

December 2019

Network OS 7.2.0c for Extreme VDX

Release Notes v3.0

9036239-00 Rev AB

Copyright Statement

© 2019, 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 |
| Preface | 10 |
| Contacting Extreme Technical Support | 10 |
| Document feedback | 10 |
| Hardware | 12 |
| New devices | 12 |
| New interface modules..... | 12 |
| Deprecated Hardware | 12 |
| Software Features | 12 |
| Deprecated Software Features | 13 |
| Software Upgrade | 13 |
| New Software Features for Network OS v7.2.0c..... | 13 |
| L2 features..... | 13 |
| New Software Features for Network OS v7.2.0b | 13 |
| L3 features..... | 13 |
| New Software Features for Network OS v7.2.0a | 14 |
| New Software Features for Network OS v7.2.0 | 14 |
| L2 Multicast features | 14 |
| L3 features..... | 14 |
| Key RFEs from customers:..... | 14 |
| IP Fabric Enhancements:..... | 14 |
| CLI Changes | 14 |
| New Commands for Network OS v7.2.0c | 14 |
| New Commands for Network OS v7.2.0b | 14 |
| New Commands for Network OS v7.2.0a | 15 |
| New Commands for Network OS v7.2.0 | 15 |
| Modified Commands for Network OS v7.2.0c | 15 |
| Modified Commands for Network OS v7.2.0b | 15 |
| Modified Commands for Network OS v7.2.0a | 15 |
| Modified Commands for Network OS v7.2.0 | 15 |

| | |
|---|----|
| Deprecated Commands for Network OS v7.2.0c | 15 |
| Deprecated Commands for Network OS v7.2.0b | 16 |
| Deprecated Commands for Network OS v7.2.0a | 16 |
| Deprecated Commands for Network OS v7.2.0 | 16 |
| API Changes | 16 |
| Newly supported standards and RFCs | 16 |
| Supported devices | 17 |
| Deprecated Devices | 17 |
| Extreme VDX 6940-144S | 17 |
| Extreme VDX 6940-36Q | 18 |
| Extreme VDX 6740 | 18 |
| Extreme VDX 6740T | 18 |
| Extreme VDX 6740T-1G | 19 |
| Extreme VDX 8770-4 and VDX 8770-8 | 19 |
| Supported Blades for VDX 8770 | 20 |
| Support for 100-Mb interfaces | 20 |
| Supported power supplies | 20 |
| Supported Optics for Network OS v7.2.0 | 23 |
| 10GBase-T Copper SFP+ | 25 |
| Image filenames | 27 |
| Upgrade/Downgrade considerations | 27 |
| Migration Path | 27 |
| Management IP connectivity | 29 |
| Firmware Installation | 30 |
| Upgrading to this Release (Best Practices) | 31 |
| Downgrading to a Previous Release | 31 |
| Upgrade/downgrade Considerations for vLAG deployments | 32 |
| Upgrade/downgrade with default configuration | 34 |
| Management Gateway IP changes | 34 |
| Management Services | 35 |
| Telnet, SSH and AAA VRF support | 35 |
| CLI Changes for Telnet, SSH, AAA | 35 |

| | |
|--|----|
| HTTP VRF support..... | 35 |
| NTP VRF support | 35 |
| CLI Changes for NTP | 35 |
| SNMP- Community string maximum length increased to 64: | 35 |
| SNMP - Support for traps during hafailover: | 36 |
| SNMP-Trap Source IP support:..... | 36 |
| SNMP context based query: | 36 |
| SNMP MIB – VLAN update | 36 |
| SNMP Trap VRF Support | 36 |
| SNMP-Trap CLI | 37 |
| SNMP – IF MIB..... | 37 |
| Sflow VRF Support..... | 37 |
| Sflow-CLI | 37 |
| Syslog VRF Support..... | 37 |
| Syslog-CLI | 37 |
| Firmware download, Copy support, Copy config..... | 37 |
| Other Management Services | 38 |
| Scalability numbers | 39 |
| Compatibility and Interoperability | 47 |
| IP Storage | 48 |
| FC/FCoE Storage | 48 |
| Network Adapters | 49 |
| Limitations and Restrictions..... | 51 |
| Command Line Interface..... | 51 |
| Line cards | 53 |
| USB | 53 |
| Licensing..... | 53 |
| VCS..... | 54 |
| VCS Node Replacement | 55 |
| Logical Chassis | 56 |
| Extreme Trunks | 56 |
| Breakout Interfaces | 56 |

| | |
|--|----|
| Dual-personality Ports | 57 |
| 1G Mode | 57 |
| vLAG..... | 57 |
| Virtual IP Address Support | 58 |
| Security, Management ACLs, Authentication, Authorization | 58 |
| SPAN & RSPAN | 59 |
| MAC Learning Considerations in VCS | 59 |
| PVLAN | 59 |
| UDLD | 60 |
| STP/DiST..... | 60 |
| IGMPv3 Snooping | 61 |
| Edge Loop Detection (ELD)..... | 61 |
| Long Distance ISL Ports | 61 |
| AMPP and Port-Profiles | 62 |
| vCenter..... | 62 |
| QoS | 63 |
| FCoE | 64 |
| FlexPorts | 65 |
| Fibre Channel | 65 |
| Access Gateway..... | 65 |
| IP Fabric | 66 |
| ND/RA | 67 |
| BFD..... | 67 |
| VRRP | 68 |
| OSPFv3 | 69 |
| BGP | 69 |
| Layer 2/Layer 3 Multicast | 69 |
| VRF..... | 69 |
| BGP-VRF | 70 |
| ACL..... | 70 |
| Policy-based Routing (PBR) | 70 |
| Inter-VRF Leaking (Static)..... | 70 |

| | |
|--|----|
| DHCP IP Helper | 71 |
| Dynamic ARP Inspection (DAI) | 71 |
| DHCP-based Firmware Download (DAD – DHCP Automatic Deployment)..... | 71 |
| Link State Tracking..... | 72 |
| OpenFlow | 72 |
| Mac Port Based Authentication | 73 |
| Uplink Switch Support | 73 |
| Layer 2 and Layer 3 ISSU on VDX 6740x | 74 |
| Extreme Vyatta Controller (BVC)..... | 74 |
| REST API | 75 |
| NetConf..... | 75 |
| VXLAN Gateway for VMware NSX | 76 |
| VF Extension using VxLAN..... | 76 |
| TCAM Profiles..... | 77 |
| Management VRF | 77 |
| Conversational MAC Learning..... | 77 |
| System level Flow based QoS..... | 77 |
| Port level Flowbased QoS | 78 |
| URPF | 78 |
| BGP Auto neighbor discovery..... | 78 |
| Non-trivial Merge | 78 |
| HA on TOR switches..... | 78 |
| Logical Chassis HA..... | 78 |
| Interoperability | 79 |
| MAPS | 79 |
| Maintenance Mode | 80 |
| BNA..... | 80 |
| Miscellaneous..... | 80 |
| Defects..... | 81 |
| TSBs - Critical Issues to Consider Prior to Installing This Network OS Release | 81 |
| Network OS v7.2.0a Caveats | 81 |
| Network OS v7.2.0 Caveats..... | 82 |

| | |
|--|-----|
| Closed with code changes for Network OS v7.2.0c..... | 85 |
| Closed with code changes for Network OS v7.2.0b | 90 |
| Closed with code changes for Network OS v7.2.0a..... | 101 |
| Closed with code changes for Network OS v7.2.0 | 109 |
| Closed with code changes for Network OS v7.1.0 | 136 |
| Closed without code changes for Network OS v7.2.0c | 194 |
| Closed without code changes for Network OS v7.2.0b | 194 |
| Closed without code changes for Network OS v7.2.0a | 194 |
| Closed without code changes for Network OS v7.2.0 | 194 |
| Closed without code changes for Network OS v7.1.0 | 204 |
| Known Issues for Network OS v7.2.0c..... | 225 |
| Known Issues for Network OS v7.2.0b | 225 |
| Known Issues for Network OS v7.2.0a..... | 225 |
| Known Issues for Network OS v7.2.0 | 225 |
| Known Issues for Network OS v7.1.0 | 241 |

DOCUMENT HISTORY

| Version | Summary of Changes | Publication Date |
|----------------|---------------------------------|-------------------------|
| 1.0 | Initial Release | June 5, 2019 |
| 2.0 | Added information for NOS-55130 | November, 2019 |
| 3.0 | Updated PVST Limits | December, 2019 |

Preface

Contacting Extreme Technical Support

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

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

- GTAC (Global Technical Assistance Center) for immediate support
- Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.
- Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- GTAC Knowledge - Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- The Hub - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but 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

Hardware

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

New devices

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

New interface modules

None

Deprecated Hardware

None

Software Features

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

Deprecated Software Features

None.

Software Upgrade

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

New Software Features for Network OS v7.2.0c

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

L2 features

- Dynamic shutdown of BFD is supported.
- Up to 500 PVST and RPVST instances are supported on 6 ports, extendable up to 34 ports.
- **Limitations:**
 - Enabling Cisco BPDU MAC on some ports and Extreme BPDU MAC on other ports not supported.
 - It is not recommended to use Hello-time of less than 2 seconds.
 - This feature is not supported on any VDX 6940 or VDX 8770 switches.

xSTP BPDUs are transmitted with Extreme Multicast OUI i.e 01.e0.52 by default. Some vendor switches don't allow packets with multicast MAC as source MAC. To inter-operate with such vendor switches, "system-id oui" command is supported.

With non-default "system-id oui" configured, convergence of VDX-VCS with Cisco Peer-switch takes more time than usual.

This command is not recommended to be used with Cisco N9K switches.

New Software Features for Network OS v7.2.0b

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

L3 features

- VDX identifies IPv4 and IPv6 MTU violation and violated packets can be trapped to software for further processing. Traps on IPv6 MTU violation was enabled from initial NOS releases. Similarly, traps for IPv4 MTU will be enabled starting NOS 7.4.0 release and backported to NOS 7.2.0b patch.

Control plane has all the necessary support in place to handle such MTU violation [trapped packets] which generates ICMP notification and sends it to source to alert about this violation hit.

Note: Extreme VDX 6740 and Extreme VDX 8770 are the only supported platform for IP MTU changes in NOS 7.2.0b release.

- IGMP snooping is supported on 4k VLANs from this release and there is no change in the total number of groups, which is 6k in all default profile.

New Software Features for Network OS v7.2.0a

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

New Software Features for Network OS v7.2.0

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

L2 Multicast features

- IGMPv3 support

L3 features

- BGP Graceful Shutdown for maintenance mode (GSHUT)
- IP Unnumbered needs automatic or manual metric for link speed
- uRPF with Strict and Loose Mode IPv4 Support

Key RFEs from customers:

- Option to Block VLAN creation when importing Port Profiles from vCenter
- MIB that would report on the ARP and IP tables for multiple VRFs
- Uplink switch support on VCS and Virtual Fabrics
- Display local and remotely learned MAC address via SNMP on VDX switches
- 100Mbps support on 6940-144S

IP Fabric Enhancements:

- Automatic BGP neighbor discovery based on LLDP
- L2 IGMPv1v2v3 snooping and IP Multicast in default-VRF support for IP Fabric with VxLAN
- Testing Connectivity across IP Fabric -Ping/Traceoute
- L2 traceroute for VxLAN

Scale Enhancements

- Egress ACL support increased to 2k

CLI Changes

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

New Commands for Network OS v7.2.0c

- Bfd-shutdown

New Commands for Network OS v7.2.0b

- None

New Commands for Network OS v7.2.0a

The following configuration commands are new in this release:

- snmp-server host 10.20.58.116 public severity-level Info **source-interface management**

New Commands for Network OS v7.2.0

The following configuration commands are new in this release:

- neighbor lf-spn-grp accept-lldp-neighbors
- advertise bgp-auto-nbr-tlv
- rpf-mode
- Uplink-switch enable
- Protected-port enable
- vcenter <vcenter-name>
- ip igmp static-group
- neighbor x.x.x.x graceful-shutdown <g-shut timer>

The following show commands are new in this release:

- show statistics rpf
- show protected-ports [vlan <vlan-id>]

Modified Commands for Network OS v7.2.0c

- None

Modified Commands for Network OS v7.2.0b

- None

Modified Commands for Network OS v7.2.0a

- None

Modified Commands for Network OS v7.2.0

- None

Deprecated Commands for Network OS v7.2.0c

- None

Deprecated Commands for Network OS v7.2.0b

- None

Deprecated Commands for Network OS v7.2.0a

- None

Deprecated Commands for Network OS v7.2.0

- None

API Changes

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

Newly supported standards and RFCs

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

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

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

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

HARDWARE SUPPORT

Supported devices

Extreme Network OS v7.2.0 supports following VDX Switches:

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

Deprecated Devices

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

- Extreme VDX 2741
- Extreme VDX 2746

Extreme VDX 6940-144S

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

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

Extreme VDX 6940-36Q

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

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

Extreme VDX 6740

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

These ports support the following:

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

Extreme VDX 6740T

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

- Available in 24, 48 and 64 port SKU.
- 3 microsecond latency for any port to port to assure rapid response for latency-sensitive applications.

- The base SKU is available with 24 10GbE Base-T ports and can be upgraded up to 48 ports via 10Gbe DPOD license of 8 ports.
- It has 4 X 40 GbE QSFP ports which can be used for uplink and VCS fabric formation.
- Each 40GbE port is capable of doing a breakout of 4 x 10GbE ports.
- Each 40GbE port is also capable of doing a FC breakout of 4*8G or 4*16G. These ports can be used to connect to the FOS switches.
- Each 40GbE port is also capable of doing an FC breakout of 4 x 8G/16G.
- Additional 4X40GbE ports can be added to base version with 2X40GbE POD license increments.
- 100Mb Support – Refer to “Support for 100-Mb interfaces” below.

Extreme VDX 6740T-1G

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

- Base version is available with 48 x 1000BASE-T ports and 2 x 40 GbE QSFP+ ports.
- 3-microsecond latency for any port to port to assure rapid response for latency-sensitive applications.
- All 48 x 1000BASE-T ports can be upgraded to 10Gbase-T port with capacity on demand license.
- Additional 2X40Gbe ports can be added to base version with 2X40Gbe POD license.
- It has 4 X 40Gbe QSFP ports which can be used for the uplink and VCS fabric formation.
- Each 40GbE port is capable of doing a breakout of 4 X 10GbE ports.
- Each 40GbE port is also capable of doing a FC breakout of 4 x 8G/16G.
- 100Mb Support – Refer to “Support for 100-Mb interfaces” below.

Extreme VDX 8770-4 and VDX 8770-8

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

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

Supported Blades for VDX 8770

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

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

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

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

Supported power supplies

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

| Part number | Description | Compatible devices |
|----------------|---------------------------|------------------------|
| XBR-ACPWR-3000 | FRU,3000W AC POWER SUPPLY | VDX 8770-4, VDX 8770-8 |
| XBR-DCPWR-3000 | FRU,3000W DC POWER SUPPLY | VDX 8770-4, VDX 8770-8 |

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

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

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

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

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

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

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

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

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

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

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

Supported Optics for Network OS v7.2.0

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

The VDX switches support following optics types listed below. The FC SFP+ optics are supported only on VDX 6740 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 |
| 100G-CFP2-LR4-10KM | 100 GbE CFP2 optic, LR4, for distances up to 10 km over SMF | | 57-1000285-01 | |
| 100G-CFP2-ER4-40KM | 100 GbE CFP2 optic, ER4, for distances up to 40 km over SMF | | 57-1000328-01 | |
| 100G-QSFP28-SR4 | 100 GbE SR4 QSFP28 optic for distances up to 100m over MMF. Supported on VDX6940-144S and VDX8770-4/8 platforms. | | 57-1000326-01 | |
| 100G-QSFP28-LR4L-2KM | 100 GbE QSFP28 optic for distances up to 2 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms. | | 57-1000329-01 | |
| 100G-QSFP28-LR4-10KM | 100 GbE QSFP28 optic for distances up to 10 km over SMF. Supported on VDX 6940-144s and VDX 8770 platforms. | | 57-1000334-01 | |

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

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

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

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

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 Methode Electronics or their partners. Its Methode 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.Extremenetworks.com

| Image Filename | Description | Supported Device or Module |
|---------------------------|---|----------------------------|
| nos7.2.0c.tar.gz | Network OS v7.2.0c for unix | NA |
| nos7.2.0c.zip | Network OS v7.2.0c for Windows | NA |
| nos7.2.0c_all_mibs.tar.gz | Network OS v7.2.0c MIBS | NA |
| NOS_7.2.0c_Release_Notes | Network OS v7.2.0c Release Notes v1.0 (PDF) | NA |
| nos7.2.0c.md5 | Network OS v7.2.0c MD5 Checksum | NA |

Upgrade/Downgrade considerations

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

Migration Path

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

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

| To \ From | 6.0.1x | 6.0.2x | 7.0.0 | 7.0.1x | 7.1.0x | 7.2.0x |
|-----------|-----------------------------------|--------|----------|----------|----------------|----------------|
| 6.0.1x | ISSU for upgrade; Coldboot for | ISSU | coldboot | coldboot | default-config | default-config |

| | | | | | | |
|---------------|----------------|---|----------------|---|---|---|
| | downgrade | | | | | |
| 6.0.2x | coldboot | ISSU for upgrade; Coldboot for downgrade | coldboot | coldboot | default-config | default-config |
| 7.0.0 | coldboot | coldboot | NA | ISSU* | coldboot | default-config |
| 7.0.1x | coldboot | coldboot | coldboot | ISSU for upgrade; Coldboot for downgrade | coldboot | default-config |
| 7.1.0x | default-config | default-config | coldboot | coldboot | ISSU for upgrade; Coldboot for downgrade | Coldboot |
| 7.2.0x | default-config | default-config | default-config | default-config | coldboot | ISSU for upgrade ⁸ ; Coldboot for downgrade |

NOTES

1. ** CFP2 to QSFP28 conversion module (PN: 80-1008646-01) Version3 downgrade to any release prior to Network OS7.0.1 will cause CRC errors on the link.
2. Only Extreme Network Advisor (BNA) v14.2.1 (available separately) and above supports Network OS v7.2.0. It is required to first upgrade to BNA v14.2.1 and then upgrade switches to Network OS v7.2.0.
3. Before downgrading to lower releases, it is recommended to disable all new features that are not supported on lower releases by using the “no” version of the CLIs. Stray configurations left out before downgrade can cause undesired behavior.
4. While upgrading chassis-based system, under stress condition (e.g. due to excessive processing load on the processor), some linecards may become faulty during firmware download. To recover, run “power off <linecard>” followed by “power on <linecard>” command.
5. You must remove the IGMP snooping static mrouter configuration from all VLANs before upgrading or downgrading from or to the Network OS 6.0.2x release.
6. Firmware download from Network OS7.0.1a to Network OS6.x or Network OS5.x with default-config option needs AG mode disabled.
7. **Limitations:

- a) After downgrading from Network OS v7.1.0 to Network OS v5.0.x FCoE devices may not log back in or FCoE configuration may be lost. To recover, reload the switch. Alternate recovery method: re-configure FCoE by removing and adding fcoeport configuration (no fcoeport/fcoeport default) on the affected interfaces.
 - b) In rare occurrence, 40G links may not come up online after upgrade to 7.1.0, need to do shut/no shut to recover
 - c) In VDX 8770 platforms, After upgrade from 6.0.2 to 7.1.0 with coldboot, SNMP V3 traps are not received for the V3host which is under Rbridge.
 - d) Dport test between VDX 6740T and VDX 6940-144S breakout link may fail in upgrade to 7.1.0 and above.
8. ISSU upgrade from NOS 7.2.0/NOS7.2.0a/NOS7.2.0ax to NOS 7.2.0b is not supported and blocked. Please refer "669062" defect below in NOS 7.2.0b closed defect list for further details.
9. ISSU upgrade from NOS 7.2.0b to NOS 7.2.0c is also not supported because of open defect.

Management IP connectivity

In regards to SNMP, firmware downgrade from Network OS v7.1.0 to v7.0.x/v6.0.x/v5.0.x that do not support "use-vrf" keyword, the host/v3host with use-vrf value as "default-vrf" or "user-defined vrf" is not supported. The host/v3host configuration should set the use-vrf value as "mgmt-vrf" before downgrade.

Also, firmware downgrade from Network OS v7.1.0 and above to v7.0.x/v6.0.x/v5.0.x with use-vrf option in host/v3host set to user-defined vrf is not supported. The host/v3host configuration should set the use-vrf value as "mgmt-vrf" or "default-vrf" before downgrade.

Firmware upgrade to Network OS v7.1.0 and above from v7.0.x/v6.0.x/v5.0.x that do not support "use-vrf" keyword will modify the host/v3host configuration to append "use-vrf" keyword with value of mgmt-vrf and all the existing host/v3host entries will be assigned to mgmt-vrf.

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

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

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

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

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

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

Firmware Installation

In logical chassis cluster mode

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

This section includes special considerations and caveats to be aware of when upgrading to or from this version of Extreme Network OS, as well as recommended migration paths to use to reach this version of Extreme Network OS.

