



December 2020

Network OS 7.4.1b for Extreme VDX

Release Notes

Copyright Statement

© 2020, 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.....	7
Preface.....	8
Contacting Extreme Technical Support.....	8
Document feedback.....	8
Overview.....	10
New devices.....	10
New interface modules.....	10
Deprecated Hardware.....	10
Software Features.....	11
Deprecated Software Features.....	11
New Software Features for Network OS v7.4.1b.....	11
New Software Features for Network OS v7.4.1a.....	11
Fix exposure to netkit-telnetd vulnerability CVE-2020-10188.....	11
SFM python script:.....	11
New Software Features for Network OS v7.4.1.....	11
DNS Timeout Service.....	11
NOS DHCPv4 Relay SAG IP Support.....	11
TCP and ICMP Filter Enhancements.....	12
Multiple Primary IPv6 Support.....	12
Line Card Behavior when BE Link errors are detected on ASIC.....	12
Password Encryption Policy: SHA-512 Support.....	13
Thermal policy change on CasatorT.....	15
CLI Changes.....	15
New Commands for Network OS v7.4.1b.....	15
New Commands for Network OS v7.4.1a.....	15
New Commands for Network OS v7.4.1.....	15
Modified Commands for Network OS v7.4.1b.....	16
Modified Commands for Network OS v7.4.1a.....	16
Modified Commands for Network OS v7.4.1.....	16
Deprecated Commands for Network OS v7.4.1b.....	16

Deprecated Commands for Network OS v7.4.1a.....	16
Deprecated Commands for Network OS v7.4.1.....	16
API Changes.....	16
Newly supported standards and RFCs	17
Software Upgrade	17
HARDWARE SUPPORT	18
Supported devices	18
Supported power supplies	22
Supported Optics for Network OS v7.4.1.....	24
Extreme Branded Optics	26
10GBase-T Copper SFP+	27
SOFTWARE UPGRADE AND DOWNGRADE.....	29
Image filenames.....	29
Upgrade/Downgrade considerations.....	29
NOS7.4.1-SHA512 upgrade/downgrade basics:	30
Migration Path	31
Management IP connectivity	34
Firmware Installation	35
Upgrading to this Release (Best Practices)	36
Downgrading to a Previous Release	36
Upgrade/downgrade Considerations for vLAG deployments.....	37
Upgrade/downgrade with default configuration	39
Management Gateway IP changes	39
Management Services.....	40
Other Management Services	43
Scalability numbers.....	44
Compatibility and Interoperability.....	53
IP Storage	53
Limitations and Restrictions.....	54
Command Line Interface.....	54
Line cards	56
USB.....	56
Licensing.....	56
Logical Chassis.....	58

Extreme Trunks	58
Breakout Interfaces.....	59
Dual-personality Ports	59
1G Mode	59
vLAG	60
Virtual IP Address Support.....	61
Security, Management ACLs, Authentication, Authorization	61
SPAN & RSPAN	62
MAC Learning Considerations in VCS.....	62
PVLAN.....	62
UDLD	63
STP/DiST	63
IGMPv3 Snooping.....	64
Edge Loop Detection (ELD)	64
Long Distance ISL Ports	64
AMPP and Port-Profiles	65
vCenter.....	66
QoS.....	66
IP Fabric.....	67
ND/RA	68
BFD	69
VRRP.....	70
OSPFv2	71
OSPFv3	71
\BGP	71
Layer 2/Layer 3 Multicast	71
VRF	71
ACL	72
Policy-based Routing (PBR).....	72
Inter-VRF Leaking (Static).....	72
DHCP IP Helper.....	73
Dynamic ARP Inspection (DAI)	73
DHCP-based Firmware Download (DAD – DHCP Automatic Deployment).....	74
Link State Tracking	74

OpenFlow	75
Uplink Switch Support.....	76
Layer 2 and Layer 3 ISSU on VDX 6740x.....	77
REST API	77
NetConf.....	77
VXLAN Gateway for VMware NSX.....	77
VF Extension using VxLAN.....	78
TCAM Profiles.....	78
Management VRF.....	79
Conversational MAC Learning.....	79
System level Flowbased QoS	79
Port level Flowbased QoS	79
URPF.....	79
BGP Auto neighbor discovery	80
Non-trivial Merge.....	80
HA on TOR switches.....	80
Logical Chassis HA	80
Interoperability	81
MAPS.....	81
Maintenance Mode.....	81
LACP and individual ports for PXE boot enhancement.....	82
Miscellaneous	82
Defects	84
TSBs - Critical Issues to Consider Prior to Installing This Network OS Release.....	84
.Closed with code changes for Network OS v7.4.1b.....	85
Closed with code changes for Network OS v7.4.1a.....	86
Closed with code changes for Network OS v7.4.1.....	87
Closed with code changes for Network OS v7.4.0a.....	92
Closed with code changes for Network OS v7.4.0.....	97
Closed without code changes for Network OS v7.4.0	148
Known Issues for Network OS v7.4.0.....	157

Document History

Version	Summary of Changes	Publication Date
1.0	Initial Release Removed versions older than Network OS v7.4.0a	December 2020

Preface

Contacting Extreme Technical Support

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

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

- GTAC (Global Technical Assistance Center) for immediate support
- Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.
- Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- GTAC Knowledge - Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- The Hub - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- Support Portal - Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

Document feedback

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

You can provide feedback in two ways:

- Use our short online feedback form at <http://www.extremenetworks.com/documentation-feedback-pdf/>
- Email us at internalinfodev@extremenetworks.com

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)

- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Overview

NOS 7.4.1 is a major software release for VDX6740, VDX6940 and VDX8770. NOS 7.4.1b is an update to the NOS 7.4.1 release.

Hardware

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

New devices

None

New interface modules

None

Deprecated Hardware

None

Software Features

For information about Network OS v7.4.0a and earlier releases, please refer to the [Network OS v7.4.0a Release Notes](#).

The following section lists new, modified, and deprecated software features for Network OS 7.4.1a. For information about which platforms support these features, refer to the *Network OS Features and Standards support Matrix*.

Deprecated Software Features

- None

New Software Features for Network OS v7.4.1b

There are no new software features for NOS7.4.1b.

New Software Features for Network OS v7.4.1a

The following software features are new or enhanced in this release:

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

SFM python script:

- This feature detect Serial Commination Interface (SCI) H/W failure on the SFM and power off the SFM slot and Generate a RASLOG message to expose the criticality. Programmable Event Manager(PEM) is used to do script-based configuration on the switch and Event-Handlers allow event-based Python script execution.

Python script will power of SFM and generate RASLOG if SFM detect SCI error and throw RASLOG 5010.

Limitations and Caveats

[New Raslog Event EM-1102: SFM2 Powered off due to SCI H/W error. Reseat SFM.](#)

Symptom -

A rare scenario Management module going to reboot loop when SFM detect SCI error
2020/09/23-01:37:54:528225, [EM-1102], 89/12, M2 | Active, CRITICAL, VDX8770-8, SFM2 Powered off due to SCI H/W error. Reseat SFM. if it is not recovered SCI error, replace SFM ,
raslog_test.c, line: 345, comp:rltd, ltime:2020/09/23-01:37:54:52799

Condition -

If the SFM detects the SCI H/W error.

Recovery -

Reseat the SFM. If it is not recovered, replace the SFM.

Event handler trigger for SFM's RASLOG 5010 and Power off SFM. It also generate RAS log EM-1102 to take necessary action to restore SFM service.

New Software Features for Network OS v7.4.1

The following software features are new or enhanced in this release:

DNS Timeout Service

- New CLI commands for DNS service provided to configure timeout and retry parameter to achieve quick turnaround/fallback to next server while resolving given domain name in case the first DNS server not reachable.

NOS DHCPv4 Relay SAG IP Support

- This feature enables DHCPv4 Relay functionality with Anycast Gateway Address (SAG) in NOS7.4.1

release.

TCP and ICMP Filter Enhancements

- This feature gives users the ability to filter packets by ICMP type and code using the IPv4 ACL function in the VDX (NOS) switch. Extended IPv4 ACL policies can be created to filter traffic based on its ICMP message type and code. Only the description of the message type will be given in the CLI. NOS 7.4.1 release supports only echo-request and echo-reply.

Multiple Primary IPv6 Support

- Allows configuration of multiple primary IPv6 addresses with different subnets, thereby the user can create IPv6 addresses with different subnets on the same L3 interfaces. NOS software gives overlap error message when user attempts to configure the same subnet IPv6 address which is already configured on the same interface or on different interfaces.

Line Card Behavior when BE Link errors are detected on ASIC

- This feature triggers blade fault or generates user visible RASLOG upon exceeding the credit loss recovery threshold during a 24 hour window. The user has the choice to use either of the two options.

The objective of this feature is to fault the blade or generate only a user visible RASLOG when credit recovery happens exceeding the credit loss recovery threshold in a 24 hour window. If credit losses detected are not resulting in traffic black holing, no action is to be taken. The user can choose whether to generate only the RASLOG or to fault the blade as well. If credit losses detected result in traffic black holing user visible RASLOG must be printed and E2E credit recovery mechanism must be triggered. If credit recovery happens more than the recovery threshold in a 24 hour window, fault the blade. By default, this threshold is set as 2. The first time that excessive credit loss and recovery is detected the LC is faulted. Status of LC shows "FAULTED". LC auto-recovery happens with a delay of 150secs. On recovery, LC status should be back to "ENABLED". In the 24 hour window, if excessive E2E credit loss and recovery is detected for the second time, LC will be faulted. Recovering LC will require the line card to be powered off and on manually.

We have added a polling routine at the Ingress ASIC driver that checks the VOQs head registers and the E2E credit every second. If the head register and E2E credit is not changing for X seconds, then it means traffic is not moving (or the lossless traffic has been back-pressured too long, which will be a false positive and customer needs to increase X). If there is a credit missing, we will initiate a credit recovery.

At the Egress ASIC driver, it monitors the number of credit recovery initiated by the Ingress ASIC driver on a per-Ingress basis. This is monitored on a per-Ingress basis because a single BE link error could cause all Ingress ASIC drivers to initiate recovery. If the number of credit recovery exceeds Y within period window of Z, then we RASLOG or fault the blade.

Programmable Event Manager(PEM) is used to do script-based configuration on the switch. 2 scripts are provided – one to generate RASLOG only and another to set the LC to faulty. RASLOG will be generated in both the case. Python CLI module in the PEM allows Python scripts to run and execute. Event-Handlers allow event-based Python script execution.

Details about the RASLOG generated :

When excessive credit recovery and loss is detected and event-handler is activated to trigger the RASLOG script:

```
2020/03/27-04:28:16, [HAWK-1005], 44579, L3/0 | Active | FFDC, ERROR, VDX8770-4, Chip6  
OID=0x43248880 excessive E2E credit lost and recovery detected.
```

When excessive credit recovery and loss is detected and event-handler is activated to trigger the script for faulting LC:

```
2020/03/26-12:34:23, [HAWK-1005], 43194, L3/0 | Active | FFDC, ERROR, VDX8770-4, Chip6
```

OID=0x43248880 excessive E2E credit lost and recovery detected.

2020/03/26-12:34:23, [EM-1034], 43196, M2 | Active, ERROR, VDX8770-4, L3 is set to faulty, rc=2006a.

2020/03/26-06:10:26, [EM-1100], 39030, M2 | Active | CFFDC, CRITICAL, VDX8770-4, Unit in L3 with ID 114 is faulted(106). 1 of 1 total attempt(s) at auto-recovery is being made. Delay is 150 seconds.

Limitations and Caveats

[New Raslog Event HWK2-1006 : E2E credit loss detection](#)

Symptom -

2020/05/08-06:44:04:874822, [HWK2-1006], 44373161/369865, L6/0 | Active | FFDC, ERROR, VDX8770-8, Chip0 OID=0x43308080 E2E credit lost detected., OID:0x0000020c, hawk2_e2e.c, line: 7846, comp:insmod, ltime:2020/05/08-06:44:04:871985

Condition -

ASIC detects normal E2E credit loss which will get replenished within microseconds.

Workaround -

HAWK/HAWK2 ASIC attempts to automatically replenish the credits lost using self-recovery logic.

Recovery -

E2E credit loss reported as part of HWK2-1006 raslog will not cause any hang situation for the traffic as ASIC always replenishes the credits within few microseconds.

[New Raslog Event HWK2-1005 : Excessive E2E credit loss detection](#)

Symptom -

2020/05/08-06:47:38:876160, [HWK2-1005], 44377197/369867, L6/0 | Active | FFDC, ERROR, VDX8770-8, Chip0 OID=0x43308080 excessive E2E credit lost and recovery detected., OID:0x0000020c, hawk2_e2e.c, line: 9135, comp:insmod, ltime:2020/05/08-06:47:38:872842

Condition -

If the ASIC hits an excessive credit loss recovery threshold within recovery window, user will be notified about the event through the raslog HWK2-1005.

Workaround -

LC attempts to auto recover from the first occurrence of excessive E2E credit loss by reloading the LC. The following raslog will be generated to report the LC level auto recovery.

2020/05/08-06:47:39:051374, [EM-1100], 44377202/369869, M2 | Active | CFFDC, CRITICAL, VDX8770-8, Unit in L6 with ID 150 is faulted(106). 1 of 1 total attempt(s) at auto-recovery is being made. Delay is 150 seconds., OID:0x43300000, em_board_lib.c, line: 2590, comp:emd, ltime:2020/05/08-06:47:39:051263

Credit recovery fault value will be reset after elapsing the recovery time window. These values can be modified by editing the script if customer is willing to change the parameter.

Recovery -

Action configured by the user which is mainly LC reload or generate raslog will be initiated. If the second instance of excessive E2E credit loss is detected after auto recovery described in the workaround section, LC will be set to Faulty. In order to recover from the LC faulty situation, user manually needs to recover it using slotpoweroff and slotpoweron.

Password Encryption Policy: SHA-512 Support

Password encryption policy

The software supports encrypting the passwords of all existing user accounts by enabling password encryption at the device level. By default, the encryption service is enabled.

The following rules apply to password encryption:

- When you enable password encryption, all existing clear-text passwords are encrypted. Subsequently, passwords that are added in clear text are stored in encrypted format.
- There are three levels of password encryption:
 - Encryption Level 0: No encryption, clear text
 - Encryption Level 7: AES-256 encryption
 - Encryption Level 10: SHA-512 salted HASH format. This is the default encryption level.
- In the following example, the testuser account password is created in clear text after password encryption is enabled. The global encryption policy overrides command-level encryption settings, and the password is stored as encrypted.

```
device(config)# service password-encryption
device(config)# do show running-config service password-encryption
service password-encryption
device(config)# username testuser role testrole desc "Test User" encryption-level 0 password hellothere
device(config)# do show running-config username
username admin password $6$mAog0c./JxVGulzy$6wFogQmek0KOEgTav.0DVYKzlvRodclUCAbipYft/DWnT5R6/
Y3qpq7V3JH1hRNvtwguLgXnzdtBDKPKaXbBg/encryption-level 10 role admin desc Administrator
username testuser password $6$78rhJxmF0zFKbhu4$0WvJVdRv7.ke07E5sL7m04stPw3XO9hgIxZ/
xArDpKCPk6eGtlCn0YBi3xRv856hoiDv8U9eAbxi6ZZNY4CiV/encryption-level 10 role testrole desc "Test User"
username user password $6$mAog0c./JxVGulzy$6wFogQmek0KOEgTav.0DVYKzlvRodclUCAbipYft/DWnT5R6/
Y3qpq7V3JH1hRNvtwguLgXnzdtBDKPKaXbBg/encryption-level 10 role user desc User
```

- When you disable the password encryption service, any new passwords added in clear text are stored as clear text on the device. Existing encrypted passwords remain encrypted.
- In the following example, the testuser account password is stored in clear text after password encryption is disabled. The default accounts, user and admin, remain encrypted.

```
device(config)# no service password-encryption
device(config)# do show running-config service password-encryption
no service password-encryption
device(config)# username testuser role testrole desc "Test User" encryption-level 0 password hellothere
enable true
device(config)# do show running-config username
username admin password $6$mAog0c./JxVGulzy$6wFogQmek0KOEgTav.0DVYKzlvRodclUCAbipYft/DWnT5R6/
Y3qpq7V3JH1hRNvtwguLgXnzdtBDKPKaXbBg/encryption-level 10 role admin desc Administrator
username testuser password hellothere encryption-level 0 role testrole desc "Test User"
username user password $6$mAog0c./JxVGulzy$6wFogQmek0KOEgTav.0DVYKzlvRodclUCAbipYft/DWnT5R6/
Y3qpq7V3JH1hRNvtwguLgXnzdtBDKPKaXbBg/encryption-level 10 role user desc User
```

- If you have passwords with encryption-level 7 on the device, then you can use the exec command **password-encryption convert-enc-to-level-10** to upgrade the passwords to encryption-level 10 (SHA-512 hash format), making the passwords more secure. After you run this command, all encryption-level 7 passwords are converted to encryption-level 10. However, if you downgrade to a release lower than NOS 7.4.1, these accounts will not be available.
- This command is available only to admin users. Any clear-text (encryption-level 0) passwords are retained as-is in the configuration database and not converted to encryption-level 10 (SHA-512 hash format). These clear-text passwords can be converted using the **service password-encryption configuration** command.
- In the following example, testuser1 has encryption-level 7, and after running the exec command, the encryption-level is changed to 10.

```

SLX# show running-config user | inc testuser
username testuser password "cONW1RQ0nTV9Az42/9uCQg==\n" encryption-level 7 role
testrole desc "Test User"
SLX# password-encryption convert-enc-to-level-10
%WARN:This operation will convert all existing user passwords to SHA-512 format.
However, the enc level 0 (clear-text) passwords, if any, will be retained as is in the
configurationdatabase. These configurations will be lost if the system is downgraded
to lower releases than SLX 20.1.1
Do you want to continue? [Y/N]y
All passwords are converted successfully.
SLX# show running-config user | inc testuser
username testuser password $6$gV7AS1DXqcGc8/ma
$MEVxe20jaBarALGhmSYw.p3oc9IXVj9xqNUGDnfNABGs.FAqwrM8EPDMvCJcZe/MsY9geY0ej01gma7mWWWTz0
encryption-level 10 role testrole desc "Test User"
SLX#

```

The exec command **password-encryption convert-enc-to-level-10** is not allowed if there is a configuration rollback in-progress.

Thermal policy change on CasatorT

The basics of the change:

CastorT basic thermal policy remains unchanged and depends on whether or not the HW is ECO or non-ECO (as well as air-flow direction). The highest PHY temperature is now the third temperature listed in "show env temp". If the max phy temp is > 70 the fan speed will go to HIGH regardless of what the normal thermal policy would indicate. If the max phy temp is < 70 the fan speed will be dictated by the normal thermal policy. If the max phy temp is > 95 a two minute shutdown will be initiated.

Summary of additions to current thermal policy:

- Max phy temp > 70 = HIGH Fan speed
- Max phy temp < 70 = fan speed dictated by regular thermal policy
- Max phy temp > 95 = two minute shutdown initiated

CLI Changes

For information about Network OS v7.4.0a and earlier releases, please refer to the [Network OS v7.4.0a Release Notes](#).

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 v7.4.1b

The following configuration commands are new in this release:

- None

New Commands for Network OS v7.4.1a

The following configuration commands are new in this release:

- None

New Commands for Network OS v7.4.1

The following configuration commands are new in this release:

- ip dns retry
- ip dns timeout
- lacp default-up

Modified Commands for Network OS v7.4.1b

The following commands have been modified in this release:

- None

Modified Commands for Network OS v7.4.1a

The following commands have been modified in this release:

- None

Modified Commands for Network OS v7.4.1

The following commands have been modified in this release:

- show ip dhcp relay option
- ip dhcp relay gateway

Deprecated Commands for Network OS v7.4.1b

The following configuration commands have been deprecated in this release:

- None

Deprecated Commands for Network OS v7.4.1a

The following configuration commands have been deprecated in this release:

- None

Deprecated Commands for Network OS v7.4.1

The following configuration commands have been deprecated in this release:

- None

API Changes

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

Newly supported standards and RFCs

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

- RFC 5280 – TLS client authenticating the server certificate
- RFC 6960 – TLS client authentication doing X.509v3 certificate revocation check dynamically using Online Certificate Status Protocol (OCSP)
- RFC 6187 - SSH authentication using X.509v3 certificates
- RFC 7474: OSPFv2 HMAC-SHA Cryptographic Authentication
- RFC 7166: Supporting Authentication Trailer for OSPFv3 instead of IPsec

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.

Software Upgrade

Note: [The Field Notice FN-2019-439 – “SFP breakout subinterfaces may be incorrectly set to administratively down in Network OS 7.4.0”](#) has been updated to state that, the fix for Defect NOS-67192 will be in Network OS 7.4.1 not Network OS 7.4.0a.

