Network OS 6.0.2h for Extreme VDX

Release Notes v4.0

Contents

DOCUMENT HISTORY

Version	Summary of Changes	Publication Date
1.0	Initial Release	April 14th, 2018
2.0	Update to Upgrade Matrix	May 21st, 2018
3.0	Update on defect with code change	November 12, 2018
4.0	Updates to Software Update process	November 2019

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 employees, but it 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:

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 http://www.extremenetworks.com/documentation-feedback-pdf/
- Email us at internalinfodev@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.

- 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

OVERVIEW

Network OS v6.0.2h introduces new features.

- None

Network OS v6.0.2g introduces new features.

- None

Network OS v6.0.2f introduces new features.

- None

Network OS v6.0.2e introduces new features.

- None

Network OS v6.0.2d introduces new features.

- None

Network OS v6.0.2c introduces new features.

- None

Network OS v6.0.2b introduces new features.

- None

Network OS v6.0.2a introduces new features.

- DHCP option-82 with text string support

Following table lists the set of new features and additional information is available below for each feature:

- None

Network OS v6.0.2b also introduces the following storage features:

- None

WHAT IS NEW IN THIS RELEASE

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

Optics	Description
40G-QSFP-SR-BIDI	40GbE QSFP+ BiDi 100m

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 v6.0.2h

- None

New software features for Network OS v6.0.2g

- None

New software features for Network OS v6.0.2f

- None

New software features for Network OS v6.0.2e

- None

New software features for Network OS v6.0.2d

- None

New software features for Network OS v6.0.2c

- None

Modified software features for Network OS v6.0.2c

The following software features have been modified in this release:

- None

Deprecated software features for Network OS v6.0.2c

The following software features are deprecated beginning in this release:

- None

New software features for Network OS v6.0.2b

- None

Modified software features for Network OS v6.0.2b

The following software features have been modified in this release:

- None

Deprecated software features for Network OS v6.0.2b

The following software features are deprecated beginning in this release:

- None

New software features for Network OS v6.0.2a

- DHCP option-82 with text string support

Modified software features for Network OS v6.0.2a

The following software features have been modified in this release:

- Enhancements in SNMP MIB functionality
- Support for configuring source-interface for SNMP traps & additional traps
- FCOE source interface improvement.

Modified software features for Network OS v6.0.2b

The following software features have been modified in this release:

- Enhancement to optimize the fan-speed of the VDX6740T & VDX6740T-1G

Deprecated software features for Network OS v6.0.2a

The following software features are deprecated beginning in this release:

- None

CLI Changes

The following section lists new, modified, and deprecated commands for this release:

New Commands for Network OS v6.0.2h

The following commands are new in this patch release:

- None

Deprecated commands for Network OS v6.0.2h

The following commands have been deprecated beginning with this release:

- show process cpu history

Note: This command is deprecated on the VDX6740, VDX6740T, and VDX6740T-1G only.

New Commands for Network OS v6.0.2g

The following commands are new in this patch release:

- None

New Commands for Network OS v6.0.2f

The following commands are new in this patch release:

- None

New Commands for Network OS v6.0.2e

The following commands are new in this patch release:

- None

New Commands for Network OS v6.0.2d

The following commands are new in this patch release:

- None

New Commands for Network OS v6.0.2c

The following commands are new in this patch release:

- None

New Commands for Network OS v6.0.2a

The following commands are new in this patch release:

- ip dhcp relay information option
- show ip dhcp relay address interface

Modified Commands for Network OS v6.0.2c

The following commands have been modified for this patch release:

- None

Deprecated commands for Network OS v6.0.2c

The following commands have been deprecated beginning with this release:

- None

New Commands for Network OS v6.0.2b

The following commands are new in this patch release:

- [no] snmp-server three-tuple-if enable

This command enable/disable 2-tuple/3-tuple format of ifDescr and ifName objects

Modified Commands for Network OS v6.0.2g

The following commands have been modified for this patch release:

- None

Modified Commands for Network OS v6.0.2f

The following commands have been modified for this patch release:

- None

Modified Commands for Network OS v6.0.2e

The following commands have been modified for this patch release:

- None

Modified Commands for Network OS v6.0.2d

The following commands have been modified for this patch release:

- None

Modified Commands for Network OS v6.0.2c

The following commands have been modified for this patch release:

- None

Modified Commands for Network OS v6.0.2b

The following commands have been modified for this patch release:

- None

Modified Commands for Network OS v6.0.2a

The following commands have been modified for this patch release:

- fcoeport
- ipv6 vrrp-suppress-interface-ra
- mode (27x40 GbE line card)
- qos drop-monitor enable
- qos random-detect traffic-class
- qos red profile
- show qos red profile
- show qos red statistics interface
- snmp-server community
- storm-control ingress
- vcs replace rbridge-id

Deprecated commands for Network OS v6.0.2g

The following commands have been deprecated beginning with this release:

None

Deprecated commands for Network OS v6.0.2f

The following commands have been deprecated beginning with this release:

- None

Deprecated commands for Network OS v6.0.2e

The following commands have been deprecated beginning with this release:

- None

Deprecated commands for Network OS v6.0.2d

The following commands have been deprecated beginning with this release:

- None

Deprecated commands for Network OS v6.0.2c

The following commands have been deprecated beginning with this release:

- None

Deprecated commands for Network OS v6.0.2b

The following commands have been deprecated beginning with this release:

- None

Deprecated commands for Network OS v6.0.2a

The following commands have been deprecated beginning with this release:

- None

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:

- None

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

The VDX products conform to the following Ethernet standards:

- IEEE 802.1D   Spanning Tree Protocol
- IEEE 802.1s   Multiple Spanning Tree
- IEEE 802.1w   Rapid reconfiguration of Spanning Tree Protocol
- IEEE 802.3ad   Link Aggregation with LACP
- IEEE 802.3ae   10G Ethernet
- IEEE 802.1Q   VLAN Tagging
- IEEE 802.1p   Class of Service Prioritization and Tagging
- IEEE 802.1v   VLAN Classification by Protocol and Port
- IEEE 802.1AB   Link Layer Discovery Protocol (LLDP)
- IEEE 802.3x   Flow Control (Pause Frames)

The following draft versions of the Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) Standards are also supported on VDX products:

- IEEE 802.1Qbb   Priority-based Flow Control
- IEEE 802.1   DCB Capability Exchange Protocol (Proposed under the DCB Task Group of IEEE 802.1 Working Group)
- FC-BB-5   FCoE (Rev 2.0)

The VDX products conform to the following Internet IETF RFCs:

- RFC 2865   Remote Authentication Dial In User Service (RADIUS)
- RFC 1112   IGMP
- RFC 2236   IGMPv2
- RFC4601   PIM-SM
- RFC2131   DHCP
- RFC 2571   Architecture for Describing SNMP Framework
- RFC 3176   sFlow
- RFC 1157   SNMPv1/v2c
- RFC4510   Lightweight Directory Access Protocol (LDAP)
- RFC 3768   Virtual Router Redundancy Protocol (VRRP)
- RFC 2328   OSPF Version 2
- RFC 1587   OSPF NSSA Option
- RFC 3101   OSPF Not-So-Stubby-Area (NSSA) Option
- RFC 1765   OSPF Database Overflow
- RFC 2154   OSPF with Digital Signatures (MD-5 Support)
- RFC 3137   OSPF Stub Router advertisement
- RFC 2460   IPv6
- RFC 5340   OSPF for IPv6
- RFC 3623   Graceful OSPF Restart
- RFC 5187   OSPFv3 Graceful Restart (Helper Only)
- RFC 4271   A Border Gateway Protocol 4 (BGP-4)
- RFC 1745   BGP – OSPF Interactions
- RFC 1997   BGP Communities Attributes
- RFC 2385   Protection of BGP Sessions via the TCP MD5 Signature options
- RFC 2439   BGP Route Flap Dampening
- RFC 4456   BGP Route Reflection
- RFC 5492   Capabilities Advertisement with BGP-4
- RFC 3065   Autonomous System Confederations for BGP
- RFC 2858   Multiprotocol Extensions for BGP-4
- RFC 2918   Route Refresh Capability for BGP-4
- RFC 4724   Graceful Restart Mechanism for BGP
- RFC 5798   VRRP Version 3 for IPv4 and IPv6
- RFC 4541   MLDv1 Snooping

- RFC 6987   OSPFv3 Stub Router Advertisement (R-bit in Router LSA not supported)

The VDX 6740x, VDX 2740 and VDX 2746 products conform to the following Fibre Channel standards:

- • FC-GS-5 ANSI INCITS 427:2007 (includes the following)
    - FC-GS-4 ANSI INCITS 387: 2004
- • FC-SP-2 INCITS 496-2012 (AUTH-A, AUTH-B1 only)
- • FC-DA INCITS TR-36: 2004 (includes the following)
    - FC-FLA INCITS TR-20: 1998
    - FC-PLDA INCIT S TR-19: 1998
- • FC-MI-2 ANSI/INCITS TR-39-2005
- • FC-PI INCITS 352: 2002
- • FC-PI-2 INCITS 404: 2005
- • FC-PI-4 INCITS 1647-D, revision 7.1 (under development)
- • FC-FS-2 ANSI/INCITS 424:2006 (includes the following)
    - FC-FS INCITS 373: 2003
- • FC-LS INCITS 433: 2007
- • MIB-FA INCITS TR-32: 2003

HARDWARE SUPPORT

Supported devices

NOS v6.0.2 supports following VDX Switches:

- VDX 6940-144S
- VDX 6940-36Q
- VDX 6740
- VDX 6740T
- VDX 6740T-1G
- VDX 8770-4
- VDX 8770-8
- VDX 2740
  [The VDX 2740 is also known as the Lenovo Flex System EN4023 10GbE scalable Switch. This platform is identified in the system as "EN4023".]
- VDX 2746

VDX 6940-144S

The 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 X 10GbE ports, providing up to 144 x 10GbE ports in a 2RU form factor. (Enabled in a future release)
- 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.
- There is no 40G breakout support.

VDX 6940-36Q

The 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.

VDX 6740

The 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. 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 fabric formation.
- Each 40GbE port is capable of doing a breakout of 4 X 10GbE ports.
- 100Mb Support – Refer to "Support for 100-Mb interfaces" sections below.

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.

- Available in 24, 48 and 64 port SKU.
- 3 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.
- It has 4 X 40 GbE QSFP ports which can be used for uplink and VCS fabric formation.
- Each 40GbE port is capable of doing a breakout of 4 x 10GbE ports.
- Each 40GbE port is also capable of doing an FC breakout of 4 x 8G/16G.
- 100Mb Support – Refer to "Support for 100-Mb interfaces" below.

VDX 6740T-1G

The 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 port can be added to base version with 2X40Gbe POD license.
- It has 4 X 40Gbe QSFP ports which can be used for the uplink and VCS fabric formation.
- Each 40GbE port is capable of doing a breakout of 4 X 10GbE ports.
- Each 40GbE port is also capable of doing a FC breakout of 4 x 8G/16G.
- 100Mb Support – Refer to "Support for 100-Mb interfaces" below.

VDX 8770-4 and VDX 8770-8

The 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 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-sensitive applications.
- Up to 384,000 MAC addresses per fabric for extensive virtualization scalability.
- More than 8000 ports in a single VCS Fabric with 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 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:

- **VDX 8770-4**: Supports up to 192 1/10 GbE ports, or 108 40 GbE ports and 24 100 GbE ports, or a combination.
- **VDX 8770-8**: Supports up to 384 1/10 GbE ports, or 216 40 GbE 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 QSFP ports.
- 100 GbE: LC6×100G line card provides up to 6 x 100 GbE CFP2 ports.

VDX 2740

- VDX blade switch for IBM Flexsystems. It is sold through IBM with part number – IBM EN4023 Ethernet Fabric Switch
- 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
- IBM specific features:
  - Feature on Demand (S/W upgrades)
  - Management integration with IBM CMM

VDX 2746

- VDX blade switch for the Hitachi BladeSymphony 2500 chassis
- 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.

- 6740T, 6740T-1G support 100 MB on front-panel ports. 6740 platform supports 100MB on front-panel ports via qualified 1G copper SFP Optics.

- Note that 100MB is NOT supported on the Management interfaces for 6740 platform series.

With Network OS v6.0.1 support for the following Layer 2 protocols on 100-Mb interfaces has been added:

- LAG/vLAG (see below)
- Port-channel redundancy group
- Virtual Fabrics (extended VLAN)
- BPDU drop
- Storm control
- Port-based and trust CoS
- QoS (all features except Priority Flow Control). The cee default command is not supported on 100-Mb interfaces. The user is advised NOT to use "class cee" policy-map bindings on 100-Mb interfaces.

Note the following additional considerations:

- Trunking (static or dynamic) is NOT supported on 100-Mb interfaces.
- 100 Mb is not supported on ISLs. Only physical and port-channel edge ports are supported.
- 100 Mb is supported for UTP connections.

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

Part number	Description	Compatible devices
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
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 exhaust airflow
- XBR-250WPSDC-F - FRU 250 W DC power supply/fan, non-port-side exhaust airflow
- XBR-250WPSDC-R - FRU 250 W DC power supply/fan, port-side exhaust airflow

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

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

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 exhaust airflow
- XBR-1100WPSAC-R - FRU 500 W AC power supply/fan, port-side exhaust airflow
- 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 v6.0.2b

- 40G-QSFP-SR-BIDI: Support for 40GbE Bi-Directional optics on VDX 6740, VDX 6740T, VDX 6940 and VDX 8770.

Supported optics for Network OS v6.0.2

- 40G-QSFP-ESR4: Hot-pluggable, industry-standard Quad Small Form-Factor Pluggable (QSFP+), MTP (MPO) 1×8 or 1×12 ribbon connector; Extended Short-Reach (eSR4), supports breakout to four 10GBASE-SR modules up to 300 m on OM3 multimode fiber and 400 m on OM4 multimode fiber.

Supported optics for Network OS v6.0.1

For a list of supported fiber-optic transceivers that are available from Exteme Networks, refer to the latest version of the 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, 2740 and 2746 switches. Breakout optics are supported only for the VDX 8770 (40G line-card), 6740/T, 2740, 2746 and 6940 platforms. The Mellanox (MAM1Q00A) optic is only supported on the VDX 8770 and 6740/T platforms. Optics/cables for the VDX 2740 are supplied by IBM, and IBM maintains the support table for them.

	FRU and Optics SKU	Description
1GbE	XBR-000190 (1-pack)	1 GbE copper
	E1MG-SX-OM (1-pack)* E1MG-SX-OM-8 (8-pack)*	1000Base-SX
	E1MG-LX-OM (1-pack)* E1MG-LX-OM-8 (8-pack)*	1000Base-LX
10GbE	10G-SFPP-SR (1-pack) 10G-SFPP-SR-8 (8-pack)	10 Gbps SR
	10G-SFPP-LR (1-pack) 10G-SFPP-LR-8 (8-pack)	10 Gbps LR (10km)
	10G-SFPP-ER (1-pack) 10G-SFPP-ER-8 (8-pack)	10 Gbps ER (40km)
	10G-SFPP-ZR	10 Gbps ZR (80km)
	10G-SFPP-TWX-0101 (1-pack) 10G-SFPP-TWX-0108 (8-pack)	1 m Twinax copper cable
	10G-SFPP-TWX-0301 (1-pack) 10G-SFPP-TWX-0308 (8-pack)	3 m Twinax copper cable
	10G-SFPP-TWX-0501 (1-pack) 10G-SFPP-TWX-0508 (8-pack)	5 m Twinax copper cable
	10Ge-SFPP-AOC-0701	10GbE SFP+ Direct Attached Active Optical Cable, 7m, 1-pack
	10Ge-SFPP-AOC-1001	10GbE SFP+ Direct Attached Active Optical Cable, 10m, 1-pack
	10G-SFPP-USR	10GE USR SFP+ optic (LC), target range 100m over MMF, 1-pack
40GbE	40G-QSFP-QSFP-C-0101	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 1m, 1-pack
	40G-QSFP-QSFP-C-0301	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 3m, 1-pack
	40G-QSFP-QSFP-C-0501	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 5m, 1-pack
	40G-QSFP-4SFP-C-0101	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 1m, 1-pack
	40G-QSFP-4SFP-C-0301	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 3m, 1-pack
	40G-QSFP-4SFP-C-0501	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 5m, 1-pack
	40G-QSFP-SR4	40 GbE SR4 optic
	40G-QSFP-SR4-INT	40 GbESR4 (4×10 Gbe SFPP break-out capable). Breakout Optical cable is not included with this optics
	40G-QSFP-ESR4	40GBase-eSR4 QSFP+ optic (MTP 1x12), 300m over MMF, (10GBASE-SR compatible, breakout), 1-pack QSPF+ can support 4*8 or 4*16G on VDX 6740T , VDX 6740T-1G, and VDX 2740
	40G-QSFP-SR4-BIDI	40GbE QSFP+ BiDi 100m
	40G-QSFP-LM4	40 GbE 140m multi-mode or 2km single-mode optic

	FRU and Optics SKU	Description
	40G-QSFP-QSFP-AOC-1001	40GbE Direct Attached QSFP+ to QSFP+ Active Optical Cable, 10m, 1-pack
	40G-QSFP-4SFP-AOC-1001	4x10GE Direct Attached QSFP+ to 4 SFP+ Active Optical Breakout Cable, 10m, 1-pack
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-LR4-10KM	100 GbE CFP2 optic, LR4, for distances up to 10 km over SMF
	100G-CFP2-SR10 (1-pack)	100 GbE CFP2 optic, SR10, for distances up to 100 m over MMF

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

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

Mellanox MAM1Q00A	Quad to Serial Small Form Factor Pluggable Adapter which can be used with following Extreme 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)

*Note: Legacy Foundry Networks branded optics are not supported

SOFTWARE UPGRADE AND DOWNGRADE

Image filenames

Download the following images from www.extremenetworks.com

Image filename	Description	Supported Device or Module
nos6.0.2h.tar.gz	Network OS v6.0.2h for Unix	NA
nos6.0.2h.zip	Network OS v6.0.2h for Windows	NA
nos6.0.2h_all_mibs.tar.gz	Network OS v6.0.2h MIBS	NA
nos6.0.2h_releasenotes_v1.0	Network OS v6.0.2h Release Notes v1.0 (PDF)	NA
nos6.0.2h.md5	Network OS v6.0.2h md5 Checksum	NA

Upgrade/downgrade considerations

Starting with Network OS v6.0.0, a Exteme Networks 4GB USB drive is the preferred storage medium for firmware installation using USB, but 2GB USB drives are still valid.

During firmware upgrade from NOS5.x to NOS6.0.2 and vise versa, if rule has been configured to creates the Role-Based Access Permissions (RBAC) permissions associated with a role then we recommend to perform any of below action else the rules will get impacted.

1. Copy running configuration to server before firmware upgrade. After upgrade, default the configuration and configure the saved configuration. OR

2. Save all configured rules in some file, remove all rules from configuration, perform firmware upgrade, after firmware upgrade reconfigure those rules from the saved file.

Migration Path

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

The "write/erase" issue fixes support ISSU firmware upgrade, and hence no cold boot is required to activate these fixes. However, please review the release notes thoroughly to check if any other fixes require coldboot or reload after the ISSU upgrade. Network OS 6.0.2h has uboot changes for new HW components and a "write erase" issue fix.

Note: Firmware download is not available for identical release numbers, such as Network OS 6.0.2a to Network OS 6.0.2a.

To / From	NOS v4.0.1_hit1 (Only for VDX 2746)	NOS v4.1.x	NOS v5.0.x	NOS v6.0.0	NOS v6.0.1x	NOS v6.0.2x
NOS v4.0.1_hit1 (Only for VDX 2746)	NA	NA	NA	NA	FWDL "*default-config*" (Config loss	FWDL "*default-config*" (Config loss
NOS v4.1.x	NA	ISSU for upgrade; Coldboot for downgrade.	FWDL with "*coldboot*"	FWDL "*default-config*" (Config loss)	FWDL "*default-config*" (Config loss)	FWDL "*default-config*" (Config loss)
NOS v5.0.x	NA	FWDL with "*coldboot*"	ISSU for upgrade; Coldboot for downgrade.	FWDL with "*coldboot*"	FWDL with "*coldboot*". (If 5.0.0->6.0.1, FWDL with "default-config" – with config loss only)	FWDL with "*coldboot*". (If 5.0.0->6.0.1, FWDL with "default-config" – with config loss only)
NOS v6.0.0	NA	FWDL "*default-config*" (Config loss)	FWDL with "*coldboot*"	NA	FWDL with "*coldboot*"	FWDL with "*coldboot*"
NOS v6.0.1x	FWDL "*default-config*" (Config loss)	FWDL "*default-config*" (Config loss)	FWDL with "*coldboot*". (If 6.0.1->5.0.0, FWDL with "default-config" – with config loss only)	FWDL with "*coldboot*"	ISSU for upgrade; Coldboot for downgrade.	ISSU for upgrade; Coldboot for downgrade.
NOS v6.0.2x	FWDL "*default-config*" (Config loss)	FWDL "*default-config*" (Config loss)	FWDL with "*coldboot*". (If 6.0.1->5.0.0, FWDL with "default-config" – with config loss only)	FWDL with "*coldboot*"	ISSU for upgrade; Coldboot for downgrade.	ISSU for upgrade; Coldboot for downgrade.

FWDL – firmware download

Note: The following defect fixes are not effective for ISSU upgrades from any 6.0.x version prior to this release. You must reboot after upgrading in order to activate these fixes:

- DEFECT000568368
- DEFECT000636649
- DEFECT000635844

- DEFECT000620617
- DEFECT000637857

Refer to the defect tables for more information.

NOTES

1. Only Extreme Network Advisor (BNA) v12.4.2 (available separately) supports NOS v6.0.1a. It is required to first upgrade to BNA v12.4.2 and then upgrade switches to Network OS v6.0.1a.
2. 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.
3. 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.
4. **Limitations:
   a) After downgrading from Network OS v6.0.2 to Network OS v5.0.x FCoE devices may not login back 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) If port-security feature is enabled, it is necessary to limit the Max OUI configuration to 13 ports to avoid switch instability during firmware download operation.
   c) If DNS is enabled on the switch, it is necessary to ensure the DNS servers are valid and reachable before executing firmware download command. If the DNS servers are not reachable, it is necessary to remove/correct the DNS configuration before executing firmware download command as the command is not blocked and may cause timeouts during firmware download operation.
   d) When downgrade from NOS6.0.2 to NOS6.0.1 is performed, if IPv6 VRRP Link local Virtual IP is not in fe80::/64 format(but in fe80::/10 format), user is expected to remove this config before performing downgrade."

Management IP connectivity

Firmware downgrade from NETWORK OS V6.0.2 to NOS6.0.0/NOS5.0.1 with default-vrf option in host/v3host use-vrf is not supported. The trap configuration use-vrf should be set to mgmt-vrf before downgrade.

Firmware upgrade to v6.0.1 will modify the 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 (5.x) OR Default VRF (4.x).

In regards to SNMP, firmware downgrade from NETWORK OS V6.0.2 to lower versions that do not support "use-vrf" keyword, the trap host/v3host configured with use-vrf option as "default vrf" is not supported. Trap configuration with use-vrf as "mgmt-vrf" needs to set before downgrade.

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.

Firmware Installation

In fabric cluster mode

- 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

- logical-chassis firmware download 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.

  - 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 Network OS can be found in the *Network OS Administrator's Guide.* This section includes special considerations and caveats to be aware of when upgrading

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

Note: Installing 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 Cluster Wide**
- Download the firmware on all the switches running Network OS v5.0.1x 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 v5.0.1x 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 coldboot option.
- 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 the Cluster**
- Download the firmware in the switch nodes one node at a time in cluster running NOS 5.0.1x 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 v5.0.1x Cluster and form Fabric cluster. Eventually, when all the nodes are upgraded, they will form one Network OS 6.0.2 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 5.0.1x config back by executing "copy file running".

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 a time.
- **Approach 2**: Static vLAGs and Dynamic vLAGs without configuration changes.

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 an error.
- 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 changes

SNMP MIB – VLAN update

During an snmpwalk or snmpgetbulk, all the VLAN interfaces are filtered out from the 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 both management/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 (default/management).

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 <management | default>

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

To disable per link TRAP under interface

[No] snmp trap link-status

SNMP – IF MIB

ifDescr and ifName in IF MIB modified to show the interface in 3 tuple format (rbridge-id/slot/port) from 2 tuple (slot/port)

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 OR magmt-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>

Syslog VRF Support

Syslog servers logging can be configured to point to syslog servers in default-vrf OR mgmt-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>
[secure [port <xxxx>]]
```

SSH Support in non-default VRFs

```
Support added for SSH in non-default VRFs (the supported number of VRFs is 6).
```

SNMP support in non-default VRFs

```
Support added for SNMP Infrastructure in non-default VRFs (the supported number of
VRFs is 6).
```

SNMP community name string