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

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 7.1.0 using the coldboot option. Principal node in a cluster should be last to be upgraded.
- After a node is upgraded, it will join the existing Network OS v7.2.0 cluster. Eventually, when all the nodes are upgraded, they will form one Network OS 7.2.0 VCS Cluster. [Note that no configuration changes are allowed during this time.]

Downgrading to a Previous Release


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

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.

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

Management Gateway IP changes

VDX Fixed-form switches (No L3 license required)

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

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

Note:

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

VDX 8770 (with L3 license/without L3 license)

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

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

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

Note:

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

Management Services

Telnet, SSH and AAA VRF support

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

CLI Changes for Telnet, SSH, AAA

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

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

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

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

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

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

HTTP VRF support

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

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

NTP VRF support

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

CLI Changes for NTP

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

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

SNMP- Community string maximum length increased to 64:

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

SNMP - Support for traps during hafailover:

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

SNMP-Trap Source IP support:

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

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

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

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

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

```
snmp-server host ip-address <community-string> source-interface management ?
```

Possible completions:

chassis-ip Use chassis IP as source address

mm-ip Use local MM IP as source address

SNMP context based query:

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

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

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

SNMP MIB – VLAN update

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

SNMP Trap VRF Support

SNMP is able to receive the packets from any VRF including mgmt-vrf/default-vrf and respond to the corresponding VRF from where the SNMP packet is received. The support is also added to send the

notification (trap) to the host/v3host configured in the switch through the vrf-name mapped with the host/v3host.

SNMP-Trap CLI

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

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

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

To disable per link TRAP under interface

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

SNMP – IF MIB

To display Interface details when linecard is powered-off

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

Sflow VRF Support

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

Sflow-CLI

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

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

Syslog VRF Support

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

Syslog-CLI

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

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

Firmware download, Copy support, Copy config

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

Other Management Services

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

SCALABILITY AND INTEROPERABILITY

Scalability numbers

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

| Network OS v7.2.0 Scalability Numbers | VDX 6740, 6740T, 6740T-1G | VDX 8770 | VDX 6940-36Q | VDX 6940-144S |
|---|---------------------------|----------|--------------|---------------|
| Maximum # of dot1Q VLANs (Virtual-Fabric Disabled) | 4096 | 4096 | 4096 | 4096 |
| Maximum # of VLANs (dot1Q + Virtual-Fabric) | 6000 | 8192 | 8192 | 8192 |
| Maximum # of Service Virtual Fabric VLANs | 2000 | 4096 | 4096 | 4096 |
| Maximum # of Transport Virtual Fabric VLANs | 1000 | 1000 | 1000 | 1000 |
| Maximum # of MAC addresses per Switch | 120000 | 256000 | 75000 | 75000 |
| Maximum # of MAC addresses per Fabric (with CML) | 512000 | 512000 | 512000 | 512000 |
| Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for VMware NSX | 8000 | N/A | 8000 | 8000 |
| Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for Virtual-Fabric Extension | 120000 | N/A | 75000 | 75000 |
| Maximum # of MAC-based Virtual-Fabric VLAN Classification per switch | 256 | 1024 | 1000 | 1000 |
| Maximum # of Classified Virtual Fabric VLANs per Trunk Interface | 2000 | 4096 | 4096 | 4096 |
| Maximum # of port profiles (AMPP) | 1000 | 1,000 | 512 | 512 |
| Maximum # of VLANs in port profiles | 3500 | 4000 | 3500 | 3500 |
| Maximum # of sites (tunnels) in Virtual-Fabric Extension | 50 | N/A | 50 | 50 |
| Maximum # of dot1q VLANs that can be attached on VxLAN GW for Virtual-Fabric Extension | 4000 | N/A | 4000 | 4000 |
| Maximum # of Virtual-Fabric (Service + Transport) VLANs that can be extended via Virtual-Fabric Extension | 2000 | N/A | 4000 | 4000 |
| Maximum # of dot1q VLANs + Virtual-Fabric VLANs enabled on edge-interfaces that can be attached to VxLAN GW and extended via Virtual-Fabric Extension | (2000+1000) | N/A | (2000+1000) | (2000+1000) |
| Max # of IGMP groups over Tunnels via Virtual-Fabric Extension | 6000 | N/A | 6000 | 6000 |
| Max # of BFD sessions over Virtual-Fabric Extension | 10 | N/A | 10 | 10 |

| Network OS v7.2.0 Scalability Numbers | VDX 6740, 6740T, 6740T-1G | VDX 8770 | VDX 6940- 36Q | VDX 6940- 144S |
|---|---------------------------------|------------|------------------|-------------------|
| Tunnels | | | | |
| Maximum # of dot1q VLANs that can be attached on VxLAN GW for VMware NSX | 2000 | N/A | 2000 | 2000 |
| Maximum # of VLANs (dot1q VLANs attached to VxLAN GW for NSX + Virtual Fabric VLANs enabled on edge-interfaces) | (2000+1,000) | N/A | (2000+1000) | (2000+1000) |
| Maximum # of VxLAN tunnels with VMware NSX | 250 | N/A | 250 | 250 |
| Maximum # of service-nodes with VMware NSX | 5 | N/A | 5 | 5 |
| Maximum # of MAC Associations for AMPP | 8000 | 4000 | 8000 | 8000 |
| Maximum # of per priority pause levels | 3 | 8 | 3 | 3 |
| Maximum # of VMware vCenters per Fabric | 4 | 4 | 4 | 4 |
| Maximum # of ELD instances in the fabric | 2000 | 2000 | 2000 | 2000 |
| Maximum # of IGMPv2v3 Snooping Interfaces supported | 512 | 512 | 512 | 512 |
| Learning rate for IGMP snooping (groups/second) | 512 | 512 | 512 | 512 |
| Maximum # of L2 (IGMPv2 Snooping) multicast groups | 6000 | 6000 | 6000 | 6000 |
| Maximum # of L2 (IGMPv3 Snooping) multicast groups | 4000 | 4000 | 4000 | 4000 |
| Maximum # of MLD Interfaces | 256 | 256 | 256 | 256 |
| Maximum # of MLD Groups | 4000 | 4000 | 4000 | 4000 |
| Learning rate for MLD snooping (groups/second) | 512 | 512 | 512 | 512 |
| # of L3 (S,G) forwarding Entries | 2000 | 2000 | 2000 | 2000 |
| # of L3 (*,G) forwarding Entries | 256 | 256 | 256 | 256 |
| # of L3 (*,G) joins per RP | 256 | NA | 256 | 256 |
| PIM Interfaces Supported | 32 | 32 | 32 | 32 |
| IGMP interfaces supported | 32 | 32 | 32 | 32 |
| Learning Rate for PIM-SM (flows/second) | 32 | 32 | 32 | 32 |
| Maximum # of L2 ACL(ingress/egress) * | 3000/120 | 12000/2000 | 6128/4960 | 6128/4960 |
| Maximum # of L3 ACL ipv4 (ingress/egress) * | 1500/120 | 12000/2000 | 3064/4750 | 3064/4750 |
| Maximum # of class-maps | 2048 | 2048 | 2048 | 2048 |
| Maximum # of policy-maps | 2048 | 2048 | 2048 | 2048 |
| Maximum # of class-maps per policy map | 50 | 50 | 50 | 50 |
| Maximum Total # of L3 ACL ipv6 (ingress/egress) * | 500/120 | 4000/2000 | 1000/500 | 1000/500 |
| Maximum # of VF/FCoE interfaces/Logins (Per switch) | 1000 | 1000 | 1000 | 1000 |
| Maximum # of Enodes/FCoE Devices per Fabric | 2000 | 2000 | 2000 | 2000 |
| Maximum # of NPIV per Port | 64 | 64 | 64 | 64 |
| Maximum # of SAN Devices (FC + FCoE) per Fabric | 3000 | 3000 | 3000 | 3000 |

| Network OS v7.2.0 Scalability Numbers | VDX 6740, 6740T, 6740T-1G | VDX 8770 | VDX 6940- 36Q | VDX 6940- 144S |
|--|---------------------------------|----------|------------------|-------------------|
| Maximum # of MSTP instance | 32 | 32 | 32 | 32 |
| Maximum # of VLAN in PVST | 500 | 500 | 500 | 500 |
| Maximum # of LAGs (Port Channels) | 64 | 288 | 144 | 144 |
| Maximum # of members in a standard LAG | 16 | 16 | 16 | 16 |
| Maximum # of members in 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 |
| 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 |

| Network OS v7.2.0 Scalability Numbers | VDX 6740, 6740T, 6740T-1G | VDX 8770 | VDX 6940- 36Q | VDX 6940- 144S |
|--|---------------------------------|----------|------------------|-------------------|
| Maximum # of FVG instances per system | 256 | 4096 | 1024 | 1024 |
| Maximum # of FVG instances per interface | 1 | 1 | 1 | 1 |
| Maximum # of routers participating in a FVG session | 32 | 32 | 32 | 32 |
| Maximum # of Gateway IP addresses per FVG instance | 1 | 1 | 1 | 1 |
| Maximum # of IPv4 routes with ECMP supported * | 8000 | 200000 | 10000 | 10000 |
| Maximum # of IPv6 routes with ECMP supported * | 1500 | 64000 | 1500 | 1500 |
| Maximum # of L3 ECMP | 16 | 32 | 32 | 32 |
| Maximum # of IPv4 interfaces per system *(Ve intf) | 2000 | 4000 | 2000 | 2000 |
| Maximum # of IPv6 interfaces per system *(Ve intf) | 512 | 4000 | 512 | 512 |
| Maximum # of VRF per node | 32 | 512 | 512 | 512 |
| Maximum # of VRFs support protocols per node | 32 | 128 | 128 | 128 |
| Maximum # of I-BGP peers | 256 | 512 | 256 | 256 |
| Maximum # of E-BGP peers | 256 | 256 | 256 | 256 |
| Maximum # of IPv4 BGP routes in HW * | 8000 | 200000 | 10000 | 10000 |
| Maximum # of IPv6 BGP routes in HW * | 1,500 | 64000 | 1500 | 1500 |
| Maximum # of IPv4 RIB (IN + OUT) Routes * | 110000 | 1300000 | 110000 | 110000 |
| Maximum # of IPv6 RIB (IN + OUT) Routes * | 110000 | 1300000 | 110000 | 110000 |
| Maximum # BGP IPv4/IPv6 Peer Group | 100 | 250 | 100 | 100 |
| Maximum # of BFD sessions per node | 100 | 100 | 100 | 100 |
| Maximum # of UDLD enabled interfaces | 64 | 384 | 144 | 108 |
| Maximum # of PVLAN domain supported | 1000 | 1000 | 1000 | 1000 |
| Maximum # of Secondary VLANs per PVLAN supported | 24 | 24 | 24 | 24 |
| Maximum # of primary VLANs per PVLAN supported in promiscuous mode | 24 | 24 | 24 | 24 |
| DHCP IP Helper Addresses per interface | 16 | 16 | 16 | 16 |
| DHCP IP Helper VE interfaces | 256 | 1,000 | 256 | 256 |
| DHCP IP Helper physical ports | 60 | 384 | 60 | 60 |
| DHCP IP Relay Addresses per Node | 2000 | 4000 | 2000 | 2000 |
| DHCP IPv6 Relay Address per Node | 2000 | 4000 | 2000 | 2000 |
| Max Number of configurable PBR route maps | 64 | 64 | 64 | 64 |
| Max Number of configurable PBR stanzas | 1024 | 1024 | 1024 | 1024 |
| Max Number of HW entries available for PBR | 512 | 8192 | 512 | 512 |
| Max Number of configurable next hops within a single PBR stanza | 128 | 128 | 128 | 128 |
| Max # of OpenFlow Active Connections | 1 | 1 | 1 | 1 |
| Max # of OpenFlow Passive Connections | 1 | 1 | 1 | 1 |
| Maximum # of OpenFlow L2 flows | 1000 | 4000 | 879 | 879 |

| Network OS v7.2.0 Scalability Numbers | VDX 6740, 6740T, 6740T-1G | VDX 8770 | VDX 6940- 36Q | VDX 6940- 144S |
|--|---------------------------------|----------|------------------|-------------------|
| Maximum # of OpenFlow L3 flows | 1000 | 4000 | 879 | 879 |
| Maximum # of Total OpenFlow GROUP | 768 | 768 | 768 | 768 |
| Maximum # of OpenFlow GROUP Type ALL | 256 | 256 | 256 | 256 |
| Maximum # of OpenFlow GROUP Type SELECT | 256 | 256 | 256 | 256 |
| Maximum # of OpenFlow GROUP Type INDIRECT | 256 | 256 | 256 | 256 |
| Max # of Buckets per GROUP ALL | 16 | 16 | 16 | 16 |
| Max # of Buckets per GROUP SELECT | 8 | 8 | 8 | 8 |
| Max # of Buckets per GROUP INDIRECT | 1 | 1 | 1 | 1 |
| Max # of ACTIONS per Bucket | 3 | 3 | 3 | 3 |
| Max # METERS | 1024 | 4096 | 1024 | 1024 |
| Maximum # of MAPS policy | 10 | 10 | 10 | 10 |
| Maximum # of MAPS rules | 250 | 250 | 250 | 250 |
| Maximum # of MAPS groups | 64 | 64 | 64 | 64 |
| Maximum # of MAC's supported for 802.1x MAC authentication | 3000 | 3000 | 3000 | 3000 |

* Parameters mentioned are applicable on specific HW profiles. Please check the Network OS documentation for the specific HW profiles.

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

IP Fabric Scalability:

| IP Fabric Scalability Numbers | VDX- 8770 | VDX-6940 | | VDX 6940-144s | | VDX 6740, VDX 6740T |
|--|--------------|----------|------|---------------|------|------------------------|
| | Spine | Spine | Leaf | Spine | Leaf | Leaf |
| VLANS extended with VxLANs (no. of tunnels * VLANS * ECMP) | NA | NA | 16k | NA | 16k | 16k |
| Software MAC entries (CML) | NA | 200k | 200k | 200k | 200k | 200k |
| Software ARP entries (Conversational ARP) | NA | 100k | 100k | 100k | 100k | 100k |
| Software ND entries (Conversational-ND) | NA | 50k | 50k | 50k | 50k | 50k |
| BGP eVPN IPv4 routes | 200k | 200k | 200k | 200k | 200k | 200k |
| BGP eVPN IPv6 routes | 64k | 2k | 2k | 2k | 2k | 2k |
| BGP eVPN MAC-IP routes | 100k | 100k | 100k | 100k | 100k | 100k |
| BGP eVPN MAC routes | 200k | 200k | 200k | 200k | 200k | 200k |
| Max # of IP Unnumbered interface | 384 | 36 | 36 | 144 | 144 | 52 |
| Max # of IP Port channel interface | 384 | 36 | 36 | 144 | 144 | 52 |
| Max # of members per IP Port-Channel | 8 | 8 | 8 | 8 | 8 | 8 |

| | | | | | | |
|--------------------------------------|----|----|----|----|----|----|
| Interface | | | | | | |
| Max # of Leaf – Spine ECMP | 16 | 16 | 16 | 16 | 16 | 16 |
| Max # of SAG addresses per interface | 64 | 64 | 64 | 64 | 64 | 64 |

HW Profile and Platform Specific Scale Numbers

Route Profile Scale:

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

| VDX 6940-36Q, VDX 6940-144S | | | | | | |
|--|---------------|----------------|--------------|-------------|----------------|-------------|
| Network OS v7.x Scalability Numbers | ROUTE PROFILE | | | | | |
| | DEFAULT | IPV4-MAX-ROUTE | IPV4-MAX-ARP | IPV4-MIN-V6 | IPV6-MAX-ROUTE | IPV6-MAX-ND |
| Maximum # of IPv4 routes with ECMP supported * | 6000 | 10000 | 10000 | 8000 | 2500 | 2500 |
| Maximum # of IPv6 routes with ECMP | 1000 | 0 | 0 | 500 | 2000 | 2000 |

| | | | | | | |
|---|-------|-------|-------|-------|-------|-------|
| supported * | | | | | | |
| Maximum # of OSPF routes * | 6000 | 10000 | 10000 | 8000 | 2500 | 2500 |
| Maximum # of OSPFv3 routes * | 1000 | 0 | 0 | 500 | 2000 | 2000 |
| Maximum # of IPv4 BGP routes in HW * | 6000 | 10000 | 10000 | 8000 | 2500 | 2500 |
| Maximum # of IPv6 BGP routes in HW * | 1000 | 0 | 0 | 500 | 2000 | 2000 |
| Maximum # of IPv4 routes in SW * | 6000 | 10000 | 10000 | 8000 | 2500 | 2500 |
| Maximum # of IPv6 routes in SW * | 1000 | 0 | 0 | 500 | 2000 | 2000 |
| Maximum # of ARP Entries * | 43000 | 49000 | 73000 | 49000 | 6000 | 6000 |
| Maximum # of IPv6 neighbor cache Entries * | 12000 | 0 | 0 | 10000 | 30000 | 30000 |

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

L2 L3 Multicast Scale:

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

| TCAM PROFILE IPV4-IPV6-MCAST | | | | |
|---|---------|----------|--------------|---------------|
| Network OS v7.x Scalability Numbers | VDX6740 | VDX-8770 | VDX-6940-36Q | VDX-6940-144S |
| Maximum # of L2 (IGMPv2 Snooping) multicast groups | 1000 | 6000 | 6000 | 6000 |
| Maximum # of MLD Groups | 512 | 4000 | 4000 | 4000 |
| Maximum # of L2 (IGMPv3 Snooping) multicast groups | 4000 | 4000 | 4000 | 4000 |
| # of L3 (S,G) forwarding Entries | 2,000 | 2,000 | 2000 | 2000 |
| # of L3 (*,G) forwarding Entries | 256 | 256 | 256 | 256 |

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

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

ACL Scale:

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

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

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

Compatibility and Interoperability

The following tables list the devices tested for IP storage, FC/FCoE storage and host adapters for VDX as of Network OS v7.2.0. This is a representative list of devices, Network OS v7.2.0 supports all standards-based devices connected to it for these types of storage.

IP Storage

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

FC/FCoE Storage

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

| Vendor | Storage Array Model | Protocol | Switch Model | Initiator |
|--------|---------------------|----------|------------------------|---|
| | XIV | FC | 6740/T | Windows 2012 R2 |
| HP | MSA2040 | FC | 6740/T | RHEL 7.0, Windows 2012, Windows 2012 R2 |
| | P10000 | FC | 6740/T | RHEL 7.0, Windows 2012, Windows 2012 R2 |
| | P6500 | FC | 6740/T | RHEL 7.0, Windows 2012, Windows 2012 R2 |
| | P6300 | FC, FCoE | 6740, 8770 (FCoE only) | RHEL 7.0, Windows 2012, Windows 2012 R2 |
| | P4330 | FC | 6740/T | RHEL 7.0, Windows 2012, Windows 2012 R2 |
| | P9500 | FC, FCoE | 6740/T | RHEL 7.0, Windows 2012, Windows 2012 R2 |

Network Adapters

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

ADDITIONAL CONSIDERATIONS

Limitations and Restrictions

Command Line Interface

- Break command is not supported. ctrl-c can be used as an alternative.
- Few commands may not display paginated output.
- For few clear and show commands “?” will not show all options for VRF. Tab completion will give all possible values.
- For certain commands (including “no” form with some commands), “?” shows unsupported additional options.
- Some CLI commands will generate an “Error:Access denied” message upon failure. This means the operation failed on the switch and may not be related to permissions.
- Tab completion and <ctrl>-c (cancel) does not work for some commands.
- Incorrect range might be displayed in the help text for some of the show commands.
- Range support is available for all the interfaces in Network OS v7.1.0. Following limitations are applicable:
 - Interface range command is supported on breakout ports of same connector. Range is not supported involving breakout ports of multiple connectors.
 - Interface range command does not support mix of regular ports and breakout ports.
 - Range command is not supported across multiple slots of the chassis.
 - Range command for rbridge-id is not supported.
 - In some instances, there could be a delay in starting of operation specified in the range command after being issued.
 - When range issued for very large subset (e.g 4k VLAN, 2k port-channels, etc.), timeout can occur or user may temporarily see switch being unresponsive or with high CPU utilization. Extreme recommends using range in smaller chunks. Especially, while configuring VLANs/VEs and Port-channels, Extreme recommends range to be less than 500.
 - Range prompt doesn't get updated when few or all of interface in that range are deleted. Therefore, user should exit from Range submode if few or all interfaces are deleted that are part of that range. New configuration performed on same range submode may give unpredictable results.
 - On a large VCS cluster, configurations performed on Range of physical interfaces and port-channels may spike high memory usage.
- System does not warn user on deleting the IP config when VRF is configured.
- If “switchport trunk allowed vlan all” is already configured on any interface, then VLAN creation using range command will be slow as each VLAN will get provisioned individually.
- Some unsupported debug commands may be seen in Network OS v7.1.0. Extreme recommends not to run them on switches:
 - Show confd-state –, for debugging purpose only.
 - Show parser dump –, for debugging purpose only.
 - Show notification stream –, for debugging purpose only.

