



March 2021

Extreme Network OS 7.2.0g for Extreme VDX devices Release Notes

9036239-04 Rev AA

Copyright Statement

© 2021, 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

Copyright Statement	2
DOCUMENT HISTORY.....	8
PREFACE.....	9
Contacting Extreme Technical Support	9
Document feedback	10
OVERVIEW	11
Important information – NOS 7.2.0e release	11
Hardware.....	11
New devices	11
New interface modules	11
Deprecated Hardware	11
Software Features	12
Deprecated Software Features	12
Software Upgrade	12
New Software Features for Network OS v7.2.0g	12
New Software Features for Network OS v7.2.0f	12
New Software Features for Network OS v7.2.0e	12
CLI Changes	13
New Commands for Network OS v7.2.0g.....	13
New Commands for Network OS v7.2.0f	13
New Commands for Network OS v7.2.0e.....	13
Modified Commands for Network OSv7.2.0g.....	13
Modified Commands for Network OSv7.2.0f	13
Modified Commands for Network OSv7.2.0e.....	13
Deprecated Commands for Network OS v7.2.0g	13
Deprecated Commands for Network OS v7.2.0f	13
Deprecated Commands for Network OS v7.2.0e	13
API Changes.....	14
Newly supported standards and RFCs.....	14
HARDWARE SUPPORT.....	15
Deprecated Devices.....	15
Supported devices	15

Extreme VDX 6940-144S	15
Extreme VDX 6940-36Q.....	15
Extreme VDX 6740	16
Extreme VDX 6740T.....	16
Extreme VDX 6740T-1G.....	16
Extreme VDX 8770-4 and VDX 8770-8	17
Supported power supplies	18
Supported Optics for NOS 7.2.0g	20
SOFTWARE UPGRADE AND DOWNGRADE	22
Image filenames	22
Upgrade/Downgrade considerations	22
Migration Path	23
Management IP connectivity	26
Firmware Installation	27
Downgrading to a Previous Release.....	27
Upgrade/downgrade with default configuration.....	31
Management Gateway IP changes.....	32
Management Services	33
Telnet, SSH and AAA VRF support.....	33
HTTP VRF support.....	33
NTP VRF support	33
SNMP- Community string maximum length increased to 64:.....	33
SNMP - Support for traps during hafailover:.....	33
SNMP-Trap Source IP support:.....	33
SNMP context based query:.....	34
SNMP MIB – VLAN update	34
SNMP Trap VRF Support.....	34
SNMP-Trap CLI.....	34
SNMP – IF MIB.....	34
Sflow VRF Support.....	34
Syslog VRF Support.....	35
Firmware download, Copy support, Copy config.....	35
Other Management Services.....	35
SCALABILITY AND INTEROPERABILITY	36

Scalability numbers	36
IP Fabric Scalability	41
HW Profile and Platform Specific Scale Numbers	42
Compatibility and Interoperability	47
IP Storage	47
FC/FCoE Storage	48
Network Adapters	49
ADDITIONAL CONSIDERATIONS	50
Limitations and Restrictions	50
Command Line Interface	50
Platform	51
Line cards	52
USB	52
Licensing	52
VCS	52
VCS Node Replacement	54
Logical Chassis	55
Extreme Trunks	55
Breakout Interfaces	55
Dual-personality Ports	56
1G Mode	56
vLAG	56
Virtual IP Address Support	56
Security, Management ACLs, Authentication, Authorization	57
SPAN & RSPAN	57
MAC Learning Considerations in VCS	58
PVLAN	58
UDLD	58
STP/DiST	58
IGMPv3 Snooping	59
Edge Loop Detection (ELD)	59
Long Distance ISL Ports	59
AMPP and Port-Profiles	60
vCenter	61

QoS.....	61
FCoE.....	62
FlexPorts.....	63
Fibre Channel	63
Access Gateway.....	63
IP Fabric.....	64
ND/RA.....	65
IPv4.....	65
BFD	65
VRRP	65
Fabric Virtual Gateway (FVG)	66
OSPFv3	66
BGP.....	66
Layer 2/Layer 3 Multicast.....	67
VRF	67
ACL	67
Policy-based Routing (PBR)	68
Inter-VRF Leaking (Static).....	68
DHCP IP Helper	68
Dynamic ARP Inspection (DAI)	68
DHCP-based Firmware Download (DAD – DHCP Automatic Deployment)	68
Link State Tracking	69
OpenFlow	69
Mac Port Based Authentication	70
Uplink Switch Support.....	70
Layer 2 and Layer 3 ISSU on VDX 6740x.....	71
Extreme Vyatta Controller (BVC).....	71
REST API	72
NetConf	72
VXLAN Gateway for VMware NSX.....	73
VF Extension using VxLAN	73
TCAM Profiles.....	74
Management VRF.....	74
Conversational MAC Learning.....	74

System level Flow based QoS.....	74
Port level Flowbased QoS.....	74
URPF.....	74
BGP Auto neighbor discovery.....	74
Non-trivial Merge.....	74
HA on TOR switches.....	75
Logical Chassis HA.....	75
Interoperability.....	75
MAPS.....	76
Maintenance Mode.....	76
BNA.....	76
Miscellaneous.....	76
DEFECTS.....	78
TSBs - Critical Issues to Consider Prior to Installing This Network OS Release.....	78
Network OS v7.2.0a Caveats.....	78
Network OS v7.2.0 Caveats.....	79
Closed with code changes for Network OS v7.2.0g.....	82
Closed with code changes for Network OS v7.2.0f.....	82
Closed with code changes for Network OS v7.2.0e.....	83
Closed with code changes for Network OS v7.2.0d.....	86
Closed with code changes for Network OS v7.2.0c.....	91
Closed with code changes for Network OS v7.2.0b.....	97
Closed with code changes for Network OS v7.2.0a.....	108
Closed with code changes for Network OS v7.2.0.....	117
Known Issues for Network OS v 7.2.0e.....	147

DOCUMENT HISTORY

Version	Summary of Changes	Publication Date
1.0	Initial Release Removed versions Network OS v7.2.0d and older	March 2021

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:

- Your Extreme Networks service contract number or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions 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

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

Important information – NOS 7.2.0e release

With the release of the NOS 7.2.0f software, the NOS 7.2.0e software and release notes have been removed from the Extreme Portal due to “FN-2020-456 VDX IPC-5024 fail messages and l2sysd and fab_vcscd unexpected reload” issue.

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

Optics Name	Extreme PN
10GE USR SFP+ Low temp	57-1000343-01
10GE SR SFP+ TAA	57-1000344-01
10GE SR SFP+ Low Temp	57-1000340-01
10GE LR SFP+ TAA	57-1000345-01
10GE LR SFP+ Low Temp	57-1000341-01
40G QSFP+ -> 4x10G LR	57-1000477-01

New interface modules

- None

Deprecated Hardware

- None

Software Features

For information about Network OS v7.2.0d and earlier releases, please refer to the [Network OS v7.2.0d Release Notes](#).

The following section lists new, modified, and deprecated software features for Network OS 7.2.0. For information about which platforms support these features, refer to the *Network OS Features and Standards support Matrix*.

Deprecated Software Features

- None.

Software Upgrade

ISSU (In Service Software Upgrade) from Network OS 7.1.0 to Network OS 7.2.0 is not supported.

New Software Features for Network OS v7.2.0g

- None

New Software Features for Network OS v7.2.0f

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

New Software Features for Network OS v7.2.0e

- 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 v7.2.0g

- None

New Commands for Network OS v7.2.0f

- None

New Commands for Network OS v7.2.0e

- None

Modified Commands for Network OS v7.2.0g

- None

Modified Commands for Network OS v7.2.0f

- None

Modified Commands for Network OS v7.2.0e

- None

Deprecated Commands for Network OS v7.2.0g

- None

Deprecated Commands for Network OS v7.2.0f

- None

Deprecated Commands for Network OS v7.2.0e

- None

API Changes

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

Newly supported standards and RFCs

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

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.

BGP Graceful Shutdown (GSHUT) Feature adheres to the below RFC and draft :

RFC 6198 - Requirements for the Graceful Shutdown of BGP Sessions draft-ietf-grow-bgp-gshut-03 - Graceful BGP session shutdown

HARDWARE SUPPORT

Deprecated Devices

Following VDX Switches are not supported starting Extreme Network OS v7.2.0:

- Extreme VDX 2741
- Extreme VDX 2746

Supported devices

Extreme Network OS v7.2.0 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 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 slots 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: LC48x1G line card provides up to 48 SFP/SFP-copper ports.
- 10 GbE: LC48x10G line card provides up to 48 SFP+ ports.
- 10 GbE-T: LC48x10GT line card provides up to 48 RJ-45 ports.
- 40 GbE: LC12x40G line card provides up to 12 x 40 GbE QSFP ports.
- 40 GbE: LC27x40G line card provides up to 27 x 40 GbE QSFP ports.
- 100 GbE: LC6x100G line card provides up to 6 x 100 GbE CFP2 ports.

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, NONPORTSIDEXHAUST	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
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-FFRU 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 NOS 7.2.0g

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 <https://cloud.kapostcontent.net/pub/a070d154-d6f1-400b-b2f0-3d039ae2f604/data-center-ethernet-optics-data-sheet?kui=Cc1YBpmqyfb2mDfw2vlq2g>.

Extreme-branded Top Level SKU	Description
10065	10/100/1000BASE-T SFP
10301	ASSY, SR SFP+ SHIPPING
10302	ASSY, LR SFP+ SHIPPING
10304	1m SFP+ Cable
10306	5m SFP+ Cable
10310	ZR SFP+ module
10051H	1000BASE-SX SFP, Hi
10052H	1000BASE-LX SFP, Hi
10070H	10/100/1000BASE-T SFP, Hi
100G-LR4-QSFP10KM	100G LR4 QSFP28 10km
100G-LR4-QSFP2KM	100G LR4 QSFP28 2km
100G-SR4-QSFP100M	100G SR4 QSFP28 100m
10G-AOC-SFP10M	10G AOC SFP+ 10m
10GB-BX10-D	10 GB, SINGLE FIBER SM, -D 10 KM
10GB-BX10-U	10 GB, SINGLE FIBER SM, -U 10 KM
10G-ER-SFP40KM-ET	10G ER SFP+ 40km Ext.Temp
10G-LR-SFP10KM-ET	10G LR SFP+ 10km Ext.Temp
10G-SR-SFP300M-ET	10G SR SFP+ 300m Ext.Temp
10G-USR-SFP100M	10G USR SFP+ 100m Hight Rx Sens
40G-AOC-QSFP100M	40G AOC QSFP+ 100m
40G-AOC-QSFP10M	40G AOC QSFP+ 10m
40G-AOC-QSFP20M	40G AOC QSFP+ 20m
40G-AOC-QSFP3M	40G AOC QSFP+ 3m
40G-AOC-QSFP5M	40G AOC QSFP+ 5m
40G-BDSR-QSFP150M	40G BiDi SR QSFP+ 150m
40G-DACA-QSFP1M	40G Active DAC QSFP+ 1m
40G-DACA-QSFP3M	40G Active DAC QSFP+ 3m
40G-DACA-QSFP4SFP1M	40G Active DAC QSFP+ to 4xSFP+ 1m
40G-DACA-QSFP4SFP5M	40G Active DAC QSFP+ to 4xSFP+ 5m
40G-DACA-QSFP5M	40G Active DAC QSFP+ 5m
40G-DACP-QSFP1M	40G Passive DAC QSFP+ 1m
40G-DACP-QSFP3M	40G Passive DAC QSFP+ 3m

Extreme-branded Top Level SKU	Description
40G-DACP-QSFP4SFP2M	40G Passive DAC QSFP+ to 4xSFP+ 2m
40G-DACP-QSFP4SFP3M	40G Passive DAC QSFP+ to 4xSFP+ 3m
40G-DACP-QSFP4SFP5M	40G Passive DAC QSFP+ to 4xSFP+ 5m
40G-DACP-QSFP5M	40G Passive DAC QSFP+ 5m
40G-DACP-QSFPZ5M	40G Passive DAC QSFP+ 0.5m
40G-ESR4-QSFP400M-NT	40G ESR4 QSFP+ 400m 10G-SR interop.
40G-LM4-QSFP160M	40G LM4 QSFP+ 160m 160m MMF. 1km SMF
40G-LR4-QSFP10KM	40G LR4 QSFP+ 10km
40G-SR4-QSFP150M	40G SR4 QSFP+ 150m
MGBIC-LC01-G	1GB SX MM, SFP, TAA