```
The maximum length of the SNMP community name string has been extended to 64
characters.
```

SNMP module removal traps

Module removal traps now follow a logic numbering scheme and have self-explaining text.

Other Management Services

Other management services like REST, Netconf, HTTP, SNMP MIB's would be available in both default & management VRF's.

HA Failover

The order of traps on HA failovers has been modified to be hafailoverstart, warmstart, and hafailoverDone.

Added new API s handleCpEvent_failover and handleFailover, to send cpstatus change trap for hafailoverStart/Done events.

The hafailoverStart traps are now part of cpstatuschangetrap in addition to warmstart trap.

SNMP MIB enhancement

SNMP MIB has been enhanced to include details about fabric ISL links, neighbor information, and logical chassis ID.

MIB Enhancement

MIB messages now include the part name.

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 v6.0.2 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
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

NOS v6.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
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	2,000	2,000	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
Maximum # of MSTP instance	32	32	32	32
Maximum # of VLAN in PVST	128	128	128	128
Maximum # of LAGs (Port Channels)	60	384	63	63
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

NOS v6.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
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	32000	1500	1500
# of OSPFv3 Interfaces	100	256	100	100
# of OSPFv3 enabled subnets	100	256	100	100
Maximum # of IPv4 routes in SW *	8000	200000	10000	10000
Maximum # of IPv6 routes in SW *	1500	32000	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	255	255
Maximum # of VRRP v3 instances per system	255	1024	255	255
Maximum # of VRRP instances per interface	32	16	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	32000	1500	1500
Maximum # of L3 ECMP	16	32	32	32

NOS v6.0.2 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940-36Q	VDX 6940-144S
Maximum # of IPv4 interfaces per system *	2000	4000	2000	2000
Maximum # of IPv6 interfaces per system *	512	4000	512	512
Maximum # of VRF per node	32	512	512	512
Maximum # of VRFs support protocols per node	32	128	32	32
Maximum # of I-BGP peers	256	512	256	256
Maximum # of E-BGP peers	64	128	64	64
Maximum # of IPv4 BGP routes in HW *	8000	200000	10000	10000
Maximum # of IPv6 BGP routes in HW *	1,500	32,000	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 Address on the system	2000	4000	2000	2000
DHCP IPv6 Relay Address	2000	4000	2000	2000
Max Number of configurable PBR route maps	64	64	64	64
Max Number of configurable PBR stanzas	1024	1024	1024	1024
Max Number of TCAMs available for PBR	512	8192	512	512
Max Number of configurable next hops within a single PBR stanza	128	128	128	128
Maximum # of OpenFlow L2 flows	1000	4000	879	879
Maximum # of OpenFlow L3 flows	1000	4000	879	879

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

**Please consult your 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.

Following tables list the TCAM profiles for the VDX switches:

VDX 8770 TCAM profile

	L2 ACL	IPv4 ACL	IPv6 ACL	PBR	QOS	Multicast	Openflow
Legacy(Default)	8192	8192	504	1024	512	2048	0
l2-ipv4-acl	8192	12288	0	2048	1024	0	0
l2-acl-qos	12288	4096	0	0	1024	0	0
ipv4-ipv6-pbr	1016	4096	4088	4096	1024	0	0
ipv4-ipv6-qos	1016	8192	4088	2048	2048	0	0
ipv4-ipv6-mcast	504	512	504	1024	2048	2048	0
l2-dai	8192	8192	504	1024	2048	2048	0
openflow	8192	8192	504	1024	512	8192	4096

VDX 6740 TCAM profile

	L2 ACL	IPv4 ACL	IPv6 ACL	PBR	QOS	Multicast	Openflow
Legacy(Default)	504	508	504	512	512	1024	0
l2-ipv4-acl	1016	1532	0	512	512	0	0
l2-acl-qos	3064	508	0	0	0	0	0
ipv4-ipv6-pbr	0	508	504	512	0	0	0
ipv4-ipv6-qos	0	508	504	0	512	0	0
ipv4-ipv6-mcast	0	508	504	0	0	1024	0
l2-dai	504	508	504	0	512	1024	0
openflow	504	508	504	0	0	0	2048

VDX 6940 TCAM profile

	L2 ACL	IPv4 ACL	IPv6 ACL	PBR	QOS	Multicast	Openflow
Legacy(Default)	504	1020	504	512	512	1024	0
l2-ipv4-acl	1528	1532	504	512	512	1024	0
l2-acl-qos	3064	1020	0	0	512	1024	0
ipv4-ipv6-pbr	0	508	504	2048	0	1024	0

ipv4-ipv6-qos	0	508	504	0	2048	1024	0
ipv4-ipv6-mcast	504	508	504	0	512	1024	0
l2-dai	504	1020	504	0	512	1024	0
openflow	504	508	504	512	512	0	2048

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 v6.0.1. This is a representative list of devices, Network OS v6.0.0 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

Vendor	Storage Array Model	Protocol	Switch Model	Initiator
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, 2740	Windows 2012 R2
	Flash 840	FC	6740/T, 2740	Windows 2012 R2
	XIV	FC	6740/T, 2740	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

Vendor	Model	Protocol	Switch Model	OS
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
Broadcom	88Y5920 IBM CN4022 2-port 10Gb Converged Adapter	Eth	2740	Windows 2008, Windows 2012
	49Y7900 IBM EN2024 4-port 1Gb Ethernet Adapter	Eth	2740	Windows 2008, Windows 2012
Emulex	00Y3306 IBM CN4054R 10Gb Virtual Fabric Adapter	Eth/FCoE	2740	Windows 2008, Windows 2012
	IBM FC EC24	Eth	2740	Windows 2008, Windows 2012
	IBM FC 1762	Eth	2740	Windows 2008, Windows 2012
	94Y5160 IBM CN4058S 8-port 10Gb Virtual Fabric Adapter	Eth/FCoE	2740	Windows 2008, Windows 2012
	00JY800 IBM CN4052 2-port 10Gb Virtual Fabric Adapter	Eth/FCoE	2740	Windows 2008, Windows 2012
	10GE BE3/BE3R LOM	Eth/FCoE	2740	Windows 2008, Windows 2012
Mellanox	IBM 90Y3466	Eth	2740	Windows 2008, Windows 2012
	IBM FC EC26	Eth	2740	Windows 2008, Windows 2012

ADDITIONAL CONSIDERATIONS

Limitations and Restrictions

DHCP option 82

- Adding ip dhcp option 82 field for L2 interface is not supported.
- Vlan ID in Circuit ID field is sent as Hex value not as string value.

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 v6.0.0. 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
    - In some instances, there could be delay in starting of range operation after issued
    - When range issued for very large subset (e.g 4k VLAN), timeout can happen or user may see temporary switch being unresponsive or high CPU. Exteme Networks recommends using range in smaller chunks.. Especially, while configuring VLANs/VEs, Exteme Networks 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.
- System does not warn user on deleting the ip config when vrf is configured.
- Redistributed connected/static routes may be shown twice as part of config.
- 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 v6.0.0. Exteme Networks 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.
    - Show features – no use.

- ▪ Show ssm –, for debugging purpose only.
- ▪ Autoupgrade command in config mode
- During "copy running-config startup-config" or "copy support" user might see occasional and temporary CPU spikes (up to ~30-40%).
- show mac-address-table command on console with include option can not be aborted with a break/ctrl-C. Use a telnet session for the same.
- Short form of MAC-Address is not supported as filter in "show running-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.
- NTP server with full length IPv6 address configuration can be used only with NTP key with less than 15 characters length.
- 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 vrf default-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 "chassis enable".
- Chassis-name is limited to 15 characters.
- VDX 6940-144S does not support breakout of 40G ports in Network OS v6.0.1 and 6.0.1a
- 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.
- System verification/ offline diagnostics tests need "chassis disable" before the test and "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 6940 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 "non 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 v3.x and above:
  - LC48×10G
  - LC12×40G
- The VDX 8770 supports following line-cards only on Network OS v4.1.2 and above:
  - 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 VDX 8770, 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, a Extreme 4GB USB drive is the preferred storage medium for performing any USB-based operations, but 2GB USB drives are still valid.

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 VDX 674x. 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 VDX 8770 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 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 VDX 8770 switches.

VCS

- VCS supports traffic traversing through maximum 9 hops of VDX switches. Thus VCS can have maximum of 9 hops in a ring topology.
- 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 it's own VCS and not join the parent cluster. In such situations, reload the node that is required to join the cluster.
- Fabric Cluster Mode:
  o 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*.
  o After a cluster reboot, Exteme Networks 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 most cases.
- "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.
- VCS for Network OS v6.0.1a:
  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

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 migrating from FC to LC.
- "show vcs" output displaying "Coordinator" indicates "Principal" node role.
- Principal priority value ranges from 1 to 128, 1 being the highest. Recommend to set higher principle priority to VDX 8770.

- 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" or exercise the default-config boot feature on the segmented switch to bring it back to the cluster.
- Number of config snapshots saved on switch is limited to 4 per rbridge ID. So on 24 node cluster, a max of 24 * 4 = 96 snapshots are possible.
- For Netconf and SNMP, user has to poll using individual node Management IP.
- Creating a snapshot with "\" in snapshot-id creates the snapshot file with incorrect name.
- There will not be any raslog to the user when replacement of a node fails.
- With large configs, while a switch is rejoining a fabric with default config, "%Error:Could not find Interface" may be printed temporarily. The switch will recover and join the fabric.
- Config changes during principle switch-overs are not supported and may segment the cluster.
- An Rbridge in bare-metal state can join a VCS with or without the pre-provisioning mapping. For the scenario without the pre-provisioning mapping, the Rbridge must be in bare-metal state and in the "Offline" state of the VCS.
- Disabling virtual-fabric may take up to 10 minutes depending on the number of ISLs and VLAN interfaces configured in the VCS.

Trunks for Network OS v6.0.1a

- The VDX 6940-36Q and VDX 6940-144s support 63 port-channels, including LACP and Extreme PO.

Extreme Trunks

- The VDX 6740, VDX 6740T, VDX 2740 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 2740, 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 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).
- The VDX 6940-36Q Extreme trunk (BTRUNK) can support up to 3 member links with a maximum throughput of 120G using 3x40G or 120G using 12x10G breakout ports in the same trunk group.
- The VDX 6940-144S Extreme trunk (BTRUNK) can support up to 3 member links with a maximum throughput of 120G 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 breakout mode.
- VDX 6940-144S does not support breakout of 40G ports in Network OS v6.0.1.
- For VDX 6740, 6740T, 2740 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, there is only SFP and no per-breakout media information. The show media command will display the same media information for all breakout interfaces. 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/2740/2746, the breakout ports are FlexPort capable, and may be configured to connect to FC switches with 4x16G breakout supported cables and optics.

Restrictions for Ports in 1G Mode and 1G Ports on VDX 6740T-1G

- 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 ISL links.
- Extreme Trunks cannot be formed with 1G. Extreme Trunks are only supported on 10G.
- 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.

vLAG

- LAGs are created with default speed of 10G. Therefore Exteme Networks recommends end user to set required speed manually based on member speed using "speed" command.
- When configuring LACP LAG between VDX & non- Exteme Networks switches it is highly recommended to enable the vLAG ignore-split on the VDX . Ignore split option is enabled 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 be the same as the gateway address for the management port of this switch.
- 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 same subnet.
- 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 VDX 8770.

Security, Management ACLs, Authentication, Authorization

- Login authentication service (aaa authentication login cli):
  - o 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 not available.
  - o Behavior of "local" option in pre-4.1.0 releases is changed to the "local-auth-fallback" option.
  - o 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 management interfaces.
- Configuring TACACS+ or RADIUS without a key is not supported. If no key is configured, the switch uses a default key of "sharedsecret".
- 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.
- 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 LDAP client:
  - o Windows 2000
  - o Windows 2003
  - o 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). Exteme Networks 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.

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 VLAN ID is greater than secondary there is an issue with config replay.

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 Exteme Networks 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 configured platform.
- 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.
- "long-distance isl" command based extended fabrics are supported only on 10G interfaces. 40G and 100G interfaces do not support "long-distance isl" command, however can extend distances for non-lossless traffic up to 2Km 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 and 10G-SFPP-ZR 80km optics.

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.

    Exteme Networks 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.
- Default port-profile configuration is changed from Network OS v4.1.0 onwards. The "switch port trunk allow VLAN all" that was present in prior version is removed. Other configuration stays the same.
- From Network OS v4.1.0 onwards, user defined port-profile-domain is introduced to control the VM mobility. Port-profile created must be explicitly associated with a profile domain.
- From Network OS v4.1.0 onwards, after upgrade, a new port-profile named UpgradedVlanProfile is auto-created. This profile has the single VLAN profile that contains the

"switch port trunk allow VLAN all". This is the configuration that is present in the default port-profile prior to Network OS v4.1.0.

- 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 trafficfailure.
- Adding/removing the auto-port-profile to the user-created domain when Service VF is enabled 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.

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.
- 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 for Network OS v6.0.1a
- On switches shipped with NOS 6.0.1 the default mode of operation is Access Gateway for the **VDX 6740, 6740T, 6740T-1G, VDX 2740, VDX 2746.**. Logical SANs have been supported since Network OS v6.0.0.
- Exteme Networks 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. **FCoE logical SAN is not supported in an FC cluster.**
- 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.

FCoE
- Adding/removing an ISL or adding/removing a VDX switch from cluster may cause that exiting VDX switch in the cluster stops generating a FCOE keep alive. And this results in FCoE port continuously flapping. To recover, re-configure FCoE by removing and adding fcoeport configuration (no fcoeport/fcoeport default) on the affected interfaces. As an alternate recovery, shut / no shut on the affected interfaces.
- On switches shipped with NOS 6.0.1x the default mode of operation is Access Gateway {see preceding section). 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 6.0x, it will be in FCF mode.
- 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.

- Exteme Networks 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.
- Exteme Networks 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.
- If FCoE priority 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 {see preceding section).
- 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.
- When an RBridge is removed from VCS cluster, it has to be manually removed from fcoe fabric-map {see preceding section).
- In NETWORK OS V6.0.1, 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, 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, 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.

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.

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 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/N ports).
- Trunking is not supported on Fibre Channel ports running at 2G or 4G speeds.
- On the VDX 6740, VDX 6740T, VDX 6740T-1G, VDX 2740 and VDX 2746 platforms Fibre Channel trunks are limited to 2 FC trunks per port group.
- Fibre Channel trunks only form with ports running at the same speed.

Access Gateway

- AG enable/disable command is moved to configuration mode in 6.0.1. From principal node AG mode can be changed on individual nodes under rbridge-id-ag configuration mode.
- All AG configurations have to be done under rbridge-id-ag sub mode. Prefix "ag" is not allowed any more.
- 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 log in.

- 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>".

ND/RA

- Proxy ND is not supported.

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 6740T platforms. Command "protocol vrrp-extended" is added to specifically enable VRRP-E.
- 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 from file.
- VRRP v4 or v6 can be enabled with VRRP-E v4 and v6 on the VDX 6940 family.
- 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 6740T
  - 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.

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 using a route-map.
- EBGP TTL Security Hack Protection is not supported.

ACL
- L2 User ACL deny rule can prevent trapping of L3 control frames.
- IPv6 ACLs at ingress are not applicable for packets with Link local source address.
- ACL Logging at egress can impact forwarding traffic at high rates.

- Counters for hard-drop ACLs may not count accurately.
- 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.
- 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 to it.
  - 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 to it.

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 ipv4 command.
  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.
- IP Services like telnet are supported only on mgmt-vrf.
- User will be able to access VDX switches only through interfaces belonging to mgmt-vrf.
- 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.
- Support added for SNMP Infrastructure in non-default VRFs (the supported number of VRFs is 6).
- Support added for SSH in non-default VRFs (the supported number of VRFs is 6).

BGP-VRF

- Local-as <num> can be configured for particular VRF under "address-family ipv4 unicast vrf <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 unicast vrf <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 the route-map.

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 of indirection.
- 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 routing loops.
- Switch management from non-management VRF by leaking route from non-management to 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.

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 DAI acls.
- ARP packets more than 190 bytes on a DAI enabled vlan will be dropped.
- ARP access-list with longer names is not effective (greater than 20 characters)

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 metal mode.

- 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 Principle node alone before turn on power for all secondary nodes.
- When the switch is in Factory default, DAD is enabled upon power up the switch
- DAD executes only if the switch configuration is the default configuration. If the configuration on the switch is not the default configuration, DAD exits.
- 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.

- [UPDATED for 6.0.1a] 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 the" track min-link" number is greater than number of uplinks, then the downlink will be shut down with a warning message

OpenFlow

- 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.
- Pre-provisioned flow-mods will not be replayed to a new slot coming online. GROUP and METER configurations will be always replayed.
- On 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.
- [UPDATED for 6.0.1a] For layer 3 rules, switch can't differentiate between tagged and untagged flows when matching against rules. This applies to all supported platforms.
- [UPDATED for 6.0.1a] Filtering options are not supported for show openflow CLIs. Show openflow commands with filter option show the complete output.
- [UPDATED for 6.0.1a] 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.
- [UPDATED for 6.0.1a] 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-default rcv-queue.
- [UPDATED for 6.0.1a] With large number of flows, "show openflow flow <>" may take 20 seconds to display packet counts.
- [UPDATED for 6.0.1a] "Module Unknown" is shown for CLI "show open resources" in VDX 6940-144S.

Exteme Networks Vyatta Controller (BVC)

- Controller does not update the config database based on the flow rejected notification/group rejected notification/meter rejected notification/delete notification/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 : User needs to delete the group/meter/flow from the config database and program them again.

- In scale scenario, few flow-mods are not programmed after reconnection. Work around : User needs to delete the missing flow-mods and program them again.
- 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 known issues/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 , even though 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 way round.
  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.
- **[UPDATED for 6.0.1a]** 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 ACL only.
- Range is not supported.
- Maximum 16 sessions supported.

VXLAN Gateway for VMware NSX

- VCS VXLAN Gateway for NSX is supported only on VDX 6740, VDX 6740T,VDX 6740T-1G, VDX 6940-36Q VDX 6940-144S and VDX 2740/2746.
- VCS VXLAN Gateway for NSX is supported only in the VCS Logical Chassis mode.
- 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 and VF Extension cannot be enabled in the same VCS fabric.
- VMware NSX vSwitch with vSphere version 5.5 (ESXi 5.5), XenServer 6.2, and KVMon 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.
- When using the command "show overlay-gateway name <name> VLAN statistics" for debugging overlay-gateway VLANs, it should be noted that the statistics information is limited to 256 VLANs (Rx) and 116 VLANs (Tx) for VDX 6740x; 256 VLANs (Rx) and 250 VLANs (Tx) for VDX 6940-36Q; and 256 vlans (Rx) and 125 vlans (Tx) for VDX6940-144S.
- When multiple VMware NSX Service Nodes are setup, only one of the node would be used for handling BUM traffic. During service node failover scenarios another Service node would be selected for handling BUM traffic, if BFD is enabled for all the Service nodes.
- ALL the VE interfaces should run VRRP-E with the same VRID and same virtual-mac to terminate the incoming packets on other VLANs.
- Tunnels egressing/ingressing through an ISL port is not supported.
- 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.
- In-band management for VCS VxLAN GW with default-vrf is not supported.
- Load balancing between multiple Service node tunnels is not supported.

VF Extension using VxLAN

- [UPDATED for 6.0.1a] 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 Chassis mode.
- 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 VCS Cluster.
- 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 not supported.
- L3 routing protocols and static routes over tunnels are not supported.
- 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 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
- [UPDATED for 6.0.1a] The TCAM profiles the user can create may not match the max scale numbers due to reserved routes/entries which are created for internal use.

Management VRF
OSPF/BGP/PIM/VRRP/VRRPe is 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.
- In-band management ports can be part of Management VRF or default VRF.
- Switch cannot be managed from leaked routes pointing to Management-VRF.
- Address family on Management VRF cannot be removed.
- Firmware download/supportsave is not supported on in-band ports. This limitation is applicable for Management vrf/default-vrf/ non-default vrf.

Conversational MAC Learning
- Conversational MAC Learning and 'Disable Source MAC Learning' cannot be enabled 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. Hence once 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 the cluster.

Element Manager Support

- The Element Manager GUI is intended for use with the VDX 2740 (Lenovo EN4023) and VDX 2746 platforms only, and may not be used with any other VDX platform.

Logical Chassis HA

- HA failover and unplanned failover is supported on VDX 8770 only.
- When the principle 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 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 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 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.
- For frames forward 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.
- "fabric neighbor-discovery" command should be entered on both the ends of the link to disable the ISL formation when Extreme switches used on both the ends.
- 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

- Only BNA v12.4.2 (to be made available separately) supports NOS v6.0.1. It is required to first upgrade to BNA v12.4.2 and then upgrade the switches to NOS v6.0.1.
- Only one MAPS policy can be active at any time.
- All MAPS thresholds, policies, rules and groups are pre-defined in NETWORK OS V6.0.1 and may not be modified.
- MAPS port level alerting in NETWORK OS V6.0.1 is not available for FlexPorts configured in Fibre Channel mode.
- MAPS configuration and monitoring is applicable per switch, so users need to apply configuration on each switch being monitored. BNA must be used for fabric-wide configurations.
- Supported on VDX 2740, 2746, 6740, 6940 and 8770 platforms.
- RX_SYM_ERR MAPS messages are displayed when breakout cable is connected on a 40G interface that is not configured for breakout.

- When linecard on the remote end of the link is powered off, MAPS generates Insertion/Removal notification for the SFPs on the local side. These can be ignored.

Miscellaneous

- 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 inbound direction.
- 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 4093,4095 and 1002 are reserved and used for internal cluster operations.
- "Clear ip route all" need to be issued once the maximum number of routes supported by a router is exceeded.
- SNMPset operation is supported for certain MIB objects in v5.0.1a.
- 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.

- It is recommended to avoid using " !" OR ";" OR ":" in the password for BNA to poll the VDX switches. It might impact the switch adversely.
- RBAC may not be able to block the base "show" or "show running" cmds, thus it is recommended to use more specific "show running-config <>" command to be blocked by rules.
- Portchannel OR breakout ports using BIDI Optics may encounter unexpected link-flaps on the link-up event & may take up to 20 seconds to stabilize.

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 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 http://my.extremenetworks.com.

TSB Issues Resolved in Network OS v6.0.2e

TSB	Summary
TSB 2017-254-A	ONMd daemon gets terminated and sends below RAS log on console when VDX would undergo an unexpected reload. `2017/03/23-11:21:26, [HASM-1200], 7957057, SW/0	Active	FFDC, WARNING, VDX6740, Detected termination of process onmd:4383`
TSB 2017-256-A	VDX may not be able to learn MAC addresses OR in certain cases, it may encounter an unexpected reload due to termination of l2agtd. VDX may encounter mac learning issues OR in certain conditions, undergo an unexpected reload. The issue may reoccur even after the reloads until the preventive workaround is applied. When vdx encounters the state where mac is not learnt anymore, below is the signature to check the existence of this defect occurrence. `sw0#debug internal l2agt 0 p0 dump globals` `is_vcs_enabled: 0` `sw0# debug internal l2agt 0 p0 dump counters all	i msg_len`  `l2agt_recv_msg_len_err : 1` The issue would persist until the layer-2 related configuration is modified.

TSB Issues Resolved in Network OS v6.0.2b

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 prevent overheating. 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 become unusable.

	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 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.

TSB 2016-234-A	When a VDX is configured with an ingress IP access-list on a VE interface with routed keyword, the access-list operations (eg. permit, deny, implicit deny, etc.) would not act upon the traffic destined to VRRP's virtual MAC address. This unexpected behavior may impact the security of the network.
	
	Any one of the following prevention mechanisms can be used as a preventive workaround.
	
	a) Configure the access-group to apply ACL rules in the egress direction to act upon the traffic routed by the VRRP enabled VE interface.
	
	`interface Ve 20`
	`  ip access-group <name> out routed`
	
	b) Configure the access-group without the routed keyword in the ingress direction.
	
	`interface Ve 20`
	`  ip access-group <name> in`
	
	c) Configure gateway as Ve IP address instead of VRRP virtual-ip on the end hosts.
	
	The defect will be fixed for VDX 6740 & VDX69XX platforms & addressed in releases 6.0.2a1 (CCE patch), 6.0.2b, 7.0.1a & higher. The fix will not be applicable for VDX8770 due to architectural differences.
	
	For VDX8770, it is recommended to avoid using the "routed" keyword OR utilize the above workaround.

TSB Issues Resolved in Network OS v6.0.2a

TSB	Summary
TSB 2016-231-A	When "aaa accounting" is configured on the VDX switches running the specified releases the user is at risk of an unexpected reload due to memory depletion. This issue affects VDX switches and occurs only when "aaa accounting" is configured and NMS keeps issuing commands OR netconf calls. This is most often seen with BNA polling VDX periodically. Recovery and preventive workarounds are provided. The recovery workaround is to reload switch gracefully or contact TAC for other alternatives for graceful recovery.

Customer should remove below configuration to prevent DCMd memory depletion:

```
sw0(config)# no aaa accounting exec default start-stop tacacs+
sw0(config)# no aaa accounting commands default start-stop tacacs+
```

Caveats

- After clearing ipv6 neighbor table with "clear ipv6 neighbor no-refresh", VDX still uses the old nd entries to send ipv6 packet out instead of refreshing ND Cache. [599575 for more info]
- Using ACLs to secure RW or RO access for SNMP works upto 20 SNMP community. Having more than that configuration followed by reload may lead to rolling reboot on the node. Workaround is to used reduced scale configuration.
- For Defect 597782-In rare scenarios Mgmt interface(eth0) and Ve interfaces may share the same mac address on 6940 .This fix is present in 6.0.2b and 7.0.1 GA release. For the fix to be effective, user need to do cold boot of the node. After the cold-boot, any traffic destined to VE interface would be dropped as mac-address of VE had changed to different value. To recover from the state, arp should be refreshed from all the connected hosts.

Firmware Upgrade/ISSU (In-Service Software Upgrade)

- After upgrading from nos6.0.0x or prior version, BGP Confederation peer configuration will be lost if confederation peer AS has local-as configured. Starting nos6.0.1x, BGP confededration peer AS configuration does not allow local-as. Before upgrading to nos6.0.1x, it is necessary to ensure the peer AS configuration does not have local AS configured. If upgrade has already been completed to 6.0.1x with local-as configured in confederation peers list, removing and re-adding the configuration without the local-as as part of the confederation peer list should be performed.
- For configurations larger than 60K lines of running-config, it is recommended to use coldboot upgrade. Performing ISSU with larger than 60K lines may lead to NOS versions on the primary and secondary partitions getting out of sync with standby. In order to recover, it is required to perform "firmware sync" to go back to 6.0.1 and then perform coldboot upgrade.
- Under rare circumstances, performing ISSU upgrade from 6.0.1 to 6.0.1a on VDX6940 and VDX6740 platforms may cause secondary partitions to not come up. When this issue occurs, the "show version" command will not display secondary partition on the switch. In order to recover, it is required to reload the node to go back to 6.0.1 and re-try ISSU or coldboot upgrade again.
- When VRRPE sessions are present on VFAB Ve interfaces and non-VFAB Ve interfaces along with Fabric Virtual-Gateway sessions in the system, ISSU can result in traffic outage for some of the VRRPe flows on VFAB Ve sessions. To recover, it is required to disable and enable the affected VRRPe session.
- After ISSU upgrade from 6.0.1 to 6.0.1a, if the peer port (40G/Breakout/1G) goes down for any reason then the local port may go down and not come up. To recover, it is required to perform shut/no shut on the local port.
- Static Route BFD, BGP BFD and OSPFv2/v3 BFD
  - If BFD timer values are not configured by user in NOS v6.0.1, and ISSU upgrade to 6.0.1a is performed, the running-config will display 50 ms, however the BFD will function at default value of 200 ms (VDX8770) and 500 ms (VDX6740x, 6940x). It is recommended to change the BFD timer values to default or higher before performing ISSU upgrade.
  - When doing ISSU it is necessary to ensure ISSU is performed on one node at a time. BFD state might get stuck if both the peers are upgraded simultaneously.

- During ISSU, some of the BFD sessions may go down. To recover disable and enable BFD at interface level or remove and re-configure BFD session. This issue is related to timing of session delete and create and can even happen with few sessions. Probability of hitting the issue is higher with scaled configuration.

1G support on VDX6940-144S

- Starting 6.0.1a, VDX6940-144S supports 1G speed on 10G interfaces. In order to enable 1G support on VDX6940-144S, it is required to upgrade to 6.0.1a using coldboot option. If ISSU upgrade is performed, it is required to reload the switch before using 1G functionality.
- During the firmware upgrade and replaying of configuration, the interface may display an IP address of 255.0.0.0/8 momentarily, however, the original IP address will be restored once the upgrade is completed. We recommend not to make any configuration changes during the upgrade.
- On upgrading from 5.0.x OR 6.0.1x to 6.0.2, the ssh and telnet service may be enabled on upgrade, even though it was disabled before the upgrade. We recommend to re-apply the ssh and telnet configuration after the upgrade is complete.

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 supported.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.

- If Transparent VLANs (TVLANs) are associated to tunnels on VTEP configured on VDX6940, reloading the switch may result in prolonged TVLAN traffic loss. In order to recover, it is required to unconfigure and reconfigure TVLANs under overlay-gateway.
- Adding and Removing RBridges under overlay-gatway may take longer than expected time if large number of VLANs are configured in the fabric.
- During ISSU, some of the BFD sessions may go down due to high scaled BFD configurations. To recover disable and enable BFD at the interface level, or remove and re-configure BFD session.

Long Distance ISL

- The "long-distance-isl" functionalilty 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 ping issues.

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 the config.
- 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 Networks 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 certain conditions".

Port Channel Scalability:

- Port-Channel supported scale for VDX 6940 has been increased from 60 to 128.

Scale limit for configs with Extreme Port-channel:

- Under certain circumstances, port-channel configured with Extreme protocol, may limit the maximum scale number to a lower value.

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 manual discovery.
- Output of show vnetwork vmpolicy command is not displaying the VM name and datacenter-id for a cloned VM. Workaround is to initiate a manual discovery.

IGMP over Tunnel:

- In a single node VCS if physical/LAG port is not provisioned for IGMP snooping enabled vlan other than the Tunnel port IGMP groups will not learn over tunnel.

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-default rcv-queue.
- With large number of flows, "show openflow flow <>" may take 20 seconds to display packet 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 VDX 6940-144S.

FTP protocol support:

- FTP protocol has been enabled for VDX 2746.

Closed with code changes for NOS 6.0.2h

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS 6.0.2h.

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:	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:	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:	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:	DEFECT000658692		
Technical Severity:	High	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	Hardware Monitoring
Reported In Release:	NOS6.0.2	Technology:	Monitoring
Symptom:	Unexpected Line Card reload while collecting SS from BNA		
Condition:	Copy Support save CLI execution can lead to this issue.		

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:	DEFECT000660509		
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:	Ping loss across vxlan tunnel with Extreme trunks		
Condition:	Port channel between 2 rbridges was a Extreme trunk and other 2 rbridges was a standard lag.		

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" .		

Closed with code changes for NOS 6.0.2g

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of January 4th, 2018 in Network OS 6.0.2g.

Defect ID:	DEFECT000582371		
Technical Severity:	High	**Probability:**	High
Product:	Extreme Network OS	**Technology Group:**	VCS
Reported In Release:	NOS7.0.0	**Technology:**	Logical Chassis
Symptom:	IPv6 traffic drop might happens some time.		
Condition:	When "clear ipv6 neighbor no-refresh" command is used.		

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:	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:	DEFECT000636696	
Technical Severity: Medium		**Probability:** Medium
Product: Extreme Network OS		**Technology Group:** Monitoring
Reported In Release: NOS6.0.1		**Technology:** Hardware Monitoring
Symptom: RAS log has "Too many interrupts (501) happened" internal error message.		
Condition: ASIC error. Console error " Too many interrupts (501)" occur, on VDX 6740 and 6940 switches. May be due to improperly seated SFP+ / cable		
Workaround: Try NOSCLI shut / no shut. If this doesn't help, try reseating or replacing SFP+ and/or cable.		
Recovery: Try NOSCLI shut / no shut. If this doesn't help, try reseating or replacing SFP+ and/or cable.		

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:	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 SLX switch.		
Recovery: Try NOSCLI shut / no shut on VDX and or SLX switch.		

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

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:	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:	DEFECT000648291	
Technical Severity: High	**Probability:** Medium	
Product: Extreme Network OS	**Technology Group:** Management	
Reported In Release: NOS7.0.1	**Technology:** CLI - Command Line Interface	
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:	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: 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: DEFECT000654900	
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: 1G Port won't come online.	
Condition: Connected 1G with 10G at other end.	

Defect ID: DEFECT000657838	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.2	**Technology:** Scripting
Symptom: Importing the pydoc module through the python interpreter on the VDX switch or chassis results in "ImportError: No module named 'pydoc'".	
Condition: The pydoc python module was not included in the firmware distribution RPM package for the VDX switch or chassis.	

Closed with code changes for NOS 6.0.2f

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of July 10, 2017 in Network OS 6.0.2f.

Defect ID: DEFECT000567860	
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: Unexpected system failover or reload.	
Condition: This may be observed when there is a repeated MAC move between interfaces resulting in continuous MAC re-learning.	

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 interface details. It provides only port-channel details.	
Condition: This issue will happen only when the number of port-channels configured is equal to or more than 70. If the number of port-channels is fewer than 70, this issue will not be encountered.	
Workaround: Total number of port-channels configured should be fewer than 70.	
Recovery: If the total number of port-channels configured exceeds 70, delete port-channels to reduce the total count to fewer than 70.	

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 a reload of an LC may cause the VDX to abruptly reload due to a software daemon termination.	

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:	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 occurs, existing FCOE hosts are logged out.		
Condition: Adding/removing an ISL or adding/removing a switch from a 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:	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:	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:	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: DEFECT000638990	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.1	**Technology:** Logical Chassis
Symptom: After a reload, a VDX device is not able to join a cluster and is stuck in the "Awaiting Rejoin" state.	
Condition: VDX has "switchport port-security" CLI configured and a reload occurs in the 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 during an upgrade.	

Defect ID: DEFECT000641617	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS6.0.2	**Technology:** IGMP - Internet Group Management Protocol
Symptom: MAC addresses replicate on all VLANs with PIM and IGMP are enabled in network.	
Condition: When PIM and IGMP are enabled in network and PIM packets are coming through ISL.	

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.	

Closed with code changes for NOS 6.0.2e

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of April 10, 2017 in Network OS 6.0.2e.

Defect ID:	DEFECT000568368		
Technical Severity: High		**Probability:** Medium	
Product: Extreme Network OS		**Technology Group:** Layer 3 Routing/Network Layer	
Reported In Release: NOS5.0.2		**Technology:** IP Addressing	
Symptom: VDX does not generate NS because of which end-end transit ping fails.			
Condition	1. ipv6 transit packets 2. ingress and egress interface is on the same Ve		
Workaround: /proc/sys/net/ipv4/icmp_redirect_forward Change it to 1 for respective vrf.			

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:	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:	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:	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:	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: 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: 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	

Closed with code changes for NOS 6.0.2d

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of March 3, 2017 in Network OS 6.0.2d.

Defect ID: DEFECT000572975	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS7.0.0	**Technology:** OpenStack Integration
Symptom: FCOE links will flap/down	
Condition: If FCOE links connected through FSB and port-channel	

Defect ID: DEFECT000580639	
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: Firmware Download operation fail with "Firmware operation (5) was aborted due to timeout" HASM-1020 RAS log.
Condition: Firmware Download operation

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: DEFECT000583435	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.0	**Technology:** Logical Chassis
Symptom: In FCR, while rebooting the non-principal switch in the fabric, the host receives traffic disruption.	
Condition: Rebooting the non-principal switch In FCR.	

Defect ID: DEFECT000597202	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.1	**Technology:** VLAN - Virtual LAN
Symptom: Under certain conditions, show vlan <vid> may output a message as "application communication failure" and later impact the switch stability to go through an unexpected reload.	
Condition: Only applicable for GVLAN's.	

Defect ID: DEFECT000600171	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	**Technology:** VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: All RBs in a VCS act as VRRP Masters	
Condition: VRRP-E packets are dropping , due to this all RBs in a VCS act as Masters	

Defect ID: DEFECT000610251	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS4.1.3	**Technology:** ICMP - Internet Control Message Protocol
Symptom: Dcmd process termination can occur.	
Condition: A script which launches multiple simultaneous "copy running-config <file>" operations can trigger the Dcmd process to terminate. Manually invoking simultaneous operations will not hit the small time window achievable by a script.	
Workaround: Ensure that the script does not invoke multiple simultaneous "copy running-config <file>" operations to a given switch.	

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: 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: 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: 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: 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: 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: 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: 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: DEFECT000627263	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.2	**Technology:** VLAN - Virtual LAN
Symptom: Multicast drops when traffic from two ingress ports is sent towards one egress port.	
Condition: Traffic from two ingress ports is sent towards one egress port.	

Defect ID: DEFECT000629513	
Technical Severity: High	**Probability:** Medium
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	

Closed with code changes for NOS 6.0.2c

Defect ID: DEFECT000507145	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.0	**Technology:** ACLs - Access Control Lists
Symptom: IP access-list in egress direction does not work for CPU originated frames.	
Condition: IP access list configuration to match CPU originated frames.	

Defect ID: DEFECT000551893	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** Configuration Fundamentals
Symptom: VCS VIP becomes unreachable on VRF mode change	
Condition: VCS VIP becomes unreachable with IP address change on the associated VE interface or VRF mode change	

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: DEFECT000562447	
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: Unexpected Reload due to inconsistent maintenance of internal data.	
Condition: Issuing the CLI "show running-config rbridge-id <id> snmp-server" when configurations contain local SNMPv3 Host	
Workaround: Global SNMPv3 Host can be used for querying instead of Local SNMPv3 Host.	

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: DEFECT000567383	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** LAG - Link Aggregation Group
Symptom: Port may not be added to static port-channel.	
Condition: Port may not be added to port-channel when the link speed is changed.	

Workaround: Use LACP instead of LAG.

Defect ID: DEFECT000568285	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.1	**Technology:** NETCONF - Network Configuration Protocol
Symptom: Password with special character like $ will not work with copy running config.	
Condition: Have password with special character like $ and do "copy running config". The issue can be reproduced easily.	
Workaround: Workaround for this issue is not to use password contains Linux Specific special character like $, \, ' and `. Password without special character listed above will work fine.	

Defect ID: DEFECT000571346	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** LAG - Link Aggregation Group
Symptom: The maximum number of port-channel can be configured on 6940-144 is 128.	
Condition: The maximum number is raised to 144 on this release.	

Defect ID: DEFECT000576383	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	**Technology:** OSPFv3 - IPv6 Open Shortest Path First
Symptom: Error message at "no area 0.0.0.0" CLI	
Condition: When the interface does not have the OSPF configuration	

Defect ID: DEFECT000577179	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS7.0.0	**Technology:** PIM - Protocol-Independent Multicast
Symptom: PIM group registration fails between VDX and teh RP, thereby VDX failing to transition to SPT	
Condition: VDX is first-hop PIM router that sends erroneous group registration messages to RP.	

Defect ID: DEFECT000577716	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS7.0.0	**Technology:** Logical Chassis
Symptom: All user commands will result in "application communication failure" error.	
Condition: This can happen only when user commands are continuously executed during principal node change.	

Defect ID: DEFECT000577873	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: show tunnel <number> does not display loopback interface number. It displays as "loopback" instead of "loopback <number>"	
Condition: Execution of "show tunnel <number>" CLI.	

Defect ID: DEFECT000577986	
Technical Severity: Medium	**Probability:** Low

Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: LAG - Link Aggregation Group
Symptom: MAC learning may not happen in a protected VLAG over a trunk port.	
Condition: This happens when a protected VLAG is failed-over or failed-back.	

Defect ID: DEFECT000579487	
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: VDX 6940-36Q using 10 GbE adapter module part number MAM1Q00A-QSA may stay offline when link partner's port is toggled up/down (i.e., administratively disabled/re-enabled, or cable removed/inserted).	
Condition: VDX 6940-36Q using 10 GbE adapter module part number MAM1Q00A-QSA may stay offline when link partner's port is toggled up/down (i.e., administratively disabled/re-enabled, or cable removed/inserted).	
Recovery: Issue shut / no shut commands on the VDX6940-36Q 10 GbE port to bring the port back online.	

Defect ID: DEFECT000583304	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: BIDI optics won't work	
Condition: Enable the BIDI optics support	

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: DEFECT000584749	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: Unexpected failover	
Condition: Un-configure more number SNMP server context configurations in a short time	

Defect ID: DEFECT000584922	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.2	Technology: SSH - Secure Shell
Symptom: Shutting Telnet server on Active and Standby partition fails	
Condition: High Availability fail over	

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: 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: 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: 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: 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: 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: 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: DEFECT000594682	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.2	**Technology:** ACLs - Access Control Lists
Symptom: SNMP walk failure in some scenarios.	
Condition: Creating IP ACL with sequence id as 0 causes this issue.	
Workaround: Avoid using sequence id 0 while creating IP ACL.	

Defect ID: DEFECT000594819	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.1	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: Switch can experience an unexpected reload with HSL kernel backtrace.	
Condition: When VXLAN tunnels are deleted and then added again.	

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: DEFECT000595233	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VPN
Reported In Release: NOS7.0.1	**Technology:** EVPN - Ethernet VPN
Symptom: In very rare scenarios after ISSU upgrade traffic drops may be observed for Tunnel terminated traffic	
Condition: ISSU upgrade is a necessary condition for this issue. But not all ISSU upgrades will results in this issue	
Workaround: Perform disruptive firmware upgrades that involve reboots	
Recovery: Rebooting the switch will recover the system.	

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: DEFECT000595877	
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: Unexpected reload of switch observed when removing a VRF configuration or removing ipv4/ipv6 address family configuration of a VRF.	
Condition: When a custom VRF is unconfigured or IPv4/IPv6 address family of a VRF is unconfigured, switch will be reloaded.	

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: 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: 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: 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: 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: 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: 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: DEFECT000600023	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.2	**Technology:** SSH - Secure Shell
Symptom: Shutting SSH server on Standby partition fails	
Condition: After High Availability fail over, we may hit the issue.	

Defect ID: DEFECT000600696

Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.2	Technology: Hardware Monitoring
Symptom: Unable to run the RTE tool for CBR2 platform	
Condition: While executing the RTE tool on CBR2 platforms.	

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: 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: 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: DEFECT000605230	
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: After ISSU upgrade the configuration "ipv6 nd prefix 2011::/64 2592000 86400 no-autoconfig" no longer works.
Condition: Internal configuration data didn't sync properly.

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: 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: 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: 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: DEFECT000612673	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.1.0	**Technology:** VLAN - Virtual LAN
Symptom: May see spurious "Too many interrupts" events.	
Condition: Other interrupts come in within a second and first one not cleared.	

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: 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: 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: 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: 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: 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: DEFECT000617313	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.1	Technology: TRILL - Transparent Interconnection of Lots of Links
Symptom: RTE capture won't work for breakout interface	
Condition: when ingress/trill port is breakout mode	

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: 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: 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: 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: 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 FC port.	
Condition: Can happen only on FC port.	

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: 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: DEFECT000625386	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.2	Technology: Logical Chassis
Symptom: Unable to login to the device and customer should do netinstall to bring up the device	
Condition: After upgrading the firmware form version nos6.0.2 to nos7.0.0, user unable to login the switch.	
Workaround: Upgrade the firmware version from nos6.0.2 to nos6.0.2c before upgrading it to nos7.0.0 or higher version.	
Recovery: Netinstall is required to recover the switch state.	

Defect ID: DEFECT000626712	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.2	Technology: PIM - Protocol-Independent Multicast
Symptom: VDX experience unexpected reload due to pimd daemon termination.	
Condition: Protocol Independent Multicast [PIM] enabling on VDX can cause memory corruption.	

Closed with code changes for NOS 6.0.2b

Defect ID: DEFECT000516373	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS4.1.2	**Technology:** Logical Chassis
Symptom: Command field in TACACS accounting log does not display the user command from VDX correctly.	
Condition: Examining Command field in TACACS accounting log.	

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: DEFECT000549157	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** Configuration Fundamentals
Symptom: An unexpected reload may occur if a user tries to remove a child zoning cfg/zone/alias object from a parent when the parent happens to have a duplicate of that child object.	
Condition: User will experience this issue if adding multiple members to a parent cfg/zone/alias object using the semicolon-separated multi-member method and creating a scenario where duplicate members are added to a parent zone object.	
Workaround: To avoid this problem in previous version, only add zone members one at a time. With the fix in place, there should be no restrictions.	

Defect ID: DEFECT000554155	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.1	**Technology:** SSH - Secure Shell
Symptom: Garbled characters seen on SSH session during login	
Condition: SSH login to management interface	

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: DEFECT000559794	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS6.0.1	**Technology:** NETCONF - Network Configuration Protocol

Symptom: Switch may reload when BNA queries with get-config when the config is large OR when BNA polls at an aggressive rate (not configured for lazy-polling).
Condition: Issue can happen when BNA does get-config for large cluster which has more than 4500+ interface config OR when BNA is polls aggressively.
Workaround: Please do not use BNA if the cluster is large (with 4200+ interfaces) & ensure it is configured with lazy-polling.

Defect ID: DEFECT000560868	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** IP Addressing
Symptom: IP directed-broadcast feature is not working as expected.	
Condition: With a regular topology, the functionality did not work as expected.	

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: 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: 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: DEFECT000566534	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: The configuration command "no snmp trap link-status" does not show up under the interface sub-mode when a show running config command is executed, even though it has been executed under the interface sub-mode. This command is available under Tengigabit ethernet, Forty Gigabit Ethernet, Gigabit Ethernet, Loopback and Ve interface submodes.	
Condition: This issue is seen only when the command "no snmp trap link-status" is executed under interface submode and show running config is then executed.	

Defect ID: DEFECT000566855	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: BGP4+ - IPv6 Border Gateway Protocol
Symptom: IPv6 prefix routes are not advertised to its BGP peer.	
Condition: When route-map is configured and deleted and configured, it might be possible the route passing via the route-map might get denied. Due to this ipv6 routes are not advertised to its neighbor.	
Workaround: Create a new route-map with different name.	

Defect ID: DEFECT000567363	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS5.0.1	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: Upon bootup, show version output may rarely show only one GOS in a 6740.	
Condition: Repeated power cycles.	
Recovery: "reload system" will recover from the issue.	

Defect ID: DEFECT000569413	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS5.0.2	Technology: NETCONF - Network Configuration Protocol
Symptom: The switch may experience an unexpected reload while repeatedly accessing get-system-uptime via netconf.	
Condition: Repeated access of get-system-uptime via netconf	
Workaround: Limit the usage of get-system-uptime calls.	

Defect ID: DEFECT000570631	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.2	Technology: LAG - Link Aggregation Group
Symptom: Continuous Ping stops for about 11-15 seconds when VDX6740T ports are gracefully shutdown through reloads/chassis disable operation. >> Over 36 channel-group in total makes 5 seconds downtime when reload. >> Over 40 channel-group in total makes 10 seconds downtime when reload. >> Over 44 channel-group in total makes 15 seconds downtime when reload.	
Condition: When vLAG members are on VDX 6740T port number 36 or higher, the vLAG failover time during port disable operation is 5 seconds or more. The failover time is 15 seconds for ports > 44. Issue is not seen if the port number is < 36. Issue is seen with copper ports on VDS 6740T/VDX6740T-56-1G and 48x10BaseT line cards.	
Workaround: None when vLAG is created with members belonging to port numbers > 36	

Defect ID: DEFECT000571163	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Syslog
Symptom: L2SS-1023 messages indicating internal MAC inconsistency within VCS which will be attempted to be auto recovered (without any user intervention/action) are unnecessarily logged with WARNING severity (instead of INFO)	

Condition:	When there is internal MAC inconsistency within VCS and L2SS-1023 is logged to indicate this. (the message is logged with WARNING severity instead of INFO)

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: DEFECT000573549	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS5.0.1	Technology: OpenStack Integration
Symptom: Two line cards went to faulty (119) status after firmware upgrade	
Condition: IPv4 and IPv6 ACLs configured with TCP/UDP Port range operators neq, gt, lt, range and applied to interfaces.	
Workaround: Remove IPv4 and IPv6 ACLs configured with TCP/UDP Port range operators neq, gt, lt, range and save it to startup configuration before firmware upgrade.	
Recovery: Remove IPv4 and IPv6 ACLs configured with TCP/UDP Port range operators neq, gt, lt, range and save it to startup configuration and reload.	

Defect ID: DEFECT000573626	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.2	Technology: AAA - Authentication, Authorization, and Accounting
Symptom: TACACS+ accounting replies does not provided the needed information to properly identify the ifDescr and/or ifname for user to take appropriate action, It returns nothing for dev/port/0 as it should be in 3 tuple format with rb-id/slot/port so users can check the appropriate device once the accounting entries is received for an events,	
Condition: TACACs+ command audit logging only contained the user command issued. In the case of "shut" or "no shut" commands issued on an iterface	
Workaround: Add modeInfoInAudit flag set to true in /etc/confd/confd.conf file. Have confd reparse the config changes using /usr/confd/bin/confd_cmd -c reload	

Defect ID: DEFECT000573869	
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: Certain flows from the routed traffic traversing VDX may not be delivered correctly to the destination.	
Condition: When the total number of ARP entries exceed 3K, and presence of frequent ARP churn. This is specific to VDX6940 only.	
Recovery: clear arp with no-refresh option would recover this issue for a short time.	

Defect ID: DEFECT000573943	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: Static Routing (IPv4)

Symptom: Ping to VRRP virtual-IPs address has failed in a VXLAN VTEP topology.
Condition: In a VXLAN VTEP topology configured with port-channel redundancy group, unexpected ping loss is observed when one of the participating nodes is reloaded.

Defect ID: DEFECT000573960	
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: Very rare scenario, customer can not delete the ACL	
Condition: After reloading the node with default-config	

Defect ID: DEFECT000574087	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VXLAN - Virtual Extensible LAN
Symptom: When VXLAN Tunnel lost its underlay path (L2/L3), there is small time frame when VXLAN tunnel management routine waits before it can declare the tunnel as operational down. This is to detect any momentary flap in the network. This time frame is called "debounce" time. This is helpful because momentary flaps and immediate tunnel events cause lot of churn in the system. Having a debounce time helps to detect if the loss of underlay reach-ability is real and not momentary.	
Condition: When "debounce" time is helpful to detect any momentary failures, it also causes delay in traffic switchover to the other node in active-active gateway and traffic loss for the "debounce time" period. In some network deployment this traffic loss may be unacceptable. Hence this defect/feature exposes a CLI for switching off debouncing logic. This will eliminate any delay between underlay path unavailability and tunnel operation state change.	

Defect ID: DEFECT000574823	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.1	**Technology:** Logical Chassis
Symptom: STP BPDUs won't pass through the fabric though we enable 'tunnel tagged-ieee-bpdu'	
Condition: After configuring the CLI 'tunnel tagged-ieee-bpdu'	

Defect ID: DEFECT000575385	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Unable to remove SNMP user under Rbridge configuration.	
Condition: When SNMP user configured at global and local but it is already deleted from Global.	
Workaround: Do not delete the Global user before deleting at local .	

Defect ID: DEFECT000575438	
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: BUM traffic received on backup port-channel of a protected group vLAG is processed and forwarded via primary port-channel (on another Rbirdge) causing a loop	
Condition: When protected vLAG feature is used without explicit active port channel configuration and Rbridges leaving & joining the fabric.	
Workaround: Explicitly configure the active & backup port channels of the protected group.	
Recovery: Activate any one specific port channel of the protected vLAG group	

Defect ID: DEFECT000575461

Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS4.1.3	Technology: IP Addressing

Symptom: Unexpected reload while deleting the VLAN.

Condition: After upgrade from NOS4.0.0b to any ver >= NOS4.1.3 with one of the interface in 'switchport-private-vlan-trunk-promiscuous' mode .

Workaround: After the upgrade, execute 'no switchport' and restore the previous configs of the interface.

Defect ID: DEFECT000575656

Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: FCoE - Fibre Channel over Ethernet

Symptom: Both the FCoE and CEE provisioning are not supposed to be allowed on the port-profile. But it is allowed.

Condition: The FCoE and CEE map provisioning was allowed on the port-profile which will create unexpected results.

Workaround: Don't try to add FCoE and CEE map provisioning in same profile.

Recovery: User has to remove either CEE or FCoE provision from the port-profile to come out of this situation.

Defect ID: DEFECT000575922

Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.2	Technology: VXLAN - Virtual Extensible LAN

Symptom: IP connectivity between hosts in two different VCS, which are connected through vxlan tunnel , drops for the duration of 3-7 seconds.

Condition: Ping/connectivity drops are observed when one of the vxlan tunnel end point Rbridge is rebooted and is in the process of joining the vxlan tunnel.This problem is only seen when VLAG/pVLAG is used as underlay, and VRRP IP is used for Layer 3 connectivity between tunnel endpoints.

Defect ID: DEFECT000576207

Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS5.0.2	Technology: Logical Chassis

Symptom: If storm control is configured on edge ports which will become ISL ports, then ISL flapping may cause unexpected reload.

Condition: Issue is seen only during ISL flapping with storm control config is applied on them.

Workaround: Remove the storm control configuration on edge ports which may get transitioned to ISL ports during cluster formation.

Defect ID: DEFECT000577094

Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS5.0.2	Technology: IPv4 Multicast Routing