- Autoupgrade command in config mode
- During “copy running-config startup-config” or “copy support” user might see occasional and temporary CPU spikes (up to ~30-40%).
- show mac-address-table command on console with include option can not be aborted with a break/ctrl-C. Use a telnet session for the same.
- Short form of MAC-Address is not supported as filter in “show running-config”.
- For IP access lists, display filtering based on sequence number alone does not work as expected.
- Certain oscmd commands may not work or give a different output under admin login
- If an alias exactly matches a partial keyword anywhere in the command line, pressing the TAB key for CLI command completion will claim that the input is invalid, and pressing the ENTER key will first replace the partial keyword with the alias expansion string. To avoid this, make sure that any partial keywords are not an exact match for an alias name.
- The authentication mode with primary & secondary sources of authentication cannot be updated to a configuration containing only the primary source. For example, the authentication mode cannot be changed from “radius local or radius local-auth-fallback” to ‘radius’. The workaround is to remove the existing configuration and then configure it to the required configuration.
- The “logging syslog server” command returns an error on the “secure” keyword. Use “secure port” to assign a nondefault port number.
- OSPFv3 on default VRF can be created without mentioning VRF name but while removing default VRF 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, enabling FibreChannel ports requires only the FCoE license to be installed and does not require any Port Upgrade license. The Port Upgrade license only controls Ethernet ports (number of ports or speed supported).
- An Integrated Routing license is NOT required on FOS-based SAN platforms running FOS 7.0.1 or above for FCR interoperability connectivity with VCS fabrics and the VDX6740x. Please refer to the FOS v7.0.1 Admin Guide documentation on configuring FOS platforms for connectivity to VDX 674x switches and VCS fabrics.
- The Layer 3 license is required on VDX8770 switches to enable Layer 3 feature set including OSPF, VRRP, BGP, VRF etc. A separate Layer 3 license is not required on VDX fixed-form factor switches as Layer 3 features are included in the default license.

VCS

- The Advanced Services License provides a single upgrade option to enable Fibre Channel over Ethernet (FCoE) and Layer 3 features on VDX8770 switches.
- Loopback connection is not supported in VCS mode. If a loopback connection is done (either using loopback plugs or port to port connections on the same switch), those interfaces become ISL interfaces.
- A node with default configuration will not join a cluster if the intermediate nodes between the node being defaulted and rest of the cluster are also undergoing reload. If the node boots up earlier than the intermediate nodes, it will form its own VCS and not join the parent cluster. In such situations, reload the node that is required to join the cluster.
- Logical Chassis Cluster Mode:
 - When a new switch is added to an existing VCS Fabric and if the new switch takes the role of principal node, the other switches in the fabric will receive the configuration of the distributed features such as Virtual IP and VM-Aware Network Automation from the newly added switch. This will cause the existing distributed configuration to be overwritten by the newly added switch in the principal role. This can be avoided by following the new switch addition procedures in *the Network OS Management Configuration Guide*.
 - After a cluster reboot, Extreme recommends to do both “show fabric all” and “show vcs” to ensure that cluster is entirely formed without any issue. User might see that ‘show vcs’ takes an additional 2-3 minutes to show all participating switches. This is an existing behavior and doesn’t affect data path functionality in most cases.
- “show fabric isl” & “show fabric trunk” may show the interfaces in random order without sorting.
- The default-configuration behavior may be different depending on the default-configuration triggers.
- The snapshot restore feature in VCS should be used to restore the local configuration and not the global configurations.
- Usage of Rbridge-range option to configure Rbridge context specific configurations is not recommended.
- Fastboot option is not recommended as a preferred method of reloading the switch.
- VCS for Network OSv7.0.1:
Note the following results for the given actions.

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

VCS Node Replacement

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

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

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

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

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

```
dpod 1/0/2
    !

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

Logical Chassis

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

Extreme Trunks

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

Breakout Interfaces

- VDX 8770 supports only static breakout of 40G ports. It is required to power OFF and ON linecard for the 40G ports on it to be converted into 10G breakout ports and vice versa.

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

Dual-personality Ports

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

1G Mode

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

vLAG

- LAGs are created with default speed of 10G. Therefore Extreme recommends end user to set required speed manually based on member speed using "speed" command.
- When configuring LACP LAG between VDX and non-Extreme switches it is highly recommended to enable the vLAG ignore-split on the VDX. Ignore split option is enabled by default.
- The port-channel interface "load-balance" is not the same as "fabric port-channel <#> load-balance"

- The port-channel interface "**load-balance**" command configures load-balancing on the actual vLAG member links (effective on Rbridges directly participating in the vLAG).
- The "**fabric port-channel <#> load-balance**" configures load-balancing on Rbridges NOT participating in the vLAG, but connecting to neighboring vLAG participating Rbridges.

Virtual IP Address Support

- A separate gateway cannot be configured for Virtual IP address. Default gateway will be the same as the gateway address for the management port of this switch.
- For VCS Virtual IP address to work correctly, the management port's IPv4 or IPv6 address should be assigned, functional and both address should be in same subnet.
- There is no Virtual MAC address associated with the Virtual IP address, physical MAC will be used.
- Chassis Virtual-IP is only supported on the VDX 8770.

Security, Management ACLs, Authentication, Authorization

- Login authentication service (aaa authentication login cli):
 - With "local" option specified as secondary authentication service, local authentication will be tried only when the primary authentication service (TACACS+/RADIUS/LDAP) is either unreachable or not available.
 - Behavior of "local" option in pre-4.1.0 releases is changed to the "local-auth-fallback" option.
 - When login authentication configuration is modified, the user sessions are not logged out. All connected user sessions can be explicitly logged out using "clear sessions" CLI.
- ACLs are not supported for egress traffic flows on management interfaces.
- Configuring TACACS+ or RADIUS without a key is not supported. If no key is configured, the switch uses a default key of "sharedsecret". The use-vrf option should be used to enter any additional parameters such as retries, timeout or key.
- Same NTP server configuration with different vrf not supported.
- There is a possibility that locked user accounts will get unlocked after a reboot if the running-config (before reboot) is different from startup-config of user accounts.
- Encrypted text (taken from running-config of any user account password with encryption turned on) should not be used as input for clear-text password for the same user. This may result in login failure of the user subsequently.
- When the ACL is applied to a management interface, only the top 256 rules will be applied if the ACL contains more than 256 rules.
- It is advised to not to apply ACL with 12k rules to management interface.

- When more than 250 rules ACL's are configured (over supported scale), they may be partially installed & effective.
- Access to ONLY the following Active Directory (AD) servers is supported by Extreme LDAP client:
 - Windows 2000
 - Windows 2003
 - Windows 2008 AD
- IPv6 RA Guard feature is not supported on VDX 8770 although the CLIs are visible.

SPAN & RSPAN

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

MAC Learning Considerations in VCS

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

PVLAN

- Following PVLAN features are not supported:
 - IGMP on PVLANS but there is no error message displayed if operator configures IGMP snooping on PVLAN
 - ARP & Routing in PVLAN domain
 - Enabling Routing in Primary and Secondary Vlan.
 - CLI to enable Local Proxy ARP on primary VLAN.
 - IP Configuration on PVLANS
 - Ve Configuration on both Primary and Secondary Vlan

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

UDLD

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

STP/DiST

- VDX does not support tunneling non-standard BPDUs and thus IEEE BPDUs (0180:C200:0000) generated as tagged packets in STP/RSTP/MSTP modes may not be tunneled successfully across VCS Fabric. However, VDX supports tunneling standards’ based BPDUs such as untagged IEEE BPDUs and tagged or untagged PVST BPDUs (0100:0CCC:CCCD). Post 3.0.1, the tagged IEEE BPDU can be tunneled across VCS fabric using command: “tunnel tagged-ieee-bpdu” under interface configuration.
- By default global spanning-tree and interface level spanning-tree will be disabled, user has to explicitly enable on the desired ports. VLAN spanning-tree state is default enabled.
- BPDU tunnel configurations are permitted only when spanning-tree is disabled in VCS.
- For Cisco proprietary Per Vlan Spanning Tree protocols (PVST and RPVST) user needs to configure Extreme switch to send BPDU on Cisco multicast destination mac address “0100.0ccc.cccd” for non-native VLANs. By default, Network OS 6.0.1 software uses Extreme "0304.0800.0700" multicast mac to send BPDU's on non-native VLANs.

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

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

