

9035850

Network OS 7.0.1 for VDX

Release Notes v5.0

Copyright Statement

© 2018, Extreme Networks, Inc. All Rights Reserved.

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries. All other names are the property of their respective owners. For additional information on Extreme Networks Trademarks please see www.extremenetworks.com/company/legal/trademarks. Specifications and product availability are subject to change without notice.

Contents

Contacting Extreme Technical Support.....	Error! Bookmark not defined.
Extreme Customers.....	Error! Bookmark not defined.
Extreme OEM customers	Error! Bookmark not defined.
Related Documentation.....	Error! Bookmark not defined.
Document feedback.....	Error! Bookmark not defined.
Hardware	11
New devices	11
New interface modules.....	11
Deprecated Hardware.....	11
Software Features.....	11
New Software Features for Network OS v7.0.1.....	11
Modified Features.....	13
CLI Changes	13
New Commands for Network OS v7.0.1	13
Modified Commands for Network OS v7.0.1.....	15
Deprecated Commands for Network OS v7.0.1.....	16
API Changes.....	16
Newly supported standards and RFCs	16
Supported devices	17
Extreme VDX 6940-144S	17
Extreme VDX 6940-36Q	17
Extreme VDX 6740	18
Extreme VDX 6740T	18
Extreme VDX 6740T-1G.....	19
Extreme VDX 8770-4 and VDX 8770-8	19
Supported Blades for VDX 8770.....	19
Extreme VDX 2741	20
Extreme VDX 2746	20
Support for 100-Mb interfaces	21
Supported power supplies	21
Supported Optics for Network OS v7.0.1.....	24
Network OS v7.0.1 Release Notes v5.0	

Image filenames.....	27
Upgrade/Downgrade considerations.....	27
Migration Path	28
Management IP connectivity	29
Firmware Installation	30
Upgrading to this Release (Best Practices)	32
Downgrading to a Previous Release	32
Upgrade/downgrade Considerations for vLAG deployments.....	33
Upgrade/downgrade with default configuration	35
Management Gateway IP changes	35
Management Services.....	36
Telnet, SSH and AAA VRF support.....	36
CLI Changes for Telnet, SSH, AAA	36
HTTP VRF support	36
NTP VRF support	36
CLI Changes for NTP	37
SNMP- Community string maximum length increased to 64:	37
SNMP - Support for traps during hafailover:	37
SNMP-Trap Source IP support:	37
SNMP context based query:.....	37
SNMP MIB – VLAN update	37
SNMP Trap VRF Support	38
SNMP-Trap CLI	38
SNMP – IF MIB	38
Sflow VRF Support.....	38
Sflow-CLI	38
Syslog VRF Support	38
Syslog-CLI	38
Firmware download, Copy support, Copy config.....	39
Other Management Services	39
Scalability numbers.....	40
Compatibility and Interoperability.....	49

IP Storage	49
FC/FCoE Storage.....	49
Network Adapters.....	50
Limitations and Restrictions.....	52
Command Line Interface.....	52
Line cards	54
USB.....	54
Licensing.....	54
VCS	55
Logical Chassis.....	56
Extreme Trunks	56
Breakout Interfaces.....	57
Dual-personality Ports	57
1G Mode	57
vLAG	58
Virtual IP Address Support.....	58
Security, Management ACLs, Authentication, Authorization	58
SPAN & RSPAN	59
MAC Learning Considerations in VCS.....	59
PVLAN.....	60
UDLD	60
STP/DiST	60
Edge Loop Detection (ELD)	61
Long Distance ISL Ports	61
AMPP and Port-Profiles	62
vCenter.....	63
QoS.....	63
FCoE	64
FlexPorts.....	65
Fibre Channel	65
Access Gateway	65
IP Fabric.....	66

ND/RA	67
BFD	67
VRRP	68
OSPFv3	69
BGP.....	69
Layer 2/Layer 3 Multicast	69
VRF	69
BGP-VRF	70
ACL	70
Policy-based Routing (PBR).....	70
Inter-VRF Leaking (Static).....	70
DHCP IP Helper.....	71
Dynamic ARP Inspection (DAI)	71
DHCP-based Firmware Download (DAD – DHCP Automatic Deployment).....	71
Link State Tracking	72
OpenFlow	72
Layer 2 and Layer 3 ISSU on VDX 6740x.....	73
Extreme Vyatta Controller (BVC)	73
Auto QoS for NAS.....	74
REST API	74
NetConf.....	75
VXLAN Gateway for VMware NSX.....	75
VF Extension using VxLAN	76
TCAM Profiles.....	76
Management VRF.....	76
Conversational MAC Learning.....	76
System level Flowbased QoS	77
Port level Flowbased QoS	77
Non-trivial Merge.....	77
HA on TOR switches	77
Logical Chassis HA	78
Interoperability	78

MAPS.....	79
Maintenance Mode.....	79
BNA	79
Miscellaneous	79
Defects	80
TSBs - Critical Issues to Consider Prior to Installing This NOS Release	80
TSB Issues Resolved in Network OS v7.0.1	81
Network OS v7.0.1 Caveats.....	81
Closed with code changes for NOS v7.0.1	85
Closed with code changes for NOS v7.0.0	100
Closed without code changes for Network OS v7.0.1	155
Closed without code changes for Network OS v7.0.0	158
Known Issues for Network OS v7.0.1.....	174
Known Issues for Network OS v7.0.0.....	180

DOCUMENT HISTORY

Version	Summary of Changes	Publication Date
1.0	Initial Release	May 25th, 2016
2.0	Limitations updated	June 15 th , 2016
3.0	NSX-V certification added	August 4th, 2016
4.0	Defect tables updated	August 24 th , 2016
5.0	Update on upgrading nodes (best practices) for the release revision Added document number	November 13 th , 2018

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 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

Extreme Network OS 7.0.1 release introduces incremental capabilities to further enable the Layer 3 Extreme® IP Fabrics, Layer 2 VCS® Data Center Fabrics, Standards (BGP-EPVN) based network Virtualization for Intra DC and DC-Interconnect solutions.

Hardware

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

New devices

None

New interface modules

None

Deprecated Hardware

None

Software Features

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

New Software Features for Network OS v7.0.1

The following software features are new in this release:

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

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

Two Factor Authentication

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

Duration login

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

IP Fabric enhancements:

Decouple IPv4 and eVPN address-family

To have separate BGP peering session between leaf and spine for underlay and overlay, to decouple overlay neighbor-ship errors from the underlay session and so the routes are not compromised.

Layer 3 VNI IMR routes

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

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

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

Peer disable AS-check

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

Hardware TCAM profiles

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

MAPS policy reapply

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

MAC-move-detect feature

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

LLDP Enhancement

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

Modified Features

The following software features have been modified in this release:

- IP Fabrics updates

See overview in the IP Fabrics Configuration Guide.

- Layer 2 MTU and IP MTU

The IP MTU, IPv6 MTU, and MTU commands now operate as global commands.

Deprecated Software Features

None

CLI Changes

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

New Commands for Network OS v7.0.1

The following configuration commands are new in this release:

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

```
clear bgp evpn l2route type inclusive-multicast
```

```
maps reapply-policy
```

```
mac-address-table mac-move detect
```

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

```
area <area_id> nssa no-redistribution
```

```
area <area_id> nssa translator-always
```

```
area <area_id> nssa translator-interval <interval_value>
```

```
snmp-server three-tuple-if enable
```

```
spanning-tree ieee-bpdu limit-vlan-flood
```

The following show commands are new in this release:

ARP:

```
show system internal arp clientlist structures
show system internal arp counters
show system internal arp interface ve 100
show system internal arp ipv4 vrf all
show system internal arp ipv6 vrf all
show system internal arp l2 clientdb
show system internal arp memstats
show system internal arp rib memstats
show system internal arp rib routes
show system internal arp static config interface ve 100
show system internal arp static config vrf all
show system internal arp summary
show system internal arp vrf all
```

RIB:

```
show system internal rib clients
show system internal rib ip adj vrf default-vrf
show system internal rib ip route vrf default-vrf
show system internal rib ipv6 adj vrf default-vrf
show system internal rib ipv6 route vrf default-vrf
show system internal ribm memory-stats
show system internal ribm vrf
```

FIB-ARP:

```
show system internal fib 0 arp arp-bum-stats
```

```
show system internal fib 0 arp counters
show system internal fib 0 arp dai-filter
show system internal fib 0 arp dai-list
show system internal fib 0 arp dai-stats
show system internal fib 0 arp dai-trust
show system internal fib 0 arp dai-usr-acl
show system internal fib 0 arp interface ve 100
show system internal fib 0 arp ipv4 vrf all
show system internal fib 0 arp ipv6 vrf all
show system internal fib 0 arp memstats
show system internal fib 0 arp nd-bum-stats
show system internal fib 0 arp summary
show system internal fib 0 arp vrf all
```

FIB-RIB:

```
show system internal fib 0 rib clients
show system internal fib 0 rib ip adj vrf default-vrf
show system internal fib 0 rib ip route vrf default-vrf
show system internal fib 0 rib ipv6 adj vrf default-vrf
show system internal fib 0 rib ipv6 route vrf default-vrf
show system internal fib 0 rib memory-stats
```

Modified Commands for Network OS v7.0.1

The following commands have been modified for this release:

```
duplicate-mac-timer <interval> max-count <count>
username <name> access-time HHMM to HHMM
ip mtu
```

ipv6 mtu

mtu

password-attributes

Deprecated Commands for Network OS v7.0.1

no nssa-translator

system tunnel replicator load-balance

API Changes

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

Newly supported standards and RFCs

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

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

None

HARDWARE SUPPORT

Supported devices

Extreme NOS v7.0.1 supports following VDX Switches:

- Extreme VDX 6940-144S
- Extreme VDX 6940-36Q
- Extreme VDX 6740
- Extreme VDX 6740T
- Extreme VDX 6740T-1G
- Extreme VDX 8770-4
- Extreme VDX 8770-8
- Extreme VDX 2741
- Extreme VDX 2746

Extreme VDX 6940-144S

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

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

Extreme VDX 6940-36Q

The Extreme 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.

Extreme VDX 6740

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

These ports support the following:

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

Extreme VDX 6740T

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

- Available in 24, 48 and 64 port SKU.
- 3 microsecond latency for any port to port to assure rapid response for latency-sensitive applications.
- The base SKU is available with 24 10GbE Base-T 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 a FC breakout of 4*8G or 4*16G. These ports can be used to connect to the FOS switches.
- Each 40GbE port is also capable of doing an FC breakout of 4 x 8G/16G.
- Additional 4X40GbE ports can be added to base version with 2X40GbE POD license increments.
- 100Mb Support – Refer to “Support for 100-Mb interfaces” below.

Extreme VDX 6740T-1G

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

- Base version is available with 48 x 1000BASE-T ports and 2 x 40 GbE QSFP+ ports.
- 3-microsecond latency for any port to port to assure rapid response for latency-sensitive applications.
- All 48 x 1000BASE-T ports can be upgraded to 10Gbase-T port with capacity on demand license.
- Additional 2X40Gbe ports can be added to base version with 2X40Gbe 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.

Extreme VDX 8770-4 and VDX 8770-8

The Extreme 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 Extreme 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 Extreme Fabric Multipathing technology, enabling the switch to serve extremely large-scale deployments with the best-possible network utilization.

Supported Blades for VDX 8770

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

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

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

Extreme VDX 2741

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

Extreme VDX 2746

- Extreme 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.

Supported power supplies

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

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

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

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

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

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

- XBR-250WPSAC-F - FRU 250 W AC power supply/fan, non-port-side exhaust airflow
- XBR-250WPSAC-R - FRU 250 W AC power supply/fan, port-side 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 v7.0.1

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

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

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

Speed	FRU and Optics SKU	Description	Part Number
	40G-QSFP-SR4-INT	40 GbE SR4 (4x10 GbE SFP break-out capable) Breakout optical cable is not included with this optics	57-1000129-01
	40G-QSFP-SR-BIDI	40 GbE QSFP+ Bi-Directional 100m optics	57-1000339-01
	40G-QSFP-ESR4	40GBase-eSR4 QSFP+ optic (MTP 1x12) 300m over MMF, (10GBASE-SR compatible, breakout), 1-pack	57-1000296-01
	40G-QSFP-ER4	40 GbE 40Km optic	57-1000327-01
	40G-QSFP-LR4	40 GbE 10Km optic	57-1000263-01
	40G-QSFP-LM4	40 GbE 140m multi-mode or 2km single-mode optic	57-1000325-01
	40G-QSFP-QSFP-AOC-1001	40GE Direct Attached QSFP+ to QSFP+ Active Optical Cable, 10m, 1-pack	57-1000306-01
	40G-QSFP-4SFP-AOC-1001	4x10GE Direct Attached QSFP+ to 4 SFP+ Active Optical Breakout Cable, 10m, 1-pack	57-1000307-01
8G FC	XBR-000163 (1-pack)	8G FC SWL	
	XBR-000164 (8-pack)		
	XBR-000153 (1-pack)	8G FC LWL	
	XBR-000172 (8-pack)		
16G FC	XBR-000174	8G FC ELWL	
	XBR-000192 (1-pack)	16G FC SWL	
	XBR-000193 (8-pack)		
	XBR-000198 (1-pack)	16G FC LWL	
FC QSFP	XBR-000199 (8-pack)		
	XBR-000245	4x8G or 4x16G FC QSFP breakout. VDX 6740T, 6740T-1G, 2740 and 2746 only (not applicable for VDX 6740).	
100GbE	100G-CFP2-SR10 (1-pack)	100 GbE CFP2 optic, SR10, for distances up to 100 m over MMF	57-1000284-01
	100G-CFP2-LR4-10KM	100 GbE CFP2 optic, LR4, for distances up to 10 km over SMF	57-1000285-01
	100G-CFP2-ER4-40KM	100 GbE CFP2 optic, ER4, for distances up to 40 km over SMF	57-1000328-01
	100G-QSFP28-SR4	100 GbE SR4 QSFP28 optic for distances up to 100m over MMF. Supported on VDX6940-144S and VDX8770-4/8 platforms.	57-1000326-01
	100G-QSFP28-LR4L-2KM	100 GbE QSFP28 optic for distances up to 2 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms.	57-1000329-01
	100G-QSFP28-LR4-10KM	100 GbE QSFP28 optic for distances up to 10 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms.	57-1000334-01

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

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

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

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

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

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

Mellanox MAM1Q00A-QSA	Quad to Serial Small Form Factor Pluggable Adapter which can be used with following 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)
CFP2 to QSFP28 conversion module (PN: 80-1008646-01)	CFP2 to QSFP28 conversion module connects the QSFP28 optic (100G optic) in a CFP2 capable port of 2/6x100G line cards in VDX8770-4/8 chassis.

*Note: Legacy Foundry Networks branded optics are not supported

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

SOFTWARE UPGRADE AND DOWNGRADE

Image filenames

Download the following images from www.Extremenetworks.com

Image Filename	Description	Supported Device or Module
nos7.0.1.tar.gz	Network OS v7.0.1 for unix	NA
nos7.0.1.zip	Network OS v7.0.1 for Windows	NA
nos7.0.1_all_mibs.tar.gz	Network OS v7.0.1 MIBS	NA
	Network OS v7.0.1 Release	NA
nos7.0.1_releasenotes_v4.0	Notes v4.0 (PDF)	
nos7.0.1.md5	Network OS v7.0.1 MD5 Checksum	NA

Upgrade/Downgrade considerations

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

Migration Path

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

To / From	5.0.2a	6.0.1a	6.0.1a2	6.0.2	7.0.0	7.0.1
5.0.2a	NA	coldboot	coldboot	coldboot	coldboot default-config	coldboot default-config
6.0.1a	coldboot	NA	ISSU	ISSU	coldboot	coldboot
6.0.1a2	coldboot	coldboot	NA	ISSU	coldboot	coldboot
6.0.2	coldboot	coldboot	coldboot	NA	coldboot	coldboot
7.0.0	coldboot default-config	coldboot	coldboot	coldboot	NAr downgrade	ISSU*
7.0.1	coldboot default-config	coldboot	coldboot	coldboot	coldboot**	NA

NOTES

1. *ISSU with only non BGP-eVPN IP Fabric configurations are supported between NOS7.0.0 to NOS7.0.1.
2. ** CFP2 to QSFP28 conversion module (PN: 80-1008646-01) Version3 downgrade to any release prior to NOS7.0.1 will cause CRC errors on the link.
3. Only Extreme Network Advisor (BNA) v14.0.1 (available separately) supports NOS v7.0.1. It is required to first upgrade to BNA v14.0.1 and then upgrade switches to Network OS v7.0.1.
4. 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.
5. 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.

6. You must remove the IGMP snooping static mrouter configuration from all VLANs before upgrading or downgrading from or to the Network OS 6.0.2x release.
7. ****Limitations:**
 - a) After downgrading from Network OS v7.0.1 to Network OS v5.0.x FCoE devices may not log back in or FCoE configuration may be lost. To recover, reload the switch. Alternate recovery method: re-configure FCoE by removing and adding fcoeport configuration (no fcoeport/fcoeport default) on the affected interfaces.
 - b) In rare occurrence, 40G links may not come up online after upgrade to 7.0.1, need to do shut/no shut to recover
 - c) In VDX 8770 platforms, After upgrade form 6.0.2 to 7.0.1 with coldboot, SNMP V3 traps are not received for the V3host which is under Rbridge.
 - d) Dport test between VDX 6740T and VDX 6940-144S breakout link may fail in upgrade to 7.0.1

Management IP connectivity

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

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

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

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

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

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

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

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

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

Firmware Installation

In fabric cluster mode

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

In logical chassis cluster mode

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

General information on installing Extreme Network OS can be found in the Extreme *Network OS Administrator's Guide*. This section includes special considerations and caveats to be aware of when upgrading to or from this version of Extreme Network OS, as well as recommended migration paths to use to reach this version of Extreme Network OS.

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

Upgrading to this Release (Best Practices)

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

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

- Download the firmware on all the switches running Network OS v6.0.2 using the coldboot option.
- After all switches complete the firmware download, they will be automatically rebooted.
- Since all nodes reboot at the same time, this procedure is service disruptive.

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

- This is the most recommended procedure for lossless upgrade. This requires servers to be dual homed.
- Download the firmware in all the odd nodes running Network OS with the coldboot option.
- After these switches complete the firmware download, they will be rebooted automatically. After they boot up, half of the cluster is now on the latest version. Traffic resumes and passes through the other half of the cluster.
- Now download the firmware in all even nodes with the 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 Extreme NOS 6.0.x using the coldboot option. Principal node in a cluster should be last to be upgraded.
- After a node is upgraded, it will join the existing Network OS v6.0.2 cluster and form Fabric cluster. Eventually, when all the nodes are upgraded, they will form one Network OS 7.0.1 VCS Cluster. The data path will remain intact in Fabric cluster. [Note that no configuration changes are allowed during this time.]

Downgrading to a Previous Release

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

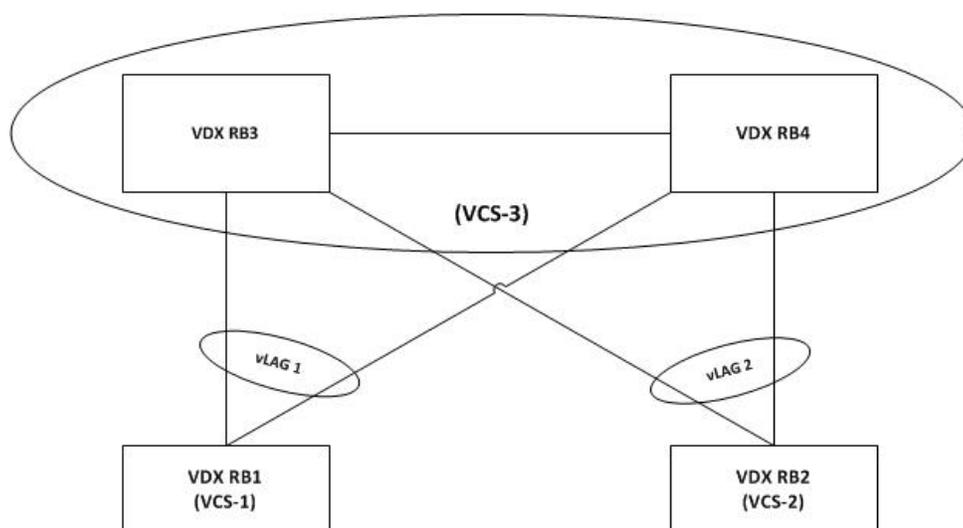
- Execute a coldboot downgrade. In FC mode, the startup-config file will be re-applied automatically. In LC mode, copy the 6.0.1x config back by executing “copy file running”.