Symptom: When a secondary IP is not configured, VDX sends option 24 in PIMV2 hello header which results in PIM adjacency failing to form on other switch.

Condition: Happens with secondary IP address not configured.

Defect ID: DEFECT000577323

Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals

Symptom: QoS Values in QoS profile are not being applied automatically.
Condition: On configuring a vCenter, the expected QoS Values in QoS profile are not applied automatically.
Workaround: There is no work-around but a recovery method is available.
Recovery: We do not recommend manually interfering in automatic vCenter configuration but here the only workaround is manually applying the expected QoS values in QoS profile.

Defect ID: DEFECT000577414	
Technical Severity: Medium	**Probability:** Low
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: A possible ECMP path cannot be established after a cluster reload since OSPF changes the forwarding address in its NSSA Type7 LSA.	
Condition: On reloading multiple OSPF neighbors at same time, OSPF routers selects a different forwarding address than existed before reload even if there is an internal Loopback address present on those nodes.	
Recovery: Interface shutdown/no-shutdown on the network corresponds to currently selected forwarding address will resolve the issue.	

Defect ID: DEFECT000577795	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.1	**Technology:** VCS Fabric
Symptom: Customer may experience unexpected reload due to Dcmd daemon termination.	
Condition: If SNMP "admin" user has configured at Global level as well as on Local level as below then BNA/user tries to configure SNMP v3 host with the same username "admin", VDX Dcmd daemon terminates. snmp-server user admin groupname snmpadmin … rbridge-id 41 snmp-server user admin groupname snmpadmin …	

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: DEFECT000577908	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.2	**Technology:** UDLD - Uni-Directional Link Detection
Symptom: After the hafailover/ISSU upgrade, ARP is not resolving behind the port-profile ports	
Condition: Below both conditions have to meet : 1. Apply the port-propfile-port 2. After hafailover/ISSU upgrade	
Recovery: Remove the port-profile-port and add it back	

Defect ID: DEFECT000577986	
Technical Severity: Medium	**Probability:** Low

Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: MAC learning may not happen in a protected VLAG over a trunk port.	
Condition: This happens when a protected VLAG is failed-over or failed-back.	

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: DEFECT000579664	
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: "waiting for pending actions to exit" warning message appears on console session and eventually VDX experience unexpected reload.	
Condition: When user query to get any running-config using netconf and if password contains a special char like ";" , VDX throws "waiting for pending actions to exist" error and DCM gets terminated eventually.	

Defect ID: DEFECT000579668	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology:** DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP relay stopped working on ve interface and DHCP Discover not being forwarded.	
Condition: Configuring gateway IP addres with DHCP Relay	

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: Commad set field on the Winows based TACACS server is empty	
Condition: 1. When TACACS server is windows based 2. Accounting is enabled	

Defect ID: DEFECT000581707	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Switch unexpectedly reloads while configuring 'ipv6 mld snooping mrouter interface GigabitEthernet' for an 'interface vlan'.	
Condition: Configuring 'ipv6 mld snooping mrouter interface GigabitEthernet' for an 'interface vlan' multiple times results in unexpected switch reload.	

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: DEFECT000581851	
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: In chassis based VDX switches, if any host is configured as a trap recipient, then the SNMP traps are received with active IP as the source IP address.	
Condition: This issue is observed only on a chassis based switches, when any host is configured as a trap recipient.	

Defect ID: DEFECT000581852	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: SNMP query ifDescr output will be partial.This is specific to VDX6940-144S platform.	
Condition: SNMPwalk for ifDescr & Ifname table.	

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: DEFECT000582731	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS6.0.2	**Technology:** IGMP - Internet Group Management Protocol
Symptom: Unexpected reload in corner case.	
Condition: configure static igmp group on Physical L3 interface	

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: DEFECT000584060

Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: IGMP - Internet Group Management Protocol
Symptom: A request [Ex: NETCONF request] made to create IGMP static group will be rejected.	
Condition: A request [Ex: NETCONF request] made to create IGMP static group prior to 6.0.1a will be rejected in 6.0.1a and the request made to create a static group in 6.0.1a will be rejected by earlier NOS.	

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: 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: DEFECT000585841	
Technical Severity: High	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.	

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: 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: DEFECT000587380	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: CLI - Command Line Interface
Symptom: "no snmp trap link-status" does not show up in show running config after executing it under an interface level.	
Condition: "no snmp trap link-status" under interface is enhanced and it got missed to keep the same in running config.	

Defect ID: DEFECT000587653	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: xSTP - Spanning Tree Protocols
Symptom: Although BPDU drop has been configured, VDX transmit STP BPDU.	
Condition: BPDU-drop doesn't work for egress even though it [bpdu-drop enable all] is configured.	

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: DEFECT000588222	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.0.0	Technology: ACLs - Access Control Lists
Symptom: Message Generic Error during removal of IP ACL from management interface.	
Condition: Unable to remove IP ACL from management interface after applying it.	
Workaround: Remove the ACL rules from backend through ipfilter commands and replay the same ACL through NOSCLI if needed.	

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: 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: 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: 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: 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: DEFECT000590771	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: High Availability

Symptom: The standby MM on 8770-8 went to faulty state. This caused switch HA failover not to work. Also seen on primary MM was DCMD daemon termination, causing primary MM to reboot.
Condition: The primary MM was booted with wrong Model ID, resulting in communication failure with secondary MM and database corruption.

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 the expected 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: 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: 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: 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: DEFECT000595653	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	**Technology:** IP Addressing
Symptom: IP Directed broadcast would not work after HA failover, but the CLI configuration may present.	
Condition: HA fail-over trigger the issue.	
Recovery: Reconfigure ip directed-broadcast	

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: 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: 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.	

Closed with code changes for NOS 6.0.2a

Defect ID: DEFECT000566819	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** VLAN - Virtual LAN
Symptom: After reloading VDX6940/VDX6940-144S, TVLAN ctags are not retained at egress & traffic may be affected.	
Condition: On reloading of VDX6740 / VDX6940.	

Defect ID: DEFECT000573258	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS5.0.1	**Technology:** Logical Chassis
Symptom: Under certain rare conditions, switch may encounter an unexpected reload due to memory exhaustion when being polled by BNA.	
Condition: When aaa accounting feature is configured, it would result in unexpected memory growth due to continuous polling of BNA.	

Workaround: Customer should have to remove below configuration that will reduce/stop DCMd memory increment.
 	
    sw0(config)# no aaa accounting exec default start-stop tacacs+
    sw0(config)# no aaa accounting commands default start-stop tacacs+	

Defect ID: DEFECT000575242	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Software Installation & Upgrade
Symptom: Show process cpu reports high CPU [~50%] even though there is minimal activity and traffic running.	
Condition: When certain timing conditions are hit on the bootup of the switch.	

Defect ID: DEFECT000576916	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: In the output of CLI "show interface status" status of a port-channel interface is sometime shown as 'Notconnected' even though some of its member interfaces are up.	
Condition: On execution of CLI "show interface status".	

Defect ID: DEFECT000577861	
Technical Severity: Medium	**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: Enhance DHCP Relay Option-82 to support <string> as sub-option field. Circuit-ID value to change from <ifindex> to <vlan-id,if_description_string>. The string is configured as Interface Description under the vlan interface for the respective layer-3 VE interface.	
Condition: When DHCP Option 82 needs to be configured with a string.	

Defect ID: DEFECT000577922	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast

Reported In Release: NOS6.0.2	Technology: PIM - Protocol-Independent Multicast
Symptom: In case of large no. of multicast groups, when traffic is sent from 2 different source, sometimes there would be traffic loss and S,G entries missing	
Condition: Around 1008 S,G entry and more than 1 source	
Workaround: Send traffic from 1 source	
Recovery: Reload the receiver DR	

Defect ID: DEFECT000578607	
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: When configuring Ve interface, the switch may encounter an unexpected reload.	
Condition: When Ve interface is created.	

Defect ID: DEFECT000578692	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.2	**Technology:** SSH - Secure Shell
Symptom: The ssh/telnet becomes enabled after upgrade from 5.0.1x or 6.0.1x to 6.0.2 even though it was blocked before upgrade	
Condition: On firmware upgrade.	

Defect ID: DEFECT000579510	
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: For VDX6740, only first 24 flexports that are converted into FC port can accept the FC connections. The end devices connected to FC ports that are beyond 24 operational flexports, may not be able to communicate to the other end devices.	
Condition: When more than 24 flexports are configured for FC ports.	

Closed with code changes for NOS 6.0.2

Defect ID: DEFECT000493809	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: The VDX response to SNMP queries with destination address as either management address/next-hop subnet to source of request.	
Condition: SNMP query for an inband IP which is not directly connected to host (i.e source of request)	
Workaround: In inband network, we can have both switch and host in same network by connecting via a physical link.	

Defect ID: DEFECT000499981	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS4.1.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: When using inband management, SNMP will report using the VCS VIP and the syslog will report the VE IP.	
Condition: Managing the switch using inband management.	

Defect ID: DEFECT000543072	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS5.0.1	**Technology:** VCS Fabric
Symptom: On upgrade from 4.x to 5.x, the IPv6 ACL rule stating "seq 1 permit ip any any" would be converted into "seq 1 permit any any".	
Condition: When upgrading from 4.x to 5.x release.	

Defect ID: DEFECT000544034	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS5.0.1	**Technology:** IPv4 Multicast Routing
Symptom: VDX last hop router may unexpectedly send Prune messages, interrupting the traffic to the receivers.	
Condition: In topologies where Multicast source is connected to the same VLAN as receivers & RP VE interface.	

Defect ID: DEFECT000549696	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Introducing a new configuration command to enable or disable SNMP traps on a per interface basis on VDX platforms.	
Condition: New feature support.	

Defect ID: DEFECT000551273	
Technical Severity: Medium	**Probability:** Low
Product: Extreme 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: "show interface status " shows all the port-channel details.

Defect ID: DEFECT000558457	
Technical Severity: Critical	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS5.0.1	**Technology:** OpenStack Integration
Symptom: Cisco UCS fabric discovery may stall if and when said unit issues a RIP_NN FC-GS query to the VDX. The RIP_NN query is marked as being obsolete by current FC Standards (FC-GS-6), thus it is rejected by the VDX. It has been observed in some cases that this reject causes the Cisco UCS product to halt any further discovery operations. Ultimately this manifests as no connectivity for any initiators entering the cluster via the Cisco UCS product.	
Condition: Affects Cisco UCS deployments with VDX.	

Defect ID: DEFECT000559106	
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: Switch will go for unexpected reload after multiple times of config/unconfig syslog with default-vrf.	
Condition: When switch is configured/unconfigured with syslog in default-vrf multiple times then issue will occur.	

Defect ID: DEFECT000561024	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.2	**Technology:** SSH - Secure Shell
Symptom: Unexpected reload	
Condition: Continuous "ssh server standby enable/disable"	

Defect ID: DEFECT000562938	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** Configuration Fundamentals
Symptom: Enhanced the mini supportsave by adding new group.	
Condition: Cannot collect the new group logs as part of mini supportsave	
Workaround: collecting full copy support	

Defect ID: DEFECT000563327	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS5.0.1	**Technology:** Hardware Monitoring
Symptom: SFPs model number 57-1000042-01 and 57-1000042-02 are not supported on 6740 switches at both 100MB and 1G.	
Condition: New feature and/or hardware support.	

Defect ID: DEFECT000563667	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS4.1.3	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom: Under certain conditions, the ha sync between the MM's may not be achieved & thus all configuration may not be synced.	

| **Condition:** When FFDC is disabled in configuration. |
| **Workaround:** Enable the FFDC in configuration. |

Defect ID: DEFECT000564080	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Configuration Fundamentals
Symptom: Unexpected Line card reload	
Condition: Due to link flap we will hit OOM	
Recovery: Shut the flapping link	

Defect ID: DEFECT000564304	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.1	**Technology:** VLAN - Virtual LAN
Symptom: Switchport configs do not have effect and all traffic destined to these affected ports will be dropped.	
Condition: Doing a switchport configs on interfaces that doesn't have sfp or cable inserted and doing firmware upgrade wont bring back the interface online when sfp or cable is inserted at later stage.	
Recovery: Doing a shut and no shut on the affected interfaces brings the interface online again.	

Defect ID: DEFECT000564313	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: The Modulename for the corresponding slots are displayed correctly in the output of entPhysicalMfgName and it s in sync with Type filed in show slot output	
Condition: Before the modulename dispalyed in entPhysicalMfgName is not in sync with show slot output Type filed.	

Defect ID: DEFECT000564964	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Request for support for Fabric Inter-switch Links (ISLs) related information in a new MIB	
Condition: This feature would be required to view the details of the ISLs.	
Workaround: There is no MIB object to get this information.	

Defect ID: DEFECT000565415	
Technical Severity: Medium	**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:** DEFECT000566943 |

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: Unable to execute 'vlan' commands, even after creating a read-write accept rule for a user defined role.	
Condition: To define RBAC support to VLAN command.	

Defect ID: DEFECT000567038	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.0.0	Technology: SSH - Secure Shell
Symptom: Request for support for SSH/Telnet in non-default VRF.	
Condition: To configure specific VRF on SSH/TELNET.	

Defect ID: DEFECT000567262	
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: Length of community string for SNMPv1 communities is increased to fit 64 character long community names.	
Condition: Extended the support to enable the customer have long SNMP community names.	

Defect ID: DEFECT000567774	
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: By default SNMP uses relevant VE IP address as a source IP address during trap packet send. We can now configure source-interface VE/Loopback to use as a source interface.	
Condition: VDX SNMP default behavior is to select relevant VE IP address as a source IP address during trap packet send. Now we can select what VE/Loopback interface to use.	

Defect ID: DEFECT000567823	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VLAN - Virtual LAN
Symptom: Sometime execution of command "show interface description" hangs. This happens when there is improper output to be processed.	
Condition: On hang of execution of "show interface description" command.	

Defect ID: DEFECT000567846	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: Licensing
Symptom: Timeouts in other CLI responses after "show license id" CLI was issued and was non-responsive.	
Condition: Issuing "show license" CLI after the WWN card had a hardware failure causes timeouts in other CLI's	
Recovery: Reload the system, or power cycle the system to recover. If needed run diagnostics to test the WWN card hardware integrity.	

Defect ID: DEFECT000568047	
Technical Severity: Medium	Probability: Low

Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol

Symptom: Hafailoverstart/ End traps will be seen as part of cpststatuschange trap during init time. The flow of traps will be hafailoverstart/ warmstart/ hafailoverend traps.	
Condition: There is no trap which indicates that hafailover has happened. In warmstart trap, there are no varbinds.	

Defect ID: DEFECT000568523	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS5.0.1	Technology: sFlow

Symptom: In rare cases, a VDX switch may experience an unexpected failover, especially when sFlow is enabled.	
Condition: This could happen after months or years after enabling sFlow.	
Workaround: Disable sFlow or lower the SFlow sample rate.	

Defect ID: DEFECT000569661	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: VLAN - Virtual LAN

Symptom: Under certain condition, the internal port may become enabled even though one of the track ports is disabled.	
Condition: On switch reload after enabling the tracking.	
Workaround: Preventive workaround is to remove the track related configuration before reload.	

Defect ID: DEFECT000570137	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals

Symptom: Syslog messages & sflow traps may not be sent when configured for default-vrf.	
Condition: Only when there is no IP address configured on Management Interface.	
Workaround: Configure dummy IP address on the management interface (in mgmt-vrf) with the IP address belonging to the same subnet as syslog / sflow host.	

Defect ID: DEFECT000570270	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS4.0.1	Technology: LAG - Link Aggregation Group

Symptom: Discrepancy in Port-channel programming causes links within the port-channels to be seen as individual link instead of one link.	
Condition: Linecard that contains all of the ISL's connecting from that switch to the fabrics multicast root must be reloaded. When this situation occurs all PO's that are residing on that switch that contain PO members on other linecards other than the linecard that reloaded will experience improper flooding of BUM	
Recovery: Bounce any affected port-channel member interfaces one by one.	

Defect ID: DEFECT000570457	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: SNMP - Simple Network Management Protocol

Symptom:	SNMP MIB walk to ifName would not return all interface list after SNMP daemon gets restarted.
Condition:	If SNMP daemon gets restarted due to some error then SNMP MIB walk to ifName would not return all interface list.
Recovery:	We can perform HA failover to recover the issue under Extreme TAC guidance.

Defect ID: DEFECT000570673	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: UDLD - Uni-Directional Link Detection
Symptom: UDLD blocks links which are connected using certain breakout ports between different Extreme devices such as between VDX and CER.	
Condition: UDLD blocks the link after detecting mismatch between locally stored vs received port numbers in UDLD PDUs.	

Defect ID: DEFECT000571172	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: Static Routing (IPv4)
Symptom: Static route with more superior mask is not installed	
Condition: When static route is configured, if route with less mask exists, then new route is not installed in RTM	

Defect ID: DEFECT000571382	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: Sflow using default-vrf for inband connection does not function correctly when there is no ip address configured for management interface (OOB port).	
Condition: When there is no ip address configured for management interface.	
Workaround: Configuring the ip address on management interface.	

Defect ID: DEFECT000571388	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.2	Technology: IPv4 Multicast Routing
Symptom: When certain corner case disruptive events occur, a full cleanup of multicast route table is triggered. During the post recovery of those events, a partial traffic loss can be observed.	
Condition: This issue is observed when a disruptive event such as "chassis disable" or "shutting down RPF interface" is executed on a switch with large number of multicast routes.	
Recovery: Reload the affected switch.	

Defect ID: DEFECT000571714	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.2	Technology: LDAP - Lightweight Directory Access Protocol
Symptom: Switch could not import LDAP certificate in FIPS mode.	
Condition: When trying to import LDAP certificate in FIPS mode, the operation fails.	

Defect ID: DEFECT000571889	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.2	Technology: SSH - Secure Shell

Symptom: A VDX switch may undergo unexpected reload while configuring "ssh server use-vrf"
Condition: When VRF name of 32 characters or more is used.
Workaround: Reduce the characters in VRF name before configuring.

Defect ID: DEFECT000571906	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.2	**Technology:** FIPS - Federal Information Processing Standards
Symptom: Switch is not ready for configuration	
Condition: Crash in Dcmd is observed while FIPS zeroize operation in progress	

Defect ID: DEFECT000571981	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: Traffic drops for some vLAGs when one of the rbridges went for reload.	
Condition: Reload of one rbridge with vLAG scale configuration.	
Recovery: Port Shut and no shut will recover	

Defect ID: DEFECT000571982	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** IPv6 Addressing
Symptom: 'ipv6 nd suppress-ra' command configuration in interface mode will throw following error. %%% Error: VRF Address Family not configured	
Condition: Interfaces which belong to non-principal switches only will not allow this configuration.	

Defect ID: DEFECT000572287	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** Software Installation & Upgrade
Symptom: Firmware download won't be successful on 6740-1G and 6940 platforms.	
Condition: Firmware download functionality is affected under all circumstances for 6740-1G & 6940 platforms.	

Defect ID: DEFECT000572524	
Technical Severity: Medium	**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 debug messages show up on the console. The fix takes care of moving debug messages to RAS infrastructure.	
Condition: Some of the FCoE debug messages show up on console. They do not cause any functional issue.	

Defect ID: DEFECT000572746	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS5.0.1	**Technology:** Hardware Monitoring
Symptom: Customer may experience unexpected reload or some daemon termination.	
Condition: We don;t compress and delete confd core files till now and due to that our switch can experience disk full issue.	
Recovery: We have to manually remove confd core files from the disk.	

Defect ID: DEFECT000572990	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.2	Technology: SSH - Secure Shell
Symptom: Request to support for SCP file transfer without using root user.	
Condition: To move configuration files between server and switch without using root user.	

Defect ID: DEFECT000573129	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.2	Technology: Logical Chassis
Symptom: Confd prints internal log message on Console.	
Condition: In normal condition customer don't see internal log on console, but in some corner cases confd prints internal log on console.	

Defect ID: DEFECT000573171	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.2	Technology: SNMP - Simple Network Management Protocol
Symptom: On VDX8770 ha failover, IPV6 v1/v2c traps may not be sent	
Condition: On failing over the MM for VDX8770	

Defect ID: DEFECT000573379	
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: Service stopped in particular server/storage where it connected to VDX's particular port.	
Condition: VDX's Back End port link loss	
Recovery: Line card reset	

Defect ID: DEFECT000573950	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: Port Mirroring
Symptom: Monitor session could not be deleted after reload, under certain conditions.	
Condition: The issue occurs when a RSPAN-VLAN is configured as destination and a port-channel, VxLAN or RBridge (system flow based QOS) is configured as a source in a monitor session configuration.	
Workaround: If destination for a monitor session is rspan-vlan, remove that monitor session from port-channel, Vxlan and Rbridge(System Flowbased QOS),if applied on any of these.	

Defect ID: DEFECT000574606	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: LAG - Link Aggregation Group
Symptom: Customer can experience unexpected VDX reload.	
Condition: VDX experiences OOM condition when the switch receive multicast traffic and failed to free the received packet buffer due to any internal error condition.	
Workaround: If VDX has configured as an RP then please remove it as that is unsupported config.	

Defect ID: DEFECT000575436

Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.2	Technology: VCS Fabric
Symptom: No error messages are displayed when the user wrongly configures "fcoeport default" command on an ISL port.	
Condition: The issue is seen when the user configures a certain non-applicable command on an ISL port.	

Defect ID: DEFECT000576481	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Hardware Monitoring
Symptom: When using Mellanox/QSA adapter and Twinax cable on VDX6740 / VDX8770, the link may fail to come up.	
Condition: After reload of the VDX switch with POST diagnostics enabled.	
Workaround: Configure to disable POST on bootup to prevent this issue: "sw0(config)# no diag post rbridge-id <> enable "	
Recovery: May need to replace the cables.	

Defect ID: DEFECT000551273	
Technical Severity: Medium	**Probability:** Low
Product: Extreme 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: "show interface status " shows all the port-channel details.	

Defect ID: DEFECT000558457	
Technical Severity: Critical	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS5.0.1	**Technology:** OpenStack Integration
Symptom: Cisco UCS fabric discovery may stall if and when said unit issues a RIP_NN FC-GS query to the VDX. The RIP_NN query is marked as being obsolete by current FC Standards (FC-GS-6), thus it is rejected by the VDX. It has been observed in some cases that this reject causes the Cisco UCS product to halt any further discovery operations. Ultimately this manifests as no connectivity for any initiators entering the cluster via the Cisco UCS product.	
Condition: Affects Cisco UCS deployments with VDX.	

Defect ID: DEFECT000561024	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS5.0.2	**Technology:** SSH - Secure Shell
Symptom: Unexpected reload	
Condition: Continuous "ssh server standby enable/disable"	

Defect ID: DEFECT000562938	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.2	**Technology:** Configuration Fundamentals
Symptom: Enhanced the mini supportsave by adding new group.	
Condition: Cannot collect the new group logs as part of mini supportsave	
Workaround: collecting full copy support	

Defect ID: DEFECT000563327	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS5.0.1	**Technology:** Hardware Monitoring
Symptom: SFPs model number 57-1000042-01 and 57-1000042-02 are not supported on 6740 switches at both 100MB and 1G.	
Condition: New feature and/or hardware support.	

Defect ID: DEFECT000564080	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Configuration Fundamentals
Symptom: Unexpected Line card reload	
Condition: Due to link flap we will hit OOM	
Recovery: Shut the flapping link	

Defect ID: DEFECT000564304	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS5.0.1	**Technology:** VLAN - Virtual LAN
Symptom: Switchport configs do not have effect and all traffic destined to these affected ports will be dropped.	
Condition: Doing a switchport configs on interfaces that doesn't have sfp or cable inserted and doing firmware upgrade wont bring back the interface online when sfp or cable is inserted at later stage.	
Recovery: Doing a shut and no shut on the affected interfaces brings the interface online again.	

Defect ID: DEFECT000564313	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: The Modulename for the corresponding slots are displayed correctly in the output of entPhysicalMfgName and it s in sync with Type filed in show slot output	
Condition: Before the modulename dispalyed in entPhysicalMfgName is not in sync with show slot output Type filed.	

Defect ID: DEFECT000564964	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Request for support for Fabric Inter-switch Links (ISLs) related information in a new MIB	
Condition: This feature would be required to view the details of the ISLs.	
Workaround: There is no MIB object to get this information.	

Defect ID: DEFECT000565415	
Technical Severity: Medium	**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: DEFECT000566943	
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: Unable to execute 'vlan' commands, even after creating a read-write accept rule for a user defined role.	
Condition: To define RBAC support to VLAN command.	

Defect ID: DEFECT000567038	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS7.0.0	**Technology:** SSH - Secure Shell
Symptom: Request for support for SSH/Telnet in non-default VRF.	
Condition: To configure specific VRF on SSH/TELNET.	

Defect ID: DEFECT000567262	
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: Length of community string for SNMPv1 communities is increased to fit 64 character long community names.	
Condition: Extended the support to enable the customer have long SNMP community names.	

Defect ID: DEFECT000567774	
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: By default SNMP uses relevant VE IP address as a source IP address during trap packet send. We can now configure source-interface VE/Loopback to use as a source interface.	
Condition: VDX SNMP default behavior is to select relevant VE IP address as a source IP address during trap packet send. Now we can select what VE/Loopback interface to use.	

Defect ID: DEFECT000567823	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS7.0.0	**Technology:** VLAN - Virtual LAN
Symptom: Sometime execution of command "show interface description" hangs. This happens when there is improper output to be processed.	
Condition: On hang of execution of "show interface description" command.	

Defect ID: DEFECT000567846	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS5.0.1	**Technology:** Licensing
Symptom: Timeouts in other CLI responses after "show license id" CLI was issued and was non-responsive.	
Condition: Issuing "show license" CLI after the WWN card had a hardware failure causes timeouts in other CLI's	
Recovery: Reload the system, or power cycle the system to recover. If needed run diagnostics to test the WWN card hardware integrity.	

Defect ID: DEFECT000568047	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS7.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: Hafailoverstart/ End traps will be seen as part of cpststatuschange trap during init time. The flow of traps will be hafailoverstart/ warmstart/ hafailoverend traps.	
Condition: There is no trap which indicates that hafailover has happened. In warmstart trap, there are no varbinds.	

Defect ID: DEFECT000568523	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS5.0.1	**Technology:** sFlow
Symptom: In rare cases, a VDX switch may experience an unexpected failover, especially when sFlow is enabled.	
Condition: This could happen after months or years after enabling sFlow.	
Workaround: Disable sFlow or lower the SFlow sample rate.	

Defect ID: DEFECT000569661	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology:** VLAN - Virtual LAN
Symptom: Under certain condition, the internal port may become enabled even though one of the track ports is disabled.	
Condition: On switch reload after enabling the tracking.	
Workaround: Preventive workaround is to remove the track related configuration before reload.	

Defect ID: DEFECT000570270	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Layer 2 Switching
Reported In Release: NOS4.0.1	**Technology:** LAG - Link Aggregation Group
Symptom: Discrepancy in Port-channel programming causes links within the port-channels to be seen as individual link instead of one link.	
Condition: Linecard that contains all of the ISL's connecting from that switch to the fabrics multicast root must be reloaded. When this situation occurs all PO's that are residing on that switch that contain PO members on other linecards other than the linecard that reloaded will experience improper flooding of BUM	
Recovery: Bounce any affected port-channel member interfaces one by one.	

Defect ID: DEFECT000570457	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.0	**Technology:** SNMP - Simple Network Management Protocol
Symptom: SNMP MIB walk to ifName would not return all interface list after SNMP daemon gets restarted.	
Condition: If SNMP daemon gets restarted due to some error then SNMP MIB walk to ifName would not return all interface list.	
Recovery: We can perform HA failover to recover the issue under Extreme TAC guidance.	

Defect ID: DEFECT000570673	
Technical Severity: Medium	**Probability:** Medium

Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: UDLD - Uni-Directional Link Detection
Symptom: UDLD blocks links which are connected using certain breakout ports between different Extreme devices such as between VDX and CER.	
Condition: UDLD blocks the link after detecting mismatch between locally stored vs received port numbers in UDLD PDUs.	

Defect ID: DEFECT000571172	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology:** Static Routing (IPv4)
Symptom: Static route with more superior mask is not installed	
Condition: When static route is configured, if route with less mask exists, then new route is not installed in RTM	

Defect ID: DEFECT000571382	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.1	**Technology:** Configuration Fundamentals
Symptom: Sflow using default-vrf for inband connection does not function correctly when there is no ip address configured for management interface (OOB port).	
Condition: When there is no ip address configured for management interface.	
Workaround: Configuring the ip address on management interface.	

Defect ID: DEFECT000571388	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** IP Multicast
Reported In Release: NOS6.0.2	**Technology:** IPv4 Multicast Routing
Symptom: When certain corner case disruptive events occur, a full cleanup of multicast route table is triggered. During the post recovery of those events, a partial traffic loss can be observed.	
Condition: This issue is observed when a disruptive event such as "chassis disable" or "shutting down RPF interface" is executed on a switch with large number of multicast routes.	
Recovery: Reload the affected switch.	

Defect ID: DEFECT000571714	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.2	**Technology:** LDAP - Lightweight Directory Access Protocol
Symptom: Switch could not import LDAP certificate in FIPS mode.	
Condition: When trying to import LDAP certificate in FIPS mode, the operation fails.	

Defect ID: DEFECT000571889	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security
Reported In Release: NOS6.0.2	**Technology:** SSH - Secure Shell
Symptom: A VDX switch may undergo unexpected reload while configuring "ssh server use-vrf"	
Condition: When VRF name of 32 characters or more is used.	
Workaround: Reduce the characters in VRF name before configuring.	

Defect ID: DEFECT000571906	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Security

Reported In Release: NOS6.0.2	Technology: FIPS - Federal Information Processing Standards
Symptom: Switch is not ready for configuration	
Condition: Crash in Dcmd is observed while FIPS zeroize operation in progress	