```

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

```

```
VDX 6740-VCS1(config-Port-channel-100)# spanning-tree bpdu-mac 0100.0ccc.cccd
VDX 6740-VCS1(config-Port-channel-100)# exit
VDX 6740-VCS1(config)#
```

- STP Interop with certain vendor switches

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

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

IGMPv3 Snooping

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

Edge Loop Detection (ELD)

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

Long Distance ISL Ports

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

- The 40G and 100G interfaces do not support “long-distance isl” command, however can extend distances for non-lossless traffic up to 40Km using standard ISLs.
- On standard ISLs, the 10G, 40G and 100G interfaces support lossless traffic up to 1Km.
- The “long-distance-isl” command will not be supported on the SO-10GE-ZR-CX, 10G-SFPP-ZR, and 10G-SFPP-ZRD-T 80km optics.
- The SO-10GE-ZR-CX, 10G-SFPP-ZR, and 10G-SFPP-ZRD-T 80km optics requires a minimum distance of 20km in order to successfully form a standard ISL connection
- To form an ISL between 10G tunable ZR optics (57-1000266-01) when initially inserting the optic and configuring "tunable sfpp channel x", please configure any channel other than 1 on both ends.

AMPP and Port-Profiles

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

vCenter

- Receiving more than five vCenter events within a span of 30 seconds, results in asset discovery getting initiated. Post discovery cluster configuration will be in sync with vCenter.
- vCenter auto-profile is automatically added/deleted to the default port-profile-domain in Service VF enabled/disabled mode.

- Modifying/editing the auto port-profiles in the default-domain is not recommended, which may cause auto-pp application failure during vCenter operation and end up in traffic failure.
- Adding/removing the auto-port-profile to the user-created domain when Service VF is enabled is not recommended which may cause auto-pp application failure during vCenter operation and end up in traffic failure.
- vCenter auto-profile does not support SERVICE VF classification.
- Output of show vnetwork vss displays the vmnic against the vSwitch even after the removal of the vnic from the vSwitch through vCenter. Recovery happens in the next auto-discovery cycle.

QoS

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

FCoE

- On switches shipped with Network OS 7.2.0, the default mode of operation is Access Gateway for the VDX 6740, 6740T, 6740T-1G,.
- Adding/removing an ISL or adding/removing a VDX switch from cluster may cause that exiting VDX switch in the cluster stops generating a FCoE keep alive. And this results in FCoE port continuously flapping. To recover, re-configure FCoE by removing and adding fcoeport configuration (no fcoeport/fcoeport default) on the affected interfaces. As an alternate recovery, shut / no shut on the affected interfaces.
- Logical SANs have been supported since Network OS v6.0.0. If user needs to enable Fibre Channel Forwarder (FCF) mode, switch needs to be configured in FCF mode. If the switch is upgraded from a lower Network OS version (pre 4.1.2 builds) to v7.0.1, it will be in FCF mode.
- Extreme recommends not having FCoE ports and Long Distance ISL ports in the same port-groups. This configuration will NOT be prevented by the CLI; however it can result in unpredictable behavior for FCoE traffic.
- If the FCoE FCMAP is changed from default to non-default, user might see that FCoE login may not happen. Please toggle the interface using “shutdown” followed by “no shutdown” to work this around.
- When an RBridge is removed from VCS cluster, it does not need to be manually removed from fcoe fabric-map.
- MAC binding for remote SANs is not honored during config replay.
- In case the FIF is multiple hops away from the FCF, it is recommended that the interfaces of the intermediate switch shouldn't be configured with the same remote san as that of the FIF/FCF
- VLAN's which are reserved for FCoE may not be used for any other purpose. This is true for both Fabric Cluster and Logical Chassis modes.
- Extreme recommends that for all LAGs with FSB, the fcoeport config must be applied on the LAG itself. And for all LAGs with directly attached CNAs, the fcoeport config must be applied on the member ports.
- Binding an enode mac to FCoE interface is not allowed in range context, as only one enode mac can be bound to one FCoE interface.
- While providing range for FCoE interfaces, it's recommended to provide the range only in ascending order. For ex: interface fcoe 1/48/11-38 is recommended, interface fcoe 1/48/38-11 is not recommended.
- FCoE traffic may not be mirrored using RSPAN. Workaround is to use SPAN.
- In use cases with FSB, it is noticed that after converting dynamic port-channel to static, hosts and targets don't see each other.
- In NETWORK OS V6.0.1 and later, up to four FCoE Vlans are supported in VDX . But on a single VDX , All member ports in a LAG have to be configured with the same FCoE Vlan. Different LAG can be configured with different FCoE Vlan.
- In NETWORK OS V6.0.1 and later, it is recommended user define different fabric-map for Remote Logical SAN and Local Logical SAN configuration. If user configures a fabric-map to work on Remote Logical SAN first and then later change the same fabric-map to become Local Logical SAN, it may cause FCoE port continuously flapping.

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

FlexPorts

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

Fibre Channel

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

Access Gateway

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

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

IP Fabric

Provisioning :

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

BGP eVPN :

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

- Tunnel creation is triggered by BGP NH installation resulting in creating more tunnels than configured which might be seen at the Border Leaf.

ARP/ND Suppression:

- ARP/ND suppression should be configured on the VLAN if IPv4/IPv6 Static Anycast Gateway is being configured.
- Upto 512 VLANs are supported with DAI hardware profile. Default hardware 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.

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, Static Anycast Gateway to be used to achieve similar functionality.

OSPFv3

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

BGP

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

Layer 2/Layer 3 Multicast

- The following PIM features are not supported in this release:
 - IP version 6
 - VRF

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

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

VRF

- Under VRF submode there is a syntax change for the address-family ipv4 command.
Old format: address-family ipv4 [max-route <value>]
New format:
address-family ipv4 unicast
max-route <value>

Note: "max-route" command is now moved to address-family submode.

- There is no provision to configure “max-routes” for default-vrf.
- There is no use case for “rd” configuration in VRF and this command will be deprecated in next release.
- On configuring VRF on an interface, all previous IP config on that interface will be deleted.
- Removing VRF address family on a non-default VRF will delete all relevant address-family configurations including the interface and protocol configuration for that VRF.

BGP-VRF

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

ACL

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

Policy-based Routing (PBR)

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

Inter-VRF Leaking (Static)

- S+ symbol routes indicates leaked routes.
- VRF route leak cascading is not supported— only one level of indirection.

- User should avoid making Static, dynamic and connected route conflict with routes in target VRF when configuring route leak.
- For bidirectional traffic with router leak, user needs to configure route leak in both direction separately.
- Route leak configuration to next hop IP on the same box on different VRF is not a valid configuration, but CLI will be accepted.
- Precaution needs to be taken when leaking default routes - this can result in routing loops.
- Switch management from non-management VRF by leaking route from non-management to management VRF is not supported.

DHCP IP Helper

- There is no HA support for DHCP relay statistics. When a switchover happens, the statistics will not be replicated to the new active MM.
- Clients may not converge in some IP Fabric environment. Care should be taken to not configure DHCP IP helper and Static Anycast Gateway on the same interface.
- Two DHCP OFFER per one DHCP DISCOVER and two DHCP ACK for single DHCP request seen IP fabric setup.
- DHCP relay doesn't work correctly with just Fabric Virtual Gateway (FVG) on the same VE interface. The workaround is to configure unique IP addresses on VE interfaces simultaneously.

Dynamic ARP Inspection (DAI)

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

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

- In order for successful version upgrade using DAD method, switch should undergo 2 reloads. For switch in factory default, there is additional reboot to cancel bare metal mode.
- If firmware download is skipped only config download is allowed.
- For dual MM chassis, dual MM must be in sync for DAD to function.
- DAD is a disruptive. ISSU is not supported.
- In FIPS mode, DAD is not supported.
- Cluster principal node failover is not supported.
- DAD over in-band is not supported. Virtual Fabrics is not supported with DAD. You must disable Virtual Fabrics before starting the DAD process in the global configuration file or in the script.
- DAD must complete and succeed on Principal node alone before turn on power for all secondary nodes.
- When the switch is in Factory default, DAD is enabled upon power up the switch
- DAD executes only if the switch configuration is the default configuration. If the configuration on the switch is not the default configuration, DAD exits.

- If the switch is in the default configuration before DAD is triggered, DHCP will remain enabled after the deployment completes. However, this setting can be overwritten by the switch-specific configuration file or the dad.py script.
- You must enable DHCP in the DCMD default configuration to ensure that the switch receives its IP address from the preconfigured DHCP server.
- The factory default DAD runs only once in a DHCP-enabled environment. Irrespective of whether this process is a success or failure, DAD will not be triggered again after a reboot or power off. You can run DAD manually using the dhcp auto-deployment enable command if required.
- Must set ztp=0 in dad configuration file since ZTP (Zero Touch Provisioning) is enabled by default.
- The "vcsmode" value in dad.conf MUST be set to "LC" regardless of whether the existing cluster is in LC or FC mode. If "vcsmode" set to "FC" value in dad.conf, the DAD request can fail.
- DAD is enabled automatically upon switch reboot when you use Network OSCLI "write erase" command.

Zero Touch Provisioning (ZTP) consideration

DAD supports up to two nodes for IP fabric in logical chassis mode

All nodes can either be powered up at the same time or enabled from the CLI. This is the key difference vs regular DAD process.

Link State Tracking

- The "track enable/disable" command can only be used to enable or disable the tracking. In order to remove tracking configuration from internal database for a given interface "track remove all" command must be used.
- When there are no uplink interfaces configured, the track disable command will remove tracking configuration from internal database and this behavior is applicable only in 6.0.1a patch and not in prior releases.
- If track min-link number is greater than the number of uplinks, then the downlink will be shutdown with a warning message.
- After toggling the line card using "power-off / on", LC related interfaces that are configured as uplink interfaces are not seen in "show track summary" cli output.

OpenFlow

- Interoperability support only with Extreme Controller aka. BVC/BSC.
- Once an interface becomes OpenFlow enabled, very limited set of conventional commands are allowed which includes some of the QoS related configuration commands. For complete list of allowed commands please refer to "*NETWORK OS V6.0.1 SDN Configuration Guide*"
- Priority-tagged frames are not supported.
- L3 Generic flows (incoming port as "*wildcard*") are not supported.
- PUSH/POP operations can only be associated with action type OFPAT_OUTPUT inside a flow-mod.
- Type of an existing GROUP cannot be changed.

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

Mac Port Based Authentication

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

Uplink Switch Support

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

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

configured should be same on protected and uplink ports.

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

Layer 2 and Layer 3 ISSU on VDX 6740x

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

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

The implementation is based on a type-1 hypervisor.

Extreme Vyatta Controller (BVC)

- Controller does not update the config database based on the flow rejected notification/group rejected notification/meter rejected notification/delete notification/hard

timeout aging notification from switch. Workaround : User needs to delete the flow from the config database and program the correct flow.

- In rare scenario, Controller sends the asynchronous messages leading to flow rejections e.g. flow-mods (associated with group/meter) are rejected after reconnection due to flow-mods being programmed before group/meter config. Work around is that the user needs to delete the group/meter/flow from the config database and program them again.
- In scale scenario, few flow-mods are not programmed after reconnection. Work around is that the user needs to delete the missing flow-mods and program them again.
- Topology/Change of interface states are not reflected correctly on BVC.
- Topology with multiple links are not reflected on BVC. BVC shows only single link between the switches.
- Refer to BVC 1.3.0 release noted for all the known issues/workaround.
- Limitations while configuring flows using BVC:

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

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

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

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

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

REST API

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

NetConf

- Netconf configuration for startup-config datastore is not supported

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

VXLAN Gateway for VMware NSX

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

VF Extension using VxLAN

- VF Extension overlay-gateway (VTEP) is supported only on the VDX 6740, VDX 6740T, VDX 6740T-1G, VDX 6940-36Q, and VDX 6940-144S.
- VF Extension overlay-gateway is supported only in the VCS Logical Chassis mode.
- VDX 8770 can be in the same VCS fabric where VF-Extension functionality is enabled.
- VxLAN Tunnels are supported over ISL links.
- VF Extension overlay-gateway can be enabled on maximum 4 Rbridges in a VCS Cluster.
- VxLAN GW for VMware NSX and VF Extension cannot be enabled in the same VCS fabric.
- Only 1 VF Extension overlay-gateway is supported in a VCS Cluster.
- Only one-to-one VLAN to VNI mapping is supported.

- Tunnel interfaces cannot be used as SPAN (Switch Port ANalyzer) destination.
- Only Ingress ACLs can be applied to tunnels.
- Ingress/Egress QoS policies cannot be applied to tunnels.
- Multicast routing between VXLAN and VLAN/VXLAN is not supported.
- L3 routing protocols and static routes over tunnels are not supported.
- Connected subnet L3 forwarding is supported over tunnels.
- Tunnels egressing/ingressing through an ISL port is supported only on VDX 6940 as a VTEP beginning with Network OS Network OS v6.0.1. Such topologies and configuration must be removed before downgrading to any version below Network OS 6.0.1.
- Fabric-Virtual-Gateway (FVG) based VTEP is not supported. CLIs for configuring FVG as VTEP are available under overlay-gateway, however these CLIs should not be used as the functionality is not available in this release.

TCAM Profiles

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

Management VRF

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

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

Conversational MAC Learning

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

System level Flow based QoS

- System Flow based QoS is not supported on the Egress direction.
- QoS can operate on either of three modes – MLS, CEE and MQC. Hence once service-policy is configured, the interface will be assumed to be in MQC mode and existing MLS and CEE commands will not be supported on the same interface. Un-configuring the policy will put the interface back to default mode which is MLS.
- For Policer, aggregation is possible only within a single chip. Hence when policer is applied on port- channel interface, multi-chip aggregation is not expected.
- SFLOW as action is not supported on Port-Channel interface.

- Any ACL that is used in Flowbased QoS class-map as a match criterion is considered as “QoS ACL” and is special in nature. Hence behavior in some aspects may differ from that of regular “User ACL”.
- System based QoS is not supported in egress direction.

Port level Flowbased QoS

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

URPF

- uRPF is not supported in Mercury

BGP Auto neighbor discovery

- BGP Auto Neighbour 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 (> 1 sec)
- HA is not supported for OpenFlow feature, however, system level ISSU is supported. For ISSU, it is recommended that the controller is disconnected first, all flows are cleared using “clear OpenFlow all” command and then perform the upgrade.

Interoperability

- In a VPC environment where the Extreme VDX side has the active LACP settings and the Cisco side has the passive settings on the vLAG, the port-channel takes over 30 seconds to come up.
Workaround: Reverse the settings and have the Extreme VDX LACP settings passive and the Cisco side set as active. The port channel will then restore after about 10 seconds.
- VDX interop with Cisco Nexus switch with ‘peer-switch’ enabled on VPC is not supported.
- When interoperating with Extreme 8000, it is recommended to set the *mac-aging* time to 0 on the VDX switch to prevent any adverse impact caused by certain errors generated by the Extreme 8000.
- ADX HA Sync packets use UDLD PDU’s which may be dropped by VDX . To enable forwarding, we recommend configuring dot1q tagging to treat UDLD packets as data-packets to be forwarded across VCS.Virtual Fabric.
- PIM-SM is not supported on Virtual Fabric on VDX8770.
- For frames forwarded on a transport fabric, ingress CTAG tagging is preserved at the egress port regardless of the egress tagging classification.
- Default-VLAN can only be configured using TRANSPORT VF IDs.
- The “no vcs virtual-fabric enable” command execution time is dependent on the number of ISLs and VLANs in the VCS.
- To allow for STP interop with certain 3rd party switches that would not accept the BPDU source MACs with default OUI. The selection can be changed using the 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

- BNA version 14.2.1 is supported with Network OS Network OS 7.2.0.
- MAPS is supported on VDX 6740, 6940 and 8770 platforms.

- MAPS port level alerting in Network OS V7.0.x is not available for Flex Ports configured in Fiber Channel mode.
- RX_SYM_ERR MAPS messages are displayed when breakout cable is connected on a 40G interface that is not configured for breakout.
- When line card on the remote end of the link is powered off, MAPS generates Insertion/Removal notification for the SFPs on the local side. These can be ignored.
- 100G SFP threshold monitoring is not supported on VDX6940-144s.

Maintenance Mode

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

BNA

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

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

Miscellaneous

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

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

Defects

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

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

Network OS v7.2.0a Caveats

None

| |
|-----------------------------------|
| Defect ID: DEFECT000660024 |
|-----------------------------------|

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

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

Network OS v7.2.0 Caveats

BFD

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

VxLAN

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

Note:

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

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

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

- On sending IGMP queries over VF_Extension Tunnel with VLAG as underlay, packets might loop over the tunnel .Queries come back from the same tunnel interface from which its egressed out.
- “show ip igmp groups interface tunnel <tunnel_id>” cli shows all IGMP interfaces instead of just the tunnel interface.
- Adding and Removing RBridges under overlay-gateway may take longer than expected time if large number of VLANs are configured in the fabric.

Long Distance ISL

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

Loopback interfaces

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

Route distribution

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

FCoE/Access Gateway

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

BNA/NetConf/REST

- Special character ‘\$’ under the custom RPC “bna-config-cmd” cannot be used for Netconf and REST API for performing copy operation.
- REST API deletion on the main resource will remove all the sub-resources under it. For Example, REST API delete Operation without specifying ACL name will remove all the ACLs in the system. Specify the ACL name in the request in order to delete particular ACL from the config.

- For large scale VCS fabrics with more than 4000 ports, querying the cluster with BNA/REST APIs may result in switch software exception. For this purpose it is not recommended to enable BNA monitoring or querying with REST APIs for large VCS fabrics.

AAA Configuration

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

Sync Failure Error

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

Mac Loop Detect Feature:

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

Port Channel Scalability:

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

AMPP/vCenter:

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

OpenFlow:

- With default rcv-queue and after coldboot group select traffic may not be correct, need to do shut/no shut on the interface. This is not observed with non-default rcv-queue.
- With large number of flows, “show openflow flow <>” may take 20 seconds to display packet counts
- Filtering options (e.g. show | include) will not work for show openflow commands. show commands will display the complete output.
- "Module Unknown" is shown for CLI "show open resources" in VDX 6940-144S.

- Hardware Profile:

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

- Copy Config command:

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

Syslog:

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

Closed with code changes for Network OS v7.2.0c

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

| | | | |
|-----------------------------|--|--------------------------|--------------------|
| Parent Defect ID: | NOS-55130 | Issue ID: | NOS-55130 |
| Severity: | P3 - Medium | | |
| Product: | Network OS | Technology Group: | Data Center Fabric |
| Reported in Release: | NOS6.0.2b | Technology: | Logical Chassis |
| Symptom: | Down-level firmware of the internal eUSB device in some VDX6940s may cause the eUSB to fail, causing the system to fail. | | |
| Condition: | Resolved with code. | | |

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

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

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

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

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

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

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

| | | | |
|-----------------------------|---|--------------------------|---|
| Parent Defect ID: | NOS-54704 | Issue ID: | NOS-66957 |
| Severity: | S3 - Medium | | |
| Product: | Network OS | Technology Group: | Management |
| Reported in Release: | NOS7.1.0b | Technology: | SNMP - Simple Network Management Protocol |
| Symptom: | Unexpected reload | | |
| Condition: | when two tlvs are received but tlv type of one is either zero or unsupported. | | |
| Workaround: | None | | |

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

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

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

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

| | | | |
|--------------------------|------------|--------------------------|------------|
| Parent Defect ID: | NOS-66850 | Issue ID: | NOS-66987 |
| Severity: | S4 - Low | | |
| Product: | Network OS | Technology Group: | Management |

| | | | |
|-----------------------------|--|--------------------|---|
| Reported in Release: | NOS7.3.0a | Technology: | SNMP - Simple Network Management Protocol |
| Symptom: | snmpget of snIfOpticalLaneMonitoringTable returns "No Such Instance" errors. | | |
| Condition: | snIfOpticalLaneMonitoringTable is not supported for 1G and 10G interfaces and it is applicable only for 40G/100G optics. | | |
| Workaround: | None | | |

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

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

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

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

| | | | |
|--------------------------|-----------|------------------|-----------|
| Parent Defect ID: | NOS-66864 | Issue ID: | NOS-67034 |
| Severity: | S2 - High | | |

| | | | |
|-----------------------------|--|--------------------------|-------|
| Product: | Network OS | Technology Group: | Other |
| Reported in Release: | NOS7.4.0 | Technology: | Other |
| Symptom: | NOS version 7.0.2b, 7.2.0b and 7.4.0 will not allow the LC48X10GT to boot up due to a missing file | | |
| Condition: | LC48X10GT will not boot up online | | |

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

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

| | | | |
|-----------------------------|--|--------------------------|--------------------|
| Parent Defect ID: | NOS-66998 | Issue ID: | NOS-67041 |
| Severity: | S3 - Medium | | |
| Product: | Network OS | Technology Group: | Data Center Fabric |
| Reported in Release: | NOS7.0.0 | Technology: | Logical Chassis |
| Symptom: | No reason at RASLOG when interface administratively down with out user config. | | |
| Condition: | When exception created on port due to ASIC fault. | | |

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

| | | | |
|-----------------------------|-------------|--------------------------|--------------------------------------|
| Parent Defect ID: | NOS-66898 | Issue ID: | NOS-67048 |
| Severity: | S3 - Medium | | |
| Product: | Network OS | Technology Group: | Layer 3 Routing/Network Layer |
| Reported in Release: | NOS7.1.0a | Technology: | BGP4+ - IPv6 Border Gateway Protocol |

| | |
|-------------------|--|
| Symptom: | Unexpected reload. |
| Condition: | When we run the "show ip bgp neighbors" after giving the "no neighbor activate" followed by "clear ip bgp neighbour" |

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

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

| | | | |
|-----------------------------|---|--------------------------|------------------------------|
| Parent Defect ID: | NOS-66324 | Issue ID: | NOS-67106 |
| Severity: | S2 - High | | |
| Product: | Network OS | Technology Group: | Layer 2 Switching |
| Reported in Release: | NOS6.0.2 | Technology: | LAG - Link Aggregation Group |
| Symptom: | Port channel will be operationally down | | |
| Condition: | This is rare scenario and can happen only when we have the ASIC resources are full. | | |

Closed with code changes for Network OS v7.2.0b

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

| | | | |
|-----------------------------|--|--------------------------|-----------------|
| Defect ID: | DEFECT000609232 | | |
| Technical Severity: | High | Probability: | High |
| Product: | Extreme Network OS | Technology Group: | Logical Chassis |
| Reported In Release: | NOS7.1.0 | Technology: | VCS |
| Symptom: | 1. Compact Flash grows and disk full issue can occur. 2. Unexpected DCMd daemon termination. | | |
| Condition: | No special condition or configuration required to hit this issue. | | |
| Workaround: | For syslog.log file growing beyond limit below is preventive workaround: Please comment below two lines from syslog config & template files | | |

| | |
|------------------|--|
| | and reread config file from root using below command [killall]. It should stop all those IO messages. Please also verify that syslog utility is working fine after the workaround applied to make sure all is fine before we try on customer production environment. /etc/syslog-ng/syslog-ng.conf /etc/syslog-ng/syslog-ng.template #destination df_kern { pipe("/var/log/kmsg"); }; #log { source(s_all); filter(f_kern); destination(df_kern); }; root > /usr/bin/killall -HUP syslog-ng |
| Recovery: | Empty /var/log/syslog.log file if it is growing beyond 1 Mb. |

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

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

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