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

Telnet, SSH and AAA VRF support

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

CLI Changes for Telnet, SSH, AAA

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

```
[no] ssh server use-vrf <vrf-name> [shutdown]
[no] telnet server use-vrf <vrf-name> [shutdown]
[no] ldap-server host <IPv4|IPv6|hostname> [use-vrf <VRF name>]
[no] tacacs-server host < IPv4|IPv6|hostname > [use-vrf <VRF name>]
[no] radius-server host < IPv4|IPv6|hostname > [use-vrf <VRF name>]
```

HTTP VRF support

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

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

NTP VRF support

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

CLI Changes for NTP

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

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

SNMP- Community string maximum length increased to 64:

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

SNMP - Support for traps during hafailover:

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

SNMP-Trap Source IP support:

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

```
[no] snmp-server host ip-address <community-string> source-interface {  
loopback number|ve vlan_id}}
```

```
[no] snmp-server v3host ip-address <username> source-interface {  
loopback number|ve vlan_id}}
```

SNMP context based query:

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

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

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

SNMP MIB – VLAN update

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

SNMP Trap VRF Support

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

SNMP-Trap CLI

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

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

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

To disable per link TRAP under interface

```
[No] snmp trap link-status
```

SNMP – IF MIB

To display Interface details when linecard is powered-off

```
[No] snmp-server offline-if enable
```

Sflow VRF Support

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

Sflow-CLI

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

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

Syslog VRF Support

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

Syslog-CLI

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

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

Firmware download, Copy support, Copy config

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

Other Management Services

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

SCALABILITY AND INTEROPERABILITY

Scalability numbers

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

NOS v7.0.1 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)

NOS v7.0.1 Scalability Numbers	VDX 6740, 6740T, 6740T- 1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
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
Maximum # of ELD instances in the fabric	2000	2000	2000	2000
Maximum # of IGMP Snooping Interfaces supported	512	512	512	512
Learning rate for IGMP snooping (groups/second)	512	512	512	512
Maximum # of L2 (IGMP Snooping) multicast groups	6000	6000	6000	6000
Maximum # of MLD Interfaces	256	256	256	256
Maximum # of MLD Groups	4000	4000	4000	4000
Learning rate for MLD snooping (groups/second)	512	512	512	512
# of L3 (S,G) forwarding Entries	2000	2000	2000	2000
# of L3 (*,G) forwarding Entries	256	256	256	256
PIM Interfaces Supported	32	32	32	32
IGMP interfaces supported	32	32	32	32
Learning Rate for PIM-SM (flows/second)	32	32	32	32
Maximum # of L2 ACL(ingress/egress) *	3000/120	12000/2000	6128/496	6128/496
Maximum # of L3 ACL ipv4 (ingress/egress) *	1500/120	12000/2000	3064/475	3064/475
Maximum # of class-maps	2048	2048	2048	2048
Maximum # of policy-maps	2048	2048	2048	2048
Maximum # of class-maps per policy map	50	50	50	50
Maximum Total # of L3 ACL ipv6 (ingress/egress) *	500/120	4000/2000	1000/500	1000/500
Maximum # of VF/FCoE interfaces/Logins (Per switch)	1000	1000	1000	1000
Maximum # of Enodes/FCoE Devices per Fabric	2000	2000	2000	2000

NOS v7.0.1 Scalability Numbers	VDX 6740, 6740T, 6740T- 1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
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)	64	288	144	144
Maximum # of members in a standard LAG	16	16	16	16
Maximum # of members in a Extreme Trunk (10G)	16	8	12	12
Maximum # of members in a Extreme Trunk (40G)	2	NA	3	3
Maximum # of members in a Extreme Trunk (100G)	NA	NA	NA	NA
Maximum # of switches in a Fabric cluster mode **	48	48	48	48
Maximum # of switches in Logical cluster mode **	48	48	48	48
Maximum # of L2 ECMP Paths	16	8	16	16
Maximum # of vLAGs in a fabric	2000	2000	2000	2000
Maximum # of member ports in a vLAG	64	64	64	64
Maximum # of nodes in a vLAG	8	8	8	8
Maximum # of member ports per vLAG per Node	16	16	16	16
Maximum # of Management ACL	256	256	256	256
Maximum # of ARP Entries *	16000	126000	72000	72000
Maximum # of OSPF areas	20	64	20	20
Maximum # of OSPF routers in a single area	64	200	64	64
Maximum # of OSPF adjacencies	100	200	100	100
Maximum # of OSPF routes *	8,000	64,000	10000	10000
# of OSPF Interfaces	100	1,000	100	100
# of OSPF enabled subnets	100	1,000	100	100
# of local subnets in a single area	100	1,000	100	100
Maximum # of OSPFv3 areas	9	9	9	9
Maximum # of OSPFv3 routers in a single area	64	200	64	64
Maximum # of OSPFv3 adjacencies	100	200	100	100
Maximum # of OSPFv3 routes *	1500	64000	1500	1500
# of OSPFv3 Interfaces	100	256	100	100
# of OSPFv3 enabled subnets	100	256	100	100
Maximum # of IPv4 routes in SW *	8000	280000	10000	10000
Maximum # of IPv6 routes in SW *	1500	64000	1500	1500

NOS v7.0.1 Scalability Numbers	VDX 6740, 6740T, 6740T- 1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
Maximum # of IPv4 static routes *	2000	40,000	2000	2000
Maximum # of IPv6 static routes *	500	20,000	500	500
Maximum # of VRRP instances per system	255	1024	512	512
Maximum # of VRRP v3 instances per system	255	1024	512	512
Maximum # of VRRP instances per interface	32	32	32	32
Maximum # of routers participating in a VRRP-E session	8	8	8	8
Maximum # of virtual IP addresses per VRRP instance	16	16	16	16
Maximum # of FVG instances per system	256	4096	1024	1024
Maximum # of FVG instances per interface	1	1	1	1
Maximum # of routers participating in a FVG session	32	32	32	32
Maximum # of Gateway IP addresses per FVG instance	1	1	1	1
Maximum # of IPv4 routes with ECMP supported *	8000	200000	10000	10000
Maximum # of IPv6 routes with ECMP supported *	1500	64000	1500	1500
Maximum # of L3 ECMP	16	32	32	32
Maximum # of IPv4 interfaces per system *(Ve intf)	2000	4000	2000	2000
Maximum # of IPv6 interfaces per system *(Ve intf)	512	4000	512	512
Maximum # of VRF per node	32	512	512	512
Maximum # of VRFs support protocols per node	32	128	128	128
Maximum # of I-BGP peers	256	512	256	256
Maximum # of E-BGP peers	256	256	256	256
Maximum # of IPv4 BGP routes in HW *	8000	200000	10000	10000
Maximum # of IPv6 BGP routes in HW *	1,500	64000	1500	1500
Maximum # of IPv4 RIB (IN + OUT) Routes *	110000	1300000	110000	110000
Maximum # of IPv6 RIB (IN + OUT) Routes *	110000	1300000	110000	110000
Maximum # BGP IPv4/IPv6 Peer Group	100	250	100	100
Maximum # of BFD sessions per node	100	100	100	100
Maximum # of UDLD enabled interfaces	64	384	144	108
Maximum # of PVLAN domain supported	1000	1000	1000	1000
Maximum # of Secondary vlans per PVLAN supported	24	24	24	24
Maximum # of primary vlans per PVLAN supported in promiscuous mode	24	24	24	24
DHCP IP Helper Addresses per interface	16	16	16	16

NOS v7.0.1 Scalability Numbers	VDX 6740, 6740T, 6740T- 1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
DHCP IP Helper VE interfaces	256	1,000	256	256
DHCP IP Helper physical ports	60	384	60	60
DHCP IP Relay Addresses per Node	2000	4000	2000	2000
DHCP IPv6 Relay Address per Node	2000	4000	2000	2000
Max Number of configurable PBR route maps	64	64	64	64
Max Number of configurable PBR stanzas	1024	1024	1024	1024
Max Number of HW entries available for PBR	512	8192	512	512
Max Number of configurable next hops within a single PBR stanza	128	128	128	128
Max # of OpenFlow Active Connections	1	1	1	1
Max # of OpenFlow Passive Connections	1	1	1	1
Maximum # of OpenFlow L2 flows	1000	4000	879	879
Maximum # of OpenFlow L3 flows	1000	4000	879	879
Maximum # of Total OpenFlow GROUP	768	768	768	768
Maximun # of OpenFlow GROUP Type ALL	256	256	256	256
Maximun # of OpenFlow GROUP Type SELECT	256	256	256	256
Maximun # of OpenFlow GROUP Type INDIRECT	256	256	256	256
Max # of Buckets per GROUP ALL	16	16	16	16
Max # of Buckets per GROUP SELECT	8	8	8	8
Max # of Buckets per GROUP INDIRECT	1	1	1	1
Max # of ACTIONS per Bucket	3	3	3	3
Max # METERS	1024	4096	1024	1024
Maximum # of MAPS policy	10	10	10	10
Maximum # of MAPS rules	250	250	250	250
Maximum # of MAPS groups	64	64	64	64



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

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

IP Fabric Scalability:

IP Fabric Scalability Numbers	VDX- 8770	VDX-6940		VDX 6940-144s		VDX 6740, VDX 6740T
	Spine	Spine	Leaf	Spine	Leaf	Leaf
VLANS extended with VxLANs (no. of tunnels * VLANs * ECMP)	NA	NA	16k	NA	16k	16k
Software MAC entries (CML)	NA	200k	200k	200k	200k	200k

Software ARP entries (Conversational ARP)	NA	100k	100k	100k	100k	100k
Software ND entries (Conversational-ND)	NA	50k	50k	50k	50k	50k
BGP eVPN IPv4 routes	200k	200k	200k	200k	200k	200k
BGP eVPN IPv6 routes	64k	2k	2k	2k	2k	2k
BGP eVPN MAC-IP routes	100k	100k	100k	100k	100k	100k
BGP eVPN MAC routes	200k	200k	200k	200k	200k	200k
Max # of IP Unnumbered interface	384	36	36	144	144	52
Max # of IP Port channel interface	384	36	36	144	144	52
Max # of members per IP Port-Channel Interface	8	8	8	8	8	8
Max # of Leaf – Spine ECMP	16	16	16	16	16	16
Max # of SAG addresses per interface	32	32	32	32	32	32

HW Profile and Platform Specific Scale Numbers

Route Profile Scale:

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

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

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

Maximum # of IPv4 BGP routes in HW *	65000	280000	198000	163000	20000	12000
Maximum # of IPv6 BGP routes in HW *	16000	2000	2000	8000	64000	12000
Maximum # of IPv4 routes in SW *	65000	280000	198000	163000	20000	12000
Maximum # of IPv6 routes in SW *	16000	2000	2000	8000	64000	12000
Maximum # of ARP Entries *	98000	40000	129000	98000	12000	20000
Maximum # of IPv6 neighbor cache Entries *	28000	2000	2000	12000	12000	65000

L2 L3 Multicast Scale :

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

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

ACL Scale:

VDX8770-4									
NOS v7.0.0 Scalability Numbers	TCAM PROFILES								
	DEFAULT	DNY-ARP-INS	IPV4-ACL	IPV4-V6-MCAST	IPV4-V6-PBR	IPV4-V6-QOS	L2-ACL-QOS	L2-IPV4-ACL	OPEN FLOW
Maximum # of L2 ACL(ingress/egress) *	16000/2000	12000/2000	512/1016	500/1000	500/1000	500/1000	32000/2000	16000/2000	12000/2000

Maximum # of L3 ACL ipv4 (ingress/egress) *	16000/2000	16000/2000	51000/2000	500/2000	8000/2000	8000/2000	5000/2000	24500/2000	12000/2000
Maximum # of L3 ACL ipv6 (ingress/egress) *	500/2000	500/2000	0/2000	500/2000	4000/2000	4000/2000	0/1000	0/2000	500/2000

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

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

Compatibility and Interoperability

The following tables list the devices tested for IP storage, FC/FCoE storage and host adapters for VDX as of Network OS v7.0.1. This is a representative list of devices, Network OS v7.0.1 supports all standards-based devices connected to it for these types of storage.

IP Storage

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



FC/FCoE Storage

Vendor	Storage Array Model	Protocol	Switch Model	Initiator
Netapp	FAS3250-cdot	FC, FCoE	6740, 8770 (FCoE only)	Windows 2012, VMWare
HDS	R800	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	R700	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	HUSVM	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	DF850	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
	DF800	FC	6740	RHEL 6.5, 7.0, SLES 11sp3, 12, Windows 2012 R2, Windows 2008 R2 SP1
EMC	CX4-120	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2
	VMAX 40K	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2, Windows 2012
	VNX-5300	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2

Vendor	Storage Array Model	Protocol	Switch Model	Initiator
	VNX-5500	FC, FCoE	6740, 8770 (FCoE only)	RHEL 6.3, 6.5, Solaris 10, Windows 2008, Windows 2008 R2, Windows 2012
	VSP	FC, FCoE	6740	RHEL 6.5, Windows 2012
IBM	DS8100	FC	6740/T	Windows 2012 R2
	Flash 840	FC	6740/T	Windows 2012 R2
	XIV	FC	6740/T	Windows 2012 R2
HP	MSA2040	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P10000	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P6500	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P6300	FC, FCoE	6740, 8770 (FCoE only)	RHEL 7.0, Windows 2012, Windows 2012 R2
	P4330	FC	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2
	P9500	FC, FCoE	6740/T	RHEL 7.0, Windows 2012, Windows 2012 R2

Network Adapters

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

Vendor	Model	Protocol	Switch Model	OS
	2672	FC	6740	RHEL 6.5, Windows 2008
	8152	FCoE	6740	ESX 5.1
	8142	FCoE	6740	Windows 2012, RHEL 6.5
	2672	FC	6740	RHEL 6.5
	2762	FC	6740	RHEL 5.1, Windows 2012

ADDITIONAL CONSIDERATIONS

Limitations and Restrictions

Command Line Interface

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

- Show notification stream –, for debugging purpose only.
 - 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.
- 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.
- 1G copper SFPs do not support exchanging flow-control settings during the auto-negotiation process. It is recommended to configure static mode of configuration of flow-control on both the ends of the desired link.
- 1G Optical ports should use the same speed config (speed auto or speed 1000) on both sides of the link for a proper link up.
- The VDX6940-36Q and VDX6940-144S requires 40 seconds between the removal and insertion of the 100G QSFP28 optics in order to establish a stable 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 6940x is supported only in auto mode.
- Multiple OIR (Online insertion and removal) of 40G LR optics when connected to ICX/FCX may cause link to remain down. Performing “shutdown” followed by “no shutdown” of the interface will recover the link.
- VDX 6740/6740T/6740T-1G/6940 platforms do not support IP fragmentation. MTU errors are reported in “show interface” as “Errors” under the “Transmit Statistics”.
- When a switch fan or PSU is removed or is faulty, switch status LED will blink green on VDX6940-144S and amber-green on VDX6940-36Q and VDX6740.
- For 6940 platform family, if all ports in a given trunk-group are used as ISLs, it is recommended to configure only 1 lossless priority on the switch.
- Logical SAN is not supported in fabric cluster mode.

Line cards

- The VDX 8770 supports following line-cards only on Network OS v4.1.2 and above:
 - LC48×10G
 - LC12×40G
 - LC48×10GT
 - LC27×40G
 - LC6×100G
- It is required to upgrade the chassis to the line-card’s supported Network OS version before plugging the line-card into the chassis.
- If there exists a configuration for a line-card on the slot of VDX8770, before inserting a new line-card of other type in the same slot, it is required to remove the configuration of the old line-card from that slot. The “no line-card” command should be used to remove the old line-card configuration from the slot where the new line-card is to be inserted. The new line card may be faulted with appropriate code if the new line-card is plugged into the slot which has configuration of a line card of other type.

USB

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

Licensing

- On VDX platforms that have Flexport FC capable interfaces, enabling FibreChannel ports requires only the FCoE license to be installed and does not require any Port Upgrade license. The Port Upgrade license only controls Ethernet ports (number of ports or speed supported).
- An Integrated Routing license is NOT required on FOS-based SAN platforms running FOS 7.0.1 or above for FCR interoperability connectivity with VCS fabrics and the VDX6740x. Please refer to the FOS v7.0.1 Admin Guide documentation on configuring FOS platforms for connectivity to VDX 674x switches and VCS fabrics.

- The Layer 3 license is required on VDX8770 switches to enable Layer 3 feature set including OSPF, VRRP, BGP, VRF etc. A separate Layer 3 license is not required on VDX fixed-form factor switches as Layer 3 features are included in the default license.
- The Advanced Services License provides a single upgrade option to enable Fibre Channel over Ethernet (FCoE) and Layer 3 features on VDX8770 switches.

VCS

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

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

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.
- 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.
- 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.
- Config snapshot cannot be restored on pizza box platform when SW1 is active.
- 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 principal switch-overs are not supported and may segment the cluster.
- Disabling virtual-fabric may take up to 10 minutes depending on the number of ISLs and VLAN interfaces configured in the VCS.



Extreme Trunks

- The VDX 6740, VDX 6740T, VDX 2741 Extreme trunk (BTRUNK) can support up to 16 member links with a maximum throughput of 160G using 16x10G ports in the same trunk group. On these platforms traffic may not be distributed evenly across all member of a trunk at lower traffic rates.
- The VDX 6740, VDX 6740T, VDX 2741, VDX 2746 and VDX 6740T-1G Extreme trunk (BTRUNK) can support up to 2x40G member links in the same trunk group for a maximum throughput of 80G. The VDX 2746 can support 4x40G.
- The VDX 8770 Extreme trunk (BTRUNK) can support up to 8 member links with a maximum throughput of 80G using 8x10G ports in the same trunk group. Full link utilization of 8 ports in a trunk group is achievable with larger packet size (>128 Bytes).
- In the VDX 6940-36Q and VDX 6940-144s, only 63 port-channels are supported including LACP and Extreme PO.
- The VDX 6940-36Q Extreme trunk (BTRUNK) can support up to a maximum 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 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 10G breakout ports and vice versa.
- VDX 6940-36 and 6940-144S supports only static breakout of 40G ports. It is required to reboot the switch for the 40G ports on it to be converted into 10G breakout ports
- For VDX 6740, 6740T, 2741 and 6740T-1G platforms, the LED state for a breakout interface is deterministic. For all other supported platforms, the LED state for a breakout interface is non-deterministic.
- In breakout mode, the ‘show media’ CLI will display the same media information for all breakout interfaces, except for temperature, Tx voltage, Tx bias current and Rx power. These parameters would be displayed on per line basis. The TX Power Field in the show media command is not supported by the 40G optics.
- On 40G native mode - Breakout configuration is not blocked. If configured on one side, other side of link won’t able be identify peer port config is breakout and link won’t be stable.
- On VDX 6740T/6740T-1G/2746, the breakout ports are FlexPort capable, and may be configured to connect to FC switches with 4x16G breakout supported cables and optics.
- On VDX6940-144S, breakout connection using non-breakout cable is not supported.

Dual-personality Ports

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

1G Mode

- RMON stats are calculated incorrectly for packet sizes 64-127 bytes.
- 1G ports cannot form ISL links. Only 10G ports can be used to form 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.

- For 1G optics used in VDX6740 and VDX6940-144S, port speed should be set to Auto on both sides. If one side is speed 1000 and other side is Auto, link may not come online.

vLAG

- LAGs are created with default speed of 10G. Therefore Extreme recommends end user to set required speed manually based on member speed using “speed” command.
- When configuring LACP LAG between VDX and non-Extreme switches it is highly recommended to enable the vLAG ignore-split on the VDX . Ignore split option 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):
 - 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.
 - Behavior of “local” option in pre-4.1.0 releases is changed to the “local-auth-fallback” option.
 - When login authentication configuration is modified, the user sessions are not logged out. All connected user sessions can be explicitly logged out using “clear sessions” CLI.
- ACLs are not supported for egress traffic flows on 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”. The use-vrf option should be used to enter any additional parameters such as retries, timeout or key.
- Same NTP server configuration with different vrf not supported.
- There is a possibility that locked user accounts will get unlocked after a reboot if the running-config (before reboot) is different from startup-config of user accounts.
- Encrypted text (taken from running-config of any user account password with encryption turned on) should not be used as input for clear-text password for the same user. This may result in login failure of the user subsequently.
- When the ACL is applied to a management interface, only the top 256 rules will be applied if the ACL contains more than 256 rules.
- It is advised to not to apply ACL with 12k rules to management interface.
- When more than 250 rules ACL’s are configured (over supported scale), they may be partially installed & effective.
- Access to ONLY the following Active Directory (AD) servers is supported by Extreme LDAP client:
 - Windows 2000
 - Windows 2003
 - Windows 2008 AD
- IPv6 RA Guard feature is not supported on VDX 8770 although the CLIs are visible.