Defect ID: DEFECT000571981	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: LAG - Link Aggregation Group
Symptom: Traffic drops for some vLAGs when one of the rbridges went for reload.	
Condition: Reload of one rbridge with vLAG scale configuration.	
Recovery: Port Shut and no shut will recover	

Defect ID: DEFECT000571982	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: IPv6 Addressing
Symptom: 'ipv6 nd suppress-ra' command configuration in interface mode will throw following error. %% Error: VRF Address Family not configured	
Condition: Interfaces which belong to non-principal switches only will not allow this configuration.	

Defect ID: DEFECT000572287	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: Software Installation & Upgrade
Symptom: Firmware download won't be successful on 6740-1G and 6940 platforms.	
Condition: Firmware download functionality is affected under all circumstances for 6740-1G & 6940 platforms.	

Defect ID: DEFECT000572524	
Technical Severity: Medium	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 debug messages show up on the console. The fix takes care of moving debug messages to RAS infrastructure.	
Condition: Some of the FCoE debug messages show up on console. They do not cause any functional issue.	

Defect ID: DEFECT000572990	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.2	Technology: SSH - Secure Shell
Symptom: Request to support for SCP file transfer without using root user.	
Condition: To move configuration files between server and switch without using root user.	

Defect ID: DEFECT000573129	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** VCS
Reported In Release: NOS6.0.2	**Technology:** Logical Chassis
Symptom: Confd prints internal log message on Console.	
Condition: In normal condition customer don't see internal log on console, but in some corner cases confd prints internal log on console.	

Defect ID: DEFECT000573171	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Management
Reported In Release: NOS6.0.2	**Technology:** SNMP - Simple Network Management Protocol
Symptom: On VDX8770 ha failover, IPV6 v1/v2c traps may not be sent	
Condition: On failing over the MM for VDX8770	

Defect ID: DEFECT000573379	
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: Service stopped in particular server/storage where it connected to VDX's particular port.	
Condition: VDX's Back End port link loss	
Recovery: Line card reset	

Defect ID: DEFECT000573950	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS7.0.0	**Technology:** Port Mirroring
Symptom: Monitor session could not be deleted after reload, under certain conditions.	
Condition: The issue occurs when a RSPAN-VLAN is configured as destination and a port-channel, VxLAN or RBridge (system flow based QOS) is configured as a source in a monitor session configuration.	
Workaround: If destination for a monitor session is rspan-vlan, remove that monitor session from port-channel, Vxlan and Rbridge(System Flowbased QOS),if applied on any of these.	

Defect ID: DEFECT000575436	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology Group:** Data Center Fabric
Reported In Release: NOS6.0.2	**Technology:** VCS Fabric
Symptom: No error messages are displayed when the user wrongly configures "fcoeport default" command on an ISL port.	
Condition: The issue is seen when the user configures a certain non-applicable command on an ISL port.	

Defect ID: DEFECT000576481	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.1	**Technology:** Hardware Monitoring
Symptom: When using Mellanox/QSA adapter and Twinax cable on VDX6740 / VDX8770, the link may fail to come up.	
Condition: After reload of the VDX switch with POST diagnostics enabled.	
Workaround: Configure to disable POST on bootup to prevent this issue:	

"sw0(config)# no diag post rbridge-id <> enable "	
Recovery: May need to replace the cables.	

Closed with code changes for NOS 6.0.1a

Defect ID: DEFECT000361772	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS2.1.0	**Technology Area:** Configuration Fundamentals
Symptom: Incomplete command "no banner" deletes the configured "incoming" banner message; "motd" banner and "login" banner are not affected.	
Condition: Run "no banner" command without sub command options such as "login" or "incoming" or "motd"	

Defect ID: DEFECT000445107	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS3.0.1	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: Network interface up/down traps may be received but a corresponding raslog will not exist.	
Condition: Intermittently under normal operating conditions.	

Defect ID: DEFECT000453916	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS4.0.0	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: After disable/enable port based sflow, sampling rate may be set to 0 instead of the actual rate when queried through SNMP SFLOW MIB.	
Condition: On SNMP query of sFlow MIB, sampling rate is not accurate.	

Defect ID: DEFECT000465655	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS4.0.0	**Technology Area:** UDLD - Uni-Directional Link Detection
Symptom: Udld protocol statistics are not getting reset.	
Condition: Udld is enabled on few ports and "clear counters all" command was run with the intention of clearing all statistics.	
Workaround: Use "clear udld statistics" command to reset all udld protocol statistics.	
Recovery: Use "clear udld statistics" command to reset all udld protocol statistics.	

Defect ID: DEFECT000518899	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Traffic Management
Reported In Release: NOS5.0.0	**Technology Area:** Rate Limiting and Shaping
Symptom: "bp-rate-limit heavy slot" config CLI is not supported in logical-chassis mode for fixed form factor switches.	
Condition: Support for ratelimiting CLIs in logical-chassis mode.	

Defect ID: DEFECT000527713	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology Area:** BGP4 - IPv4 Border Gateway Protocol
Symptom: Static route leak is not installing in ribm after removing a dynamic route leak for the same prefix.	
Condition: Configuring route leak - special case of removing a dynamic route leaked for a prefix conflicting with a static	
Recovery: Reapply the configuration.	

Defect ID: DEFECT000532520	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS5.0.1	**Technology Area:** PIM - Protocol-Independent Multicast
Symptom: Upon receiving Assert winner on upstream, downstream PIM router sends (*,G) join to Assert loser and not to the Assert winner. This leads to unnecessary state changes in the LAN.	
Condition: VDX switch is not adhering protocol guideline mentioned in RFC 4601.	

Defect ID: DEFECT000535440	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology Area:** ICMP- Internet Control Message Protocol
Symptom: IPv4 ICMP rate-limiting on Mgmt interface, when set to max value does not work as expected.	
Condition: When IPv4 ICMP rate-limiting on Mgmt interface is set to max value of "4294967295" milliseconds	

Defect ID: DEFECT000537193	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	**Technology Area:** IP Addressing
Symptom: Static route and route leak configuration accepts wrong nexthop format.for IP address	
Condition: Configuration of static route and route leak commands	

Defect ID: DEFECT000540852	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	**Technology Area:** IP Addressing
Symptom: On firmware upgrade, chassis reload, or power-cycle of a linecard, error messages indicating FFDC & Software Verify errors may be seen without any functional impact.	
Condition: While doing a firmware upgrade, chassis reload, or power-cycle of a linecard	

Defect ID: DEFECT000542893	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Configuration Fundamentals
Symptom: While connected to the serial console port, pressing "Ctrl-Shift-6" then "x" causes only the serial console to become unresponsive. You can still telnet to the switch from another window.	
Condition: The unresponsive behavior is seen with keys "Ctrl-Shift-6" then "x".	
Recovery: Either "Ctrl-Shift-6" then "b" or "Ctrl-Shift-6" then "c" will recover the console session.	

Defect ID: DEFECT000544185	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS6.0.0	**Technology Area:** IPv4 Multicast Routing

Symptom: User does a
      "show mcagt group routes slot 0" --
      shows stale multicast (S,G) Cache entries.	

Condition: Do ISSU twice without doing a reload.	

Defect ID: DEFECT000549696	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: There is no configuration command available to enable/disable snmp traps on a per interface basis on VDX platforms.	
Condition: There is no mechanism to enable/disable snmp traps on a per-interface basis.	

Defect ID: DEFECT000549853	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: After removing the last track config using 'no track enable' , the tracked link still remains in the track database.	
Condition: When 'no track enable' is issued and no track related configuration is present, the downlink is still not removed from track summary.	
Workaround: The CLI cmd 'track remove all' inside the downlink interface will remove the interface from tracking.	
Recovery: Issue 'track remove all' on the interface to remove it completely from tracking.	

Defect ID: DEFECT000550826	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Traffic Management
Reported In Release: NOS5.0.1	**Technology Area:** Rate Limiting and Shaping
Symptom: ICMP rate limiting it won't work as expected for VE interface	
Condition: VE interface applied with ICMP rate limiting configuration	

Defect ID: DEFECT000552066	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS5.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: "show mac-address-table port-profile" information is not captured in support save.	
Condition: When support save is collected on a switch	

Defect ID: DEFECT000552067	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	**Technology Area:** BGP4 - IPv4 Border Gateway Protocol
Symptom: Router will install the received default route in the routing table even though it has 'default- information-originate always' command configured.	
Condition: Router receives default routes from its peer and also 'default-information-originate always' command is configured on the Router.	

Defect ID: DEFECT000552570	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: SA MAC address of IP routed frames do not match the egress physical interface MAC	
Condition: IP routing based on L3TTP OpenFlow flow-mods.	

Defect ID: DEFECT000554372	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: Some of the port-security static macs may not get aged out.	
Condition: port-security sticky configuration is removed from a port which has more than 2000 source MAC's learnt.	

Defect ID: DEFECT000554493	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** VCS
Reported In Release: NOS5.0.2	**Technology Area:** Logical Chassis
Symptom: Switch encounters an unexpected reload if an IPv6 ACL rule is removed using a long form of the ACL rule.	
Condition: The issue is encountered when an ACL rule is removed using the long form of the ACL rule.	

Defect ID: DEFECT000555772	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS4.1.3	**Technology Area:** Hardware Monitoring
Symptom: VDX 10gT port (possibly 40gT) may result in the port stuck in offline state.	
Condition: Repeated cable pull/re-insert on VDX 10gT port (possibly 40gT) may result in the port stuck in offline state. Higher probability of hitting this when cabled to HP server with the Intel NIC x520.	
Workaround: Shut/no-shut recovers it.	
Recovery: Shut/no-shut recovers it.	

Defect ID: DEFECT000555882	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** Configuration Fundamentals
Symptom: On the VDX2740/ VDX2746 platforms breakout port QSFP LED stays green on dynamic breakout.	
Condition: After dynamic breakout of QSFP ports on admin down LED may remain green.	
Workaround: Reload system will recover the LED state correctly.	

Defect ID: DEFECT000556079	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: Qlogic and Emulex CNAs connected ports on VDX 6940 as a FIF and fcoe provisioned with remote logical SANs fail to re-login after remapping the VF_Port to AG N_Port ports manually.	
Condition: Get FCoE logins from either Qlogic and Emulex CNA connected to VDX 6940 as FIF. Ports connected to the CNA should be fcoe provisioned with remote logical SAN. Remap the VF_Ports to N_Port on AG manually. Logins are lost and never come back.	
Workaround: do a "no fcoeport" on the CNA connected Ethernet port on FIF before VF port to N_Port remapping and "fcoeport <fabric-map" again on the CNA connected Ethernet port on FIF after VF_Port to N_Port remapping is done.	
Recovery: "no fcoeport" and "fcoeport <fabric-map" on the CNA connected Ethernet port on FIF after VF_Port to N_Port remapping is done should recover the logins.	

Defect ID: DEFECT000556094	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** VCS
Reported In Release: NOS6.0.1	**Technology Area:** Logical Chassis
Symptom: On a cluster with large configuration (more than 100K lines of configuration), HA might go out of sync on secondary nodes during cluster formation.	
Condition: Cluster formation gets triggered on a cluster with large configuration of 100K lines or more.	
Recovery: On secondary nodes where HA is out of sync, executing following command will bring HA back to sync: "ha sync start"	

Defect ID: DEFECT000556654	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS4.1.3	**Technology Area:** Hardware Monitoring
Symptom: VDX 10Gb port LED remains green and online in noscli after only inserting TX strand of the fiber cable into the link partner's RX SFP side.	
Condition: VDX 10Gb port LED remains green and online in noscli under the following condition: By only inserting TX strand of the fiber cable into the link partner's RX SFP side.	
Workaround: Make sure both TX and RX strands of the fiber are clean and connected.	

Defect ID: DEFECT000556655	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS4.1.3	**Technology Area:** IGMP - Internet Group Management Protocol
Symptom: "show ip igmp interface vlan vlanId" can show IGMP querier as "Functionality Disabled" in some instances.	
Condition: In a logical-chassis, VLAN having active ports in multiple nodes and no active ports in at least one node and IGMP snooping enabled on the VLAN.	

Defect ID: DEFECT000557197	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** BGP4 - IPv4 Border Gateway Protocol
Symptom: System allows configuring ECMP static route leak across multiple destination VRFs when it should not.	
Condition: Configuring static route leak across VRFs.	

Defect ID: DEFECT000557683	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS4.1.3	**Technology Area:** Hardware Monitoring
Symptom: Repeated cable pull/re-insert of VDX port 10gT connected to Intel NIC may result in temporary port flapping. After 3 or 4 link flaps, then port will be online.	
Condition: Repeated cable pull/re-insert of VDX port 10gT connected to Intel NIC may result in temporary port flapping. After 3 or 4 link flaps, then port will be online.	
Workaround: Wait several seconds for port to stop flapping, and port will come online.	
Recovery: Wait several seconds for port to stop flapping, and port will come online.	

Defect ID: DEFECT000557718	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: FCoE provisioning for the Interfaces on principal node will not be retained and which will result the FCoE device login failures.	
Condition: While restoring configuration from a file containing FCoE configuration, if the operation is aborted in between by pressing Ctrl-C	
Workaround: Avoid aborting configuration replay triggered due to copy <file> running-config and allow command to execute to completion.	
Recovery: Execute copy <file> running-config again with valid configuration file.	

Defect ID: DEFECT000558106	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** Software Upgrade
Symptom: For a large Logical Chassis cluster, ISL link may be toggled after fabric wide ISSU (with "rbridgeid all" option).	
Condition: Execute "firmware download logical-chassis scp rbridge-id all" command to download firmware for a large Logical Chassis cluster.	
Workaround: Avoid fabric-wide ISSU firmware upgrade in a large cluster.	

Defect ID: DEFECT000558159	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Security
Reported In Release: NOS5.0.0	**Technology Area:** Security Vulnerability
Symptom: Enhancement request to provide a script to selectively restrict SSH traffic to certain VRF's.	
Condition: When need to restrict SSH to only a specific VRF	

Defect ID: DEFECT000558165	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** OSPF - IPv4 Open Shortest Path First
Symptom: OSPFv2 neighborship is not formed.	
Condition: OSPFv2 is configured on the VE interface and change VRF from default-VRF to user defined VRF.	
Workaround: Remove and reconfigure Ve interface and VRF	
Recovery: Remove VRF	

Defect ID: DEFECT000558202	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: FVG Load balancing may not work under following scale conditions	
Condition: 4k FVG sessions are attached to the VCS fabric node, followed by detach FVG sessions in the interface Ve range command mode.	

Defect ID: DEFECT000558266	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: Status of interface is showing as active in "show fcoe interface ethernet" for local logical san on Ag box, after following operations	
Condition: Interface is port-profile-port domain config. In that domain's profile, fcoeport with remote san is removed and fcoeport with local san is added.	
Workaround: Remove port-profile-port domain configuration from the interface. Then change the fcoe provisioning in profile. Then add port-profile-port domain back on the interface.	

Defect ID: DEFECT000558562	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: Port-channel and its member port has contradictory port-profile-port configurations	
Condition: 1. Configure an interface as port-profiled with a san, for example sana 2. Create a port-channel and configure it as port-profiled with fcoe sub profile with a different san, for example sanb 3. Add interface as part of port-channel	
Workaround: Remove port-profile config from interface before adding the interface to port-profiled port-channel Or Remote port-profile config from port channel before adding a port-profiled interface into it.	
Recovery: Remove port-profile-port config from either interface or port-channel	

Defect ID: DEFECT000558668	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: When min-link (minimum link) is configured as more than the number of uplinks present, the downlink is brought down immediately even if all the uplinks are online.	
Condition: When min-link is configured as more than the number of uplinks present, the downlink is brought down immediately even if all the uplinks are online.	
Workaround: configure min-link less than or equal to the number of uplinks configured.	
Recovery: configure min-link as 0.	

Defect ID: DEFECT000558794	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: Physical interface transmit statistics increment for OpenFlow enabled interfaces even when interface is in admin down state.	
Condition: Traffic hits an OpenFlow based data path entry and try to go out on an OpenFlow enabled interface in admin down state.	

Defect ID: DEFECT000558891	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: "fcoe_ha_send_eth_port_config: fcoe_map_name_to_profile failed" message coming up on console	
Condition: 1. Add vCenter 2. Create a VMkernel and initiate CDP trigger on the vSwitch of the host that is connected to the Extreme switch 3. Verify that interface connected to host becomes port-profile-port . 4. Delete the VMkernel	
Workaround: Stop CDP before deleting the VMKkernel	

Defect ID: DEFECT000558915	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** VCS
Reported In Release: NOS5.0.2	**Technology Area:** Logical Chassis
Symptom: Unexpected switch reboot is encountered when an IPv4 ACL is associated with SNMP community.	
Condition: Unexpected switch reboot is encountered when an IPv4 ACL is associated with SNMP community.	

Defect ID: DEFECT000559194	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS6.0.1	**Technology Area:** Port Mirroring
Symptom: The output of the MAPS command "show maps policy name <policy>" is not aligned	
Condition: MAPS needs to be enabled and "show maps policy name <policy>" should be executed.	

Defect ID: DEFECT000559275	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS6.0.1	**Technology Area:** IPv4 Multicast Routing
Symptom: In "show ip pim mcache" output, some multicast mcache entries shows same interface as out going interface as well as incoming interface(RPF),	
Condition: Issue can be observed on scaled multicast configuration on a router where IGMP host and Multicast source are reachable via same interface. Flapping multicast boundary configuration multiple times on all OIF interfaces will produce this issue.	
Workaround: Clearing inconsistent multicast mcache entries (clear ip pim mcache x.x.x.x) will fix the inconsistency.	

Defect ID: DEFECT000559371	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: Interface are shown as active in "show fcoe interface ethernet", when configured with remote san through port-profile-port domain config , though there is no FCF group associated to the fabric map	

Condition: If user tries below steps
    1. User creates a fabric-map, then creates an fcf-group inside it, assign some rb-ids as fcf and fif
    2. Delete the fcf-group created above
    3. Add this fcoe provisioning with this fabric-map into a port-profile
    4. Apply the port-profile on the interface through a port-profile domain
    the interface will be wrongly shown as Active, though fabric-map doesn't have fcf-group in it.	

Defect ID: DEFECT000559390	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** xSTP - Spanning Tree Protocols
Symptom: Changing Spanning-tree mode from MSTP to RPVST may result in showing wrong Spanning-tree status on port channel.	

Condition: 1) Enable MSTP
    2) Disable MSTP
    3) Now enable RVPST
    4) Issue "show spanning-tree interface port-channel <po>"	

Recovery: Shut/no shut port-channel	

Defect ID: DEFECT000559540	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Traffic Management
Reported In Release: NOS6.0.1	**Technology Area:** QoS - Quality of Service
Symptom: Traffic which is expected to be rate-limited at 200mbps is getting rate limited at 1.6Gig	
Condition: The problem happens with port speed 40G or more.	
Workaround: Per port rate limit has a minimum rate of 1.6 Gbps with port speeds of 40G or more. Potential workaround is to rate limit all eight tx queues instead, where the minimal rate is 40 mbps.	

Defect ID: DEFECT000559629	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS6.0.1	**Technology Area:** IGMP - Internet Group Management Protocol
Symptom: L2 multicast traffic is not getting forwarded under certain conditions.	
Condition: When IGMP snooping enabled and restrict unknown multicast configured	

Defect ID: DEFECT000559675	
Technical Severity: Low	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** Configuration Fundamentals
Symptom: Raslog message DAD-1323 displayed after DAD is unexpectedly terminated.	
Condition: DAD expects python script used for configuration to be named as dad.py, it fails otherwise	

Defect ID: DEFECT000559684	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: On VDX6940 device, enabling Conversational MAC Learning (CML) feature may result in flooding VxLAN L2 Extension traffic.	
Condition: 1) Enable VxLAN L2 Extension tunnel involving VDX6940 device 2) Enable CML 3) Send bidirectional traffic across the tunnel 4) Observe traffic flooding even though the MAC addresses are already learnt	
Workaround: Do not enable CML while using VxLAN L2 Extension tunnels	
Recovery: Disable and re-enable CML	

Defect ID: DEFECT000559780	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: Configuration of tracking feature may be rejected after using an interface for tracking and clearing it.	
Condition: After an interface is configured with track config and then remove all the track configuration, the interface is still left in track database. As a result, if this interface is to be configured by another track interface as the upstream, the configuration will be rejected incorrectly.	
Workaround: 'track remove all' can be issued to remove the interface from tracking permanently.	
Recovery: Issue 'track remove all'	

Defect ID: DEFECT000559831	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: Output of the "show openflow resources" is not property aligned, all details are captured.	
Condition: "show openflow resources" output not align to format.	
Workaround: no	

Defect ID: DEFECT000559847	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: Openflow CLI "do show openflow flow" displays all the flows, including flows not configured.	
Condition: OpenFlow configured on the node.	
Recovery: No	

Defect ID: DEFECT000559853	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** NTP - Network Time Protocol
Symptom: NTP with IPv6 key string having a length greater than 15 ASCII characters throws error but configurations except the key is updated in config.	
Condition: Key String of length greater than 15 ASCII characters configured.	
Workaround: Configure a key string of size less than 15 characters.	

Defect ID: DEFECT000559861	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** NTP - Network Time Protocol
Symptom: NTP with IPv6 doesn't honor key string having a length of more than 15 ascii characters	
Condition: Key String of length more than 15 ascii characters configured	
Workaround: Configure a key string of size less than 15 characters	

Defect ID: DEFECT000559865	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: In some instances, error is thrown on configuring/removing the openflow controller even though command is successfully processed.	
Condition: Configuring/removing openflow controller with only name and no other parameters.	
Workaround: If error is thrown, check the running-config by issuing the command "show running-config openflow-controller" to ensure that controller is successfully configured/unconfigured.	

Defect ID: DEFECT000559868	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** LAG - Link Aggregation Group
Symptom: 100MB static port-channel is not coming up	
Condition: 100MB static port-channel may not come up when the speed is changed from 1G to 100MB.	
Recovery: .Delete and re-add port-channel.	

Defect ID: DEFECT000559902	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS6.0.1	**Technology Area:** Port Mirroring
Symptom: When linecard is powered-off , BLADE_STATE==IN rule is triggered instead of OFF rule.	
Condition: When line card is powered off.	

Defect ID: DEFECT000559920	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** Configuration Fundamentals
Symptom: Switch Compact Flash fills up and results in high CPU utilization when running multiple REST requests.	
Condition: Running REST request with wrong password.	
Recovery: Take a backup of "/var/confd/log/ localhost.access" and empty the file contents.	

Defect ID: DEFECT000559962	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS6.0.1	**Technology Area:** IPv4 Multicast Routing
Symptom: PIM will not get enabled on an interface due to which multicast packets will not get forwarded properly out of the interface.	
Condition: If PIM Sparse mode is enabled on an interface before an IP Address is configured on the interface	
Workaround: Disable and enable PIM Sparse mode on the interface.	

Defect ID: DEFECT000559976	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: While unconfiguring passive mode using "no passive" command, command line help indicates that "no-ssl" is required but command is accepted without that.	
Condition: Openflow logical instance is configured with passive no-ssl option	
Workaround: "no passive" will remove the config of all the options associated with "passive" command like "passive no-ssl ip-address xx.xx.xx.xx port 33".	

Defect ID: DEFECT000560025	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Security
Reported In Release: NOS6.0.1	**Technology Area:** User Accounts and Passwords
Symptom: Detected termination of process Dcmd message is observed when Role Based Access rule is created for specific VLAN interface number(for example: rule <rule-number> role <role-name> command interface vlan <VLAN-number>).	
Condition: Role Based Access rule is created with vlan interface number	
Workaround: Create a Role Based Access rule for entire VLAN command (for example: rule <rule-number> role <role-name> command interface vlan)	

Defect ID: DEFECT000560037	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS5.0.2	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: Introducing support 3-tuple interface for ifDescr and ifname object of IF MIB. Rbridge ID is now part of the interface name.	
Condition: ifDescr and ifname object of IF MIB will have the 3-tuples.	

Defect ID: DEFECT000560117	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VXLAN - Virtual Extensible LAN
Symptom: On VDX6940, configuring large number of VEs even in admin down state, may adversely affect BFD sessions over VxLAN tunnels.	
Condition: 1) Create more than 1K VE interfaces, most of them in admin down state 2) Configure VxLAN L2 Extension Tunnels 3) Enable BFD for VxLAN Tunnels	
Workaround: Do not configure large number of VE interfaces that will not be used and consume system resources	
Recovery: Remove VE interfaces in admin down state to free up system resources	

Defect ID: DEFECT000560158	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** OSPF - IPv4 Open Shortest Path First
Symptom: DCM daemon may terminate in a very rare scenario.	
Condition: When user tries to delete the ospf config from two different management sessions at the same time.	
Recovery: Refrain from deleting OSPF configuration from two management sessions at the same time.	

Defect ID: DEFECT000560199	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: VLAN membership lost after ISL port becomes an edge port.	
Condition: Port becoming an edge from ISL.	
Workaround: Execute 'no switchport' and reconfigure VLANs.	

Defect ID: DEFECT000560221	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: CLI session might freeze if openflow related show commands are executed from multiple sessions simultaneously.	
Condition: Executing Openflow related show commands from multiple sessions simultaneously.	
Workaround: Avoid executing show openflow commands from more than one session	
Recovery: Current CLI session need to be aborted and new CLI session need to be started to login to the switch.	

Defect ID: DEFECT000560258	
Technical Severity: Low	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS6.0.1	**Technology Area:** Port Mirroring
Symptom: Output of the "show maps policy detail" command is not aligned	
Condition: Execution of the command "show maps policy detail" will display the output as not aligned.	

Defect ID: DEFECT000560450	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** VRRPv3 - Virtual Router Redundancy Protocol Version 3
Symptom: Node with lowest RBridge id does not become ARP responder.	
Condition: FVG (Fabric Virtual Gateway) sessions are configured and attached to a node for that session. Either no tracking object is present, or if present, it is in the UP state. After issuing "clear ip fabric-virtual-gateway interface ve <vlan id>" for this session, the device will become an ARP responder if it was not the ARP responder earlier, irrespective of its RBridge id.	
Workaround: Track an object in the "Down" state with its track priority as 5 for an FVG session.	

Defect ID: DEFECT000560552	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** Configuration Fundamentals
Symptom: An error message may appear when the Switch is booting up "pem0 Fail to connect to WaveServer ipaddr=127.0.0.1 status=655364"	
Condition: On Switch boot up scenarios.	
Recovery: This is not a functional issue and the Switch automatically recovers from this condition. No recovery is required.	

Defect ID: DEFECT000560644	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** VCS
Reported In Release: NOS6.0.1	**Technology Area:** TRILL - Transparent Interconnection of Lots of Links
Symptom: Unexpected reload when using fcping diagnostic tool.	
Condition: When using the fcping diagnostic tool in an unstable VCS fabric it is possible to see an unexpected reload.	

Defect ID: DEFECT000560711	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: FCoE login may be continuously flapping when a fabric-map is changing from remote logical SAN configuration to local logical SAN configuration.	
Condition: Configuring a fabric-map to be remote logical and later modifying the same fabric-map to be local logical.	

Defect ID: DEFECT000560826	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS6.0.1	**Technology Area:** IPv4 Multicast Routing
Symptom: The multicast traffic may not be sent to receiver.	
Condition: When ISSU or ha failover is done and the route to rendezvous point is not learnt.	
Recovery: shut/no shut the port connected to the rendezvous point (RP).	

Defect ID: DEFECT000560834	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Security
Reported In Release: NOS6.0.0	**Technology Area:** Security Vulnerability
Symptom: A high rate of REST API requests may cause VDX switch to run out of disk space & eventually unexpected reload	
Condition: A sustained high rate of REST API requests	
Recovery: Emptying the contents of /etc/fabos/last_login file	

Defect ID: DEFECT000560844	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Security
Reported In Release: NOS6.0.1	**Technology Area:** ACLs - Access Control Lists
Symptom: Application of 'ip arp inspection filter' fails when the filter name is of longer length (greater than 20 characters)	
Condition: Configuring ip arp inspection filters	

Defect ID: DEFECT000560853	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** IP Multicast
Reported In Release: NOS6.0.1	**Technology Area:** IPv4 Multicast VLAN Traffic Reduction
Symptom: Multicast traffic does not get forwarded when switch is supposed to Layer 2-forward the traffic i.e. either VLAN does not have PIM enabled or switch is not PIM DR on that VLAN	
Condition: When copying running configuration to startup configuration and then reloading the switch, sometimes this issue is seen.	
Recovery: Globally disabling and enabling IGMP Snooping helps to recover from the error state. Disabling and enabling IGMP Snooping on the affected VLAN also helps to recover from the error state.	

Defect ID: DEFECT000560889	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: Passive controller details are shown with command "show openflow controller" even after removing the configuration	
Condition: Provisioned a passive configuration and removing the passive controller.	

Defect ID: DEFECT000560915	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** Configuration Fundamentals
Symptom: Cannot enable the Trunk from Element Manager and the status will be shown as disabled even when the port is enabled.	
Condition: When the fabric trunk configuration is enabled from Element Manager	

Defect ID: DEFECT000560990	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** BGP4 - IPv4 Border Gateway Protocol
Symptom: BGP route may point to leaked route even though nexthop is reachable in the same VRF.	
Condition: BGP route nexthop is pointing to leaked route first. Then BGP nexthop is resolved by same VRF route. Even though nexthop is now resolved in the same VRF, BGP route will be pointing to the leak route nexthop.	

Defect ID: DEFECT000561018	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: On a port channel, trying to add VLANs in comma separated format may throw the following error: %%Error: Platform hardware limitation or resource limit reached.	
Condition: 1) Create port channel 2) Try adding VLANs using comma separated format	
Workaround: Add VLANs using range command instead. Please see example below: Instead of the following command: switchport trunk allowed vlan add 10,100,101,102,103,104,105,106,107,108,109,110 Please use the following command: switchport trunk allowed vlan add 10,100-110	

Defect ID: DEFECT000561135	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VXLAN - Virtual Extensible LAN
Symptom: "show statistics access-list overlay-gateway" does not display any statistics for ingress traffic over tunnel.	
Condition: Execution of CLI command "show statistics access-list overlay-gateway"	

Defect ID: DEFECT000561207	
Technical Severity: Medium	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** FCoE - Fibre Channel over Ethernet
Symptom: AG-FCF interface is configured as port-profile-port where port-profile default is fcoe enabled with default fcoe-map. Later when the AG is made as a part of the remote fcoe fabric map. The port which is port-profiled is displayed as "Active" (should display as "Inactive")	
Condition: Switch is in AG mode 1) port-profile default 2) fcoe-profile, fcoeport default 3) on interface towards CNA, port-profile-port 4) create a remote fcoe-fabric-map "sana" 5) add this rb-id as fcf-id	

Defect ID: DEFECT000561209	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: snmp-server host/v3host running-config will not display the use-vrf option.	
Condition: When switch was upgraded from NOS6.0.0 to NOS6.0.1 with host/v3host configuration.	
Workaround: configure the use-vrf option manually.	

Defect ID: DEFECT000561260	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: Using VLAN range, valid errors printed for failed VLAN creation is also shown for VLANs that were created successfully.	
Condition: Few VLANs are reserved for specific uses, like the FCoE VLAN. During VLAN creation through range, these specific VLANs will fail to be provisioned.	

Defect ID: DEFECT000561283	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Monitoring
Reported In Release: NOS6.0.1	**Technology Area:** Port Mirroring
Symptom: MAPS RASLOG indicating high current is logged for 100G LR optics	
Condition: MAPS needs to be enabled and 100G LR optics need be present in the system	

Defect ID: DEFECT000561308	
Technical Severity: High	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Other
Reported In Release: NOS6.0.1	**Technology Area:** Other
Symptom: Unexpected system reload seen with OpenFlow feature enabled on VDX 8770 platform.	