Please contact Extreme GTAC Support for the latest status for Network OS 7.4.1.

HARDWARE SUPPORT

Supported devices

Extreme Network OS v7.4.1 supports following VDX Switches:

- ExtremeSwitching VDX 6740-48
- ExtremeSwitching VDX 6740T
 - ExtremeSwitching VDX 6740T-64
 - ExtremeSwitching VDX 6740T-1G
- ExtremeSwitching VDX 6940-144S
- ExtremeSwitching VDX 6940-36Q
- ExtremeSwitching VDX 8770
 - ExtremeSwitching VDX 8770-4
 - ExtremeSwitching VDX 8770-8

Deprecated Devices

- None

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 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 8770-4 and VDX 8770-8

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

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

Supported Blades for VDX 8770

The flexible, modular switch design offers interconnection with other Extreme switches, traditional Ethernet switch infrastructures, and direct server connections. Modular 4-slot and 8-slot chassis options are available to match the switch to the needs of the organization. These include:

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

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

- 1 GbE: LC48×1G line card provides up to 48 SFP/SFP-copper ports.
- 10 GbE: LC48×10G line card provides up to 48 SFP+ ports .
- 10 GbE-T: LC48×10GT line card provides up to 48 RJ-45 ports .
- 40 GbE: LC12×40G line card provides up to 12 x 40 GbE QSFP ports.
- 40 GbE: LC27×40G line card provides up to 27 x 40 GbE QSFP ports.
- 100 GbE: LC6×100G line card provides up to 6 x 100 GbE CFP2 ports.

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, TRILL, PFC configuration are NOT supported on 100 Mbit interfaces.
- Examples for 100 Mbit/s usage are as follows:
 - 100 Mbit/s Host device requirement with IPv4/v6 Connectivity.

Supported power supplies

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

Part number	Description	Compatible devices
XBR-ACPWR-3000	FRU,3000W AC POWER SUPPLY	VDX 8770-4, VDX 8770-8
XBR-DCPWR-3000	FRU,3000W DC POWER SUPPLY	VDX 8770-4, VDX 8770-8
XBR-250WPSAC-F	FRU,250W,ACPS/FAN,NONPORTSIDE EXHAUST	VDX 6740
XBR-250WPSAC-R	VDX 6740 AC RTF PWR SUPPLY FAN	VDX 6740
XBR-250WPSDC-F	FRU,250W,DCPS/FAN,NONPORTSIDE EXHAUST	VDX 6740
XBR-250WPSDC-R	FRU,250W,DCPS/FAN,PORT SIDE EXHAUST	VDX 6740
XBR-500WPSAC-F	FRU 500W ACPS	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
XBR-500WPSAC-R	FRU 500W ACPS	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
RPS9DC+E	FRU,500W DC PSU PORT SIDE EXHAUST	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
RPS9DC+I	FRU,500W,DCPS/FAN,NONPORTSIDE EXHAUST	VDX 6740T, VDX 6740T-1G, VDX 6940-36Q
XBR-1100WPSAC-R	FRU,1100W PSAC,PORTSIDE EXHAUST AF	VDX 6940-144S
XBR-1100WPSAC-F	FRU,1100W PSAC,NON-PORT SIDE EXHAUST AF	VDX 6940-144S
XBR-1100WPSDC-01-R	FRU 1100W DCPS,PORTSIDE EXHAUST	VDX 6940-144S
XBR-1100WPSDC-01-F	FRU 1100W DCPS,NON PORTSIDE EXHAUST	VDX 6940-144S

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

ordered for redundancy. The VDX 8770-8 system ships with a minimum of 3 PSU and additional PSU may be ordered for redundancy:

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

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

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

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

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

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

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

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

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

Supported Optics for Network OS v7.4.1

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

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

Speed	FRU and Optics SKU	Description	Part Number
1GbE	XBR-000190 (1-pack)	1 GbE copper	57-1000042-01
	E1MG-SX-OM (1-pack)*	1000Base-SX	33211-100
	E1MG-SX-OM-8 (8-pack)*		
	E1MG-LX-OM (1-pack)*	1000Base-LX	33210-100
	E1MG-LX-OM-8 (8-pack)*		
10GbE	10G-SFPP-SR (1-pack)	10 Gbps SR	57-0000075-01
	10G-SFPP-SR-8 (8-pack)		
	10G-SFPP-LR (1-pack)	10 Gbps LR (10km)	57-0000076-01
	10G-SFPP-LR-8 (8-pack)		
	10G-SFPP-ER (1-pack)	10 Gbps ER (40km)	57-0000085-01
	10G-SFPP-ER-8 (8-pack)		
	10G-SFPP-ZR	10 Gbps ZR (80km)	57-1000180-01
	10G-SFPP-ZRD-T	10 Gbps tunable DWDM SFP+ (80km)	57-1000266-01
	10G-SFPP-TWX-0101 (1-pack)	1m Twinax copper cable	58-1000026-01
	10G-SFPP-TWX-0108 (8-pack)		
	10G-SFPP-TWX-0301 (1-pack)	3m Twinax copper cable	58-1000027-01
	10G-SFPP-TWX-0308 (8-pack)		
	10G-SFPP-TWX-0501 (1-pack)	5m Twinax copper cable	58-1000023-01
	10G-SFPP-TWX-0508 (8-pack)		
	10G-SFPP-AOC-0701	10GbE SFP+ Direct Attached Active Optical Cable, 7m, 1-pack	57-1000273-01
	10G-SFPP-AOC-1001	10GbE SFP+ Direct Attached Active Optical Cable, 10m, 1-pack	57-1000274-01
	10G-SFPP-USR	10GE USR SFP+ optic (LC), target range 100m over MMF, 1-pack	57-1000130-01
	10G-SFPP-BXD-S	10GBase-BXD 10GE Bidi Downstream	57-1000349-01
	10G-SFPP-BXU-S	10GBase-BXU 10GE Bidi Upstream	57-1000348-01
	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	
40GbE	40G-QSFP-QSFP-C-0101	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 1m, 1-pack	58-0000041-01
	40G-QSFP-QSFP-C-0301	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 3m, 1-pack	58-0000042-01
	40G-QSFP-QSFP-C-0501	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 5m, 1-pack	58-0000043-01
	40G QSFP+ -> 4x10G LR		57-1000477-01

Speed	FRU and Optics SKU	Description	Part Number
	40G-QSFP-4SFP-C-0101	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 1m, 1-pack	58-0000051-01
	40G-QSFP-4SFP-C-0301	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 3m, 1-pack	58-0000052-01
	40G-QSFP-4SFP-C-0501	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 5m, 1-pack	58-0000053-01
	40G-QSFP-SR4	40 GbE SR4 optic	57-1000128-01
	40G-QSFP-SR4-INT	40 GbE SR4 (4x10 GbE SFPP break-out capable) Breakout optical cable is not included with this optics	57-1000129-01
	40G-QSFP-SR-BIDI	40 GbE QSFP+ Bi-Directional 100m optics	57-1000339-01
	40G-QSFP-ESR4	40GBase-eSR4 QSFP+ optic (MTP 1x12) 300m over MMF, (10GBASE-SR compatible, breakout), 1-pack	57-1000296-01
	40G-QSFP-ER4	40 GbE 40Km optic	57-1000327-01
	40G-QSFP-LR4	40 GbE 10Km optic	57-1000263-01
	40G-QSFP-LM4	40 GbE 140m multi-mode or 2km single-mode optic	57-1000325-01
	40G-QSFP-QSFP-AOC-1001	40GE Direct Attached QSFP+ to QSFP+ Active Optical Cable, 10m, 1-pack	57-1000306-01
	40G-QSFP-4SFP-AOC-1001	4x10GE Direct Attached QSFP+ to 4 SFP+ Active Optical Breakout Cable, 10m, 1-pack	57-1000307-01
	8G FC	XBR-000163 (1-pack) XBR-000164 (8-pack)	8G FC SWL
XBR-000153 (1-pack) XBR-000172 (8-pack)		8G FC LWL	
XBR-000174		8G FC ELWL	
16G FC		XBR-000192 (1-pack) XBR-000193 (8-pack)	16G FC SWL
	XBR-000198 (1-pack) XBR-000199 (8-pack)	16G FC LWL	
	FC QSFP	XBR-000245 4x8G or 4x16G FC QSFP breakout. VDX 6740T, 6740T-1G only (not applicable for VDX 6740).	
100GbE	100G-CFP2-SR10 (1-pack)	100 GbE CFP2 optic, SR10, for distances up to 100 m over MMF	57-1000284-01
	100G-CFP2-LR4-10KM	100 GbE CFP2 optic, LR4, for distances up to 10 km over SMF	57-1000285-01
	100G-CFP2-ER4-40KM	100 GbE CFP2 optic, ER4, for distances up to 40 km over SMF	57-1000328-01
	100G-QSFP28-SR4	100 GbE SR4 QSFP28 optic for distances up to 100m over MMF. Supported on VDX6940-144S and VDX8770-4/8 platforms.	57-1000326-01
	100G-QSFP28-LR4L-2KM	100 GbE QSFP28 optic for distances up to 2 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms.	57-1000329-01
	100G-QSFP28-LR4-10KM	100 GbE QSFP28 optic for distances up to 10 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms.	57-1000334-01

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

Extreme Branded Optics

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-ADPT-CFP2-QSFP	100G CFP2 to QSFP28 adapter
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
10G-AOC-SFP7M	10G AOC SFP+ 7m
10GB-BX10-D	10 GB, SINGLE FIBER SM, -D 10 KM
10GB-BX10-U	10 GB, SINGLE FIBER SM, -U 10 KM
10G-DACA-SFP1M	10G Active DAC SFP+ 1m
10G-DACA-SFP3M	10G Active DAC SFP+ 3m
10G-DACA-SFP5M	10G Active DAC SFP+ 5m
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-QSFP5M	40G Active DAC QSFP+ 5m
40G-DACP-QSFP1M	40G Passive DAC QSFP+ 1m
40G-DACP-QSFP4SFP1M	40G Passive DAC QSFP+ to 4xSFP+ 1m
40G-DACP-QSFP4SFP2M	40G Passive DAC QSFP+ to 4xSFP+ 2m
40G-DACP-QSFP4SFP3M	40G Passive DAC QSFP+ to 4xSFP+ 3m

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
40G-DACP-QSFP4SFP5M	40G Passive DAC QSFP+ to 4xSFP+ 5m
40G-DACP-QSFP5M	40G Passive DAC QSFP+ 5m

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

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

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

10GBase-T Copper SFP+

The 10GBase-T Copper SFP+ optic is supported on VDX 6740, 6940-144S and 8770. Please note that this optic is not Extreme part, and must be purchased from Method Electronics or their partners. Its Method part number is SP7051-BRCD.

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

Mellanox MAM1Q00A-QSA	Quad to Serial Small Form Factor Pluggable Adapter which can be used with following Extreme P/Ns: 10G-SFPP-SR (10G SR) 10G-SFPP-USR (10G USR) 10G-SFPP-LR (10G LR) 10G-SFPP-ER (10G ER) 10G-SFPP-AOC-0701 (10G AOC 7m) 10G-SFPP-AOC-1001 (10G AOC 10m) 10G-SFPP-TWX-0101 (10G 1m Twinax cable) 10G-SFPP-TWX-0301 (10G 3m Twinax cable) 10G-SFPP-TWX-0501 (10G 5m Twinax cable)
-----------------------	---

CFP2 to QSFP28 conversion module
(PN: 80-1008646-01)

CFP2 to QSFP28 conversion module connects the QSFP28 optic (100G optic) in a CFP2 capable port of 2/6x100G line cards in VDX8770-4/8 chassis.

*Note: Legacy Foundry Networks branded optics are not supported

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

SOFTWARE UPGRADE AND DOWNGRADE

Image filenames

Download the following images from www.extremeportal.force.com

Image Filename	Description	Supported Device or Module
nos7.4.1b.tar.gz	Network OS v7.4.1b for unix	NA
nos7.4.1b.zip	Network OS v7.4.1b for Windows	NA
nos7.4.1b_all_mibs.tar.gz	Network OS v7.4.1b MIBS	NA
nos7.4.1b-releasenotes	Network OS v7.4.1b Release Notes(PDF)	NA
nos7.4.1b.md5	Network OS v7.4.1b 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.

NOS7.4.1-SHA512 upgrade/downgrade basics:

SHA512 NOS741 Upgrade scenario:

Upgrading NOS from NON-SHA512 to SHA512 release:

Cold boot	Default and Non-default user passwords that were created in prior NOS versionw will preserve and it should be in MD5 format. Execute below CLI to convert the existing MD5 format to SHA 512 format. CLI: sw0# password-encryption convert-enc-to-level-10
Default Config	Default user's password will auto convert into SHA512 format. Non-default users will not exist.

SHA512 NOS741 Downgrade scenario:

Downgrading NOS from SHA512 to NOS-SHA512 release:

- Downgrade will be blocked in the below scenarios (warning message in place to notify the user for user action to continue with FWDL cmd):
 - o If the system exists with non-default users with SHA512 format/encryption level 10 password.
 - o If the user changed the Default users(admin & user) password to SHA512 format/encryption level 10 password.
- User should remove the non-default users and reset the default user password to factory default (Use the new CLI sw0# execute-script factory_reset_defaultuser_passwords.sh).
- See the following examples of **Downgrading**.

Downgrade cases:

Case 1: When the customer has configured sha512 passwords for both default and non-default users.

```
sw0#  
sw0# firmware download scp host 10.6.46.51 user fvt password pray4green directory  
/buildsjc/sre_nos/SQA/nos/nos7.4.0/nos7.4.0a  
Performing system sanity check...
```

Downgrade is not allowed due to the below reason(s):

- Non-default users configured with password encryption-level 10. Encryption-level 10 passwords are not supported in the firmware being downloaded.

Please remove the non-default user configuration by using "sw0(config)# no username <user-name>"

- Reset default user's (admin & user) password to factory-default by using "sw0# execute-script factory_reset_defaultuser_passwords.sh".

Pls. try the downgrade after correcting the same.

The preinstall script failed.

```
sw0#
```

Case 2: When the customer has configured sha512 passwords for the default user accounts.

```
sw0#
sw0# firmware download scp host 10.6.46.51 user fvt password pray4green directory
/buildsjc/sre_nos/SQA/nos/nos7.4.0/nos7.4.0a
Performing system sanity check...

Downgrade is not allowed due to the below reason:
- Reset default user's (admin & user) password to factory-default by using "sw0# execute-
script factory_reset_defaultuser_passwords.sh".

Pls. try the downgrade after correcting the same.

The preinstall script failed.
sw0#
```

Case 3: When the customer has configured sha512 passwords for non-default user accounts only.

```
sw0#
sw0# firmware download scp host 10.6.46.51 user fvt password pray4green directory
/buildsjc/sre_nos/SQA/nos/nos7.4.0/nos7.4.0a
Performing system sanity check...

Downgrade is not allowed due to the below reason:
- Non-default users configured with password encryption-level 10. Encryption-level 10
passwords are not supported in the firmware being downloaded.
  Please remove the non-default user configuration by using "sw0(config)# no username
<user-name>"

Pls. try the downgrade after correcting the same.

The preinstall script failed.
sw0#
```

Known issues:

1. SHA512 password encryption for "root" user is not covered with the current image.

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	7.0.0	7.0.1x	7.1.0x	7.2.0x	7.3.0x	7.4.0x	7.4.1	7.4.1x
7.0.0	NA	ISSU*	coldboot	default-config	default-config	default-config	default-config	default-config
7.0.1x	coldboot	ISSU	coldboot	default-config	default-config	default-config	default-config	default-config
7.1.0x	coldboot	coldboot	ISSU	coldboot	default-config	default-config	default-config	default-config
7.2.0x	default-config	default-config	coldboot	ISSU	coldboot	default-config	default-config	default-config
7.3.0x	default-config	default-config	default-config	coldboot	ISSU	coldboot	default-config	default-config
7.4.0x	default-config	default-config	default-config	default-config	coldboot	ISSU	coldboot	coldboot
7.4.1	default-config	default-config	default-config	default-config	default-config	coldboot	NA	ISSU
7.4.1x	default-config	default-config	default-config	default-config	default-config	coldboot	ISSU	NA

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. Before downgrading to lower releases, it is recommended to disable all new features that are not supported on lower releases by using the “no” version of the CLIs. Stray configurations left out before downgrade can cause undesired behavior.
3. While upgrading chassis based system, under stress condition (e.g. due to excessive processing load on the processor), some linecards may become faulty during firmware download. To recover, run “power off <linecard>” followed by “power on <linecard>” command.
4. 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.
5. Firmware download from Network OS7.0.1a to Network OS6.x or Network OS5.x with default-config option needs AG mode disabled.
6. **Limitations:
 - a) Rarely, 40G links may not come up online after upgrade to 7.1.0, you need to do *shut/no shut* to recover
 - b) In VDX 8770 platforms, after upgrading from 6.0.2 to 7.1.0 with coldboot, SNMP V3 traps are not received for the V3host which is under Rbridge.
 - c) Dport test between VDX 6740T and VDX 6940-144S breakout link may fail in upgrade to 7.1.0 and above.
7. Nos7.3.0aa can only be upgraded using coldboot from any earlier versions.
8. Coldboot or ISSU upgrade works fine from NOS7.3.0 to NOS 7.3.0aa.
9. ISSU downgrade from NOS 7.4.0a to NOS 7.4.0 is not supported because of open defects.
10. Future upgrades to 7.4.1 will be through patches. Upgrades will be from 7.4.1 to 7.4.1x through ISSU.

Management IP connectivity

With respect 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.

Upgrading to this Release (Best Practices)

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

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

- Download the firmware on all the switches running Network OS 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 7.2.0 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 v7.3.0 cluster. Eventually, when all the nodes are upgraded, they will form one Network OS 7.3.0 VCS Cluster. [Note that no configuration changes are allowed during this time.]

Downgrading to a Previous Release

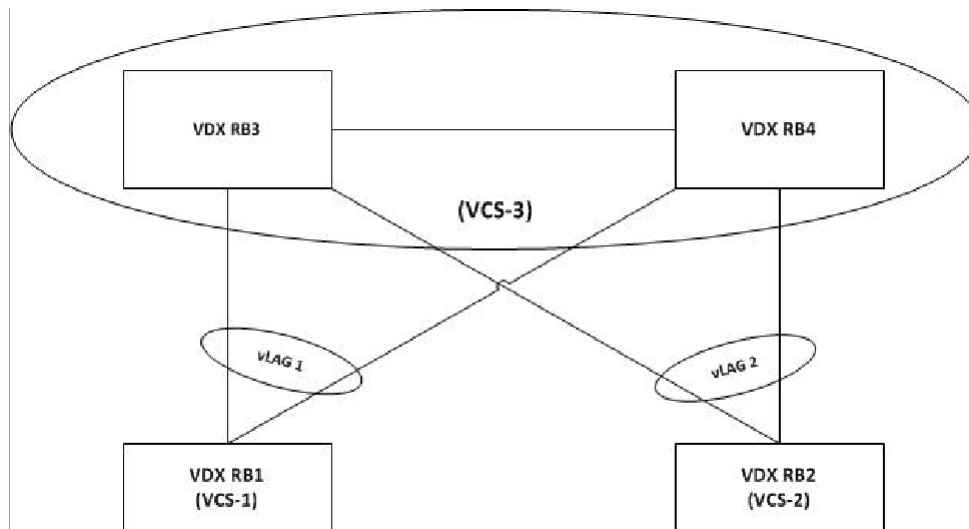
- In normal circumstances, the SW/0 partition is Active. When an ISSU performed, the SW/1 partition becomes active. In order to ensure config is retained during coldboot downgrade, it is important to have SW/0 partition Active before downgrade. The SW/0 partition can be made Active by reloading the switch 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.

Upgrade/downgrade Considerations for vLAG deployments

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

- **Approach 1:** Graceful shutdown of vLAG ports on one node at a time.
- **Approach 2:** Static vLAGs and Dynamic vLAGs without configuration changes.

vLAG deployment upgrade Illustration



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

Step 1: 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.