SPAN & RSPAN

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

MAC Learning Considerations in VCS

- Under rare circumstances, end user might see mac address sync up issues on few nodes of a cluster (where 1 or more MAC addresses might be missing in some nodes). Extreme recommends to do “clear mac-address-table dynamic” in such cases.
- Static mac addresses will be displayed even when interfaces are down. This may cause blackholing of the traffic.

- Under certain conditions, MAC addresses may not be learnt even though ARP's may be learnt for those same MAC addresses.
- Under certain conditions, multicast traffic destined for static multicast address will flood on to other VLANs.

PVLAN

- Following PVLAN features are not supported:
 - IGMP on PVLANs but there is no error message displayed if operator configures IGMP snooping on PVLAN
 - ARP & Routing in PVLAN domain
 - Enabling Routing in Primary and Secondary Vlan.
 - CLI to enable Local Proxy ARP on primary VLAN.
 - IP Configuration on PVLANs
 - Vx Configuration on both Primary and Secondary Vlan
 - 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.
 - In Logical Chassis mode source macs may not learn on PVLAN configured ports, after deleting some of the secondary vlans for which the traffic is not flowing.

UDLD

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

STP/DiST

- VDX does not support tunneling non-standard BPDUs and thus IEEE BPDUs (0180:C200:0000) generated as tagged packets in STP/RSTP/MSTP modes may not be tunneled successfully across VCS Fabric. However, VDX supports tunneling standards' based BPDUs such as untagged IEEE BPDUs and tagged or untagged PVST BPDUs (0100:0CCC:CCCD). Post 3.0.1, the tagged IEEE BPDU can be tunneled across VCS fabric using command: “tunnel tagged-ieee-bpdu” under interface configuration.
- In Fabric Cluster mode, global spanning-tree configurations (STP enable, STP Vlan configurations, STP over vLAG configurations) have to be performed in all the switches in

VCS at the same time. For example, to run spanning-tree, it has to be enabled on all the switches including switches that don't have any edge ports.

- By default global spanning-tree and interface level spanning-tree will be disabled, user has to explicitly enable on the desired ports. VLAN spanning-tree state is default enabled.
- BPDU tunnel configurations are permitted only when spanning-tree is disabled in VCS.
- For cisco proprietary Per Vlan Spanning Tree protocols (PVST and RPVST) user needs to configure Extreme switch to send BPDU on Cisco multicast destination mac address "0100.0ccc.cccd" for non-native VLANs. By default, NOS 6.0.1 software uses Extreme "0304.0800.0700" multicast mac to send BPDU's on non-native VLANs.

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

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

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

Edge Loop Detection (ELD)

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

Long Distance ISL Ports

- Long distance ISL configuration ("long-distance isl" command) is not allowed if CEE Map/fcoeport is configured on any edge ports in the same port group.
- CEE Map modification is not allowed when long distance ISL is configured.
- A maximum of three PFCs can be supported on a long distance ISL 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.
- The “long-distance isl” command based extended fabrics are supported only on 10G interfaces.
- The 40G and 100G interfaces do not support “long-distance isl” command, however can extend distances for non-lossless traffic up to 40Km using standard ISLs.
- On standard ISLs, the 10G, 40G and 100G interfaces support lossless traffic up to 1Km.
- The “long-distance-isl” command will not be supported on the SO-10GE-ZR-CX, 10G-SFPP-ZR, and 10G-SFPP-ZRD-T 80km optics.
- The SO-10GE-ZR-CX, 10G-SFPP-ZR, and 10G-SFPP-ZRD-T 80km optics requires a minimum distance of 20km in order to successfully form a standard ISL connection

AMPP and Port-Profiles

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

vCenter

- Receiving more than five vCenter events within a span of 30 seconds, results in asset discovery getting initiated. Post discovery cluster configuration will be in sync with vCenter.
- vCenter auto-profile is automatically added/deleted to the default port-profile-domain in Service VF enabled/disabled mode.
- Modifying/editing the auto port-profiles in the default-domain is not recommended, which may cause auto-pp application failure during vCenter operation and end up in traffic failure.
- 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.
- Output of show vnetwork vss displays the vmnic against the vSwitch even after the removal of the vmnics from the vSwitch through vCenter. Recovery happens in the next auto-discovery cycle.

QoS

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

- The CLI “qos trust cos” is not applicable in VCS mode. However, “show qos int” will show as cos is trusted on ports on which “cos-mutation” or “cee default” config is applied.
- Configuring an interface with a nondefault DSCP-to-traffic class-map is allowed. However, configuring an interface with a nondefault CoS-to-traffic class-map is not supported.

FCoE

- On switches shipped with NOS7.0.1, the default mode of operation is Access Gateway for the VDX 6740, 6740T, 6740T-1G, VDX 2741, VDX 2746. The VDX 2741 was supported in 5.0.x (not supported in 6.0.x) and is upgradable to 7.0.1.
- Logical SANs have been supported since Network OS v6.0.0. If user needs to enable Fibre Channel Forwarder (FCF) mode, switch needs to be configured in FCF mode. If the switch is upgraded from a lower NOS version (pre 4.1.2 builds) to v7.0.1, it will be in FCF mode.
- FCoE logical SAN is not supported in an FC cluster. 
- Extreme recommends not having FCoE ports and Long Distance ISL ports in the same port-groups. This configuration will NOT be prevented by the CLI; however it can result in unpredictable behavior for FCoE traffic.
- If the FCoE FCMAP is changed from default to non-default, user might see that FCoE login may not happen. Please toggle the interface using “shutdown” followed by “no shutdown” to work this around.
- When an RBridge is removed from VCS cluster, it does not need to be manually removed from fcoe fabric-map. 
- MAC binding for remote SANs is not honored during config replay.
- In case the FIF is multiple hops away from the FCF, it is recommended that the interfaces of the intermediate switch shouldn't be configured with the same remote san as that of the FIF/FCF
- VLAN's which are reserved for FCoE may not be used for any other purpose. This is true for both Fabric Cluster and Logical Chassis modes.
- Extreme recommends that for all LAGs with FSB, the fcoeport config must be applied on the LAG itself. And for all LAGs with directly attached CNAs, the fcoeport config must be applied on the member ports.
- Binding an enode mac to FCoE interface is not allowed in range context, as only one enode mac can be bound to one FCoE interface.
- While providing range for FCoE interfaces, it's recommended to provide the range only in ascending order. For ex: interface fcoe 1/48/11-38 is recommended, interface fcoe 1/48/38-11 is not recommended.
- FCoE traffic may not be mirrored using RSPAN. Workaround is to use SPAN.
- In use cases with FSB, it is noticed that after converting dynamic port-channel to static, hosts and targets don't see each other.
- In NETWORK OS V6.0.1 and later, up to four FCoE Vlan's are supported in VDX . But on a single VDX , All member ports in a LAG have to be configured with the same FCoE Vlan. Different LAG can be configured with different FCoE Vlan.
- In NETWORK OS V6.0.1 and later, it is recommended user define different fabric-map for Remote Logical SAN and Local Logical SAN configuration. If user configures a fabric-map to

work on Remote Logical SAN first and then later change the same fabric-map to become Local Logical SAN, it may cause FCoE port continuously flapping.

- In NETWORK OS V6.0.1 and later, when FCoE CNA connect through VDX 6940-36Q/VDX 6940-144S to a Remote Logical SAN, if user performs certain operation in AG switch (e.g. N-port failover, VF-port remapping, fcmmap change etc), FCoE CNA may fail to login. The workaround is to do shut and no shut on the FCoE port on which FCoE CNA is connected.
- LC27x40G and LC12x40G linecards do not respond to incoming Ethernet pause (802.3x) and PFC frames in native 40G interface mode. This does not impact throttling of incoming traffic in response to PFC or Pause frames transmitted by the linecards. In order for LC27x40G and LC12x40G linecards to respond to pause frames, it is required to configure the 40G interfaces in breakout mode and use 40G breakout optics with regular native 40G cables. This is required for FCoE traffic.

FlexPorts

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

Fibre Channel

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

Access Gateway

- The switch can be operated as Fibre Channel Forwarder (FCF) by disabling Access Gateway mode.

- AG does not bridge the VCS and SAN fabrics because hosts connected to the AG switch are registered in the SAN name-server only. Therefore, all zoning operations for AG are done on the SAN fabric.
- At least one N-port must be online in order for FCoE devices to 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>”.
- Cannot configure the default Logical SAN to the interfaces in the FCF-group switch.

IP Fabric

BGP eVPN :

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

ARP/ND Suppression:

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

Conversational ARP:

- It is recommended to enable both Conversational-ARP and Conversational-MAC together.

Static Anycast Gateway:

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

ND/RA

- Proxy ND is not supported.

IPv4

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

BFD

- Static Route BFD, BGP BFD and OSPFv2/v3 BFD
 - For Single HOP BFD sessions configured with source IP as secondary IP is not supported, since significance of Source IP in BFD configuration is only to determine on which interface BFD session should be started and hence interfaces' Secondary IP is not used as source in BFD PDU.
 - BFD is not supported on leaked routes.
 - BFD for multi-HOP BFD neighbor reachable via multiple paths with equal cost(ECMP) will not be supported since BFD requires BFD session to be created for the neighbor for each ECMP path.
 - BFD is not supported for OSPFv2 & OSPFv3 virtual links.
 - For single hop BFD sessions, BFD will consider the interval values that are configured on interface, and not the non-default values that are configured within the global command.
- BFD for VxLAN tunnels

- BFD session may not come online or may flap if VCS cluster is in transient state during reload, vLAG failover, fabric split, chassis disable/enable and such scenarios. It is required to have a stable VCS cluster in order for BFD sessions on VxLAN tunnels to work as expected.
- BFD parameters are not configurable on VCS VxLAN Gateway for Type NSX. The parameters are derived from NSX controller.

VRRP

- VRRP and VRRP-E cannot be enabled together on VDX 6740 and VDX 6740T platforms.
- 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 VDX 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.
- Fabric Virtual Gateway (FVG) is not applicable in IP Fabric environment, Static Anycast Gateway to be used to achieve similar functionality.

OSPFv3

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

BGP

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

Layer 2/Layer 3 Multicast

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

VRF

- Under VRF submode there is a syntax change for the address-family 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.

- Removing VRF address family on a non-default VRF will delete all relevant address-family configurations including the interface and protocol configuration for that VRF.

BGP-VRF

- Local-as <num> can be configured for particular VRF under "address-family ipv4 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.

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.
- ACL with "Routed" keyword functions only for VE/Router Port MACs. It does not work for VRRP Routed.
 - Work-around: Apply default mode ACLs (No "routed" keyword).
- For Private VLANs, Egress ACLs on Primary VLAN is applied only for all traffic which ingresses primary VLAN i.e.
 - If the traffic ingresses from Primary VLAN but gets translated to Secondary VLAN at egress, ACL on primary VLAN at egress is still applicable 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.

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.
- Clients may not converge in some IP Fabric environment. Care should be taken to not configure DHCP IP helper and Static Anycast Gateway on the same interface.
- Two DHCP OFFER per one DHCP DISCOVER and two DHCP ACK for single DHCP request seen IP fabric setup

Dynamic ARP Inspection (DAI)

- The ARPs learnt on trusted ports would be deleted when DAI is enabled or DAI filter changed.
- Static ARPs not permitted by DAI filter would be promoted to active state. Administrator is responsible for configuring static ARPs in sync with 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 Principal 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.
- When there are no uplink interfaces configured, the track disable command will remove tracking configuration from internal database and this behavior is applicable only in 6.0.1a patch and not in prior releases.
- If track min-link number is greater than the number of uplinks, then the downlink will be shutdown with a warning message.
- After toggling the line card using "power-off / on", LC related interfaces that are configured as uplink interfaces are not seen in "show track summary" cli output.

OpenFlow

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

- Uncontrolled failover on 6740 and 6940 would need power-cycle to recover hardware resources for the feature to continue to work.
- Pre-provisioned flow-mods will not be replayed to a new slot coming online. GROUP and METER configurations will be always replayed.
- On the Extreme VDX 8770, queue statistics should be interpreted as wire-vlan (COS) priority statistics.
- Actual number of supported flow-mods (L2/L3) may be less since MAX scale values include per port default miss entries, and single LLDP entry is needed for topology discovery. This applies to all supported platforms.
- For layer 3 rules, switch can't differentiate between tagged and untagged flows when matching against rules. This applies to all supported platforms.
- Filtering options are not supported for show openflow CLIs. Show openflow commands with filter option show the complete output.
- For the port based flow mod, if the ivid reference is active, egress tagging is not cleared. The new flow mod will not be installed If the previous flow mod has created the egress tagging behavior. This case has to be handled by work-around flow mods or take the port off from openflow and bring it back.
- With default rcv-queue and after coldboot, group select traffic may not be correct, need to do shut/no shut on the interface. This issue is not there with non-default rcv-queue.
- With large number of flows, "show openflow flow <>" may take 20 seconds to display packet counts.
- "Module Unknown" is shown for CLI "show open resources" in VDX 6940-144S.
- Openflow is not supported on Lag/vlag or port-channel interface.

Layer 2 and Layer 3 ISSU on VDX 6740x

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

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

The implementation is based on a type-1 hypervisor.

Extreme 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 is that the 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 is that the 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.
- 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.
- On a large cluster (of 32 nodes or more) and with scaled up configuration, it is recommended to query configuration using rBridge ID filter. In extreme scenario, querying cluster wide configuration without specifying rbridge ID filter might cause switch to run out of memory.
- Maximum 16 sessions supported.

VXLAN Gateway for VMware NSX

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

VF Extension using VxLAN

- VF Extension overlay-gateway (VTEP) is supported only on the VDX 6740, VDX 6740T, VDX 6740T-1G, VDX 6940-36Q, and VDX 6940-144S.
- VF Extension overlay-gateway is supported only in the VCS Logical 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 with NOS v6.0.1. Such topologies and configuration must be removed before downgrading to any version below NOS 6.0.1.
- Fabric-Virtual-Gateway (FVG) based VTEP is not supported. CLIs for configuring FVG as VTEP are available under overlay-gateway, however these CLIs should not be used as the functionality is not available in this release.

TCAM Profiles

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

Management VRF

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

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

Conversational MAC Learning

- Conversational MAC Learning and 'Disable Source MAC Learning' cannot 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
<code>/hardware/flexport <interface tuple>/type fibre-channel</code>	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.
<code>/rbridge-id <#>/vrf <name></code>	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.

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

HA on TOR switches

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

Logical Chassis HA

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

Interoperability

- In a VPC environment where the Extreme 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 Extreme 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 on VDX8770.
- For frames forwarded on a transport fabric, ingress CTAG tagging is preserved at the egress port regardless of the egress tagging classification.
- Default-VLAN can only be configured using TRANSPORT VF IDs.
- The “no vcs virtual-fabric enable” command execution time is dependent on the number of ISLs and VLANs in the VCS.
- The virtual-fabric resource allocation are platform dependent as follows:
 - VDX 8770 – no limitation
 - VDX 6740/6740T/6740T-1G – uses TCAM table
 - VDX 6940-36Q – virtual-fabric transport and service VLANs use TCAM and EXM table respectively.

MAPS

- BNA version 14.0.1 is supported with NOS7.0.1
- MAPS is supported on VDX 2741, 2746, 6740, 6940 and 8770 platforms.
- MAPS port level alerting in NOS V7.0.1 is not available for Flex Ports configured in Fiber Channel mode.
- RX_SYM_ERR MAPS messages are displayed when breakout cable is connected on a 40G interface that is not configured for breakout.
- When line card on the remote end of the link is powered off, MAPS generates Insertion/Removal notification for the SFPs on the local side. These can be ignored.
- 100G SFP threshold monitoring is not supported on VDX6940-144s.

Maintenance Mode

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

BNA

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

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

Miscellaneous

- Extreme 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 4087-4095 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.
- 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 reseal the card.
- On rare instances of ISSU, HA failover, line-card may turn faulty. Workaround is to reset the line-card.
- PCAP utility is not supported on standby MM on VDX 8770.
- Please make sure to not have large no of unreachable tacacs+ accounting server configured, else it might cause unit to reboot. This issue is hit only with large config (4K VLAN etc and 20K lines or config).
- Configuration of more than one In-band management port on a single switch is not recommended.
- Under certain stress conditions the 'copy support' command might time out for some modules. In such cases, it is recommended to retry 'copy support' with a higher timeout multiplier value.
- It is highly recommended to copy the configuration file to running-config and then save the running-config to startup-config, instead of directly copying the external configuration file to startup-config, especially when using fabric distributed features such as Zoning, VM Aware Network Automation and Virtual IP.

Defects

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

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

TSB Issues Resolved in Network OS v7.0.1

TSB	Summary
TSB 2016-236-A	<p>A component within the 10G Base-T power circuitry can overheat causing the switch to power off and be unusable.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>Upgrading to one of the specified firmware versions or any higher version of the Extreme Network OS will provide an increased fan speed and airflow to prevent any overheating. The nominal power consumption of the switch will only be increased by 20W with this change, and the maximum power consumption is not affected by this change.</p>

Network OS v7.0.1 Caveats

BFD

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

VxLAN

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

Note:

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

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

- On occurrence of events that may bring down the tunnel on an R-Bridge, there could be few seconds of traffic interruption due to a default de-bounce-timer which is set to 2 secs, this could delay the fail-over of the traffic to redundant path. A debug command "show system internal tnlmgr de-bounce-timer 0 0" can be utilized to reduce the traffic impact, however, the command settings are not persistent across reloads.
- On sending IGMP queries over VF_Extension Tunnel with VLAG as underlay, packets might loop over the tunnel .Queries come back from the same tunnel interface from which its egressed out.
- "show ip igmp groups interface tunnel <tunnel_id>" cli shows all IGMP interfaces instead of just the tunnel interface.
- Adding and Removing RBridges under overlay-gateway may take longer than expected time if large number of VLANs are configured in the fabric.

Long Distance ISL

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

Loopback interfaces

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

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

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.

OpenFlow:

Network OS v7.0.1 Release Notes v5.0

- 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.
- Hardware Profile:
 - When modifying the route-table profile type and maximum-path using the hardware-profile command, the user should only change one parameter at a time. Otherwise the maximum-path setting will be incorrect. If the issue already occurred, the user can re-run the command to set the maximum-path with the correct value.
- Copy Config command:
 - In VDX6940-144S, 100G mode configuration replay can fail when executing "copy <file> running-config" if DPOD license is not reserved. To work around this issue, the user can manually reserve the license and then run “copy <file> running-config”.

Syslog:

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

Closed with code changes for NOS 7.0.1

Closed with code changes for NOS v7.0.1

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

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

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

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

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

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

Closed with code changes for NOS 7.0.1

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

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

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

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

Defect ID: DEFECT000578258	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: IP Fabric
Symptom: Traffic loss may be observed for destination subnets under non-default VRF advertised over BGP-EVPN using L3VNI.	
Condition: Leaf nodes extending the VRF over BGP-EVPN are not advertising any prefix route.	
Workaround: One of following options may be chosen: <ul style="list-style-type: none"> (1) Redistribute connected routes under VRF into BGP VRF. (2) Configure static IP route and redistribute into BGP VRF. (3) Configure network or static-network under BGP VRF instance. 	

Defect ID: DEFECT000579234	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VLAN - Virtual LAN
Symptom: Multicast traffic destined for static multicast address, will flood, if the mac is configured on remote node of VCS.	

Closed with code changes for NOS 7.0.1

Condition: Static multicast MAC is configured in a remote node within a VCS, with no local interface part of the group.