Condition: Controller installs more than 1K flows or Group buckets having same port to VLAN binding. Firmware upgrade or any change in active slots say slot poweroff, Slot poweron, reload.	
Workaround: Port and VLAN in a flow-mod match or in action or in Group bucket action could be seen as new binding of the port to VLAN. Limit the number of such bindings for a port and a VLAN always within 1K.	

Defect ID: DEFECT000561762	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	**Technology Area:** VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: When both IPv4 and IPv6 VRRP sessions are enabled on same interface, data traffic for one of the protocols will get dropped.	
Condition: On VDX 6940 platforms, When both IPv4 and IPv6 VRRP sessions are enabled on the same interface.	
Workaround: Disable either IPv4 or IPv6 VRRP session based on priority for the required protocol type to work.	

Defect ID: DEFECT000562447	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS5.0.2	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: Unexpected Reload due to inconsistent maintenance of internal data.	
Condition: Issuing the CLI "show running-config rbridge-id <id> snmp-server" when configurations contain local SNMPv3 Host	
Workaround: Global SNMPv3 Host can be used for querying instead of Local SNMPv3 Host.	

Defect ID: DEFECT000562609	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Security
Reported In Release: NOS5.0.2	**Technology Area:** SSH - Secure Shell
Symptom: VDX may reload unexpectedly when executing scripts that continuously login and logout of a switch via SSH using correct password and using an incorrect password,	
Condition: The reload was seen in affected releases only with an intensive scripted login.	

Defect ID: DEFECT000563290	
Technical Severity: High	**Probability:** High
Product: Extreme Network OS	**Technology:** Management
Reported In Release: NOS5.0.2	**Technology Area:** SNMP - Simple Network Management Protocol
Symptom: Snmpwalk on in-band default-vrf and mgmt-vrf with v1, v2, v3 will not work, when both IPv4 and IPv6 Acl are associated to snmp community or user.	
Condition: When switch is configured in in-band management(default-vrf/mgmt-vrf) with SNMP community/v3 user associated with IPv4/IPv6 ACLs.	

Defect ID: DEFECT000565659	
Technical Severity: Medium	**Probability:** Medium
Product: Extreme Network OS	**Technology:** Layer 2 Switching
Reported In Release: NOS6.0.1	**Technology Area:** VLAN - Virtual LAN
Symptom: System reload may occur after sustained conditions where the MAC Consistency Check feature would detect and try to correct an inconsistency.	
Condition: Any condition that can cause long term inconsistencies in MAC learning. For example, L2 MAC learning loops that persist for more than 25 hours.	
Workaround: Disable MAC Consistency Check feature	

Closed with code changes for NOS 6.0.1

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of July 29, 2015 in Network OS 6.0.1.

Defect ID: DEFECT000443595	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS3.0.0	**Technology Area:** IGMP
Symptom: The command "show ip igmp groups detail" may not show updated information (uptime & last reported values) about the learnt groups.	
Condition: This issue is observed only in the show command output & no functionality is impacted.	

Defect ID: DEFECT000443927	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS4.0.0	**Technology Area:** Switch Management
Symptom: The commands which comes under the "rbridge-id" field are not accounted for in the accounting log file in the server. It only logs accounting up to "rbridg-id" and any following executed commands will not be made aware to the user.	
Condition: Under normal operating conditions.	

Defect ID: DEFECT000453568	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS4.0.0	**Technology Area:** Syslog
Symptom: When a FC QSFP with non-ethernet supported speed is inserted in a 40G port, the incompatible SFP RASLOG is not displayed.	
Condition: When a FC QSFP with non-ethernet supported speed is inserted in a 40G port	

Defect ID: DEFECT000466315	
Technical Severity: Low	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS3.0.1	**Technology Area:** IP Route Management
Symptom: The command "no debug all" does not clear ARP debugs	
Condition: This is seen after issuing the "no debug all" command when ARP debug is enabled.	

Defect ID: DEFECT000481607	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS4.1.0	**Technology Area:** TRILL
Symptom: ECMP does not utilize all equal cost links although 10G, 40G, 100G links have equal cost.	
Condition: Mixed bandwidth connections to another RBridge, like 100G and 10G in parallel.	

Defect ID: DEFECT000495669	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.1	**Technology Area:** VLAN
Symptom: Under certain conditions, Switch may continue to push the locally learnt MAC's to NSX-controller even after they are aged out. Additionally, may not be able to receive updates from controller	
Condition: This issue can be observed if switch learns more than 10K MACs behind the VLANs which are also attached to overlay-gateway (via "attach vlan <vid>" command).	

Defect ID: DEFECT000497425	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.1	**Technology Area:** VLAN
Symptom: On a Virtual Fabric enabled switch, VLAN tagged IGMP membership reports will be learnt on port configured as untagged.	
Condition: Switch with virtual-fabric and IGMP snooping enabled; and switchport configured with "switchport mode trunk-no-default-native" and "native-vlan-untagged"	

Defect ID: DEFECT000498001	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.1	**Technology Area:** VLAN
Symptom: With Virtual Fabric enabled mode, if ctag list is greater than 1023, then the configuration is not applied, but it can be seen in "show running config" command.	
Condition: If the user tries to configure ctag list greater than 1023 on Virtual Fabric enabled switch.	

Defect ID: DEFECT000498510	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.1	**Technology Area:** VLAN
Symptom: Incorrect auto port-profile names getting created in switch.	
Condition: Creating a dvSwitch along with dvport-groups while VCenter discovery is in progress, then datacenter-id is not added to the name of the port-profiles.	
Workaround: Create dvSwitch first without dvport-groups option and then create dvport-groups.	
Recovery: After dvSwitch with dvPgs add event, when incorrect port-profiles are created perform the following: 1) Run manual discovery (vnetwork vcenter <vCenter-name> discovery) or wait for next discovery run. 2) Run "vnetwork reconcile vcenter <vCenter-name>" command. This will remove all incorrect additional auto-port-profiles from the switch.	

Defect ID: DEFECT000499125	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS4.0.1	**Technology Area:** OSPF (IPv4)
Symptom: When "auto-cost reference-bandwidth" is modified to 100000 and "ip ospf cost 1" is configured on a VE interface the line is not seen in the running config.	
Condition: Changing OSPF reference bandwidth to a higher value and the default cost on a VE interface has changed to "10" as per "sh ip ospf int". Next, configure "ip ospf cost 1" on Ve and "sh ip os int" correctly reflects it, but nothing is seen in the running config. When removing it with "no ip ospf cost 1" the running config is still the same but the cost is back to 10 now.	

Defect ID: DEFECT000500412	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.1	**Technology Area:** SSH - Secure Shell
Symptom: The SSH daemon must not permit user environment settings.	
Condition: Normal operating conditions.	

Defect ID: DEFECT000500414	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Other
Reported In Release: NOS4.1.1	**Technology Area:** Other
Symptom: The SSH daemon must have capability to limit connections to a single session.	
Condition: For secure environments.	

Defect ID: DEFECT000503440	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** Switch Management
Symptom: Unexpected error while removing the IP ACL on Management interface.	
Condition: After Reload/Ha failover we can hit "Generic error" while removing the IPACL on Management interface.	

Defect ID: DEFECT000506572	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.0	**Technology Area:** OSPFv3 (IPv6)
Symptom: Some IPv6 global addresses like feba::feba/128 which matches fe80::/10 are advertised as link local address in link LSA.	
Condition: When global IPv6 addressing is used in fe80::/10 range.	

Defect ID: DEFECT000511360	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Other
Reported In Release: NOS4.1.2	**Technology Area:** Other
Symptom: TACACS server encrypted password gets defaulted.	
Condition: When the length of the plain-text password is more ~22 character while loading back the backup config from external server.	

Defect ID: DEFECT000512515	
Technical Severity: Low	**Probability:** Low
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS4.1.1	**Technology Area:** IP Route Management
Symptom: Community list that has a dash ("-") does not work when matched in Route Map.	
Condition: Community list having a name that has dash "-"	
Workaround: Community list shouldn't have "-" in the name.	

Defect ID: DEFECT000516882	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.0	**Technology Area:** Switch Management
Symptom: User authentication fails REST request for write operation.	
Condition: This issue happens when user is being authenticated from external server , not local on system.	
Workaround: The workaround for this issue is to create local user and do not use remote authentication.	

Defect ID: DEFECT000517228	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS4.1.2	**Technology Area:** PIM
Symptom: SPT or RPT bit does not get set.	
Condition: This issue was seen in scaled setup after chassis disable/enable on first-hop router.	
Recovery: Perform "clear ip pim mcache" to relearn the entries to set the bit correctly	

Defect ID: DEFECT000518033	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS4.1.2	**Technology Area:** Logical Chassis
Symptom: DCMD panic when 'show running-configuration' command is repeated executed over a period of days, continuously.	
Condition: Repeated monitoring of the output of the command 'show running-configuration' can cause this issue.	

Defect ID: DEFECT000518129	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.0	**Technology Area:** Fabric Build
Symptom: Software VERIFY() error can be observed on the console in relation to ISL interfaces, following slot-power-off and HA failover operations.	
Condition: There is a timing case where the ISL information replicated on the standby MM is not fully in sync with the active MM when a linecard having ISL connections is powered off, and then followed by HA failover.	
Recovery: NOS will recover automatically.	

Defect ID: DEFECT000520539	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.0	**Technology Area:** Static Routing (IPv4)
Symptom: The range command "int Ve 2-4080" times out.	
Condition: With virtual-fabric enabled, attempt to configure large numbers of SVI interfaces using the range command.	
Workaround: Create VEs in smaller chunks of 500 at a time.	

Defect ID: DEFECT000520643	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.0.1	**Technology Area:** VLAN
Symptom: Code upgrade may fail to complete on VDX switches when invalid DNS is configured in the switches.	
Condition: When DNS is directly configured or is supplied via DHCP.	
Recovery: Configure a bogus dns server in rbridge mode and delete it. This will remove any stale dns config in the system acquired via DHCP / NOS CLI. conf t # rbridge 101 sw0(config-rbridge-id-101)# ip dns domain-name test123 Warning : Name Server also needs to be configured for successful operation sw0(config-rbridge-id-101)# no ip dns domain-name test123	

Defect ID: DEFECT000521050	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.2	**Technology Area:** IEEE 802.1s MSTP
Symptom: Inconsistent MSTP VLAN information between the running config and "show spanning-tree mst-config" command after VLAN configuration for MSTP instance.	
Condition: Issue will be hit when the length of VLAN configuration for MSTP instance exceeded the limit of 500 characters.	

Defect ID: DEFECT000523618	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.0	**Technology Area:** Component
Symptom: 10GbT links may not come online on chassis reload	
Condition: This issue can occur due to timing issues when the speed auto negotiation is stuck.	
Recovery: Toggle the port via a shut/no-shut	

Defect ID: DEFECT000523851	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: Request to support SNMP with additional MIB for eSR4 optics, including channel monitoring support for QSFP port in 40G mode.	
Condition: Monitoring the optics health using SNMP	

Defect ID: DEFECT000524743	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.0	**Technology Area:** VLAN
Symptom: TVLAN Traffic will be forwarded over dot1q edge ports in a misconfigured VCS.	
Condition: The following misconfiguration is needed: 1) When TVLAN is configured and extended over tunnels with certain ctags in one VCS 2) In other VCS where the tunnel terminates, no TVLAN is configured, but dot1q VLANs are configured which are same as ctags for TVLAN in other VCS.	
Workaround: Fix the configuration to match TVLANs.	

Defect ID: DEFECT000524852	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.0	**Technology Area:** Logical Chassis
Symptom: The switches go back to the default configuration when powered on. Running configuration does not persist.	
Condition: When the Logical chassis cluster has only VDX2740 switches and all the switches are powered down at the same time.	
Workaround: User needs to run below command on principal switch before the entire cluster goes down or after every configuration change in cluster. # "vcs auto-upgrade-config"	

Defect ID: DEFECT000525323	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** Switch Management
Symptom: When a user adds description to a destination mirror port for the VDX siwtch, an error message is displayed indicating the command is not allowed. However, in the output of "show interface status", the description is shown.	
Condition: When the interface is a destination mirror/session port and user add description for same.	

Defect ID: DEFECT000526450	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** NetCONF
Symptom: In Certain conditions, L2traceroute RPC doesn't give proper data	
Condition: When user query L2TraceRoute through NetConf RPC	

Defect ID: DEFECT000526682	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: SNMP query is displaying temperature, RX & TX power and current values for the offline/admin down interface with low warning.	
Condition: SNMP query	

Defect ID: DEFECT000526683	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: SNMP query for breakout ports displays each lane's information(Rx & TX power, temperature and current values) 4 times instead of one time.	
Condition: Snmp query for breakout ports.	

Defect ID: DEFECT000526684	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.2	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: SNMP query for temperature,RX power,Bias current displays same values for all the four lanes in Non-breakout mode.	
Condition: SNMP query	

Defect ID: DEFECT000526918	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.3	**Technology Area:** Optics
Symptom: Unexpected output for the CLI show interface status by displaying "SFP Absent" instead of displaying "connnected/Not connected".	
Condition: This can hit when the copper ports for VDX 6710/VDX 6740T are used.	

Defect ID: DEFECT000527525	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS5.0.0	**Technology Area:** PIM
Symptom: PIM interface does not support neighbor filtering capability. As a result, black listed neighbors cannot be filtered.	
Condition: PIM neighbors cannot be filtered out. leading to some functional problems in deployments where administrator wants to blacklist rogue PIM neighbors.	

Defect ID: DEFECT000529011	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS4.1.3	**Technology Area:** VRRP & VRRP-E (IPv4)
Symptom: Enhancement request for support of Unicast ARP request using a CLI knob.	
Condition: When the partner device only support unicast ARP request & not the usual broadcast ARP request.	

Defect ID: DEFECT000529802	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Switch Management
Symptom: Management connectivity to VCS virtual IP address is getting lost. Management Module or Chassis IP connectivity has no issues.	
Condition: Issue is seen only with Single node VCS cluster having ISSU fail-over for VCS virtual ip address	
Workaround: Cold boot will recover the VCS virtual IP address connectivity	

Defect ID: DEFECT000529874	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Security
Reported In Release: NOS4.1.2	**Technology Area:** Receive ACLs
Symptom: After reload, SSH access to the switch is allowed even when SSH service is configured to be disabled ("ssh server shutdown" command).	
Condition: The issue is observed after a reload, if the saved configuration includes an extended ACL policy (enforced on Ethernet management port) containing an ACL rule with sequence number greater than 100.	
Workaround: Configure ACL rules with sequence number less than 100.	
Recovery: Issue command to start and shutdown the server : "no ssh server shutdown" followed by "ssh server shutdown" command.	

Defect ID: DEFECT000530011	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Firmware Upgrade/ Downgrade
Symptom: Configuration may be lost during downgrade from 5.0.1 to 5.0.0 on TOR switches (VDX 6740, VDX 6740T, VDX 2740, etc)	
Condition: On TOR switches running with 5.0.1 when GOS1 is the Active GOS.	
Workaround: Reboot the switch before downgrading from 5.0.1 to 5.0.0 with the coldboot option.	

Defect ID: DEFECT000530495	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.2	**Technology Area:** Component
Symptom: Performance degradation for FibreChannel traffic between FibreChannel ports of the VDX 6730 and an FCR switch.	
Condition: Degradation of FibreChannel link throughput occurs due to unrecoverable transmit credit loss. Unrecoverable credit loss occurs when there are faulty or marginal optics in the system, and either primitives on the link or corrupted and are not interpreted correctly.	

Defect ID: DEFECT000530557	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: When VDX 8770 generates a WLV-1003 raslog, Line card number is not part of the string. BNA consumes only the string and is unable to discern the line card number.	
Condition: When 8770 port goes faulty, a WLV-1003 message is generated.	

Defect ID: DEFECT000530570	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Firmware Upgrade/ Downgrade
Symptom: VRF configuration doesn't apply properly on management interface	
Condition: Upgrade from 4.x to higher releases can cause this issue	
Recovery: Re-apply the VRF config on management interface	

Defect ID: DEFECT000530775	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS4.1.2	**Technology Area:** Fabric Build
Symptom: The functionality of option "rbridge-id" in CLI "vcs" varies based on its usage in the CLI. If rbridge-id comes between "vcs" and "logical-chassis", it is for local node to change its rbridge-id. It will take only one rbridge-id for example, "sw0# no vcs rbridge-id 104 logical-chassis enable" will change local node RBridge to 104 and change LC to FC mode And if RBridge-id comes after logical-chassis, it is meant for that RBridge and will change the cluster mode for that RBridge. for example "sw0# no vcs logical-chassis enable rbridge-id 104" will change the VCS mode of RBridge 104 from FC to LC As this is creating confusion, new CLI option "set-rbridge-id" is introduced in CLI "vcs" which will be used to change the existing RBridge-id of the local node.	
Condition: Not applicable as it is done for usability enhancement.	

Defect ID: DEFECT000531087

Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** Conversational MAC Learning

Symptom: Dynamically learnt MAC Addresses are not shown in the output of CLI "show mac-address-table interface <interface-type> <interface-name>".

Condition: This issue is seen when CLI "show mac-address-table" is executed for a specified interface provided as input.

Workaround: Execute the CLI "show mac-address-table". This will display the information about all the MAC addresses including dynamically learnt MAC Addresses.

Defect ID: DEFECT000531793

Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS5.0.1	**Technology Area:** PIM

Symptom: VDX switch does not trigger PIM assert after receiving data packet on out going interface of (*,G) which results in duplicate traffic on the LAN.

Condition: Absence of PIM assert leads to duplicate data in the LAN as the assert loser does not prune off it's outgoing interface (OIF)

Workaround: Fix RPF path in the neighbor router so that every PIM neighbor on the LAN elects same RPF upstream.

Defect ID: DEFECT000531974

Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Switch Management

Symptom: When REST is used, an error response "401 Resource not found" is returned during the deletion of the port-profile domain" resource.

Condition: This issue is encountered after configuring the port-profile domain resource for the given interface, and while deleting the same using DELETE method

---

Defect ID: DEFECT000532218

Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.1	**Technology Area:** OSPF (IPv4)

Symptom: When doing the SNMP Set operation on the MIB object ospfStubRouterAdvertisement by setting it to the value "doNotAdvertise", the Set operation fails though no error is returned.

Condition: This behavior is seen only when doing SNMP Set operation on the ospfStubRouterAdvertisement MIB object to the value "doNotAdvertise".

Defect ID: DEFECT000532476

Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** VCS
Reported In Release: NOS6.0.0	**Technology Area:** Logical Chassis

Symptom: Unexpected reload due to Dcmd process on conversion from local only to distributed mode.

Condition: When principal node goes out of cluster during the end of cluster formation.

Defect ID: DEFECT000532520	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS5.0.1	**Technology Area:** PIM
Symptom: Upon receiving Assert winner on upstream, downstream PIM router sends (*,G) join to Assert loser and not to the Assert winner. This leads to unnecessary state changes in the LAN.	
Condition: VDX switch is not adhering protocol guideline mentioned in RFC 4601.	

Defect ID: DEFECT000532620	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.3	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: SNMP returns incorrect speed value for FAN FRU in faulty state.	
Condition: Issue was seen during SNMP query on swSensorTable with faulty FAN FRU.	

Defect ID: DEFECT000533226	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.2	**Technology Area:** MAC ACLs
Symptom: On scaled setups, continuing to configure additional MAC ACLs may result in exhausting TCAM resources. User notification now available using RASLOG.	
Condition: Solution available in NOS6.0.1 Release for this RFE.	

Defect ID: DEFECT000533273	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** FCoE
Symptom: Devices logged into same FCoE Logical SAN but through a different FCF Group will toggle if the FCF of the other FCF group is rebooted.	
Condition: When there are multiple FCFs(Access Gateway) configured for the same remote FCoE Logical SAN(fabric-map) and there are devices logged in to the remote FCoE Logical SAN through different FCFs and one of the FCF goes for a reboot, the devices logged in through another FCF will toggle.	
Workaround: Use one FCF-Group per Fabric-map	
Recovery: Devices would automatically recover after the toggle.	

Defect ID: DEFECT000533424	
Technical Severity: Low	**Probability:** Low
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS5.0.1	**Technology Area:** Copy Support
Symptom: support save does not capture command output for processes CPU usage, details about VCS and history about firmware download such as: show process cpu top show vcs detail show firmwaredownloadhistory	
Condition: Running copy support	

Defect ID: DEFECT000533508	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: On running a snmpwalk periodically, memory leak is observed in the snmpd process and eventually the switch will experience an unexpected reload.	
Condition: The memory leak and unexpected reload are seen only when the TCP and UDP MIBs are being queried.	

Defect ID: DEFECT000533602	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.3	**Technology Area:** Component
Symptom: Power cycling the VDX 6740 switch resulted in FLASH Monitor reporting Marginal status after 70 reboots.	
Condition: Continuous Power cycling	
Recovery: After rolling reboot detection, cleanup cfg and reboot switch.	

Defect ID: DEFECT000533986	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** VCS
Reported In Release: NOS4.1.2	**Technology Area:** Logical Chassis
Symptom: In the output for 'show rbridge-running-config rbridge-id <>' where rbridge-id refers to non-principal switch, the spanning tree configuration under a port-channel will not be displayed.	
Condition: This issue occurs only when rbridge-id refers to a non-principal switch.	
Workaround: To view configuration on port-channel the following command can be used "show running-config interface Port-channel".	

Defect ID: DEFECT000534883	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** Switch Management
Symptom: The user is allowed to configure an in-band Virtual IP address for a VE interface that is not in a subnet of any of the primary IP addresses configured for the VE interface.	
Condition: Configuring an inband Virtual IP address for a VE interface that is not in a subnet of any of the primary IP addresses configured for the VE interface.	
Recovery: The out-of subnet Virtual IP address can be deleted or changed to an IP address that is in the subnet of one of the primary IP addresses configured for the VE interface.	

Defect ID: DEFECT000535163	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS4.1.3	**Technology Area:** Other IPv4
Symptom: User currently not allowed to manually configure multiple DHCP Relay IP addresses and select the desired Gateway IP address.	
Condition: When clients are to be assigned IP addresses from a specific range and the user wants a Gateway IP address other than the lowest IP address configured on the DHCP relay interface.	

Defect ID: DEFECT000535306	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** VLAN
Symptom: CLI shows Unicast packets received with CRC errors separately in two counters, whereas SNMP MIB query will show both in single counter.	
Condition: 1) Some Unicast packets received with CRC error 2) CLI show command and SNMP MIB query are done, which provide different counter values	
Workaround: CLI show command is accurate and accounts for CRC error packets separately.	

Defect ID: DEFECT000535703	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** Link Aggregation
Symptom: Port-Channel (with LACP short timeout) can have interface flap	
Condition: User tries to clear a large number of MACs (> 70K)	

Defect ID: DEFECT000535705	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.2	**Technology Area:** Optics
Symptom: Link between VDX and Linksys is unable to come online.	
Condition: Connecting VDX link to Cisco SF 102-24 unmanaged switch will end up in link down state.	

Defect ID: DEFECT000536150	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** OSPF (IPv4)
Symptom: Virtual Link for OSPFv2 is not forming.	
Condition: After failover to standby OSPFv2 virtual link with MD5 configuration might not become full.	

Defect ID: DEFECT000536599	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS4.1.3	**Technology Area:** IP Route Management
Symptom: Certain incoming TCP/UDP traffic will be dropped when transiting VCS fabric	
Condition: When incoming traffic with DSCP value is set as 63 & traversing VDX 6740	
Workaround: Use DSCP value other than 63.	

Defect ID: DEFECT000536995	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS4.1.3	**Technology Area:** Syslog
Symptom: When ACLs are configured to restrict access to a switch from outside, then following raslog messages may be seen repeatedly on the console: " \<snip>unauthorized host with ip address \<> tries to establish connection using UDP breach at port \<>"	
Condition: This occurs when the ACLs that are configured to restrict access to a switch from the outside are violated.	
Workaround: Configure like the example below for any \<message>. " logging raslog message SEC-3039 suppress". Even so, sometimes these messages could be seen for the Standby OS.	

Defect ID: DEFECT000537532	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS5.0.1	**Technology Area:** PIM
Symptom: The SPT bit in the multicast entry is not set after the shortest path switchover completes.	
Condition: Upon reloading the system or performing chassis disable/enable	

Defect ID: DEFECT000537636	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS5.0.1	**Technology Area:** PIM
Symptom: The data encapsulation will not happen for some of the source traffic.	
Condition: Only seen with 8-11 static RP with prefix-list.	

Defect ID: DEFECT000537925	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Security
Reported In Release: NOS6.0.0	**Technology Area:** Zoning
Symptom: When performing a copy-file-running operation where the downloaded file contains a wildcard '*' character appended to the enabled-configuration cfg-name field, the zoning configuration does not get activated.	
Condition: A user will see this issue occur if they perform a copy-running-to-file operation while having an enabled-configuration when there is also an open zone transaction or if the defined and enabled zone configurations are mismatched. Either of these two scenarios will result in a wildcard character appended to the enabled-configuration cfg-name field in the saved file. If the user later tries to perform a copy-file-to-running operation with this uploaded file, they will experience the described customer symptom.	
Workaround: To avoid this condition, users need to ensure that there is no wildcard character appended to the enabled-configuration cfg-name field before performing a copy-running-to-file operation.	
Recovery: To recover, the user will need to manually enable the zone configuration after the copy-file-to-running operation completes.	

Defect ID: DEFECT000538582	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.0	**Technology Area:** CLI
Symptom: Non existing Breakout Interface Te X:1 is visible in show ip interface brief command	
Condition: User may see non existing breakout port TenGigi X:1 after breakout and un-breakout of FortyGig interface.	
Workaround: User can avoid using 36th Forty gig interface for breakout and un-breakout	
Recovery: system reload will solve the issue	

Defect ID: DEFECT000538655	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** IP Route Management
Symptom: Rib manager process may terminate under stress conditions	
Condition: Repeated slotpower off/on operation on a 10G linecard.	

Defect ID: DEFECT000538660	
Technical Severity: Low	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** SFLOW
Symptom: Currently, user is not able to specify the desired source IP address to be used for sFlow traffic.	
Condition: When the physical Mgmt IP address is not desired for sFlow traffic.	

Defect ID: DEFECT000539017	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** VLAN
Symptom: l2traceroute will fail if the remote MAC address is learnt on the same switch.	
Condition: l2traceroute command is issued with remote MAC addresses learnt locally.	
Workaround: Do not use l2traceroute for remote MAC addresses which are learnt locally on the same switch.	

Defect ID: DEFECT000539172	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.0	**Technology Area:** CLI
Symptom: LLDP advertise configurations under LLDP Profile are not shown in the running configuration.	
Condition: LLDP advertised configurations under LLDP Profile are not shown in running configuration. When running config is saved to remote host, these configuration will be skipped. LLDP advertised configuration under LLDP profiles will be lost on downloading the configuration back from remote host.	

Defect ID: DEFECT000539337	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.2	**Technology Area:** Component
Symptom: Clock drift seen on VDX 6740, VDX 6740-T, and VDX 6940-36Q of up to 3 seconds.	
Condition: Using the set clock, local clock commands without NTP.	
Workaround: Use NTP.	
Recovery: Periodically synch with NTP.	

Defect ID: DEFECT000539560	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: After upgrading from 4.1.x to 5.0.x, when the command "snmp-server enable trap" is issued, it may not be shown in the running-config. Thus "no" form of this command also does not work.	
Condition: Upgrade from 4.1.x to 5.0.x.	
Workaround: Default the config & reload to see the configuration "snmp-server enable trap" in the default-config.	

Defect ID: DEFECT000540199	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** Overlay Gateway
Symptom: If the same MAC addresses are detected over multiple VTEP tunnels, then the switch may learn those MAC addresses as static & dynamic entries.	
Condition: When the same MAC address flaps between multiple VTEP tunnels continuously.	

Defect ID: DEFECT000540363	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** CLI
Symptom: Beacon enable on an interface can cause switch to reload.	
Condition: Issuing of 'beacon enable interface' command.	

Defect ID: DEFECT000540462	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Security
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Authentication
Symptom: A non-default admin password is not synced with a newly joined secondary node upon bootup in a logical-chassis cluster.	
Condition: This occurs only when the admin password is changed to non-default password in a logical-chassis cluster and a secondary node joins the fabric with default-config.	
Recovery: Use the default admin password for the affected node, or reload the affected node once.	

Defect ID: DEFECT000540495	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.3	**Technology Area:** Switch Management
Symptom: Sometimes when SFP is plugged in and cable is not connected, show interface status is shown as connected.	
Condition: This issue is seen with SFP plugged in and cable not connected . In the case of the SFP not being plugged in, sfpabsent is displayed correctly.	
Workaround: Output can be verified with other cmd 'show ip interface <>	

Defect ID: DEFECT000540563	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Firmware Upgrade/ Downgrade
Symptom: Inband Telnet & SSH connectivity is lost when upgrade from 4.x to 5.x image.	
Condition: When telnet / SSH are being utilized inband	
Workaround: Access the switch via Out-of-band management ports.	
Recovery: Access the switch via console & configure the Layer 3 interfaces in management VRF for inband access.	

Defect ID: DEFECT000540752	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Traffic Management
Reported In Release: NOS5.0.1	**Technology Area:** Traffic Policies
Symptom: Can not remove the speed 1000 from the breakout port configuration.	
Condition: Manually configure the speed to 1000.	
Recovery: Remove the breakout config for the affected port and re-configure the breakout config again.	

Defect ID: DEFECT000540851	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Traffic Management
Reported In Release: NOS5.0.1	**Technology Area:** Traffic Policies
Symptom: Ping recovery takes 2+ Seconds when a node in the fabric is rebooted.	
Condition: Happens in scaled VLAN environments.	
Recovery: Ping recovers automatically in 2+ seconds	

Defect ID: DEFECT000541178	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** BGP4 (IPv4)
Symptom: BGP IPv6 route might be missing when route is filtered using route-map tag value.	
Condition: Static IPv6 route-map filter based on tag is not filtering route correctly.	

Defect ID: DEFECT000541426	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** FCoE
Symptom: There is a mismatch between "show fcoe login" and "show fcoe devices" output	
Condition: There are FCOE devices connected to a Slot of FIF (VDX-8770) and logged into a FCF (e.g. a VDX6740 with Access Gateway enabled). When the Slot of FIF goes faulty, FCOE device logins are not removed from FCF. So the FCOE device info in "show fcoe login" and "show fcoe devices" does not match.	
Workaround: Toggle the FCOE interface on the FCF where the device is logged into.	

Defect ID: DEFECT000541622	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** FCoE
Symptom: FCoE devices logged into remote logical SAN are not cleared on the FCF after a slot power off occurred in FIF where FCoE hosts are connected.	
Condition: FCoE devices are connected to FIF (VDX-8770) and logged into a FCF (e.g. a VDX 6740 with Access Gateway enabled). Then perform a slot power off on FIF, all FCoE devices are connected to the powered off slot are supposed to be cleaned up in FCF but they are not.	
Recovery: Toggle the FCoE interface on the FCF where the device is logged into.	

Defect ID: DEFECT000541649	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS4.1.2	**Technology Area:** IPv4 Multicast Switching
Symptom: Multicast packets with source IP as 0.0.0.0 is not forwarded by VDX. If any application is sending multicast packets with source IP as 0.0.0.0, then this packet drop may cause application failure.	
Condition: Any multicast packet carrying source IP 0.0.0.0 will face this issue.	

Defect ID: DEFECT000542100	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS6.0.0	**Technology Area:** Logical Chassis
Symptom: One of the node in the cluster might fail to join the cluster and goes for an unexpected reboot due to an exception.	
Condition: Short duration of around 30 seconds on secondary node when it is in the process of rejoining the cluster. In this short duration, if Fabric daemon queries for information in database, it might encounter an exception.	
Recovery: Secondary node rejoins the cluster after a reload.	

Defect ID: DEFECT000542160	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.3	**Technology Area:** Optics