| | | | |
|----------------------------|--------------------|--------------------------|-----------------|
| Defect ID: | DEFECT000647840 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | Logical Chassis |

| | | | |
|-----------------------------|---|--------------------|-----|
| Reported In Release: | NOS6.0.2 | Technology: | VCS |
| Symptom: | System may undergo unexpected reload | | |
| Condition: | Media removal while media data is reading | | |
| Workaround: | shut/ no shut media removed interface | | |

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

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

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

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

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

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

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

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

| | | | |
|-----------------------------|--------------------|--------------------------|--------------------------------------|
| Defect ID: | DEFECT000658079 | | |
| Technical Severity: | Low | Probability: | Medium |
| Product: | Extreme Network OS | Technology Group: | OSPF - IPv4 Open Shortest Path First |
| Reported In Release: | NOS7.0.1 | Technology: | Layer 3 |

| | | | |
|-------------------|---|--|-----------------------|
| | | | Routing/Network Layer |
| Symptom: | Route summarization does not happen even after configuring it on the device. | | |
| Condition: | This issue is seen when configured route summarization prefix triggers OSPF Appendix E calculation with the existing Type 3 LSAs. | | |

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

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

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

| | | | |
|----------------------------|--------------------|--------------------------|-----------------|
| Defect ID: | DEFECT000659781 | | |
| Technical Severity: | Medium | Probability: | High |
| Product: | Extreme Network OS | Technology Group: | Logical Chassis |

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

| | | | |
|-----------------------------|--|--------------------------|--------------------|
| Defect ID: | DEFECT000659860 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | Logical Chassis |
| Reported In Release: | NOS7.0.1 | Technology: | Data Center Fabric |
| Symptom: | Unexpected reload. | | |
| Condition: | This occurs when a physical port is added to a port-channel after an ISSU upgrade was performed (Before VDX reloaded since upgrade). | | |

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

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

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

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

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

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

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

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

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

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

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

| | | | |
|-----------------------------|--|--------------------------|--------------------|
| Defect ID: | DEFECT000662379 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | Logical Chassis |
| Reported In Release: | NOS7.3.0 | Technology: | Data Center Fabric |
| Symptom: | VDX device will startup with the default configuration and also reports the same using below RAS log. [DCM-3053], SW/0 Active, INFO, VDX-VCS, Dcmd database corruption was detected. The system will startup with the default | | |

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

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

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

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

| | | | |
|-----------------------------|--------------------|--------------------------|--------------------|
| Defect ID: | DEFECT000663500 | | |
| Technical Severity: | Medium | Probability: | High |
| Product: | Extreme Network OS | Technology Group: | YANG |
| Reported In Release: | NOS7.0.1 | Technology: | Network Automation |

| | | | |
|-------------------|---|--|-------------------|
| | | | and Orchestration |
| Symptom: | NOS fails to un-escape special characters in passwords received via Netconf XML for config backup upload. | | |
| Condition: | Special characters used in password | | |

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

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

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

| | | | |
|-----------------------------|--|--------------------------|-------------------------------|
| Defect ID: | DEFECT000665913 | | |
| Technical Severity: | High | Probability: | High |
| Product: | Extreme Network OS | Technology Group: | Static Routing (IPv4) |
| Reported In Release: | NOS7.3.0 | Technology: | Layer 3 Routing/Network Layer |
| Symptom: | BGP route not cleared in secondary node of VCS cluster | | |
| Condition: | Remove ve interface | | |
| Recovery: | clear bgp neighbor | | |

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

| | | | |
|-----------------------------|--|--------------------------|--|
| Defect ID: | DEFECT000666699 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | ICMP - Internet Control Message Protocol |
| Reported In Release: | NOS7.4.0 | Technology: | Layer 3 Routing/Network Layer |
| Symptom: | VDX doesn;t generate ICMP notification for IP MTU violation [trapped packets]. | | |
| Condition: | IP MTU violation | | |

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

| | | | |
|-----------------------------|--|--------------------------|--|
| Defect ID: | DEFECT000668234 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | OSPFv3 - IPv6 Open Shortest Path First |
| Reported In Release: | NOS7.1.0 | Technology: | Layer 3 Routing/Network Layer |
| Symptom: | Device experienced sudden reload due to DCMd daemon termination. | | |
| Condition: | Execution of "no ipv6 ospf cost" CLI command. | | |

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

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

Closed with code changes for Network OS v7.2.0a

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

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

| | | | |
|----------------------------|--------------------|--------------------------|--|
| Defect ID: | DEFECT000579904 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | AAA - Authentication, Authorization, and |

| | | | |
|-----------------------------|--|--------------------|------------|
| | | | Accounting |
| Reported In Release: | NOS5.0.2 | Technology: | Security |
| Symptom: | Command set field on the Windows based TACACS server is empty | | |
| Condition: | 1. When TACACS server is windows based 2. Accounting is enabled | | |

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

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

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

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

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

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

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

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

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

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

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

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

| | | | |
|----------------------------|--------------------|--------------------------|---------------------|
| Defect ID: | DEFECT000646908 | | |
| Technical Severity: | Medium | Probability: | Medium |
| Product: | Extreme Network OS | Technology Group: | Hardware Monitoring |

| | | | |
|-----------------------------|---|--------------------|------------|
| Reported In Release: | NOS7.1.0 | Technology: | Monitoring |
| Symptom: | The source IP for SNMP traps is not deterministic. | | |
| Condition: | When VCS virtual IP address is configured and SNMP traps are enabled. | | |

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

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

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

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

| | | | |
|----------------------------|--------------------|--------------------------|--|
| Defect ID: | DEFECT000648164 | | |
| Technical Severity: | Medium | Probability: | Medium |
| Product: | Extreme Network OS | Technology Group: | SNMP - Simple Network Management |

| | | | |
|-----------------------------|---|--------------------|------------|
| | | | Protocol |
| Reported In Release: | NOS6.0.2 | Technology: | Management |
| Symptom: | SNMP responding on VCS IPv6 instead of management IPv6 address. | | |
| Condition: | Both MM/Chassis IPv6 and virtual vcs IPv6 addresses are configured. | | |
| Workaround: | Have the management IPv6 configured latest | | |

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

| | | | |
|-----------------------------|--|--------------------------|----------------------------|
| Defect ID: | DEFECT000648357 | | |
| Technical Severity: | Critical | Probability: | High |
| Product: | Extreme Network OS | Technology Group: | Configuration Fundamentals |
| Reported In Release: | NOS7.2.0 | Technology: | Management |
| Symptom: | REST POST/PUT/PATCH configuration fails and errors out | | |
| Condition: | Issue in REST POST/PUT/PATCH methods if payload has space. | | |
| Workaround: | DO not add space in payload | | |

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

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

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

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

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

| | | | |
|-----------------------------|--|--------------------------|---|
| Defect ID: | DEFECT000650040 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | RAS - Reliability, Availability, and Serviceability |
| Reported In Release: | NOS6.0.2 | Technology: | Monitoring |
| Symptom: | Suddenly edge port admin down without user/admin action. | | |
| Condition: | CRC error hits to threshold limits. | | |
| Recovery: | Amin no shut | | |

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

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

| | |
|-------------------|---------------------------------------|
| Condition: | Chassis IP and VCS IP not configured. |
|-------------------|---------------------------------------|

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

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

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

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

| | |
|--------------------|--------------------------------------|
| Workaround: | Avoid the CLI during cluster re-join |
|--------------------|--------------------------------------|

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

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

| | | | |
|-----------------------------|--|--------------------------|--------------------------------|
| Defect ID: | DEFECT000655415 | | |
| Technical Severity: | Medium | Probability: | Low |
| Product: | Extreme Network OS | Technology Group: | VXLAN - Virtual Extensible LAN |
| Reported In Release: | NOS7.0.1 | Technology: | Layer 2 Switching |
| Symptom: | PBR is applied to only some flows, when it's configured on Ve that terminated VxLAN. | | |
| Condition: | PBR configuration on Ve that terminated VxLAN. | | |

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

Closed with code changes for Network OS v7.2.0

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

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

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

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

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

| | |
|--------------------------------------|--|
| Defect ID: DEFECT000592902 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS6.0.2 | Technology: SNMP - Simple Network Management Protocol |

| |
|--|
| Symptom: False SNMP traps observed for link status of management interface. |
| Condition: Very rare scenario to hit this false status traps |
| Recovery: Restart SNMP process. |

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

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

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

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

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

| | |
|--|---|
| Defect ID: DEFECT000604338 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.0.0 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: After the VDX switch was reloaded, the configured MSTP hello time was ignored and uses default value of 2 seconds. The value configured in the system does not change, NOS simply ignores it and uses default value. | |
| Condition: The configured the MSTP hello time is not persistent after VDX switch was reloaded. | |
| Recovery: The configured MSTP hello-time is persistent after the VDX switch was reloaded. That is if MSTP hello-time was configured for example set to 5. The value was applied after the VDX was reloaded. | |

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

| |
|-----------------------------------|
| Defect ID: DEFECT000609676 |
|-----------------------------------|

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

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

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

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

| | |
|------------------------------------|--|
| Defect ID: DEFECT000612699 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |

| | |
|--|---|
| Reported In Release: NOS6.0.2 | Technology: LAG - Link Aggregation Group |
| Symptom: Unexpected reload of VDX | |
| Condition: In rare cases, deleting PO or a reload of an LC may cause the VDX to abruptly reload due to a software daemon termination. | |

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

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

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

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

| | |
|-----------------------------------|--------------------------|
| Defect ID: DEFECT000617887 | |
| Technical Severity: High | Probability: High |

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

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

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

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

'show running-config interface management'

| | |
|---|--------------------------------------|
| Defect ID: DEFECT000621402 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.0.1 | Technology: Inband Management |
| Symptom: Telnet access to VDX is blocked via default-VRF and user defined VRF. | |
| Condition: Firmware install with "no-activate" option. | |
| Workaround: Activate the firmware which was installed with "no-activate" option. | |

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

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

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

| | |
|-----------------------------------|--|
| Defect ID: DEFECT000623805 | |
|-----------------------------------|--|

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

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

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

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

| | |
|-----------------------------------|--|
| Defect ID: DEFECT000624872 | |
|-----------------------------------|--|

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

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

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

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

| | |
|--------------------------------------|--|
| Defect ID: DEFECT000625670 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: Software Installation & Upgrade |

| |
|---|
| Symptom: On rare occasions a SW error may be seen during HA synchronization. |
| Condition: Rare occurrence that is not linked to a specific change. |

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

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

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

| | |
|--------------------------------------|---|
| Defect ID: DEFECT000626555 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS7.1.0 | Technology: PIM - Protocol-Independent Multicast |

| |
|--|
| Symptom: Multicast functionality daemon "PIMd" goes down with memory leak |
| Condition: PIM enable configuration. |

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

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

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

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

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

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

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

| | |
|--|---|
| Defect ID: DEFECT000628238 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.2.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: VDX does not connect to NSX6.3.0 EA controllers. | |
| Condition: VDX to NSX6.3.0 EA controllers connection. | |

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

Condition: "RD auto" configuration through netconf.

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

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

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

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

| | |
|--------------------------------------|--|
| Defect ID: DEFECT000631386 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: SNMP - Simple Network Management Protocol |

| |
|--|
| Symptom: Unsupported snmp configuration like snmp v3 user exist in FIPS mode. |
| Condition: FIPS mode has unsupported snmp configuration |

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

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

| | |
|--|------------------------------------|
| Defect ID: DEFECT000631759 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.2.0 | Technology: Logical Chassis |
| Symptom: 100Mbps speed can be set on VDX6940-144S with CuSFP P/N 57-1000042-02, 57-1000042-01. | |
| Condition: Configuring 100Mbps speed on VDX6940-144S with CuSFP P/N 57-1000042-02, 57-1000042-01. | |

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

| | |
|-----------------------------------|--------------------------|
| Defect ID: DEFECT000632419 | |
| Technical Severity: Medium | Probability: High |

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

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

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

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

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

| |
|-----------------------------------|
| Defect ID: DEFECT000633831 |
|-----------------------------------|

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

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

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

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

| | |
|---|---|
| Defect ID: DEFECT000634192 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.0.1 | Technology: IP Fabric |
| Symptom: Tunnel terminated packets that generate the ARP's may keep looping across VCS nodes in the same VCS | |

Condition: VDX puts incorrect source information cause the issue

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

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

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

| | |
|---|---|
| Defect ID: DEFECT000634370 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: Configuration Fundamentals |
| Symptom: show running-config rbr route-map output is being sorted based on the action (permit or deny) and then the sequence number instead of just sequence number. | |
| Condition: Execution of "show running-config rbr route-map" CLI. | |

| | |
|------------------------------------|-------------------------------------|
| Defect ID: DEFECT000634372 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |

| | |
|--|---|
| Reported In Release: NOS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: route map is allowing to change the existing action without any warning/error. | |
| Condition: Modifying route map action. | |

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

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

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

| | |
|---|--|
| Defect ID: DEFECT000634904 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS6.0.2 | Technology: VLAN - Virtual LAN |
| Symptom: show vlan brief and show vlan XXX command outputs may show inconsistent data. | |
| Condition: Around 4000 vlans are associated with switchport | |
| Recovery: | |

| |
|-----------------------------------|
| Defect ID: DEFECT000635101 |
|-----------------------------------|

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

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

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

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

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

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

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

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

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

| | |
|---|--|
| Defect ID: DEFECT000636651 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS6.0.2 | Technology: IPv6 Addressing |
| Symptom: IPv6 Ping does not work sometimes | |

Condition: When VDX not trying to do IPv6 ND for unknown addresses and dropping packets instead, when both source and destination subnets are configured on the same VE interface

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

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

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

| | |
|---|---|
| Defect ID: DEFECT000637538 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.1.0 | Technology: IP Fabric |
| Symptom: In a IP fabric topology, the leaf (or border leaf) does not forward L3 packets (e.g. ICMP) destined to hosts on to peer leaf nodes. This causes the end to end L3 traffic between leaf nodes to fail. | |
| Condition: The prefix route learnt over evpn-bgp was incorrectly programed in Linux kernel, thereby forwarding traffic to wrong destination. | |
| Recovery: Issuing following command at NOS should help to recover from the problem state. clear arp no-refresh vrf <vrf-name> | |

Defect ID: DEFECT000637684

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

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

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

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

| | |
|-------------------------------------|--------------------------|
| Defect ID: DEFECT000638081 | |
| Technical Severity: Critical | Probability: High |

| | |
|---|---|
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS6.0.2 | Technology: DHCP - Dynamic Host Configuration Protocol |
| Symptom: traffic may flood and result in high latency. | |
| Condition: layer 2 loop exists in the topology resulting in high number of interrupts to VDX | |

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

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

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

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

| |
|-----------------------------------|
| Defect ID: DEFECT000640567 |
|-----------------------------------|

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

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

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

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

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

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

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

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

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

| |
|-----------------------------------|
| Defect ID: DEFECT000644087 |
|-----------------------------------|

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

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

Closed with code changes for Network OS v7.1.0

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as November 22, 2016 in Network OS v7.1.0.

| | |
|--|--|
| Defect ID: DEFECT000440702 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS3.0.1 | Technology: Static Routing (IPv4) |
| Symptom: User will be allowed to configure virtual interfaces for all VLANs except the default VLAN 1 | |
| Condition: Configuring virtual interfaces for VLANs | |

| | |
|--|--|
| Defect ID: DEFECT000443595 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS3.0.0 | Technology: IGMP - Internet Group Management Protocol |
| Symptom: The command "show ip igmp groups detail" may not show updated information (uptime & last reported values) about the learnt groups. | |
| Condition: This issue is observed only in the show command output & no functionality is impacted. | |

| | |
|---|--|
| Defect ID: DEFECT000471058 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS3.0.1 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: When SNMP with IPv6 is configured & SNMP Manager sends 2 get-requests divided in separate fragments, then "ICMPv6 Destination Unreachable" is returned from the switch. | |
| Condition: IPv6 & SNMP configuration | |

| | |
|--|---|
| Defect ID: DEFECT000490740 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS4.0.1 | Technology: Configuration Fundamentals |
| Symptom: Setting 'deny ip any any' ACL does not prevent telnet access via management port as expected. | |
| Condition: Setting following ACL does not prevent telnet access via management port as expected: sw0# show access-list ip ACL001 in ip access-list ACL001 on Management 1/0 at Ingress (From User) seq 10 deny ip any any (Active) seq 20 permit tcp any any (Active) | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000518780 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS4.1.3 | Technology: VMWare |
| Symptom: NSM crashes while configuring Port Group name with 128 or more character in vCenter. | |
| Condition: PG is configured and soon after that NSM Daemon crash is observed. | |
| Workaround: So not configure a port-group name with length more than 127 characters. | |
| Recovery: Once a NSM crash is seen, immediately delete the PG which was more than 128 characters long. Wait for switch to reboot. | |

| | |
|---|---|
| Defect ID: DEFECT000525575 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS4.1.3 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: If there are multiple subnets configured on an interface, PIM First Hop Router mechanism only works for the subnet with the highest IP address. | |
| Condition: PIM with multiple subnets on an interface. | |

| | |
|--|---|
| Defect ID: DEFECT000532352 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.1 | Technology: OSPFv3 - IPv6 Open Shortest Path First |
| Symptom: Show IPv6 OSPF neighbor is showing some special characters in the o/p. | |
| Condition: show IPv6 OSPF neighbor is showing some special characters in the o/p when there is no output to be printed. | |

| | |
|---|---|
| Defect ID: DEFECT000535663 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.1 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: VDX unexpectedly reloads when it connects to an NSX controller when a high scale of MAC addresses are injected over tunnel. It continues to reload unless booted with factory defaults. | |
| Condition: VDX connects to NSX controller and discovers more than 70000 MACs. | |
| Recovery: Boot with factory default configuration. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000548633 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.1 | Technology: Logical Chassis |
| Symptom: IPv6 clients will not be able to fetch IPv6 addresses from DHCP server | |
| Condition: The issue is observed after re-configuration of IPv6 address on Ve interface | |
| Workaround: Toggle the Ve interface by "shut" followed by "no shut" | |

| | |
|---|--|
| Defect ID: DEFECT000549174 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS5.0.1 | Technology: ACLs - Access Control Lists |
| Symptom: Continuous reload of switch may be observed on VDX6740 when static MACs are configured on physical interface | |
| Condition: The issue is observed in a Logical chassis cluster consisting of VDX6740 and VDX8770, if static MACs are configured on the physical interface of a LC of VDX8770. Note that issue is not observed when static MACs are configured on port-channels. Under these conditions if one of the VDX6740s are reloaded (or any action resulting in config replay), the VDX6740 that is reloaded might go into a continuous reload. | |
| Workaround: Remove any static MACs configured on the VDX8770 pointing to physical interfaces. Static MACs towards port-channels will not result in this issue. | |

| | |
|---|--|
| Defect ID: DEFECT000551918 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.0 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: Termination of Pimd process when multicast routes are aging out | |
| Condition: Issue is seen when PIM has learnt multiple group ranges for RP. | |

| | |
|---|--|
| Defect ID: DEFECT000555460 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.1 | Technology: IP Addressing |
| Symptom: 'ICMP unreachable are always sent' displayed in the configuration even when disabled in the configuration | |
| Condition: Default ICMP unreachable is not set | |

| | |
|--|---|