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 ha failover:

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

SNMP-Trap Source IP support:

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

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

```
loopback number | ve vlan_id}}
```

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

```
loopback number | ve vlan_id}}
```

```
snmp-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 v7.4.0 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

Network OS v7.4.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
Max # of BFD sessions over Virtual-Fabric Extension Tunnels	10	N/A	10	10
Maximum # of dot1q VLANs that can be attached on VxLAN GW for VMware NSX	2000	N/A	2000	2000
Maximum # of VLANs (dot1q VLANs attached to VxLAN GW for NSX + Virtual Fabric VLANs enabled on edge-interfaces)	(2000+1,000)	N/A	(2000+1000)	(2000+1000)
Maximum # of VxLAN tunnels with VMware NSX	250	N/A	250	250
Maximum # of service-nodes with VMware NSX	5	N/A	5	5
Maximum # of MAC Associations for AMPP	8000	4000	8000	8000
Maximum # of per priority pause levels	3	8	3	3
Maximum # of VMware vCenters per Fabric	4	4	4	4
Maximum # of ELD instances in the fabric	2000	2000	2000	2000
Maximum # of IGMPv2v3 Snooping Interfaces supported	4000	4000	4000	4000
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) *	1500/1000	12000/2000	3064/2000	3064/2000
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 interfaces/Logins (Per switch)	1000	1000	1000	1000
Maximum # of Enodes Devices per Fabric	2000	2000	2000	2000

Network OS v7.4.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
Maximum # of NPIV per Port	64	64	64	64
Maximum # of MSTP instance	32	32	32	32
Maximum # of VLAN in PVST	128	128	128	128
Maximum # of LAGs (Port Channels)	64	288	144	144
Maximum # of members in a standard LAG	16	16	16	16
Maximum # of members in a Extreme Trunk (10G)	16	8	12	12
Maximum # of members in a Extreme Trunk (40G)	2	NA	3	3
Maximum # of members in a Extreme Trunk (100G)	NA	NA	NA	NA
Maximum # of switches in 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 v7.4.0 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	16	16	16	16
Maximum # of FVG multiple subnets in Session	32	32	32	32
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	512	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
Max Number of HW entries available for PBR	512	8192	512	512

Network OS v7.4.0 Scalability Numbers	VDX 6740, 6740T, 6740T-1G	VDX 8770	VDX 6940- 36Q	VDX 6940- 144S
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 Interface	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

ROUTE PROFILE						
Network OS v7.x Scalability Numbers	DEFAULT	IPV4-MAX-ROUTE	IPV4-MAX-ARP	IPV4-MIN-V6	IPV6-MAX-ROUTE	IPV6-MAX-ND
Maximum # of IPv4 routes with ECMP supported *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 routes with ECMP supported *	1000	0	0	500	2000	2000
Maximum # of OSPF routes *	6000	10000	10000	8000	2500	2500
Maximum # of OSPFv3 routes *	1000	0	0	500	2000	2000
Maximum # of IPv4 BGP routes in HW *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 BGP routes in HW *	1000	0	0	500	2000	2000
Maximum # of IPv4 routes in SW *	6000	10000	10000	8000	2500	2500
Maximum # of IPv6 routes in SW *	1000	0	0	500	2000	2000
Maximum # of ARP Entries *	43000	49000	73000	49000	6000	6000
Maximum # of IPv6 neighbor cache Entries *	12000	0	0	10000	30000	30000

VDX 8770						
Network OS v7.x Scalability Numbers	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	51000/20000	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	DN Y-AR P-INS P	IPV4 -ACL	IPV4-V6-MCAS T	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 and host adapters for VDX as of Network OS v7.4.0. This is a representative list of devices, Network OS v7.4.0 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

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 v7.4.0. 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.

Licensing

- On VDX platforms that have Flexport FC capable interfaces does not require any Port Upgrade license. The Port Upgrade license only controls Ethernet ports (number of ports or speed supported).
- An Integrated Routing license is NOT required on FOS-based SAN platforms running FOS 7.0.1 or above for FCR interoperability connectivity with VCS fabrics and the VDX6740x. Please refer

to the FOS v7.0.1 Admin Guide documentation on configuring FOS platforms for connectivity to VDX 674x switches and VCS fabrics.

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

VCS

- Loopback connection is not supported in VCS mode. If a loopback connection is done (either using loopback plugs or port to port connections on the same switch), those interfaces become ISL interfaces.
- A node with default configuration will not join a cluster if the intermediate nodes between the node being defaulted and rest of the cluster are also undergoing reload. If the node boots up earlier than the intermediate nodes, it will form its own VCS and not join the parent cluster. In such situations, reload the node that is required to join the cluster.
- Logical Chassis Cluster Mode:
 - When a new switch is added to an existing VCS Fabric and if the new switch takes the role of principal node, the other switches in the fabric will receive the configuration of the distributed features such as Virtual IP and VM-Aware Network Automation from the newly added switch. This will cause the existing distributed configuration to be overwritten by the newly added switch in the principal role. This can be avoided by following the new switch addition procedures in *the Network OS 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

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 need to be rebooted for the dual personality mode change to take effect.
- Configuring 40GbE dual personality interface in 100GbE mode would result in the other two 40GbE interface in the port-group being disabled.

1G Mode

- RMON stats are calculated incorrectly for packet sizes 64-127 bytes.
- 1G ports cannot form ISL links. Only 10G ports can be used to form ISL links.
- Extreme Trunks cannot be formed with 1G. Extreme Trunks are only supported on 10G.
- A LAG cannot be created between 1G and 10G ports.
- 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
- Ve Configuration on both Primary and Secondary Vlans
- AMPP on PVLANS
- In case of MSTP if a primary VLAN is added to the instance automatically secondary VLAN also added to the instance.
- When the operator wants to delete the host association on a host port recommended to use “no switchport” rather than “no switchport private-VLAN host-association”. This is applicable only when the host port is untagged. When the host port is tagged both the commands can be used.
- Primary VLAN ID needs to be lower than the secondary VLAN IDs. If primary VLANID is greater than secondary there is an issue with config replay.
- In Logical Chassis mode source macs may not learn on PVLAN configured ports, after deleting some of the secondary VLANs for which the traffic is not flowing.

UDLD

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

STP/DiST

- VDX does not support tunneling non-standard BPDUs and thus IEEE BPDUs (0180:C200:0000) generated as tagged packets in STP/RSTP/MSTP modes may not be tunneled successfully across VCS Fabric. However, VDX supports tunneling standards’ based BPDUs such as untagged IEEE BPDUs and tagged or untagged PVST BPDUs (0100:0CCC:CCCD). Post 3.0.1, the tagged IEEE BPDU can be tunneled across VCS fabric using command: “tunnel tagged-ieee-bpdu” under interface configuration.
- 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 port 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 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.

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 CLI’s added to assist in MTU configuration across all interfaces for a switch using single CLI. This allows quick setting of 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 profile supports upto 32 VLANs.
- ARP/ND suppression feature is supported only on VDX 6740, 6940, 6940-144s platforms.

Conversational ARP:

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

Static Anycast Gateway:

- ARP/ND suppression should be configured on the VLAN if IPv4/IPv6 Static Anycast Gateway is being configured.
- Static Anycast Gateway address/static Anycast MAC configuration should be identical for a given VLAN across leaf nodes in IP Fabric.
- IP services/protocols cannot be enabled on an interface where only Static Anycast Gateway address is configured.
- VRRP/VRRP-E configuration should be identical for a given VLAN across leaf nodes in IP Fabric. But it is recommended to use Static Anycast Gateway.
- All VLANs having Static Anycast Gateway configuration should be extended into eVPN on a vLAG pair.
- 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.
- Multipath BFD for Unnumbered ECMP
 - Each ECMP link part of Multipath BFD will take up one BFD session in addition there will be one primary session created . The overall BFD scale is consumed accordingly.
- Defect ID: NOS-67539 [SQA] Static route BFD for FE80::5 and neighbor-address FE80::2 is down for ve 11 link-local address Defect ID: NOS-67539 [SQA] Static route BFD for FE80::5 and neighbor-address FE80::2 is down for ve 11 link-local address
 - Symptom: For a 2 node Logical-chassis Cluster , BFD Session will be down for IPv6 Link-Local address on VE Interface.
 - Condition: BFD should be enabled on a VE interface with IPV6 Link-Local address on both the nodes, with ipv6 static route BFD configured.
 - Workaround: No workaround
 - Recovery: Cannot recover in 7.4.1 image. Need to load previous release for ipv6 link-local address

VRRP

- VRRP and VRRP-E cannot be enabled together on VDX 6740 and VDX 6740T platforms.
- IPv6 and IPv4 VRRP sessions cannot be configured with the same VRRP group-ID on the same Layer3 interface.
- If an IPv6 VRRP session is configured with only global VIP address without Link-Local VIP, VIP configuration will fail for that session during download of configuration from file.
- VRRP v4 or v6 can be enabled with VRRP-E v4 and v6 on the VDX 6940 family.
- VRRP v4 and v6 cannot be enabled together on an interface on the VDX 6940 family.
- “show vrrp summary” and “show ipv6 vrrp summary” will display all sessions in default vrf. In earlier 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, StaticAnycast Gateway to be used to achieve similar functionality.

OSPFv2

- Appendix-e processing for NSSA is not supported on ABR for type7 to type5 translated routes.

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.
- 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

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

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.
- Vlans with the same internal vlan mapping can be used in and both the vlans are treated as different vlans. Traffic from one doesn't flood to the other
- 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.
- 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 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, and Layer 3 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.

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.

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 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 number of routes due

- to reserved routes/entries which are created for internal use.
- Use count field is added to show the number of entries currently in use.

Management VRF

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

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

Conversational MAC Learning

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

System level Flowbased QoS

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

Port level Flowbased QoS

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

URPF

- uRPF is not supported in VDX8770.

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 (> 1sec)
- 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 below command 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 EXM table respectively.

MAPS

- MAPS is supported on VDX 6740, 6940 and 8770 platforms.
- RX_SYM_ERR MAPS messages are displayed when breakout cable is connected on a 40G interface that is not configured for breakout.
- When 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.

LACP and individual ports for PXE boot enhancement

- PXE client uses one of its interfaces like eth0 , eth1 for initial DHCP discovery communication with the PXE server. This interface MAC will be learned in our VDX PO and the same interface/MAC needs to be used at the PXE client side for completing the PXE boot sequence
- During Pre-boot stage, when user configures LACP default-up the IF state in the running config remains in “no shut” state even though they are brought down by the PXE mechanism. If HA is triggered this state is changing to “shut” state. User needs to check the interface state in show running once the PXE boot is completed and move the Interface state to “no shut”

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 3seconds 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.
- 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.
- The fix for DEFECT659781 reduces the number of writes to the CF in order to reduce the occurrence of CF corruption and CPU usage history information will not be available on VDX 6740, and VDX 6740-T.

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 extremenetworks.com . Note that TSBs are generated for all Extreme platforms and products, so not all TSBs apply to Network OS-based platforms).

Closed with code changes for Network OS v7.4.1b

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS v7.4.1b.

Parent Defect ID:	NOS-68001	Issue ID:	NOS-68001
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.4.1a
Technology Group:	Other	Technology:	Other
Symptom:	IPC-5024 fail message and Unexpected reload.		
Condition:	MAC learnt on VxLAN tunnel. (OR) MAC consistency check with default interval of 300sec.		
Workaround:	Increase the mac consistency-check interval and aging interval time.		

Closed with code changes for Network OS v7.4.1a

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS v7.4.1a.

Parent Defect ID:	NOS-53113	Issue ID:	NOS-67868
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.0.1c
Technology Group:	Network Automation and Orchestration	Technology:	YANG
Symptom:	NOS fails to un-escape special characters in passwords received via Netconf XML for config backup upload.		
Condition:	Special characters used in password		
Workaround:	None		

Parent Defect ID:	NOS-67121	Issue ID:	NOS-67869
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.		
Workaround:	None		

Parent Defect ID:	NOS-67795	Issue ID:	NOS-67870
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-67866	Issue ID:	NOS-67873
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-67828	Issue ID:	NOS-67885
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.4.1
Technology Group:	Other	Technology:	Other
Symptom:	Edge port with optic 10306 is remain Admin Down		
Condition:	When peer server reload and brought up online		
Recovery:	Reseat optics OR reload the VDX		

Parent Defect ID:	NOS-67830	Issue ID:	NOS-67886
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.1.0b3

Technology Group:	Other	Technology:	Other
Symptom:	SFM fault and sometime unexpected reload.		
Condition:	Transient H/W error detected in SFM's SCI bus. 2020/02/04-20:59:46:269977, [EM-5010], 1000964/0, M1 Active, WARNING, VDX8770-8, S3: FRU type or I2C problems (-7), FRU will be faulted as UNKNOWN, OID:0x431c0000, object.c, line: 3979, comp:emd, ltime:2020/02/04-20:59:46:269762		
Recovery:	Re-seat/Replace faulted SFM		

Parent Defect ID:	NOS-67805	Issue ID:	NOS-67892
Severity:	S3 - Medium		
Product:	Network OS	Reported in Release:	NOS7.4.1
Technology Group:	Management	Technology:	Software Installation & Upgrade
Symptom:	Unable to form the VCS & 40G ISL port fault		
Condition:	When we have only 40G ISL with DAC cable & Upgrade to 7.4.1 firmware .		
Workaround:	Replace DAC cable by using Non DAC QSFP optics		
Recovery:	Replace DAC cable by using Non DAC QSFP optics		

Closed with code changes for Network OS v7.4.1

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS v7.4.1.

Parent Defect ID:	NOS-67447	Issue ID:	NOS-67458
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.0.2c	Technology:	ACLs - Access Control Lists
Symptom:	"show access-list interface port-channel XX" will not show the proper output about the active/Inactive status. No functional impact, it is a show command issue.		
Condition:	It is show command issue , when ACL applied under VLAG.		

Parent Defect ID:	NOS-67454	Issue ID:	NOS-67464
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.1.0b	Technology:	AMPP - Automatic Migration of Port Profiles
Symptom:	Unexpected reload		
Condition:	While activating auto port-profile on interface		

Parent Defect ID:	NOS-67387	Issue ID:	NOS-67471
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-53094	Issue ID:	NOS-67473
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:	when we use "sw0(config)# logging raslog message ARP-1038 suppress"		
Workaround:	Avoid using "sw0(config)# logging raslog message ARP-1038 suppress"		

Parent Defect ID:	NOS-67208	Issue ID:	NOS-67474
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:	When 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-66997	Issue ID:	NOS-67491
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	ARP - Address Resolution Protocol
Symptom:	Repeated IPAD-1001-log entries even though there is no change to mgmt interface status		
Condition:	mgmt-vrf default route is resolved via inband interface		
Workaround:	do not configure default route with nh pointing to inband interface		

Parent Defect ID:	NOS-67109	Issue ID:	NOS-67495
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-67503	Issue ID:	NOS-67506
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.3.0a	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP trap fails		

Condition:	With SNMP version v2
-------------------	----------------------

Parent Defect ID:	NOS-66881	Issue ID:	NOS-67508
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:	Instead of executing parallel query we can execute the query sequentially		

Parent Defect ID:	NOS-66898	Issue ID:	NOS-67509
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	BGP4+ - IPv6 Border Gateway Protocol
Symptom:	Unexpected reload.		
Condition:	When we run the "show ip bgp neighbors" after giving the "no neighbor activate" followed by "clear ip bgp neighbour"		
Workaround:	None		

Parent Defect ID:	NOS-66998	Issue ID:	NOS-67510
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-67475	Issue ID:	NOS-67511
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.2.0c_CVR	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		

Parent Defect ID:	NOS-67014	Issue ID:	NOS-67513
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:	1.Ping between the 2 VCS fails. 2.BGP convergence is not triggered.		

Condition:	1.Ping between the 2 VCS fails upon reload of one of the node in VCS Topology. 2.BGP convergence is not triggered when port between the VCS is shut.
-------------------	---

Parent Defect ID:	NOS-54700	Issue ID:	NOS-67517
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:	Unable to get the IP address through DHCP		
Condition:	For Linux command "dhclient eth0" after upgrading from nos6.x to nos7.x		
Workaround:	keep the broadcast flag set with dhclient command: "dhclient -B eth0"		

Parent Defect ID:	NOS-67405	Issue ID:	NOS-67529
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0a	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-67192	Issue ID:	NOS-67534
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.4.0	Technology:	Hardware Monitoring
Symptom:	QSFP breakout interface does not link up and port hard fault		
Condition:	ADA serial number based QSFP breakout port does not link up while another end of interface connects to Non VDX device.		

Parent Defect ID:	NOS-67133	Issue ID:	NOS-67542
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.3.0aa	Technology:	CLI - Command Line Interface
Symptom:	"show version" output displayed error during the copy support collect.		
Condition:	There is no functional impact. "show version" o/p at supportsave logs is not proper.		
Workaround:	Verify the "show version" at switch instead of supportsave logs.		

Parent Defect ID:	NOS-67666	Issue ID:	NOS-67666
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.2.0c	Technology:	SSH - Secure Shell
Symptom:	Unable to login via SSH using the user accounts.		
Condition:	After downgrade from NOS7.4.x to NOS7.2.x.		

Parent Defect ID:	NOS-67674	Issue ID:	NOS-67674
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.0.2b	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-67689
Severity:	S2 - High		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.3.0a	Technology:	IGMP - Internet Group Management Protocol
Symptom:	Unexpected reload		
Condition:	Periodic IGMP join and leave messages		

Parent Defect ID:	NOS-67692	Issue ID:	NOS-67692
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.2.0c	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		

Parent Defect ID:	NOS-67461	Issue ID:	NOS-67698
Severity:	S2 - High		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.3.0a	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.		
Workaround:	Use less than 1000 IGMP join/leave request per second		

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

Parent Defect ID:	NOS-67543	Issue ID:	NOS-67723
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management

Reported in Release:	NOS7.2.0c	Technology:	NTP - Network Time Protocol
Symptom:	Unexpected reload		
Condition:	When ntpdate (with NTP config) process spawn's multiple times.		

Parent Defect ID:	NOS-67724	Issue ID:	NOS-67724
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.1.0b3	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-67742	Issue ID:	NOS-67742
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.4.1	Technology:	User Accounts & Passwords
Symptom:	Admin user can get the root privileges		
Condition:	when user try to use python from admin login		

Closed with code changes for Network OS v7.4.0a

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS v7.4.0a.

Parent Defect ID:	NOS-67162	Issue ID:	NOS-67165
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.3.0a	Technology:	User Accounts & Passwords
Symptom:	"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-67203	Issue ID:	NOS-67215
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.2.0b	Technology:	Hardware 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.		

Parent Defect ID:	NOS-67110	Issue ID:	NOS-67247
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-67248
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 cause inconsistency between CLI output and SNMP/XMC view.		
Workaround:	None		

Parent Defect ID:	NOS-66866	Issue ID:	NOS-67250
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-67017	Issue ID:	NOS-67251
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching

Reported in Release:	NOS6.0.2ha	Technology:	LAG - Link Aggregation Group
Symptom:	Momentary traffic black-hole is observed		
Condition:	When we shut the peer LAG port or reload the peer device		
Workaround:	None		

Parent Defect ID:	NOS-66324	Issue ID:	NOS-67252
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 and can happen only when the ASIC resources are full.		
Workaround:	None		

Parent Defect ID:	NOS-67019	Issue ID:	NOS-67254