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

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

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

Defect ID: DEFECT000582010	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: ARP - Address Resolution Protocol
Symptom: Under rare conditions, some of the hosts may lost IP connectivity with the VDX switch acting as a layer-3 gateway.	
Condition: This would occur if the MAC to the IP association of a VDX learnt ARP changes. ie. For the same IP address, the MAC changes from say Mac1 to Mac2.	
Recovery: "clear arp no-refresh" would clean the ARP table and recover from the problem state.	

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

Closed with code changes for NOS 7.0.1

The tunnel terminated IGMP frames sent to other nodes are trapped and flood back on the vlan by control path. These packets can loop back to source node.
Recovery: Shut down the tunnel

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

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

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

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

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

Defect ID: DEFECT000585392

Closed with code changes for NOS 7.0.1

Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: CLI - Command Line Interface
Symptom: BNA unable to discover NOS switch, when NOS switch is connected to a FCR.	
Condition: 1. Have a setup with VCS connected to FCR 2. Install BNA 14.0.1 3. Shift to IP tab. 4. In Discovery dialog, add IP of one of the VCS switches and click ok. 5. Observe the device is not discovered and shows "Discovery Failed" message.	

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

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

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

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

Closed with code changes for NOS 7.0.1

Defect ID: DEFECT000585960	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: VCS Fabric
Symptom: 40G Interface is administratively (or) protocol down with FFDC raslogs	
Condition: Admin operations on 40G Interface.	
Recovery: Shut/no-shut both interfaces on either side of the link	

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

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

Defect ID: DEFECT000586178	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: Non-existent port-channel shows up in "show fcoe interface ethernet"	
Condition: <ol style="list-style-type: none"> 1. Create a port-channel 2. Add members to it and make it fcoe-provisioned 3. Delete the port-channel 	
Workaround: Remove FCOE provisioning from port-channel before deleting it	

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

Closed with code changes for NOS 7.0.1

Defect ID: DEFECT000586338	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: IN VDX 6940-144S, link flap occurs on 40 G ISL ports if breakout configuration mis-matched with any adjacent ports.	
Condition: In VDX 6940-144S, a 40 G Port with breakout QSFP is not configured as breakout may cause adjacent 40 G port to flap, whereas its peer port is configured as 40G breakout.	
Workaround: In VDX 6940-144S, configure 40 G port as breakout if the peer port is configured as 40G breakout. After that, the link flap on the port will stop.	

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

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

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

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

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

Defect ID: DEFECT000587615	
-----------------------------------	--

Closed with code changes for NOS 7.0.1

Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: The SNMP V3 traps may not be received for the SNMP v3 host recipients configured under rbridge mode.	
Condition: The trap may not be received after upgrade from NOS6.0.1a to NOS7.0.0 with cold boot option	

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

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

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

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

Closed with code changes for NOS 7.0.1

Recovery: Issue chassis disable then chassis enable.

Defect ID: DEFECT000587804	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: IP Fabric
Symptom: Even though there are no matching EVPN import route-targets configured under VRF, imported EVPN routes are present in BGP VRF table.	
Condition: EVPN import route-target(s) is/are removed while matching routes are present in BGP-EVPN and imported into BGP VRF table.	
Workaround: Issuing "clear bgp evpn neighbors all soft in" command should cleanup the routes which are still imported in BGP VRF instance after matching EVPN import route-targets are removed.	

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

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

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

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

Defect ID: DEFECT000588238	
Technical Severity: High	Probability: High

Closed with code changes for NOS 7.0.1

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

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

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

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

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

Defect ID: DEFECT000588918	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.1	Technology: VCS Fabric
Symptom: Customer encountered an unexpected VDX6740 reload:RAS logs & stack trace for the reset as below:2016/02/16-01:02:57, [SEC-1203], 795596, SW/0 Active, INFO, NINMUM03-DC1-R107-NSPL-RTR-049, Login information: Login successful via TELNET/SSH/RSH. IP Addr: A.B.C.D.2016/02/16-01:03:13, [SEC-3022], 795597, SW/0 Active, INFO, NINMUM03-DC1-R107-	

Closed with code changes for NOS 7.0.1

<p>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*****</p>
<p>Condition: When high rate of TFTP ip directed broadcast packets are sent destined to known subnets.</p>

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

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

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

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

Defect ID: DEFECT000590808	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management

Closed with code changes for NOS 7.0.1

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 un hiding and configuring them. Even the copy running to file was not having the configuration after copy command was executed after un hiding.	
Condition: Config commands under hide group "debug" and "foscmd" have to be executed after un hiding respective hide group. Post this, executing "show running config" will not show these un hidden 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.	

Closed with code changes for NOS 7.0.1

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 'ves replace' operation.
Recovery: Rebooting the affected node forces it to refresh the effective priority value for the multi-cast tree root priority. Alternatively, explicitly changing the priority to a different value and then setting it back to the original desired value causes the priority to be updated. However, setting the root priority to a different value may affect the multi-cast tree formation depending on the temporary priority specified.

Defect ID: DEFECT000592617	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.2	Technology: xSTP - Spanning Tree Protocols
Symptom: IEEE BPDU Local VLAN tunnel CLI allowed to be configured when protocol spanning tree is already configured or vice versa.	
Condition: When both STP protocol and IEEE BPDU Local VLAN tunnel CLI are enabled at the same time.	

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

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

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

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

Defect ID: DEFECT000594815

Closed with code changes for NOS 7.0.1

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

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

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

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

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

Defect ID: DEFECT000596496	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.2	Technology: Logical Chassis
Symptom: Protocol spanning-tree configuration will not be allowed even after removing the "spanning-tree ieee-bpdu limit-vlan-flood" and "tunnel tagged-ieee-bpdu" configuration.	

Closed with code changes for NOS 7.0.1

Condition: When all the switches in the VCS are configured with "spanning-tree ieee-bpdu limit-vlan-flood" and one or more switches are removed from VCS.
Recovery: Copy running configuration to remote. Reload the switch with default configuration and copy back the running configuration.

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

Defect ID: DEFECT000600386	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: When EVPN related configuration is done in RBridge range mode, the nodes might reboot	
Condition: When EVPN related configuration is done in RBridge range mode	
Workaround: Instead of using RBridge range mode, use RBridge mode to configure the individual rbridges	
Recovery: If issue is encountered, reboot the RBridges to recover	

Closed with code changes for Network OS v7.0.0

Closed with code changes for NOS v7.0.0

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as May 3, 2016 in Network OS v7.0.0.

Defect ID: DEFECT000485838	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS4.1.0	Technology: Configuration Fundamentals
Symptom: When applying an access-list with wildcard mask to the management interface it does not correctly permit/deny traffic.	
Condition: When ACL rule is added with wildcard mask	
Workaround: Add ACL rule without the wildcard mask.	

Defect ID: DEFECT000489898	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS4.1.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: '<cr>' appears as possible command completion after 'no table-map'	
Condition: Unconfigure table-map configuration	
Workaround: Use 'no table-map <route-map-name>' as in 'no table-map map!'	

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: 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 .	

Closed with code changes for Network OS v7.0.0

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000509978	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.0	Technology: Configuration Fundamentals
Symptom: On VDX8770 Platforms with large scale configuration, the support save operation may time out	
Condition: 1) Configure large number of VEs and VRFs e.g. 2K VEs, 1000 VRFs 2) issue the support save cmd	

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: DEFECT000518899	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: NOS5.0.0	Technology: 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: DEFECT000522928	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS5.0.0	Technology: VCS Fabric
Symptom: When user enters "police" command under "class-map" mode, the following error is reported and cmd exits to config mode: 'syntax error: incomplete path: 'police-priority-map''	
Condition: Observed when the user attempts to enter the "police" command under class-map mode in config mode.	

Defect ID: DEFECT000523737	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.0	Technology: AAA - Authentication, Authorization, and Accounting
Symptom: syslog messages are not generated for authentication failure in REST	
Condition: Login with REST interface(EM)	
Recovery: Run "show logging auditlog" for failed login attempts	

Defect ID: DEFECT000525991	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: BGP session will get terminated	
Condition: BGP update message with error	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000527713	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: 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: DEFECT000528456	
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: When running the Silver Creek IP MIB suite on the VDX switch, many errors are noticed. These errors are related to Set operations on IP MIB objects.	
Condition: The issue is seen when running Silver Creek IP MIB suite on VDX switch because VDX does not support many tables/objects in the IP MIB.	

Defect ID: DEFECT000530174	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.1	Technology: VLAN - Virtual LAN
Symptom: There is no option for port-channel in the “show mac-address count conversational” CLI command	
Condition: When attempting to issue "show mac-address count conversational" CLI command	

Defect ID: DEFECT000530654	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS4.1.2	Technology: Logical Chassis
Symptom: There is no “show” command to track the TCAM utilization for various protocols	
Condition: Normal operating conditions, mostly with scale configurations.	

Defect ID: DEFECT000531647	
Technical Severity: Medium	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: When traffic generator sends a DD packet and retransmits the same packet after RouterDeadInterval seconds, SeqNumberMismatch event is not generated.	
Condition: SeqNumberMismatch event will be generated when traffic generator sends a DD packet and retransmits it after RouterDeadInterval.	

Defect ID: DEFECT000532608	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS5.0.1	Technology: PIM - Protocol-Independent Multicast
Symptom: VDX PIM router does not have upstream join suppression capability. As a result, even after seeing neighbor sending join upstream, switch does not suppress its scheduled periodic join leading to bandwidth wastage.	

Closed with code changes for Network OS v7.0.0

Condition: VDX PIM enabled router's upstream join suppression behavior is not according to RFC 4601 which leads to unnecessary protocol traffic in the LAN.
--

Defect ID: DEFECT000537193	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	Technology: 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: DEFECT000538299	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.0	Technology: ACLs - Access Control Lists
Symptom: When ACL rules with IPv6 address masks are configured on ethernet management ports, the IPv6 mask addresses are displayed incorrectly in the output of “show access-list ipv6” command.	
Condition: The behavior is observed only when ACL rules with IPv6 address masks are configured on ethernet management ports.	
Workaround: Command “show running-config ipv6 access-list” can be used instead to identify the mask addresses.	

Defect ID: DEFECT000538468	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS4.0.1	Technology: Configuration Fundamentals
Symptom: CONFD and DCMD timeout seen during the copy support operation.	
Condition: FTP server connectivity lost can lead into this issue.	
Workaround: Make sure FTP server is reachable at all times.	
Recovery: Retry the copy support operation once FTP server is reachable.	

Defect ID: DEFECT000539407	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
Symptom: USB enable command will fail.	
Condition: This behavior is observed if a new user session is launched or started in between usb enable and disable command execution.	
Workaround: The user is warned with RAS-3009 message, when issue symptoms are observed so that, before the issue can occur, the user can apply the workaround. The workaround is that, User needs to logout of all user sessions and re-login once RAS-3009 log is observed.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000539623	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.0	Technology: Hardware Monitoring
Symptom: VCS fabric formation protocol will not work when two VDX switches are cabled using a standard, non-breakout 40 Gbe cable. The "show fabric isl" command will show "protocol incomplete" for the 4 x 10 Gbe links.	
Condition: The problem only occurs if one of the VDX switches is a VDX 6740 (or 6740T or 6740T-1G), and it is cabled to another VDX switch using a standard, non-breakout 40 Gbe cable, and then both switches are dynamically reconfigured for 4 x 10 Gbe mode using the "sfp breakout" command.	
Workaround: After "sfp breakout" command on both VDX switches, save the configuration using the "copy running-configuration startup-configuration" command. Then reboot both VDX switches.	
Recovery: Save the "sfp breakout" configuration using the "copy running-configuration startup-configuration" command. Then reboot both VDX switches.	

Defect ID: DEFECT000539755	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
Symptom: User will see an error message while enabling break out port from Element Manager and will fail to enable.	
Condition: Enabling break out port from Element Manager for VDX2741 device	

Defect ID: DEFECT000542893	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: 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: DEFECT000543010	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS5.0.1	Technology: VCS Fabric
Symptom: On abrupt removal of an RBridge from VCS, there may be few ms of traffic loss for layer-2 traffic.	
Condition: When an RBridge is removed from the fabric either via reload OR powering off the switch	

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

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000545530	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Syntax error is displayed upon entering the command 'no maximum-paths'.	
Condition: Configure 'maximum-paths <value>'. Unconfigure using command 'no maximum-paths'	
Workaround: Include the maximum paths value also in the 'no' form of command as in 'no maximum-paths <value>'.	

Defect ID: DEFECT000545642	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP v6 Client bind rate is low	
Condition: Auto provisioning DHCP v6 clients using DHCP v6 relay	

Defect ID: DEFECT000547565	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.0	Technology: IGMP - Internet Group Management Protocol
Symptom: Multicast flooding for a very small amount of time - in the range of 15 to 20 microseconds.	
Condition: IGMP is configured and an IGMP receiver leaves a group.	

Defect ID: DEFECT000548196	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: ICMP - Internet Control Message Protocol
Symptom: Switch is not sending ICMPv6 Redirect.	
Condition: Traffic forwarding in suboptimal path	

Defect ID: DEFECT000548905	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: ' ' appears as possible command completion after 'no maximum-paths'	
Condition: Unconfigure maximum-paths configuration	
Workaround: Use 'no maximum-paths <value>' as in 'no maximum-paths 5'	

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.	

Closed with code changes for Network OS v7.0.0

Condition: New feature support.
--

Defect ID: DEFECT000549853	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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: DEFECT000550271	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS4.1.3	Technology: FCoE - Fibre Channel over Ethernet
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: DEFECT000550826	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: NOS5.0.1	Technology: 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: 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: DEFECT000552066	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.1	Technology: 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	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000552067	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	Technology: 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: DEFECT000552365	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: BFD on DGW flaps with disruptive activities on the ecmp path.	
Condition: If OSPF peers have different default SPF throttle timer values for OSPF.	
Workaround: OSPFv2 SPF timer values can be configured to (0,0,0) using OSPF config CLI (timers throttle spf 0 0 0) under router ospf [vrf <name>] mode.	

Defect ID: DEFECT000552369	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BFD - BiDirectional Forwarding Detection
Symptom: Tunnel BFD session can flap when the Tunnel is operationally down in any of the RBridges.	
Condition: When Tunnel is using vLAG as underlay and the RBridge doing BFD frame transmission leaves the vLAG.	

Defect ID: DEFECT000552570	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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: DEFECT000552923	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Sysmon
Symptom: When the MAPS feature is enabled, the Fabric Watch feature will be automatically disabled without any user intervention. Even though Fabric Watch is disabled its configuration is displayed in the running-config. The Fabric Watch config will not be effective.	
Condition: Seen when the MAPS feature is enabled on the switch.	

Defect ID: DEFECT000553037	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: Help does not show SFTP support available, but SFTP is supported for config upload.	

Closed with code changes for Network OS v7.0.0

Condition: SFTP Functionality is present. Help is not updated. This is only a documentation issue.

Defect ID: DEFECT000553242	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: PortChannel interface SNMP Stats are not reset after executing 'clear counter all'	
Condition: 'clear counter all' doesn't reset SNMP stats of PortChannel interface	

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: DEFECT000554298	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Hardware Monitoring
Symptom: Configuring speed command on 1G interface for VDX2746 may cause the port to go down and remain in down state.	
Condition: Configure speed on 1G interface using the CLI command "speed 1000".	
Workaround: Do not configure speed manually on 1G interfaces	
Recovery: Remove the speed command	

Defect ID: DEFECT000554372	
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: 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 Group: VCS
Reported In Release: NOS5.0.2	Technology: 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: DEFECT000554506	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: OSPFv3 - IPv6 Open Shortest Path First
Symptom: OSPFv3 max-metric on-startup option for wait-for-bgp will not work.	
Condition: OSPFv3 max-metric on-startup option for wait-for-bgp if configured will not work.	
Workaround: OSPFv3 max-metric on-startup option with timer can be used instead	

Closed with code changes for Network OS v7.0.0

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000554793	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.1	Technology: ACLs - Access Control Lists
Symptom: The output display for "show ip arp inspection vlan" is not in ascending order	
Condition: execute show ip arp inspection vlan command	

Defect ID: DEFECT000554905	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: When applying access-list configuration, the user may not be able to notice when secondary switches run out of TCAM space & access-list configuration is not successful.	
Condition: When secondary switch TCAM gets full, the Raslog is not generated on Principal switch.	

Defect ID: DEFECT000554989	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: IP Addressing
Symptom: Under certain rare traffic conditions, the VDX 6740 may experience an unplanned reload with the ARPD termination.	
Condition: When high rate of certain pattern of traffic is sent towards a high scale of hosts.	

Defect ID: DEFECT000555058	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: BGP session might take 3 up to minutes for peer to reach establish state.	
Condition: maxas-limit is configured and then "clear ip bgp neighbor all" is triggered.	

Defect ID: DEFECT000555211	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: DAD request fails and DAD-1306 ERROR shows on the console.	
Condition: The "vesmode" value in dad.conf is set to "FC"	
Workaround: The "vesmode" value in dad.conf has to be set to "LC" regardless of the existing cluster mode.	

Defect ID: DEFECT000555470	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: An interrupted CLI command may still get committed even though it is not entered.	
Condition: If a CLI command has multiple inputs where last input is a leaf prompting the user and when the last input is not provided then command will still be accepted.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000555589	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: "DHCP Auto-Deployment node succeed" message seen before user configured file is fully executed	
Condition: If ZTP configuration file includes non ZTP related node's configuration.	
Workaround: Remove non ZTP related nodes configuration.	

Defect ID: DEFECT000555772	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS4.1.3	Technology: 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.	
Recovery: Shut/no-shut recovers it.	

Defect ID: DEFECT000555882	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000555954	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: BFD sessions flap if session transmit and receive interval are configured less than the non-default value of 500 ms.	
Condition: Session interval parameter configured with less than recommended.	
Workaround: Use BFD transmit and receive intervals greater than or equal to default 500 ms.	
Recovery: Reconfigure BFD session intervals greater than or equal to default 500 ms.	

Defect ID: DEFECT000556079	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000556094	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: 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 Group: Monitoring
Reported In Release: NOS4.1.3	Technology: Hardware Monitoring
Symptom: VDX 10gb port LED lights up green when only TX strand is inserted.	
Condition: This issue is seen when only TX strand of the fiber cable is inserted 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 Group: IP Multicast
Reported In Release: NOS4.1.3	Technology: 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: DEFECT000556923	
Technical Severity: Critical	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.0	Technology: ACLs - Access Control Lists
Symptom: VDX is unexpectedly reloaded when multiple configuration REST requests are sent at the same time.	
Condition: When VDX switch is bombarded with REST requests from multiple servers continuously over prolonged period, it might go for an unexpected reload	
Workaround: Do not send multiple REST requests at the same time	

Defect ID: DEFECT000557061	
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: "switchport access vlan <vlan-id> mac-group <mac-group-id> " cannot be deleted through the REST URI	
Condition: When the URI for the DELETE request contains the keys for the resource without the payload	
Workaround: If the keys are specified in the payload of the REST request, the config can be deleted.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000557197	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000557575	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: unexpected reload of system.	
Condition: When there are many unique BGP attribute and command "no router bgp" or "clear ip bgp neighbor all" are triggered	

Defect ID: DEFECT000557683	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS4.1.3	Technology: Hardware Monitoring
Symptom: VDX ports connecting to certain Intel NIC cards, temporarily flap under certain corner case conditions.	
Condition: Repeated cable pull/re-insert of VDX port 10gT connected to Intel NIC may result in temporary port flapping.	
Workaround: No human intervention is needed to workaround this issue. The port will come online after a few seconds.	
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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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 Group: Management
Reported In Release: NOS6.0.1	Technology: Software Installation & 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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000558165	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: 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 Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000558213	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Hardware Monitoring
Symptom: "Vendor PN" info in "show media interface" command may not be refreshed.	
Condition: On VDX6940: 1) Replace 40G QSFP media with a different type 2) Issue "show media interface forty <interface>" command	
Workaround: Shut/no shut of the interface may display right info	

Defect ID: DEFECT000558224	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: Logins from Qlogic Converged Network Adapter connected to VDX VDX6940-36Q and VDX6940-144S does not come back up after certain operations.	
Condition: Logins from Qlogic Converged Network Adapter does not come back up after certain operations 1. After shutting VF ports, 2. After changing fcmmap 3. After changing N Port to F Port mapping on AG	
Workaround: Do a "no fcoeport" on the interface connected to Qlogic Adapter before doing the below operations 1. After shutting VF ports, 2. After changing fcmmap 3. After changing N_Port to F_Port mapping on AG and "fcoeport <fabric-map>" on interface connected to Qlogic ADapter afterwards to avoid the situation	
Recovery: Do a "no fcoeport" and "fcoeprt <fabric-map>" on the interface connected to Qlogic Adapter to recover.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000558266	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: Port-channel and its member port has contradictory port-profile-port configurations	
Condition: <ol style="list-style-type: none"> 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 Remove 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 Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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: DEFECT000558692	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: ICMP - Internet Control Message Protocol
Symptom: IPv4 Pings between VDX switches in VCS Fabric failing.	
Condition: IPv4 Pings between VDX switches in VCS Fabric failing, though ARP, MAC learned correctly after HA Failover or ISSU upgrade.	
Workaround: Coldboot the VDX switch.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000558794	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: "fcoe ha send eth port config: fcoe map name to profile failed" message coming up on console	
Condition: <ol style="list-style-type: none"> 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 Group: VCS
Reported In Release: NOS5.0.2	Technology: 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: 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: DEFECT000559136	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: Redundant PORT-1014 online messages may be displayed for the same port.	
Condition: A Link Reset (LR) on a FC link may cause redundant PORT-1014 online messages.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000559141	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: Logical port numbers shown instead of physical port numbers in [CBR-1014] and [CBR-1029] raslog messages used to notify credit loss on a Fiber Channel port	
Condition: Seen on credit loss on a FC port usually due to a noisy link. No functionality is broken.	
Workaround: Workaround is not needed and the switch software automatically takes corrective action (Link Reset) to restore the credits	

Defect ID: DEFECT000559194	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: 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 Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000559280	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: IP Addressing
Symptom: FFDC & PIMD termination may be triggered on reboot of a VDX switch enabled with PIM sparse mode	
Condition: When one of the VDX in the cluster enabled with PIM sparse-mode is rebooted.	