| Defect ID: DEFECT000556411 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: FCoE - Fibre Channel over Ethernet |
| Symptom: The RASLOG FCPH-1003 generated on console wrongly shows interface type as "Fi" instead of "Fcoe" with wrong tuple information. Functionality is not broken, only port type in raslog is printed wrongly. | |
| Condition: When logins with Duplicate WWN are attempted on multiple ports at same time with Ethernet port being the port on which second login is attempted. | |

| | |
|--|--------------------------------------|
| Defect ID: DEFECT000560160 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: Inband Management |
| Symptom: Pings to VE interfaces are slow after upgrade | |
| Condition: In rare upgrade scenarios. | |
| Recovery: Shut/No shut the affected interface where the ping is seen slow or have a reload. | |

| | |
|--|--|
| Defect ID: DEFECT000560868 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.1 | Technology: IP Addressing |
| Symptom: IP directed-broadcast feature is not working as expected. | |
| Condition: With a regular topology, the functionality did not work as expected. | |

| | |
|---|--|
| Defect ID: DEFECT000562543 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: IP ACL for SNMP community and v3 user config lost after loading the config back to running-config from back-up config | |
| Condition: When we do config upload of running configuration with SNMP IP ACL's applied on SNMP community/ v3 users. | |

| | |
|--|--|
| Defect ID: DEFECT000562722 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Traffic Management |
| Reported In Release: Network OS5.0.2 | Technology: Rate Limiting and Shaping |
| Symptom: ipv6 icmpv6 rate-limiting does not work per interface | |
| Condition: The above situation occurs under two conditions 1. More than one ipv6 interface 2. Different rate-limiting value configured When both of the above conditions met, then the recently configured rate-limiting value applied to all interfaces | |

| | |
|---|---|
| Defect ID: DEFECT000562737 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS4.0.1 | Technology: OpenStack Integration |
| Symptom: SNMP trap of topology change will be sent from the switch, when switchport configuration is done on an interface where spanning-tree is shutdown. | |
| Condition: Topology change trap will be observed, when switchport configuration is done on an interface in spanning-tree shutdown state. | |

| | |
|---|--|
| Defect ID: DEFECT000566249 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Traffic Management |
| Reported In Release: Network OS6.0.1 | Technology: Rate Limiting and Shaping |
| Symptom: IPv4 and IPv6 ICMP rate-limiting does not work after HA failover. | |
| Condition: HA is required for this problem and executing CLI will fix issue. | |

| | |
|---|--|
| Defect ID: DEFECT000567339 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: ARP - Address Resolution Protocol |
| Symptom: debug arp packet output shows destination mac address of ARP request as ffff:fff:fff, instead of 0000:0000:0000 | |
| Condition: debug arp packet command is executed | |

| | |
|--|--|
| Defect ID: DEFECT000568542 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.2 | Technology: Static Routing (IPv4) |
| Symptom: static route and next-hop gateway on a virtual Ethernet interface should work fine | |
| Condition: proxy ND ipv6 not supported | |

| | |
|--|---|
| Defect ID: DEFECT000570331 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS5.0.2 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: Traffic loss will be seen if there is a change in RP(Rendezvous Point) address configuration | |
| Condition: if there is a change in RP(Rendezvous Point) address configuration in the running system | |

| | |
|---|--|
| Defect ID: DEFECT000570673 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: UDLD - Uni-Directional Link Detection |
| Symptom: UDLD blocks links which are connected using certain breakout ports between different Extreme devices such as between VDX and CER. | |
| Condition: UDLD blocks the link after detecting mismatch between locally stored vs received port numbers in UDLD PDUs. | |

| | |
|---|---|
| Defect ID: DEFECT000572393 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.1 | Technology: OSPFv3 - IPv6 Open Shortest Path First |
| Symptom: Not able to remove ospf area 0.0.0.0 in ipv6 ospf router user vrf | |
| Condition: "%Error" message while remove ospf area 0.0.0.0 in ipv6 ospf router under user vrf eventhough the interface doesn't have OSPF configuration | |

| | |
|---|--|
| Defect ID: DEFECT000573107 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS6.0.1 | Technology: ACLs - Access Control Lists |
| Symptom: When we applied IP ACL on SNMP community/user configuration, then wildcard subnet mask on IP ACL is not working on SNMP. But subnet mask on IP ACL is working fine on SNMP. | |
| Condition: When we have wildcard subnet mask on IP ACL applied for SNMP configuration, then we will observe this issue. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000573422 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.2 | Technology: Port Mirroring |
| Symptom: SPAN session fails when the destination port is deleted and rejoins the cluster. | |
| Condition: When an interface(s) of a node is set as a destination in a SPAN session and the node is removed from the VCS. Removal of the node leaves the SPAN session in un-defined state, causing problems with new SPAN destination configuration. | |
| Workaround: Remove the SPAN sessions that were monitoring the interfaces whose node was removed from VCS, and create sessions. | |

| | |
|--|---|
| Defect ID: DEFECT000574438 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: OpenStack Integration |
| Symptom: VDX 6940-144s 40 GbE ports 97 - 108 may flap after changing frame size and/or nframes using noscli "diag dport setargs" command. | |
| Condition: VDX 6940-144s 40 GbE ports 97 - 108 may flap after changing framesize and/or nframes noscli "diag dport setargs" command. | |
| Workaround: Shut/no shut the port if persistent link flapping occurs. | |
| Recovery: Shut/no shut the port if persistent link flapping occurs. | |

| | |
|--|--|
| Defect ID: DEFECT000574645 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: The "fruStatusChanged" trap may be received more than once while power ON the line cards in VDX8770-4 and VDX8770-8 platforms. | |
| Condition: The "fruStatusChanged" trap may be received more than once while power ON the line cards. | |

| | |
|--|--|
| Defect ID: DEFECT000576391 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: The HA failover start trap may not be seen for the HA fail over event, for user defined VRF in VDX-6740 platforms. | |
| Condition: The HA failover start trap may not be seen for user defined VRF, for HA fail over event. | |

| | |
|---|---|
| Defect ID: DEFECT000577171 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: OpenStack Integration |
| Symptom: The Network OSCLI command "show openflow interface" does not reflect the actual operating speed of the OpenFlow interface | |
| Condition: If the interface speed has been manually changed to something else which is not same as suggested via interface name | |

| | |
|--|---|
| Defect ID: DEFECT000577563 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS6.0.2 | Technology: OpenStack Integration |
| Symptom: On upgrade to nos5.0.2b release, certain interfaces may encounter following error message & will remain admin-down: <Date-Timestamp>, [NSM-1028], 7885, M2 Active DCE, ERROR, <Host-name>, Incompatible SFP transceiver for interface TenGigabitEthernet 1/1/17 is detected. | |
| Condition: Issue encounters on upgrade / downgrade of VDX8770 to nos5.0.2b release using following interfaces: <ul style="list-style-type: none"> Linecard 48X10G: <ul style="list-style-type: none"> (a) Links using Copper SFP's (1G) configured with "speed 1000" (b) Links using Fiber Optics (1G and 10G) configured with "speed" command (c) Links using Twinax cables (10G) Linecard 48X1G: <ul style="list-style-type: none"> (d) Links using Copper SFP's (1G) (e) Links using Fiber Optics (1G only) Linecard 6X100G: <ul style="list-style-type: none"> (f) Links using 100G Optics | |

| | |
|---|---|
| Defect ID: DEFECT000577822 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS5.0.1 | Technology: OpenStack Integration |
| Symptom: Errors [crc, encoding...] on 8G links. | |
| Condition: The issue is only seen on 8G links to 3Par storage devices using 16G SFPs | |
| Workaround: Changing the SFP to 8G SFP and running at 8G speed the issue was not seen. | |

| | |
|---|--|
| Defect ID: DEFECT000577928 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: MAPS - Monitoring and Alerting Policy Suite |
| Symptom: Slot number need to be verified while creating groups on VDX8770. | |
| Condition: Currently the API which converts slot/port to port index doesn't throw error while creating group on VDX8770. | |

| | |
|--|---|
| Defect ID: DEFECT000578258 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: IP Fabric |
| Symptom: Traffic loss may be observed for destination subnets under non-default VRF advertised over BGP-EVPN using L3VNI. | |
| Condition: Leaf nodes extending the VRF over BGP-EVPN are not advertising any prefix route. | |
| Workaround: One of following options may be chosen: (1) Redistribute connected routes under VRF into BGP VRF. (2) Configure static IP route and redistribute into BGP VRF. (3) Configure network or static-network under BGP VRF instance. | |

| | |
|---|--------------------------------------|
| Defect ID: DEFECT000578492 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: Inband Management |
| Symptom: "show interface status" incorrectly shows the status of some port channels as "connected" even though they are down due to all member interfaces down. No functional impact except the incorrect display. | |
| Condition: Rarely happens due to a corner case condition for some port channels | |
| Workaround: "show interface" command can be used to see the correct status. | |

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

| | |
|--|------------------------------------|
| Defect ID: DEFECT000579138 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: Very rare case chassis name set CLI fails. | |
| Condition: After upgrade to 6.0.1 | |
| Recovery: Reload and re-apply the CLI | |

| | |
|---|---|
| Defect ID: DEFECT000579176 | |
| Technical Severity: High | Probability: Medium |
| 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: BFD may not work over Layer 3 Port Channels when the gateway address and nexthop pointing to port channel overlap | |
| Condition: Running BFD with Layer 3 Port Channels | |

| | |
|--|--|
| Defect ID: DEFECT000579234 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VLAN - Virtual LAN |
| Symptom: Multicast traffic destined for static multicast address, will flood, if the mac is configured on remote node of VCS. | |
| Condition: Static multicast MAC is configured in a remote node within a VCS, with no local interface part of the group. | |

| | |
|---|---|
| Defect ID: DEFECT000579664 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS5.0.1 | Technology: OpenStack Integration |
| Symptom: "waiting for pending actions to exit" warning message appears on console session and eventually VDX experience unexpected reload. | |
| Condition: When user query to get any running-config using netconf and if password contains a special char like ";", VDX throws "waiting for pending actions to exist" error and DCM gets terminated eventually. | |

| | |
|--|--|
| Defect ID: DEFECT000579695 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: IP Addressing |
| Symptom: SNMP walk on the IP-FORWARD-MIB may throw an error | |
| Condition: SNMP operations on IP-FORWARD-MIB | |

| | |
|--|--|
| Defect ID: DEFECT000579835 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP Set operation on Port-channel interface for the ipv6RouterAdvertLinkMTU object in ipv6RouterAdvertTable in RFC 4293 will fail. The SNMP Set operation will fail for other objects also in this table. | |
| Condition: The SNMP Set operation will fail only on port-channel interfaces. | |
| Workaround: In order to configure the ipv6RouterAdvertLinkMTU, the corresponding CLI "ipv6 nd mtu" needs to be used. | |

| | |
|--|--|
| Defect ID: DEFECT000579904 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS5.0.2 | Technology: AAA - Authentication, Authorization, and Accounting |
| Symptom: Command set field on the Windows based TACACS server is empty | |
| Condition: 1. When TACACS server is windows based 2. Accounting is enabled | |

| | |
|--|--|
| Defect ID: DEFECT000580478 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VLAN - Virtual LAN |
| Symptom: Sometimes, SFP removal messages are displayed incorrectly even though the media is present, when a chassis disable is executed after failover or ISSU. | |
| Condition: Media presence check is incorrect on the new active partition after failover or ISSU | |
| Recovery: 'no shut' on the interface would make the correct Media presence state consistent. | |

| | |
|--|--|
| Defect ID: DEFECT000581205 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: In rare case, SNMPv3 traps will not be received when any host is configured as SNMPv3 trap recipient under RBridge mode. | |
| Condition: Configure SNMPv3 host under RBridge mode. | |
| Recovery: reconfigure the specific SNMPv3 host config under RBridge mode. | |

| | |
|---|---|
| Defect ID: DEFECT000581259 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: IP Fabric |
| Symptom: Even though overlay-gateway configuration is deactivated, BGP discovered dynamic tunnels are still present. Traffic loss will be observed if remote Leaf nodes send traffic over dynamic tunnels. | |
| Condition: Overlay-gateway configuration is deactivated using "no activate" command. | |
| Workaround: Avoid deactivating the overlay-gateway using "no activate" command. Instead detach the RBridge from overlay gateway configuration. | |

| | |
|---|--|
| Defect ID: DEFECT000582010 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.2 | Technology: ARP - Address Resolution Protocol |
| Symptom: Under rare conditions, some of the hosts may lost IP connectivity with the VDX switch acting as a layer-3 gateway. | |
| Condition: This would occur if the MAC to the IP association of a VDX learnt ARP changes. ie.. For the same IP address, the MAC changes from say MAc1 to Mac2. | |
| Recovery: "clear arp no-refresh" would clean the ARP table and recover from the problem state. | |

| | |
|---|--|
| Defect ID: DEFECT000582119 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS7.0.0 | Technology: IGMP - Internet Group Management Protocol |
| Symptom: The tunnel terminated IGMP frames sent to other nodes can loop back to the source node. The CPU generated IGMP frames are not getting source suppressed in active-active gateway. | |
| Condition: This happens in specific tunnel topology with multicast root RBridge and BUM forwarder. The tunnel terminated IGMP frames sent to other nodes are trapped and flood back on the VLAN by control path. These packets can loop back to source node. | |
| Recovery: Shut down the tunnel. | |

| | |
|---|---|
| Defect ID: DEFECT000582797 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Configuration Fundamentals |
| Symptom: Switch does not responds to any user inputs. | |
| Condition: Using 'PUT' request to update BGP configuration instead of using 'PATCH' request (when router BGP has neighbor associated to a peer group) causes switch not to respond to further user commands. | |
| Workaround: Use PATCH request instead of PUT to update any configuration as documented in Extreme Network OS REST API Guide. | |

| | |
|--|--|
| Defect ID: DEFECT000582847 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: On a VDX 6940, when a snmpwalk is done on ifTable in the IF MIB, only the first 36 interfaces are fetched. Remaining interfaces are not retrieved. | |
| Condition: This issue is seen only on VDX 6940. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000583123 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.2 | Technology: Logical Chassis |
| Symptom: There is a time delay(debounce-timer delay) of approximately 1 sec between underlay network down and tunnel down because of which traffic impact may occur for this debounce-timer duration . | |
| Condition: The above mentioned time delay happens whenever tunnel goes down. Now customer is provided with the following knob to suppress the debounce-timer delay. [no] system tunnel suppress-debounce | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000583245 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: Sysmon |
| Symptom: The event-handler trigger-mode options "on-first-instance" and "only-once" incorrectly allows one more execution of an action to be triggered. | |
| Condition: Activate an event-handler with trigger-mode set to "on-first-instance" or "only-once". | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000583349 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Licensing |
| Symptom: In VDX6940-144S, 100G mode configuration replay fails when executing "copy <file> running-config" because DPOD license has not been reserved. | |
| Condition: This issue will happen only if the DPOD license has not been reserved. | |
| Workaround: Manually reserve the DPOD license and then run "copy <file> running-config" | |
| Recovery: Manually reserve the DPOD license and then run "copy <file> running-config" | |

| | |
|---|--|
| Defect ID: DEFECT000583626 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: Error message is not thrown, if more than the maximum number of SNMP V3 users configured with both global and local configurations combined together. | |
| Condition: If the SNMP V3 users configured more than the maximum number of supported, for both global and local configurations combined together then the error message is not thrown. | |

| | |
|---|---|
| Defect ID: DEFECT000584215 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: IEEE BPDU packets are flooded from one VF to another, in the absence of "spanning-tree ieee-bpdu limit-vlan-flood" configuration. | |
| Condition: IEEE BPDU packet are received at the ingress port of a switch configured with VFs. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000584364 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: Syslog |
| Symptom: User can configure the user defined VRF in cluster, though the user defined VRF is not configured on all the RBridge's. | |
| Condition: In cluster though the VFR is not configured on all the RBridge's, it is allowing to configure syslog-server on user defined VRF. | |

| | |
|---|--|
| Defect ID: DEFECT000584668 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.2 | Technology: IP Addressing |
| Symptom: The access-list configured for the interface with "Routed" keyword may not act upon the traffic flows on ingress destined towards the VRRP-E VIP. | |
| Condition: Applicable for the traffic destined to VRRP-E Virtual-mac coming on ingress & access-list using "Routed" keyword. | |
| Workaround: The traffic flows destined towards VE MAC will not be impacted & thus hosts can be configured to point to VE IP address as default gateway. Alternatively, remove the "routed" keyword in the access-list. | |

| | |
|--|--|
| Defect ID: DEFECT000584709 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VLAN - Virtual LAN |
| Symptom: Physical or port-channel is not added back to normal VLAN in a particular sequence. | |
| Condition: Physical or port-channel is not added back to normal VLAN after changing a private VLAN to a normal VLAN on a primary VLAN | |
| Workaround: Delete private VLAN and create the same again instead of changing the type on a private VLAN. | |
| Recovery: | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000584733 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: Management GUI |
| Symptom: Firmware downgrade from Network OS6.x to Network OS5.x using BNA display below error in BNA even though firmware download process is successful on VDX. Download Failed: (Unknow Error Code: 1) Other Errors: Firmware Image download reboot operation has timed out | |
| Condition: Firmware downgrade from Network OS6.x to Network OS5.x using BNA cause the issue. | |

| | |
|--|---------------------------------------|
| Defect ID: DEFECT000584922 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS6.0.2 | Technology: SSH - Secure Shell |
| Symptom: Shutting Telnet server on Active and Standby partition fails | |
| Condition: High Availability fail over | |

| | |
|---|---|
| Defect ID: DEFECT000585043 | |
| 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+ - IPv6 Border Gateway Protocol |
| Symptom: When multi-hop BFD session is created, default BFD interval will be shown for loopback interface in show BFD output | |
| Condition: Default BFD interval will be shown for loopback interface in show bfd output | |

| | |
|--|--|
| Defect ID: DEFECT000585337 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: NTP - Network Time Protocol |
| Symptom: Configuring more than one NTP server with same IP address but different VRF name will always use the first VRF name in the list to sync with NTP server. Also removing one of the NTP server entry fails to remove it completely | |
| Condition: This issue is introduced as part of NTP VRF support. Issue will be seen with more than one NTP server with same IP address and different VRF name. | |
| Workaround: Remove all NTP servers and configure the desired servers again | |

| | |
|---|---|
| Defect ID: DEFECT000585392 | |
| 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: BNA unable to discover Network OS switch, when Network OS switch is connected to a FCR. | |
| Condition: 1. Have a setup with VCS connected to FCR 2. Install BNA 14.0.1 3. Shift to IP tab. 4. In Discovery dialog, add IP of one of the VCS switches and click okay. 5. Observe the device is not discovered and shows "Discovery Failed" message. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000585422 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: IPv6 traffic loss since neighbor resolution fails, as neighbor solicitation packet forward might fail if it is switched over VDX6940. | |
| Condition: IPv6 traffic that ingress on a specific set of ports which are serviced by odd numbered RTEs on VDX6940 platform will not be forwarded properly. | |
| Workaround: Use ports that are part of trunk group 3 and 4 for VDX6940-144s or Use ports that are part of trunk group 1 and 4 for VDX6940-36Q platform. | |

| | |
|---|---|
| Defect ID: DEFECT000585445 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Some 40 GbE ports on VDX 6940-144S may not come online after cold boot. | |
| Condition: Some 40 GbE ports on VDX 6940-144S may not come online after cold boot. | |
| Workaround: Execute noscli command shut / no shut on the 40 GbE port to bring it online. | |
| Recovery: Execute noscli command shut / no shut on the 40 GbE port to bring it online. | |

| | |
|---|---|
| Defect ID: DEFECT000585634 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: VCS Fabric |
| Symptom: On consecutive reboot certain ports may be shown administratively down. | |
| Condition: Multiple switch reboots. | |
| Recovery: Enable the port to administratively up. | |

| | |
|---|--|
| Defect ID: DEFECT000585723 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: MIB walk for IP Forwarding MIB will return with an error with unnumbered interfaces. | |
| Condition: MIB walk of IP Forwarding MIB and has ECMP routes with unnumbered/L3 VNI interfaces will lead to error. | |

| | |
|---|---|
| Defect ID: DEFECT000585818 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Configuration Fundamentals |
| Symptom: REST GET request fails to retrieve an empty leaf | |
| Condition: When 'cli-show-no' annotation is used along with the cli-run-template to customize the display of an empty leaf in the 'show running config', the same might not be retrieved through the REST GET request. | |

| | |
|--|---|
| Defect ID: DEFECT000585895 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: OpenStack Integration |
| Symptom: VDX 6940 and 6940-144S ports may not come online after converting from 40 GbE to breakout mode 4 x 10 GbE. | |
| Condition: VDX 6940 and 6940-144S ports may not come online after converting from 40 GbE to breakout mode 4 x 10 GbE. | |
| Workaround: Execute noscli shut / no shut commands | |
| Recovery: Execute noscli shut / no shut commands | |

| | |
|---|--|
| Defect ID: DEFECT000585903 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: MAPS - Monitoring and Alerting Policy Suite |
| Symptom: IPMAPS Custom policy modifications are not reflected. | |
| Condition: IPMAPS Custom policy modifications are not dynamically reflected. | |
| Workaround: Revert to default policy, and then reapply custom policy. | |
| <p>Run CLI "enable policy <policy_name> actions <actions_list>" then we can re-enable the same policy to reflect the changes made.</p> <p>Here actions list can be same as what was already configured.</p> | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000585927 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.1 | Technology: Port Mirroring |
| Symptom: Mirrored VXLAN packets outer header was getting removed while going out on destination mirror port. This gives misleading information when validating the VXLAN mirroring. The data path traffic goes out properly but mirrored copy has the outer header stripped only in VXLAN frames. | |
| Condition: The VXLAN packets outer header is not handled properly and causing the stripped packet to go out on destination mirror port | |
| Recovery: This is not functional data path issue, but mirrored information shows wrong details. | |

| | |
|--|---|
| Defect ID: DEFECT000585960 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: VCS Fabric |
| Symptom: 40G Interface is administratively (or) protocol down with FFDC raslogs | |
| Condition: Admin operations on 40G Interface. | |
| Recovery: Shut/no-shut both interfaces on either side of the link | |

| | |
|--|--|
| Defect ID: DEFECT000585970 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: VRRPv2 - Virtual Router Redundancy Protocol Version 2 |
| Symptom: On VDX 8770 switch, maximum VRRPv2 and VRRPv3 sessions supported on an interface are 16 in Network OS6.0.x. This limit got increased to 32 in Network OS7.0.0. Firmware downgrade from Network OS7.0.0 to Network OS6.0.x need to be blocked in case if more that 16 sessions are present on an interface. | |
| Condition: Issue can be seen if more that 16 VRRPv2 and VRRPv3 sessions are configured on an interface and firmware is downgraded from Network OS7.0.0 to Network OS6.0.x. In this case only 16 sessions will get enabled and rest will be disabled. | |
| Workaround: As a workaround user should delete/unconfigure more than 16 VRRPv2/VRRPv3 sessions present on an interface in Network OS7.0.0 before downgrading it to Network OS6.0.x. | |

| | |
|--|--|
| Defect ID: DEFECT000586001 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: ARP - Address Resolution Protocol |
| Symptom: IPv4 DHCP relay packets forwarded through a VxLAN tunnel is trapped but not forwarded unless ARP is forcefully resolved. | |
| Condition: Running DHCP Relay in IP Fabric EVPN. | |
| Workaround: Resolve ARP forcefully. | |

| | |
|---|---|
| Defect ID: DEFECT000586125 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: ISL between VDX8770 and VDX6940 could be flapping continuously. | |
| Condition: One side of the link connects to VDX8770 from VDX 6940 | |
| Workaround: Toggle the flapping ISL link by doing "shutdown" and "no shutdown" of the ports. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000586178 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Non-existent port-channel shows up in "show fcoe interface ethernet" | |
| Condition: 1. Create a port-channel 2. Add members to it and make it fcoe-provisioned 3. Delete the port-channel | |
| Workaround: Remove FCOE provisioning from port-channel before deleting it | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000586205 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: Syslog |
| Symptom: Syslog server not working via inband under def-vrf, mgmt.-vrf and user-vrf, all having different IP address | |
| Condition: Configure inband under def-vrf, mgmt.-vrf and user-vrf for syslog server. | |

| | |
|---|--|
| Defect ID: DEFECT000586230 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: AMPP - Automatic Migration of Port Profiles |
| Symptom: System goes for unexpected reboot | |
| Condition: More than 8000 profiled and non-profiled mac addresses are learnt on profiled port and all these flows are moving across the ports. | |

| | |
|---|---|
| Defect ID: DEFECT000586252 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Configuration Fundamentals |
| Symptom: Physical or port-channel is not added back to normal VLAN in a particular sequence. | |
| Condition: Physical or port-channel is not added back to normal VLAN after changing a private VLAN to a normal VLAN on a primary VLAN. | |
| Workaround: Delete private VLAN and create the same again instead of changing the type on a private VLAN. | |

| | |
|--|--|
| Defect ID: DEFECT000586577 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.1 | Technology: Multi-VRF |
| Symptom: VDX switch can go for unexpected reload after configuring no vrf. | |
| Condition: If static route leaks exist in system and we configure "no vrf" command. | |

| | |
|---|--|
| Defect ID: DEFECT000586614 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VLAN - Virtual LAN |
| Symptom: 100G port-channel displays some junk values after HA failover some times. | |
| Condition: 100G port-channel may not show proper statistics after HA failover. | |
| Workaround: .Clear the counters on 100G port-channel. | |
| Recovery: .shut/no-shut on port-channel followed by clearing counters. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000586638 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: VDX6740 and VDX6740-T are unable to transmit control plane traffic. As a side effect, "Fabric Hello Timeout: Inter-switch Frame Delivery problem" may be seen in any of the nodes in the VCS cluster | |