SOFTWARE UPGRADE AND DOWNGRADE

Image filenames

Download the following images from www.Extremenetworks.com

Image Filename	Description	Supported Device or Module
nos7.2.0g.tar.gz	Network OS v7.2.0g for UNIX	NA
nos7.2.0g.zip	Network OS v7.2.0g for Windows	NA
nos7.2.0g_all_mibs.tar.gz	Network OS v7.2.0gMIBS	NA
NOS_7.2.0g_Release_Notes	Network OS v7.2.0g Release Notes v1.0 (PDF)	NA
nos7.2.0g.md5	Network OS v7.2.0g MD5 Checksum	NA

Upgrade/Downgrade considerations

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

Migration Path

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

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

To From	6.0.1x	6.0.2x	7.0.0	7.0.1x	7.0.2x	7.1.0x	7.2.0a 7.2.0b	7.2.0c 7.2.0d 7.2.0e 7.2.0f 7.2.0g
6.0.1x	ISSU for upgrade; coldboot for downgrade.	ISSU	Coldboot	coldboot	coldboot	default-config	default-config	default-config
6.02x	coldboot	ISSU for upgrade; coldboot for downgrade.	coldboot	coldboot	coldboot	default-config	default-config	default-config
7.0.0	coldboot	coldboot	NA	ISSU*	coldboot	coldboot	default-config	default-config
7.0.1x	coldboot	coldboot	coldboot	ISSU for upgrade; coldboot for downgrade.	coldboot	coldboot	default-config	default-config

To	6.0.1x	6.0.2x	7.0.0	7.0.1x	7.0.2x	7.1.0x	7.2.0a 7.2.0b	7.2.0c 7.2.0d 7.2.0e 7.2.0f 7.2.0g
From								
7.0.2x	coldboot	coldboot	coldboot	coldboot	coldboot	coldboot	default-config	default-config
7.1.0x	default-config	default-config	coldboot	coldboot	coldboot	ISSU for upgrade; coldboot for downgrade.	coldboot	coldboot
7.2.0a 7.2.0b	default-config	default-config	default-config	default-config	default-config	coldboot	coldboot	coldboot
7.2.0c 7.2.0d 7.2.0e 7.2.0f 7.2.0g	default-config	default-config	default-config	default-config	default-config	coldboot	coldboot	ISSU

NOTES

1. ** CFP2 to QSFP28 conversion module (PN: 80-1008646-01) Version3 downgrade to any release prior to Network OS7.0.1 will cause CRC errors on the link.
2. Only Extreme Network Advisor (BNA) v14.2.1 (available separately) and above supports Network OS v7.2.0. It is required to first upgrade to BNA v14.2.1 and then upgrade switches to Network OS v7.2.0.
3. Before downgrading to lower releases, it is recommended to disable all new features that are not supported on lower releases by using the “no” version of the CLIs. Stray configurations left out before downgrade can cause undesired behavior.
4. While upgrading chassis-based system, under stress condition (e.g. due to excessive processing load on the processor), some linecards may become faulty during firmware download. To recover, run “power off <linecard>” followed by “power on <linecard>” command.
5. You must remove the IGMP snooping static mrouter configuration from all VLANs before upgrading or downgrading from or to the Network OS 6.0.2x release.
6. Firmware download from Network OS7.0.1a to Network OS6.x or Network OS5.x with default-config option needs AG mode disabled.
7. **Limitations:
 - a. After downgrading from Network OS v7.1.0 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 7.1.0, need to do shut/no shut to recover
 - c. In VDX 8770 platforms, After upgrade from 6.0.2 to 7.1.0 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 7.1.0 and above.
8. ISSU upgrade from NOS 7.2.0/NOS7.2.0a/NOS7.2.0ax to NOS 7.2.0b is not supported and blocked. Please refer “669062” defect below in NOS 7.2.0b closed defect list for further details.
9. ISSU upgrade from NOS 7.2.0b to NOS 7.2.0c is also not supported because of open defect.
10. ISSU upgrade is supported from NOS 7.2.0c to NOS 7.2.0d only.
11. To upgrade to NOS 7.2.0d from NOS 7.2.0b or earlier use coldboot/default-config.

Management IP connectivity

- In regards to SNMP, firmware downgrade from Network OS v7.1.0 to v7.0.x/v6.0.x/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 v7.1.0 and above to v7.0.x/v6.0.x/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 Network OS v7.1.0 and above from v7.0.x/v6.0.x/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 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.

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 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.

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 v7.1.0 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 Network OS 7.1.0 using the coldboot option. Principal node in a cluster should be last to be upgraded.
- After a node is upgraded, it will join the existing Network OS v7.2.0 cluster. Eventually, when all the nodes are upgraded, they will form one Network OS 7.2.0 VCS 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)
 - Execute a coldboot downgrade.

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: 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: 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 Network OS 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 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 LC mode, configuration is automatically saved and is persistent.

NOTE: FC mode is not supported in Network OS 7.1.0 and above.

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 hfailover:

Cpstatuschange trap will be triggered during hfailover with cpLastEvent as hfailoverstart and hfailoverdone to notify that hfailover is started and hfailover 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-server host ip-address <community-string> source-interface management ?
```

Possible completions:

```
chassis-ip Use chassis IP as source address
mm-ip      Use local MM IP as source address
```

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.

Network OS v7.2.0 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)	512000	512000	512000	512000
Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for VMware NSX	8000	N/A	8000	8000
Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for Virtual-Fabric Extension	120000	N/A	75000	75000
Maximum # of MAC-based Virtual-Fabric VLAN Classification per switch	256	1024	1000	1000
Maximum # of Classified Virtual Fabric VLANs per Trunk Interface	2000	4096	4096	4096
Maximum # of port profiles (AMPP)	1000	1,000	512	512
Maximum # of VLANS in port profiles	3500	4000	3500	3500
Maximum # of sites (tunnels) in Virtual-Fabric Extension	50	N/A	50	50
Maximum # of dot1q VLANs that can be attached on VxLAN GW for Virtual-Fabric Extension	4000	N/A	4000	4000
Maximum # of Virtual-Fabric (Service + Transport) VLANs that can be extended via Virtual-Fabric Extension	2000	N/A	4000	4000
Maximum # of dot1q VLANs + Virtual-Fabric VLANs enabled on edge-interfaces that can be attached to VxLAN GW and extended via Virtual-Fabric Extension	(2000+1000)	N/A	(2000+1000)	(2000+1000)
Max # of IGMP groups over Tunnels via Virtual-Fabric Extension	6000	N/A	6000	6000
Max # of BFD sessions over Virtual-Fabric Extension	10	N/A	10	10

Network OS v7.2.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
Tunnels				
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 IGMPv2v3 Snooping Interfaces supported	512	512	512	512
Learning rate for IGMP snooping (groups/second)	512	512	512	512
Maximum # of L2 (IGMPv2 Snooping) multicast groups	6000	6000	6000	6000
Maximum # of L2 (IGMPv3 Snooping) multicast groups	4000	4000	4000	4000
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
# of L3 (*,G) joins per RP	256	NA	256	256
PIM Interfaces Supported	32	32	32	32
IGMP interfaces supported	32	32	32	32
Learning Rate for PIM-SM (flows/second)	32	32	32	32
Maximum # of L2 ACL(ingress/egress) *	3000/120	12000/2000	6128/496	6128/496
Maximum # of L3 ACL ipv4 (ingress/egress) *	1500/120	12000/2000	3064/475	3064/475
Maximum # of class-maps	2048	2048	2048	2048
Maximum # of policy-maps	2048	2048	2048	2048
Maximum # of class-maps per policy map	50	50	50	50
Maximum Total # of L3 ACL ipv6 (ingress/egress) *	500/120	4000/2000	1000/500	1000/500
Maximum # of VF/FCoE interfaces/Logins (Per switch)	1000	1000	1000	1000
Maximum # of Enodes/FCoE Devices per Fabric	2000	2000	2000	2000
Maximum # of NPIV per Port	64	64	64	64

Network OS v7.2.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
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	500	500	500	500
Maximum # of LAGs (Port Channels)	64	288	144	144
Maximum # of members in a standard LAG	16	16	16	16
Maximum # of members in an Extreme Trunk (10G)	16	8	12	12
Maximum # of members in an Extreme Trunk (40G)	2	NA	3	3
Maximum # of members in an Extreme Trunk (100G)	NA	NA	NA	NA
Maximum # of switches in Logical cluster mode **	48	48	48	48
Maximum # of L2 ECMP Paths	16	8	16	16
Maximum # of vLAGs in a fabric	2000	2000	2000	2000
Maximum # of member ports in a vLAG	64	64	64	64
Maximum # of nodes in a vLAG	8	8	8	8
Maximum # of member ports per vLAG per Node	16	16	16	16
Maximum # of Management ACL	256	256	256	256
Maximum # of ARP Entries *	16000	126000	72000	72000
Maximum # of OSPF areas	20	64	20	20
Maximum # of OSPF routers in a single area	64	200	64	64
Maximum # of OSPF adjacencies	100	200	100	100
Maximum # of OSPF routes *	8,000	64,000	10000	10000
# of OSPF Interfaces	100	1,000	100	100
# of OSPF enabled subnets	100	1,000	100	100
# of local subnets in a single area	100	1,000	100	100
Maximum # of OSPFv3 areas	9	9	9	9
Maximum # of OSPFv3 routers in a single area	64	200	64	64
Maximum # of OSPFv3 adjacencies	100	200	100	100
Maximum # of OSPFv3 routes *	1500	64000	1500	1500
# of OSPFv3 Interfaces	100	256	100	100
# of OSPFv3 enabled subnets	100	256	100	100
Maximum # of IPv4 routes in SW *	8000	280000	10000	10000
Maximum # of IPv6 routes in SW *	1500	64000	1500	1500
Maximum # of IPv4 static routes *	2000	40,000	2000	2000
Maximum # of IPv6 static routes *	500	20,000	500	500
Maximum # of VRRP instances per system	255	1024	512	512
Maximum # of VRRP v3 instances per system	255	1024	512	512
Maximum # of VRRP instances per interface	32	32	32	32

Network OS v7.2.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
Maximum # of routers participating in a VRRP-E Session	8	8	8	8
Maximum # of virtual IP addresses per VRRP instance	16	16	16	16
Maximum # of FVG instances per system	256	4096	1024	1024
Maximum # of FVG instances per interface	1	1	1	1
Maximum # of routers participating in a FVG session	32	32	32	32
Maximum # of Gateway IP addresses per FVG instance	1	1	1	1
Maximum # of IPv4 routes with ECMP supported *	8000	200000	10000	10000
Maximum # of IPv6 routes with ECMP supported *	1500	64000	1500	1500
Maximum # of L3 ECMP	16	32	32	32
Maximum # of IPv4 interfaces per system *(Ve intf)	2000	4000	2000	2000
Maximum # of IPv6 interfaces per system *(Ve intf)	512	4000	512	512
Maximum # of VRF per node	32	512	512	512
Maximum # of VRFs support protocols per node	32	128	128	128
Maximum # of I-BGP peers	256	512	256	256
Maximum # of E-BGP peers	256	256	256	256
Maximum # of IPv4 BGP routes in HW *	8000	200000	10000	10000
Maximum # of IPv6 BGP routes in HW *	1,500	64000	1500	1500
Maximum # of IPv4 RIB (IN + OUT) Routes *	110000	1300000	110000	110000
Maximum # of IPv6 RIB (IN + OUT) Routes *	110000	1300000	110000	110000
Maximum # BGP IPv4/IPv6 Peer Group	100	250	100	100
Maximum # of BFD sessions per node	100	100	100	100
Maximum # of UDLD enabled interfaces	64	384	144	108
Maximum # of PVLAN domain supported	1000	1000	1000	1000
Maximum # of Secondary VLANs per PVLAN supported	24	24	24	24
Maximum # of primary VLANs per PVLAN supported in promiscuous mode	24	24	24	24
DHCP IP Helper Addresses per interface	16	16	16	16
DHCP IP Helper VE interfaces	256	1,000	256	256
DHCP IP Helper physical ports	60	384	60	60
DHCP IP Relay Addresses per Node	2000	4000	2000	2000
DHCP IPv6 Relay Address per Node	2000	4000	2000	2000
Max Number of configurable PBR route maps	64	64	64	64
Max Number of configurable PBR stanzas	1024	1024	1024	1024

Network OS v7.2.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
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
Maximum # of OpenFlow GROUP Type ALL	256	256	256	256
Maximum # of OpenFlow GROUP Type SELECT	256	256	256	256
Maximum # 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
Maximum # of MAC's supported for 802.1x MAC Authentication	3000	3000	3000	3000

*Parameters mentioned are applicable on specific HW profiles. Please check the Network OS documentation 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	8	8	8	8	8	8
Max # of Leaf – Spine ECMP	16	16	16	16	16	16
Max # of SAG addresses per interface	64	64	64	64	64	64

HW Profile and Platform Specific Scale Numbers

Route Profile Scale:

VDX 6740, 6740T, 6740T						
Network OS v7.x 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						
Network OS v7.x 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 *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 routes with ECMP	1000	0	0	500	2000	2000

supported *						
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						
Network OS v7.x 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 *	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				
Network OS v7.x Scalability Numbers	VDX6740	VDX-8770	VDX-6940-36Q	VDX-6940-144S
Maximum # of L2 (IGMPv2 Snooping) multicast groups	1000(openflow)	6000	6000	6000
Maximum # of MLD Groups	0	512	512	512
# 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				
Network OS v7.x Scalability Numbers	VDX6740	VDX-8770	VDX-6940-36Q	VDX-6940-144S
Maximum # of L2 (IGMPv2 Snooping) multicast groups	1000	6000	6000	6000
Maximum # of MLD Groups	512	4000	4000	4000
Maximum # of L2 (IGMPv3 Snooping) multicast groups	4000	4000	4000	4000
# of L3 (S,G) forwarding Entries	2,000	2,000	2000	2000
# of L3 (*,G) forwarding Entries	256	256	256	256

NOTE: IGMPV3 snooping configurations should use TCAM PROFILE IPV4-IPV6-MCAST

NOTE: IGMPv3 scale on VDX6940 is 4,000 entries shared between PIM (2000 entries max) and IGMPv3 (4000 max, with no PIM). First Come First Serve basis.

ACL Scale:

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

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

VDX6740									
Network OS v7.x 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 v7.2.0. This is a representative list of devices, Network OS v7.2.0 supports all standards- based devices connected to it for these types of storage.

IP Storage

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

FC/FCoE Storage

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

Network Adapters

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

ADDITIONAL CONSIDERATIONS

Limitations and Restrictions

Command Line Interface

- Break command is not supported. ctrl-c can be used as an alternative.
- Few commands may not display paginated output.
- For few clear and show commands “?” will not show all options for VRF. Tab completion will give all possible values.
- For certain commands (including “no” form with some commands), “?” shows unsupported additional options.
- Some CLI commands will generate an “Error:Access denied” message upon failure. This means the operation failed on the switch and may not be related to permissions.
- Tab completion and <ctrl>-c (cancel) does not work for some commands.
- Incorrect range might be displayed in the help text for some of the show commands.
- Range support is available for all the interfaces in Network OS v7.1.0. Following limitations are applicable:
 - Interface range command is supported on breakout ports of same connector. Range is not supported involving breakout ports of multiple connectors.
 - Interface range command does not support mix of regular ports and breakout ports.
 - Range command is not supported across multiple slots of the chassis.
 - 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 v7.1.0. 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 cannot 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 user needs to enter "no ipv6 router ospf vrf default-vrf".
- The “show ip interface ve xx” displays “ICMP unreachable are always sent” even though it is disabled.

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.

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.
- The requirement for a Brocade branded USB is removed, and all branded USBs are now supported. The user must create a /Brocade directory in order for the USB to function.

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.

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 Management Configuration 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 Network OSv7.0.1:

Note the following results for the given actions.

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

VCS Node Replacement

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

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

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

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

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