Symptom: During an extended power cycle test overnight on a 6740, one of the 40-gig port would stay admin down.	
Condition: The 6740 is connected with external loop back cables and is power-cycled for extended periods.	
Workaround: If admin down is seen, cable + SFP pull [followed by push] would bring up the link after issuing a 'no shut' on the down port.	

Defect ID: DEFECT000543001	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Firmware Upgrade/ Downgrade
Symptom: Downgrade of a firmware on a 6740 from NOS 6.0.0d to NOS 5.0.1b might leave out these two versions one in SW/0 and other in SW/1. Firmware recovery may not work.	
Condition: Downgrade of a firmware on a 6740 from NOS 6.0.0d to NOS 5.0.1b might leave out these two versions, one in SW/0 and other in SW/1.	

Defect ID: DEFECT000543066	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** LLDP & LLDP-Med
Symptom: The LLDP "show lldp neighbor detail" command output continues to display the System Capabilities of the other end long after the other end has stopped including the System Capabilities information in the LLDP PDU that it sends to us.	
Condition: The LLDP daemon at the other end had included the System Capabilities information in the LLDP PDU that it sent to us and then it stopped including the System Capabilities information in the LLDP PDU.	
Recovery: The "clear lldp neighbor" command can be executed to clear the cached neighbor info. The System Capabilities of the other end will no longer be displayed in the "show lldp neighbor detail" command output.	

Defect ID: DEFECT000543082	
Technical Severity: Low	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** DHCP (IPv4)
Symptom: Even though DHCP Relay and Gateway addresses are not configured, the show command output will display the gateway address being used on that interface.	
Condition: Entering the "show ip dhcp relay gateway interface" command before configuring relay address.	

Defect ID: DEFECT000543234	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.0	**Technology Area:** Component
Symptom: ISL port may take upwards of 5 minutes to come online	
Condition: Repeated reload of the system.	
Recovery: Wait for ~6 minutes or shut/no shut both interfaces	

Defect ID: DEFECT000543239	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** AAA
Symptom: Certain commands like "router <xyz>" are not available to be configured for RBAC access rule as they are considered as admin-only commands.	
Condition: When command "router <>" needs to be assigned to a specific user using an access-rule.	
Workaround: Configure the user with the admin role to offer the permissions.	

Defect ID: DEFECT000543301	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** Switch Management
Symptom: Intermittently, some interfaces may go down	
Condition: After repeated chassis disable and chassis enable.	
Recovery: Shut/no shut on the down interface.	

Defect ID: DEFECT000543496	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** VLAN
Symptom: l2traceroute command will not function properly for few reserved GVLANs	
Condition: 1) GVLAN Enabled 2) l2traceroute command used for VLANs 8186 through 8191	

Defect ID: DEFECT000543532	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** Switch Management
Symptom: The range support for breakout index in fibre-channel interface is not present.	
Condition: The below command cannot be executed for FC breakout ports configurations for range of breakout ports. sw0(config)# interface FibreChannel 10/0/50:1-4	
Workaround: The FC breakout ports can be configured individually. sw0(config)# interface FibreChannel 10/0/50:1	

Defect ID: DEFECT000543534	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** IEEE 802.1s MSTP
Symptom: There will be a traffic loss of 20s when VCS is changed from non-root bridge to root bridge with 32 MSTP instance and 1000 VLANs	
Condition: When MSTP bridge is changed from non-root bridge to root bridge.	

Defect ID: DEFECT000543635	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** VLAN
Symptom: Configuring VFAB VLAN with name will throw error, if the VLAN is not provisioned already	
Condition: 1) GVLAN is configured, but not provisioned yet 2) Assign name using CLI command "name <vlan name>"	
Workaround: Assign the VLAN name after the VLAN is provisioned. Warning message will still be shown, which can be ignored.	

Defect ID: DEFECT000543721	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** VLAN
Symptom: source MAC address state in show mac-address-table output is shown as inactive instead of Remote.	
Condition: source MAC learning on port-channel with no local members.	
Recovery: Issue clear mac-address-table dynamic command	

Defect ID: DEFECT000543870	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** VCS
Reported In Release: NOS4.1.3	**Technology Area:** Fabric Build
Symptom: In an extended power cycle test, 40-gig port on a 6740 may become administratively down.	
Condition: In a single 6740 connected with external loop back links, an extended power-cycle test was run.	
Workaround: Execute 'no shut' and pull out and push-in the QSFP from the offending port to bring up the port.	

Defect ID: DEFECT000543914	
Technical Severity: Medium	Probability: High
Product: NOS	Technology: Monitoring/RAS
Reported In Release: NOS5.0.1	Technology Area: Syslog
Symptom: When "logging raslog message <msgId> suppress" is configured to suppress a specific message Id, raslog messages from Standby MM sent to Active MM or messages from Standby GOS sent to Active GOS can be seen on (active) console.	
Condition: This behavior will be seen when raslog suppression is configured.	
Workaround: There is no functional impact. The unnecessary raslog messages can be ignored.	

Defect ID: DEFECT000543915	
Technical Severity: High	Probability: Medium
Product: NOS	Technology: Layer 2
Reported In Release: NOS6.0.0	Technology Area: FCoE
Symptom: FCoE logins are not coming up after changing FCoE VLAN	
Condition: For a FCoE fabric-map, both priority and VLAN can be changed from default/auto-configured values. When priority is changed before the VLAN is changed, this issue can happen.	
Workaround: First change FCoE VLAN and then change FCoE priority in a given FCoE fabric-map	
Recovery: Change FCoE priority to default and then back to desired priority to recover.	

Defect ID: DEFECT000543958	
Technical Severity: High	Probability: High
Product: NOS	Technology: Layer 3
Reported In Release: NOS6.0.0	Technology Area: VRRP & VRRP-E (IPv4)
Symptom: VE configuration through File Replay resets the VE admin status to default status("shutdown) if admin status is not specifically mentioned in the file being replayed.	
Condition: File replay of VE interface doesn't have admin status	
Workaround: Always have desired admin state (shutdown/no shutdown) with configuration of VE interface in the file which is going to be replayed. This always remains there if file is created by using valid configuration from a switch. Do not alter VE configuration manually in the file by deleting it's admin state.	
Recovery: In case this issue occurs please configure admin state as "no shutdown".	

Defect ID: DEFECT000544031	
Technical Severity: Medium	Probability: Medium
Product: NOS	Technology: System
Reported In Release: NOS5.0.0	Technology Area: Optics
Symptom: Port LED exhibits inconsistent behavior after ISSU on VDX 6740T	
Condition: ISSU	

Defect ID: DEFECT000544034	
Technical Severity: High	Probability: Medium
Product: NOS	Technology: IP Multicast
Reported In Release: NOS5.0.1	Technology Area: IPv4 Multicast Routing
Symptom: VDX last hop router may unexpectedly send Prune messages, interrupting the traffic to the receivers.	
Condition: In topologies where Multicast source is connected to the same VLAN as receivers & RP VE interface.	

Defect ID: DEFECT000544081	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Build
Symptom: After loading the config from flash to running-config , we could see some IPv6 config loss	
Condition: When loading IPv6/ICMPv6 configuration from flash to running-config.	

Defect ID: DEFECT000544085	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Build
Symptom: After loading the config from flash to running-config , some ipv6 icmpv6 config may be lost.	
Condition: When loading IPv6/ICMPv6 configuration from flash to running-config.	

Defect ID: DEFECT000544255	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.3	**Technology Area:** Component
Symptom: Some 40G Interfaces may fail to link up.	
Condition: This condition can be observed after multiple consecutive power cycles.	
Recovery: shut/no shut interface on both local and remote sides	

Defect ID: DEFECT000544256	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** VCS
Reported In Release: NOS4.1.3	**Technology Area:** Fabric Build
Symptom: Some Forty Gig Interfaces may take time to come online	
Condition: If a device is undergoing Multiple power cycles this condition may appear	
Recovery: shut/no shut interface on local and remote sides	

Defect ID: DEFECT000544403	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Traffic Management
Reported In Release: NOS5.0.1	**Technology Area:** Routing
Symptom: When a NOS switch connecting to FOS backbone is rebooted, hosts experience errors in LUN discovery. Hosts report error message such as LUNs failed.	
Condition: The issue is seen on a 24-node NOS cluster connecting to a FOS backbone.	

Defect ID: DEFECT000544647	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS5.0.1	**Technology Area:** Copy Support
Symptom: Support-save core files generated due to support-save on a VDX 6740.	
Condition: Taking support-save	

Defect ID: DEFECT000544875	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** VCS
Reported In Release: NOS6.0.0	**Technology Area:** AMPP
Symptom: Interface connected to a vmnic of ESXi host will not revert back to port-profile mode after adding and removing swithport config on the interface.	
Condition: This can occur when the interface which has been put in port-profile mode as part of vCenter orchestration is put into switch-port configuration and later the switch-port configuration is removed from it. This can happen by manually changing the configurations on the interface(rare scenario) in automated vCenter orchestration .	
Workaround: Since it is a manual intervention, shut followed by a no-shut on the interface can be done after removing the switch-port configurations.	
Recovery: Shut followed by a no-shut on the interface would put back the interface in port-profile mode.	

Defect ID: DEFECT000544928	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS5.0.0	**Technology Area:** Syslog
Symptom: RASLOG messages are not logged when terminal monitor is enabled for Telnet or SSH sessions.	
Condition: When terminal monitor is enabled for Telnet or SSH sessions	

Defect ID: DEFECT000545032	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Security
Reported In Release: NOS6.0.1	**Technology Area:** Security Vulnerability
Symptom: User is blocked if AAA user with role as "root" tries to login to the switch.	
Condition: When a AAA user with root privileges is created.	

Defect ID: DEFECT000545298	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.1	**Technology Area:** Other IPv4
Symptom: IP addressing with /31 mask is unsupported. Even though the command is accepted, applications such as ping do not work.	
Condition: While using the /31 subnet, ping and such applications would not work.	

Defect ID: DEFECT000545378	
Technical Severity: Low	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** NTP - Network Time Protocol
Symptom: Enhancement request to support allowing users to configure Source IP for the NTP messages.	
Condition: When required to source the NTP messages with non-interface management IP addresses.	

Defect ID: DEFECT000545417	
Technical Severity: High	Probability: Medium
Product: NOS	Technology: VCS
Reported In Release: NOS5.0.1	Technology Area: Logical Chassis
Symptom: Switch may exhibit FFDC RAS LOG message while joining the cluster after the downgrade from NOS5.0.1a to NOS5.0.1b.	
Condition: This issue may be seen in a 3-node logical chassis cluster while downgrading the cluster from NOS5.0.1a to NOS5.0.1b version.	

Defect ID: DEFECT000545470	
Technical Severity: High	Probability: Medium
Product: NOS	Technology: Layer 2
Reported In Release: NOS5.0.1	Technology Area: Overlay Gateway
Symptom: A link shut/no-shut can cause the BUM traffic to flood from one tunnel to another.	
Condition: This could happen in a full mesh topology of tunnels with VTEP IP's as VRRP-E gateway addresses.	

Defect ID: DEFECT000545670	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** Other IPv4
Symptom: Error is thrown when loading the saved config from a file that has IPv6 DHCP relay configuration on interface configured with IPv6 DHCP relay address	
Condition: Running-config has two entries for the DHCP relay address. When the running-config is copied and replayed on a different node, this error is seen.	

Recovery: The file generated from the command "copy running-config <file-name>" will have the IPv6 DHCP relay configuration in 2 lines as given below
 	
interface Ve 2184
  ipv6 dhcp relay address 2000::200:168:25:101 use-vrf default-vrf
  ipv6 dhcp relay address 2000::200:168:25:101 interface Ve 2184
 	
To recover from the error, re-configure the IPv6 DHCP relay configuration in a single line as shown below after loading the configuration from file completes.
 	
  ipv6 dhcp relay address 2000::200:168:25:101 use-vrf default-vrf interface ve 2184	

Defect ID: DEFECT000545751	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** BGP4+ (IPv6)
Symptom: BGP debug message for other IPv6 address is seen, while BGP debug is enabled for specific IPv6 address.	
Condition: BGP debug is enabled for specific IPv6 address.	

Defect ID: DEFECT000545815	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** DHCP (IPv4)
Symptom: Invalid configurations will be successful, but will not work.	
Condition: Invalid configuration of dhcp relay or gateway address, such as configuring gateway address before relay is configured, more than allowed number of relay/gateway address per interface, etc.	
Workaround: Configure with valid dhcp relay or gateway address.	
Recovery: Remove the relay and gateway addresses that are shown in the running configuration; configure correctly.	

Defect ID: DEFECT000545858	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.3	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: Enhance sysDescr OID to display switch name + model name + version.	
Condition: Enhance sysDescr OID to display switch name + model name + version.	

Defect ID: DEFECT000546039	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.0	**Technology Area:** Component
Symptom: "QSFP tx enable failed" Warning Message may be seen followed by Port Fault	
Condition: Repeated chassis disable and chassis enable commands are being executed	
Recovery: Reload the line card (LC).	

Defect ID: DEFECT000546254	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Security
Reported In Release: NOS5.0.1	**Technology Area:** Receive ACLs
Symptom: Applying ACL on Ethernet management interface fails with "% Error: Internal Error" error message	
Condition: The issue is observed when an attempt is made to apply ACL policy on Ethernet management interface with the following conditions: a. An ACL policy is already enforced on the Ethernet management interface. b. The new ACL policy and enforced policy names differ only in letter cases (e.g. "TEST001" vs "Test001").	
Workaround: Create ACLs with distinctive policy names (e.g. with different letters, numbers etc.).	
Recovery: On Ethernet management interface: a. Remove the enforced (active) ACL with "no ip access-group <ACL_name>" CLI. b. Enforce the new ACL policy (unique policy names) using "ip access-group <ACL_name> in".	

Defect ID: DEFECT000546366	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** Optics
Symptom: Enhancement request for the support for LM4 optics.	
Condition: Attempting to use LM4 optics.	

Defect ID: DEFECT000546441	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Build
Symptom: "show support" command might time out on a standby partition of the VDX 6740.	
Condition: Issuing the "show support" command	

Defect ID: DEFECT000546723	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** VMWare
Symptom: Output of show vnetwork vms is not updated after migrating a VM from one host to another using the migrate option.	
Condition: This scenario happens when a Virtual Machine is live migrated from one host to another host.	
Workaround: Drag and drop method worked.	
Recovery: Do a manual vCenter discovery or wait for the next scheduled discovery for the CLI output to get corrected.	

Defect ID: DEFECT000546937	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.0	**Technology Area:** Switch Management
Symptom: Unexpected reload when downgrading to 5.x code from 6.x. caused by dcm daemon.	
Condition: Downgrade from 6.x code to 5.x when the startup config still contains flexport on breakout ports configuration which is not supported in 5.0.1	
Workaround: Ensure that startup config doesn't contain flexport on breakout ports configuration before downgrade.	

Defect ID: DEFECT000547186	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** Optics
Symptom: A Line card may turn into FAULTY (51) state after copy default to start-up is issued followed by a system reload on VDX-8770	
Condition: When copy default to startup config is issued followed by a system reload.	

Defect ID: DEFECT000547210	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Build
Symptom: FFDC RAS LOG may be seen when "copy default-config startup-config" command is issued.	
Condition: Issuing a "copy default-config startup-config" command can result in FFDC log message.	

Defect ID: DEFECT000547271	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** Other IPv4
Symptom: Ping RTT delay when CPU load is high.	
Condition: When CPU intensive features such as ACL log and SFLOW are configured on VDX 6940. Can also be triggered intermittently due to internal housekeeping operations such as statistics being collected on the Switch	

Defect ID: DEFECT000547313	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** Virtual Fabrics
Symptom: "show mac" command displays MAC Address learnt on VF Extension Tunnel Interface even after MAC age-out timer expiry.	
Condition: MAC move from Physical interface to VF Extension Tunnel interface.	
Workaround: Issue "clear mac dynamic" command.	

Defect ID: DEFECT000547658	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** Conversational MAC Learning
Symptom: unexpected reload of VDX6740 while downgrading NOS.	
Condition: This may occur when the number of configured port security OUI (switchport port-security oui) is on more than 13 ports and NOS version is downgraded.	
Workaround: If the number of port security OUI's (switchport port-security oui) configured is fewer than 13 ports issue will not be observed during downgrade.	

Defect ID: DEFECT000547740	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** ARP
Symptom: VDX responding to ARP request for an IP address that does not belong to that subnet, even though proxy-arp has been disabled.	
Condition: ARP response handling for an IP address that does not belong to that subnet	

Defect ID: DEFECT000548284	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Traffic Management
Reported In Release: NOS5.0.1	**Technology Area:** QoS - Quality of Service
Symptom: Packet drops with "qos rcv-queue limit 8000" configuration.	
Condition: There is congestion (due to flooding), and one or a few queues consumed all buffers.	
Workaround: Clear mac address table takes care of the flooding. reduce the rcv_queue limit to 2000. this would limit packet drops on the congested data path only.	
Recovery: Remove any congestion. Reduce the rcv_queue limit.	

Defect ID: DEFECT000548822	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** Overlay Gateway
Symptom: Traffic fails when a VM residing on a hypervisor in one VCS CLuster is migrated to a remote VCS Cluster over VXLAN tunnels.	
Condition: Happens when the VM is moved to a destination hypervisor host connected to a non-VTEP RBridge in the destination VCS cluster.	

Workaround: Two workarounds:
    1) Create separate VCS clusters for Aggregation layer and Access layer and configure VTEP-Gateway on aggregation switches if hosts are connected to Access (Top-of-Rack) switches.
    OR
    2) Connect hosts to VTEP gateway RBridges.	

Recovery: Perform "clear mac-address table dynamic" on the destination VCS cluster after VM migration.	

Defect ID: DEFECT000548911	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Security
Reported In Release: NOS5.0.1	**Technology Area:** Port Security
Symptom: Switch may become unresponsive during heavy CPU usage due to certain IPv6 traffic. Error on console message: 1. Network OS is not ready 2. nf_conntrack table full : dropping packet 3. nf_conntrack: expectation table full 4. nf_ct_tftp: dropping packetIN= OUT=vlan...	
Condition: There are certain conditions that may cause the IPv6 traffic to be sent to the CPU.	
Workaround: Identify the traffic flows coming to the CPU and try to see if we can avoid those flows by topology change.	

Defect ID: DEFECT000549216	
Technical Severity: Low	**Probability:** Low
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.1	**Technology Area:** Other IPv4
Symptom: IP addressing with /127 mask is unsupported. Even though the command is accepted, the applications such as ping do not work.	
Condition: While using the /127 subnet, ping and such applications would not work.	

Defect ID: DEFECT000549484	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS4.1.3	**Technology Area:** Component
Symptom: Incorrect flash partition in VDX6740-1G-T. "df -h" output showing each partition 3.6GB instead of 1.8GB	
Condition: Net-install switch	

Defect ID: DEFECT000549632	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** Overlay Gateway
Symptom: MAC addresses learnt on tunnel won't be relearnt after issuing 'clear mac-address-table dynamic' command.	
Condition: VxLAN traffic will get affected.	
Workaround: Delete MAC addresses learnt on tunnel before issuing "clear mac-address-table dynamic' command.	

Defect ID: DEFECT000549638	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS6.0.1	**Technology Area:** High Availability
Symptom: Host name changes is not updated to standby MM.	
Condition: Manual host name changes and firmware download.	

Defect ID: DEFECT000549909	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Switch Management
Symptom: The customer gets an error when trying to configure any IP in 10.0.0.0/24 range on 8770 management interface VDX8770-2(config-Management-3/1)# ip add 10.0.0.31/24 Cannot set IPv4 Address Likewise, "chassis virtual-ip" cannot be set in this range.	
Condition: This issue can be seen in any affected release.	
Workaround: Not use 10.0.0.x/24 on management port	

Defect ID: DEFECT000550010	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.3	**Technology Area:** Switch Management
Symptom: switch might experience an unexpected reload during boot process, if inband virtual-ip is configured.	
Condition: The switch is in the fabric cluster mode and inband virtual-ip is configured.	

Defect ID: DEFECT000550064	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.1	**Technology Area:** BGP4 (IPv4)
Symptom: IP addressing with /127 mask is unsupported. Even though the command is accepted, the applications such as ping do not work.	
Condition: While using the /127 subnet, ping and such applications would not work.	

Defect ID: DEFECT000550132	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS4.0.1	**Technology Area:** Component
Symptom: Under very rare conditions, unexpected reload may occur indicating mlt_garbage_collector inconsistency	
Condition: In steady state when Memory tracking tool (MLT) encounters corruption.	

Defect ID: DEFECT000550177	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.1	**Technology Area:** Extended Fabrics
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: DEFECT000550500	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** CLI
Symptom: Switch may reload on execution of "show fabric ecmp group" command with source rbridge-id other than the rbridgde-id on which it is executed.	
Condition: Execution of "show fabric ecmp group" command with source rbridge-id other than the rbridgde-id on which it is executed.	
Workaround: Execute "show fabric ecmp group" command on source rbridge-id to avoid this issue.	

Defect ID: DEFECT000550550	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.1	**Technology Area:** VLAN
Symptom: Tunnels packets get flooded	
Condition: Shut down the edge port(s).	
Recovery: issue 'clear mac-address-table dynamic'.	

Defect ID: DEFECT000550765	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.1	**Technology Area:** Other IPv4
Symptom: After reload, switch replies to an arp from different subnet even though the proxy is disabled on that interface.	
Condition: Disabling proxy ARP is not persistent across switch reload.	
Recovery: Toggle the proxy-arp setting once the switch is reloaded.	

Defect ID: DEFECT000550967	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** CLI
Symptom: No Link up occurs when a connected VDX 6740T interface is configured with speed 100.	
Condition: VDX 6740T interface with "speed 100" configuration is connected to peer having "forced" speed 100.	

Defect ID: DEFECT000551157	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** VLAN
Symptom: After completion of ISSU, issuing a "show vlan brief" may adversely impact the switches in the VCS	
Condition: When the upgrade is done using ISSU & vlan names were configured in the earlier releases.	
Recovery: Reload the switch after ISSU upgrade, before issuing "show vlan brief" command.	

Defect ID: DEFECT000551169	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** BGP4 (IPv4)
Symptom: Some Routes may not be removed from the OSPF routing table and RTM table even after that neighbor from which it is learned is removed	
Condition: When an OSPF neighbor is shutdown	
Recovery: Clearing ip routes from the route-table would help refresh the routing entries.	

Defect ID: DEFECT000551418	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: When many VLANs are created on the switch and a snmpwalk operation is done on the IF MIB, the snmpwalk operation takes a long time to complete.	
Condition: This issue is seen only when snmpwalk is done on the IF MIB on a switch which has many VLANs created.	

Defect ID: DEFECT000551496	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** AMPP
Symptom: Some hosts using static port-profile may lose connectivity to network and show up in NotForwarding (NF) state in sh mac output	
Condition: When vcenter is enabled & disabled multiple times.	
Recovery: Depending upon the flavor of issue hit, port flap OR clearing mac table OR port unconfig & config OR switch reload would recover the issue.	

Defect ID: DEFECT000551559	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Monitoring/RAS
Reported In Release: NOS5.0.1	**Technology Area:** High Availability
Symptom: In rare cases during ISSU upgrade in 5.0.1x releases, we may see switch stuck with "application communication failure" message.	
Condition: Seen in ISSU upgrade from 5.0.1b to 5.0.1c	
Workaround: Reload will bring up the switch.	

Defect ID: DEFECT000551789	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** User Accounts
Symptom: User configured user names won't work after node rejoins into the cluster	
Condition: This can happens when the node got segmented and rejoined into the cluster	
Workaround: Remove the usernames and re-add	

Defect ID: DEFECT000551796	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** Component
Symptom: Switch may undergo unexpected reload.	
Condition: 'show port-profile status' command is executed	

Defect ID: DEFECT000551833	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Logical Chassis
Symptom: The ACTIVE Management module may be seen to panic and perform a failover to STDBY Management module upon executing the "show fabric ecmp group" command.	
Condition: When "show fabric ecmp group src-rbridge # dest-rbridge #" command executed where src-rbridge is not local domain then it causes panic.	

Workaround: Don't run "show fabric ecmp group src-rbridge # dest-rbridge #" command executed with src-rbridge not being local domain.

Defect ID: DEFECT000551912	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** AMPP
Symptom: show mac-add table command is not showing the correct output when we have the vlan configured with description.	
Condition: vlans configured with the description.	

Defect ID: DEFECT000551980	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Traffic Management
Reported In Release: NOS4.1.2	**Technology Area:** Routing
Symptom: A user initiated cluster node removal can cause I/O traffic disruptions longer than 1 second. The duration of the disruption can be impacted by many factors including the size of the cluster, number of connections to the node being removed, and the overall control plane load within the cluster.	
Condition: Multiple routes to the node being removed exist in the cluster. Multiple links to the node being removed exist in the cluster. User initiates a node removal via system reload or chassis disable.	
Recovery: Routes and services within the cluster will converge without any user intervention.	

Defect ID: DEFECT000552038	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** AMPP
Symptom: Under rare conditions, port-profiled source MAC's may not learn on new interface even if the traffic-stream has moved.	
Condition: When port-profiled MAC stream moves from one interface to another	

Defect ID: DEFECT000552158	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** Firmware Upgrade/ Downgrade
Symptom: VRF configuration doesn't apply properly on management interface	
Condition: Upgrade from 4.x to higher releases can cause this issue	
Recovery: Re-apply the VRF config on management interface	

Defect ID: DEFECT000552506	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** FCoE
Symptom: Traffic loss will be observed on FCoE LAG ports.	
Condition: FCoE default configuration applied on LAG port and switch is rebooted resulting in configuration replay.	
Recovery: Remove and add FCoE configuration on LAG port	

Defect ID: DEFECT000552726	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.0	**Technology Area:** VLAG
Symptom: Port-Channel showing the status as "Not connected" when one or more ports part of Port-channel are in that state.	
Condition: One or more ports in vLAG have interface status as "Not connected".	
Workaround: "show interface status" shows status of all the ports that are part of port-channel. If port channel status indicate "Notconnected", check for status of all the ports that are part of the port-channel.	

Defect ID: DEFECT000552736	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Traffic Management
Reported In Release: NOS5.0.1	**Technology Area:** BUM
Symptom: Added support for to capture the frames for the VDX2740 platform.	
Condition: Added support for to capture the frames.	

Defect ID: DEFECT000552816	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS4.0.0	**Technology Area:** CLI
Symptom: Enhancement request to support TFTP protocol for firmware upgrade process.	
Condition: During firmware download process.	

Defect ID: DEFECT000552832	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS4.1.3	**Technology Area:** VRRP & VRRP-E (IPv4)
Symptom: Ping to VRRP VIP will fail in some corner conditions	
Condition: This can happen after the VRRP master failover. After failover, the VRRP master node may learn an ARP with its own IP, but the MAC of VRRP backup node. This results in ping failures.	
Recovery: "clear arp ip <ip address> no-refresh"	

Defect ID: DEFECT000552949	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: SNMP MIB query on specific community string fails.	
Condition: It is a rare occurrence when community string was added and removed in a random sequence followed by reload/ node rejoin.	
Workaround: Reconfigure the same community.	

Defect ID: DEFECT000553165	
Technical Severity: High	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.1.3	**Technology Area:** IEEE 802.1d STP
Symptom: STP root port may not be correctly designated thus causing unexpected STP states. Additionally, PVST may go into FWDing state on vlag when principal fail-over occurs on uncontrolled failover	
Condition: After reload of principle node	

Defect ID: DEFECT000553370	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: SNMP MIB query on specific community string fails and host recipient associated with that specific community will not receive any SNMP trap.	
Condition: It is a rare occurrence when community string was added and removed in a random sequence followed by reload/ node rejoin/upgrade.	
Workaround: Reconfigure the same community.	

Defect ID: DEFECT000553435	
Technical Severity: Critical	**Probability:** Low
Product: NOS	**Technology:** Security
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Authentication
Symptom: In some rare cases a VDX switch may reload with these messages on the console. rascq full Fail to enqu drop rreq ty=305h cop=0h ctxt=4h reqh=137c8fh	
Condition: This would be seen during bootup when the available diskspace is minimal.	
Recovery: Remove any large (accumulated) files on the switch with the help of TAC.	

Defect ID: DEFECT000553561	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** Optics
Symptom: On unplugging and plugging back the TX side of the cable on 1G SFP, link did not come up again on the VDX Switch.	
Condition: Unplugging and plugging a cable from a 1G SFP may cause the link to not come back up.	
Recovery: shut/no shut will recover the port.	

Defect ID: DEFECT000553692	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.1	**Technology Area:** Component
Symptom: Enhancements made to support DC power supplies for 6740 & 6740-T platforms. 6740-T platforms now supports 500W DC power supply and 6740 platform now supports 250W DC power supply.	
Condition: 6740 & 6740-T platforms can run DC power supply units.	

Defect ID: DEFECT000553721	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.1	**Technology Area:** Fibre Channel Ports
Symptom: FCOE/FC hosts will not be able to discover FCOE/FC targets over E-Port or N-Port trunks if more than two trunks are present in a 8-port FC port connector group.	
Condition: Some FC port trunks will not work when more than two trunk masters are present on an 8-port connector group. When there are more than two unique trunk masters the masters the third and fourth trunk will not reliable pass Rx frames to destinations. All 8-ports are still allowed to trunk in any combination as long as only two unique masters are present. This issue is caused by problems with switch hardware specific to FC-port trunking. This issue does not affect Ethernet port trunking.	
Workaround: If the user does not configure more than two trunks in a connector group (8 trunkable FC ports) then the issue will not be seen. However if the configuration ever allowed more than two trunk masters and the customer disabled some of the trunks it is still vulnerable since the active trunks may still use the problematic trunk hardware. It is recommended that the user configure the switch so a max of two trunks are possible in the group and then disable/re-enable all ports in the group if the issue is suspected.	
Recovery: If the configuration ever allowed more than two trunk masters and the user disabled some of the trunks it is still vulnerable since the active trunks may still use the problematic trunk state. It is recommended that the user configure the switch so max of two trunks are possible in the group and then disable/re-enable all ports in the group if the issue is suspected.	

Defect ID: DEFECT000553787	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** CLI
Symptom: Enhancement Request to log show commands in TACACS accounting	
Condition: When TACACS accounting is enabled.	

Defect ID: DEFECT000554216	
Technical Severity: Low	**Probability:** Low
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Logical Chassis
Symptom: Unhide the useful "vcs auto-shut vlag" command	
Condition: Unhide the useful "vcs auto-shut vlag" command	

Defect ID: DEFECT000554231	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** SNMPv2, SNMPv3 & MIBs
Symptom: Observe authorizationError for specific user during SNMP walk.	
Condition: This issue may hit when any SNMP group configuration was removed and added continuously in different sequence followed by node rejoin/reload.	
Workaround: Reconfigure the same community or user.	

Defect ID: DEFECT000554322	
Technical Severity: Low	Probability: Low
Product: NOS	Technology: Management
Reported In Release: NOS5.0.2	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: SNMP query & trap support for default VRF.	
Condition: SNMP query & trap support for default VRF.	

Defect ID: DEFECT000554855	
Technical Severity: Medium	Probability: Low
Product: NOS	Technology: Management
Reported In Release: NOS5.0.1	Technology Area: SFLOW
Symptom: Feature enhancement to support sflow collector per vrf.	
Condition: Sflow collector configuration per vrf basis.	

Defect ID: DEFECT000554933	
Technical Severity: Critical	Probability: Medium
Product: NOS	Technology: Layer 2
Reported In Release: NOS5.0.1	Technology Area: IEEE 802.1s MSTP
Symptom: The switch may encounter traffic loss when configuring "no spanning-tree shutdown" on a port-channel to rectify configuration across PO members	
Condition: When certain member switches of PO are not configured with the same configuration.	
Recovery: Flapping the port-channel interface on both the switches would recover the traffic losses.	

Defect ID: DEFECT000555418	
Technical Severity: High	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS4.1.3	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: When a snmpwalk operation is done on the ipAddrTable, VE interfaces that have IP addresses assigned to them will not show up.	
Condition: One or more VE interfaces with IP address assigned to them and snmpwalk operation is done on ipAddrTable.	

Defect ID: DEFECT000555450	
Technical Severity: Medium	Probability: High
Product: NOS	Technology: System
Reported In Release: NOS5.0.1	Technology Area: CLI
Symptom: show ip route command with "longer" option shows number of routes as 0 even when routes are present with the specified prefix. For example, sw0# sh ip route 192.168.1.0/29 longer Total number of IP routes: 0	