| Condition: ISSU or HA failover on VDX6740 or VDX6740-T | |
| Workaround: Perform another HA Failover from the Network OSCLI | |

| | |
|---|---|
| Defect ID: DEFECT000586771 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Configuration Fundamentals |
| Symptom: Secondary to primary association fails in a particular sequence. | |
| Condition: Cannot associate secondary vlan to primary vlan, once its type is changed from community to isolated or vice versa. | |
| Workaround: Delete secondary vlan, create, configure secondary type and associate again. | |

| | |
|--|---|
| Defect ID: DEFECT000586852 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: CLI - Command Line Interface |
| Symptom: Chassis enable may fails after ISSU followed by chassis disable in a scaled environment. | |
| Condition: Chassis enable fails sometimes in a scaled environment with ISSU followed by chassis disable because of slow responding clients. | |
| Workaround: Not to try ISSU followed by chassis disable/enable | |
| Recovery: Reload the switch | |

| | |
|---|---|
| Defect ID: DEFECT000586856 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: BGP add path is not showing up all the 4 available paths. | |
| Condition: Sometimes when the RR is reloaded the BGP add path is not showing up all the 4 available paths. | |

| | |
|---|---|
| Defect ID: DEFECT000586891 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: CLI - Command Line Interface |
| Symptom: The CLI command 'show debug ipv6 nd' is not provided to the user | |
| Condition: When the user tries to issue the CLI command 'show debug ipv6 nd' he won't see an option for 'nd' | |

| | |
|--|---|
| Defect ID: DEFECT000586973 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.0.0 | Technology: LDAP - Lightweight Directory Access Protocol |
| Symptom: LDAP authentication is not working | |
| Condition: LDAP authenticaion is not working via inband default and non-default-vrf | |

| | |
|--|------------------------------|
| Defect ID: DEFECT000587139 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Metro VCS |
| Symptom: During VDX6940-144S switch bootup, FFDC core is seen. | |
| Condition: Slow processing of switch during bootup stage leads to core file generation. | |
| Workaround: Reload the switch again. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000587154 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS4.1.3 | Technology: Logical Chassis |
| Symptom: Switch may reload when BNA queries with get-config when the config is large OR when BNA polls at an aggressive rate (not configured for lazy-polling). | |
| Condition: Issue can happen when BNA does get-config for large cluster which has more than 4500+ interface config OR when BNA is polls aggressively. | |
| Workaround: Please do not use BNA if the cluster is large (with 4200+ interfaces) & ensure it is configured with lazy-polling | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000587170 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: Continuous occurrence of ECC correctable errors | |
| Condition: This is very rare scenario to occur. | |

| | |
|---|--|
| Defect ID: DEFECT000587208 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: An unexpected reboot of Dcmd might occur while collecting support-save data in a scenario where OSPFv2 config on Loopback interface is missing from protocol while config present in running config | |
| Condition: collecting support-save while OSPFv2 config on Loopback interface is missing from protocol although config is present in running config | |
| Workaround: Remove Loopback interface and collect support-save | |

| | |
|--|--|
| Defect ID: DEFECT000587276 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.2 | Technology: Hardware Monitoring |
| Symptom: Blocked or stopped fan units may not show up as faulty. | |
| Condition: This was a defect in the original release of this product. | |

| | |
|---|---|
| Defect ID: DEFECT000587380 | |
| Technical Severity: Low | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.1 | Technology: CLI - Command Line Interface |
| Symptom: “no snmp trap link-status” does not show up in show running config after executing it under an interface level. | |
| Condition: “no snmp trap link-status” under interface is enhanced and it got missed to keep the same in running config. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000587419 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: Syslog |
| Symptom: Ipv6 syslog-server not working via inband def-vrf and user-vrf. When multiple server configured as default-vrf or user-defined vrf. | |
| Condition: Was deferred from 7.0.0 but fixed in 7.0.1. | |

| | |
|--|--|
| Defect ID: DEFECT000587463 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Software Installation & Upgrade |
| Symptom: At time of Firmwaredownload, the switch may fail to come back to an online state | |
| Condition: Issue can happen when the inetd file needs to be updated with a large configuration change | |
| Recovery: An additional reload will recover the switch. | |

| | |
|---|--|
| Defect ID: DEFECT000587615 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: The SNMP V3 traps may not be received for the SNMP v3 host recipients configured under RBridge mode. | |
| Condition: The trap may not be received after upgrade from Network OS6.0.1a to Network OS7.0.0 with cold boot option | |

| | |
|---|--|
| Defect ID: DEFECT000587617 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: Static leaked VRF route can not be imported into BGP RIB-IN and can't advertise via eBGP/iBGP using network/static command. | |
| Condition: Advertise static leaked VRF route via BGP. | |
| Workaround: Use "redistribute static" command to leak the static VRF route into BGP RIB-IN and then can advertise it via eBGP/iBGP. | |

| | |
|--|---|
| Defect ID: DEFECT000587637 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: NETCONF - Network Configuration Protocol |
| Symptom: Netconf RPC "get-interface-detail" does not provide physical interfaces details. It provides only port-channel details. | |
| Condition: This issue will happen only when number of port-channels configured are equal to or more than 70. If number of port-channels are less than 70, this issue will not be encountered. | |
| Workaround: Total number of port-channels configured should be less than 70. | |
| Recovery: If total number of port-channels configured are exceeding 70, delete few port-channels to reduce the total count to be less than 70. | |

| | |
|--|---|
| Defect ID: DEFECT000587654 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: OSPF - IPv4 Open Shortest Path First |
| Symptom: The ECMP configuration in the hardware profile shows incorrect values. | |
| Condition: This will only happen when a user changes both route-table profile type and maximum-path at the same time using the hardware-profile command | |
| Workaround: The user can change the route-table profile type and maximum-path one at a time. | |
| Recovery: The user can re-run the hardware-profile command to set the maximum-path with the correct value. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000587704 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Traffic is not being forwarded and a member of lag shows up on one side and but show lag/negotiation failed on other side even though the link is shown in LLDP neighbor. | |
| Condition: LAG hasat least two members and after performing ISSU/reboot sometimes member port goes into LAG negotiating/failed. | |
| Recovery: Shut./no shut on the member port/LAG. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000587767 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Possible for Edge port interfaces to stay inactive after chassis enable command. | |
| Condition: This issue can occur in releases prior to Network OS 7.0. If multiple attempts to issue the chassis enable command is failed and the command is retried, it is possible that the configuration replay will be blocked after the chassis enable succeeds. | |
| Recovery: Issue chassis disable then chassis enable. | |

| | |
|--|---|
| Defect ID: DEFECT000587804 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: IP Fabric |
| Symptom: Even though there are no matching EVPN import route-targets configured under VRF, imported EVPN routes are present in BGP VRF table. | |
| Condition: EVPN import route-target(s) is/are removed while matching routes are present in BGP-EVPN and imported into BGP VRF table. | |
| Workaround: Issuing "clear bgp evpn neighbors all soft in" command should cleanup the routes which are still imported in BGP VRF instance after matching EVPN import route-targets are removed. | |

| | |
|--|---|
| Defect ID: DEFECT000587828 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BFD - BiDirectional Forwarding Detection |
| Symptom: BFD Session on a tunnel interface, could go down when one of the egress vLAG member node is disabled. | |
| Condition: This issue could occur only if the underlay egress interface of the tunnel is a vLAG interface and also the vLAG membership exists only in remote RBridge(s), i.e., vLAG membership not exists in RBridges doing tunnel termination. | |
| Recovery: Disabling and re-enabling of the complete vLAG interface will bring up the BFD session. | |

| | |
|---|---|
| Defect ID: DEFECT000587925 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: FCoE - Fibre Channel over Ethernet |
| Symptom: Syslog daemon generates a silent core file as it is restarted to reload configuration. There is no crash or loss of traffic in this case. | |
| Condition: Defect exists in previous releases of Network OS. Core file is generated due to SIGTERM signal received by syslog instead of SIGHUP. | |

| | |
|--|---|
| Defect ID: DEFECT000587984 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: OpenStack Integration |
| Symptom: VDX6940-144S 100GbE port may not come online after fastboot in Network OS release 7.0.0. | |
| Condition: VDX6940-144S 100GbE port may not come online after fastboot Network OS release 7.0.0. | |
| Workaround: "shut" then "no shut" the 100GbE port to bring it online. | |
| Recovery: "shut" then "no shut" the 100GbE port to bring it online. | |

| | |
|--|--|
| Defect ID: DEFECT000588001 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: AMPP - Automatic Migration of Port Profiles |
| Symptom: Traffic may flood though the source mac was seen behind profiled port | |
| Condition: Port-profile is configured on a VLAG and 'clear-mac-address table' command is executed more than 10 times in short interval. | |

| | |
|--|---|
| Defect ID: DEFECT000588041 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BFD - BiDirectional Forwarding Detection |
| Symptom: BFD session state for tunnels do not come up | |
| Condition: When tunnel toggles on one RBridge (which has least RBridge ID) in the VTEP, BFD session state will not come up for that tunnel. | |
| Workaround: It is recommended to remove and add the RB-X in overlay gateway configuration. | |

| | |
|--|---|
| Defect ID: DEFECT000588062 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: OSPFv3 - IPv6 Open Shortest Path First |
| Symptom: OSPFv6 is not coming up when IPSEC authentication is configured on the peers but inbound and outbound are not updated with IPsec authentication value. | |
| Condition: IPsec authentication to be configured on the switch over OSPFv3. When IPsec authentication key is configured, it will take key roll-over time for the key to get updated on the box. | |
| Workaround: The neighbor ship will automatically come up after the key rollover time | |

| | |
|--|---|
| Defect ID: DEFECT000588178 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: OpenStack Integration |
| Symptom: Interface remains protocol down after speed change | |
| Condition: speed change config performed on an interface which is not in protocol up state. | |
| Recovery: shut/no-shut the interface | |

| | |
|--|---|
| Defect ID: DEFECT000588190 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: IP Fabric |
| Symptom: Aggregate route(s) configured under BGP VRF instance are not exported into BGP-EVPN. | |
| Condition: BGP VRF address-family is removed and added back. | |
| Workaround: Remove the aggregate route configuration under BGP VRF instance and configure it again. | |

| | |
|---|--|
| Defect ID: DEFECT000588238 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: 'Invalid value' error is thrown for 'vni add' command under 'rbridge-id evpn-instance' mode. | |
| Condition: Configure 'vni add' command under 'rbridge-id evpn-instance' mode. If the value falls in 10000000-15999999 range. | |
| Workaround: 1. Use a VNI range in 'vni add' command that is less than 10000000-15999999. 2. Use 'vni <vni-number>' CLI under 'rbridge-id evpn-instance' mode. | |

| | |
|---|---|
| Defect ID: DEFECT000588241 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Configuration Fundamentals |
| Symptom: breakout lane 1 will not go admin down with incompatible media | |
| Condition: A QSFP28 media is present in the cage when breakout is configured. | |
| Workaround: Remove the QSFP28 media from the cage before executing the breakout CLI. | |
| Recovery: Manually 'shut' first lane on the cage. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000588265 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Management GUI |
| Symptom: Gradual increase of memory for process DCMd when REST requests are sent continuously for a long period of time. | |
| Condition: Continuous REST requests are sent to the switch to retrieve operational information such as 'show interface', 'show running-config', etc. | |
| Recovery: If standby management module is available, please perform HA failover when DCMd memory consumption crosses 600MB.. | |

| | |
|---|---|
| Defect ID: DEFECT000588323 | |
| Technical Severity: Critical | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: CLI - Command Line Interface |
| Symptom: Switchport configuration may fails in a particular sequence. | |
| Condition: After having a L3 configuration on a physical interface and upgrading to a particular build mentioned in the defect description switchport configuration may fails. | |
| Workaround: Save and reboot with default configuration. | |
| Recovery: | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000588333 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Switch will go for an unexpected reload when trying to enter into RBridge ID context in configuration mode. | |
| Condition: if any special character other than "-" or "," is given for specifying the range of RBridge ID will cause an exception and will trigger switch reload. E.g. : rbrdige 9*1 | |
| Workaround: For specifying the RBridge ID range, use only '-' or ',' and do not use other special characters | |

| | |
|---|--|
| Defect ID: DEFECT000588451 | |
| Technical Severity: High | Probability: Low |
| 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: IPv6 traffic may not forward when it received on tunnel | |
| Condition: When bigger VNI like 10000000 configured as l3vni | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000588463 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: For IPv6/v4 Virtual-IP Address stops responding to pings after protocol change from VRRP to VRRP-E. | |
| Condition: Protocol is changed from VRRP to VRRP-E with some enabled sessions, and vice-versa. | |
| Recovery: Reload system. | |

| | |
|--|---|
| Defect ID: DEFECT000588519 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: When the RBridge responsible for Multicast distribution over VXLAN Tunnels is powered off, there is a multi-second delay before the multicast stream changes to the standby RBridge. | |
| Condition: Issue when the RBridge responsible for multicast distribution is powered off or the ISL cables are physically disconnected. | |

| | |
|---|--|
| Defect ID: DEFECT000588610 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: Multi-VRF |
| Symptom: On removing VRF, clean up of ipv4 prefixes from EVPN database may lead to unexpected restart of routing component (ribmgr daemon) | |
| Condition: Removing VRF in a controller less IP Fabric environment | |

| | |
|---|--|
| Defect ID: DEFECT000588644 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VLAN - Virtual LAN |
| Symptom: Unwanted message on the console. | |
| Condition: When removing VLAN's under port-channel or under overlay. | |
| Recovery: | |

| | |
|--|--|
| Defect ID: DEFECT000588647 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.0.0 | Technology: AAA - Authentication, Authorization, and Accounting |
| Symptom: Allocated memory is not freed while log-in to the switch via web browser | |
| Condition: Login to the switch using web browser | |

| | |
|---|---|
| Defect ID: DEFECT000588730 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS6.0.2 | Technology: NETCONF - Network Configuration Protocol |
| Symptom: When querying the VDX netconf server an invalid yang model "ietf-netconf-notifications-ann" is advertised. | |
| Condition: This issue will show up when trying to view the mounted netconf capabilities for a VDX mounted with Extreme SDN Controller (BSC). | |

| | |
|---|---|
| Defect ID: DEFECT000588764 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: LAG - Link Aggregation Group |
| Symptom: Error RASLOG related to vLAGs is issued erroneously. The RASLOG indicates vLAG that spans Rbridges is connected to different end devices. The raslog is issued on Rbridges that are not connected to the vLAG and only occurs when running in new vlag-commit-mode disable configuration. | |
| Condition: Can occur when change vlag-commit-mode to disable when there are a large number of active vLAGs. | |

| | |
|---|---|
| Defect ID: DEFECT000588822 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS6.0.1 | Technology: TRILL - Transparent Interconnection of Lots of Links |
| Symptom: An ISL (Inter Switch Link) flap is seen on VDX6940. | |
| Condition: This can be seen due to un-handled internal memory parity error interrupts. | |

| | |
|---|--|
| Defect ID: DEFECT000588823 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS7.0.0 | Technology: IGMP - Internet Group Management Protocol |
| Symptom: MCASTSS process termination may be seen in standby partition of VDX8770-8. | |
| Condition: When lag or tunnel is configured on active and standby MM has just come up. When active MM sends a LAG dump, the standby may see MCASTSS process termination. | |
| Workaround: LAG configuration should not be done before send dump is complete. | |

| | |
|--|---|
| Defect ID: DEFECT000588837 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Configuration Fundamentals |
| Symptom: when min-link value is equal to number of online uplinks, the downlink will go lst down. | |
| Condition: When min-link configured is equal or greater than the number of online uplinks. | |
| Workaround: Config min-link with one less than the number of online uplinks | |
| Recovery: Removing min-link config will bring the downlink online. | |

| | |
|---|---|
| Defect ID: DEFECT000588918 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS6.0.1 | Technology: VCS Fabric |
| Symptom: Customer encountered an unexpected VDX6740 reload:RAS logs & stack trace for the reset as below: 2016/02/16-01:02:57, [SEC-1203], 795596, SW/0 Active, INFO, NINMUM03-DC1-R107-NSPL-RTR-049, Login information: Login successful via TELNET/SSH/RSH. IP Addr: A.B.C.D. 2016/02/16-01:03:13, [SEC-3022], 795597, SW/0 Active, INFO, NINMUM03-DC1-R107-NSPL-RTR-049, Event: logout, Status: success, Info: Successful logout by user [admin]. 2016/02/16-06:38:11, [HSL-1012], 795598, INFO, VDX6740T-1G, Interface lo is link up 2016/02/16-06:38:11, [HSL-1012], 795599, INFO, VDX6740T-1G, Interface eth0 is link up 2016/02/16-06:38:11, [HSL-1012], 795600, INFO, VDX6740T-1G, Interface eth1 is link up 2016/02/16-06:38:11, [HSL-1012], 795601, INFO, VDX6740T-1G, Processor reloaded - Software Fault:Kernel Panic. 2016/02/16-06:38:11, [HASM-1004], 795602, WARNING, VDX6740T-1G, The last reboot is due to Kernel Panic in kernel .Network OSCLI show support: Tue Feb 16 09:25:17 IST 2016 ***** | |
| Condition: When high rate of TFTP ip directed broadcast packets are sent destined to known subnets. | |

| | |
|--|--|
| Defect ID: DEFECT000589277 | |
| 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: VLAN is unprovisioned by BGP in the case of unexpected flapping of IMR EVPN routes from the peer node. | |
| Condition: Tunnel discovery option is enabled in BGP EVPN address-family. | |
| Workaround: Delete and re-create the provisioned VLAN. | |

| | |
|--|--|
| Defect ID: DEFECT000589286 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.2 | Technology: UDLD - Uni-Directional Link Detection |
| Symptom: Link of 1G Copper SFP comes up too early during the power-cycle on VDX 6940-144S | |
| Condition: Power-cycle on VDX 6940-144S with 1G Copper SFP | |

| | |
|---|--|
| Defect ID: DEFECT000589893 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.1 | Technology: Hardware Monitoring |
| Symptom: Request for Enhancement to optimize the fan speed to achieve better temperature distribution for the VDX 6740T & VDX6740T-1G switches | |
| Condition: Applies only to the VDX6740-T-R & VDX6740T-1G-R switches running port-side exhaust fans | |

| | |
|--|---|
| Defect ID: DEFECT000589911 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS5.0.2 | Technology: VCS Fabric |
| Symptom: Data loss is seen when an ISL port is flapped in a VCS that is employing VXLAN to connect to the remote data center VCS fabric. | |
| Condition: Flapping ISL link in a VCS fabric connecting to remote data center network using VXLAN/VTEP technology, would incur 1 to 2 seconds of data loss. | |

| | |
|--|--|
| Defect ID: DEFECT000589967 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS7.0.0 | Technology: IGMP - Internet Group Management Protocol |
| Symptom: error message seen on console while trying to configure Query-Interval on L3 physical interface Or PO. The queries will be generated at default time interval = 125 sec. | |
| Condition: This issue is seen when user try to configure Query-Interval on PO/Physical interface being in shut state. The config will not be applied as long as interface is in "protocol-down state" | |
| Workaround: Customer should bring the interface in "protocol up" state before applying Query-Interval config. Once the interface is up, Config will succeed. | |

| | |
|---|--|
| Defect ID: DEFECT000590101 | |
| Technical Severity: High | 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: Configuring remote-as or configuring "neighbor x.x.x.x capability additional path" under BGP VRF instance using Netconf can cause unexpected reload. | |
| Condition: When using Netconf to configure remote-as or configuring "neighbor x.x.x.x capability additional path" under the VRF instance and BGP VRF instance is not present then can cause unexpected reload. | |
| Workaround: For any configuration under BGP VRF instance using Netconf first configure BGP VRF instance then configure under BGP VRF instance. | |

| | |
|--|--|
| 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 fosecmd 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 "fosecmd" 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 untagged 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 a 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. | |

| | |
|--|------------------------------------|
| 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. <ol 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: <ol style="list-style-type: none"> 1. Configure multiple TACACS+/RADIUS/LDAP servers(max 5) 2. Remove few server entries 3. Add those servers entries back | |
| Workaround: Remove all Server entries and configure those servers back in the desired order. | |

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

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

| | |
|--|------------------------------------|
| Defect ID: DEFECT000600057 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.1 | Technology: Logical Chassis |
| Symptom: Switch might not rejoin the cluster when reloaded using 'fastboot' command. | |
| Condition: Reloading switch using 'fastboot' command on VDX6940 and VDX6740 platforms when SW1 partition is active might lead to this issue. | |
| Workaround: Reload the switch using 'reload' command which is more graceful way of reloading. | |
| Recovery: Bring the switch which failed to join the cluster to default configuration using command 'copy default-config startup-config'. On reload, switch rejoins the cluster and regains older configuration. | |

| | |
|--|---|
| Defect ID: DEFECT000600066 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: 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. | |

Closed With Code Changes for Network OS 7.2.0b

| | |
|--|------------------------------------|
| 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: DEFECT00060227 | |
| 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 can not 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"> - R Bridges - Ports (physical and virtual) - Devices that appear in the Name Server <p>The maximum assignable port indexes are listed here by platform type:</p> <table style="margin-left: 40px;"> <tr> <td>Chassis-based systems (Director class)</td> <td>: 1800</td> </tr> <tr> <td>VDX 6740/VDX 6740T/VDX 6740T-1G</td> <td>: 1200</td> </tr> <tr> <td>VDX 6940</td> <td>: 1312</td> </tr> </table> <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 | | Chassis-based systems (Director class) | : 1800 | VDX 6740/VDX 6740T/VDX 6740T-1G | : 1200 | VDX 6940 | : 1312 |
| Chassis-based systems (Director class) | : 1800 | | | | | | |
| VDX 6740/VDX 6740T/VDX 6740T-1G | : 1200 | | | | | | |
| VDX 6940 | : 1312 | | | | | | |
| 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 Make interface up again: 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 can not find particular port channel interface. | |
| Condition: This issue can occur during configuration file replay in which a port channel can be configured as tracking interface. | |

| | |
|---|--|
| Defect ID: DEFECT000615646 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: 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-MACfor multicast-mac address floods the packet after removing static-ARP and static-MAC entry and re-configuring. | |
| Condition: Static multicast MAC configured as static ARP. | |

| | |
|---|---|
| Defect ID: 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 a Extreme device. | |

| | |
|---|--|
| Defect ID: DEFECT000617919 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: 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. | |
| Condition: The issue is observed for below sequence of steps: 1. Configure OSPF authentication key on interface using CLI. 2. Save running configuration using command: copy running-config flash://<file-name> 3. Remove configured OSPF authentication key using CLI. 4. Replay saved configuration by using command: copy flash://<file-name> running-config It is observed that OSPF authentication key is not applied after step-4 though it was expected to be applied on the interface. | |
| Workaround: After config-replay fails to configure OSPF authentication key on the interface, it is possible to configure authentication key using CLI. | |

| | |
|---|---|
| Defect ID: DEFECT000620617 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: 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 Release: Network OS7.1.0 | Technology: VCS Fabric |
| Symptom: Any packet transmitted from CPU gets dropped on FC port. | |
| Condition: Can happen only on FC port. | |

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

| | |
|--|---|
| Defect ID: DEFECT000624701 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS6.0.0 | 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 Release: Network OS7.1.0 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: Multicast functionality daemon "PIMd" goes down with memory leak. | |
| Condition: On enabling PIM, for every 60 seconds there is a memory leak of 5K bytes. | |
| Workaround: Do not enable PIM on router. | |
| Recovery: Disable PIM on router and reboot the router. Do not enable PIM after reboot. | |

Closed without code changes for Network OS v7.2.0c

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change in Network OS v7.2.0c.

- NONE

Closed without code changes for Network OS v7.2.0b

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change in Network OS v7.2.0b.

- NONE

Closed without code changes for Network OS v7.2.0a

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of February 15th, 2018 in Network OS v7.2.0a.

- NONE

Closed without code changes for Network OS v7.2.0

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of July 10, 2017 in Network OS v7.2.0.

| | |
|-----------------------------------|-----------------------------------|
| Defect ID: DEFECT000472972 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |

Closed With Code Changes for Network OS 7.2.0b

| | |
|---|--|
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS3.0.1 | Technology: ARP - Address Resolution Protocol |
| Symptom: ARP Packet capture gets enabled for all VE interfaces even when the user has enabled it on a single VE interface. | |
| Condition: This issue is seen while enabling ARP PCAP on a single VE interface. | |

| | |
|---|--|
| Defect ID: DEFECT000510114 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS4.1.2 | Technology: VLAN - Virtual LAN |
| Symptom: In VDX 6740, when we have different load balancing scheme configured on the port channel, we see unexpected results with respect to load balance. | |
| Condition: If we have different load balancing schemes applied on VDX 6740, the latest configured value will take effect on the switch. | |
| Workaround: Use the same LB scheme for all PO in VDX 6740. | |
| Recovery: Re-configure the same LB scheme wherever required. | |

| | |
|--|--|
| Defect ID: DEFECT000546702 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS5.0.1 | Technology: ACLs - Access Control Lists |