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.2b	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP walk failure		
Condition:	Reload of the switch having IP ACL for SNMP community config will result in SNMP walk failure		

Parent Defect ID:	NOS-55113	Issue ID:	NOS-67258
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-67061	Issue ID:	NOS-67259
Severity:	S2 - High		

Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.2a	Technology:	Configuration Fundamentals
Symptom:	Linecard reset by the detection of Multibit Parity error condition show in RAS Log. “[HWK2-5651], 14604207/0, L4/0 Active, ERROR, VDX8770-8, CHIPO: Interrupt: DMC dmc mbit err.”		
Condition:	Linecard reset.		
Workaround:	None		

Parent Defect ID:	NOS-67225	Issue ID:	NOS-67266
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-67221	Issue ID:	NOS-67274
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 a particular VDX.		

Parent Defect ID:	NOS-67189	Issue ID:	NOS-67281
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.		

Parent Defect ID:	NOS-53106	Issue ID:	NOS-67288
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)		

Workaround:	None
--------------------	------

Parent Defect ID:	NOS-67290	Issue ID:	NOS-67299
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-67223	Issue ID:	NOS-67301
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 or sometimes unexpected reload of the switch		
Condition:	Doing a Port-Channel config on ISL port which is in shut state may cause fabric disconnection or unexpected reload of the switch.		
Workaround:	Always do Port-channel config on an interface in link UP state		
Parent Defect ID:	NOS-67182	Issue ID:	NOS-67303
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.		

Closed with code changes for Network OS v7.4.0

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change in Network OS v7.4.0.

Parent Defect ID:	NOS-38127	Issue ID:	NOS-38127
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Traffic Management
Reported in Release:	NOS5.0.2c2	Technology:	Rate Limiting and Shaping
Symptom:	<p>VDX normally does not reject service-policy configuration when the interface has storm-control settings:</p> <pre>sw0# show run int ten 1/0/1 interface TenGigabitEthernet 1/0/1 switchport switchport mode trunk switchport trunk allowed vlan add 10 switchport trunk tag native-vlan spanning-tree shutdown fabric isl enable fabric trunk enable storm-control ingress broadcast limit-percent 1 storm-control ingress multicast limit-percent 8 storm-control ingress unknown-unicast limit-percent 1 no shutdown !</pre> <p>sw0# conf t Entering configuration mode terminal sw0(config)# int ten 1/0/1 sw0(conf-if-te-1/0/1)# service-policy in vlan10PLC %%Error: Policer can't be enabled as storm control is enabled. sw0(conf-if-te-1/0/1)#</p> <p>However, a customer mentions that, they can tweak configuring both on the same interface through a port-channel. They would like to know the expected traffic behavior under this condition. An example below:</p> <pre>sw0# show run int po 10</pre>		

	<pre> interface Port-channel 10 no vlag ignore-split speed 1000 service-policy in vlan10PLC switchport switchport mode trunk switchport trunk allowed vlan add 10 switchport trunk tag native-vlan spanning-tree shutdown no shutdown ! sw0# conf t Entering configuration mode terminal sw0(config)# int ten 1/0/1 sw0(config-if-te-1/0/1)# no switchport 2018/08/27-01:05:39, [NSM-1010], 1203, SW/0 Active DCE, INFO, sw0, InterfaceMode changed from L2 to None for interface TenGigabitEthernet 1/0/1. sw0(config-if-te-1/0/1)# channel-group 10 mode on type standard 2018/08/27-01:05:46, [NSM-1004], 1204, SW/0 Active DCE, INFO, sw0, Interface Port-channel 10 is created. 2018/08/27-01:05:46, [NSM-1017], 1205, SW/0 Active DCE, INFO, sw0, Interface TenGigabitEthernet 1/0/1 is added on interface Port- channel 10. 2018/08/27-01:05:46, [NSM-1002], 1206, SW/0 Active DCE, INFO, sw0, Interface TenGigabitEthernet 1/0/1 is protocol down. 2018/08/27-01:05:46, [NSM-1003], 1207, SW/0 Active </pre>
Condition:	<p>VDX normally does not reject service-policy configuration when the interface has storm-control settings.</p> <p>Configuring both together conflicts the expected behavior.</p>
Workaround:	Remove storm-control config from each interface before applying service policy in the portchannel

Parent Defect ID:	NOS-47588	Issue ID:	NOS-47588
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS6.0.2a	Technology:	LAG - Link Aggregation Group
Symptom:	Traffic disruption in the cluster due to unresponsive rbridge		
Condition:	In rare conditions, the ISLs stay up on an unresponsive rbridge.		

Parent Defect ID:	NOS-47745	Issue ID:	NOS-47745
Severity:	S3 - Medium		

Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS6.0.2b	Technology:	Logical Chassis
Symptom:	Management cluster/VCS goes offline when ISL between two nodes goes down even though the connectivity could have been established through the other nodes' ISL.		
Condition:	It happens rarely when the new link/connectivity is established slowly.		

Parent Defect ID:	NOS-47754	Issue ID:	NOS-47754
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS6.0.2b	Technology:	Logical Chassis
Symptom:	System may undergo unexpected reload		
Condition:	Media removal while media data is being read		
Workaround:	shut/ no shut media removed interface		

Parent Defect ID:	NOS-47790	Issue ID:	NOS-47790
Severity:	S1 - Critical		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS6.0.2b	Technology:	Management GUI
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" .		

Parent Defect ID:	NOS-47956	Issue ID:	NOS-47956
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS6.0.2e	Technology:	Hardware Monitoring
Symptom:	Unexpected Line Card reload while collecting SS from BNA		
Condition:	Copy Support save CLI execution can lead to this issue.		

Parent Defect ID:	NOS-47966	Issue ID:	NOS-47966
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS6.0.2e	Technology:	Hardware 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]'		

Parent Defect ID:	NOS-47970	Issue ID:	NOS-47970
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS6.0.2e	Technology:	ACLs - Access Control Lists
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		

Parent Defect ID:	NOS-48005	Issue ID:	NOS-48005
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS6.0.2g	Technology:	Hardware Monitoring
Symptom:	No ASIC parity error messages in RASlog.		
Condition:	Switch did not go to faulty state even though there were parity errors.		

Parent Defect ID:	NOS-52059	Issue ID:	NOS-52059
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.0	Technology:	VRRPv2 - Virtual Router Redundancy Protocol Version 2
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.		

Parent Defect ID:	NOS-52732	Issue ID:	NOS-52732
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.1	Technology:	Logical Chassis
Symptom:	Cannot configure IPv6 (/126) address on VIP for VRRP-E.		
Condition:	Configuring IPv6 address (/126) for VRRP-E		

Parent Defect ID:	NOS-53076	Issue ID:	NOS-53076
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.0.1c	Technology:	Hardware 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.		

Parent Defect ID:	NOS-53087	Issue ID:	NOS-53087
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.0.1c	Technology:	VXLAN - Virtual Extensible LAN
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.		

Parent Defect ID:	NOS-53091	Issue ID:	NOS-53091
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.0.1c	Technology:	Syslog
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.		

Parent Defect ID:	NOS-53092	Issue ID:	NOS-53092
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.1c	Technology:	OSPF - IPv4 Open Shortest Path First
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.		

Parent Defect ID:	NOS-53101	Issue ID:	NOS-53101
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.1c	Technology:	Logical Chassis
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).		

Parent Defect ID:	NOS-53104	Issue ID:	NOS-53104
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.1c	Technology:	Logical Chassis
Symptom:	unexpected core files fills up disk.		
Condition:	"show logging raslog rbridge-id" CLI execution for multiple rbridge at the same time.		

Parent Defect ID:	NOS-53110	Issue ID:	NOS-53110
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.1c	Technology:	Static Routing (IPv4)
Symptom:	BGP route not cleared when VE interface is shut.		
Condition:	VE interface shut		

Parent Defect ID:	NOS-53138	Issue ID:	NOS-53138
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:	Expired TLS certificate should not be imported to device.		

Parent Defect ID:	NOS-53143	Issue ID:	NOS-53143
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.0.2	Technology:	Software Installation & Upgrade
Symptom:	Lost configuration during upgrade tsd terminated with core dump		
Condition:	Upgrade from 6.0.2e to 7.0.2		
Workaround:	None		

Parent Defect ID:	NOS-53165	Issue ID:	NOS-53165
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.0.2a	Technology:	Hardware Monitoring
Symptom:	Started supporting Extreme optics on VDX devices.		
Condition:	Extreme optics qualification on VDX devices.		

Parent Defect ID:	NOS-54586	Issue ID:	NOS-54586
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.1.0	Technology:	Configuration Fundamentals
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 stays down.		

Parent Defect ID:	NOS-54598	Issue ID:	NOS-54598
Severity:	S4 - Low		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.1.0a	Technology:	SSH - Secure Shell
Symptom:	'ssh-server' CLI is unable to configure options such as cipher, mac, kex ...etc		
Condition:	FIPS mode is enabled		

Parent Defect ID:	NOS-54637	Issue ID:	NOS-54637
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	VPN
Reported in Release:	NOS7.1.0a	Technology:	EVPN - Ethernet VPN
Symptom:	GARP Doesn't flood to hosts to updated their ARP cache irrespective of whether ARP suppression is enabled/disabled.		
Condition:	Ipfabric environment where L2VPN is enabled.		

Parent Defect ID:	NOS-54643	Issue ID:	NOS-54643
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.1.0a	Technology:	Logical Chassis
Symptom:	Port does not come online on VDX 6740-T platform		
Condition:	Port didn't come online when the peer server is CentOS was rebooted multiple times.		

Parent Defect ID:	NOS-54648	Issue ID:	NOS-54648
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	BFD - BiDirectional Forwarding Detection
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 bfdd off sw0:FID128:root> chroot /mnt chkconfig bfdd off		

Parent Defect ID:	NOS-54710	Issue ID:	NOS-54710
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.1.0b1	Technology:	TACACS & TACACS+
Symptom:	Tacacs accounting functionality does not work properly.		
Condition:	In VCS cluster node rejoin operation can cause this issue.		

Parent Defect ID:	NOS-55119	Issue ID:	NOS-55119
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0	Technology:	CLI - Command Line Interface
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.		

Parent Defect ID:	NOS-55127	Issue ID:	NOS-55127
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP polling for cpStatus OID returns incorrect values.		
Condition:	When SNMP get/walk request done for cpStatus OID.		

Parent Defect ID:	NOS-55768	Issue ID:	NOS-55768
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.3.0	Technology:	Logical Chassis
Symptom:	Consistent DB corruption on sudden power outage		
Condition:	Sudden Power Cycle the device can cause the issue.		
Workaround:	<p>We can use below workaround for planned outage or power cycle.</p> <p>NOS7.1.0 and Later Releases:</p> <ol style="list-style-type: none"> Execute chassis power-cycle-db-shutdown command through NOS CLI. 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. <p>Any release prior to NOS 7.1.0: Root level command. root> shutdowncmdb 2018/09/20-20:58:29 : shutdowncmdb : Shutting Down Database (New)</p>		

Parent Defect ID:	NOS-55895	Issue ID:	NOS-55895
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0a	Technology:	Static Routing (IPv4)
Symptom:	BGP route not cleared in secondary node of VCS cluster		
Condition:	Remove ve interface		

Parent Defect ID:	NOS-56056	Issue ID:	NOS-56056
Severity:	S1 - Critical		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS6.0.2c	Technology:	VLAN - Virtual LAN
Symptom:	Traffic impact or packet loss between directly connected hosts.		
Condition:	Traffic impact or packet loss between directly connected hosts.		

Parent Defect ID:	NOS-56057	Issue ID:	NOS-56057
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS6.0.2c	Technology:	CLI - Command Line Interface
Symptom:	"show system" CLI execution does not 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.		

Parent Defect ID:	NOS-56100	Issue ID:	NOS-56100
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.1.0	Technology:	Hardware 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.		
Workaround:	None		

Parent Defect ID:	NOS-56142	Issue ID:	NOS-56142
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.3.0	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SW-MIB capability will be set to 'NO' in the snmp config.		
Condition:	When we upgrade from lower versions to the version greater than NOS7.2.0, we will observe "SW-MIB: NO" in the snmp config.		
Workaround:	None		

Parent Defect ID:	NOS-56156	Issue ID:	NOS-56156
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	ICMP - Internet Control Message Protocol
Symptom:	VDX does not generate ICMP notification for IP MTU violation [trapped packets].		
Condition:	IP MTU violation		

Parent Defect ID:	NOS-66321	Issue ID:	NOS-66321
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:	None		

Parent Defect ID:	NOS-66864	Issue ID:	NOS-66864
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		

Defect ID: DEFECT000590108	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS4.1.3	Technology: IP Addressing
Symptom: ACL may not work as expected	
Condition: ACL rule configured with /8 (255.0.0.0) mask	
Workaround: Need to apply the specific ACL	

Defect ID: DEFECT000590465	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.1	Technology: Configuration Fundamentals
Symptom: channel-group configurations for port-channel member interfaces are lost upon reload.	
Condition: VDX replays configuration through file [startup-config] when configuration has been defaulted and it causes channel-group configuration lost.	

Defect ID: DEFECT000590478	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: Network OS5.0.2	Technology: IPv4 Multicast Routing
Symptom: mcasgt process termination	
Condition: The issue is seen when multicast routes are added and deleted from the system, which leaves some amount of memory leak, which grows over time and causes a system crash.	
Workaround: Yes	

Defect ID: DEFECT000590517	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.0	Technology: VXLAN - Virtual Extensible LAN
Symptom: VLAN flooding on a tunnel will not work.	
Condition: VLAN flooding on the tunnel will not work.	
Workaround: Run the following command - "tunnel replicator bum-VLANs redistribute" in exec mode of Network OSCLI.	
Recovery: VLAN's will be distributed to the available SN tunnels.	

Defect ID: DEFECT000590808	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: Hidden commands under debug and foscmd hide group were not shown as part of show running config even after un hiding and configuring them. Even the copy running to file was not having the configuration after copy command was executed after un hiding.	
Condition: Config commands under hide group "debug" and "foscmd" have to be executed after un hiding respective hide group. Post this, executing "show running config" will not show these un hidden configurations.	

Defect ID: DEFECT000591179	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS4.1.3	Technology: VLAN - Virtual LAN
Symptom: VDX incorrectly sends packets as un tagged over trunk port.	
Condition: Network OS4.1.3b can hit the issue.	

Defect ID: DEFECT000591223	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: This is an enhancement that introduces a new CLI under rbridge-id sub-mode to configure the behavior of some IF-MIB attributes: ifName and ifDescr. If this knob is configured to 3-tuple, then the above 2 objects will be of 3-tuple format. Else, they will be of 2-tuple format. These 2 attributes will also be in the same format during Link Up/Down Trap generation.	
Condition: This is applicable only for ifName and ifDescr attributes of IF MIB and the linkUp/Down traps.	

Defect ID: DEFECT000591225	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS6.0.1	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: SNMP IP ACL config mismatch between the Frontend & Backend database.	
Condition: Reload with default config will retain the IP ACL data for SNMP community string.	

Defect ID: DEFECT000591256	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.0	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: ECMP paths through some VE interfaces might not be calculated in a multi node scenario consisting of VLAGs after Port-channel flap.	
Condition: This issue can happen when a self originated max-age network LSA is received after Port-channel flap and there is a delay in reforming OSPFv2 adjacency among VLAG end points causing the network LSA to contain only one of the neighbor info for some time.	
Workaround: Issue "no shutdown" of port-channel interface after neighbor nodes have flushed the max-age Network LSA.	
Recovery: Issue "shutdown/no-shutdown" on VE interfaces which are missing in the ECMP nexthop list. Also can be recovered by issuing "clear ip ospf all".	

Defect ID: DEFECT000591616	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: Switch goes for an unexpected reload with the REST request.	
Condition: When the switch is pounded with the REST requests from multiple concurrent sessions simultaneously and continuously over a long period of time.	
Workaround: As far as possible, send REST requests to configure the switch from one session only. Multiple sessions can be used for retrieving information from the switch with GET requests.	

Defect ID: DEFECT000591700	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: Network OS6.0.1	Technology: QoS - Quality of Service
Symptom: BUM traffic has higher latency compare to data traffic.	
Condition: BUM traffic use store and forward method and data traffic use cut through method.	

Defect ID: DEFECT000592128	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.1	Technology: IP Addressing

Symptom: Software Fault: A rare memory corruption issue in the tty driver caused Kernel Panic and rebooting of the switch.
Condition: The issue was introduced in the 2.6.34 kernel and the same was addressed by an open source fix in the tty driver.

Defect ID: DEFECT000592256	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: Network OS4.1.3	Technology: VCS Fabric
Symptom: Downlink ports take long time to come online with the latest FW (10.6) of Hitachi 520X blade LOM.	
Condition: The issue was introduced after the FW upgrade of Hitachi 520X blade LOM.	
Recovery: Upgrade to the new Network OS version.	

Reported In Release:	Network OS4.1.3	Technology: VCS Fabric
Symptom: Downlink ports take long time to come online with the latest FW (10.6) of Hitachi 520X blade LOM.		
Condition: The issue was introduced after the FW upgrade of Hitachi 520X blade LOM.		
Recovery: Upgrade to the new Network OS version.		

Defect ID: DEFECT000592398		
Technical Severity: Medium	Probability: Medium	
Product: Extreme Network OS	Technology Group: VCS	
Reported In Release:	Network OS7.0.0	Technology: Logical Chassis
Symptom: During multi-cast tree formation, a RBridge with a configured root priority level may not take effect for the tree's formation. Instead, the configured RBridge behaves as though it has a default or lowest priority configuration. However, when displaying the running configuration, it shows the expected tree root priority configuration.		
Condition: Following an operation where a RBridge boots up with a default configuration, and then downloads it's configuration from the active cluster, a non-default setting for the RBridge's multi-cast root priority may not take affect. This may happen such as after a 'vcs replace' operation.		
Recovery: Rebooting the affected node forces it to refresh the effective priority value for the multi-cast tree root priority. Alternatively, explicitly changing the priority to a different value and then setting it back to the original desired value causes the priority to be updated. However, setting the root priority to a different value may affect the multi-cast tree formation depending on the temporary priority specified.		

Defect ID: DEFECT000592617		
Technical Severity: Medium	Probability: Medium	
Product: Extreme Network OS	Technology Group: Layer 2 Switching	
Reported In Release:	Network OS5.0.2	Technology: xSTP - Spanning Tree Protocols
Symptom: IEEE BPDU Local VLAN tunnel CLI allowed to be configured when protocol spanning tree is already configured or vice versa.		
Condition: When both STP protocol and IEEE BPDU Local VLAN tunnel CLI are enabled at the same time.		

Defect ID: DEFECT000592647		
Technical Severity: Medium	Probability: Low	
Product: Extreme Network OS	Technology Group: Management	
Reported In Release:	Network OS5.0.1	Technology: NTP - Network Time Protocol
Symptom: Timezone set might fail		
Condition: Particular timezone related files got corrupted. It is very rare scenario to hit.		
Recovery: Delete the failed timezone file under /usr/share/zoneinfo/ .		

Configure the timezone .

Defect ID: DEFECT000592669	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.0	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: "[no]bfd-shutdown" and "bfd-interval<>" commands are not available under port-channel interfaces.	
Condition: Unable to configure BFD commands on L3 port-channel	

Defect ID: DEFECT000592874	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS6.0.1	Technology: Hardware Monitoring
Symptom: In very rare scenario they can observe interface flap	
Condition: Due to excessive symbol errors	

Defect ID: DEFECT000593245	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: Multi-VRF
Symptom: Ping Round-Trip-Times fluctuate between 4 and 16 ms.	
Condition: Happens in 6.0.2a and later releases.	

Defect ID: DEFECT000593285	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS4.1.3	Technology: Static Routing (IPv4)
Symptom: Move Ethernet cable from one VDX to another which causes the PC to loose connectivity.	
Condition: Moving management eth cables between VDX can cause the issue.	
Recovery: Clear mac-address-table will recover the condition.	

Defect ID: DEFECT000593611	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS7.0.1	Technology: MAPS - Monitoring and Alerting Policy Suite

Symptom: While deleting policy with REST API, actually it is success but an error was thrown. It has been fixed and check-in.
Condition: While deleting policy with REST API, actually it is success but an error was thrown. It has been fixed and check-in.

Defect ID: DEFECT000593960	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: With 3-tuple format configured for ifDescr and ifName, the linkUp/Down traps generated still contain ifDescr var-bind in 2-tuple format.	
Condition: This is related to ifDescr var-bind in the linkUp/Down trap only.	

Defect ID: DEFECT000594223	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: Software Installation & Upgrade
Symptom: TFTP server/service was enabled by default.	
Condition: Any device from outside can try to connect VDX using TFTP and VDX burn its resources unnecessary.	