Defect ID: DEFECT000559281	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: Invalid DAD config flag - "scriptcfgflag" value is accepted and treated as 1.	
Condition: If the value of DAD config flag - "scriptcfgflag" is set to the value greater than 2, it will be treated as 1.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000559286	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: BFD sessions flaps	
Condition: BFD session flaps are observed when CPU utilization is very high due to system configuration during chassis enable followed by reload.	
Workaround: Use higher BFD packet transmit and receive interval	

Defect ID: DEFECT000559310	
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: On power-cycle Port channel link between VDX6940-144S 10G and VDX 8770 breakout may observe continuous link down messages.	
Condition: VDX6940-144S power-cycle and 10G link down message.	
Workaround: Interface admin shut/no shut on peer port or this port can recover this.	

Defect ID: DEFECT000559371	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: xSTP - Spanning Tree Protocols
Symptom: Changing Spanning-tree mode from MSTP to RVPST 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: DEFECT000559399	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS5.0.1	Technology: IPv6 Multicast VLAN Traffic Reduction
Symptom: High CPU utilixzation casued by IPv6 MLD packet looping in the VCS	
Condition: Mirroring (SPAN) where source and destination mirror ports are on different rbridges causes the loop.	

Closed with code changes for Network OS v7.0.0

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000559540	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000559566	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: Flow-mods having a group action along with PUSH/POP are rejected by switch	
Condition: Controller tries to install a Flow-mod having group and PUSH/POP in action.	
Workaround: VLAN translation for each of the member ports within the group buckets have to be updated by programming a low priority L2TTP flow-mod. Later original flow-mod can be added without an explicit PUSH/POP in action.	

Defect ID: DEFECT000559629	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: 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 Group: Management
Reported In Release: NOS6.0.1	Technology: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: VLAN - Virtual LAN
Symptom: On VDX6940 device, enabling Conversational MAC Learning (CML) feature may result in flooding VxLAN L2 Extension traffic.	
Condition: <ol style="list-style-type: none"> 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	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000559754	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.0	Technology: ACLs - Access Control Lists
Symptom: The tab auto-completion may display incorrect listing of available commands for IPv4 extended acl and MAC extended acl.	
Condition: When tab-completion method is used to fill in the rest of the command.	

Defect ID: DEFECT000559767	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.1	Technology: VCS Fabric
Symptom: After upgrade from firmware 5.0.1b to 6.0.1, 40Gb port may not be in ISL up state.	
Condition: After upgrade from firmware 5.0.1b to 6.0.1, 40Gb port may not be in ISL up state.	
Recovery: Shut/no shut the port	

Defect ID: DEFECT000559780	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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: 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: DEFECT000559831	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: Output of the "show openflow resources" is not property aligned, all details are captured.	
Condition: "show openflow resources" output not align to format.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000559847	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: Openflow CLI "do show openflow flow" displays all the flows, including flows not configured.	
Condition: OpenFlow configured on the node.	

Defect ID: DEFECT000559853	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: 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 Group: Management
Reported In Release: NOS6.0.1	Technology: 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 Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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 Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Port Mirroring
Symptom: When linecard is powered-off , BLADE STATE==IN rule is triggered instead of OFF rule.	

Closed with code changes for Network OS v7.0.0

Condition: When line card is powered off.
--

Defect ID: DEFECT000559920	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: 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 Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: 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 Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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 172.16.32.33 port 33"	

Defect ID: DEFECT000560008	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: The command "copy support sep group ?" does not display list of group names.	
Condition: User did "copy support sep group ?" command and there is no display of the list of group names.	
Recovery: User to use "show copy-support group-list" command to find out all possible string options for the group names.	

Defect ID: DEFECT000560025	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.1	Technology: User Accounts & 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)	

Closed with code changes for Network OS v7.0.0

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000560036	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: Unexpected reload on system under scale configuration	
Condition: With bgp & ospf scaled config, on 5 or more node setup, if upgrading firmware on all nodes while pumping 5K routes on all nodes in the fabric, there may be an unexpected reload	
Workaround: With scaled config and more node setup. While upgrading the boxes, reboot all boxes at the same time.	

Defect ID: DEFECT000560112	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS4.1.3	Technology: LAG - Link Aggregation Group
Symptom: Interface remains in Administratively down state with reason code "(Insufficient resources for Port-channel member)".	
Condition: This issue will occur when a member interface of a port-channel in the following scenarios: multiple shutdown/no shutdown of the interface or multiple remove/insert of a breakout interface	

Defect ID: DEFECT000560117	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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 Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000560221	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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 Group: Monitoring
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000560342	
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: show ip route showing the output partially for the withdrawn ipv4 route in VDX 6940 144S	
Condition: run show ip route	

Defect ID: DEFECT000560450	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: 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 Group: Management
Reported In Release: NOS6.0.1	Technology: 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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000560557	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: While restoring configuration from a file by executing 'copy <file> running-config, if it is aborted in between using Ctrl+C command, an exception might be reported on a console for a process ConfigCmd. Exception is seen at the time of process exit and will not impact switch functionality.	
Condition: Aborting 'copy <file> running-config' command in between using Ctrl+C command.	
Workaround: Avoid aborting 'copy <file> running-config' command.	
Recovery: No recovery is needed as this process will be spawned again for next copy command	

Defect ID: DEFECT000560623	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Software Installation & Upgrade
Symptom: On firmware downgrade from nos6.0.1 to nos6.0.0 or NOS5.0.1c or previous NOS version, syslog-server related configuration may not be retained.	
Condition: This would occur only if same IP address has been configured for syslog server on both mgmt-vrf and default-vrf as below: Example: logging syslog-server 10.20.53.134 use-vrf mgmt-vrf logging syslog-server 10.20.53.134 use-vrf default-vrf	
Workaround: Take configuration backup & then remove the entry for the syslog server pointing to default-vrf.	
Recovery: Add configuration of syslog server pointing to mgmt-vrf.	

Defect ID: DEFECT000560644	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000560738	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: IP Addressing
Symptom: Transient IPv6 Ping failure between VDX 6940 144S and Spirent Host	

Closed with code changes for Network OS v7.0.0

Condition: Executing IPv6 ping

Defect ID: DEFECT000560826	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: 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 Group: Security
Reported In Release: NOS6.0.0	Technology: 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 Group: Security
Reported In Release: NOS6.0.1	Technology: 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 Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000560855	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS5.0.1	Technology: IPv6 Multicast VLAN Traffic Reduction
Symptom: When RSPAN MLD packets enter VCS and MLD is disabled on the switch, the packets are trapped to MLD daemon if TRILL TTL is decremented to 0, and also flooded.	
Condition: When Virtual Fabric is disabled, RSPAN MLD packets are flooded.	
Workaround: Enable Virtual Fabric or Disable RSPAN.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000560889	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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 Group: Management
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000560932	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: Configurations which are not supported on 100 Mb interface are getting supported on 100 Mb port-channel	
Condition: Port-channel is configured with 100 MB speed.	
Recovery: Remove configuration on port-channel related to Layer 2 and Layer 3 protocols set with 100MB speed	

Defect ID: DEFECT000560990	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: 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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000561018	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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: 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: DEFECT000561095	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Software Installation & Upgrade
Symptom: Cluster wide firmware download status is reported as "Failed" in CLI session when downgrading from 6.0.1 to older releases, even if the download actually passed. There is no loss of functionality.	
Condition: This incorrect status may be displayed in CLI session when a cluster wide firmware download is attempted from 6.0.1 to any prior releases.	
Recovery: No loss of functionality is observed, other than the incorrect status message in CLI session. So no recovery is required.	

Defect ID: DEFECT000561112	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Unexpected reload of system can occur when more than supported 128 VRF BGP instances are created.	
Condition: Configure beyond 128 VRF instances	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000561120	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.1	Technology: Security Vulnerability
Symptom: Request to enhance the behavior to restrict traffic to TCP port 8008 towards the switch CPU.	
Condition: Under all conditions.	
Workaround: Configure access-list to restrict traffic towards this port	

Defect ID: DEFECT000561135	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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 Group: Management
Reported In Release: NOS6.0.1	Technology: 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 Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000561283	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: 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: DEFECT000561285	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Port Mirroring
Symptom: MAPS RASLOGs for Insertion/Removal of SFP are logged when the remote link is reset due to line card power-off	
Condition: MAPS needs to be enabled.	

Defect ID: DEFECT000561308	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
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: DEFECT000561517	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Software Installation & Upgrade
Symptom: Switch undergoes unexpected reload during firmware upgrade if inband VIP is configured and Switch is in Fabric Cluster (FC) mode.	
Condition: Inband VIP is configured and Switch is in FC mode	
Workaround: Remove inband VIP configuration if Switch is in FC mode before upgrade and then add back after upgrade.	

Defect ID: DEFECT000561632	
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: Can encounter an unexpected reload if use range command to remove 64 or more port-channels. Ex- sw0(config)# no interface Port-channel 64-125	
Condition: Configure 64 or more Port Channels on a single RBridge and remove them all at once.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000561734	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: Source MAC address may not be learnt on the profiled port	
Condition: access profile with access vlan 1 if applied on the interface and switch is upgraded using ISSU	
Recovery: Use no port-profile-port followed by port-profile-port on the impacted interface.	

Defect ID: DEFECT000561737	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: Source MAC address may not be learnt on the profiled port.	
Condition: When trunk profile is applied on the interface and switch is upgraded using ISSU.	
Recovery: no port-profile-port followed by port-profile-port on the impacted interface.	

Defect ID: DEFECT000561738	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.1	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: show mac-address-table port-profile may show less number of source MAC's than existing.	
Condition: show mac-address-table port-profile is executed after upgrading the switch	

Defect ID: DEFECT000562010	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: Configuration Fundamentals
Symptom: In rare circumstances, the last_login file on the switch may become very large and cause other stability issues on the switch due to low disk space.	
Condition: This has been seen with some script based logins or when using REST API to login.	

Defect ID: DEFECT000562050	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS4.0.1	Technology: Configuration Fundamentals
Symptom: Persisting configuration might fail due to low disk space and thereby eventually causing switch to go for unexpected reload.	
Condition: High disk usage and uncleared log files might cause disk to run out of space.	
Workaround: Monitor for High Disk Utilization RAS log and cleanup unwanted files periodically.	
Recovery: Switch comes up with default configuration and if it was part of Logical Cluster, will recover the configuration from the cluster.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000562053	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS4.0.1	Technology: OpenStack Integration
Symptom: Error message printed while applying the command "switchport" to interface configuration.	
Condition: "Committing to Underlying Data Store failed" error message printed while applying the command "switchport" to interface configuration during interface flap.	

Defect ID: DEFECT000562306	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: Configuration Fundamentals
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: 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: DEFECT000562609	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.2	Technology: 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: DEFECT000562703	
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: VDX switch ports could go "admin-down" state when removing/re-inserting cables.	
Condition: Some times a port can go Admin Down when the cable is physically removed from the port.	
Workaround: Operator should execute "shut" then "no shut" noscli command to bring the port back online.	
Recovery: Operator should execute "shut" then "no shut" noscli command to bring the port back online.	

Closed with code changes for Network OS v7.0.0

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: DEFECT000563056	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: NOS5.0.1	Technology: QoS - Quality of Service
Symptom: VDX 6740 may log errors indicating remaining system memory is below 100MB. Switch may undergo unexpected reload after some time.	
Condition: Switch is configured with FC & FCoE ports in a VCS cluster	

Defect ID: DEFECT000563290	
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: Sntpwalk 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: 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: DEFECT000563574	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS4.1.3	Technology: Logical Chassis
Symptom: If multiple interfaces are configured with the same DHCP Server Relay IP and one of them having local-subnet-broadcast address for that Server, clients may not be able to obtain an IP address from the DHCP Server after reboot.	
Condition: If the switch get rebooted with configuration of multiple VE interfaces configured with the same DHCP server Relay IP address and on that if any of them having local-subnet-broadcast address for that server, then we will see this issue.	
Workaround: Use the Server IP instead of local-subnet-broadcast address for the client VE interfaces.	
Recovery: The issue will be recovered by either doing anyone of the option: - Re-configuring the "dhcp relay" on the client VE interfaces (or) - Doing chassis disable and enable, if we have multiple client interfaces in problematic state.	

Closed with code changes for Network OS v7.0.0

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: DEFECT000563769	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS4.1.2	Technology: IPv4 Multicast Routing
Symptom: Under heavy scale of multicast configuration causing OLIST to be updated often, the multicast daemon may consume more memory than expected, thus impacting the switch stability.	
Condition: With high number of link-flaps, causing OLIST for a group to be updated frequently.	

Defect ID: DEFECT000563958	
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 few iterations of config/unconfig snmp community/v3users with IPv4/IPv6 ACLs.	
Condition: when switch is configured with SNMP IP ACLs on the Community/v3users.	

Defect ID: DEFECT000563969	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS5.0.2	Technology: OpenStack Integration
Symptom: The CLI "vcenter <name> discovery ignore-delete-all-response always " will be enabled by default when vcenter name is configured.	
Condition: When you delete the port-profile in vcenter, then vcenter sends a message to VDX to delete that "auto-port-profile" from the config. However, with this setting, when u delete from vcenter, it will not delete from the VDX.	

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	

Closed with code changes for Network OS v7.0.0

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: DEFECT000564629	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS5.0.2	Technology: Metro VCS
Symptom: On reload of one of the VTEP switches, certain traffic flows will be impacted.	
Condition: When the tunnelled VLAN is part of a non-default spanning tree instance (e.g. pvst,rpvt).	

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: DEFECT000565453	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.0	Technology: VCS Fabric
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 Network OS v7.0.0

Defect ID: DEFECT000565588	
Technical Severity: Medium	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 neighbor detects authentication failure after the VDX boots up.	
Condition: OSPF authentication is configured on a VDX switch and is reloaded.	

Defect ID: DEFECT000565659	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: 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	

Defect ID: DEFECT000566084	
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: Line cards on a 8770-4 or 8770-8 chassis could report faulty.	
Condition: A duplicate MAC address in a VCS cluster could increase the memory usage and make the line cards on 8770-4 or 8770-8 chassis report faulty.	

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: 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: 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.	

Closed with code changes for Network OS v7.0.0

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: DEFECT000567283	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.2	Technology: Configuration Fundamentals
Symptom: Under stress conditions, if the running-config is repeatedly run either via CLI, BNA, Netconf, the switch may encounter unexpected memory growth & undergo a reload.	
Condition: When repeatedly the running-config is requested either via CLI, BNA OR Netconf calls.	

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: 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: 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.	

Closed with code changes for Network OS v7.0.0

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: DEFECT000569138	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS5.0.2	Technology: Syslog
Symptom: SEC-1206 messages indicating last login attempt are truncated when received and logged in a syslog server	
Condition: Logging of the SEC-1206 messages indicating last login attempt	

Defect ID: DEFECT000569413	
Technical Severity: Medium	Probability: Medium

Closed with code changes for Network OS v7.0.0

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: 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: To recover from this state, it is recommended to flap individual port-channel members on the other linecards, one by one (non-disruptive). In addition, please note that if the firmware with the fix is installed on the switch using the ISSU method, the above described workaround should still be applied in order to remove the incorrect hardware programming.	

Defect ID: DEFECT000570402	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS5.0.1	Technology: IPv4 Multicast Routing
Symptom: Under certain conditions, the multicast video streams going over a shared tree might encounter duplication of packets, thus impacting the video quality momentarily.	
Condition: When SPT-Threshold is set to infinity.	

Closed with code changes for Network OS v7.0.0

Workaround: Set the SPT threshold to 1.

Closed with code changes for Network OS v7.0.0

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: 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: DEFECT000570803	
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: On the VDX6740 only 60 port-channels are supported instead of 64 port-channels.	
Condition: When trying to configure 64 Port-channels on VDX 6740.	

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: 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	

Closed with code changes for Network OS v7.0.0

Closed with code changes for Network OS v7.0.0

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: 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: Low
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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000573250	
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: Customer can experience disk full issue.	
Condition: Core files in "/" and "/core_files" where we don't have any file compression mechanism. So eventually VDX can run out of disk space due to uncompressed core files in those dir.	
Recovery: Manually remove core files from "/" and "/core_files" if there is any.	

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: 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: 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.	

Closed with code changes for Network OS v7.0.0

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 interface	
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: 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: 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: DEFECT000574207	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VXLAN - Virtual Extensible LAN
Symptom: Tunnel terminated packets have IPv6 hop limit values that are same as the underlay outer header hop limit.	
Condition: For VxLAN traffic terminated at the Switch VTEP.	

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'	

Closed with code changes for Network OS v7.0.0

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: 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 Rbridge) 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: DEFECT000575735	
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: After HA failover SNMP daemon is restarted.	
Condition: SNMP daemon restarted after HA failover, since vrfname passed from SNMP end was invalid.	

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.	

Closed with code changes for Network OS v7.0.0

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: 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: DEFECT000577036	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: Raslog and Email are not generated when the threshold is hit for security violations.	
Condition: Configure VDX switches to generate Raslog and Email when the threshold is hit for security violations.	

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: 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.	

Closed with code changes for Network OS v7.0.0

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.	
<pre>snmp-server user admin groupname snmpadmin ... rbridge-id 41 snmp-server user admin groupname snmpadmin ...</pre>	

Defect ID: DEFECT000577854	
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: 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: 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 :	
<ol style="list-style-type: none"> 1. Apply the port-profile-port 2. After hafailover/ISSU upgrade 	
Recovery: Remove the port-profile-port and add it back	

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	

Network OS v7.0.1 Release Notes v5.0

Closed with code changes for Network OS v7.0.0

Recovery: Reload the receiver DR

Defect ID: DEFECT000577944	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: The previous operational statistics will carry over when an event-handler is deactivated and then the same event-handler is reactivated on the same R-bridge. The stale operational statistics will be seen on reactivation.	
Condition: 1) An activated event-handler on a given R-bridge must have been triggered and the action run at least once. 2) Deactivation and reactivation of the same event-handler on the same R-bridge will temporarily show stale operational statistics until the event-handler is triggered.	

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: DEFECT000578236	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VLAN - Virtual LAN
Symptom: Resource check fails for few vlans during configuring on ports.	
Condition: Due to internal database cleanup issue, stale entries are present in software. This will prevent programming hardware for few vlans.	

Defect ID: DEFECT000578472	
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: Not able to unconfigure using "no map fport interface fcoe <RB/SLOT/PORT-PORT>".	
Condition: Trying to unconfigure VF-Ports in the N-Port mapping using range command.	

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.	

Closed with code changes for Network OS v7.0.0

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.	