```
dpod 1/0/2
    !

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

Logical Chassis

- Configurations are not auto preserved on mode transitions (between Fabric Cluster and Logical Chassis mode). Please follow the mode transition procedure as outlined in the Network OS Management Configuration Guide.
- 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 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 and VDX 6740T-1G Extreme trunk (BTRUNK) can support up to 2x40G member links in the same trunk group for a maximum throughput of 80G.
- The VDX 8770 Extreme trunk (BTRUNK) can support up to 8 member links with a maximum throughput of 80G using 8x10G ports in the same trunk group. Full link utilization of 8 ports in a trunk group is achievable with larger packet size (>128 Bytes).
- 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.

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 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 be able to identify peer port config is breakout and link won't be stable.
- On VDX 6740T/6740T-1G, 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 needs 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 logincli):
 - 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 Vlans.
 - CLI to enable Local Proxy ARP on primary VLAN.
 - IP Configuration on PVLANS
 - Vx Configuration on both Primary and Secondary Vlans

UDLD

- AMPP on PVLANS
- In case of MSTP if a primary VLAN is added to the instance automatically secondary VLAN also added to the instance.
- When the operator wants to delete the host association on a host port recommended to use “no switchport” rather than “no switchport private-VLAN host-association”. This is applicable only when the host port is untagged. When the host port is tagged both the commands can be used.
- Primary VLAN ID needs to be lower than the secondary VLAN IDs. If primary VLANID is greater than secondary there is an issue with config replay.
- In Logical Chassis mode source macs may not learn on PVLAN configured ports, after deleting some of the secondary VLANs for which the traffic is not flowing.
 - The UDLD protocol is not supported on the members of an 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.
- 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, Network OS 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)#
```

- STP Interop with certain vendor switches

To allow for STP interop with certain 3rd party switches that would not accept the BPDU source MAC's with default OUI. The selection can be changed using the below command now:

```
system-id oui <01.e0.52 | 00.e0.52> (under stp configuration)
```

IGMPv3 Snooping

- IPv4 PIM is not supported on IGMPv3 enabled VLAN (No error is displayed when user tries to enable PIM on IGMPv3 enabled VLAN or vice-versa).
- When user is enabling IGMPv3 snooping, the feature restrict-unknown-multicast needs to be enabled on the same VLAN.

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
- To form an ISL between 10G tunable ZR optics (57-1000266-01) when initially inserting the optic and configuring "tunable sfpp channel x", please configure any channel other than 1 on both ends.

AMPP and Port-Profiles

- Port-profile status does not reflect the remote interface info in VCS 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 Network OS 7.2.0, the default mode of operation is Access Gateway for the VDX 6740, 6740T, 6740T-1G,.
- Adding/removing an ISL or adding/removing a VDX switch from cluster may cause that exiting VDX switch in the cluster stops generating a FCOE keep alive. And this results in FCoE port continuously flapping. To recover, re-configure FCoE by removing and adding fcoeport configuration (no fcoeport/fcoeport default) on the affected interfaces. As an alternate recovery, shut / no shut on the affected interfaces.
- 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 Network OS version (pre 4.1.2 builds) to v7.0.1, it will be in FCF mode.
- 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 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 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

Provisioning :

- A new CLI has been introduced in 7.0.1a under Rbridge mode that allows the user to disable the ISL capability of all interfaces in the switches using single command. Specific interfaces that needs ISL capability can be enable the functionality using “no” form of command under interface mode.
fabric neighbor-discovery disable (under Rbridge mode)
- Similarly, there are new CLIs added to assist in MTU configuration across all interfaces for a switch using single CLI. This allows quickly setting the jumbo frame capability across the switch for Vxlan / Storage traffic.

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.
- 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 profiles 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 avLAG pair.
- In 7.0.1a, the scale support for SAG been increased from 32 to 64 under each interface.

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 defaultvrf. In earlier Network OS 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

Traffic duplication is seen on Last hop router on shared RP tree initially when new source traffic starts for about 40 seconds in scale scenarios.

- Static or Dynamic RP Candidate is not supported on VDX 8770

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.
- DHCP relay doesn't work correctly with just Fabric Virtual Gateway (FVG) on the same VE interface. The workaround is to configure unique IP addresses on VE interfaces simultaneously.

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 Network OSCLI "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.

Mac Port Based Authentication

- For Mac Auth Bypass to work, user should configure 'dot1x reauthentication' followed by 'dot1x reauthMax “3 or more”'.

Uplink Switch Support

- STP should not be enabled on uplink ports
- Transparent vlans are not supported on protected and uplink ports.
- User VLANs with same internal (7168-8191)VLAN mapping should not be used
- Same 1k VLANs should be used across rbridges (though configuring different VLANs are supported)
- Virtual Fabric should be enabled in the switch in order to enable uplink-switch feature using the global CLI.
- VLANs 7168-8191 would be reserved internally when the global CLI is executed and these VLANs are not allowed to be created by the user using the CLI.
- Without enabling the feature using the global CLI, enabling protected port configuration on interface level will not work and throws an error.

scenario	# protected ports	# dot1q vlans	# gvlan
min ports, dot1q vlans	1	1024	0
max ports, dot1q vlans	46	80	0
min ports, gvlan	1	0	634
max ports, gvlan	46	0	80
min ports, mixed vlans	1	496	512
max ports, mixed vlans	46	32	40

- ☒ The VLAN /VF configured should be same on protected and uplink ports.
 - By default, all switchports are in unprotected mode which is same as uplink port mode.
 - No new CLI is needed to distinguish an uplink port, since by default all switchports are in uplink port mode.
 - Enabling protected port configuration is not allowed without any VLAN(s) configured on the interface.
 - At least one uplink port should be present in order to have a protected port configuration.
 - In case of VCS one uplink port should be present for each r-bridge.
 - In case of a vLAG, each node of vLAG should have at least one uplink port in order to have successful protected configuration on vLAG.
 - VDX6740 scaling limitations

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.

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 Network OS Network OS v6.0.1. Such topologies and configuration must be removed before downgrading to any version below Network OS 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 Flow based 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.

URPF

-  uRPF is not supported in Mercury

BGP Auto neighbor discovery

- BGP Auto Neighbor Discovery is only supported for IPv4 in default VRF. VE and MULTI HOP supported is also not available

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.

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.
- To allow for STP interop with certain 3rd party switches that would not accept the BPDU source MACs with default OUI. The selection can be changed using the belowcommand now:

system-id oui <01.e0.52 | 00.e0.52> (under stp configuration)
- 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 EXMtable respectively.

MAPS

- BNA version 14.2.1 is supported with Network OSNetwork OS 7.2.0.
- MAPS is supported on VDX 6740, 6940 and 8770 platforms.
- MAPS port level alerting in Network OS V7.0.x 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 only 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.
- It is recommended to keep same values for Global MTU & Interface value as due to a known defect, change in Global MTU may impact the interface MTU too.

DEFECTS

TSBs - Critical Issues to Consider Prior to Installing This Network OS Release

Technical Support Bulletins (TSBs) are produced to provide detailed information about high priority defects or issues present in Network OS 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 Network OS-based platforms).

Network OS v7.2.0a Caveats

None

Defect ID: DEFECT000660024	Defect ID: DEFECT000660024
Technical Severity: Critical	Probability: Critical
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.2.0	Technology: IP Fabric
Symptom: Could not learn mac from node which is part of vcs.	
Condition: MAC moved from LAG to Edge port	

Defect ID: DEFECT000660361	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.2.0	Technology: IP Fabric
Symptom: DCM daemon termination	
Condition: ISSU upgrade from NOS7.2.0 to NOS7.2.0a (timing condition).	

Network OS v7.2.0 Caveats

Kernel Space

Kernel space memory high consumption and high CPU followed by Out of Memory might happen if MAC scale is more than 25K and those MAC are learned through Vxlan tunnel.

Recovery: Reduce MAC scale.

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 tnImgr 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:

- 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 Network OS v7.2.0g

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as March 2021 in Network OS v7.2.0g.

Parent Defect ID:	NOS-68139	Issue ID:	NOS-68139
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.2.0f
Technology Group:	Other	Technology:	Other
Symptom:	IPC-5024 fail message and high CPU for L2sys and unexpected failover.		
Condition:	During EthNs MAC checksum update and MAC update in high scale		
Workaround:	Increase the MAC age-out time and disable or increase the interval of MAC consistency check.		

Closed with code changes for Network OS v7.2.0f

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as January 2021 in Network OS v7.2.0f.

Parent Defect ID:	NOS-67752	Issue ID:	NOS-67754
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.1.0b3
Technology Group:	Management	Technology:	Management GUI
Symptom:	Unexpected reload.		
Condition:	When we receive multiple REST queries continuously.		

Parent Defect ID:	NOS-67887	Issue ID:	NOS-67887
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.2.0c
Technology Group:	Other	Technology:	Other
Symptom:	Unexpected reload.		
Condition:	High rate of software assisted layer 3 forwarding of traffic, causing connection tracking table to fill up.		

Parent Defect ID:	NOS-67936	Issue ID:	NOS-67936
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.2.0c
Technology Group:	Management	Technology:	Management GUI
Symptom:	Unexpected reload.		
Condition:	When we receive multiple REST queries continuously.		

Parent Defect ID:	NOS-67866	Issue ID:	NOS-67989
Severity:	S4 - Low		
Product:	Network OS	Reported in Release:	NOS7.0.2c

Technology Group:	Security	Technology:	Security Vulnerability
Symptom:	VDX is vulnerable if telnet connection enabled.		
Condition:	VDX is vulnerable if telnet connection enabled.		
Workaround:	Disable the telnet feature and use SSH for secure login to switch		

Parent Defect ID:	NOS-67933	Issue ID:	NOS-67998
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.2.0a1
Technology Group:	Security	Technology:	SSH - Secure Shell
Symptom:	Unexpected reload.		
Condition:	Providing the empty password while importing the SSH key.		

Parent Defect ID:	NOS-68002	Issue ID:	NOS-68002
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.2.0e
Technology Group:	Other	Technology:	Other
Symptom:	IPC-5024 fail message and Unexpected reload.		
Condition:	MAC consistency check with default interval of 300sec.		
Workaround:	Increase the mac consistency-check interval and aging interval time.		

Parent Defect ID:	NOS-67895	Issue ID:	NOS-68035
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.1.0b
Technology Group:	Management	Technology:	Management GUI
Symptom:	Unexpected reload.		
Condition:	When we receive multiple REST queries continuously.		

Closed with code changes for Network OS v7.2.0e

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as July 2020 in Network OS v7.2.0e.

Parent Defect ID:	NOS-67405	Issue ID:	NOS-67418
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.2.0a
Technology Group:	Management	Technology:	SNMP - Simple Network Management Protocol
Symptom:	dot1qPvid MIB variable is always 0 for trunk port		
Condition:	When particular interface configured as trunk port.		

Parent Defect ID:	NOS-67447	Issue ID:	NOS-67448
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.0.2c
Technology Group:	Security	Technology:	ACLs - Access Control Lists

Symptom:	It is a show command issue. "show access-list interface port-channel XX" will not show the proper output about the active/Inactive status. No functional impact,
Condition:	When ACL applied under VLAG.

Parent Defect ID:	NOS-67454	Issue ID:	NOS-67455
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.1.0b
Technology Group:	Data Center Fabric	Technology:	AMPP - Automatic Migration of Port Profiles
Symptom:	Unexpected reload		
Condition:	While activating auto port-profile on interface		

Parent Defect ID:	NOS-67461	Issue ID:	NOS-67462
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.3.0a
Technology Group:	IP Multicast	Technology:	IGMP - Internet Group Management Protocol
Symptom:	Unable to process the IGMP join/leave messages after hitting the "memory allocation error" RASLOG message.		
Condition:	When they have more than 1000 IGMP join / leave request per second.		
Recovery:	Reload system		

Parent Defect ID:	NOS-67475	Issue ID:	NOS-67477
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.2.0c_CVR
Technology Group:	Layer 2 Switching	Technology:	VXLAN - Virtual Extensible LAN
Symptom:	Vxlan tunnel loopback IP is not reachable between two tunnel end point		
Condition:	While reloading one or more than one tunnel destination node		
Workaround:	No active overlay-gateway and then active overlay-gateway		
Recovery:	1 .Deactivate overlay-gateway config and then active overlay-gateway 2. Reload tunnel node		

Parent Defect ID:	NOS-67503	Issue ID:	NOS-67505
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.3.0a
Technology Group:	Management	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP trap fails		
Condition:	With SNMP version v2		

Parent Defect ID:	NOS-67543	Issue ID:	NOS-67545
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.2.0c
Technology Group:	Management	Technology:	NTP - Network Time Protocol
Symptom:	Unexpected reload		
Condition:	When ntpdate (with NTP config) process spawn's multiple times.		

Parent Defect ID:	NOS-67674	Issue ID:	NOS-67676
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.0.2b
Technology Group:	Monitoring	Technology:	MAPS - Monitoring and Alerting Policy Suite
Symptom:	Continuous RASLOG, SNMP, EMAIL notification for Rx_sym_err.		
Condition:	Whenever any symbol errors are logged		

Parent Defect ID:	NOS-67689	Issue ID:	NOS-67691
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.3.0a
Technology Group:	IP Multicast	Technology:	IGMP - Internet Group Management Protocol
Symptom:	Unexpected reload		
Condition:	Periodic IGMP join and leave messages		

Parent Defect ID:	NOS-67692	Issue ID:	NOS-67694
Severity:	S2 - High		
Product:	Network OS	Reported in Release:	NOS7.2.0c
Technology Group:	Layer 3 Routing/Network Layer	Technology:	OSPF - IPv4 Open Shortest Path First
Symptom:	OSPF route not installed in the IP routing table		
Condition:	When RFC 1583 compatibility is not configured, and during route calculation for type-5 LSA, we are not considering INVALID ASBR route as best ASBR route		
Recovery:	Clear OSPF route		

Parent Defect ID:	NOS-67701	Issue ID:	NOS-67703
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.1.0b3
Technology Group:	Layer 2 Switching	Technology:	VXLAN - Virtual Extensible LAN
Symptom:	Unexpected reload.		
Condition:	When multiple Tunnel-Id creation/deletion operation done frequently.		

Parent Defect ID:	NOS-67724	Issue ID:	NOS-67725
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.1.0b3

Technology Group:	Monitoring	Technology:	Hardware Monitoring
Symptom:	The output of "show environment temp" not displays any values to monitor temperature.		
