ISSU upgrade and a new IP address is alllocated via DHCP for a connected host.	
Workaround: Execute "clear arp ip <IP address>" for the old IP address of host.	
Recovery: Execute "clear arp ip <IP address>" for the old IP address of host.	

Defect ID: DEFECT000600591	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface
Symptom: Logs are dumped on the screen, when there is a read failure on SFPs connected to the port.	
Condition Accessing information about the SFPs inserted in the ports.	
Recovery: Disable the port and re-enable it.	