Condition: When "longer" option is used in the "Show ip route" command.	
Workaround: Use the "show ip route" command with filters.	

Defect ID: DEFECT000555517

Technical Severity: High	Probability: High
Product: NOS	Technology: Layer 2
Reported In Release: NOS6.0.1	Technology Area: VLAG

Symptom: Interface STP State may be moved to designated discarding state after port-channel member flap

Condition: When port-channel is flapped.

Defect ID: DEFECT000556023

Technical Severity: High	Probability: High
Product: NOS	Technology: Monitoring/RAS
Reported In Release: NOS5.0.0	Technology Area: Syslog

Symptom: Enhancement request to allow user to configure syslog server per VRF

Condition: When syslog server needs to be added to default VRF.

Defect ID: DEFECT000556088

Technical Severity: High	Probability: High
Product: NOS	Technology: Layer 2
Reported In Release: NOS5.0.1	Technology Area: MAC ACLs

Symptom: ELD may not detect layer-2 loops for VDX6740 & VDX8770

Condition: When layer-2 loops are created

Workaround: May use Spanning-tree instead for detecting & blocking loops.

Defect ID: DEFECT000556136

Technical Severity: Critical	Probability: Medium
Product: NOS	Technology: Layer 3
Reported In Release: NOS5.0.1	Technology Area: Other IPv4

Symptom: Under rare scenarios, VDX6740 may encounter unexpected reloads while processing certain UDP packets.

Condition: When UDP packets are punted to the CPU (eg. VRRP-E hellos).

Defect ID: DEFECT000556254

Technical Severity: High	Probability: High
Product: NOS	Technology: Management
Reported In Release: NOS6.0.1	Technology Area: Access Gateway

Symptom: Port group policy is disabled after performing ISSU on the switch and enabling Access Gateway feature.

Condition: When the switch goes through ISSU upgrade and Access Gateway feature is enabled, the port group policy isn't getting enabled.

Defect ID: DEFECT000556317

Technical Severity: High	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS4.1.3	Technology Area: Firmware Upgrade/ Downgrade

Symptom: "application communication failure" error is observed when the switch is coming up. The switch is unresponsive for a while before customer sees a daemon crash.

Condition: "copy running-config startup-config" is triggered when switch is coming up and before HA is in sync.

Workaround: Trigger "copy running-config startup-config" only after HA is in sync. This can be checked using "show ha" command

Recovery: Reboot the system.

Defect ID: DEFECT000556657	
Technical Severity: High	**Probability:** Low
Product: NOS	**Technology:** System
Reported In Release: NOS5.0.1	**Technology Area:** Optics
Symptom: Enhancement request to support 100MB speed on Copper SFP for VDX-6740.	
Condition: When using Copper SFP (specifically 57-1000042-01).	

Defect ID: DEFECT000556966	
Technical Severity: Low	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.1	**Technology Area:** Virtual Fabrics
Symptom: CPU gets interrupts because of a port ASIC I/O read DMA channel is stuck and consequently it's usage becomes high.	
Condition: Occurence is unexpected hence our software is made robust to handle such scenarios.	
Workaround: Power cycle the affected line card.	

Defect ID: DEFECT000558105	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS5.0.1	**Technology Area:** NTP - Network Time Protocol
Symptom: NTP Authentication Key with non-alphanumeric characters is not accepted.	
Condition: NTP Authentication Key with non-alphanumeric characters not accepted	

Defect ID: DEFECT000558159	
Technical Severity: Medium	**Probability:** Medium
Product: NOS	**Technology:** Security
Reported In Release: NOS5.0.0	**Technology Area:** Security Vulnerability
Symptom: Enhancement request to provide a script to selectively restrict SSH traffic to certain VRF's.	
Condition: When need to restrict SSH to only a specific VRF	

Defect ID: DEFECT000559620	
Technical Severity: Medium	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS6.0.1	**Technology Area:** Switch Management
Symptom: Power cycle operations can in rare scenarios can cause Switch Database (DB) corruptions. Sometimes the file system can get into such a state where DB integrity check itself gets stuck. To address this the SW implements a 10 minute timer during which the DB integrity check should complete. If not the Switch SW assumes that the DB is corrupted and proceeds to cleanup steps.	
Condition: Time has been reduced to 8 min in 6.0.1 release after testing with large configuration.	

Defect ID: DEFECT000559812	
Technical Severity: Medium	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS5.0.0	**Technology Area:** ACLs (IPv4)
Symptom: Deploying IPv4 ACLs using OpenStack was slow.	
Condition: With this 6.0.1 release, should be able to use IPV4 ACLs using OpenStack, netconf as performance issues are addressed.	

Defect ID: DEFECT000560072	
Technical Severity: High	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.1	**Technology Area:** CLI

Symptom:	While restoring configuration from a file by executing 'copy <file> running-config, an exception might be reported on a console for a process ConfigCmd. Exception is usually seen at the time of process exit and will not impact switch functionality	
Condition:	Restoring configuration from a file by executing copy <file> running-config.	
Recovery:	No recovery is needed as this process will be spawned again for next copy command.	

Defect ID: DEFECT000560266		
Technical Severity: High	**Probability:** High	
Product: NOS	**Technology:** System	
Reported In Release: NOS5.0.1	**Technology Area:** Component	
Symptom: After HA failover, snmp v3 walk may result in authorization errors.		
Condition: The issue is applicable only for view and group configuration for SNMPv3. User/community/host/v3host configuration will not have any impact during hafailover.		
Recovery: Remove & Re-add the group / view after HA failover.		

Defect ID: DEFECT000560831		
Technical Severity: High	**Probability:** Medium	
Product: NOS	**Technology:** System	
Reported In Release: NOS6.0.1	**Technology Area:** Component	
Symptom: Openflowd daemon crash causing system to reboot or uncontrolled failover.		
Condition: System running low on memory.		

Closed without code changes for NOS 6.0.2g

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of January 4th, 2018 in Network OS 6.0.2g.

None.

Closed without code changes for NOS 6.0.2f

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of July 10, 2017 in Network OS 6.0.2f.

None.

Closed without code changes for NOS 6.0.2e

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of April 10, 2017 in Network OS 6.0.2e.

None.

Closed without code changes for NOS 6.0.2d

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of March3, 2017 in Network OS 6.0.2d.

None.

Closed without code changes for NOS 6.0.2c

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of November 11, 2016 in Network OS 6.0.2c.

None.

Closed without code changes for NOS 6.0.2b

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of May 20, 2016 in Network OS 6.0.2b.

None.

Closed without code changes for NOS 6.0.2a

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of November 30, 2015 in Network OS 6.0.2a.

None.

Closed without code changes for NOS 6.0.2

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of November 30, 2015 in Network OS 6.0.2.

None.

Closed without code changes for NOS 6.0.1

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of July 29, 2015 in Network OS 6.0.1.

Defect ID: DEFECT000365558	**Technical Severity:** Low
Reason Code: Already Fixed in Release	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS2.1.0	**Technology Area:** Telnet

Symptom: While trying to login as 'admin', user might encounter the following response -
       login: admin
       Password:
       Error :Success. Network OS is not ready. Please login after sometime.	

Condition: Logging in following an upgrade as admin.	
Recovery: Power cycle the affected switch.	

Defect ID: DEFECT000389383	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** Medium
Product: NOS	**Technology:** Management
Reported In Release: NOS2.1.1_sp	**Technology Area:** Switch Management
Symptom: When switch reboots, Login Accounting Stop packets are not sent to the console.	
Condition: Switch reboot	

Defect ID: DEFECT000445592	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** Medium
Product: NOS	**Technology:** System
Reported In Release: NOS4.0.0	**Technology Area:** CLI
Symptom: Deleting one or more VLANs, Port-channels, VE interfaces, or Loopback interfaces using Range command within range submode does not update the Range string displayed in the NOS prompt submode.	
Condition: Creating and deleting multiple VLANs, port-channels, VE interfaces, or Loopback interfaces into range submode.	

Defect ID: DEFECT000475754	**Technical Severity:** Medium
Reason Code: Will Not Fix	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.0.0	**Technology Area:** VLAN
Symptom: Unable to undo Reserved-VLAN configuration using 'no' form of same config command.	
Condition: Resetting the default value of the Reserved-VLAN configuration using the 'no' form of the command.	
Workaround: Use same configuration command 'reserved-vlan' with default range for that particular hardware platform.	

Defect ID: DEFECT000476325	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Low
Product: NOS	**Technology:** Management
Reported In Release: NOS4.1.0	**Technology Area:** Switch Management
Symptom: VDX 6740 may observe "application communication failure" when a command is issued.	
Condition: Virtual-fabric scale configuration with more than 2K virtual-fabric entries.	

Defect ID: DEFECT000487152	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Low
Product: NOS	**Technology:** Layer 2

Reported In Release: NOS4.1.0	**Technology Area:** VLAN
Symptom: Spanning-tree related information of a Virtual Fabric is not getting removed from the interface even after unconfiguring the Virtual Fabric from that interface.	
Condition: Configure and unconfigure Virtual Fabric VLAN on a spanning-tree edge port.	
Recovery: Use "no spanning-tree vlan xx" command for unconfiguring the spanning-tree related features enabled for a VF after unconfiguring VF from the interface.	

Defect ID: DEFECT000491047	**Technical Severity:** Low
Reason Code: Will Not Fix	**Probability:** Low
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS4.0.1	**Technology Area:** VLAN
Symptom: The output of command "show mac-address-table port-channel <poId>" is not sorted according to VLANs when data is shown from different RBridges.	
Condition: User will observe this issue on execution of command "show mac-address-table port-channel <poId>".	

Defect ID: DEFECT000493408	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** Medium
Product: NOS	**Technology:** IP Multicast
Reported In Release: NOS4.0.1	**Technology Area:** IPv4 Multicast Routing
Symptom: Multicast traffic does not recover after fourth iteration of "shut / no shut Ve interfaces", or second iteration of chassis disable/enable.	
Condition: After multiple iterations of Chassis/Disable/enable, shut/no shut and ISSU	

Defect ID: DEFECT000512265	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** Medium
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.0	**Technology Area:** MAC ACLs
Symptom: If Source MAC learning is disabled on an interface for a VLAN, then DA MACs are also not learned on that interface	
Condition: When both source MAC learning is disabled and CML is enabled on 6740	
Workaround: Do not configure source MAC learning disable and CML enable together on 6740	

Defect ID: DEFECT000514569	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS5.0.0	**Technology Area:** FCoE
Symptom: FEC-ENABLE is inactive even though both the sides have FEC enabled.	
Condition: VDX6740 in AG mode with 16Gbps N-Port connections configured with FEC enabled.	

Defect ID: DEFECT000515063	**Technical Severity:** Medium
Reason Code: Not Reproducible	**Probability:** Low
Product: NOS	**Technology:** Other
Reported In Release: NOS4.1.0	**Technology Area:** Other
Symptom: "user" entries differ from audit log entries in LC mode	
Condition: Comparing supportsave and audit log information.	

Defect ID: DEFECT000516464	Technical Severity: High
Reason Code: Not Reproducible	Probability: High
Product: NOS	Technology: Layer 3
Reported In Release: NOS5.0.0	Technology Area: Extended Fabrics
Symptom: When VXLAN traffic translating a VLAN to VNI is egressing out of a port, that port needs to be part of the VLAN. Removing the port out of the VLAN will not immediately stop the traffic but may take a few seconds depending on the load on the switch.	
Condition: The issue is observed under following conditions: 1) 2 port-channels configured and spanning-tree enabled and port-channel is in trunk mode. Only one port-channel is active at a time. 2) Change the mode on active port-channel from trunk to access and configured same VLAN.	

Defect ID: DEFECT000519985	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: NOS	Technology: Layer 3
Reported In Release: NOS5.0.0	Technology Area: BGP4 (IPv4)
Symptom: When configuring 'default-originate' with route-map in neighbor command, the default route is always originated for the neighbor.	
Condition: With conditional default-origination, default route should be generated only when route-map matching prefix is present in the IP routing table. Whereas in this case, irrespective of whether route-map matching prefix is present in the IP routing table or not default route is originated to the neighbor.	

Defect ID: DEFECT000523999	Technical Severity: High
Reason Code: Not Reproducible	Probability: Medium
Product: NOS	Technology: Other
Reported In Release: NOS5.0.0	Technology Area: Other
Symptom: While doing copy default start-up, observing VERIFY logs and DCMD also not initializing properly when 6740 is coming up in LC mode.	
Condition: When performing copy default start-up in 6740 in LC mode.	

Defect ID: DEFECT000524664	Technical Severity: High
Reason Code: Not Reproducible	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS5.0.0	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: Copy support save operation may fail due to MAPS module timeout. During the support save the MAPS module is failing due to the IPC failed with return code "Timeout".	
Condition: This issue can be seen while copying support save in the Logical Cluster mode.	

Defect ID: DEFECT000525084	Technical Severity: High
Reason Code: Not Reproducible	Probability: Low
Product: NOS	Technology: Management
Reported In Release: NOS5.0.0	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: While performing In Service Software Upgrade on VDX 6740, SNMP process restart was observed	
Condition: ISSU on VDX 6740.	

Defect ID: DEFECT000525683	Technical Severity: High
Reason Code: Not Reproducible	Probability: Medium
Product: NOS	Technology: System
Reported In Release: NOS5.0.0	Technology Area: CLI
Symptom: Switch stopped responding abruptly and one daemon (Confd) went into a reboot state.	
Condition: Traffic via various interfaces CLI, Netconf, and REST gateway.	
Recovery: Rebooting the switch.	

Defect ID: DEFECT000525777	Technical Severity: Medium
Reason Code: Already Fixed in Release	Probability: Medium
Product: NOS	Technology: Layer 2
Reported In Release: NOS5.0.0	Technology Area: VLAN
Symptom: TVLAN with a range of ctag may not be enabled on an edge port on doing switchport trunk allowed VLAN add if the TVLAN is already mapped to a VNI and extended on a site under overlay gateway.	
Condition: Occurs when site extend configurations with TVLAN are done before edge port configurations for the same TVLAN.	
Workaround: Remove both extend command under site and switchport trunk on edgeport. Apply the switchport trunk allowed VLAN add on the edgeport. After that apply the extend command under site.	

Defect ID: DEFECT000526824	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS4.1.2	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: An snmpget on any attribute in the IfTable of IF-MIB, based on an if-index that is similar to port-id on the front panel port will fail.	
Condition: This issue is seen only when snmpget operation is done on attributes in ifTable of IF MIB based on an erroneous ifIndex.	
Workaround: The workaround is to do a snmpwalk which will retrieve all the attributes of the ifTable along with their indices. Then a snmpget can be done based on that ifindex that is retrieved.	

Defect ID: DEFECT000528704	Technical Severity: High
Reason Code: Not Reproducible	Probability: High
Product: NOS	Technology: System
Reported In Release: NOS4.1.3	Technology Area: CLI
Symptom: It is cosmetic issue.Few of configurations related to icmp won't appear under running config.	
Condition: After loading the configuration from external server.	

Defect ID: DEFECT000529927	Technical Severity: High
Reason Code: Design Limitation	Probability: High
Product: NOS	Technology: Monitoring/RAS
Reported In Release: NOS4.1.2	Technology Area: PCAP
Symptom: If debug commands and terminal monitor are enabled on the console, then after some time, the switch may experience an unexpected reload due to termination of process Dcmd.	
Condition: 1.If the rate of the incoming packets to the CPU is more than 30 pkts/sec or more 2. When terminal monitor is enabled on console	
Workaround: Enable terminal monitor on telnet session instead on console	

Defect ID: DEFECT000530296	Technical Severity: High
Reason Code: Not Reproducible	Probability: Low
Product: NOS	Technology: Layer 2
Reported In Release: NOS5.0.1	Technology Area: FCoE
Symptom: With large configuration, switch may fail on ISSU from 5.0.0 to 5.0.1 release. The user may notice a FCoE failure during the ISSU.	
Condition: This problem is particularly seen when there are FCoE configurations and a ISSU is attempted from 5.0.0 to 5.0.1.	
Recovery: The switch reboots after the failure and comes back up automatically.	

Defect ID: DEFECT000535864	Technical Severity: High
Reason Code: Design Limitation	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS4.1.3	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: The MIB object dot1qVlanCurrentEgressPorts in the Q-BRIDGE-MIB when queried does not contain information related to port-channel interfaces in a VLAN. Also, physical interfaces that are present in slots 5 to 8 on VDX 8770 are not shown.	
Condition: This issue is seen only when port-channel interfaces are part of a VLAN that is active and/or physical interfaces that are in slots 5 to 8 are part of a VLAN that is active.	

Defect ID: DEFECT000536309	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: High
Product: NOS	Technology: Monitoring/RAS
Reported In Release: NOS5.0.1	Technology Area: RMON - Remote Network Monitoring
Symptom: RMON statistics cannot be configured on a ten gigabit Ethernet interface. An attempt to do so results in the error: % Error: Statistics index already enabled on another interface.	
Condition: This error occurs after break-out is de-configured on an interface and then RMON statistics are attempted to be configured on that interface.	

Defect ID: DEFECT000537283	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: NOS	Technology: System
Reported In Release: NOS4.1.3	Technology Area: CLI
Symptom: Switch is not using the configured DHCP Gateway IP Address	
Condition: When multiple IP addresses are configured on a VE interface, DHCP gateway IP address does not use the newly configured DHCP IP address.	

Defect ID: DEFECT000537306	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS4.1.3	Technology Area: DHCP (IPv4)
Symptom: CLI "ip dhcp relay gateway <IPv4 address>" does not throw error in case CLI executed more than once with different address. This should throw error by saying that gateway address is already configured if CLI is executed more than once.	
Condition: This scenario occur in case when user executes CLI "ip dhcp relay gateway <IPv4 address>" more than once.	
Workaround: No work around is required as it does not overwrite the gateway address configured first time. This is just a cosmetic issue.	

Defect ID: DEFECT000537703	Technical Severity:	High
Reason Code: Will Not Fix	Probability: High	
Product: NOS	Technology: Layer 2	
Reported In Release: NOS6.0.0	Technology Area: VLAN	
Symptom: Port-Profile with VLAN classification doesn't get applied on the profiled port.		
Condition: conflicting Port-Profiles are added to a single profile domain.		
Recovery: Remove conflicting Port-Profiles from profile domain before deleting UpgradedVlanProfile from default domain.		

Defect ID: DEFECT000537750	Technical Severity:	High
Reason Code: Will Not Fix	Probability: High	
Product: NOS	Technology: System	
Reported In Release: NOS4.1.3	Technology Area: CLI	
Symptom: No error message is displayed while removing the last relay address from an interface.		
Condition: Trying to remove the last relay address from an interface if gateway is configured.		

Defect ID: DEFECT000539170	Technical Severity:	Medium
Reason Code: Already Fixed in Release	Probability: Medium	
Product: NOS	Technology: System	
Reported In Release: NOS4.1.3	Technology Area: Component	
Symptom: Power cycle test places any of the 40-gig i/fs in admin down state.		
Condition: Continuous power-cycling causing the issue		

Defect ID: DEFECT000539176	Technical Severity:	Medium
Reason Code: Already Fixed in Release	Probability: Medium	
Product: NOS	Technology: System	
Reported In Release: NOS4.1.3	Technology Area: Component	
Symptom: VDX 6740 with loopback connections may bring one of the 40G interfaces as admin down.		
Condition: Continuous power-cycling of VDX 6740 over an extended period of time.		

Defect ID: DEFECT000539911	Technical Severity:	High
Reason Code: Not Reproducible	Probability: High	
Product: NOS	Technology: Layer 3	
Reported In Release: NOS6.0.0	Technology Area: OSPF (IPv4)	
Symptom: "show ip ospf database grace-link-state" command will show Grace LSA even after Grace LSA aged-out.		
Condition: Multiple HA failover.		
Recovery: Issue "clear ip ospf all".		

Defect ID: DEFECT000540539	Technical Severity:	High
Reason Code: Not Reproducible	Probability: Low	
Product: NOS	Technology: Layer 2	
Reported In Release: NOS6.0.0	Technology Area: Link Aggregation	
Symptom: Port-channel Redundancy Group failover can take longer than expected time in scaled environment.		
Condition: Triggering port-channel-redundancy failover on a port-channel with greater than 1.8K VLANs configured.		

Defect ID: DEFECT000541040	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: NOS	**Technology:** Layer 2
Reported In Release: NOS6.0.0	**Technology Area:** FCoE
Symptom: "show fcoe interface ethernet" shows the Ethernet interface, even after it has been converted the flexport to FC port	
Condition: It is seen when Ethernet port has FCoE provisioning configuration before the conversion to FC port.	
Workaround: Remove FCoE provisioning configuration from Ethernet port before converting the flexport to FC port	
Recovery: Move the flexport back to Ethernet port, remove the FCoE provisioning configuration and convert it back to Ethernet	

Defect ID: DEFECT000543072	**Technical Severity:** High
Reason Code: Will Not Fix	**Probability:** High
Product: NOS	**Technology:** VCS
Reported In Release: NOS5.0.1	**Technology Area:** Fabric Build
Symptom: On upgrade from 4.x to 5.x, the IPv6 ACL rule stating "seq 1 permit ip any any" would be converted into "seq 1 permit any any".	
Condition: When upgrading from 4.x to 5.x release.	

Defect ID: DEFECT000543581	**Technical Severity:** High
Reason Code: Design Limitation	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** BGP4 (IPv4)
Symptom: Minimal amount of data traffic loss is observed via BGP learned routes from Password enabled BGP Peers with lower restart threshold limit configured.	
Condition: Execute Hafailover/upgrade, configured lower threshold limit in BGP Graceful Restart timer with Password enabled BGP peers	

Defect ID: DEFECT000543744	**Technical Severity:** Medium
Reason Code: Design Limitation	**Probability:** High
Product: NOS	**Technology:** Layer 3
Reported In Release: NOS6.0.0	**Technology Area:** OSPFv3 (IPv6)
Symptom: Disabling OSPFv3 protocol without the VRF option results in error being displayed.	
Condition: This issue is observed when trying to disable OSPFv3 protocol without the VRF option.	
Workaround: Use the command with VRF option; "no ipv6 router ospf vrf default-vrf"	

Defect ID: DEFECT000545584	**Technical Severity:** High
Reason Code: Not Reproducible	**Probability:** High
Product: NOS	**Technology:** System
Reported In Release: NOS6.0.0	**Technology Area:** Optics
Symptom: Under rare conditions, an unsupported OR bad cable / media could impact the switch functionality with high rate of link-flaps. The switch may not be able to form new ISL links when impacted.	
Condition: With bad cable / media that would inject high rate of link-flaps.	
Recovery: Shutdown the impacted interface.	

Defect ID: DEFECT000545878	Technical Severity: High
Reason Code: Not Reproducible	Probability: High
Product: NOS	Technology: Layer 2
Reported In Release: NOS6.0.0	Technology Area: VLAG
Symptom: Observe traffic loss when egress traffic rate approaches the total aggregate bandwidth of the Extreme LAG port-channel.	
Condition: This is due to the ASIC fill and spill model for traffic forwarding on Extreme trunks. Applicable only when 6740 or 6940 is egressing traffic.	

Defect ID: DEFECT000546764	Technical Severity: High
Reason Code: Not Reproducible	Probability: Medium
Product: NOS	Technology: System
Reported In Release: NOS6.0.0	Technology Area: Optics
Symptom: On VDX 2740 CBR-5047 FFDC errors may be observed along with some Front-ends port going link down	
Condition: Reload operation	
Workaround: shut/no shut of the port	

Defect ID: DEFECT000547511	Technical Severity: High
Reason Code: Not Reproducible	Probability: Low
Product: NOS	Technology: Management
Reported In Release: NOS6.0.0	Technology Area: Access Gateway
Symptom: Fibrechannel port (N_Port) on Access Gateway enabled switch may not come online after disabling and enabling Access Gateway mode.	
Condition: If switch which is in Access Gateway mode, has multiple Fibrechannel ports (N_Port) and these ports are connected to different fabrics.	
Workaround: Keep all Fibrechannel ports in shutdown state before changing Access Gateway mode.	
Recovery: Toggling the port.	

Defect ID: DEFECT000548674	Technical Severity: High
Reason Code: Not Reproducible	Probability: Medium
Product: NOS	Technology: Layer 2
Reported In Release: NOS6.0.0	Technology Area: VLAG
Symptom: "VxLAN tunnel may not come online due to mac not getting resolved for next hop. When the traffic with 60k or more macs is sent over the tunnel, and reload is issued on one of the RBridge, tunnel may not come online.	
Condition: Reloading a switch while large number of flows are over a tunnel.	
Workaround: " clear mac-address-table dynamic " resolves the issue.	

Defect ID: DEFECT000549358	Technical Severity: Medium
Reason Code: Feature/Function Not Supported	Probability: High
Product: NOS	Technology: Security
Reported In Release: NOS4.1.3	Technology Area: Receive ACLs
Symptom: The following error message may be seen on VDX-6710, 6720, 6730 series of switches when a standard access-list is applied on the management port. sw0(config)# interface Management 120/0 sw0(config-Management-120/0)# ip access-group 10 in ERROR: Failed to enforce new iptables rules iptables-restore: line 57 failed	
Condition: The error message is seen only in NOS 4.1.3a1 release.	
Workaround: The message is cosmetic and can be ignored.	

Defect ID: DEFECT000549960	Technical Severity: High
Reason Code: Already Fixed in Release	Probability: High
Product: NOS	Technology: System
Reported In Release: NOS4.1.3	Technology Area: CLI
Symptom: A VDX 8770 RBridge may reload unexpectedly while copying a configuration file from flash to running.	
Condition: This is in seen in affected releases when port-group <ifname> mode performance command is used.	
Workaround: Avoid the port-group <ifname> command.	

Defect ID: DEFECT000552254	Technical Severity: Medium
Reason Code: Already Fixed in Release	Probability: Medium
Product: NOS	Technology: IP Multicast
Reported In Release: NOS4.1.3	Technology Area: IGMP
Symptom: When RSPAN is configured and egress IGMP packets are mirrored back into the VCS, IGMP query packets congest the control packet queue to CPU causing OSPF packets to drop.	
Condition: When RSPAN is configured and egress IGMP packets are mirrored back into the VCS, IGMP query packets congest the control packet queue to CPU causing OSPF packets to drop.	
Workaround: The issue does not happen when RSPAN is disabled.	

Defect ID: DEFECT000552415	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: NOS	Technology: Management
Reported In Release: NOS4.1.2	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: NOS version prior to 5.0 will not generate swEvent trap (indicating a RASLOG being logged) for interface up/down events	
Condition: Switch running NOS < 5.0	

Defect ID: DEFECT000556335	Technical Severity: Medium
Reason Code: Feature/Function Not Supported	Probability: High
Product: NOS	Technology: Management
Reported In Release: NOS5.0.1	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: SNMPv3 user created in earlier release may lose the "read" permissions after upgrade from 4.x to 5.x release	
Condition: On upgrade from 4.x to 5.x release without the fix	
Recovery: Re-assign the permissions manually for the required users	

Defect ID: DEFECT000557010	Technical Severity: High
Reason Code: Feature/Function Not Supported	Probability: Medium
Product: NOS	Technology: Layer 2
Reported In Release: NOS4.1.3	Technology Area: FCoE
Symptom: VDX 6730 cannot run 6.x firmware, which is expected behavior.	
Condition: VDX 6730 cannot run 6.x firmware, which is expected behavior. Hence, this issue will not be seen by customer.	

Defect ID: DEFECT000557049	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: NOS	Technology: Layer 3
Reported In Release: NOS6.0.1	Technology Area: BGP4 (IPv4)
Symptom: Layer 3 IPv6 traffic is getting dropped after ISSU	
Condition: ISSU may lead to IPv6 traffic drop	

Defect ID: DEFECT000559091	Technical Severity: Medium
Reason Code: Already Fixed in Release	Probability: Low
Product: NOS	Technology: Layer 3
Reported In Release: NOS5.0.1	Technology Area: VRRP & VRRP-E (IPv4)
Symptom: Under certain rare conditions, the hosts on a specific vlan connected to the TOR VDX6740 switches could not access their VRRP VIP gateway configured on spine switches.	
Condition: When hosts connected directly to VDX6740 TOR layer-2 switches & spine configured with VRRP-E.	
Recovery: VRRP failover OR unconfiguring & configuring the VE / VRRP interface would recover the issue.	

Defect ID: DEFECT000560037	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: High
Product: NOS	Technology: Management
Reported In Release: NOS5.0.2	Technology Area: SNMPv2, SNMPv3 & MIBs
Symptom: Introducing support 3-tuple interface for ifDescr and ifname object of IF MIB.	
Condition: ifDescr and ifname object of IF MIB will have the 3-tuples.	

Defect ID: DEFECT000560942	Technical Severity: High
Reason Code: Already Fixed in Release	Probability: Medium
Product: NOS	Technology: System
Reported In Release: NOS6.0.1	Technology Area: Component
Symptom: Switch will be in Faulty state during boot-up and all ports will be offline.	
Condition: Issue will only be observed on firmware version 5.x when the 40G ports have external traffic coming in during diagnostics post test that is run on switch boot-up.	
Workaround: 1. Disable diagnostics POST test using command "configure" and "no diag post rbridgeid # enable". 2. Stop external traffic on 40G ports/Remove 40G connections till switch boot-up is complete and switch is online.	

Known Issues for NOS 6.0.2g

This section lists known issues as of January 4th, 2018 in Network OS 6.0.2g.

None.

Known Issues for NOS 6.0.2f

This section lists known issues as of July 10, 2017 in Network OS 6.0.2f.

None.

Known Issues for NOS 6.0.2e

This section lists known issues as of April 10, 2017 in Network OS 6.0.2e.

None.

Known Issues for NOS 6.0.2d

This section lists known issues as of March3, 2017 in Network OS 6.0.2d.

None.

Known Issues for NOS 6.0.2c

This section lists known changes as of November 11, 2016 in Network OS 6.0.2c.

Defect ID: DEFECT000620197	
Technical Severity: High	**Probability:** Low
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: 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>	
Workaround: After config-replay fails to configure OSPF authentication key on the interface, it is possible to configure authentication key using CLI.	

Defect ID: DEFECT000570268	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Software Installation & Upgrade
Reported In Release: NOS7.0.0	**Technology:** Management
Symptom: Disruptive firmware upgrade (coldboot) will fail.	
Condition: Number of snmp communities associated with ipv4/ipv6 ACL configurations is greater than 20.	
Workaround: Limit the number of snmp communities associated withipv4/ipv6 ACL configurations to less than 20.	

Defect ID: DEFECT000596774	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** IP Addressing
Reported In Release: NOS5.0.2	**Technology:** Layer 3 Routing/Network Layer
Symptom: Switch reloads with termination of ribmgr daemon	
Condition: Static route is leaked to multiple VRF's.	
Workaround: Do not configure a static route more than once with the next-hop belonging to different VRF's. If mgmt-vrf has default route, delete default route and reload the VDX.	

Defect ID: DEFECT000624805	
Technical Severity: High	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** IP Addressing
Reported In Release: NOS7.0.1	**Technology:** Layer 3 Routing/Network Layer
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: MHigh	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** IP Adressing
Reported In Release: NOS7.0.1	**Technology:** Layer 3 Routing/Network Layer
Symptom: "show ip in ve <>" does not show "ip icmp address mask" enabled/disabled status.	
Condition: Execution of "show ip in ve <>" CLI.	
Workaround: None	

Defect ID: DEFECT000605923	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** FCoE - Fibre Channel over Ethernet
Reported In Release: NOS7.0.1	**Technology:** Layer 2 Switching
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.	

Known Issues for NOS 6.0.2b

This section lists known issues as of May 20, 2016 in Network OS 6.0.2b.

None

Known Issues for NOS 6.0.2a

This section lists known issues as of November 30, 2015 in Network OS 6.0.2a.

None

Known Issues for NOS 6.0.2

This section lists known issues as of November 30, 2015 in Network OS 6.0.2.

None

Known Issues for NOS 6.0.1

This section lists open software defects with Critical, High, and Medium Technical Severity as of November 30, 2015 in Network OS 6.0.1.

Defect ID: DEFECT000486078	
Technical Severity: Medium	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Network Automation and Orchestration
Reported In Release: NOS4.0.0	**Technology:** OpenStack Integration
Symptom: When packet capture is enabled & the rate of the packets coming to the CPU is higher than 30 pkts per second & terminal monitor is enabled, it may cause a switch to unexpectedly reboot.	
Condition: Occurs only when terminal monitor is enabled and either packet capture utility or debug is enabled for any component.	
Workaround: Disable terminal monitor on telnet sessions when PCAP is enabled.	

Defect ID: DEFECT000553915	
Technical Severity: High	**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: DEFECT000561651	
Technical Severity: Low	**Probability:** Low
Product: Extreme Network OS	**Technology Group:** Monitoring
Reported In Release: NOS6.0.0	**Technology:** RAS - Reliability, Availability, and Serviceability
Symptom: Spelling of the word 'display' in 'rasman' command' help text should be corrected.	
Condition: Help text when 'rasman' command is executed.	

Defect ID: DEFECT000562447	
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: Unexpected Reload due to inconsistent maintenance of internal data.	
Condition: Issuing the CLI "show running-config rbridge-id <id> snmp-server" when configurations contain local SNMPv3 Host	
Workaround: Global SNMPv3 Host can be used for querying instead of Local SNMPv3 Host.	

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: DEFECT000563295	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: ARP - Address Resolution Protocol
Symptom: With certain 3rd party devices, VDX learns ARP entries from a host belonging to a different subnet due to their GARP replies. This can potentially impact the traffic towards that source.	
Condition: When the device sends out GARP reply packet with source IP on different subnet than L3 interface IP.	

Defect ID: DEFECT000572168	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.2	Technology: IPv4 Multicast Routing
Symptom: If a switch has large number, greater than 500, multicast route entries and a disruptive event such as "chassis disable" or shut down of RPF interface occurs which triggers a full cleanup of multicast route table; post recovery a partial traffic loss can be observed.	
Condition: This issue is observed when a disruptive event such as "chassis disable" or "shutting down RPF interface" is executed on a switch with large number of multicast routes.	
Recovery: Reload the switch.	

Defect ID: DEFECT000578967	
Technical Severity: High	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: DEFECT000572746	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS5.0.1	Technology: Hardware Monitoring
Symptom: Customer may experience unexpected reload or some daemon termination.	
Condition: We don;t compress and delete confd core files till now and due to that our switch can experience disk full issue.	
Recovery: We have to manually remove confd core files from the disk.	

Defect ID: DEFECT000574606	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: LAG - Link Aggregation Group
Symptom: Customer can experience unexpected VDX reload.	
Condition: VDX experiences OOM condition when the switch receive multicast traffic and failed to free the received packet buffer due to any internal error condition.	
Workaround: If VDX has configured as an RP then please remove it as that is unsupported config.	

Defect ID: DEFECT000578967	
Technical Severity: High	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