Condition:	When we have all the Line Card's (which has 7 temperature sensors per LC) inserted. Ex: 6X100G, 27X40G and 48X10GT		

Parent Defect ID:	NOS-67795	Issue ID:	NOS-67809
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.1.0b
Technology Group:	Layer 3 Routing/Network Layer	Technology:	ICMP - Internet Control Message Protocol
Symptom:	Traceroute successful for shut interface		
Condition:	Traceroute from remote node		

Parent Defect ID:	NOS-67121	Issue ID:	NOS-67550
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.2.0
Technology Group:	Management	Technology:	Management GUI
Symptom:	Unexpected reload.		
Condition:	High rate of ENS MAC update.		

Closed with code changes for Network OS v7.2.0d

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as December, 2019 in Network OS v7.2.0d.

Parent Defect ID:	NOS-55113	Issue ID:	NOS-63315
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.2.0	Technology:	VLAN - Virtual LAN
Symptom:	Physical and port-channel interfaces configured for both uplink-switch protected-port and switchport mode trunk-no-default-native cannot send or receive frames.		
Condition:	Physical and port-channel interfaces configured for both uplink-switch protected-port and switchport mode trunk-no-default-native		
Workaround:	Don't configure "uplink-switch protected-port" and "switchport mode trunk-no-default-native" on the same interface. If one already has, recovery requires one to first remove "uplink-switch protected-port", remove all switchport settings with "no switchport", and then re-add all switchport settings.		

Parent Defect ID:	NOS-54700	Issue ID:	NOS-67025
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer

Reported in Release:	NOS7.1.0b	Technology:	DHCP - Dynamic Host Configuration Protocol
Symptom:	DHCP offer is not received		
Condition:	For Linux command "dhclient eth0" after upgrading from nos6.x to nos7.x		
Workaround:	Have the Broadcast flag set with dhclient command(dhclient -B eth0) used		

Parent Defect ID:	NOS-66881	Issue ID:	NOS-67049
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.2.0	Technology:	AAA - Authentication, Authorization, and Accounting
Symptom:	PAM authentication fails for curl/REST query using TACACS users		
Condition:	When we run the curl/REST queries simultaneously with multiple sessions		
Workaround:	Use the local users or avoid running queries simultaneously with multiple sessions.		

Parent Defect ID:	NOS-53109	Issue ID:	NOS-67050
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.1c	Technology:	VCS Fabric
Symptom:	NPIV device not able to login		
Condition:	After upgrading from 6.x to 7.x		

Parent Defect ID:	NOS-66997	Issue ID:	NOS-67067
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	ARP - Address Resolution Protocol
Symptom:	Repeated IPAD-1001-log entries even though there is no change to mgmt interface status		
Condition:	mgmt-vrf default route is resolved via inband interface		
Workaround:	None		

Parent Defect ID:	NOS-67019	Issue ID:	NOS-67198
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.2b	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP walk fail		
Condition:	Reload of the switch having IP ACL for SNMP community config will result in SNMP walk failure		

Parent Defect ID:	NOS-67208	Issue ID:	NOS-67244
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management

Reported in Release:	NOS6.0.1	Technology:	CLI - Command Line Interface
Symptom:	"show access-list ip" output is not showing as expected.		
Condition:	In a scenario, if we remove ACL from one of the management interface in cluster, it is impacting the "show access-list ip" output on other nodes.		

Parent Defect ID:	NOS-67014	Issue ID:	NOS-67308
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.2.0a	Technology:	OSPF - IPv4 Open Shortest Path First
Symptom:	Ping between the 2 vcs fails upon reload of RB 120 in a multi vcs topology. Second symptom is BGP convergence is not triggered when port between the VCS is shut.		
Condition:	Ping failure upon reload of one of the boxes results in packet being sent through the ISL link to the destination. ISL link encapsulation is different from normal interface. Reload box loses the isl encap details due to ecmp path availability post reboot causing the problem.		
Workaround:	Ecmp path and once of the path is isl on the device that undergoes reload. Upon reload, if isl link is chosen from traffic to be sent the issue would occur.		

Parent Defect ID:	NOS-53094	Issue ID:	NOS-67319
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.1c	Technology:	ARP - Address Resolution Protocol
Symptom:	Unexpected reload		
Condition:	Logging raslog message ARP-1038 suppress configuration		
Workaround:	Avoid using logging raslog message ARP-1038 suppress		

Parent Defect ID:	NOS-67221	Issue ID:	NOS-67331
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.2.0c	Technology:	IGMP - Internet Group Management Protocol
Symptom:	Multicast traffic is not forwarded through ISL links.		
Condition:	When the Group Specific Query is not reachable to particular VDX.		

Parent Defect ID:	NOS-67189	Issue ID:	NOS-67334
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.2.0a	Technology:	VCS Fabric
Symptom:	Unexpected reload.		
Condition:	There is a race condition in hrtimer_enqueue_reprogram during the unlock/lock sequence. It is a rare scenario to hit.		

Parent Defect ID:	NOS-67184	Issue ID:	NOS-67335
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.1c	Technology:	DHCP - Dynamic Host Configuration Protocol
Symptom:	Admin guide correction for CLI "sw0(conf-if-te-1/0/1)# ip dhcp relay gateway address".		
Condition:	"Address" keyword is used in IP DHCP relay gateway.		
Workaround:	Don't use "Address " Keyword in "sw0(conf-if-te-1/0/1)# ip dhcp relay gateway address".		

Parent Defect ID:	NOS-67290	Issue ID:	NOS-67337
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.1.0b3	Technology:	VCS Fabric
Symptom:	FSPF-1013 log with severity warning		
Condition:	Fabric having more than 16 paths to reach a RB.		

Parent Defect ID:	NOS-67182	Issue ID:	NOS-67338
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.1.0b3	Technology:	LAG - Link Aggregation Group
Symptom:	Unexpected Line card reload		
Condition:	Having lacp default-up config for a Port-channel might cause out of memory condition on a Line card.		

Parent Defect ID:	NOS-67223	Issue ID:	NOS-67339
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.1.0b3	Technology:	LAG - Link Aggregation Group
Symptom:	Fabric disconnection / Unexpected reload		
Condition:	Applying Port-Channel config on ISL port which is in shut state .		

Parent Defect ID:	NOS-67109	Issue ID:	NOS-67340
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.2b	Technology:	CLI - Command Line Interface
Symptom:	Un expected reload.		
Condition:	switch reloaded after issuing show startup-config in multiple sessions.		
Workaround:	Avoid parallel request of show startup-config		

Parent Defect ID:	NOS-67110	Issue ID:	NOS-67348
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0a	Technology:	SNMP - Simple

			Network Management Protocol
Symptom:	SNMP output for dot1qVlanStaticEgressPorts and dot1qVlanStaticUntaggedPorts is not accurate. That is, it's not matching that of 'show vlan brief' output. This is causing XMC Device view to show inaccurate data for 'Static Egress ports' and 'Static Untagged Ports' in 802_1Q VLAN Table.		
Condition:	Occurs always since, the SNMP output has some extra bits set.		
Workaround:	None		

Parent Defect ID:	NOS-67111	Issue ID:	NOS-67349
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0a	Technology:	SNMP - Simple Network Management Protocol
Symptom:	Internal VLANs 4093 and 4095 are displayed on SNMP walk of dot1qVlan (in Q-BRIDGE-MIB) and in XMC Device view of 802_1Q VLAN Table.		
Condition:	These internal VLANs are not displayed in 'show vlan brief' CLI output and thus causing inconsistency between CLI output and SNMP/XMC view.		
Workaround:	None		

Parent Defect ID:	NOS-67225	Issue ID:	NOS-67359
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.3.0aa	Technology:	LAG - Link Aggregation Group
Symptom:	Unexpected reload		
Condition:	During SNMP walk and collecting "show running-config"		

Parent Defect ID:	NOS-67162	Issue ID:	NOS-67360
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.3.0a	Technology:	User Accounts & Passwords
Symptom:	Error message "syntax error: password has a bad length/size."		
Condition:	When configuring the user-name password have the cipher text length >40 characters.		
Workaround:	Use the user-password length less than 16 characters		

Parent Defect ID:	NOS-67387	Issue ID:	NOS-67388
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.2.0c	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	Unexpected reload		
Condition:	While doing the snmpwalk/getnext for BGP peer IP address status.		
Workaround:	Avoid SNMP get for BGP Peer IP address.		

Parent Defect ID:	NOS-67402	Issue ID:	NOS-67403
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS5.0.1a	Technology:	AMPP - Automatic Migration of Port Profiles
Symptom:	Unexpected reload		
Condition:	Remove/re-add vCenter config with max character 80 switch port-profile .		

Closed with code changes for Network OS v7.2.0c

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as June 05, 2019 in Network OS v7.2.0c.

Parent Defect ID:	NOS-53138	Issue ID:	NOS-57662
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.0.2	Technology:	HTTP/HTTPS
Symptom:	HTTPS will be enabled if expired TLS certificate and key is imported to device using scpuser credentials. HTTPs should not be enabled if the certificate is expired.		
Condition:	When expired TLS certificate is imported to device using scpuser credentials, HTTPS can be enabled even with expired TLS certificate.		
Workaround:	N/A		

Parent Defect ID:	NOS-53106	Issue ID:	NOS-66241
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.0.1c	Technology:	VLAN - Virtual LAN
Symptom:	Microsoft SQL Clustering failover failed. The primary and backup server of MS SQL are connected via VDX IP fabric.		
Condition:	When we do the MS SQL fail over(Moving primary to secondary)		

Parent Defect ID:	NOS-54706	Issue ID:	NOS-66884
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Network Automation and Orchestration
Reported in Release:	NOS7.1.0b	Technology:	NETCONF - Network Configuration Protocol
Symptom:	Switch reloaded on executing vcenter REST query.		
Condition:	Reload happens on processing vcenter REST query without any data field.		
Workaround:	Avoid sending vcenter REST query without any data field.		

Parent Defect ID:	NOS-55932	Issue ID:	NOS-66885
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0b	Technology:	CLI - Command Line Interface
Symptom:	Copy running-config failed thru ftp.		
Condition:	If the config filename has the special characters, the copy running-config failed thru ftp.		
Workaround:	None		

Parent Defect ID:	NOS-54632	Issue ID:	NOS-66923
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.1.0a	Technology:	CLI - Command Line Interface
Symptom:	Rbridge misspelled as a Rbridge.		
Condition:	The Rbridge misspelled as a Rbridge in help string while executing the firmware download default-config in privilege mode.		
Workaround:	None		

Parent Defect ID:	NOS-66927	Issue ID:	NOS-66944
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.1.0a	Technology:	IPv4 Multicast Routing
Symptom:	switch reloaded after running "show pim ip mcache"		
Condition:	Running the show pim pi mcache		
Workaround:	None		

Parent Defect ID:	NOS-66321	Issue ID:	NOS-66948
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Network Automation and Orchestration
Reported in Release:	NOS7.2.0a1	Technology:	OpenStack Integration
Symptom:	VDX6740T-1G using NOS 7.2.0b will show blinking green and amber port LED when port is offline but in "no shutdown" state.		
Condition:	VDX6740T-1G using NOS 7.2.0b will show blinking green and amber port LED when port is offline but in "no shutdown" state.		
Workaround:	No workaround		

Parent Defect ID:	NOS-54704	Issue ID:	NOS-66957
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.1.0b	Technology:	SNMP - Simple Network Management Protocol

Symptom:	Unexpected reload
Condition:	when two tlvs are received but tlv type of one is either zero or unsupported.
Workaround:	None

Parent Defect ID:	NOS-66866	Issue ID:	NOS-66959
Severity:	S2 – High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS6.0.2f	Technology:	AAA - Authentication, Authorization, and Accounting
Symptom:	Unexpected Reload		
Condition:	When TACACS authorization fails on re-try		

Parent Defect ID:	NOS-66780	Issue ID:	NOS-66960
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0b	Technology:	Static Routing (IPv4)
Symptom:	"sh ip route" shows in-correct route entries.		
Condition:	IP address(subnet) update followed by VRF config change from user-defined to Default-vrf.		

Parent Defect ID:	NOS-66900	Issue ID:	NOS-66962
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS6.0.2h	Technology:	LAG - Link Aggregation Group
Symptom:	Some of the mac addresses will not be learnt after firmware upgrade.		
Condition:	On firmware upgrade.		
Workaround:	None		

Parent Defect ID:	NOS-48024	Issue ID:	NOS-66976
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS6.0.2h	Technology:	SNMP - Simple Network Management Protocol
Symptom:	Unexpected system reload		
Condition:	Reload triggered when polling ipRouteTable (1.3.6.1.2.1.4.21) SNMP		
Workaround:	None		

Parent Defect ID:	NOS-66850	Issue ID:	NOS-66987
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.3.0a	Technology:	SNMP - Simple Network Management Protocol
Symptom:	snmpget of snIfOpticalLaneMonitoringTable returns "No Such Instance" errors.		
Condition:	snIfOpticalLaneMonitoringTable is not supported for 1G and 10G interfaces and it is applicable only for 40G/100G optics.		
Workaround:	None		

Parent Defect ID:	NOS-53027	Issue ID:	NOS-67000
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.1b	Technology:	Configuration Fundamentals
Symptom:	Username is displayed incorrectly in "show logging auditlog".		
Condition:	When logged in with a username and makes some changes.		
Workaround:	None		

Parent Defect ID:	NOS-66919	Issue ID:	NOS-67005
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0b	Technology:	SNMP - Simple Network Management Protocol
Symptom:	Customer won't be able to view VE interface related info in ifTable and corresponding stats.		
Condition:	None		
Workaround:	None		

Parent Defect ID:	NOS-67008	Issue ID:	NOS-67009
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.3.0aa	Technology:	User Accounts & Passwords
Symptom:	Unable to change the root password in root prompt.		
Condition:	After a successful login to the root prompt, user unable to changes the root password.		