Defect ID: DEFECT000594682	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS6.0.2	Technology: ACLs - Access Control Lists
Symptom: SNMP walk failure in some scenarios.	
Condition: Creating IP ACL with sequence id as 0 causes this issue.	
Workaround: Avoid using sequence id 0 while creating IP ACL.	

Defect ID: DEFECT000594815	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.1	Technology: VLAN - Virtual LAN
Symptom: The execution of command "show vlan brief" will cause the box to reboot.	
Condition: This issue may be seen when all the following conditions are met. <ul style="list-style-type: none"> 1. There are more than 40 nodes in a Logical Chassis. 2. VFAB is enabled on the cluster. 3. There are 10 VLAN's configured. 4. There are more than 1000 ports configured on each VLAN. 5. show-vlan-brief was executed. 	

Workaround: Instead of "show vlan brief", the user can execute "show interface trunk" to check the vlan-port configurations.

Defect ID: DEFECT000594819	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.1	Technology: VXLAN - Virtual Extensible LAN
Symptom: Switch can experience an unexpected reload with HSL kernel backtrace.	
Condition: When VXLAN tunnels are deleted and then added again.	

Defect ID: DEFECT000594867	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.2	Technology: VXLAN - Virtual Extensible LAN
Symptom: Ingress policers not limiting traffic on VDX6740.	
Condition: Ingress policers do not work correctly when traffic needs to be encapsulated for example heading into a VxLAN tunnel.	

Defect ID: DEFECT000595049	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.1	Technology: Access Gateway
Symptom: If Access Gateway (AG) configuration commands are executed through REST API interface, even though command is successfully executed, HTTP error will be reported.	
Condition: This is cosmetic error issue due to different return code used in access gateway.	
Recovery: This is not actual error. Command will be executed successfully. Check the running config to confirm the command.	

Defect ID: DEFECT000595071	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS4.1.3	Technology: VLAN - Virtual LAN
Symptom: 'show interface trunk' won't display proper VLAN information in output.	
Condition: If 'switchport trunk native-vlan' configured above 2047 and there is no VLAN configured below 2047.	
Workaround: Configure at least one VLAN below 2047 & associate with any physical interface.	

Defect ID: DEFECT000595233	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VPN

Reported In Release:	Network OS7.0.1	Technology: EVPN - Ethernet VPN
Symptom: In very rare scenarios after ISSU upgrade traffic drops may be observed for Tunnel terminated traffic		
Condition: ISSU upgrade is a necessary condition for this issue. But not all ISSU upgrades will results in this issue		
Workaround: Perform disruptive firmware upgrades that involve reboots		
Recovery: Rebooting the switch will recover the system.		

Defect ID: DEFECT000595395		
Technical Severity: High		Probability: High
Product: Extreme Network OS		Technology Group: Layer 3 Routing/Network Layer
Reported In Release:	Network OS6.0.2	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: IP DHCP Relay is not working properly when enabled on VRRP-E master interface		
Condition: Operating IP DHCP Relay together with VRRP-E		
Workaround: toggle the VE interface		

Defect ID: DEFECT000595653		
Technical Severity: Medium		Probability: High
Product: Extreme Network OS		Technology Group: Layer 3 Routing/Network Layer
Reported In Release:	Network OS5.0.2	Technology: IP Addressing
Symptom: IP Directed broadcast would not work after HA failover, but the CLI configuration may present.		
Condition: HA fail-over trigger the issue.		
Recovery: Reconfigure IP directed-broadcast		

Defect ID: DEFECT000595709		
Technical Severity: High		Probability: High
Product: Extreme Network OS		Technology Group: VCS
Reported In Release:	Network OS7.0.1	Technology: Logical Chassis
Symptom: System reloads on VDX8770.		
Condition: This occurs with 512 or more VRRP sessions enabled and "debug vrrp packets" is turned on.		
Workaround: "debug vrrp packets" should not be turned on in a scaled environment.		

Defect ID: DEFECT000595754		
Technical Severity: High		Probability: Medium
Product: Extreme Network OS		Technology Group: Layer 3 Routing/Network Layer

Reported In Release:	Network OS6.0.2	Technology: IPv6 Addressing
Symptom: Disabling autoconfig (autonomous address-configuration flag) for an IPv6 prefix in Network OS 6.0.2 has no impact on router-advertisement.		
Condition: Disabling autoconfig		

Defect ID: DEFECT000595877		
Technical Severity: High		Probability: High
Product: Extreme Network OS		Technology Group: Layer 3 Routing/Network Layer
Reported In Release:	Network OS7.0.1	Technology: Multi-VRF
Symptom: Unexpected reload of switch observed when removing a VRF configuration or removing ipv4/ipv6 address family configuration of a VRF.		
Condition: When a custom VRF is unconfigured or IPv4/IPv6 address family of a VRF is unconfigured, switch will be reloaded.		

Defect ID: DEFECT000595980		
Technical Severity: High		Probability: High
Product: Extreme Network OS		Technology Group: VCS
Reported In Release:	Network OS6.0.2	Technology: Logical Chassis
Symptom: When tunnel tagged-ieee-bpdu is enabled on any of the interface, protocol spanning-tree is allowed to be configured.		
Condition: Tunnel tagged-ieee-bpdu configured before configuring protocol spanning tree.		

Defect ID: DEFECT000596257		
Technical Severity: High		Probability: High
Product: Extreme Network OS		Technology Group: Management
Reported In Release:	Network OS6.0.2	Technology: Software Installation & Upgrade
Symptom: After reload, though the uplink interface is down, the downlink tracking interface is still up.		
Condition: All the downlinks interface are brought up , irrespective of the uplink interface state after reboot.		

Defect ID: DEFECT000596280		
Technical Severity: Medium		Probability: Medium
Product: Extreme Network OS		Technology Group: Layer 3 Routing/Network Layer
Reported In Release:	Network OS5.0.2	Technology: IP Addressing
Symptom: Unable to delete an ACL.		
Condition: When ACL is associated to the management interface of one or more switches in the VCS and the switch gets removed from VCS.		

Defect ID: DEFECT000596480	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: Configuration Fundamentals
Symptom: On execution of CLI "track remove all" complete Link State Tracking (LST) configuration should get removed from a port. In case of port-channel interface protocol daemon is not clearing the LST configuration hence it is displayed in output of show command.	
Condition: Execution of "track remove all" CLI for a port-channel interface for which Link State Tracking (LST) configuration is present.	
Workaround: As a workaround user can remove the configuration one by one by executing respective 'no' CLIs.	

Defect ID: DEFECT000596496	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS6.0.2	Technology: Logical Chassis
Symptom: Protocol spanning-tree configuration will not be allowed even after removing the "spanning-tree ieee-bpdu limit-vlan-flood" and "tunnel tagged-ieee-bpdu" configuration.	
Condition: When all the switches in the VCS are configured with "spanning-tree ieee-bpdu limit-vlan-flood" and one or more switches are removed from VCS.	
Recovery: Copy running configuration to remote. Reload the switch with default configuration and copy back the running configuration.	

Defect ID: DEFECT000596708	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: Unicast frame counter maybe displayed incorrectly while there are Multicast and Broadcast traffic concurrently.	
Condition: This is due to HW implementation of the statistics counters. Counters are displayed correctly once traffic is idle.	

Defect ID: DEFECT000596720	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.2	Technology: CLI - Command Line Interface
Symptom: When IPv6 nd prefix is configured with a prefix flag(no-autoconfig/no-onlink/offlink) enabled and if the same prefix is updated later with different lifetime values, then the already configured prefix flag will not be present in the running configuration of that prefix.	

Condition: This issue happens when an IPv6 prefix configuration is updated with lifetime values provided a prefix flag(no-autoconfig/no-onlink/offlink) was already configured.
Workaround: NA

Defect ID: DEFECT000596781	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: IPv6 Addressing
Symptom: Lifetime configuration value of VE interface IPv6 nd prefix is reset to infinite.	
Condition: Doing "shutdown" and "no shutdown" configuration on the VE interface	

Defect ID: DEFECT000596868	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: The global MTU value cannot be deleted through REST API.	
Condition: Issue happens when the user tries to delete the global mtu using the DELETE request through the REST interface.	
Workaround: Using the PATCH request with the default value as a work around. The effect of this is same as deleting the config.	

Defect ID: DEFECT000596932	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS5.0.1	Technology: LAG - Link Aggregation Group
Symptom: Interfaces may not join into Dynamic LAG.	
Condition: Static lag creation before dynamic LAG.	
Workaround: Configuring dynamic LAG first and then static	
Recovery: Delete the static LAGs and re-add the same.	

Defect ID: DEFECT000597053	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: Network OS5.0.1	Technology: VCS Fabric
Symptom: In rare scenario, VDX can send packets with TTL=0. Which can cause the connectivity issues.	
Condition: VxLAN packets terminated on VDX6940 & BUM forwarder on other ISL partner.	
Recovery: Configure static MAC address for the specific IP address.	

Defect ID: DEFECT000597104

Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: Under rare scenarios of leaking routes between VRF's, the switch may get reloaded due to "termination of process ribmgr"	
Condition: When leaking routes from one VRF to another & presence of those same routes in target VRF as connected routes.	
Workaround: Reconfigure to avoid leaking routes between VRF's OR ensure that the leaked routes are not present in target VRF as local routes.	

Defect ID: DEFECT000597782	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.2	Technology: VLAN - Virtual LAN
Symptom: The management MAC and one of the VE MACs may conflict.	
Condition: This is a software defect that has affected the VDX6940-36Q and VDX6940-144S since their release.	

Defect ID: DEFECT000597954	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: MAC routes are dampened at lesser number of moves than the max-moves threshold configured	
Condition: 1. 2-Node VCS leaf in the topology with a misconfiguration of having different AS numbers on the rbridges belonging to the same VCS. 2. Mac move happening between a 2 node VCS leaf and any other leaf	

Defect ID: DEFECT000598328	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: when switch is warm recovered (failover state), there may be traffic impact on some tunnels. tunnel traffic may not get terminated, and there will be traffic loss.	
Condition: warm recovery may cause it under heavy load conditions. it doesn't happen always, but likelihood of happening is more under heavy loaded setups.	
Workaround: cold reboot is needed to recover,	

Recovery: powercycle the switch

Defect ID: DEFECT000598345	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: Network OS5.0.0	Technology: Rate Limiting and Shaping
Symptom: slow learning of hosts ARP entries in 6740 platform	
Condition: In rare scenarios when there is a sudden burst of routed traffic.	

Defect ID: DEFECT000598508	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS5.0.2	Technology: VXLAN - Virtual Extensible LAN
Symptom: VxLAN tunnel unexpectedly went down and did not recover resulting in a ping loss between end hosts even though tunnel was up on other participating RBs in the fabric.	
Condition: One of the VxLAN tunnel endpoint RBridge is rebooted.	

Defect ID: DEFECT000598524	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: Network OS6.0.2	Technology: Rate Limiting and Shaping
Symptom: rte_cap_acl debug tool won't work for 6940 platforms	
Condition: Enable the rte_cap_acl tool support for 6940 platforms	

Defect ID: DEFECT000598641	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.1	Technology: VXLAN - Virtual Extensible LAN
Symptom: Customer might experience unexpected reload of the system.	
Condition: This is seen on updating certain set of configuration,	

Defect ID: DEFECT000598657	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS5.0.1	Technology: SSH - Secure Shell
Symptom: Unexpected reload.	
Condition: Rare scenario where remote host IP becomes NULL.	

Defect ID: DEFECT000598663

Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS5.0.2	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: DCMd daemon terminated and sudden reload occurred.	
Condition: If customer has big cluster and actively executing CLI commands through script or monitoring tools [BNA] then Principal node receives too many message to handle and it hit this issue.	
Workaround: Please reduce any command execution frequency.	

Defect ID: DEFECT000598878	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: Configuration Fundamentals
Symptom: A stale default-route gets applied in the running configuration of the secondary nodes in cluster environment during configuration replay.	
Condition: The issue arises when secondary nodes disconnect and re-join the cluster provided DHCP is enabled.	

Defect ID: DEFECT000598972	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: Switch might go for an unexpected reload when any configuration update is performed on a range of interfaces.	
Condition: On a large cluster with scaled up configurations, performing any configuration on a range of interfaces by entering into interface range sub-mode might cause switch to run out of memory and thereby causing it to reload.	
Workaround: Required configuration update can be made on individual interfaces one at a time instead of performing it on a range of interfaces. Configuration update on multiple interfaces can still be performed by using comma (,) as separators instead of hyphen (-) when specifying the range. For ex, to shutdown interfaces 1 to 5, use "interface te 1,2,3,4,5" instead of "interface te 1-5".	

Defect ID: DEFECT000599289	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS7.0.1	Technology: ACLs - Access Control Lists
Symptom: Applying Access Control List (ACL's) with 12K rules on management interface takes more than 3 minutes to enforce it.	

Condition: When Access Control List (ACL's) is configured with 12K rules.

Defect ID: DEFECT000599306	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: Vrf information is missing for some interfaces while displaying output of "show ip interface brief" command.	
Condition: This issue is seen, then "show ip interface brief" is executed repeatedly in multiple terminals.	
Workaround: If "show ip interface brief" executed from multiple terminals, then it not should be executed too quickly. Let the command output display completed on one terminal before starting on other terminal.	

Defect ID: DEFECT000599778	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS7.0.1	Technology: TACACS & TACACS+
Symptom: LDAP/RADIUS/TACACS+ server configurations are not displayed in the same order in which they were added.	
Condition: 1. Configure multiple TACACS+/RADIUS/LDAP servers(max 5) 2. Remove few server entries 3. Add those servers entries back	
Workaround: Remove all Server entries and configure those servers back in the desired order.	

Defect ID: DEFECT000599835	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS7.1.0	Technology: AAA - Authentication, Authorization, and Accounting
Symptom: Switch with ACL policy that has 12K rules and is enforced to the management interface causes switch to reload	
Condition: Re-sequence the ACL policy which has 12K rules	

Defect ID: DEFECT000599897	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS6.0.2	Technology: Port Mirroring
Symptom: Control frame loss as the SPANed ARP frames trapping on intermediate node	
Condition: If we have huge ARP traffic coming on to a SPANed port which has SPAN destination in other RBridge	

Workaround: Make the SPAN session local to that RBridge and remove the SPAN in VCS session
Recovery: Make the SPAN session local to that RBridge and remove the SPAN in VCS session

Defect ID: DEFECT000600002	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Network OS7.0.1 Release:	Technology: Metro VCS
Symptom: When an optic is removed/ and inserted back too quickly, there is an VERIFY message on console that indicate due to media data reading failure.	
Condition: This VERIFY is not needed, since there is a retry to read media data. The media data will be successful after retry in this case.	

Defect ID: DEFECT000600022	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Network OS7.0.1 Release:	Technology: Metro VCS
Symptom: When VDX 8770 is in chassis-disabled state, the far-end 100 GbE link partners using QSFP28 optics may see intermittent link flaps. After VDX 8770 is chassis-enabled, there is a low probability that the 100 GbE port may not come online.	
Condition: When VDX 8770 is in chassis-disabled state, the far-end 100 GbE link partners using QSFP28 optics may see intermittent link flaps. After VDX 8770 is chassis-enabled, there is a low probability that the 100 GbE port may not come online.	
Recovery: Execute "shut" on 100 GbE link partner port connected to VDX 8770 to stop the port from flapping intermittently. After the VDX 8770 is chassis-enabled, execute "no shut" on the 100 GbE link partner to re-enable the port.	

Defect ID: DEFECT000600023	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Network OS6.0.2 Release:	Technology: SSH - Secure Shell
Symptom: Shutting SSH server on Standby partition fails	
Condition: After High Availability fail over, we may hit the issue.	

Defect ID: DEFECT000600057	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Network OS7.0.1 Release:	Technology: Logical Chassis
Symptom: Switch might not rejoin the cluster when reloaded using 'fastboot' command.	
Condition: Reloading switch using 'fastboot' command on VDX6940 and VDX6740 platforms when SW1 partition is active might lead to this issue.	

Workaround: Reload the switch using 'reload' command which is more graceful way of reloading.
Recovery: Bring the switch which failed to join the cluster to default configuration using command 'copy default-config startup-config'. On reload, switch rejoins the cluster and regains older configuration.

Defect ID: DEFECT000600066	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP IPv4 Relay forwarded DISCOVER packet is not getting forwarded through remote leaf node in BGP-EVPN IP Fabric.	
Condition: While deploying DHCP Relay in BGP-EVPN IP Fabric.	
Recovery: Disable "conversational-arp".	

Defect ID: DEFECT000600169	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: IP MTU configuration is not working for VE interface when IP address or L3 VNI association is not present.	
Condition: When IP MTU is configured, it is not applied on the VE interface.	
Workaround: Configure IP MTU followed by the configuration of the IP address.	

Defect ID: DEFECT000600185	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: Network OS7.0.1	Technology: OpenStack Integration
Symptom: When VDX-8770 is in chassis-disable state, the "show media" command will not show 100 GbE ports.	
Condition: When VDX-8770 is in chassis-disable state, the "show media" command will not show 100 GbE ports.	
Workaround: After the chassis is enabled using "chassis enable" command, "show media" will show the 100 GbE ports.	
Recovery: After the chassis is enabled using "chassis enable" command, "show media" will show the 100 GbE ports.	

Defect ID: DEFECT000600377	
Technical Severity: High	Probability: Low

Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: The SNMP walk may fail and SNMPV3 trap may not be received for the user configured under RBridge.	
Condition: The SNMP walk may fail and SNMPV3 trap may not be received only for the SNMPV3 user configured under RBridge after upgrade from 7.0.0 to 7.0.1.	
Recovery: Reconfigure the user under RBridge after the successful upgrade from 7.0.0 to 7.0.1.	

Defect ID: DEFECT000600579	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.1	Technology: VXLAN - Virtual Extensible LAN
Symptom: Unexpected reload of switch on performing ISSU or ha-failover.	
Condition: With VxLAN tunnel and BFD configured, ISSU from any version prior to 7.0.1 may result in unexpected reload of switch.	
Workaround: Issue is fixed 7.0.1 and hence in 7.1.0 also. No workaround needed.	
Recovery: reload of switch	

Defect ID: DEFECT000600591	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: Logs are dumped on the screen, when there is a read failure on SFPs connected to the port.	
Condition: Accessing information about the SFPs inserted in the ports.	
Recovery: Disable the port and re-enable it.	

Defect ID: DEFECT000600696	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS6.0.2	Technology: Hardware Monitoring
Symptom: Unable to run the RTE tool for CBR2 platform	
Condition: While executing the RTE tool on CBR2 platforms.	

Defect ID: DEFECT000601145	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: Network OS7.0.1	Technology: IP Fabric
Symptom: If issuing fastboot on a chassis system some modules may crash on the MM that is about to reboot. This should have no impact on functionality.	

Condition: A change was made to the reboot procedures that was not propagated to fastboot on chassis systems.

Defect ID: DEFECT000601146	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS7.0.1	Technology: Logical Chassis
Symptom: When user does an upgrade we may see system go down and come up with the same old version.	
Condition: This could happen when the ISSU notification to the standby keeps failing.	

Defect ID: DEFECT000601715	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.0	Technology: CLI - Command Line Interface
Symptom: When copying files to/from VDX switch to TFTP server we are seeing errors when 'use-vrf' option is specified.	
Condition: Copying files to/from VDX switch to TFTP server.	

Defect ID: DEFECT000601917	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: Multi-VRF
Symptom: Change of MAC address of a host connected to VCS is not updated across user defined VRF's in ARP table.	
Condition: Incorrect MAC address will be replied for an ARP request.	

Defect ID: DEFECT000601985	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.0	Technology: xSTP - Spanning Tree Protocols
Symptom: VDX switches running in a VCS cluster may encounter CIST Spanning-tree interoperability problem with certain Juniper switches where BPDU's sourced by the VDX may be dropped by the partner.	
Condition: When VDX running in VCS cluster running distributed CIST spanning-tree & VDX switches are configured as spanning-tree root.	
Workaround: Change the spanning-tree root to partner switch.	

Defect ID: DEFECT000602062	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS5.0.2	Technology: Access Gateway
Symptom: Console logs appear when snmpwalk is performed.	
Condition: When snmpwalk is performed for community/user associated with IPv6 ACL.	

Defect ID: DEFECT000602227	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.0	Technology: SNMP - Simple Network Management Protocol
Symptom: SNMP OID 1.3.6.1.2.1.17.1.3 displays 'No such instance' in output	