Defect ID: DEFECT000580464	
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: The VDX switch unexpectedly reloads and gives the following warning message on the console: "AF4 get handle error, maybe full, please check" "cbr2_lpm_handle_is_in_used() af4_hdl_0x0_error"	
Condition: Exceeding the TCAM scale for IP HOST routes.	
Recovery: Reduce the IP HOST routes to within the maximum supported scale.	

Defect ID: DEFECT000581189	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: In some rare case, when IPv6 host is configured as trap recipient for SNMP v1/v2c. Traps may not receive in the receiver.	
Condition: Configure IPv6 address as trap recipient for SNMP v1/2c followed by a switch reload.	
Recovery: Reconfigure the host with same ipv6 address.	

Defect ID: DEFECT000581603	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: NTP - Network Time Protocol
Symptom: NTP clients are not able to sync with NTP servers. "show ntp status" will always show active NTP server is "LOCL"	
Condition: After ISSU upgrade the NTP clients are not able to sync with NTP server.	

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.	

Closed with code changes for Network OS v7.0.0

Defect ID: DEFECT000582026	
Technical Severity: Critical	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Some L3 traffic flows experience flooding	
Condition: Where the IP Address assignments across VRFs is not randomized then some entries may experience hash collisions and not get programmed in hardware.	
Workaround: Ensure that IP Address assignments are not contiguous across VRFs and is more randomized.	

Defect ID: DEFECT000582065	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: MAC address table is not in sync across the rbridges. Auto recovery fails.	
Condition: Issue can be seen with ARP scale configuration	

Defect ID: DEFECT000582759	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VPN
Reported In Release: NOS7.0.0	Technology: EVPN - Ethernet VPN
Symptom: Unexpected SW exception may be seen in environments when the same IP subnet exists in VRFs where the gap in the VRF numbers is multiples of 64. (eg, VRF 64 and VRF128.)	
Condition: When the same subnets are seen on two or more VRFs and those VRFs have VRF number with (n * 64) gaps.	
Workaround: Avoid using the VRFs having the (n* 64) gaps when the same subnets have to be on different VRFs.	

Defect ID: DEFECT000583239	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: RMON - Remote Network Monitoring
Symptom: In Chassis, when any host is configured as trap recipient then SNMP traps are received with active IP as source IP address.	
Condition: This issue is observed only in chassis based switches like 8770, when any host is configured as trap recipient.	

Defect ID: DEFECT000585195	
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: A flow match with ICMPv4 type of value 0x0 and/or ICMPv4 code of value 0x0 is treated as "dont care" field in the TCAM entry.	
Condition: A flow with non-zero value of ICMPv4 type and/or ICMPv4 code is hitting an entry programmed with match for ICMPv4 type of value 0x0 and/or ICMPv4 code of value 0x0.	

Closed without code changes for Network OS v7.0.1

Closed without code changes for Network OS v7.0.1

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

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

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

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

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

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

Defect ID: DEFECT000583324	Technical Severity: High
-----------------------------------	---------------------------------

Closed without code changes for Network OS v7.0.1

Reason Code: Not Reproducible	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: FCoE - Fibre Channel over Ethernet
Symptom: ISL fails to come up due to Trunking Error	
Condition: When port is enabled between VDX8770 LC48x10G and VDX 6940 4x10G breakout interfaces	
Recovery: Issue shut followed by no shut on the port.	

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

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

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

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

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

Closed without code changes for Network OS v7.0.1

Workaround: Use the vcs-rbridge-context NETCONF RPC to set an rbridge filter prior to issuing the NETCONF get-config request. This will limit the get-config results to the rbridge context set and reduce memory usage of the Dcmd process.

Closed without code changes for Network OS v7.0.0

Closed without code changes for Network OS v7.0.0

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

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

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

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

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

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

Closed without code changes for Network OS v7.0.0

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

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

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

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

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

Closed without code changes for Network OS v7.0.0

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

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

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

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

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

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

Closed without code changes for Network OS v7.0.0

Condition: If customer is planning to use OUI based port security on ports.

Defect ID: DEFECT000529660	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
Symptom: Web Element Manager session gets logged out before 30 minutes.	
Condition: Web Element Manager session will be logged out before 30 minutes in the following conditions:	
<ul style="list-style-type: none"> 1) If there is an unauthorized request sent to switch (or) 2) Internal bad server error 	

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

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

Defect ID: DEFECT000546734	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.1	Technology: RA Guard - Router Advertisement Guard
Symptom: Router Advertisement packets will be accepted even though RA Guard is configured on LAG and vLAG interfaces.	
Condition: <ul style="list-style-type: none"> 1) Configure Port Channel 2) Enable RA Guard 3) Router Advertisement packets will still be accepted 	
Workaround: Shut/no shut Port Channel	

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

Closed without code changes for Network OS v7.0.0

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

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

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

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

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

Closed without code changes for Network OS v7.0.0

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

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

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

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

Defect ID: DEFECT000557518	Technical Severity: High
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Virtual Fabrics
Symptom: In virtual-fabric enabled mode, port-profile-port configuration on an interface can take upto 5 minutes only the first time.	
Condition: If the system is in virtual-fabric enabled mode, first time configuration of "port-profile-port" or "port-profile-port domain <domainName>" , provided the port-profile-domain has "switchport trunk allowed vlan all" config	
Workaround: No workaround	

Closed without code changes for Network OS v7.0.0

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

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

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

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

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

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000559589	Technical Severity: High
Reason Code: Not Reproducible	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Hardware Monitoring
Symptom: 10Gb port may be stuck in offline state.	
Condition: After power cycle or reboot of VDX switch, 10Gb port may be stuck in offline state.	
Recovery: Shut/no shut the port.	

Defect ID: DEFECT000559631	Technical Severity: High
Reason Code: Will Not Fix	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: FCoE login does not work	
Condition: On changing the LAG config from local SAN to remote SAN on LAG.	

Defect ID: DEFECT000559741	Technical Severity: High
Reason Code: Not Reproducible	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Access Gateway
Symptom: FCoE logins won't come through LAG.	
Condition: Changing LAG from Local SAN to Remote SAN.	

Defect ID: DEFECT000559806	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: "show openflow flow" does not display Flowids in sorted order.	
Condition: User tries to view flow-mods installed in the switch.	

Defect ID: DEFECT000559907	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: VXLAN - Virtual Extensible LAN
Symptom: Copying configuration from file takes long time.	
Condition: The configuration file contains overlay-gateway with more than 1000 "map vlan" and 500 "site" configurations.	

Defect ID: DEFECT000560127	Technical Severity: High
Reason Code: Feature/Function Not Supported	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS5.0.1	Technology: OpenStack Integration
Symptom: Traffic loss can be observed between linecards going from one LC to another.	
Condition: Traffic loss can be observed between linecards [Ex: LC # 2] with another linecard [LC #1] being installed.	

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000560280	Technical Severity: High
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: Multi-VRF
Symptom: Under conditions of VRF scale and operation, sometimes Address Family may not be instantiated on creating a VRF	
Condition: Configuring a VRF in complex network topologies	

Defect ID: DEFECT000560607	Technical Severity: Medium
Reason Code: Not Reproducible	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP client binding does not happen on one of the VDX 6740s in the cluster, when clients are configured on a tagged vlan. DHCP Relay statistics show that discover packets are not received at the relay.	
Condition: DHCP Clients configured on tagged vlan.	

Defect ID: DEFECT000560620	Technical Severity: Medium
Reason Code: Not Reproducible	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: IP Addressing
Symptom: Ping fail between VDX 6940 and VDX 8770.	
Condition: Executing ping command	
Workaround: Atleast one L3 interface needs to be enabled in a VRF for a route to be added in the kernel or stack	

Defect ID: DEFECT000560681	Technical Severity: High
Reason Code: Design Limitation	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: FCoE Host in a Local Logical SAN may not be able to see the FC Target when VDX 6740 where the FC Target is connected to, is dynamically changed from the Remote Logical SAN to the Local Logical SAN configuration.	
Condition: When a VDX6740 is reconfigured from a Remote Logical SAN to a Local Logical SAN, FC Target connected through the VDX6740 will become un-accessible in VCS.	
Workaround: Reboot the VDX6740 that has been moving from Remote Logical SAN to Local Logical SAN.	
Recovery: To prevent hitting this issue, user can do delete fcf-group, change VDX6740 from AG mode to non-AG mode, and reboot VDX6740. Then proceed on modifying the fabric-map configuration from Remote Logical SAN to Local Logical SAN.	

Defect ID: DEFECT000560802	Technical Severity: Medium
Reason Code: Already Fixed in Release	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: The switch may accept the new fibrechannel or FCoE logins even if there is Duplicate WWN detection. This fabric login policy for Duplicate WWN is not supported in NOS6.0.0.	
Condition: When the firmware downgrade from NOS6.0.1 to NOS6.0.0 happens with the non-default fabric login policy i.e new-login for DuplicateWWN is configured on the switch.	
Workaround: move the fabric login policy to old-login for Duplicate WWN (default mode) and then downgrade the switch to NOS6.0.0	

Network OS v7.0.1 Release Notes v5.0

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000560931	Technical Severity: High
Reason Code: Not Reproducible	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: Delay in OSPF session establishment in VDX6940 144S on IP address delete and reconfigure.	
Condition: Running OSPF after IP address reconfiguration	

Defect ID: DEFECT000561037	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: Openflow related hardware resources may not get de-allocated in a clean manner resulting in inconsistent behavior in data-path forwarding. New configurations may also fail to get programmed and failures will not be reported back to controller.	
Condition: If ISSU was performed with controller driven flows installed in the system.	
Workaround: Controller needs to be disconnected first, clear all installed flows using "clear openflow all" command and then trigger ISSU.	

Defect ID: DEFECT000561046	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: Internal Server Error is returned as the status even though the intended functionality works fine when the rbridge-id is changed through REST request.	
Condition: When the rbridge-id is modified through REST request.	
Workaround: Modify the rbridge-id through CLI.	

Defect ID: DEFECT000561179	Technical Severity: High
Reason Code: Will Not Fix	Probability: Low
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: NOS6.0.1	Technology: QoS - Quality of Service
Symptom: Modular QoS CLI (MQC) policy config will be lost on a port in downgrade from NOS6.0.1 to NOS5.0.1 if MQC has shaper config on lossless priority.	
Condition: Per port scheduler configuration (i.e MQC policy having scheduler and shaper configuration applied in egress direction on a port).	
Workaround: Before downgrading to NOS5.0.1, make sure that active MQC policy should not have shaper config on lossless priority.	
Recovery: Remove shaper config lossless priority in MQC policy.	

Defect ID: DEFECT000561210	Technical Severity: High
Reason Code: Design Limitation	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: Unsupported features are getting configured on 100 MB interface if the link comes up with 100 MB using autonegotiation.	
Please check the NOS 6.0.1 Release Notes for features that are supported with 100MB	
Condition: When speed is configured as an 'auto' on an interface which is connected to an 100 MB peer link.	

Closed without code changes for Network OS v7.0.0

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000561274	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: Static Routing (IPv4)
Symptom: Under conditions of scale (512 VRFs), few routes may go missing in VDX 6940-144S after an ISSU upgrade.	
Condition: ISSU in VDX 6940 144S	

Defect ID: DEFECT000561304	Technical Severity: High
Reason Code: Not Reproducible	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.1	Technology: VCS Fabric
Symptom: VDX 6940 platforms in FC mode may undergo an unexpected reload during firmware upgrade or downgrade operations if there is security related configuration - tacacs server,radius sever ip, aaa	
Condition: Switch is in FC mode, Security Configuration (tacacs server,radius sever ip,aaa) is present and Switch undergoes a firmware upgrade/downgrade.	
Workaround: Remove “ tacacs server,radius sever ip,aaa” on all the vcs nodes.	
Recovery: Powercycle will recover the Switch.	

Defect ID: DEFECT000561506	Technical Severity: High
Reason Code: Already Fixed in Release	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Configuration Fundamentals
Symptom: VDX6940 and VDX6740 may take longer to boot up in 6.0.1 than in 6.0.0. (the delay is about 1 minute with default configuration).	
Condition: When the reload command is executed.	

Defect ID: DEFECT000561605	Technical Severity: High
Reason Code: Already Fixed in Release	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS4.1.3	Technology: NTP - Network Time Protocol
Symptom: tsd module termination and VDX abrupt reload	
Condition: Fail to reach NTP server due to network reachability issue	
Recovery: Fix NTP server reachability issue.	

Defect ID: DEFECT000561706	Technical Severity: High
Reason Code: Already Fixed in Release	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.1	Technology: VLAN - Virtual LAN
Symptom: If static MAC is configured for mulicast MAC address, the packets destined to the MAC address will be flooded to all member ports in VLAN	
Condition: Configuring a static muticast MAC address will result in this issue.	

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000561713	Technical Severity: High
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: When snmpwalk operation is done on TCP MIB (RFC 4022), the operation may become very slow and may experience timeouts.	
Condition: This issue is seen when snmpwalk is done only on TCP MIB.	
Workaround: The snmpwalk should be done with a timeout set to at least 3 seconds. This will help in avoiding the timeout during the snmpwalk operation on the TCP MIB.	

Defect ID: DEFECT000562896	Technical Severity: Medium
Reason Code: Feature/Function Not Supported	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS5.0.1	Technology: Logical Chassis
Symptom: "no fabric isl enable" and "no fabric trunk enable" configurations on ports 0/33 to 0/48 got reverted back to "fabric isl enable" and "fabric trunk enable"	
Condition: Upgrade from no4.1.2a1 to nos5.0.1d	
Workaround: 1. Before upgrade, save the running config to a file 2. Upgrade to 5.0.1d 3. Copy default to startup 4. Replay saved config from the file	
Recovery: After upgrade, manually configure "fabric isl enable" and "fabric trunk enable" on ports 0/33 to 0/48	

Defect ID: DEFECT000562941	Technical Severity: Medium
Reason Code: Feature/Function Not Supported	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.2	Technology: AAA - Authentication, Authorization, and Accounting
Symptom: when upgrading from 4.1.2a to 5.0.2 with tacacs-server config	
Condition: This happens only upgrades from 4.x to 5.x with tacacs-server configuration	
Workaround: copy the config before upgrade, upgrade with default-config and replay the config	

Defect ID: DEFECT000563273	Technical Severity: High
Reason Code: Feature/Function Not Supported	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS4.1.3	Technology: Configuration Fundamentals
Symptom: High CPU usage is seen on 6730 platform.	
Condition: High memory usage due to memory leak which eventually result in CPU stuck at 100%.	
Recovery: 1. Login into Switch as root 2. Capture the "top" output 3. Run: echo 10240 > /proc/sys/vm/min_free_kbytes 4. Issue command: sysctl -q -p /etc/sysctl.conf	

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000563673	Technical Severity: High
Reason Code: Feature/Function Not Supported	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS3.0.1	Technology: Configuration Fundamentals
Symptom: We will observe unexpected reload of VDX 6720 platform switch.	
Condition: In 6720 platform switch, when packet size greater than 1536 is received on the management port, then we will observe this issue.	
Workaround: Setting MTU size of eth0 as 9512 will avoid this issue. But it should be set after every reload of switch, as MTU change is not persisted after reload.	

Defect ID: DEFECT000564101	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.0	Technology: Security Vulnerability
Symptom: Using restrict_ssh script on VDX 6740x platforms will not do a GOS sync for the SSHd config file.	
Condition: Using restrict_ssh on VDX 6740s platforms will have a GOS sync issue.	

Defect ID: DEFECT000564347	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: NOS5.0.2	Technology: Rate Limiting and Shaping
Symptom: After upgrade to 5.0.2a-bld-02, we may observe some traffic forwarding issue with 1Gigabit connections on 8770 platform	
Condition: When we have 8770 platform with 1Gigabit connections, we will observe this issue.	

Defect ID: DEFECT000564498	Technical Severity: Medium
Reason Code: Not Reproducible	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS4.1.3	Technology: Port Mirroring
Symptom: "show interface status" command shows incorrect status for internal VDX ports in the switch.	
Condition: If we have internal embedded ports in the VDX switch, then we will observe this issue.	

Defect ID: DEFECT000565415	Technical Severity: Medium
Reason Code: Will Not Fix	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: Added new enums in Fruclass to indicate the slotnames correctly 87 sfm(12), 88 lineCard(13), 89 managementModule(14) 90 }	
Condition: The slotnames are not insync with CLI slotnames	

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000565954	Technical Severity: Medium
Reason Code: Not Reproducible	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: Principal node of a logical cluster will not reach 'ready' state to accept any configuration commands. When any configuration command is issued, it displays error as 'Cluster formation is in progress'.	
Condition: Cluster reformation is triggered due to node joining or leaving the fabric or due to firmware upgrade. Cluster formation might not complete due to underlying communication issues.	
Recovery: Remove the node from the cluster either by issuing reload command or by issuing 'chassis disable' command. This will trigger some other node in the cluster to become principal. Once the cluster is stable, bring back this switch to the cluster by issuing 'chassis enable'.	

Defect ID: DEFECT000569750	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.2	Technology: Software Installation & Upgrade
Symptom: When upgrading VDX switches from 5.0.2x to 6.0.1x, certain rules configured for the RBAC feature, would change to a different set of commands.	
Condition: When upgrading switches from 5.0.x to 6.0.x & RBAC rules are configured.	
Workaround: It is recommended to take a backup of the rules before the upgrade to 6.x from 5.x so that it can be restored after upgrade.	

Defect ID: DEFECT000577381	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS4.1.3	Technology: VRRPv3 - Virtual Router Redundancy Protocol Version 3
Symptom: Under rare conditions, certain hosts lose IP connectivity with selected few devices when hosts connected on VDX6710/20/30 configured as Layer-3 gateways for the host.	
Condition: Usually when the set of MAC addresses are toggled between one interface to another on the switch.	
Recovery: Clearing the ARP on the VDX switches restores the connectivity.	

Defect ID: DEFECT000578967	Technical Severity: High
Reason Code: Will Not Fix	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.2	Technology: PIM - Protocol-Independent Multicast
Symptom: On rare occasion, 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	

Closed without code changes for Network OS v7.0.0

Defect ID: DEFECT000585841	Technical Severity: High
Reason Code: Already Fixed in Release	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.2	Technology: VXLAN - Virtual Extensible LAN
Symptom: Customer uses pVLAG between VDX 6740 and third party vendor switch for underlay of L2 extension tunnel. After pVLAG member PO's go through fail over, the L2 extension traffic drops for upto 10 seconds.	
Condition: The issue occurs when the primary member port-channel of a port-channel group goes through a fail over. Under such circumstances, a 10 second tunnel traffic is lost.	

Known Issues for Network OS v7.0.1

This section lists open software defects with Critical, High, and Medium Technical Severity as of May 25, 2016 in Network OS v7.0.1.

Defect ID: DEFECT000590114	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.1	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: If user configures two AMPP port-profiles, one of them configured with access VLAN x and other configured with trunk VLAN x, Then ,In that case , It will not be shown as conflicting in "show port-profile name <pp1-name> name <pp2-name> validate" command output.	
Condition: When user creates 2 port-profiles, one port-profile with access VLAN x and other port-profile with trunk VLAN x and executes "show port-profile name <pp1-name> name <pp2-name> validate" CLI.	

Defect ID: DEFECT000591616	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: CLI - Command Line Interface
Symptom: Switch goes for an unexpected reload with the REST request.	
Condition: When the switch is pounded with the REST requests from multiple concurrent sessions simultaneously and continuously over a long period of time.	
Workaround: As far as possible, send REST requests to configure the switch from one session only. Multiple sessions can be used for retrieving information from the switch with GET requests.	

Defect ID: DEFECT000592879	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: Configuration Fundamentals
Symptom: After LC power on/off in VDX8770, uplink interfaces from the LC are missing on show track summary output.	
Condition: When Link State Tracking (LST) configuration is present on a linecard, after slot power off/on the uplink configuration will be lost.	
Workaround: Uplinks need to be reconfigured again after slot power on.	

Defect ID: DEFECT000594276	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: Under a high scale of VCS nodes, the configuration applied for a range of interfaces across the VDX nodes may cause principal node to encounter an unexpected reload.	
Condition: When issuing a configuration command under an interface range in a large cluster (32+ nodes)	
Workaround: Avoid using interface range option in large clusters (32+ nodes) & instead configure the interfaces individually.	

Defect ID: DEFECT000594793	
Technical Severity: Medium	Probability: Medium

Known Issues for Network OS v7.0.1

Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: Software Installation & Upgrade
Symptom: System may display: "qman_recovery_exit_local: DEBUG: the FQID 516 has dest_wq as chaqman_recovery_exit_local: DEBUG: the WQ lengths for pool channel of portal 1 on cpu1 are: 0:0:0:0:0:0:0"	
Condition: This bug appears when partitions are switched with heavy traffic.	
Recovery: Reboot the system.	