Workaround:	We can use /bin/passwd to the change the root password by bypassing the ACTIVE CP validation.		

Parent Defect ID:	NOS-47973	Issue ID:	NOS-67011
Severity:	S1 - Critical		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS6.0.2e	Technology:	VCS Fabric
Symptom:	Unexpected reload.		
Condition:	There are certain conditions that may cause the IPv6 traffic to be sent to the CPU.		

Parent Defect ID:	NOS-66864	Issue ID:	NOS-67034
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Other
Reported in Release:	NOS7.4.0	Technology:	Other
Symptom:	NOS version 7.0.2b, 7.2.0b and 7.4.0 will not allow the LC48X10GT to boot up due to a missing file		
Condition:	LC48X10GT will not boot up online		

Parent Defect ID:	NOS-66828	Issue ID:	NOS-67037
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS6.0.2b	Technology:	CLI - Command Line Interface
Symptom:	Firmware download failure when using the "manual" option in the CLI firmware download command.		
Condition:	Firmware download failure when using "manual" option.		

Parent Defect ID:	NOS-66955	Issue ID:	NOS-67039
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS6.0.2b	Technology:	CLI - Command Line Interface
Symptom:	ZTP not supported two version upgrade. Also, default-config flag in DAD file is not worked.		
Condition:	When firmware upgrade using DAD or ZTP with 'defaultconfig' flag in the DAD config file.		
Workaround:	None		

Parent Defect ID:	NOS-66998	Issue ID:	NOS-67041
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.0	Technology:	Logical Chassis
Symptom:	No reason at RASLOG when interface administratively down without user config.		
Condition:	When exception created on port due to ASIC fault.		

Parent Defect ID:	NOS-36884	Issue ID:	NOS-67043
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS5.0.1	Technology:	VLAN - Virtual LAN
Symptom:	"show interface description" and "show interface trunk" do not display all port-channels.		
Condition:	Multi-node cluster with Port channels configured and most of the port- channels have ports from non-principal nodes.		
Workaround:	None		

Parent Defect ID:	NOS-66898	Issue ID:	NOS-67048
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	BGP4+ - IPv6 Border Gateway Protocol
Symptom:	Unexpected reload.		
Condition:	When we run the "show ip bgp neighbors" after giving the "no neighbor activate" followed by "clear ip bgp neighbor"		

Parent Defect ID:	NOS-66325	Issue ID:	NOS-67079
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.1.0b	Technology:	Port Mirroring
Symptom:	Flow Based SPAN does not work with ip access list.		
Condition:	If the traffic is IP traffic, VDX do not do Egress SPAN.		
Workaround:	None		

Parent Defect ID:	NOS-67061	Issue ID:	NOS-67084
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.2a	Technology:	Configurations
Symptom:	Linecard reset by the detection of Multibit Parity error condition show in RAS Log. "[HWK2-5651], 14604207/0, L4/0 Active, ERROR, VDX8770-8, CHIP0: Interrupt: DMC dmc mbit err."		
Condition:	Linecard reset.		
Workaround:	None		

Parent Defect ID:	NOS-66324	Issue ID:	NOS-67106
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS6.0.2	Technology:	LAG - Link Aggregation Group
Symptom:	Port channel will be operationally down		

Condition:	This is rare scenario and can happen only when we have the ASIC resources are full.
-------------------	---

Closed with code changes for Network OS v7.2.0b

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as December 07, 2018 in Network OS v7.2.0b.

Defect ID:	DEFECT000609232		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	Logical Chassis
Reported In Release:	NOS7.1.0	Technology:	VCS
Symptom:	<ol style="list-style-type: none"> 1. Compact Flash grows and disk full issue can occur. 2. Unexpected DCMd daemon termination. 		
Condition:	No special condition or configuration required to hit this issue.		
Workaround:	<p>For syslog.log file growing beyond limit below is preventive workaround:</p> <p>Please comment below two lines from syslog config & template files and reread config file from root using below command [killall]. It should stop all those IO messages. Please also verify that syslog utility is working fine after the workaround applied to make sure all is fine before we try on customer production environment.</p> <pre> /etc/syslog-ng/syslog-ng.conf /etc/syslog-ng/syslog-ng.template #destination df_kern { pipe("/var/log/kmsg"); }; #log { source(s_all); filter(f_kern); destination(df_kern); }; root > /usr/bin/killall -HUP syslog-ng </pre>		
Recovery:	Empty /var/log/syslog.log file if it is growing beyond 1 Mb.		

Defect ID:	DEFECT000630220		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	IGMP - Internet Group Management Protocol
Reported In Release:	NOS6.0.2	Technology:	IP Multicast
Symptom:	Multicast packet drop on mrouter port for short amount (< 1sec) of time.		
Condition:	When igmp snooping is enabled and the last locally connected receiver leaves.		

Defect ID:	DEFECT000646316		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	ARP - Address Resolution Protocol

Reported In Release:	NOS6.0.2	Technology:	Layer 3 Routing/Network Layer
Symptom:	Unexpected reload of switch.		
Condition:	Removing L3 configs (in specific IPv4 addresses) and defaulting the config for VDX.		

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

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

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

Defect ID:	DEFECT000648702		
Technical Severity:	High	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	OSPFv3 - IPv6 Open Shortest Path First
Reported In Release:	NOS7.2.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	Observed DCMD termination which deleting non existing CLI "ipv6_ospf_auth" under ipv6 ospf via Netconf.		
Condition:	Deletion of non existing CLI "ipv6_ospf_auth".		

Defect ID:	DEFECT000649945		
Technical Severity:	Medium	Probability:	High
Product:	Extreme Network OS	Technology Group:	Configuration Fundamentals
Reported In Release:	NOS7.1.0	Technology:	Management
Symptom:	VDX 1G port on "auto/auto" does not come up when remote is set to "100/full".		
Condition:	When remote is set to "100/full", the VDX 1G link istays down.		

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

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

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

Defect ID:	DEFECT000657950		
Technical Severity:	Medium	Probability:	High
Product:	Extreme Network OS	Technology Group:	Syslog
Reported In Release:	NOS7.0.1	Technology:	Monitoring

Symptom:	After adding a VDX to an existing VCS using "vcs replace", the newly added VDX is unable to send messages to a remote syslog server.
Condition:	The newly added or reconnected VDX will be able to see "logging syslog-server" settings in "show run", but it will not be able to send syslog messages to that remote server
Workaround:	After this issue has occurred on a newly added non-principal, it is possible to recover by removing and re-applying the "logging syslog-server" setting on the VCS principal rbridge.

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

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

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

Defect ID:	DEFECT000658974		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	Configuration Fundamentals
Reported In Release:	NOS7.2.0	Technology:	Management

Symptom:	Default-config operations [copy default-config startup-config, FWDL with default-config, netInstall] does not preserve the DHCP configuration on management interface
Condition:	Performing default-config operations [copy default-config startup-config, FWDL with default-config, netInstall].

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

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

Defect ID:	DEFECT000659860		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	Logical Chassis
Reported In Release:	NOS7.0.1	Technology:	Data Center Fabric
Symptom:	Unexpected reload.		

Condition:	This occurs when a physical port is added to a port-channel after an ISSU upgrade was performed (Before VDX reloaded since upgrade).
-------------------	--

Defect ID:	DEFECT000660509		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	LAG - Link Aggregation Group
Reported In Release:	NOS6.0.2	Technology:	Layer 2 Switching
Symptom:	Ping loss across vxlan tunnel with extreme trunks		
Condition:	Port channel between 2 rbridges was an extreme trunk and other 2 rbridges was a standard lag.		

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

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

Defect ID:	DEFECT000661265		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	BFD - BiDirectional Forwarding Detection
Reported In Release:	NOS7.1.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	Unexpected reload on Line card		
Condition:	When Packet with Destination port 213 is reached to BFD agent Daemon at Line card.		
Workaround:	Disable the BFD agent daemon at Line card by executing the below commands at MM. sw0:FID128:root> chkconfig bfd off sw0:FID128:root> chroot /mnt chkconfig bfd off		

Recovery:	It will auto recover.
------------------	-----------------------

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

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

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

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

Defect ID:	DEFECT000661717		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	Hardware Monitoring
Reported In Release:	NOS7.1.0	Technology:	Monitoring
Symptom:	After configuring the inband ve interface in the mgmt-vrf as SNMP trap source, the agent_addr coming up as 0.0.0.0 after a switch reloads.		
Condition:	Inbound ve used as snmp trap source and SNMPv1 used.		

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

Defect ID:	DEFECT000662247		
Technical Severity:	High	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	Static Routing (IPv4)
Reported In Release:	NOS7.0.1	Technology:	Layer 3 Routing/Network Layer
Symptom:	BGP route not cleared when VE interface is shut.		
Condition:	VE interface shut		
Recovery:	clear bgp neighbor		

Defect ID:	DEFECT000662379		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	Logical Chassis
Reported In Release:	NOS7.3.0	Technology:	Data Center Fabric
Symptom:	VDX device will startup with the default configuration and also reports the same using below RAS log. [DCM-3053], SW/O Active, INFO, VDX-VCS, Dcmd database corruption was detected. The system will startup with the default configuration for this database.		
Condition:	Sudden Power Cycle the device can cause the issue.		

Workaround:	<p>We can use below workaround for planned outage or power cycle. NOS7.1.0 and Later Releases:</p> <ol style="list-style-type: none"> 1. Execute chassis power-cycle-db-shutdown command through NOS CLI. 2. Reload the switch after the below RASLOG: [DCM-1015], SW/0 Active, INFO, VDX6740T, Switch is prepared for power-cycle. No CLIs will work henceforth. Reload or power cycle to make switch fully functional. Any release prior to NOS 7.1.0: Root level command. root> shutdowncmdb 2018/09/20-20:58:29 : shutdowncmdb : Shutting Down Database (New)
--------------------	--

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

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

Defect ID:	DEFECT000663418		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	ACLs - Access Control Lists
Reported In Release:	NOS6.0.2	Technology:	Security
Symptom:	Switch rebooted multiple times due security daemon termination during firmware upgrade from 5.0.1d to 6.0.2e.		
Condition:	Firmware upgrade from 5.0.1d to 6.0.2e		

Defect ID:	DEFECT000663500		
Technical Severity:	Medium	Probability:	High
Product:	Extreme Network OS	Technology Group:	YANG

Reported In Release:	NOS7.0.1	Technology:	Network Automation and Orchestration
Symptom:	NOS fails to un-escape special characters in passwords received via Netconf XML for config backup upload.		
Condition:	Special characters used in password		

Defect ID:	DEFECT000663965		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	Hardware Monitoring
Reported In Release:	NOS7.0.2	Technology:	Monitoring
Symptom:	Started supporting Extreme optics on VDX devices.		
Condition:	Extreme optics qualification on VDX devices.		

Defect ID:	DEFECT000665059		
Technical Severity:	Medium	Probability:	High
Product:	Extreme Network OS	Technology Group:	SNMP - Simple Network Management Protocol
Reported In Release:	NOS7.2.0	Technology:	Management
Symptom:	SNMP polling for cpStatus OID returns incorrect values.		
Condition:	When SNMP get/walk request done for cpStatus OID.		

Defect ID:	DEFECT000665593		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	VRRPv2 - Virtual Router Redundancy Protocol Version 2
Reported In Release:	NOS7.0.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	Unable to ping some VRRP-E VIP address.		
Condition:	There is no external trigger. The internal FIB (Forwarding Information Base) was out of sync with ARPD data base.		

Defect ID:	DEFECT000665913		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	Static Routing (IPv4)
Reported In Release:	NOS7.3.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	BGP route not cleared in secondary node of VCS cluster		

Condition:	Remove ve interface
Recovery:	clear bgp neighbor

Defect ID:	DEFECT000666629		
Technical Severity:	Medium	Probability:	High
Product:	Extreme Network OS	Technology Group:	CLI - Command Line Interface
Reported In Release:	NOS6.0.2	Technology:	Management
Symptom:	"show system" CLI execution doesn't display 'Burned In MAC' of secondary MM in VDX 8770. Ex: Burned In MAC : MM1 [50:EB:1A:xx:xx:xx], MM2 []		
Condition:	"show system" CLI execution.		

Defect ID:	DEFECT000666699		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	ICMP - Internet Control Message Protocol
Reported In Release:	NOS7.4.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	VDX does not generate ICMP notification for IP MTU violation [trapped packets].		
Condition:	IP MTU violation		

Defect ID:	DEFECT000667646		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	IGMP - Internet Group Management Protocol
Reported In Release:	NOS7.2.0	Technology:	IP Multicast
Symptom:	IGMP Snooping can be enabled only on 512 with previous releases. From this release onward IGMP snooping can be enabled on 4000 vlans.		
Condition:	IGMP Snooping can be enabled only on 512 with previous releases.		

Defect ID:	DEFECT000668234		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	OSPFv3 - IPv6 Open Shortest Path First