Condition: snmpwalk for SNMP OID 1.3.6.1.2.1.17.1.3	

Defect ID: DEFECT000602239	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS5.0.2	Technology: ACLs - Access Control Lists
Symptom: VDX experience unexpected reload after configuring permit statement on standard ACL applied to management interface.	
Condition: Configuration of permit statement on standard ACL applied to management interface.	
Workaround: NA	

Defect ID: DEFECT000602579	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.1	Technology: High Availability
Symptom: ISSU may fail and result in standby GOS not booting up.	
Condition: This can happen in some corner case where bootenv cannot be accessed during ISSU.	
Workaround: reboot the system	
Recovery: reboot the system	

Defect ID: DEFECT000602722	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: Network OS7.1.0	Technology: OpenStack Integration

Symptom: VDX 8770 6x100 GbE port may show RX_SYM_ERR after link is administratively flapped due to excessive mac-move detection on the port. In this case, RX_SYM_ERR will not affect traffic after link is up.
Condition: VDX 8770 6x100 GbE port may show RX_SYM_ERR after link is administratively flapped due to excessive mac-move detection on the port. In this case, RX_SYM_ERR will not affect traffic after link is up.
Workaround: RX_SYM_ERR dashboard statistics can be cleared via Network OSCLI.
Recovery: RX_SYM_ERR dashboard statistics can be cleared via Network OSCLI.

Defect ID: DEFECT000602751	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.0	Technology: xSTP - Spanning Tree Protocols
Symptom: User tries firmware downgrade and will hit error message as, User need to clean the config and then only downgrade can be done.	
Condition: When "system-oui" configuration is done under "protocol spanning-tree" configuration mode and subsequently, a downgrade is done.	
Workaround: User needs to remove the config with "no system-oui" command under "protocol spanning-tree" mode.	

Defect ID: DEFECT000602764	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.0	Technology: xSTP - Spanning Tree Protocols
Symptom: After spanning tree system OUI feature enabled and then disabled, the firmware download is failed.	
Condition: Doing spanning tree system OUI enable and disable. Then performing the firmware download.	

Defect ID: DEFECT000603443	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.1	Technology: LAG - Link Aggregation Group
Symptom: Changing LACP timeout option in VDX can cause LACP PDUs to be sent at short intervals when neighboring device is cisco Nexus 7k. Changing LACP timeout option from long to short and again to long in both the devices can cause this behavior.	
Condition: LACP timeout option in VDX internally remains as short though configuration is shown as long.	

Defect ID: DEFECT000603778	
Technical Severity: Medium	Probability: Medium

Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: IPv6 Addressing
Symptom: When both "IPv6 vrrp-suppress-interface-ra" and "IPv6 VRRP VIP" are configured the IPv6 RA response to the IPv6 RS contains link-local address instead of the VIP address.	
Condition: Configure both "IPv6 vrrp-suppress-interface-ra" and "IPv6 VRRP VIP"	

Defect ID: DEFECT000604049	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: Network OS7.0.1	Technology: VCS Fabric
Symptom: Potential for Name Server fail-over and recovery to the Standby Control Processor if the overall scale of the VCS cluster exceeds the limit described within the "Conditions for Publication" section.	
<p>Condition: The maximum number of elements within a cluster cannot exceed 32,767 prior to having this modification to increase scale to 80,000. Entities that contribute to this count are:</p> <ul style="list-style-type: none"> - RBridges - Ports (physical and virtual) - Devices that appear in the Name Server <p>The maximum assignable port indexes are listed here by platform type:</p> <p>Chassis-based systems (Director class) : 1800 VDX 6740/VDX 6740T/VDX 6740T-1G : 1200 VDX 6940 : 1312</p> <p>For example, one Director-class RBridge accounts for 1 (for the RBridge itself) + 1800 (maximum assignable port indexes) + <FC/FCoE device count>. Thus, if we have 500 devices, this would translate to 1 + 1800 + 500 = 2301 (of the total allowable 32767). Here are some sample combinations in terms of RBridge composition within a cluster, where a cluster-wide FC/FCoE device count is presumed to be 3000:</p> <ul style="list-style-type: none"> • 16 Directors • 14 Directors + 3 VDX 6940/ 3 VDX 6740 • 12 Directors + 5 VDX 6940 / 6 VDX 6740 • 8 Directors + 11 VDX 6940 	
Workaround: Limit cluster composition in a manner compliant with the maximum values described within the "Conditions for Publication" section.	
Recovery: Limit cluster composition in a manner compliant with the maximum values described within the "Conditions for Publication" section.	

Defect ID: DEFECT000604054	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.1.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Loopback interfaces are showing bogus IP MTU value, when global MTU is configured.	
Condition: Execution of "show ip interface lo <ID>" when global MTU is configured.	

Defect ID: DEFECT000604131	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.0	Technology: Multi-VRF
Symptom: If local route exists from a route source and a leaked route is added from the same route source for the prefix, the routing table is updated with the new leaked route.	
Condition: Issue is seen if dynamic route leak is configured with prefixes matching the local prefixes.	
Workaround: There should not be overlap between local and leaked prefixes	

Defect ID: DEFECT000604714	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.1.0	Technology: SNMP - Simple Network Management Protocol
Symptom: In some rare case, SNMP v1/v2c query with specific community string may not respond.	
Condition: Configure more than one community and do reload or node rejoin.	
Recovery: Reconfigure community string	

Defect ID: DEFECT000604743	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: BGP static networks are not advertised to peers.	
Condition: static-network route is configured.	

Defect ID: DEFECT000605042	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: CLI - Command Line Interface
Symptom: 'snmp-server' command doesn't update the values of 3 input parameters.	

Condition: Inputting all the 3 parameters contact, location and sys-descr on a single line of execution.
Workaround: Configure each of the input parameter separately.

Defect ID: DEFECT000605230	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: IPv6 Addressing
Symptom: After ISSU upgrade the configuration "ipv6 nd prefix 2011::/64 2592000 86400 no-autoconfig" no longer works.	
Condition: Internal configuration data didn't sync properly.	

Defect ID: DEFECT000605476	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: The route advertised by eBGP peer is not installed in the routing table.	
Condition: This issue occurs only in the self-referencing scenario i.e. when route prefix overlaps with the prefix of next hop from where the route is received.	
Workaround: Isolate the bgp peering in a different subnet so that their prefix does not overlap with the routes being advertised between them	

Defect ID: DEFECT000605776	
Technical Severity: Low	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.2	Technology: CLI - Command Line Interface
Symptom: New script will help to clear all the counters with single command	
Condition: It is an enhancement	
Workaround: Use the individual commands to clear the counters	

Defect ID: DEFECT000605899	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: Network OS4.1.3	Technology: Logical Chassis
Symptom: Radius Client connections via fabric to Radius Server failing.	
Condition: This is observed when VDX6740 receives the IP packets with DSCP 63 (0x3F) from the Radius Clients	

Defect ID: DEFECT000605923

Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.1	Technology: FCoE - Fibre Channel over Ethernet
Symptom: FCoE VLAN creation and subsequent fabric map may fail.	
Condition: When more than 64 ports are configured with 'switchport trunk allowed vlan all' configuration and tried to create FCoE VLAN.	
Workaround: Do not configure more than 64 ports with 'switchport trunk allowed vlan all' configuration while creating an FCoE VLAN.	
Recovery: Remove 'switchport trunk allowed vlan all' configuration if it is configured on more than 64 interfaces and try creating FCoE VLAN and fabric map.	

Defect ID: DEFECT000605998	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS5.0.2	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: just once "clear ip ospf all" is done in RB01, a tremendous "LSA flush rcvd Type:5" message lasted to pop up forever and network got unstable.	
Condition: In huge scale OSPF setups when there are more than 10 neighbors and OSPF peer has to retransmit an LSA to all these neighbors this issue is seen as each neighbor is added to retransmit queue multiple times	
Workaround: Decrease the number of LSAs and neighbors	

Defect ID: DEFECT000606064	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.1.0	Technology: Software Installation & Upgrade
Symptom: "/bin/cat: /etc/time.conf: No such file or directory:" is seen during firmware download	
Condition: Firmware download	

Defect ID: DEFECT000608321	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.1.0	Technology: Software Installation & Upgrade
Symptom: Firmware upgrade is initiated on a node which has default-config configured	
Condition: Warning message is given when default config is configured and a non-default option is provided during firmware download	

Defect ID: DEFECT000608446	
Technical Severity: High	Probability: High

Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS7.0.1	Technology: Logical Chassis
Symptom: VDX generates FFDC core file and throws Software 'verify' error on console.	
Condition: Execution of copy default-config startup-config from VCS primary node.	
Workaround: NA	
Recovery: LC gets automatically recovered.	

Defect ID: DEFECT000608811	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS7.1.0	Technology: Logical Chassis
Symptom: An unexpected reload of the switch can occur.	
Condition: If a Network OSCLI show command is left paginated and not exited out or completed within a week's time frame, then an unexpected reload of the switch can occur.	
Workaround: Use "terminal length 0" to turn off show command pagination.	

Defect ID: DEFECT000608838	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: SNMP - Simple Network Management Protocol
Symptom: SFP Interface goes into administratively down state. Ex: [NSM-1028], 5673/2457, SW/0 Active DCE, ERROR, <hostname>, Incompatible SFP transceiver for interface TenGigabitEthernet 1/0/45 is detected	
Condition: Execution of "[no] snmp trap link-status" command on an un-tunable SFP interface.	
Workaround: Please do not disable "snmp trap link-status" which is enabled by default on all interface.	
Recovery: Enable tunable-optics configuration and then disable it on impacted interface as below: <pre>tunable-optics sfpp channel 1 no tunable-optics sfpp channel</pre> <p>Make interface up again:</p> <pre>no shutdown</pre>	

Defect ID: DEFECT000608995	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: ARP - Address Resolution Protocol

Symptom: Traffic to/from DHCP host is not routed when the DHCP IP is assigned to a new host. The ARP for such host does not age out when age out timer expires.
Condition: DHCP Server is sending ACK packets to relay agent even when the client address is known. Mostly seen with Windows DHCP server.
Workaround: Modify DHCP server settings so that it will send reply directly to dhcp client when client IP is present in the received DHCP message.

Defect ID: DEFECT000610081	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Release: Network OS7.1.0	Technology: IGMP - Internet Group Management Protocol
Symptom: Traffic leaks on one of the ports of vLAG (port-channel) when a Layer 2 Static IGMP group configuration on the specific vLAG is removed. The port is still showing in the Layer 3 PIM Mcache Outgoing interface list. The problem shows up in one of the remote nodes of VCS, which happens to be the DR.	
Condition: During the cleanup of IGMP static group configuration removal for vLAG interface on the VLAN the information is not getting conveyed to the PIM protocol.	
Workaround: Avoid IGMP static group configuration on a vLAG	
Recovery: Disable PIM and enabling it again on the VE.	

Defect ID: DEFECT000610145	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.2	Technology: Management GUI
Symptom: Controllers which rely on multipart-reply(flow-stats) to validate/mark the flow as installed/added may get confused and may try to delete the flow again and again	
Condition: Affects the flow-mods where reserved ports are part of action set	

Defect ID: DEFECT000610510	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS5.0.1	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: OSPF routes are uninstalled from one or more VRF's, causing traffic disruption. Router LSA's do not refresh.	
Condition: Occurs when there are many OSPF session across many VRF's, with total OSPF routes exceeding 1500.	
Recovery: Flap OSPF neighbor sessions.	

Defect ID: DEFECT000610816	
Technical Severity: High	Probability: High

Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS5.0.2	Technology: LAG - Link Aggregation Group
Symptom: VDX throws FVCS-1005 RASLOG message followed by an unexpected reboot.	
Condition: The user may experience this issue when attempting to change or undo the active Port Channel in a Redundancy Group using the 'no port-channel <PortChannel ID> active' command.	
Workaround: When changing the active Port Channel in a Redundancy Group, it is best to avoid using the 'no port-channel <PortChannel ID> active' command. It is advisable to delete the Redundancy Group and recreate it when wanting to change the Active Port Channel in a Redundancy Group.	

Defect ID: DEFECT000610937	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS5.0.1	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: Gateway for default route obtained through DHCP remains in running configuration under mgmt-vrf even after deleting DHCP config and reloading the switch.	
Condition: Invalid gateway for default route may appear after reloading the switch.	

Defect ID: DEFECT000611059	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS7.1.0	Technology: Logical Chassis
Symptom: VDX experience unexpected reload due to DCMd daemon termination.	
Condition: When Principal fail-over occurs, secondary nodes DB transaction cleanup fails on standby partition due to timing condition.	

Defect ID: DEFECT000611400	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Traffic Management
Reported In Release: Network OS7.1.0	Technology: Rate Limiting and Shaping
Symptom: Switch can go for a reboot when the slot values are provided well outside the permissible range in the 'bp-rate-limit command'. Permissible range for slot is '0-16'	
Condition: The issue is seen only when the command is executed by providing the slot values well outside the permissible range.	
Workaround: Ensure that slot values are provided only in the valid range '0-16'	
Recovery: Remove any of the slot values provided outside the permissible range of '0-16'	

Defect ID: DEFECT000611576

Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS5.0.2	Technology: Logical Chassis
Symptom: Getting "% Error: VLAN string length(1139) is more than maximum length 1023" on reboot.	
Condition: VDX with allowed vlan configuration string length more than 1023 can hit the issue at boot up & configuration replay time.	

Defect ID: DEFECT000611680	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS4.1.3	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: Few OSPF LSA's are not removed from Database when they are withdrawn at the source.	
Condition: When the number of routes are highly scaled and large number of routes are redistributed from BGP and OSPF. VRRP -E has VIP configured same as physical interface IP's.	
Workaround: When these anomalies are removed, OSPF LSA's will be flushed properly.	

Defect ID: DEFECT000611688	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS7.1.0	Technology: Hardware Monitoring
Symptom: VDX 6940 and 6940-144S may show CRC errors on ports cabled with QSFP 40 GbE active copper transceiver cables.	
Condition: VDX 6940 and 6940-144S may show CRC errors on ports cabled with QSFP 40 GbE active copper transceiver cables.	
Workaround: Replace QSFP 40 GbE active copper transceiver cables with optical QSFP transceivers. Then reboot the switch.	
Recovery: Replace QSFP 40 GbE active copper transceiver cables with optical QSFP transceivers. Then reboot the switch.	

Defect ID: DEFECT000612673	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.1.0	Technology: VLAN - Virtual LAN
Symptom: May see spurious "Too many interrupts" events.	
Condition: Other interrupts come in within a second and first one not cleared.	

Defect ID: DEFECT000612821	
Technical Severity: Medium	Probability: Low

Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: VRRPv2 - Virtual Router Redundancy Protocol Version 2
Symptom: VRRP-1002 raslog message is not displayed.	
Condition: When Master to backup change happens .	

Defect ID: DEFECT000612967	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS7.1.0	Technology: Security Vulnerability
Symptom: Shutting down SSH server does not close all existing SSH login sessions	
Condition: Shutdown SSH server	
Recovery: Close all existing login sessions using "clear sessions" command, please note this command will close telnet sessions as well.	

Defect ID: DEFECT000613777	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS6.0.2	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP request packets are dropped on VDX, and are not relayed to DHCP server(s).	
Condition: This affects only DHCP request packets with option-82. For example, an intermediate layer 2 node may have inserted option 82 in the packet and then forwarded to the VDX.	
Workaround: A workaround script is available to disable option-82 check on VDX	
Recovery: A workaround script can be used to recover from this issue	

Defect ID: DEFECT000614353	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS7.0.1	Technology: Hardware Monitoring
Symptom: After inserting a media 'SFP transceiver for interface XYZ is inserted' RASLOG is missing.	
Condition: a media/SFP insertion	

Defect ID: DEFECT000614390	
Technical Severity: Critical	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS5.0.1	Technology: ICMP - Internet Control Message Protocol

Symptom: Very rarely we could see 5% of the ICMP replies are dropped in software and random interval.
Condition: The issue can be happened when we have ARP requests from 1000 different hosts at the rate of 25 ARP's/sec, and at the same time pinging VE or VRRP IP on the same SVI at 1 ICMP/sec

Defect ID: DEFECT000614988	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: IPv6 Addressing
Symptom: "ipv6 nd prefix" CLI command displays incorrect default value for lifetime and preferred lifetime parameter.	
Condition: Execution of "ipv6 nd prefix" CLI.	

Defect ID: DEFECT000615075	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS6.0.2	Technology: Licensing
Symptom: LED on unlicensed and shutdown VDX 40G ports are slow blinking amber after boot. Expected behavior is off since it is unlicensed.	
Condition: After reload, the single QSFP amber LED should only blink slow amber when all the 4 internal links/ports are offline and the port has a 40G Port Upgrade license reservation; otherwise it should be turned off (ie, no color/black).	

Defect ID: DEFECT000615165	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: IPv6 Addressing
Symptom: "ipv6 nd prefix <IPv6> no-autoconfig" config can get lost.	
Condition: Config-replay from backup configuration file when "ipv6 nd prefix <IPv6> no-autoconfig" is configured with valid and preferred life time default values.	

Defect ID: DEFECT000615176	
Technical Severity: Medium	Probability: Low
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS5.0.2	Technology: CLI - Command Line Interface
Symptom: CLI command "show support" will not show few core files from system daemon crash on usual place which is /core_files	

Condition: When there is crash by one of management daemon then core file doesn't get saved on regular system path /core_files

Defect ID: DEFECT000615242	
Technical Severity: Medium	Probability: High
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS6.0.2	Technology: AMPP - Automatic Migration of Port Profiles
Symptom: MACs on Linux Virtual Machines with VMWare Tools installed may not get programmed on VDX.	
Condition: When VMWare Tools are installed on Virtual Machines, Both IPV4 and IPV6 address gets reported from Vmware to VDX. VDX is unable to handle very long IP Strings and ignores such vnics(MACs)	
Workaround: Either disable IPV6 on the Virtual Machines or don't install VMware tools on the Virtual Machines	
Recovery: Disable IPV6 on Virtual Machines or remove VMware tools and re-run the discovery cycle	

Defect ID: DEFECT000615380	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS5.0.2	Technology: DHCP - Dynamic Host Configuration Protocol
Symptom: DHCP packets will be dropped in the box where DHCP Relay is configured.	
Condition: DHCP Relay listens on standard well-known BOOTPS and BOOTPC ports (i.e. 67 and 68). If any other ports are used for communication between DHCP Client and DHCP Server can cause the issue.	
Workaround: As a workaround, use standard BOOTPS and BOOTPC (i.e. 67 and 68) UDP ports for communication between DHCP Relay and DHCP Server.	
Recovery: Use of standard BOOTPS and BOOTPC (i.e. 67 and 68) UDP ports for communication between DHCP Relay and DHCP Server will recover the system.	

Defect ID: DEFECT000615564	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.2	Technology: LAG - Link Aggregation Group
Symptom: If a port channel interface is configured as tracking interface for an interface which exists before this port channel interface in output of "show running-config" then during replay of this configuration file will cause the issue. It throws the error that it cannot find particular port channel interface.	
Condition: This issue can occur during configuration file replay in which a port channel can be configured as tracking interface.	

Defect ID: DEFECT000615646	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: IPv6 Addressing
Symptom: Prefix is advertised in the IPv6 RA messages even though it is configured with "no-advertise" option.	
Condition: Prefix is configured using "ipv6 nd prefix" with "no-advertise"	
Workaround: Do not configure prefix if it should not be present in IPv6 RA messages.	

Defect ID: DEFECT000615651	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: IPv6 Addressing
Symptom: ipv6 nd prefix <prefix> with "off-link" option does not work.	
Condition: execution of ipv6 nd prefix <prefix> CLI with "off-link" option	
Workaround: NA	

Defect ID: DEFECT000616035	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.1.0	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: OSPF adjacency is stuck in ex-start state for some of the Ve interfaces.	
Condition: When OSPF is configured on a SAG enabled interface and the interface is reconfigured, during the re-convergence, some of the OSPF sessions could be stuck in ex-start state.	

Defect ID: DEFECT000616334	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Release: Network OS7.1.0	Technology: IP Fabric