Defect ID: DEFECT000595709	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.1	Technology: Logical Chassis
Symptom: System reloads on VDX8770.	
Condition: This occurs with 512 or more VRRP sessions enabled and "debug vrrp packets" is turned on.	
Workaround: "debug vrrp packets" should not be turned on in a scaled environment.	

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

Defect ID: DEFECT000596480	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: Configuration Fundamentals
Symptom: On execution of CLI "track remove all" complete Link State Tracking (LST) configuration should get removed from a port. In case of port-channel interface protocol daemon is not clearing the LST configuration hence it is displayed in output of show command.	
Condition: Execution of "track remove all" CLI for a port-channel interface for which Link State Tracking (LST) configuration is present.	
Workaround: As a workaround user can remove the configuration one by one by executing respective 'no' CLIs.	

Defect ID: DEFECT000596868	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: CLI - Command Line Interface
Symptom: The global MTU value cannot be deleted through REST API.	
Condition: Issue happens when the user tries to delete the global mtu using the DELETE request through the REST interface.	
Workaround: Using the PATCH request with the default value as a work around. The effect of this is same as deleting the config.	

Defect ID: DEFECT000598965	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Local configuration related to global configuration may not restore on "config snapshot restore".	

Known Issues for Network OS v7.0.1

Condition: Customer using snapshot feature may see issues when running "attached rbridge-id add 1" missing from running-config after "vcs config snapshot restore rbridge-id <rb-id> snapshot-id <snapshot-id>".
Workaround: Customer should configure the missing configurations again.

Defect ID: DEFECT000598972	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: CLI - Command Line Interface
Symptom: Switch might go for an unexpected reload when any configuration update is performed on a range of interfaces.	
Condition: On a large cluster with scaled up configurations, performing any configuration on a range of interfaces by entering into interface range sub-mode might cause switch to run out of memory and thereby causing it to reload.	
Workaround: Required configuration update can be made on individual interfaces one at a time instead of performing it on a range of interfaces. Configuration update on multiple interfaces can still be performed by using comma (,) as separators instead of hyphen (-) when specifying the range. For ex, to shutdown interfaces 1 to 5, use "interface te 1,2,3,4,5" instead of "interface te 1-5".	

Defect ID: DEFECT000599203	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: The SNMP IPV4 traps may not be received through in-band interface.	
Condition: The SNMP traps may not be received through in-band interface after upgrade from 6.0.1 to 7.0.1.	
Workaround: Configure source-interface in the SNMP host / v3host recipients.	

Defect ID: DEFECT000599289	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.0.1	Technology: ACLs - Access Control Lists
Symptom: Applying Access Control List (ACL's) with 12K rules on management interface takes more than 3 minutes to enforce it.	
Condition: When Access Control List (ACL's) is configured with 12K rules.	

Defect ID: DEFECT000599778	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.0.1	Technology: TACACS & TACACS+
Symptom: LDAP/RADIUS/TACACS+ server configurations are not displayed in the same order in which they were added.	
Condition: <ol style="list-style-type: none"> 1. Configure multiple TACACS+/RADIUS/LDAP servers(max 5) 2. Remove few server entries 3. Add those servers entries back 	
Workaround: Remove all Server entries and configure those servers back in the desired order.	

Defect ID: DEFECT000600022	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.1	Technology: Metro VCS

Known Issues for Network OS v7.0.1

Symptom: When VDX 8770 is in chassis-disabled state, the far-end 100 GbE link partners using QSFP28 optics may see intermittent link flaps. After VDX 8770 is chassis-enabled, there is a low probability that the 100 GbE port may not come online.
Condition: When VDX 8770 is in chassis-disabled state, the far-end 100 GbE link partners using QSFP28 optics may see intermittent link flaps. After VDX 8770 is chassis-enabled, there is a low probability that the 100 GbE port may not come online.
Recovery: Execute "shut" on 100 GbE link partner port connected to VDX 8770 to stop the port from flapping intermittently. After the VDX 8770 is chassis-enabled, execute "no shut" on the 100 GbE link partner to re-enable the port.

Defect ID: DEFECT000600057	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.1	Technology: Logical Chassis
Symptom: Switch might not rejoin the cluster when reloaded using 'fastboot' command.	
Condition: Reloading switch using 'fastboot' command on VDX6940 and VDX6740 platforms when SW1 partition is active might lead to this issue.	
Workaround: Reload the switch using 'reload' command which is more graceful way of reloading.	
Recovery: Bring the switch which failed to join the cluster to default configuration using command 'copy default-config startup-config'. On reload, switch rejoins the cluster and regains older configuration.	

Defect ID: DEFECT000600066	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP IPv4 Relay forwarded DISCOVER packet is not getting forwarded through remote leaf node in BGP-EVPN IP Fabric.	
Condition: While deploying DHCP Relay in BGP-EVPN IP Fabric.	
Recovery: Disable "conversational-arp".	

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

Defect ID: DEFECT000600185	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS7.0.1	Technology: OpenStack Integration
Symptom: When VDX-8770 is in chassis-disable state, the "show media" command will not show 100 GbE ports.	
Condition: When VDX-8770 is in chassis-disable state, the "show media" command will not show 100 GbE ports.	
Workaround: After the chassis is enabled using "chassis enable" command, "show media" will show the 100 GbE ports.	

Known Issues for Network OS v7.0.1

Recovery: After the chassis is enabled using "chassis enable" command, "show media" will show the 100 GbE ports.

Defect ID: DEFECT000600197	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: "show running-config overlay-gateway <name> vlan <vlan number>" throws error "% No entries found " even when VLAN is present.	
Condition: This happens only when a filter is specified after "overlay-gateway <name>". Otherwise command works fine when no filter is specified.	
Workaround: Instead of using the filter, use the " include <string>" for filters like following: "show running-config overlay-gateway <name> include "vlan <vlan-number>"	

Defect ID: DEFECT000600230	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.1	Technology: IP Fabric
Symptom: "show running-config rbridge-id evpn-instance <vni -name> vni add <vni-range>" throws an error message.	
Condition: Customer doing show running configuration with VNI range in EVPN instance.	
Workaround: Use the following command: "show running-config rbridge-id evpn-instance vni add" .	

Defect ID: DEFECT000600377	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: The SNMP walk may fail and SNMPV3 trap may not be received for the user configured under RBridge.	
Condition: The SNMP walk may fail and SNMPV3 trap may not be received only for the SNMPV3 user configured under rbridge after upgrade from 7.0.0 to 7.0.1.	
Recovery: Reconfigure the user under RBridge after the successful upgrade from 7.0.0 to 7.0.1.	

Defect ID: DEFECT000600385	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VLAN - Virtual LAN
Symptom: Duplicate ARP entries are observed.	
Condition: This can happen after an ISSU upgrade and a new IP address is allocated via DHCP for a connected host.	
Workaround: Execute "clear arp ip <IP address>" for the old IP address of host.	
Recovery: Execute "clear arp ip <IP address>" for the old IP address of host.	

Defect ID: DEFECT000600591	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: CLI - Command Line Interface
Symptom: Logs are dumped on the screen, when there is a read failure on SFPs connected to the port.	

Known Issues for Network OS v7.0.1

Condition: Accessing information about the SFPs inserted in the ports.

Recovery: Disable the port and re-enable it.

Known Issues for Network OS v7.0.0

This section lists open software defects with Critical, High, and Medium Technical Severity as of May 3, 2016 in Network OS v7.0.0.

Defect ID: DEFECT000408109	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS3.0.0	Technology: VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: On VRRP session for ISL interface track command should not be allowed as this affects priority	
Condition: VRRP Track command for VRRP session on ISL port.	
Workaround: Dont enable VRRP track on ISL port.	

Defect ID: DEFECT000420768	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS3.0.0	Technology: OpenStack Integration
Symptom: User may not be able to see the current session privileges when he is authenticated through AAA using "Show users".	
Condition: Show users command didn't display user's role when user is authenticated via RADIUS authentication.	

Defect ID: DEFECT000451282	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS4.0.0	Technology: User Accounts & Passwords
Symptom: Changing the encryption level of a given username to level 7 with 'no service password-encryption' shows success RASLOG, but does not take effect in the config.	
Condition: When service password encryption is turned off, try to change existing users encryption level to 7.	

Defect ID: DEFECT000456601	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS4.0.0	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: "max-metric" config is cleared and not retained after some add/delete operations.	
Condition: On max metric configuration, if clear ospf is executed and ospf networks are removed and added again, max-metric config is cleared and is not retained.	
Workaround: Reconfigure max-metric after clear operation.	

Defect ID: DEFECT000458128	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS4.0.0	Technology: Syslog
Symptom: Netconf login information not available in auditlog	
Condition: Applications logging into the switch using Netconf.	

Known Issues for Network OS v7.0.0

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000471058	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS3.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: When SNMP with IPv6 is configured & SNMP Manager sends 2 get-requests divided in separate fragments, then "ICMPv6 Destination Unreachable" is returned from the switch.	
Condition: IPv6 & SNMP configuration	

Defect ID: DEFECT000490740	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS4.0.1	Technology: Configuration Fundamentals
Symptom: Setting 'deny ip any any' ACL does not prevent telnet access via management port as expected.	
Condition: Setting following ACL does not prevent telnet access via management port as expected: sw0# show access-list ip ACL001 in ip access-list ACL001 on Management 1/0 at Ingress (From User) seq 10 deny ip any any (Active) seq 20 permit tcp any any (Active)	

Defect ID: DEFECT000510114	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS4.1.2	Technology: VLAN - Virtual LAN
Symptom: In VDX 6740, when we have different load balancing scheme configured on the port channel, we see unexpected results with respect to load balance.	
Condition: If we have different load balancing schemes applied on VDX 6740, the latest configured value will take effect on the switch.	
Workaround: Use the same LB scheme for all PO in VDX 6740.	
Recovery: Re-configure the same LB scheme wherever required.	

Defect ID: DEFECT000517329	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Nexthop change using outbound route-map is not allowed for EBGp neighbor connection.	
Condition: When Route-map with set-nexthop is used as outbound policy for bgp neighbor.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000519785	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.0	Technology: AAA - Authentication, Authorization, and Accounting
Symptom: When "aaa authentication" command is tried with atleast one parameter same as previous command(example:aaa authentication radius local --> aaa authentication tacacs+ local), Customer may not be able to set the correct aaa mode.	
Condition: The authentication mode with primary & secondary sources of authentication cannot be updated to a configuration containing only the primary source and configuration with primary & secondary sources of authentication, the primary mode alone cannot be modified.	
Workaround: When "aaa authentication" command is tried with at least one parameter same as previous command(example: aaa authentication radius local --> aaa authentication tacacs+ local). Need to remove existing configuration and then configure the required configuration. example : (config)# do show running-config aaa authentication aaa authentication login radius local (config)# no aaa authentication login (config)# aaa authentication login tacacs+ local (config)# do show running-config aaa authentication aaa authentication login tacacs+ local	

Defect ID: DEFECT000521573	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.0	Technology: ACLs - Access Control Lists
Symptom: When copying a file to the running config that contains MAC ACL or IPv4 ACL commands, if some of the keywords are abbreviated, the commands may be treated as invalid, or the commands may execute more slowly than if they were not abbreviated.	
Condition: Abbreviating any of the keywords "mac access-list", "ip access-list", "seq", or, if "seq" is not present, "permit", "deny", or "hard-drop", would result in the slowness issue rather than the invalid command issue. The slowness issue would become more pronounced if the file being copied contains a large number of MAC ACL or IPv4 ACL commands.	
Workaround: Avoid using abbreviated keywords in files being copied to the running config.	

Defect ID: DEFECT000523640	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.0	Technology: VMWare
Symptom: Traffic is not working through ports connected to servers after doing chassis disable/enable on the VDX.	
Condition: The issue is seen only after chassis disable followed by enable.	
Recovery: After chassis enable, Do a "shut" followed by "no shut" on the Port-Channels/Physical Interface connected to servers or do a "no port-profile-port" followed by "profile-port" on the Port-Channels/Physical interface.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000524630	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: During connection collision, switch is not closing the recent connection request if previous connection in established state.	
Condition: Second connection attempt is made after BGP peer is established.	
Recovery: The second connection request will get established automatically in case of collision.	

Defect ID: DEFECT000525575	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS4.1.3	Technology: PIM - Protocol-Independent Multicast
Symptom: If there are multiple subnets configured on an interface, PIM First Hop Router mechanism only works for the subnet with the highest IP address.	
Condition: PIM with multiple subnets on an interface.	

Defect ID: DEFECT000528408	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: New BGP connection is not accepted under following conditions.	
Condition: When the remote BGP identifier is greater than the local BGP identifier and there is a connection collision.	

Defect ID: DEFECT000529345	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
Symptom: User cannot see all the ISL port information in the port connectivity tab.	
Condition: Attempt to view ISL connection details from port connectivity tab.	

Defect ID: DEFECT000530965	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: When the Helper Router exits helper mode on a given network segment, it should re-originate its LSAs based on the current state of its adjacency to the restarting router over the segment which is not happening.	
Condition: The helper router, on the GR mode exit, is supposed to originate the router/network LSAs based on the current adjacency state. This behavior is not yet implemented.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000533582	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS6.0.0	Technology: AAA - Authentication, Authorization, and Accounting
Symptom: No uniformity in alignment and display format in the output of show access-list ip across interfaces	
Condition: Customers using the show access-list and expecting same display format across different interfaces.	

Defect ID: DEFECT000533953	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: When a SNMP Set request is attempted on the object ospfVirtIfRtrDeadInterval, this object can be set to a value greater than 65535.	
Condition: This issue is seen only on doing a SNMP Set request on the MIB object ospfVirtIfRtrDeadInterval.	

Defect ID: DEFECT000535440	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: 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: DEFECT000535663	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.1	Technology: VXLAN - Virtual Extensible LAN
Symptom: VDX unexpectedly reloads when it connects to an NSX controller when a high scale of MAC addresses are injected over tunnel. It continues to reload unless booted with factory defaults.	
Condition: VDX connects to NSX controller and discovers more than 70000 MACs.	
Recovery: Boot with factory default configuration.	

Defect ID: DEFECT000536442	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
Symptom: Management IP access list with large number of rules takes around 1 min to enforce the policy on management interface.	
Condition: Management IP access list with large number of rules.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000538887	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS4.1.2	Technology: Syslog
<p>Symptom: Under normal operations, following false-positive raslog message may be seen:</p> <p style="padding-left: 40px;">"[FW-3120], ..., WARNING, ... Interface<>,IFGViolation Errors, is above high boundary(High=100, Low=5). Current value is <> Error(s)/minute."</p> <p style="padding-left: 40px;">There is no functional impact.</p>	
Condition: When there are no interface errors incrementing in the output of "show interface".	

Defect ID: DEFECT000540858	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS4.1.0	Technology: NETCONF - Network Configuration Protocol
<p>Symptom: Configuration and get config via NetConf are slow compared to CLI and sometime show high percentage of CPU utilization</p>	
Condition: Applicable for all netconf commands.	

Defect ID: DEFECT000541060	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
<p>Symptom: With conditional default-origination, default route should be generated only when the route-map matching prefix is present in the IP routing table. When irrespective of whether route-map matching prefix is present in the IP routing table or not, default route is originated to the neighbor.</p>	
Condition: On configuring 'default-originate' with route-map in neighbor command	

Defect ID: DEFECT000541202	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
<p>Symptom: When customer configures high and low threshold values and the actual value is below the low threshold, a fabric watch raslog is displayed showing in between high and low threshold instead of below threshold.</p>	
Condition: While using fabric watch module to monitor memory usage.	

Defect ID: DEFECT000541449	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.0	Technology: BGP4+ - IPv6 Border Gateway Protocol
<p>Symptom: Peer group configuration is not accepting update-source of IPv6 address</p>	
Condition: Peer group configuration with update-source of IPv6 address	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000543579	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.1	Technology: IP Addressing
Symptom: Switch may reload due to low on memory when DHCP relay address and DHCP gateway CLIs are repeatedly used to configure and unconfigure.	
Condition: Repeated configure/un-configure of DHCP relay address and DHCP gateway CLI may lead to unexpected switch reload due to increased memory consumption.	

Defect ID: DEFECT000545603	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Configuration Fundamentals
Symptom: Extended ACL with permit TCP rule does not block ICMP frames.	
Condition: Extended ACL with permit TCP rule does not block ICMP frames.	
Workaround: Add another rule to deny ICMP frames.	
Recovery: Add another rule to deny ICMP frames.	

Defect ID: DEFECT000546702	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.1	Technology: ACLs - Access Control Lists
Symptom: When user tries to login with wrong credentials using default-vrf, debug messages are seen on console.	
Condition: When user tries to login with wrong credentials using default-vrf, debug messages are seen on console.	

Defect ID: DEFECT000547747	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: VMWare
Symptom: Event notification is not received when port-groups with special characters are created on a vSwitch . As a result , the corresponding port-profile will not be applied to that interface.	
Condition: Use of special characters for port-groups on a vSwitch.	

Defect ID: DEFECT000548727	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: VMWare
Symptom: Port-Groups having the same name (regular , VMKernel , Distributed Port Groups) can get deleted when the vCenter user/administrator deletes any of the above port-groups sharing the same name,	
Condition: This scenario can happen when Port-groups belonging to different class(regular , VMKernel, Distributed Port Groups) use the same name within a data-center.	
Workaround: Avoid using the same name for different class of port-groups. vCenter's recommended names with prefixes like dvpg , pg , VMKernel is good to follow.	
Recovery: Rename the port-groups from different class to use different names.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000548981	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: Configuration Fundamentals
Symptom: Access to the switch via management port (out of band) for all IPv6 protocols fails	
Condition: This issue is observed when a IPv6 ACL is configured on the management interface. For example "ipv6 permit any any"	
Workaround: Configure another matching entry in the ACL for permit based on protocol. Example "permit tcp any any"	

Defect ID: DEFECT000550658	
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: Momentary traffic loss is observed on NSX controller managed tunnels during VCS cluster formation.	
Condition: VCS has an overlay-gateway configuration with two or more RBridges attached and tunnel configurations are discovered by connecting to NSX controller. One of the RBridges attached to overlay-gateway is rebooted or upgraded via ISSU protocol. Traffic loss is observed on tunnels when such rebooted or upgraded RBridge rejoins VCS cluster.	

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

Defect ID: DEFECT000552520	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.1	Technology: Static Routing (IPv4)
Symptom: Memory leak observed with repeated addition/deletion of VRFs using an automated script	
Condition: Adding and deleting VRFs repetitively	
Workaround: Remove routes before deleting the VRF	

Defect ID: DEFECT000553066	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: Under certain condition, multiple "Invalid InterfaceId." log messages will be seen on the console.	
Condition: The symptom will be seen when a node joins the fabric and the fabric starts to rebuild.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000553426	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS6.0.1	Technology: Logical Chassis
Symptom: "This command is not supported on this product." message displayed	
Condition: When trying to execute a dpod command on a VDX 8770 (principal node in VCS cluster) for a non-existent rbridge-id.	

Defect ID: DEFECT000554319	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: Switch does not generate a ColdStart trap on the VE interface configured in mgmt-vrf.	
Condition: When switch is configured with VE interface in mgmt-vrf, then we will observe this issue.	

Defect ID: DEFECT000554351	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: xSTP - Spanning Tree Protocols
Symptom: On the VDX8770 Platform, executing certain Spanning Tree show commands may display the following warning message: % Warning: Output Incomplete, VCS is in transient state This is a cosmetic issue and these messages can be safely ignored. There is no impact to functionality.	
Condition: 1) Configure MSTP 2) Issue "show spanning-tree brief" or "show spanning-tree mst detail" CLI command	

Defect ID: DEFECT000554472	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: Configuration Fundamentals
Symptom: When booting logical chassis with its default config, the principal Rbridge ends up with "system-monitor MM threshold" "down-threshold 0" while non-principal Rbridges end up with "system-monitor MM threshold" "down-threshold 2".	
Condition: This issue can be observed when using multiple VDX6740 switches in a logical chassis and booting from the default config.	
Recovery: Reconfigure the "system-monitor MM threshold" to make it sync.	