Reported In Release:	NOS7.1.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	Device experienced sudden reload due to DCMd daemon termination.		
Condition:	Execution of "no ipv6 ospf cost" CLI command.		

Defect ID:	DEFECT000669001		
Technical Severity:	High	Probability:	Low
Product:	Extreme Network OS	Technology Group:	OpenStack Integration
Reported In Release:	NOS6.0.2	Technology:	Network Automation and Orchestration
Symptom:	Unexpected reload.		
Condition:	Due to low memory condition, when links flap during ISL formation.		
Workaround:	Replace/reseat optics/cable if ISL link flap persists		

Defect ID:	DEFECT000669062		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	Software Installation & Upgrade
Reported In Release:	NOS7.2.0	Technology:	Management
Symptom:	Firmware Download [sanity check] fails with error message as "ISSU is not supported to the target firmware version. Please specify coldboot option in the command-line for download."		
Condition:	Firmware Download from NOS 7.2.0 / NOS 7.2.0a to NOS 7.2.0b using ISSU will fail.		
Workaround:	Please use coldboot firmware download option.		

Closed with code changes for Network OS v7.2.0a

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of February 15th, 2018 in Network OS v7.2.0a.

Defect ID:	DEFECT000550982		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	SNMP - Simple Network Management Protocol
Reported In Release:	NOS5.0.1	Technology:	Management
Symptom:	Switch management port does not generate a ColdStart trap if a Management port is configured to acquire the IP address via DHCP.		

Condition:	when switch is configured to acquire IP address via DHCP, then we will observe this issue.
Workaround:	If IP is configured statically, the issue will not happen.

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

Defect ID:	DEFECT000596658		
Technical Severity:	High	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	Logical Chassis
Reported In Release:	NOS7.0.1	Technology:	VCS
Symptom:	Traffic getting dropped indefinitely after reload.		
Condition:	Due to /32 route functionality the packets are getting trapped twice (on local and remote leaf).		

Defect ID:	DEFECT000615778		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	Configuration Fundamentals
Reported In Release:	NOS6.0.2	Technology:	Management
Symptom:	snmp-server ? displays all Possible completions, here "view" display as "view Define an SNMPv2 MIB view" which is incorrect as it is also applicable to SNMP v3.		
Condition:	snmp-server ? displays all Possible completions		

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

Defect ID:	DEFECT000636143		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	LAG - Link Aggregation Group
Reported In Release:	NOS6.0.2	Technology:	Layer 2 Switching
Symptom:	Cosmetic issue. some of fields (actor system id, Receive link count, Transmit link count, Individual and ready) won't display properly at "show port-channel detail nomore" output.		
Condition:	Rare scenario. Execution of "show port-channel detail nomore".		

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

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

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

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

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

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

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

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

Defect ID:	DEFECT000646908		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	Hardware Monitoring
Reported In Release:	NOS7.1.0	Technology:	Monitoring
Symptom:	The source IP for SNMP traps is not deterministic.		
Condition:	When VCS virtual IP address is configured and SNMP traps are enabled.		

Defect ID:	DEFECT000647389		
Technical Severity:	Medium	Probability:	High
Product:	Extreme Network OS	Technology Group:	IP Addressing
Reported In Release:	NOS6.0.2	Technology:	Layer 3 Routing/Network Layer
Symptom:	CLI prohibits user from adding multiple /31 subnets under L3 interfaces.		
Condition:	Configuring multiple /31 subnets under L3 interfaces.		

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

Defect ID:	DEFECT000647433		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	IP Fabric
Reported In Release:	NOS7.1.0	Technology:	Data Center Fabric
Symptom:	L2 VNI and tunnel IP value in the BGP route update is set to "zero".		

Condition:	In IP fabric topology, when a route-map with set condition is applied to evpn peer.
-------------------	---

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

Defect ID:	DEFECT000648164		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	SNMP - Simple Network Management Protocol
Reported In Release:	NOS6.0.2	Technology:	Management
Symptom:	SNMP responding on VCS IPv6 instead of management IPv6 address.		
Condition:	Both MM/Chassis IPv6 and virtual vcs IPv6 addresses are configured.		
Workaround:	Have the management IPv6 configured latest		

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

Defect ID:	DEFECT000648357		
Technical Severity:	Critical	Probability:	High
Product:	Extreme Network OS	Technology Group:	Configuration Fundamentals
Reported In Release:	NOS7.2.0	Technology:	Management
Symptom:	REST POST/PUT/PATCH configuration fails and errors out		

Condition:	Issue in REST POST/PUT/PATCH methods if payload has space.
Workaround:	DO not add space in payload

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

Defect ID:	DEFECT000648729		
Technical Severity:	Medium	Probability:	Medium
Product:	Extreme Network OS	Technology Group:	OSPF - IPv4 Open Shortest Path First
Reported In Release:	NOS7.2.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	OSPF vulnerabilities CVE-2017-3224, CVE-2017-3752, CVE-2017-6770		
Condition:	Existing code has above vulnerabilities in OSPF.		

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

Defect ID:	DEFECT000649847		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	Access Gateway
Reported In Release:	NOS6.0.1	Technology:	Management
Symptom:	VDX experiences unexpected reload due to memory leak in AG daemon.		
Condition:	When AG mode is enabled.		

Defect ID:	DEFECT000650040		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	RAS - Reliability, Availability, and Serviceability

Reported In Release:	NOS6.0.2	Technology:	Monitoring
Symptom:	Suddenly edge port admin down without user/admin action.		
Condition:	CRC error hits to threshold limits.		
Recovery:	Amin no shut		

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

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

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

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

Recovery:	May try 1gb if link partner supports it. Recommend upgrade VDX firmware for fix.
------------------	--

Defect ID:	DEFECT000652749		
Technical Severity:	High	Probability:	High
Product:	Extreme Network OS	Technology Group:	BGP4 - IPv4 Border Gateway Protocol
Reported In Release:	NOS7.2.0	Technology:	Layer 3 Routing/Network Layer
Symptom:	BGP neighbor entries are not created as expected.		
Condition:	BGP Auto neighbor discovery using LLDP on IP Unnumbered interfaces.		

Defect ID:	DEFECT000652894		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	Logical Chassis
Reported In Release:	NOS6.0.2	Technology:	Data Center Fabric
Symptom:	Unexpected reload.		
Condition:	Execution of CLI(vcs replace rbridge-id) during the cluster re-join.		
Workaround:	Avoid the CLI during cluster re-join		

Defect ID:	DEFECT000653244		
Technical Severity:	Medium	Probability:	Low
Product:	Extreme Network OS	Technology Group:	CLI - Command Line Interface
Reported In Release:	NOS6.0.2	Technology:	Management
Symptom:	Displaying generic error message.		
Condition:	When ftp and sftp fails, displaying common error message.		

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

Defect ID:	DEFECT000655415		
Technical Severity:	Medium	Probability:	Low

Product:	Extreme Network OS	Technology Group:	VXLAN - Virtual Extensible LAN
Reported In Release:	NOS7.0.1	Technology:	Layer 2 Switching
Symptom:	PBR is applied to only some flows, when it's configured on Ve that terminated VxLAN.		
Condition:	PBR configuration on Ve that terminated VxLAN.		

Defect ID:	DEFECT000659451		
Technical Severity:	Low	Probability:	Low
Product:	Extreme Network OS	Technology Group:	Scripting
Reported In Release:	NOS7.2.0	Technology:	Network Automation and Orchestration
Symptom:	A new IP Fabric Underlay and Overlay configuration Automation python script is introduced.		
Condition:	Automate the IP Fabric configuration using a single command/script.		

Closed with code changes for Network OS v7.2.0

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of July 26, 2017 in Network OS v7.2.0.

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
Symptom: Peer group configuration is not accepting update-source of IPv6 address	
Condition: Peer group configuration with update-source of IPv6 address	

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: DEFECT000584534	
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: 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: DEFECT000592597	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.0	Technology: Software Installation & Upgrade
Symptom: Allowing for N+2 version upgrade with default config.	
Condition: It's a RFE to allow upgrade of N+2 version.	

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

Defect ID: DEFECT000593537	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	Technology: IP Addressing
Symptom: Host ARP is learnt even when host IP subnet does not match to VE IP subnet.	
Condition: Host is connected to a VLAN where the Ve IP subnet is different than the host IP subnet.	
Workaround: Disable proxy ARP on VE	

Defect ID: 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: DEFECT000601293	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: VXLAN - Virtual Extensible LAN
Symptom: COS Priority tag frames egressed as Untagged frames	
Condition: Over VxLAN tunnel COS Priority tag frames are egressed as untagged frames.	

Defect ID: DEFECT000602319	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS5.0.2	Technology: LAG - Link Aggregation Group
Symptom: BPDU packets creates loop in network and the network can become unstable.	
Condition: NOS5.x with Protected vLAG BPDU received on Backup PO were flooded out from Active PO.	

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

Defect ID: DEFECT000604338	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.0	Technology: xSTP - Spanning Tree Protocols

Symptom: After the VDX switch was reloaded, the configured MSTP hello time was ignored and uses default value of 2 seconds. The value configured in the system does not change, NOS simply ignores it and uses default value.
Condition: The configured the MSTP hello time is not persistent after VDX switch was reloaded.
Recovery: The configured MSTP hello-time is persistent after the VDX switch was reloaded. That is if MSTP hello-time was configured for example set to 5. The value was applied after the VDX was reloaded.

Defect ID: DEFECT000606090	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS5.0.2	Technology: Logical Chassis
Symptom: Sudden reload due to software daemon termination	
Condition: In a large scale cluster, if the Rbridges are not reachable and vLAG config is done, software lockup happens which leads to software daemon termination.	
Workaround: Monitor for the state of fab_mct if the state is unavailable then we need to reload the switch to recover <pre> /fabos/cliexec/fab_vcs mct_state /fabos/cliexec/fab_vcs fcs_state /fabos/cliexec/fab_vcs vrrp_state /fabos/cliexec/fab_vxlan vxlan_state </pre> <p>If the state is stuck then the workaround is to reload/Chassis disable and enable</p>	
Recovery: the workaround is to reload/Chassis disable and enable	

Defect ID: DEFECT000609676	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.1.0	Technology: Management GUI
Symptom: Raslog for thermal shutdown cancellation may indicate an incorrect temperature.	
Condition: Extensive changes were made to the thermal policy in this release.	
Workaround: This is cosmetic only.	

Defect ID: DEFECT000611018	
Technical Severity: High	Probability: Low

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

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

Defect ID: DEFECT000611625	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS7.1.0	Technology: OpenStack Integration
Symptom: VDX 6740 4x10 GbE port may go offline after firmware upgrade to NOS Release 7.1.0 due to unstable 4x10 GbE link. User may see FFDC excessive interrupts on the problem port resulting in the port going offline.	
Condition: VDX 6740 4x10 GbE port may go offline after firmware upgrade to NOS Release 7.1.0 due to unstable 4x10 GbE link. User may see FFDC excessive interrupts on the problem port resulting in the port going offline.	
Workaround: Perform shut / no shut to bring the port back online.	
Recovery: Perform shut / no shut to bring the port back online.	

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

Defect ID: DEFECT000614000	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS4.1.3	Technology: Logical Chassis

Symptom: After invoke "copy support" CLI, the CLI will block, not return to user prompt for more than 2 hours, then the switch will go reboot.
Condition: The failure shows on early NOS release, and is very rare to happen. So far it never failure the same in the field.r

Defect ID: DEFECT000615746	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS5.0.1	Technology: ACLs - Access Control Lists
Symptom: VDX experience unexpected reload after configuring permit statement on standard ACL applied to management interface.	
Condition: Configuration of permit statement on standard ACL applied to management interface.	

Defect ID: DEFECT000617197	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.1.0	Technology: Logical Chassis
Symptom: Warning logs with tag 'FFDC' is displayed during cluster wide reload.	
Condition: Cluster reload.	

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

Defect ID: DEFECT000617887	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.0.1	Technology: IP Fabric
Symptom: On upgrade from 6.x to 7.x, one of the python CLI libraries may not be carried forward with all the changes & might impact some of the python scripts.	
Condition: Under certain unknown condition, when upgrading from 6.x to 7.x.	
Recovery: Copy CLI.py file manually to restore the script function.	

After restoration, test using below CLI & it should appear as below with "splitlines()" instead of "split()":