Symptom: L3 traffics are not forwarded correctly.	
Condition: The environment have lots of flows which generate more than 3K hash results and some hash values are shared by 2 or more flows.	
Workaround: Reduce the total flows or consider re-arrange the private subnet prefix if there are private subnet.	
Recovery: Clear the host table.	

Defect ID: DEFECT000616345

Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS6.0.2	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: In some cases, the problem manifests itself as kernel panic occurs due to "Out Of Memory" condition. In other cases, control plane traffic is unable to egress on some of the ports.	
Condition: The problem is known to happen only with 10G edge ports.	

Defect ID: DEFECT000616987	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: BFD - BiDirectional Forwarding Detection
Symptom: BFD session is not switched over to other available links if existing BFD session is deleted and added back.	
Condition: BFD session remains in INIT state thus causing registered protocols with BFD to converge in longer duration.	

Defect ID: DEFECT000617049	
Technical Severity: Medium	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.2	Technology: VLAN - Virtual LAN
Symptom: Static-MAC for multicast-mac address floods the packet after removing static-ARP and static-MAC entry and re-configuring.	
Condition: Static multicast MAC configured as static ARP.	

Defect ID: DEFECT000617313	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: VCS
Reported In Release: Network OS7.0.1	Technology: TRILL - Transparent Interconnection of Lots of Links
Symptom: RTE capture won't work for breakout interface	
Condition: when ingress/trill port is breakout mode	

Defect ID: DEFECT000617646	
Technical Severity: Low	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS7.0.1	Technology: Syslog
Symptom: Under certain conditions Syslog message giving source IP as MGMT-interface IP, when the reachability is via inband	

Condition: 1. When both inband and OOB both IP's are present and active in MGMT-VRF
2. Syslog server is connected through inband

Defect ID: DEFECT000617886	
Technical Severity: Critical	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS6.0.1	Technology: VLAN - Virtual LAN
Symptom: VDX experience unexpected reload due to Out-Of-Memory condition. Also some of the ports are unable to transmit.	
Condition: Known to happen with 10G ports that have copper-pigtail connector. And the link-partner is not an Extreme device.	

Defect ID: DEFECT000617919	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: IPv6 Addressing
Symptom: Unable to configure update-source for IPv6 interface, it throws syntax error: "xx/x/101" is an invalid value.	
Condition: Configure update-source for IPv6 interface which is greater than 99.	

Defect ID: DEFECT000618317	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Monitoring
Reported In Release: Network OS7.1.0	Technology: RAS - Reliability, Availability, and Serviceability
Symptom: Termination of raslogd process after upgrading from 7.0.1 to 7.1.0	
Condition: In cluster environment after updating firmware.	
Recovery: Raslogd will restart automatically.	

Defect ID: DEFECT000618713	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: Network OS7.1.0	Technology: OpenStack Integration
Symptom: The VDX6940-144S 10G passive cable (1m and 3m) interfaces do not display the interface "link down" RASLOG message when the corresponding 10G interface on the remote end is shut down	
Condition: Shutting down 10G interfaces when remote switch is a VDX6940-144S connected with 10G passive cables (1m and 3m)	
Workaround: Shut the 10G interface on the local interface	

Recovery: Shut the 10G interface on the local interface

Defect ID: DEFECT000619405	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: Network OS6.0.2	Technology: OpenStack Integration
Symptom: CRC errors when using 40g DAC (direct attach copper) cable with VDX6940	
Condition: 40g DAC (direct attach copper) cable with VDX6940	

Defect ID: DEFECT000619425	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.1.0	Technology: VXLAN - Virtual Extensible LAN
Symptom: Traffic loss on Port-channel interface	
Condition: If Global MTU is smaller than Port-channel MTU or Global MTU is configured and un-configured, user may see traffic loss on port-channel interface.	
Workaround: Configure MTU same as port-channel on Port-channel member interfaces	

Defect ID: DEFECT000619467	
Technical Severity: Low	Probability: Medium
Product: Extreme Network OS	Technology Group: Management
Reported In Release: Network OS7.0.1	Technology: Inband Management
Symptom: ZR optics are undetected and shows data access errors when connected to edge ports other than xx/x/1	
Condition: ZR optic connected to edge ports other than xx/x/1	
Workaround: Connect ZR optic on first interface.	
Recovery: Connect ZR optic on first interface and reseal the other interface ZR optic.	

Defect ID: DEFECT000619719	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Security
Reported In Release: Network OS7.0.1	Technology: SSH - Secure Shell
Symptom: Telnet/ssh for default-vrf enables though user configured as disabled.	
Condition: If node disconnected and re-joined to the fabric after "no telnet server use-vrf default-vrf" OR "no ssh server use-vrf default-vrf"	
Workaround: Disable Telnet/ssh using "telnet server use-vrf default-vrf shutdown" or "ssh server use-vrf default-vrf shutdown".	

Recovery: After node rejoins the fabric, to disable the telnet/ssh, execute the CLIs "telnet server use-vrf default-vrf shutdown" or "no telnet server use-vrf default-vrf" for telnet and "ssh server use-vrf default-vrf shutdown" or "no ssh server use-vrf default-vrf".

Defect ID: DEFECT000620197	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: OSPF - IPv4 Open Shortest Path First
Symptom: Configuration of OSPF authentication key is not applied when done using config-replay.	
<p>Condition: The issue is observed for below sequence of steps:</p> <ol style="list-style-type: none"> 1. Configure OSPF authentication key on interface using CLI. 2. Save running configuration using command: copy running-config flash://<file-name> 3. Remove configured OSPF authentication key using CLI. 4. Replay saved configuration by using command: copy flash://<file-name> running-config <p>It is observed that OSPF authentication key is not applied after step-4 though it was expected to be applied on the interface.</p>	
Workaround: After config-replay fails to configure OSPF authentication key on the interface, it is possible to configure authentication key using CLI.	

Defect ID: DEFECT000620617	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Release: Network OS7.0.1	Technology: xSTP - Spanning Tree Protocols
Symptom: VDX6940 device may see traffic loss if HA failover or ISSU operation is performed from nos7.0.1 to nos7.0.1a release.	
Condition: 1) RSTP is configured 2) HA failover or ISSU is performed	
Recovery: Disable/enable spanning-tree protocol on the interface	

Defect ID: DEFECT000620922	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.1.0	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: 'neighbor peer-group shutdown generate-rib-route' command doesn't generate rib-out for peers in the group.	
Condition: When peer-group is used to shut neighbors and generate rib-out, it doesn't generate rib-out.	
Workaround: Configure command per peer, for ribout generation as a work around	

Defect ID: DEFECT000621212	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.0.1	Technology: BGP4 - IPv4 Border Gateway Protocol
Symptom: Routes are advertised to peers with AS that's present in AS path segment of the route, though enable-peer-as-check is configured	
Condition: When 4 byte AS number support is enabled for the BGP sessions, the issue shall be seen	
Workaround: Disable 4 byte ASN support if possible	
Recovery: Upgrade to latest firmware or disable 4 byte ASN support to recover	

Defect ID: DEFECT000622620	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: Network OS7.0.1	Technology: OpenStack Integration
Symptom: 4 x 10 GbE breakout ports between VDX 6940-36Q and VDX 6740-1G may flap.	
Condition: 4 x 10 GbE breakout ports between VDX 6940-36Q and VDX 6740-1G may flap.	
Workaround: Perform shut / no shut on ports to stop the flapping.	
Recovery: Perform shut / no shut on ports to stop the flapping.	

Defect ID: DEFECT000622750	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Layer 3 Routing/Network Layer
Reported In Release: Network OS7.1.0	Technology: IPv6 Addressing
Symptom: When the user updates an IPv6 prefix with preferred lifetime alone, valid lifetime changes to default value.	
Condition: The issue happens only when the user updates the preferred lifetime value to an already configured IPv6 prefix with valid and preferred lifetime.	

Defect ID: DEFECT000623309	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Network Automation and Orchestration
Reported In Release: Network OS6.0.1	Technology: OpenStack Integration
Symptom: CRCs occur on VDX 6940-36Q when DAC (direct attached copper) cable is used with DELL NIC server, and DELL NIC server is running traffic towards VDX 6940-36Q while the VDX switch is booting up.	

Condition: CRCs occur on VDX 6940-36Q when DAC (direct attached copper) cable is used with DELL NIC server, and DELL NIC server is running traffic towards VDX 6940-36Q while the VDX switch is booting up.
Workaround: Reboot VDX switch, with DAC cabled ports administratively down state (save port configuration as "no shut", then reboot); or stop all traffic coming into DAC cabled ports from DELL NIC servers.
Recovery: Reboot VDX switch, with DAC cabled ports administratively down state (save port configuration as "no shut", then reboot); or stop all traffic coming into DAC cabled ports from DELL NIC servers.

Defect ID: DEFECT000623711	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: Data Center Fabric
Reported In Network OS7.1.0 Release:	Technology: VCS Fabric
Symptom: Any packet transmitted from CPU gets dropped on FC port.	
Condition: Can happen only on FC port.	

Defect ID: DEFECT000624394	
Technical Severity: High	Probability: Low
Product: Extreme Network OS	Technology Group: Layer 2 Switching
Reported In Network OS7.1.0 Release:	Technology: VLAN - Virtual LAN
Symptom: Continuous ASIC errors causes chip fault.	
Condition: Heavy ASIC activity can cause the issue.	

Defect ID: DEFECT000624701	
Technical Severity: High	Probability: Medium
Product: Extreme Network OS	Technology Group: Security
Reported In Network OS6.0.0 Release:	Technology: Security Vulnerability
Symptom: Network OS/SLX kernel (Network OS/Host/TPVM) are all vulnerable. User can overwrite the etc/password with root access.	
Condition: CVE-2016-5195 - kernel > 2.6.22 can hit this Dirty COW issue.	

Defect ID: DEFECT000625527	
Technical Severity: High	Probability: High
Product: Extreme Network OS	Technology Group: IP Multicast
Reported In Network OS7.1.0 Release:	Technology: PIM - Protocol-Independent Multicast
Symptom: Multicast functionality daemon "PIMd" goes down with memory leak.	
Condition: On enabling PIM, for every 60 seconds there is a memory leak of 5K bytes.	
Workaround: Do not enable PIM on router.	
Recovery: Disable PIM on router and reboot the router. Do not enable PIM after reboot.	

Closed without code changes for Network OS v7.4.0

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change in Network OS v7.4.0.

Parent Defect ID:	NOS-34097	Issue ID:	NOS-34097
Reason Code:	Feature/Function Not Supported	Severity:	S3 - Medium
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS5.0.0	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	Nexthop change using outbound route-map is not allowed for EBGp neighbor connection.		
Condition:	When Route-map with set-nexthop is used as outbound policy for BGP neighbor.		

Parent Defect ID:	NOS-38054	Issue ID:	NOS-38054
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS5.0.2b1	Technology:	Management GUI
Symptom:	Unexpected reload of standby management module in VDX8770.		
Condition:	Reloading of standby management module without any user intervention.		

Parent Defect ID:	NOS-38056	Issue ID:	NOS-38056
Reason Code:	Will Not Fix	Severity:	S3 - Medium
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS5.0.2b1	Technology:	Logical Chassis
Symptom:	HA sync fails between active and standby management modules in VDX 8770 because of cluster.configuration and VCS.configurations are not synchronized.		
Condition:	HA sync fails occasionally between active and standby management modules.		

Parent Defect ID:	NOS-47833	Issue ID:	NOS-47833
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer

Reported in Release:	NOS6.0.2b1	Technology:	BFD - BiDirectional Forwarding Detection
Symptom:	Some of the BFD session over Ve interface will be seen as Down state.		
Condition:	One of the system for the BFD session is dropping the packet, resulting in DOWN state.		
Workaround:	Workaround is to do one of the following: - shut / no shut of the interface - un-config/ config of OSPF BFD.		

Parent Defect ID:	NOS-47864	Issue ID:	NOS-47864
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS6.0.2c	Technology:	VXLAN - Virtual Extensible LAN
Symptom:	VDX 6940 can undergo unexpected reload during upgrade from NOS6.0.2c to NOS7.0.1b		
Condition:	VDX6940 is upgraded from 6.0.2c to 7.0.1b		

Parent Defect ID:	NOS-48003	Issue ID:	NOS-48003
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS6.0.2g	Technology:	Logical Chassis
Symptom:	System may undergo unexpected reload because of kernel panic		
Condition:	This may be seen when VCS id of the neighbor node is changed from primary node		

Parent Defect ID:	NOS-48024	Issue ID:	NOS-48024
Reason Code:	Will Not Fix	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-49089	Issue ID:	NOS-49089
Reason Code:	Will Not Fix	Severity:	S3 - Medium
Product:	Network OS	Technology Group:	Monitoring

Reported in Release:	NOS7.0.0	Technology:	MAPS - Monitoring and Alerting Policy Suite
Symptom:	device connectivity config should be consistent on all the links in the port-channel		
Condition:	port-channel members configured as different type NAS, iSCSI		
Workaround:	Configure all members to be in same type.		

Parent Defect ID:	NOS-52514	Issue ID:	NOS-52514
Reason Code:	Insufficient Information	Severity:	S3 - Medium
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.0b	Technology:	IP Fabric
Symptom:	Host ARP is learnt even when host IP subnet does not match to VE IP subnet.		
Condition:	Host is connected to a VLAN where the Ve IP subnet is different than the host IP subnet.		
Workaround:	Disable proxy ARP on VE		

Parent Defect ID:	NOS-52929	Issue ID:	NOS-52929
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.0.1a	Technology:	ACLs - Access Control Lists
Symptom:	"show access-list ip" CLI will list only local node access-list configuration.		
Condition:	Different access-lists are configured on the management interfaces across the cluster.		
Workaround:	"show access-list rbridge-id" or "show access-list interface" CLI can be used to display the access list of desired RBridge/interface.		

Parent Defect ID:	NOS-53060	Issue ID:	NOS-53060
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.0.1c	Technology:	IP Fabric
Symptom:	Traffic forwarding issue seen between two node dual homed Leaf witch in IP Fabric topology.		
Condition:	When we remove one of the nodes from the two node VCS Leaf.		

Parent Defect ID:	NOS-53592	Issue ID:	NOS-53592
Reason Code:	Will Not Fix	Severity:	S2 - High

Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0	Technology:	ICMP - Internet Control Message Protocol
Symptom:	VDX does not update its own CurHopLimit.		
Condition:	when the device has been configured to advertise a different AdvCurHopLimit value.		
Workaround:	Currently 2 separate commands exist to achieve needed functionality ipv6 nd reachable-time <millisec> and ipv6 nd cache expire time <secs> ipv6 nd hoplimit <hlimit> and set proc entry.		

Parent Defect ID:	NOS-54297	Issue ID:	NOS-54297
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.1.0	Technology:	IP Fabric
Symptom:	Some MAC addresses learnt via BGP are not seen in mac-address-table		
Condition:	When "mac-learning protocol bgp" for sites are frequently toggled, some MAC addresses are not seen in the BGP EVPN table.		

Parent Defect ID:	NOS-54456	Issue ID:	NOS-54456
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.1.0	Technology:	IGMP - Internet Group Management Protocol
Symptom:	IGMPv2 report will be sent back on same VxLAN tunnel where the report was received from if the tunnel is terminated on TRILL ports.		
Condition:	VxLAN is terminated on TRILL port on VDX6940.		
Workaround:	VxLAN tunnel is terminated on edge ports that are non-TRILL Ports.		

Parent Defect ID:	NOS-54460	Issue ID:	NOS-54460
Reason Code:	Will Not Fix	Severity:	S3 - Medium
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.1.0	Technology:	Logical Chassis
Symptom:	When the "show ip int brief" CLI is executed on a VDX8770 switch, the output under the column "Protocol" does not contain the reason for a particular interface to be in state "down".		
Condition:	When the "show ip int brief" CLI is executed on a VDX8770 switch.		

Parent Defect ID:	NOS-54479	Issue ID:	NOS-54479
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.1.0	Technology:	IP Fabric
Symptom:	Switch experience Out Of Memory (OOM) condition and reboots		
Condition:	Using Scaled Configurations		

Parent Defect ID:	NOS-54606	Issue ID:	NOS-54606
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	IPv6 Addressing
Symptom:	VDX sends neighbor advertisement(NA) message in response to neighbor solicitation(NS) even after the auto-configured link local IPv6 address has been rejected due to duplicated address detected (DAD).		
Condition:	This behavior is not compliant with RFC4862(clause 5.4.5).		

Parent Defect ID:	NOS-54610	Issue ID:	NOS-54610
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.1.0a	Technology:	Security Vulnerability
Symptom:	Switch allows Non-admin user to execute certain operational commands even though it is denied by RBAC Rule.		
Condition:	With view privileges the user is able to execute certain operational commands.		

Parent Defect ID:	NOS-54612	Issue ID:	NOS-54612
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.1.0a	Technology:	PIM - Protocol- Independent Multicast
Symptom:	"BSR-candidate interface" and "RP-candidate interface" configuration is lost during configuration replay from external server.		
Condition:	Configuration replay from external serve		

Parent Defect ID:	NOS-54657	Issue ID:	NOS-54657
Reason Code:	Already Implemented	Severity:	S3 - Medium

Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a	Technology:	OSPFv3 - IPv6 Open Shortest Path First
Symptom:	Device experienced sudden reload due to DCMd daemon termination.		
Condition:	Execution of "no ipv6 ospf cost" CLI command.		

Parent Defect ID:	NOS-54683	Issue ID:	NOS-54683
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0a1	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	IP Forwarding table shows the stale route entry learned from eBGP source even though the egress interface is in the down state.		
Condition:	BGP advertise/learn Prefix route(x.x.x.x/32) matches exactly with BGP peer address (x.x.x.x).		

Parent Defect ID:	NOS-54688	Issue ID:	NOS-54688
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.1.0b	Technology:	IPv6 Addressing
Symptom:	IPv6 nd is responding unexpectedly		
Condition:	During shutdown/no shutdown scenario		

Parent Defect ID:	NOS-54694	Issue ID:	NOS-54694
Reason Code:	Feature/Function Not Supported	Severity:	S2 - High
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.1.0b	Technology:	Configuration Fundamentals
Symptom:	Sometimes zoning CFG fails to enable		
Condition:	This error is seen when same name is given for both Zoning CFG and member of CFG		

Parent Defect ID:	NOS-54918	Issue ID:	NOS-54918
Reason Code:	Already Implemented	Severity:	S2 - High

Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.2.0	Technology:	ACLs - Access Control Lists
Symptom:	Unable to see security violation raslog messages. No functional impact.		
Condition:	Enforcing ACL with permit rules and then changing rule as deny.		

Parent Defect ID:	NOS-55025	Issue ID:	NOS-55025
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Network Automation and Orchestration
Reported in Release:	NOS7.2.0	Technology:	Scripting
Symptom:	Under rare conditions, the script may not provide the next hop with the required string.		
Condition:	This occurs when the "show ip route detail" command parsing does not yield results.		

Parent Defect ID:	NOS-55100	Issue ID:	NOS-55100
Reason Code:	Already Implemented	Severity:	S3 - Medium
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP sysName query returns hostname instead of FQDN.		
Condition:	When SNMP sysName OID is queried.		

Parent Defect ID:	NOS-55579	Issue ID:	NOS-55579
Reason Code:	Will Not Fix	Severity:	S4 - Low
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	BFD - BiDirectional Forwarding Detection
Symptom:	Multipath BFD session will not come up		
Condition:	Multipath BFD provisioned on non-default VRF		
Workaround:	Use default VRF for multipath BFD always.		

Parent Defect ID:	NOS-55930	Issue ID:	NOS-55930
Reason Code:	Question Answered	Severity:	S2 - High
Product:	Network OS	Technology Group:	Management

Reported in Release:	NOS7.2.0b	Technology:	Software Installation & Upgrade
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.		

Parent Defect ID:	NOS-56053	Issue ID:	NOS-56053
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS6.0.2a2	Technology:	TRILL - Transparent Interconnection of Lots of Links
Symptom:	ELD fails to work as expected with speeds lower than 1G when ports from same VCS cluster (different switches and same switch) are connected.		
Condition:	Loop is detected on ELD enabled links when speed on link changed from 10G or 1G to 100Mbps. Note: ELD is not supported on 100MB.		

Parent Defect ID:	NOS-66261	Issue ID:	NOS-66261
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS6.0.2d	Technology:	Configuration Fundamentals
Symptom:	After reload "show ip route vrf mgmt-vrf" showing routes when management port is in shutdown state		
Condition:	Reloading the switch with routes contained in mgmt.-vrf		

Parent Defect ID:	NOS-66264	Issue ID:	NOS-66264
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.0.0	Technology:	PIM - Protocol-Independent Multicast
Symptom:	Excess amount of traffic seen momentarily, during the HA failover of one of the VCS node, which is acting as FHR + LHR for one of the multicast stream.		
Condition:	If a router is FHR and LHR both, and there happens to be only one path between RP and this router. Assert scenario is hit with duplicate traffic from Source and RP.		

Parent Defect ID:	NOS-66265	Issue ID:	NOS-66265
Reason Code:	Will Not Fix	Severity:	S4 - Low
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.0.0	Technology:	Logical Chassis
Symptom:	Introducing a check to verify every time if port-channel count has exceeded 4K or not will bring down the performance. It is already documented that 4K VLAG's are supported.		
Condition:	User is allowed to configure more than 4K port-channels.		

Parent Defect ID:	NOS-66270	Issue ID:	NOS-66270
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.0.1c	Technology:	User Accounts & Passwords
Symptom:	Configuration of invalid encrypted password for existing user with encryption level as 7 it is getting accepted without throwing error.		
Condition:	VDX switch allows to change password as invalid encrypted password for existing user.		