Defect ID: DEFECT000555059	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.2	Technology: SNMP - Simple Network Management Protocol
Symptom: SNMP walk should display Interfaces even when linecard is powered-off.	
Condition: When a linecard is powered off, the SNMP walk does not display and reference the interface. Fixed it as part of this defect.	

Known Issues for Network OS v7.0.0

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 unreachable are always sent' displayed in the configuration even when disabled in the configuration	
Condition: Default ICMP unreachable is not set	

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

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

Defect ID: DEFECT000558082	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: RMON - Remote Network Monitoring
Symptom: when MAPS rules are triggered and MAPS is configured to generate e-mails, they are not generated	
Condition: MAPS has to be enabled with one of the default policies and e-mail action has to be enabled on the switch. Switch IP address needs to be configured as an IPv6 address.	
Workaround: Use IPv4 address for Switch IP if MAPs email action is desired	

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

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000558937	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.1	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: Sometimes, MAC addresses are shown twice in the output of 'show mac port-profile' command.	
Condition: 'show mac port-profile' command is issued	
Workaround: Re-issue to command to see a refreshed display	

Defect ID: DEFECT000560092	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS6.0.1	Technology: OpenStack Integration
Symptom: Unable to execute connector command due to config in error state.	
Condition: User did "copy default-config startup-config" on part of the cluster while the rest of the cluster is segmented. When cluster joined, it dynamically replayed back the connector config received from new principal at rejoin. This caused its connector configuration being partially saved to database on the local switch.	
Recovery: Recovery steps: <ol style="list-style-type: none"> 1. From principal switch, remove this secondary node from cluster by using "no vcs logical-chassis enable rbridge-id XX default-config" with XX being the secondary node RBridgeid. 2. On the secondary node, do "write erase" to clean up its config. 3. Restore the secondary switch vcs mode, vcsid and RBridgeId and the necessary ISL and see it rejoin the cluster. 	

Defect ID: DEFECT000560160	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.1	Technology: Inband Management
Symptom: Pings to VE interfaces are slow after upgrade	
Condition: In rare upgrade scenarios.	
Recovery: Shut/No shut the affected interface where the ping is seen slow or have a reload.	

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: 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.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000562214	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS4.1.3	Technology: VLAN - Virtual LAN
Symptom: Source MAC may not get learnt on port channel configured with primary VLAN	
Condition: When Secondary VLANs, which are associated with other Primary VLAN are deleted	

Defect ID: DEFECT000562427	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS5.0.2	Technology: Logical Chassis
Symptom: When ISSU upgrade from 5.0.1b->5.0.2 with sflow collector config, does not allow to modify the existing configs	
Condition: ISSU upgrade with sflow configs	
Workaround: Remove the sflow config before the ISSU upgrade and the reconfig again	

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: DEFECT000562672	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.2	Technology: SNMP - Simple Network Management Protocol
Symptom: When we try to associate 6th unique IPv4/IPv6 ACL to SNMP community/v3 user, we observe DB sync issue between front and backend. And also we observe unexpected reload of SNMP daemon when we try to associate 6th unique IPv4/IPv6 ACL to SNMP community/user.	
Condition: When we associate 6th unique IPv4/IPv6 ACL to SNMP community/user.	

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.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000563295	
Technical Severity: Medium	Probability: Low
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: DEFECT000564701	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.1	Technology: Syslog
Symptom: syslog-server secure port option can not be used without using the use-vrf option.	
Condition: CLI usability issue where it shows use-vrf and secure options for syslog-server, where as secure option can be used along with use-vrf option only.secure port can not be set alone for syslog server.	

Defect ID: DEFECT000565277	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.2	Technology: Configuration Fundamentals
Symptom: Under rare condition, when ACL is applied under management interface, it may not take effect and may display "Generic Error" when trying to remove the access-list.	
Condition: When applied under management interface	

Defect ID: DEFECT000565913	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: Software Installation & Upgrade
Symptom: Nodes of same cluster after reboot of all the nodes form cluster islands.	
Condition: This can happen when some nodes of a topology come up first and form their own cluster, and do not join the cluster formed by the rest of the nodes.	

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

Defect ID: DEFECT000567346	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS5.0.2	Technology: SNMP - Simple Network Management Protocol
Symptom: dot1qTpFdbTable of Q-BRIDGE-MIB does not give any output when queried using a SNMP Get/Get-Next or Get-Bulk request.	
Condition: This issue is specific to dot1qTpFdbTable of Q-BRIDGE-MIB.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000569125	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: Hardware Monitoring
Symptom: On VDX2746 platform, NOSCLI command "show sfp" displays incorrect speed for 1G port as 10G port. For e.g. sw0# show sfp inc 44 0 44 SFP n/a 10G	
Condition: Plug in 1G media and run the NOSCLI cmd "show sfp"	

Defect ID: DEFECT000569319	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: VLAN - Virtual LAN
Symptom: The MAC's gets learned on native VLAN for the shortest duration during RSTP root rbridge toggle. There will be loop during RSTP root rbridge toggle and the macs are getting classified to native vlan causing the issue. The macs will be aged out after age time.	
Condition: The MAC's will be learned on native VLAN during RSTP toggle. There will be loop during RSTP root rbridge toggle and the macs are getting classified to native vlan causing the issue. The MAC's will be aged out after age time.	
Workaround: Clearing FDB table will recover the mac tables.	

Defect ID: DEFECT000569501	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.1	Technology: High Availability
Symptom: During ISSU, the original active partition on a LC may not boot up.	
Condition: It is because of a rare race condition during the boot up process.	
Recovery: run power-off and power-on on the LC	

Defect ID: DEFECT000570086	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: NTP - Network Time Protocol
Symptom: Removal of NTP server is failing in the cluster.	
Condition: In rare conditions when multiple NTP servers are configured, removal of NTP server is failing.	
Workaround: Delete the NTP server second time if it failed on first time.	

Defect ID: DEFECT000570284	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.1	Technology: xSTP - Spanning Tree Protocols
Symptom: Root port of MSTP Instances will become Master port after adding a VLAN to an existing instances after MSTP-I convergence	
Condition: Addition of VLAN to an already converged MSTP Instances.	
Workaround: Shutdown/no shutdown of MSTP	

Known Issues for Network OS v7.0.0

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: 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: DEFECT000573649	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: MAPS - Monitoring and Alerting Policy Suite
Symptom: Clock change on the switch is not getting reflected on switch.	
Condition: Clock change on the switch is not getting reflected on MAPS Dashboard.	

Defect ID: DEFECT000574438	
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 6940-144s 40 GbE ports 97 - 108 may flap after changing framesize and/or nframes using noscli "diag dport setargs" command.	
Condition: VDX 6940-144s 40 GbE ports 97 - 108 may flap after changing framesize and/or nframes noscli "diag dport setargs" command.	
Workaround: Shut/no shut the port if persistent link flapping occurs.	
Recovery: Shut/no shut the port if persistent link flapping occurs.	

Defect ID: DEFECT000574645	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.2	Technology: SNMP - Simple Network Management Protocol
Symptom: The "fruStatusChanged" trap may be received more than once while power ON the line cards in VDX8770-4 and VDX8770-8 platforms.	
Condition: The "fruStatusChanged" trap may be received more than once while power ON the line cards.	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000576391	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: The HA failover start trap may not be seen for the HA fail over event, for user defined VRF in VDX-6740 platforms.	
Condition: The HA failover start trap may not be seen for user defined VRF, for the the HA fail over event.	

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

Defect ID: 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: DEFECT000577800	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: MAPS - Monitoring and Alerting Policy Suite
Symptom: device connectivity config should be consistent on all the links in the port-channel	
Condition: port-channel members configured as different type NAS, iSCSI	
Workaround: Configure all members to be in same type.	

Known Issues for Network OS v7.0.0

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

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

Defect ID: DEFECT000578640	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: Few millisecond packet loss observed while collecting supportsave when multiple data stream with high line rate is running.	
Condition: Collecting supportsave when multiple stream of very high rate data is running.	

Defect ID: DEFECT000578730	
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: 'Message Generic Error' will be encountered when configuring "extend vlan add <vlan string>" command with large number of vlans on VDX6940 device.	
Condition: Adding vlans with <vlan string> exceeding 1000 characters for overlay gateway using layer2 extension.	
Workaround: Split "extend vlan add <vlan string>" into multiple commands so that <vlan string> does not exceed 1000 characters in a given command.	

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

Known Issues for Network OS v7.0.0

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

Defect ID: 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: DEFECT000579695	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: IP Addressing
Symptom: SNMP walk on the IP-FORWARD-MIB may throw an error	
Condition: SNMP operations on IP-FORWARD-MIB	

Defect ID: DEFECT000579835	
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: SNMP Set operation on Port-channel interface for the ipv6RouterAdvertLinkMTU object in ipv6RouterAdvertTable in RFC 4293 will fail. The SNMP Set operation will fail for other objects also in this table.	
Condition: The SNMP Set operation will fail only on port-channel interfaces.	
Workaround: In order to configure the ipv6RouterAdvertLinkMTU, the corresponding CLI "ipv6 nd mtu" needs to be used.	

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	

Known Issues for Network OS v7.0.0

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

Defect ID: DEFECT000581045	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: The network user might see PIM hello packets getting routed thus different networks becoming PIM neighbors with each other.	
Condition: When switch is acting as a L2 PO and connecting two L3 POs then this issue appears.	

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

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

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

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000581284	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: Introducing a check to verify everytime if port-channel count has exceeded 4k or not will bring down the performance. It is already documented that 4k vlags are supported.	
Condition: User is allowed to configure more than 4k port-channels.	
Workaround: No workaround. Not applicable	

Defect ID: DEFECT000582010	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: ARP - Address Resolution Protocol
Symptom: Under rare conditions, some of the hosts may lost IP connectivity with the VDX switch acting as a layer-3 gateway.	
Condition: This would occur if the MAC to the IP association of a VDX learnt ARP changes. ie. For the same IP address, the MAC changes from say Mac1 to Mac2.	
Recovery: "clear arp no-refresh" would clean the ARP table and recover from the problem state.	

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

Defect ID: DEFECT000582737	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: Syslog
Symptom: Syslog is not working in default/non-default-vrf via inband when same syslog server IP address is configured in mgmt-vrf and default/non-default-vrf.	
Condition: This is a limitation in open source software Syslog-NG. Every remote server configured should have a unique IP address.	
Workaround: Configure unique syslog server IP address for each VRF.	

Defect ID: DEFECT000582797	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: Configuration Fundamentals
Symptom: Switch does not responds to any user inputs.	
Condition: Using 'PUT' request to update BGP configuration instead of using 'PATCH' request (when router BGP has neighbor associated to a peer group) causes switch not to respond to further user commands.	

Known Issues for Network OS v7.0.0

Workaround: Use PATCH request instead of PUT to update any configuration as documented in Extreme Network OS REST API Guide.

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: DEFECT000583324	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: FCoE - Fibre Channel over Ethernet
Symptom: ISL fails to come up due to Trunking Error	
Condition: When port is enabled between VDX8770 LC48x10G and VDX 6940 4x10G breakout interfaces	
Recovery: Issue shut followed by no shut on the port.	

Defect ID: DEFECT000583349	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: Licensing
Symptom: In VDX6940-144S, 100G mode configuration replay fails when executing "copy <file> running-config" because DPOD license has not been reserved.	
Condition: This issue will happen only if the DPOD license has not been reserved.	
Workaround: Manually reserve the DPOD license and then run "copy <file> running-config"	
Recovery: Manually reserve the DPOD license and then run "copy <file> running-config"	

Defect ID: DEFECT000583626	
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: Error message is not thrown, if more than the maximum number of SNMP V3 users configured with both global and local configurations combined together.	
Condition: If the SNMP V3 users configured more than the maximum number of supported, for both global and local configurations combined together then the error message is not thrown.	

Defect ID: DEFECT000584099	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: xSTP - Spanning Tree Protocols
Symptom: On a MSTP root port, traffic recovery takes nearly 60 seconds	
Condition: MSTP root port is brought UP within a minute after shutting down.	

Defect ID: DEFECT000584212	
Technical Severity: High	Probability: Low

Known Issues for Network OS v7.0.0

Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VLAN - Virtual LAN
Symptom: An edge port may take a long time (up to 23 seconds) to start forwarding traffic.	
Condition: Edge port is toggled when connected to 10G interface.	

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: DEFECT000584254	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS7.0.0	Technology: Hardware Monitoring
Symptom: The output of noscli "show media" has transceiver field 0.	
Condition: This transceiver byte 0 should not affect any functionality of the optic. This field can be program to either 0 or non zero value to indicate the speed of the optic.	
Workaround: Ignore the 0 value of transceiver field. It should not affect the functionality of the optic	

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

Defect ID: DEFECT000584534	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: BGP4+ - IPv6 Border Gateway Protocol
Symptom: A transport-vlan has been configured with multiple ctags on a port, however only a single ctag is seen in the command, "show vlan brief"	
Condition: This issue is seen when the following configuration is done in the following order. 1. configuring a vni of a vlan in the evpn-instance. 2. configuring the corresponding vlan as a transport-vlan (TVlan).	
To avoid the issue, do the configuration in the order #2 & #1.	

Defect ID: DEFECT000584634	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.2	Technology: VCS Fabric
Symptom: 40G port will notice frequent online and offline events if one side is configured as breakout and the other side is not	
Condition: Failure to issue breakout on a QSFP 40G port, which is supposed to work in 4X10G mode.	

Known Issues for Network OS v7.0.0

Known Issues for Network OS v7.0.0

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

Defect ID: DEFECT000585008	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: BGP4+ - IPv6 Border Gateway Protocol
Symptom: when config apply error happens, User doesnt know which line of the config had the issue.	
Condition: Upon config replay on VDX devices	

Defect ID: DEFECT000585015	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: traffic may flood for the non-profiled macs even if the global knob is disabled.	
Condition: 'no allow non-profiled-macs' is configured.	
Workaround: configure and remove 'allow non-profiled-macs' again.	

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

Defect ID: DEFECT000585352	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.2	Technology: VXLAN - Virtual Extensible LAN
Symptom: The VxLAN traffic drops when the underlay VLAG interface links go through the state change. It is further observed that VTEP is not learning the MAC addresses of hosts located across the VXLAN tunnel.	
Condition: The problem occurs when VLAG is VXLAN underlay network AND Loopback IP is configured as VTEP IP.	
Workaround: VLAG as undelay and Loopback IP as VTEP IP is not supported. Hence, please use VRRP(-E) IP as VTEP IP when underlay network is comprised of VLAG(s).	

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000585392	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: CLI - Command Line Interface
Symptom: BNA unable to discover NOS switch, when NOS switch is connected to a FCR.	
Condition: 1. Have a setup with VCS connected to FCR 2. Install BNA 14.0.1 3. Shift to IP tab. 4. In Discovery dialog, add IP of one of the VCS switches and click ok. 5. Observe the device is not discovered and shows "Discovery Failed" message.	

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

Defect ID: DEFECT000585516	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: CLI - Command Line Interface
Symptom: LST configured on LAG span on multiple rbridge allowed and hard to track upstream/downstream.	
Condition: On configuration of LST on a port channel interface across multiple rbridge.	
Workaround: Avoid enabling LST across multiple rbridges.	

Defect ID: DEFECT000585634	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: VCS Fabric
Symptom: On consecutive reboot certain ports may be shown administratively down.	
Condition: Multiple switch reboots.	
Recovery: Enable the port to administratively up.	

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

Known Issues for Network OS v7.0.0

Defect ID: DEFECT000585818	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: Configuration Fundamentals
Symptom: REST GET request fails to retrieve an empty leaf	
Condition: When 'cli-show-no' annotation is used along with the cli-run-template to customize the display of an empty leaf in the 'show running config', the same might not be retrieved through the REST GET request.	

Defect ID: DEFECT000585895	
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 and 6940-144S ports may not come online after converting from 40 GbE to breakout mode 4 x 10 GbE.	
Condition: VDX 6940 and 6940-144S ports may not come online after converting from 40 GbE to breakout mode 4 x 10 GbE.	
Workaround: Execute noscli shut / no shut commands	
Recovery: Execute noscli shut / no shut commands	

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

Defect ID: DEFECT000585960	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: VCS Fabric
Symptom: 40G Interface is administratively (or) protocol down with FFDC raslogs	
Condition: Admin operations on 40G Interface.	
Recovery: Shut/no-shut both interfaces on either side of the link	

Known Issues for Network OS v7.0.0

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

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

Defect ID: DEFECT000586125	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: ISL between VDX8770 and VDX6940 could be flapping continuously.	
Condition: One side of the link connects to VDX8770 from VDX 6940	
Workaround: Toggle the flapping ISL link by doing "shutdown" and "no shutdown" of the ports.	

Defect ID: DEFECT000586178	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: Non-existent port-channel shows up in "show fcoe interface ethernet"	
Condition: <ol style="list-style-type: none"> 1. Create a port-channel 2. Add members to it and make it fcoe-provisioned 3. Delete the port-channel 	
Workaround: Remove FCOE provisioning from port-channel before deleting it	

Defect ID: DEFECT000586230	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.0	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: System goes for unexpected reboot	
Condition: More than 8000 profiled and non-profiled mac addresses are learnt on profiled port and all these flows are moving across the ports.	

Known Issues for Network OS v7.0.0

Known Issues for Network OS v7.0.0

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

Defect ID: DEFECT000586338	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: Logical Chassis
Symptom: IN VDX 6940-144S, link flap occurs on 40 G ISL ports if breakout configuration mis-matched with any adjacent ports.	
Condition: In VDX 6940-144S, a 40 G Port with breakout QSFP is not configured as breakout may cause adjacent 40 G port to flap, whereas its peer port is configured as 40G breakout.	
Workaround: In VDX 6940-144S, configure 40 G port as breakout if the peer port is configured as 40G breakout. After that, the link flap on the port will stop.	

Defect ID: DEFECT000586771	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: Configuration Fundamentals
Symptom: Secondary to primary association fails in a particular sequence.	
Condition: Cannot associate secondary vlan to primary vlan, once its type is changed from community to isolated or vice versa.	
Workaround: Delete secondary vlan, create, configure secondary type and associate again.	

Defect ID: DEFECT000586790	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: BGP4+ - IPv6 Border Gateway Protocol
Symptom: Using RBridge range configuration command, even after BGP VRF instance is deleted, configuration under BGP VRF instance is allowed and can cause BGP daemon termination, HA failover, and or reboot of the switch.	
Condition: Using Rbridge range configuration command, BGP VRF instance is removed using "no address-family vrf .." command. And immediately after, without exiting from the configuration mode, another command applicable under (obsolete) BGP VRF instance configuration mode is issued.	
Workaround: After removing the BGP VRF instance while using RBridge range command, exit the obsolete configuration mode using "top" , "end", or "exit" commands.	

Defect ID: DEFECT000586951	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.0.0	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCPv4 server OFFER packet is dropped by DHCPv4 relay when configured with subnet broadcast address.	
Condition: Running DHCP Relay in IP Fabric EVPN.	

Known Issues for Network OS v7.0.0

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

Defect ID: DEFECT000587566	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VXLAN - Virtual Extensible LAN
Symptom: Some VLANs that the user is expecting to be provisioned in a BGP EVPN instance get unprovisioned.	
Condition: The issue is seen when the user uses some common VNIs between BGP EVPN instance and some physical ports and then the user removed the last port from all the VLANs.	
Workaround: The user should avoid configuring common VNIs between BGP EVPN and physical ports.	
Recovery: Deleting those VLANs and configuring them back again, will fix the issue.	

Defect ID: DEFECT000587615	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: The SNMP V3 traps may not be received for the SNMP v3 host recipients configured under rbridge mode.	
Condition: The trap may not be received after upgrade from NOS6.0.1a to NOS7.0.0 with cold boot option	

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

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

Known Issues for Network OS v7.0.0

Recovery: If total number of port-channels configured are exceeding 70, delete few port-channels to reduce the total count to be less than 70.

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

Defect ID: DEFECT000587804	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.0	Technology: IP Fabric
Symptom: Even though there are no matching EVPN import route-targets configured under VRF, imported EVPN routes are present in BGP VRF table.	
Condition: EVPN import route-target(s) is/are removed while matching routes are present in BGP-EVPN and imported into BGP VRF table.	
Workaround: Issuing "clear bgp evpn neighbors all soft in" command should cleanup the routes which are still imported in BGP VRF instance after matching EVPN import route-targets are removed.	

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

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

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

Known Issues for Network OS v7.0.0

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

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