```
sw0:FID128:root> grep -A 2 get_output
/etc/fabos/Dcmd/python/CLI.py def get_output(self):
    return (self.output.splitlines())
```

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

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

Defect ID: DEFECT000619578	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.1.0	Technology: ACLs - Access Control Lists
Symptom: Unexpected reload.	
Condition: When below command is used for viewing the enforced IP ACL's: "show access-list interface Management <interface id> in " Example: "show access-list interface Management 1/0 in"	
Workaround: Use below commands to view the enforced policy ' show running config ip(v6) access-list' 'show running-config interface management'	

Defect ID: DEFECT000621402	
Technical Severity: High	Probability: High

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

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

Defect ID: DEFECT000622093	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.1.0	Technology: Logical Chassis
Symptom: Line card goes into fault state. one of the line card go to faulty state.	
Condition: After firmware download from 6.0.2c to 7.1.0 one of the line cards go to faulty state with reason code 119. What should the customer do now?	

Defect ID: DEFECT000623066	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: When "AG" was enabled, the SNMP traps for login/logout are not sent out.	
Condition: The issue is happened when "AG" was enabled.	
Workaround: Disable "AG" mode.	

Defect ID: DEFECT000623805	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring

Reported In Release: NOS6.0.2	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: I2traceroute request failure via REST API	
Condition: After 1000 successful I2traceroute request, any subsequent request for I2traceroute will fail.	
Recovery: Reload of the switch recovers from this state	

Defect ID: DEFECT000624561	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS7.1.0	Technology: OpenStack Integration
Symptom: Rebooting host connected to VDX 6940 4x10 GbE breakout port may cause one or more 4x10 GbE ports to become unstable, which could result in port faulting with FFDC excessive interrupts on the port(s).	
Condition: Rebooting host connected to VDX 6940 4x10 GbE breakout port may cause one or more 4x10 GbE ports to become unstable, which could result in port faulting with FFDC excessive interrupts on the port(s).	
Workaround: Perform shut / no shut on the port to bring it back online	
Recovery: Perform shut / no shut on the port to bring it back online.	

Defect ID: DEFECT000624714	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Exceedingly rare (not reproduced so far) error during failover that will cause the system to become faulted.	
Condition: Exceedingly rare error during failover.	
Recovery: Reboot faulted system.	

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

Defect ID: DEFECT000624872	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.1.0	Technology: Zoning
Symptom: In certain cases when joining RBridges together, if one of the RBridges has an empty zone configuration with a default zone mode set to All Access and the other RBridge being joined has an effective zone configuration but a default zone mode set to No Access, this will result in a zone conflict and the cluster will not form.	
Condition: This issue can be seen when joining RBridges together that have mismatched default zoning policies.	
Recovery: Once in this state, to recover, the customer will need to change the default zone policies so that they match across the RBridges that are being joined and then reattempt joining them together.	

Defect ID: DEFECT000624921	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.1.0	Technology: MAC Port-based Authentication
Symptom: After executing "no switchport" on a physical interface with "dot1x mac-auth-enable", dot1x mac-auth-bypass cannot be configured.	
Condition: This scenario occurs only if "dot1x mac-auth-enable" is configured while executing "no switchport"	
Workaround: Remove "dot1x mac-auth-enable" configuration before doing "no switchport"	

Defect ID: DEFECT000625243	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	Technology: IP Addressing
Symptom: "show ip in ve <>" does not show "ip icmp address mask" enabled/disabled status.	
Condition: Execution of "show ip in ve <>" CLI.	

Defect ID: DEFECT000625263	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.1.0	Technology: VLAN - Virtual LAN
Symptom: system may go for unexpected reload	
Condition: mac-authentication is enabled with more than 1500 source streams getting authenticated with layer 2 loop existing in the network. Here also loop detection is disabled.	

Defect ID: DEFECT000625670	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.1.0	Technology: Software Installation & Upgrade
Symptom: On rare occasions a SW error may be seen during HA synchronization.	
Condition: Rare occurrence that is not linked to a specific change.	

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

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

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

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

Defect ID: DEFECT000626685	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS7.1.0	Technology: PIM - Protocol-Independent Multicast
Symptom: User may observe Mcastssd process termination which leads to system reload	
Condition: When a PIM enabled node becomes the local RP by configuration and user removes the RP configuration or if the RP is changing based on the RP set message while BSR running, we may see this crash.	

Defect ID: DEFECT000626825	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: NOS7.1.0	Technology: OpenStack Integration
Symptom: VDX 6740 and 6940 may have linkup issues and/or CRCs after dynamically configuring 4x10GbE breakout mode.	
Condition: VDX 6740 and 6940 may have linkup issues and/or CRCs after dynamically configuring 4x10GbE breakout mode. Probability is increased if transceiver is changed from optical to copper or vice versa.	
Workaround: Reload switch to recover the links.	
Recovery: Reload switch to recover the links.	

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

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

Defect ID: DEFECT000627507	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.2	Technology: IP Fabric
Symptom: VDX can experience unexpected reload due to a daemon termination.	
Condition: When any telnet session is in middle of some CLI through pipe option and stays there for 7 days, VDX daemon will terminate all existing socket sessions and try to reconnect. Distributed module is not able to clear the stuck session and reconnect fails as a result after 1 hour it will get terminated.	

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

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

Defect ID: DEFECT000628238	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching

Reported In Release: NOS7.2.0	Technology: VXLAN - Virtual Extensible LAN
Symptom: VDX does not connect to NSX6.3.0 EA controllers.	
Condition: VDX to NSX6.3.0 EA controllers connection.	

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

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

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

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

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

Defect ID: DEFECT000631386	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.1.0	Technology: SNMP - Simple Network Management Protocol
Symptom: Unsupported snmp configuration like snmp v3 user exist in FIPS mode.	
Condition: FIPS mode has unsupported snmp configuration	

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

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

Defect ID: DEFECT000631759	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.2.0	Technology: Logical Chassis

Symptom: 100Mbps speed can be set on VDX6940-144S with CuSFP P/N 57-1000042-02, 57-1000042-01.
Condition: Configuring 100Mbps speed on VDX6940-144S with CuSFP P/N 57-1000042-02, 57-1000042-01.

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

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

Defect ID: DEFECT000632769	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.1.0	Technology: Security Vulnerability
Symptom: Executing show commands(listed in defect) is allowed by default(for non-admin roles)	
Condition: Executing show commands(listed in defect) for non-admin roles	

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

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

Defect ID: DEFECT000633740	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.0.1	Technology: LAG - Link Aggregation Group
Symptom: Port-channels in VDX switch may become stand alone and they will not be part of port channel redundancy group. This can cause a loop between VDX switches and other devices connected to edge ports in VCS.	
Condition: A loop can occur in VCS environment.	

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

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

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

Defect ID: DEFECT000634191	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.1.0	Technology: VLAN - Virtual LAN
Symptom: Both VDXes in VCS were not learning and MAC addresses and just flooding traffic, even when static mac entries were configured. Crash occur followed by rolling reboot.	
Condition: The issue was noticed after usually after a firmware upgrade	
Workaround: The problem was resolved when both switches reloaded together or Disable MAC-consistency Feature	

Defect ID: DEFECT000634192

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

Defect ID: DEFECT000634220	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.1.0	Technology: LAG - Link Aggregation Group
Symptom: When a protected group's configured active port channel comes back up, on rare occasions, the newly recovered active port channel also sends protected port MAC movement notification frames for MAC addresses learned on protected member port channels.	
Condition: protected port-channel undergoes switchover	

Defect ID: DEFECT000634222	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS7.1.0	Technology: LAG - Link Aggregation Group
Symptom: MAC movement notification frames are not sent.	
Condition: When a PRG configured active port channel recovers, if any physical interface other than the numerically lowest recovers first.	

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

Defect ID: DEFECT000634370	
Technical Severity: Medium	Probability: Medium

Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.1.0	Technology: Configuration Fundamentals
Symptom: show running-config rbr route-map output is being sorted based on the action (permit or deny) and then the sequence number instead of just sequence number.	
Condition: Execution of "show running-config rbr route-map" CLI.	

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

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

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

Defect ID: DEFECT000634769	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.2	Technology: Configuration Fundamentals
Symptom: SCP file transfer fails	
Condition: Using double quotes ("") for file names with spaces causes SCP to fail on certain servers.	

Workaround: Use file names without spaces

Defect ID: DEFECT000634904	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.2	Technology: VLAN - Virtual LAN
Symptom: show vlan brief and show vlan XXX command outputs may show inconsistent data.	
Condition: Around 4000 vlans are associated with switchport	
Defect ID: DEFECT000635101	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: NOS6.0.1	Technology: IGMP - Internet Group Management Protocol
Symptom: Memory leak in igmpd.	
Condition: When debug igmpd command is enabled and leads to error condition "Illegal multicast group address".	

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

Defect ID: DEFECT000635411	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS7.1.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Routing Administrative distance not honored on 6940 border leaf using EVPN and BGP	
Condition: Routing Protocol running on 6940 border leaf using EVPN and BGP.	

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

Defect ID: DEFECT000635760	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS6.0.2	Technology: Logical Chassis
Symptom: Spanning tree state is shown as disabled on Port channel	
Condition: VLAN is added newly to the port-channel after firmware upgrade	

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

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

Condition: Inserting Extreme optic with part number 33210-100 on VDX running can hit the error.	

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

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

Defect ID: DEFECT000636651	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS6.0.2	Technology: IPv6 Addressing
Symptom: IPv6 Ping does not work sometimes	
Condition: When VDX not trying to do IPv6 ND for unknown addresses and dropping packets instead, when both source and destination subnets are configured on the same VE interface	

Defect ID: DEFECT000636696	
Technical Severity: Medium	Probability: Medium

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

Defect ID: DEFECT000636734	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.1.0	Technology: Security Vulnerability
Symptom: Openssl vulnerability CVE-2016-2177, CVE-2016-2180, CVE-2016-2183, CVE-2016-2182	
Condition: Openssl vulnerability	

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

Defect ID: DEFECT000637538	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.1.0	Technology: IP Fabric

Symptom: In a IP fabric topology, the leaf (or border leaf) does not forward L3 packets (e.g. ICMP) destined to hosts on to peer leaf nodes. This causes the end to end L3 traffic between leaf nodes to fail.
Condition: The prefix route learnt over evpn-bgp was incorrectly programed in Linux kernel, thereby forwarding traffic to wrong destination.
Recovery: Issuing following command at NOS should help to recover from the problem state. clear arp no-refresh vrf <vrf-name>

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

Defect ID: DEFECT000637753	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS7.1.0	Technology: Configuration Fundamentals
Symptom: VDX returns "<html><body><h1>401 authentication needed</h1></body></html>" response.	
Condition: Proxied REST configuration request using persistent connection Authentication-Token.	

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

Defect ID: DEFECT000637857	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: NOS6.0.2	Technology: VLAN - Virtual LAN
Symptom: VDX do not learn MAC addresses.	
Condition: VDX stops learning MAC addresses when specific configuration exist and it has reached to certain number of lines in configuration.	
Workaround: <ol style="list-style-type: none"> 1. configure "mac-address-table consistency-check suppress" 2. reload the affected switches or 1. configure dummy vlan 2. Assign it to interface/Po 3. reload 	
Recovery: The same above steps works we may expect unexpected reloads during this	

Defect ID: DEFECT000638081	
Technical Severity: Critical	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: traffic may flood and result in high latency.	
Condition: layer 2 loop exists in the topology resulting in high number of interrupts to VDX	

Defect ID: DEFECT000638990	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: NOS7.0.1	Technology: Logical Chassis
Symptom: After a reload, a VDX device is not able to join a cluster and is stuck in the "Awaiting Rejoin" state.	
Condition: VDX has "switchport port-security" CLI configured and a reload occurs in the VCS cluster.	

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

Defect ID: DEFECT000639403	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: NOS6.0.2	Technology: Port Mirroring
Symptom: Monitor traffic may not appear on the monitor ports	
Condition: Two monitor sessions are created with same destination but different source and first monitor session is removed.	

Defect ID: DEFECT000639524	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.1.0	Technology: Logical Chassis
Symptom: 1. Out of memory(OOM) and subsequent switch reload. 2. toamd process termination and subsequent switch reload	
Condition: VDX supposed to receive the packet or short length (<64 byte) packet though UDP port 35512	
Workaround: ACL program to block UDP port 35512	

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

Defect ID: DEFECT000641075	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: NOS7.1.0	Technology: ACLs - Access Control Lists
Symptom: Management Interface drops ICMP & UDP for permitted hosts.	
Condition: Standard ACL applied inbound on Management Interface drops ICMP & UDP for permitted hosts.	

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

Defect ID: DEFECT000642278	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.2	Technology: SNMP - Simple Network Management Protocol
Symptom: Snmpwalk of in-band Ve interfaces fails with timeout error.	
Condition: SNMP packets are discarded when the ingress interface differs from the egress interface.	

Defect ID: DEFECT000642711	
Technical Severity: Critical	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: NOS5.0.2	Technology: BFD - BiDirectional Forwarding Detection
Symptom: BGP & BFD is not working as expected.	
Condition: Bringing down the loopback interface and VDX has "bfd interval 300 min-rx 300 multiplier 3" configuration under interface.	

Defect ID: DEFECT000642983	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: NOS7.1.0	Technology: IP Fabric

Defect ID: DEFECT000643124	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration

Reported In Release: NOS7.2.0	Technology: OpenStack Integration
Symptom: VDX 6740 or 6940 40g port connected to SLX 9xxx 40g port, and then rebooting or power-cycling the SLX switch. Upon reboot, the 40g link may not come online	
Condition: VDX 6740 or 6940 40g port connected to SLX 9xxx 40g port, and then rebooting or power-cycling the SLX switch. Upon reboot, the 40g link may not come online	
Workaround: Try NOSCLI shut / no shut on VDX and or SLX switch.	
Recovery: Try NOSCLI shut / no shut on VDX and or SLX switch.	

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

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

Defect ID: DEFECT000644087	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: NOS6.0.0	Technology: Software Installation & Upgrade
Symptom: VDX6940-2U with new DRAM may encounter machine-checks errors during or after FWDL and can cause unexpected switch reload.	
Condition: VDX6940-2U with new HW component [DRAM]	
Workaround: When a newer HW component has been detected, firmware download is being blocked by pre-install script to avoid degrading the system performance of BR-VDX6940-144S platform.	

Recovery: Perform FWDL to a newer NOS version where new uboot change does exist.

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

Known Issues for Network OS v 7.2.0e

This section lists open software defects with Critical, High, and Medium Technical Severity as of July 16, 2020 in Network OS v 7.2.0e.

Symptom	SSH connection fails after upgrading from NOS 7.2.0e to NOS 7.3.0x
Workaround	Use TELNET to login.