Parent Defect ID:	NOS-66274	Issue ID:	NOS-66274
Reason Code:	Will Not Fix	Severity:	S2 - High
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.1.0	Technology:	MAPS - Monitoring and Alerting Policy Suite
Symptom:	MAPS raslog/email is not generated when rule is triggered when CRC counters got incremented after an unexpected system reload.		
Condition:	Issue is seen after unexpected reload of switch.		

Parent Defect ID:	NOS-48024	Issue ID:	NOS-66855
Reason Code:	Already Implemented	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		

Known Issues for Network OS v7.4.0

This section lists open software defects with Critical, High, and Medium Technical Severity in Network OS v7.4.0.

Parent Defect ID:	NOS-48022	Issue ID:	NOS-48022
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Network Automation and Orchestration
Reported in Release:	NOS6.0.2h	Technology:	OpenStack Integration
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		

Parent Defect ID:	NOS-53113	Issue ID:	NOS-53113
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Network Automation and Orchestration
Reported in Release:	NOS7.0.1c	Technology:	YANG
Symptom:	NOS fails to un-escape special characters in passwords received via Netconf XML for config backup upload.		
Condition:	Special characters used in password		
Workaround:	None		

Parent Defect ID:	NOS-53133	Issue ID:	NOS-53133
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.2	Technology:	IP Addressing
Symptom:	IP direct-broadcast is not working for 40G port		
Condition:	Not able to enable ip directed-broadcast config for 40G physical interface		

Parent Defect ID:	NOS-53151	Issue ID:	NOS-53151
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.0.2	Technology:	VRRPv2 - Virtual Router Redundancy Protocol Version 2

Symptom:	After VRRP master do "copy running startup" and "reload system", main module may reload because of panic in kernel
Condition:	In a dual main module setup with scaled VRRP setup, copy running config and system was reloaded with 'reload system' command.

Parent Defect ID:	NOS-54895	Issue ID:	NOS-54895
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.2.0	Technology:	IP Fabric
Symptom:	With 512 VRFs, rp_filter error logs may be seen on reload system		
Condition:	Scaling to 512 VRFs		

Parent Defect ID:	NOS-54939	Issue ID:	NOS-54939
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.2.0	Technology:	IP Fabric
Symptom:	When a user try to unconfigure "export route-map" config under a VRF, while using rbridge range, the error is seen.		
Condition:	When a user enters into rbridge range and try to unconfigure the VRF "export route-map" configure, the error occurs.		
Workaround:	A user can go to the specific rbridge and try to unconfigure the config.		

Parent Defect ID:	NOS-55028	Issue ID:	NOS-55028
Severity:	S2 - High		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.2.0	Technology:	IGMP - Internet Group Management Protocol
Symptom:	System reboot/reload is observed. It would also affect the traffic forwarding until the system comes up.		
Condition:	The issue is only seen, when IGMPv3 reports are received with Exclude mode for Multicast Source address, in a VLAN domain. Issue is usually observed on a high scale scenario, with around 1000 IGMPv3 Multicast Group addresses joined in a VLAN domain.		

Parent Defect ID:	NOS-55098	Issue ID:	NOS-55098
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.2.0	Technology:	ARP - Address Resolution Protocol
Symptom:	When uRPF is enabled, some packets are not forwarded		

Condition:	NULL route is configured for the source and uRPF is enabled
-------------------	---

Parent Defect ID:	NOS-55113	Issue ID:	NOS-55113
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-55139	Issue ID:	NOS-55139
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.2.0	Technology:	VLAN - Virtual LAN
Symptom:	L2 agent t crashes, while disabling protected port configuration on ports of castor switch.		
Condition:	When Virtual fabric resource limit [4004] is reached and when protected port configuration is tried, it is failing but, due to inconstant state L2 agent crashes.		
Workaround:	Do not try to apply protected port configuration beyond available resource limit on a castor switch.		

Parent Defect ID:	NOS-55243	Issue ID:	NOS-55243
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.2.0	Technology:	VLAN - Virtual LAN
Symptom:	Switch panics during cleanup of PVLAN configuration.		
Condition:	When below steps are tried as part of PVLAN configurations, switch panics during cleanup of the configuration. STEP 1. Configure Vp as primary, Vi as isolated, Vc as community vlan STEP 2. Associate Vp to Vi & Vc on primary vlan Vp. STEP 3. Configure A1A as trunk promiscuous port, A2A as trunk isolated, A3A as trunk community, A4A as trunk PVLAN port. STEP 4. Try enabling IGMP snooping on secondary vlans Vi & Vc.		

	STEP 5. Enable PVST/RPVST globally. STEP 6. Try configuring bridge priority for vlan Vi & Vc STEP 7. Now disable spanning-tree globally on all nodes in cluster. STEP 8. Try creating ve interface corresponding to secondary VLANs Vi & Vc.
Workaround:	.

Parent Defect ID:	NOS-55380	Issue ID:	NOS-55380
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0	Technology:	SNMP - Simple Network Management Protocol
Symptom:	SNMP MIB counters (a) ifOutUcastPkts (b) ifHCOUcastPkts (c) ifHCInUcastPkts showing incorrect values		
Condition:	When user sends Multicast/Broadcast L2 traffic, SNMP MIB counters (a) ifOutUcastPkts (b) ifHCOUcastPkts (c) ifHCInUcastPkts showing incorrect values		
Workaround:	User can use CLI to get the accurate values for (a) ifOutUcastPkts (b) ifHCOUcastPkts (c) ifHCInUcastPkts		

Parent Defect ID:	NOS-55384	Issue ID:	NOS-55384
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.2.0	Technology:	ACLs - Access Control Lists
Symptom:	Observing "Detected termination of process secd".		
Condition:	Enabling and disabling operation of DHCPconfiguration in a sequential order with ACL configuration.		

Parent Defect ID:	NOS-55403	Issue ID:	NOS-55403
Severity:	S2 - High		
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.2.0	Technology:	Logical Chassis
Symptom:	Traffic disruption for some of the Multicast routes may be observed.		
Condition:	Issue can be seen when multicast routes are scaled to maximum supported by PIM protocol.		

Parent Defect ID:	NOS-55405	Issue ID:	NOS-55405
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.2.0	Technology:	Logical Chassis

Symptom:	The port learned via IGMPv2 (*,G) mode will not receive the traffic in specific scenario
Condition:	When the same Multicast group is learned on 2 different ports, one port in IGMPv2 (*,G) mode another is inIGMPv3 (S,G) mode.

Parent Defect ID:	NOS-55429	Issue ID:	NOS-55429
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.2.0	Technology:	ACLs - Access Control Lists
Symptom:	No functional impact. Unable to see security violation raslog messages.		
Condition:	Configuring deny rule for IPV6 host.		

Parent Defect ID:	NOS-55451	Issue ID:	NOS-55451
Severity:	S1 - Critical		
Product:	Network OS	Technology Group:	Data Center Fabric
Reported in Release:	NOS7.2.0a	Technology:	IP Fabric
Symptom:	In rare scenario the MAC is not updated properly in VCS.		
Condition:	When there is a single link connection to the end host.		

Parent Defect ID:	NOS-55490	Issue ID:	NOS-55490
Severity:	S2 - High		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.2.0aa	Technology:	IGMP - Internet Group Management Protocol
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.		

Parent Defect ID:	NOS-55577	Issue ID:	NOS-55577
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	In IP Fabric topology, traffic may sometimes not get forwarded for VRF leaked routes.		
Condition:	L3VNI routes are leaked across VRFs.		

Parent Defect ID:	NOS-55610	Issue ID:	NOS-55610
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	OSPFv3 - IPv6 Open Shortest Path First
Symptom:	Termination of ospf6d daemon when continuous BFD flaps are observed for longer period of time.		
Condition:	Continuous BFD flaps in a scaled configuration scenario leading to OOM for ospf6d daemon		

Parent Defect ID:	NOS-55654	Issue ID:	NOS-55654
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	ARP - Address Resolution Protocol
Symptom:	MAC is UnResolved in show arp command		
Condition:	HA failover and clear arp no-refresh multiple times can lead to this.		

Parent Defect ID:	NOS-55680	Issue ID:	NOS-55680
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	IP Addressing
Symptom:	DHCP traffic sent from VLAG port-channel on VLAN is flooded back on same VLAG port-channel from other peer inside VCS cluster.		
Condition:	DHCP traffic is sent from VLAG port-channel on VLAN		

Parent Defect ID:	NOS-55684	Issue ID:	NOS-55684
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	BFD - BiDirectional Forwarding Detection
Symptom:	Multipath BFD session is not coming online/UP.		
Condition:	After un-provisioning and provisioning loopback interface.		

Workaround:	Un-provision and provision Multipath BFD configuration after provisioning loopback interface.
--------------------	---

Parent Defect ID:	NOS-55733	Issue ID:	NOS-55733
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.3.0	Technology:	VLAN - Virtual LAN
Symptom:	Unexpected switch reload.		
Condition:	During the shut operation on protected group interfaces.		

Parent Defect ID:	NOS-55788	Issue ID:	NOS-55788
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.3.0	Technology:	HTTP/HTTPS
Symptom:	When user imports a certificate using 'crypto ca import' command, certificate import may not succeed always.		
Condition:	When openssl fails to verify the certificate being imported, it leads to this issue.		
Workaround:	Retry 'crypto ca import' and it may import the certificate successfully.		

Parent Defect ID:	NOS-55816	Issue ID:	NOS-55816
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.3.0	Technology:	VLAN - Virtual LAN
Symptom:	Switch may undergo unexpected reload.		
Condition:	In scaled scenario, when protected port is disabled on a port.		
Workaround:	.		

Parent Defect ID:	NOS-55820	Issue ID:	NOS-55820
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.3.0	Technology:	LAG - Link Aggregation Group
Symptom:	Partial traffic drop when member port from static LAG is removed.		
Condition:	Issue is seen when Multicast VLAG load balancing (ip igmp snooping vlag-load-balancing, ipv6 mld snooping vlag-load-balancing) is enabled for a VLAN, which has IGMP/MLD member on Static LAG, and one of the LAG member port is later removed.		
Workaround:	Shut/no-shut on any existing member port of the Static LAG, or shut/no-shut on Po interface, will recover the traffic drop.		

Parent Defect ID:	NOS-55828	Issue ID:	NOS-55828
--------------------------	-----------	------------------	-----------

Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.3.0	Technology:	BFD - BiDirectional Forwarding Detection
Symptom:	Multipath BFD is does not become up		
Condition:	This can happen sometimes when BGP session is established between router ports with Unnumbered configuration.		

Parent Defect ID:	NOS-55905	Issue ID:	NOS-55905
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.4.0	Technology:	CLI - Command Line Interface
Symptom:	[NOS 7.4.0] RESTAPI - Vlan name configuration is failed while doing through RESTAPI		
Condition:	Getting error when will configure "name" (vlan) using REST.		

Parent Defect ID:	NOS-56134	Issue ID:	NOS-56134
Severity:	S2 - High		
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.3.0	Technology:	TRILL - Transparent Interconnection of Lots of Links
Symptom:	ARP is not resolved for the IP Fabric Gateway IP address		
Condition:	Same IP address is configured for both IP Fabric Gateway and Loopback interface.		

Parent Defect ID:	NOS-66171	Issue ID:	NOS-66171
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	DHCP - Dynamic Host Configuration Protocol
Symptom:	DHCP off log are not expected when dhcp is enable on management interface.		
Condition:	When user enable DHCP for management interface		

Parent Defect ID:	NOS-66172	Issue ID:	NOS-66172
--------------------------	-----------	------------------	-----------

Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.4.0	Technology:	CLI - Command Line Interface
Symptom:	Getting error while applying vlan group to the ports when restoring the backup config from external server		
Condition:	Applying Vlan group to the ports is not working when restoring the same from external server.		
Workaround:	.		

Parent Defect ID:	NOS-66196	Issue ID:	NOS-66196
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	Static Routing (IPv6)
Symptom:	IPv6 default gate ip is removed from mgmt int after giving shut/no shut on mgmt interface.		
Condition:	If static IPv6 address and default gateway are configured for management interface, after shut/no shut mgmt ip address operation, IPv6 gateway address is deleted in 'show ipv6 route vrf mgmt-vrf'		
Workaround:	User needs re-configure IPv6 default gateway to resume Ipv6 management connectivity.		

Parent Defect ID:	NOS-66200	Issue ID:	NOS-66200
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	ICMP - Internet Control Message Protocol
Symptom:	IPv6 MTU on management interface is not updated accordingly when IPv4 MTU is changed.		
Condition:	Change ipv4 MTU on management interface.		
Workaround:	Change IPv6 MTU accordingly.		

Parent Defect ID:	NOS-66204	Issue ID:	NOS-66204
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer

Reported in Release:	NOS7.4.0	Technology:	IP Addressing
Symptom:	IP Mroute configuration automatically gets applied to Management-VRF, after a sequence VCS cluster disruptions.		
Condition:	When Management interface IP address is assigned by DHCP, removing it and configuring Static IP address. If the Node is reloaded, the IP Mroute configured in the system also gets applied to management-vrf.		

Parent Defect ID:	NOS-66224	Issue ID:	NOS-66224
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.4.0	Technology:	VLAN - Virtual LAN
Symptom:	Configuration lost is not expected for port-channel		
Condition:	Switchport configuration is lost when copied from an external server to switch for a port-channel.		
Workaround:	.		

Parent Defect ID:	NOS-66250	Issue ID:	NOS-66250
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.4.0	Technology:	SNMP - Simple Network Management Protocol
Symptom:	For breakout ports, there is a mismatch between the output of "show media optical-monitoring" and SNMP MIB bsciOptMonInfoTable, w.r.t. Bias current and Rx power.		
Condition:	The user is using breakout ports.		

Parent Defect ID:	NOS-66271	Issue ID:	NOS-66271
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.0.2	Technology:	ACLs - Access Control Lists
Symptom:	'Message Generic Error' displayed in the CLI		
Condition:	Very rare to hit, during execution of 'no ip access-group <ACL_NAME>'		
Parent Defect ID:	NOS-66283	Issue ID:	NOS-66283
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Monitoring
Reported in Release:	NOS7.2.0	Technology:	Hardware Monitoring
Symptom:	Extra characters may appear in the output of "show media" in Date-Code field for some interfaces.		

Condition:	For some new optics Date-code field in the output of "show media" command may contain some extra characters.
-------------------	--

Parent Defect ID:	NOS-66284	Issue ID:	NOS-66284
Severity:	S2 - High		
Product:	Network OS	Technology Group:	VPN
Reported in Release:	NOS7.2.0	Technology:	EVPN - Ethernet VPN
Symptom:	System reboot/reload is observed. It would affect the traffic forwarding until system comes up.		
Condition:	The issue is seen only when Candidate RP is configured with more than 200 group range prefixes. Not a typical scenario.		

Parent Defect ID:	NOS-66290	Issue ID:	NOS-66290
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 2 Switching
Reported in Release:	NOS7.2.0	Technology:	VLAN - Virtual LAN
Symptom:	"Error: Vlan has only one member interface" will be thrown during configuration restoration from external FTP server.		
Condition:	Protected ports configuration fails during configuration restoration from external FTP server.		
Workaround:	Need to re-apply protected-port configuration after configuration restoration.		

Parent Defect ID:	NOS-66291	Issue ID:	NOS-66291
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.2.0	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	"show bgp evpn l3vni all-vrf" shows same VRF information two times.		
Condition:	running "show bgp evpn l3vni all-vrf"		

Parent Defect ID:	NOS-66295	Issue ID:	NOS-66295
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.2.0a	Technology:	CLI - Command Line Interface
Symptom:	Error message is seen when configuring the vrf under loopback interface.		
Condition:	Sometimes when configuring the VRF under Loopback interface.		
Workaround:	.		

Parent Defect ID:	NOS-66315	Issue ID:	NOS-66315
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	ICMP - Internet Control Message Protocol
Symptom:	ICMP type 3 code 4 messages are not generated on IPv4 MTU violation		
Condition:	The router is part of an IP Fabric. The functionality is not supported for IP Fabric.		

Parent Defect ID:	NOS-66768	Issue ID:	NOS-66768												
Severity:	S3 - Medium														
Product:	Network OS	Technology Group:	Other												
Reported in Release:	NOS7.4.0	Technology:	Other												
Symptom:	<p>when use these special characters ` \$ & (') / < \ " > ? ; in password with vrf , "copy config" command will fail, as these characters are not supported.</p> <p>RB94# show version Network Operating System Software Network Operating System Version: 7.4.0_bld07 Copyright (c) 2017-2018 Extreme Networks, Inc. Firmware name: 7.4.0_bld07 Build Time: 16:23:15 Jan 16, 2019 Install Time: 08:54:24 Jan 17, 2019 Kernel: 2.6.34.6 BootProm: 1.0.1 Control Processor: e500mc with 4096 MB of memory</p> <table border="1"> <thead> <tr> <th>Slot</th> <th>Name</th> <th>Primary/Secondary Versions</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>SW/0</td> <td>NOS</td> <td>7.4.0_bld07 7.4.0_bld07</td> <td>STANDBY</td> </tr> <tr> <td>SW/1</td> <td>NOS</td> <td>7.4.0_bld07 7.4.0_bld07</td> <td>ACTIVE*</td> </tr> </tbody> </table>			Slot	Name	Primary/Secondary Versions	Status	SW/0	NOS	7.4.0_bld07 7.4.0_bld07	STANDBY	SW/1	NOS	7.4.0_bld07 7.4.0_bld07	ACTIVE*
Slot	Name	Primary/Secondary Versions	Status												
SW/0	NOS	7.4.0_bld07 7.4.0_bld07	STANDBY												
SW/1	NOS	7.4.0_bld07 7.4.0_bld07	ACTIVE*												

	RB94# RB94# copy running-config ftp://sk:zxz\$\#abc@10.20.232.225//home/sk/test use-vrf mgmt-vrf Please check the username or password. RB94# copy running-config ftp://sk:zxz\$\#abc@10.20.232.225//home/sk/test/rb94.cfg RB94# copy running-config scp://sk:zxz\$\#abc@10.20.232.225//home/sk/test.cfg use-vrf mgmt-vrf Please check the username or password. RB94# copy running-config scp://sk:zxz\$\#abc@10.20.232.225//home/sk/test.cfg RB94# copy running-config ftp://sk:zxz\$\#abc@10.20.232.225//home/sk/test/rb94.cfg use-vrf mgmt-vrf Please check the username or password. ---> getting error while uploading the config using mgmt-vrf
Condition:	When use these special characters ` \$ & (') / < \ " > ? ; in password for copy config command with vrf.
Workaround:	Avoid using these special characters ` \$ & (') / < \ " > ? ; in password for the copy config command with vrf.

Parent Defect ID:	NOS-66795	Issue ID:	NOS-66795
Severity:	S2 - High		
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.4.0	Technology:	Logical Chassis
Symptom:	Observed panic core files after configuring principal priority.		
Condition:	Panic files are observed when tried to configure principal priority on a logical chassis.		

Parent Defect ID:	NOS-66841	Issue ID:	NOS-66841
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	ICMP - Internet Control Message Protocol
Symptom:	Switches in same VCS will always be able to ping each other		
Condition:	When acl hard-drop is configured on VE, switch will not be able to ping switches in other VCS as expected, but will be able to ping switches in same VCS.		
Parent Defect ID:	NOS-66849	Issue ID:	NOS-66849
Severity:	S2 - High		

Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.4.0	Technology:	PIM - Protocol-Independent Multicast
Symptom:	The pim bsr command (bsr-candidate interface <interface> mask <mask-len>) replay is failing.		
Condition:	This issue happens when pim bsr command is configured and saved, the system is rebooted.		
Workaround:	After reboot, perform no pim bsr and re-configure the command.		

Parent Defect ID:	NOS-66874	Issue ID:	NOS-66874
Severity:	S2 - High		
Product:	Network OS	Technology Group:	IP Multicast
Reported in Release:	NOS7.4.0	Technology:	IPv4 Multicast Routing
Symptom:	PIM Mcache entries for directly connected source, does not get populated, for those source interfaces where IP address was removed and replied after HA failover.		
Condition:	Issue is seen only with an unlikely sequence of steps. IP address is removed from an interface where Multicast source is present, and then HA Failover is performed. IP address is configured back on that same interface after systems comes up. PIM fails to program SG entries for sources on that interface.		
Workaround:	Workaround is to disable & enable PIM-SM in the affect interface after HA failover.		

Parent Defect ID:	NOS-66875	Issue ID:	NOS-66875
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	DHCP - Dynamic Host Configuration Protocol
Symptom:	When HA failover is initiated, DHCP relay is not able to relay packets properly		
Condition:	HA entries dump from Active to Standby for L3 DHCP relay agent is not happening properly. Due to which, if HA failover is initiated L3 DHCP relay may not work properly .		
Workaround:	Re-configuring L3 DHCP Relay address after HA failover will create database properly. L3 DHCP relay will work properly after that.		
Parent Defect ID:	NOS-66891	Issue ID:	NOS-66891
Severity:	S2 - High		

Product:	Network OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	NOS7.4.0	Technology:	IPv6 Addressing
Symptom:	When a new path is added to the existing IPv6 route, newly added path is not taken for traffic forwarding in hardware.		
Condition:	Static configuration of new IPv6 route path to existing IPv6 route.		

Parent Defect ID:	NOS-66909	Issue ID:	NOS-66909
Severity:	S3 - Medium		
Product:	Network OS	Technology Group:	Management
Reported in Release:	NOS7.4.0	Technology:	VMWare
Symptom:	On migration a VM from one EXSi server to another this migration is not shown in output of "show vnetwork vms vcenter ESXi1" command. Both of the EXSi servers are managed by same vcenter.		
Condition:	The issue is seen on migration of a VM from one EXSi server to other.		

Parent Defect ID:	NOS-66961	Issue ID:	NOS-66961
Severity:	S2 - High		
Product:	Network OS	Technology Group:	VCS
Reported in Release:	NOS7.4.0	Technology:	Logical Chassis
Symptom:	on the physical interface reload delay timer is kicking in with slot power off and on.		
Condition:	reload delay timer is configured on the physical interface and slot power off and power on is triggered		
Workaround:	shut/noshut on the interface to cancel the timer		

Parent Defect ID:	NOS-67013	Issue ID:	NOS-67013
Severity:	S2 - High		
Product:	Network OS	Technology Group:	Security
Reported in Release:	NOS7.4.0	Technology:	TACACS & TACACS+
Symptom:	Audit log might show the wrong username when the user successfully logs in		
Condition:	When a user tries to login using telnet and it fails and from the same session the user successfully logs in then audit log shows username of the failed login instead of successful login		
Workaround:	When login fails then try to login from a new session		