| Symptom: When user tries to login with wrong credentials using default-vrf, debug messages are seen on console. | |
| Condition: When user tries to login with wrong credentials using default-vrf, debug messages are seen on console. | |

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

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

| | |
|---|---|
| Defect ID: DEFECT000568674 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: NOS5.0.2 | Technology: OpenStack Integration |
| Symptom: Customer would see references to IPv4 address when running the "ping ipv6" command. This is a cosmetic issue and won't affect the ping functionality. | |
| Condition: That would happen when running the "ping ipv6" command. | |

| | |
|--|---|
| Defect ID: DEFECT000577571 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS5.0.2 | Technology: IPv4 Multicast Routing |
| Symptom: Configuration: L3 PIM protocol enabled in a scaled topology with 760 sources and are learnt on an interfaces. In addition VRRPE is also enabled. Symptom: When the specific interface is disabled and enabled back, high CPU utilization is seen for PIM, MCASTSS daemons on the system. In addition, learning of new forwarding entries is delayed by 5 minutes. | |
| Condition: The PIM protocol is busy after the interface is re-enabled. The protocol is busy in processing the route updates within the system. | |
| Workaround: Do not disable the interface. | |
| Recovery: The system is stable 5 minutes after the interface is enabled. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000580922 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS5.0.2 | Technology: sFlow |
| Symptom: sFlow samples goes out of switch with SRC-IP as management IP instead of Inband IP configured. | |
| Condition: When 2 sFlow collectors are configured with same IP and different VRFs | |

| | |
|---|---|
| Defect ID: DEFECT000584172 | Technical Severity: Low |
| Reason Code: Design Limitation | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS5.0.2 | Technology: CLI - Command Line Interface |
| Symptom: When using special characters in password with the 'certutil import ssh' command, error message are thrown and it fails to configure. | |
| Condition: Special characters in password can cause the issue. | |
| Workaround: Please use back slash (\) when use special character in password. | |

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

| | |
|--|---|
| Defect ID: DEFECT000585008 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: When config apply error happens, user doesn't know which line of the config had the issue. | |
| Condition: Upon config replay on VDX devices. | |

| | |
|--|---|
| Defect ID: DEFECT000586790 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: Using RBridge range configuration command, even after BGP VRF instance is deleted, configuration under BGP VRF instance is allowed and can cause BGP daemon termination, HA failover, and or reboot of the switch. | |

Closed With Code Changes for Network OS 7.2.0b

| |
|---|
| Condition: Using RBridge range configuration command, BGP VRF instance is removed using "no address-family vrf .." command. And immediately after, without exiting from the configuration mode, another command applicable under (obsolete) BGP VRF instance configuration mode is issued. |
| Workaround: After removing the BGP VRF instance while using RBridge range command, exit the obsolete configuration mode using "top", "end", or "exit" commands. |

| | |
|---|--|
| Defect ID: DEFECT000589210 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS5.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP traps may not be received for SNMP-v1/v2/v3 hosts configured with IPv6 address. | |
| Condition: This issue is observed when IPv6 address is configured as trap recipient and VCS virtual IPv6 address is removed in the switch. | |
| Workaround: VCS virtual IPv6 should be configured to receive IPv6 traps. | |
| Recovery: VCS virtual IPv6 should be configured to receive IPv6 traps. | |

| | |
|---|--|
| Defect ID: DEFECT000591398 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS5.0.2 | Technology: IPv6 Addressing |
| Symptom: IPv6 ping timeout option did not work properly. | |
| Condition: Execution of IPv6 ping. | |

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

| | |
|---|--|
| Defect ID: DEFECT000594793 | Technical Severity: Medium |
| Reason Code: Not Reproducible | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.0.1 | Technology: Software Installation & Upgrade |
| Symptom: System may display: "qman_recovery_exit_local: DEBUG: the FQID 516 has dest_wq as chaqman_recovery_exit_local: DEBUG: the WQ lengths for pool channel of portal 1 on cpu1 are: 0:0:0:0:0:0:0" | |
| Condition: This bug appears when partitions are switched with heavy traffic. | |
| Recovery: Reboot the system. | |

| | |
|--|--|
| Defect ID: DEFECT000595199 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS6.0.2 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: Chassis disable may fail, when same is tried with scale configuration. | |
| Condition: When the scale configuration is present and chassis enable did not complete, subsequent chassis disable command may fail due to processing of time consuming events. | |

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

| | |
|---|---|
| Defect ID: DEFECT000596775 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS6.0.2 | Technology: CLI - Command Line Interface |
| Symptom: When the user configures IPv6 RA interval with the default value 600, the running-config shows the default RA value without suppressing it. | |
| Condition: The issue is seen by the user every time the RA interval is configured with the default value. | |

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

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000598248 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS7.0.1 | Technology: Port Mirroring |
| Symptom: Span on tunnel is not working after ha failover. | |
| Condition: Span on tunnel after ha failover | |
| Recovery: Unconfigure and configure back monitor session will resolve the issue. | |

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

| | |
|---|---|
| Defect ID: DEFECT000600230 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.0.1 | Technology: IP Fabric |
| Symptom: "show running-config rbridge-id evpn-instance <vni -name> vni add <vni-range>" throws an error message. | |
| Condition: Customer doing show running configuration with VNI range in EVPN instance. | |
| Workaround: Use the following command: "show running-config rbridge-id evpn-instance vni add" . | |

| | |
|-----------------------------------|-----------------------------------|
| Defect ID: DEFECT000606036 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |

| | |
|--|--|
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS6.0.2 | Technology: Software Installation & Upgrade |
| Symptom: System reload happen occasionally(very rare occurrence) at the time of firmware upgrade | |
| Condition: system reload could happen at the time of firmware upgrade in a switch having more no of user names and roles. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000607522 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: Access Gateway |
| Symptom: FFDC level log may be seen when loading config. | |
| Condition: This issue can be seen when a port is being monitored while the configuration is changing. | |

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

| | |
|---|------------------------------------|
| Defect ID: DEFECT000612521 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.1.0 | Technology: Logical Chassis |
| Symptom: Unexpected reload of switch | |
| Condition: Unexpected reload of switch while taking supportsave when ismd and ssmc core files present. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000612542 | Technical Severity: High |
| Reason Code: Design Limitation | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS7.0.1 | Technology: sFlow |
| Symptom: When IPv6 address is not configured on management port, sFlow sampled packets are sent OOB management interface with inband Ve interface IP as source IP. | |
| Condition: IPv6 address is not configured on management port. | |

| | |
|--|---|
| Defect ID: DEFECT000612933 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: Audit log is not updated with the user login/logout information | |
| Condition: This applicable only when the user accesses the device through the REST interface. | |

| | |
|---|--|
| Defect ID: DEFECT000615424 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS7.1.0 | Technology: MAC Port-based Authentication |
| Symptom: Interface does not come up after admin shut and no shut operation | |

Closed With Code Changes for Network OS 7.2.0b

| |
|---|
| Condition: Interface is enabled with MAC-authentication-bypass and host is not directly connected to switch. |
| Workaround: Remove the mac authentication configuration |

| | |
|--|---|
| Defect ID: DEFECT000615778 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS6.0.2 | Technology: Configuration Fundamentals |
| Symptom: snmp-server ? displays all Possible completions, here "view" display as "view Define an SNMPv2 MIB view" which is incorrect as it is also applicable to SNMP v3. | |
| Condition: snmp-server ? displays all Possible completions | |

| | |
|---|--|
| Defect ID: DEFECT000616966 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.0.1 | Technology: VLAN - Virtual LAN |
| Symptom: In extremely rare case, Kernel panic can be seen when VDX6940 is in idle state. | |
| Condition: This has been seen only once and has not been reproducible. The following config was running: | |
| <ul style="list-style-type: none"> - L2/L3 node with virtual-fabric turned on. - 200 vlans with ipv4 vrrp-e & 100 with v6. - In total box had 200 vlans. | |
| Recovery: Reboot the box. | |

| | |
|--|---|
| Defect ID: DEFECT000617284 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS6.0.2 | Technology: BFD - BiDirectional Forwarding Detection |
| Symptom: Unassociated BFD session for the IP address change operation may result in BFD session Down. | |
| Condition: BFD packet not reaching the system, resulting in BFD Session going Down. | |
| Workaround: BFD packet reception is affected on the system when the IP address on an unrelated interface is removed. Not definitive on the workaround for this problem. | |
| Recovery: BFD session recovers itself after going down and comes back to UP state. | |

| | |
|---|---|
| Defect ID: DEFECT000617830 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: VDX6940 takes longer to come up after a reload operation. | |
| Condition: This can happen when there are loads of configuration done on the switch. | |

| | |
|--|--|
| Defect ID: DEFECT000618553 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.1.0 | Technology: VLAN - Virtual LAN |
| Symptom: SPAN configuration is not successful with the error ""% Error: Destination port cannot have 802.1x configuration on it." | |
| Condition: Dot1x is configured and removed on an interface and now this interface is made as SPAN destination | |

| | |
|------------------------------------|-------------------------------------|
| Defect ID: DEFECT000620205 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |

Closed With Code Changes for Network OS 7.2.0b

| | |
|---|-----------------------------------|
| Reported In Release: NOS5.0.2 | Technology: Management GUI |
| Symptom: VDX-6740T Interface doesn't linkup as 1G by default and it comes up as 100Mb. | |
| Condition: VDX-6740T Interface linkup as 100Mb by default when other device has SEMI-CROSS LINK. | |
| Workaround: Configuration of speed 1000 on both side can make link 1G. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000620577 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS6.0.2 | Technology: Logical Chassis |
| Symptom: The output of "show interface description" command for port channel is not displayed in sorted order as per port channel interface number. | |
| Condition: The issue is seen in case of multi node cluster. | |

| | |
|--|---|
| Defect ID: DEFECT000620878 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.1.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: If disk is full due to too many core files, firmware download may not be successful on Draco-T device. | |
| Condition: Draco-T device already has limited disk space due to too many core files. | |
| Workaround: Delete old core files to free up space and fwdnld will be fine. | |
| Recovery: Delete old core files to free up space and fwdnld will be fine. | |

| | |
|---|--|
| Defect ID: DEFECT000621191 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.0.1 | Technology: Software Installation & Upgrade |
| Symptom: The standby GOS is unable to boot up during ISSU. | |
| Condition: It is due to a rare QMAN initialization issue during the GOS boot up process. | |
| Recovery: The switch will need to be rebooted for recovery. | |

| | |
|---|--|
| Defect ID: DEFECT000621696 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.0.1 | Technology: Software Installation & Upgrade |
| Symptom: Firmware download fails on standby GOS with error code 26. | |
| Condition: It can happen due to a rare network connectivity issue between the active and standby GOS partitions. | |
| Recovery: Firmware download will be aborted and filesystems will be recovered automatically. | |

| | |
|---|---|
| Defect ID: DEFECT000622864 | Technical Severity: High |
| Reason Code: Feature/Function Not Supported | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: Configuration Fundamentals |
| Symptom: Unexpected reload of switch while taking support save | |
| Condition: After diag test, chassis enable command will fail, switch may go through unexpected reload while taking supportsave | |
| Workaround: Reboot the switch after diag test before trying any other command | |
| Recovery: reboot the switch | |

| | |
|------------------------------------|---------------------------------|
| Defect ID: DEFECT000624075 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |

Closed With Code Changes for Network OS 7.2.0b

| | |
|---|------------------------------------|
| Reported In Release: NOS6.0.2 | Technology: Logical Chassis |
| Symptom: Disruptive firmware upgrade (coldboot) will fail. | |
| Condition: Number of SNMP communities associated with IPv4/IPv6 ACL configurations is greater than 20. | |
| Workaround: Limit the number of SNMP communities associated with IPv4/IPv6 ACL configurations to less than 20. | |

| | |
|--|---|
| Defect ID: DEFECT000624729 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: NOS7.1.0 | Technology: OpenStack Integration |
| Symptom: VDX 6940 40 GbE port may go offline after upgrade from NOS Release 7.0.0x to 7.1.0. | |
| Condition: VDX 6940 40 GbE port may go offline after firmware upgrade from NOS Release 7.0.0x to 7.1.0. | |
| Workaround: Perform shut / no shut on the problem 40 GbE port to bring the port back online. | |
| Recovery: Perform shut / no shut on the problem 40 GbE port to bring the port back online. | |

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

| | |
|---|--|
| Defect ID: DEFECT000627872 | Technical Severity: Medium |
| Reason Code: Design Limitation | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS6.0.2 | Technology: VLAN - Virtual LAN |
| Symptom: Link won't come online with Auto-negotiation enabled | |
| Condition: When the other end with 1G Intel NIC | |
| Recovery: VDX port should be in MASTER mode . Delivered the PEM script to do the same. | |

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

| | |
|--|--|
| Defect ID: DEFECT000630676 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS5.0.2 | Technology: Software Installation & Upgrade |
| Symptom: NOS BNA 14.0.1 and 14.0.2 throws the error "Firmware image download reboot operation has timed out", even the FW downgrade was successful. | |
| Condition: Firmware Download through BNA on VDX running in FC Cluster mode. | |

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

Condition: Execution of "show ip interface brief" CLI

| | |
|--|---|
| Defect ID: DEFECT000634013 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS7.1.0 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: Switch may crash with PIMD termination after executing "clear ip pim mcache" or "config/unconfig rp-address". | |
| Condition: With L3 multicast traffic running, execute "clear ip pim mcache" or "config/unconfig rp-address" in loop via script for more than 30 mins. | |

| | |
|---|---|
| Defect ID: DEFECT000637145 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.1.0 | Technology: IP Fabric |
| Symptom: Debug IPF script, ingress RT check against EVPN route is not matching in the script | |
| Condition: When debug IPF script is run on border leaf for L3 routed interface. | |

| | |
|--|---|
| Defect ID: DEFECT000638872 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: Switch throws false alarm, "Configuration not allowed when link speed set to 100Mbps" during configuration replay from external server | |
| Condition: when port-channel is configured on 100mb speed configured port | |

Closed without code changes for Network OS v7.1.0

This section lists software defects with Critical, High, and Medium Technical Severity closed without a code change as of November 22, 2016 in Network OS v7.1.0.

| | |
|---|--|
| Defect ID: DEFECT000386298 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS2.1.1_sp | Technology: NTP - Network Time Protocol |
| Symptom: Unexpected reload when SYN flood attack is exercised on the switch | |
| Condition: SYN flood attack on switch | |
| Workaround: Set the threshold for the burst rate of TCP traffic by using "tcp burstrate" command | |

| | |
|--|--|
| Defect ID: DEFECT000408109 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS3.0.0 | Technology: VRRPv2 - Virtual Router Redundancy Protocol Version 2 |
| Symptom: On VRRP session for ISL interface track command should not be allowed as this affects priority | |
| Condition: VRRP Track command for VRRP session on ISL port. | |
| Workaround: Dont enable VRRP track on ISL port. | |

| | |
|---|---|
| Defect ID: DEFECT000420768 | Technical Severity: Medium |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS3.0.0 | Technology: OpenStack Integration |
| Symptom: User may not be able to see the current session privileges when he is authenticated through AAA using "Show users". | |
| Condition: Show users command didn't display user's role when user is authenticated via RADIUS authentication. | |

| | |
|--|---|
| Defect ID: DEFECT000431087 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS3.0.0 | Technology: OpenStack Integration |
| Symptom: set "deny any" to telnet connection of ACL. By this setting, VDX denied all telnet connection. But the VDX didn't record raslog message of telnet-violation even if I connected to the VDX from denied PC over 2 times. Network OS 3.0.0a | |
| Condition: none | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000451282 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS4.0.0 | Technology: User Accounts & Passwords |
| Symptom: Changing the encryption level of a given username to level 7 with 'no service password-encryption' shows success RASLOG, but does not take effect in the config. | |
| Condition: When service password encryption is turned off, try to change existing users encryption level to 7. | |

| | |
|---|---|
| Defect ID: DEFECT000456601 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS4.0.0 | Technology: OSPF - IPv4 Open Shortest Path First |
| Symptom: "max-metric" config is cleared and not retained after some add/delete operations. | |
| Condition: On max metric configuration, if clear OSPF is executed and OSPF networks are removed and added again, max-metric config is cleared and is not retained. | |
| Workaround: Reconfigure max-metric after clear operation. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000458128 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS4.0.0 | Technology: Syslog |
| Symptom: Netconf login information not available in auditlog | |
| Condition: Applications logging into the switch using Netconf. | |

| | |
|--|--|
| Defect ID: DEFECT000519785 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS5.0.0 | Technology: AAA - Authentication, Authorization, and Accounting |
| Symptom: When "aaa authentication" command is tried with atleast one parameter same as previous command(example:aaa authentication radius local --> aaa authentication tacacs+ local), Customer may not be able to set the correct aaa mode. | |
| Condition: The authentication mode with primary & secondary sources of authentication cannot be updated to a configuration containing only the primary source and configuration with primary & secondary sources of authentication, the primary mode alone cannot be modified. | |
| Workaround: When "aaa authentication" command is tried with at least one parameter same as previous command(example: aaa authentication radius local --> aaa authentication tacacs+ local). Need to remove existing configuration and then configure the required configuration. example : (config)# do show running-config aaa authentication aaa authentication login radius local (config)# no aaa authentication login (config)# aaa authentication login tacacs+ local (config)# do show running-config aaa authentication aaa authentication login tacacs+ local | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000521573 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS5.0.0 | Technology: ACLs - Access Control Lists |
| Symptom: When copying a file to the running config that contains MAC ACL or IPv4 ACL commands, if some of the keywords are abbreviated, the commands may be treated as invalid, or the commands may execute more slowly than if they were not abbreviated. | |
| Condition: Abbreviating any of the keywords "mac access-list", "ip access-list", "seq", or, if "seq" is not present, "permit", "deny", or "hard-drop", would result in the slowness issue rather than the invalid command issue. The slowness issue would become more pronounced if the file being copied contains a large number of MAC ACL or IPv4 ACL commands. | |
| Workaround: Avoid using abbreviated keywords in files being copied to the running config. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000523640 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.0 | Technology: VMWare |
| Symptom: Traffic is not working through ports connected to servers after doing chassis disable/enable on the VDX. | |
| Condition: The issue is seen only after chassis disable followed by enable. | |
| Recovery: After chassis enable, Do a "shut" followed by "no shut" on the Port-Channels/Physical Interface connected to servers or do a "no port-profile-port" followed by "profile-port" on the Port-Channels/Physical interface. | |

| | |
|---|--|
| Defect ID: DEFECT000524630 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.0 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: During connection collision, switch is not closing the recent connection request if previous connection in established state. | |
| Condition: Second connection attempt is made after BGP peer is established. | |
| Recovery: The second connection request will get established automatically in case of collision. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000526209 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS4.1.3 | Technology: VMWare |
| Symptom: Series of De-Association failure messages may appear when a chassis disable followed by a chassis enable command is executed on the cluster. | |
| Condition: On a live vCenter connection, chassis disable followed by chassis enable is the condition under which this error messages may be observed. | |
| Workaround: Delete the vCenter entry before performing the chassis disable and add back the vCenter entry after performing the chassis enable command. | |
| Recovery: No recovery required, as there is no effect in the functionality of port-profile application. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000528408 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.1 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: New BGP connection is not accepted under following conditions. | |
| Condition: When the remote BGP identifier is greater than the local BGP identifier and there is a connection collision. | |

| | |
|--|---|
| Defect ID: DEFECT000529345 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.0 | Technology: Configuration Fundamentals |
| Symptom: User cannot see all the ISL port information in the port connectivity tab. | |
| Condition: Attempt to view ISL connection details from port connectivity tab. | |

| | |
|--|--|
| Defect ID: DEFECT000529743 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS4.1.3 | Technology: Software Installation & Upgrade |
| Symptom: LC may take longer to boot up if system is under heavy load. It will result in LC reboot timeout and FFDC. | |
| Condition: It happens during ISSU. | |
| Recovery: The user can run "power-off" and "power-on" to recover the LC. | |

| | |
|--|---|
| Defect ID: DEFECT000530965 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: High |
| 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: When the Helper Router exits helper mode on a given network segment, it should re-originate its LSAs based on the current state of its adjacency to the restarting router over the segment which is not happening. | |
| Condition: The helper router, on the GR mode exit, is supposed to originate the router/network LSAs based on the current adjacency state. This behavior is not yet implemented. | |

| | |
|---|--|
| Defect ID: DEFECT000533582 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS6.0.0 | Technology: AAA - Authentication, Authorization, and Accounting |
| Symptom: No uniformity in alignment and display format in the output of show access-list IP across interfaces | |
| Condition: Customers using the show access-list and expecting same display format across different interfaces. | |

| | |
|---|--|
| Defect ID: DEFECT000533953 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: When a SNMP Set request is attempted on the object ospfVirtIfRtrDeadInterval, this object can be set to a value greater than 65535. | |
| Condition: This issue is seen only on doing a SNMP Set request on the MIB object ospfVirtIfRtrDeadInterval. | |

| | |
|--|---|
| Defect ID: DEFECT000536442 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.0 | Technology: Configuration Fundamentals |
| Symptom: Management IP access list with large number of rules takes around 1 min to enforce the policy on management interface. | |
| Condition: Management IP access list with large number of rules. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000538887 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS4.1.2 | Technology: Syslog |
| Symptom: Under normal operations, following false-positive raslog message may be seen: <p style="text-align: center;">"[FW-3120], ..., WARNING, ... Interface<>,IFGViolation Errors, is above high boundary(High=100, Low=5). Current value is <> Error(s)/minute."</p> <p style="text-align: center;">There is no functional impact.</p> | |
| Condition: When there are no interface errors incrementing in the output of "show interface". | |

| | |
|---|---|
| Defect ID: DEFECT000540858 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS4.1.0 | Technology: NETCONF - Network Configuration Protocol |
| Symptom: Configuration and get config via NetConf are slow compared to CLI and sometime show high percentage of CPU utilization. | |
| Condition: Applicable for all netconf commands. | |

| | |
|--|--|
| Defect ID: DEFECT000541060 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.0 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: With conditional default-origination, default route should be generated only when the route-map matching prefix is present in the IP routing table. When irrespective of whether route-map matching prefix is present in the IP routing table or not, default route is originated to the neighbor. | |
| Condition: On configuring 'default-originate' with route-map in neighbor command | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|---|
| Defect ID: DEFECT000541202 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.0 | Technology: Configuration Fundamentals |
| Symptom: When customer configures high and low threshold values and the actual value is below the low threshold, a fabric watch raslog is displayed showing in between high and low threshold instead of below threshold. | |
| Condition: While using fabric watch module to monitor memory usage. | |

| | |
|---|---|
| Defect ID: DEFECT000545603 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.0 | Technology: Configuration Fundamentals |
| Symptom: Extended ACL with permit TCP rule does not block ICMP frames. | |
| Condition: Extended ACL with permit TCP rule does not block ICMP frames. | |
| Workaround: Add another rule to deny ICMP frames. | |
| Recovery: Add another rule to deny ICMP frames. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000547747 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.0 | Technology: VMWare |
| Symptom: Event notification is not received when port-groups with special characters are created on a vSwitch . As a result , the corresponding port-profile will not be applied to that interface. | |
| Condition: Use of special characters for port-groups on a vSwitch. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000548727 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: VMWare |
| Symptom: Port-Groups having the same name (regular , VMKernel , Distributed Port Groups) can get deleted when the vCenter user/administrator deletes any of the above port-groups sharing the same name, | |
| Condition: This scenario can happen when Port-groups belonging to different class(regular , VMKernel, Distributed Port Groups) use the same name within a data-center. | |
| Workaround: Avoid using the same name for different class of port-groups. vCenter's recommended names with prefixes like dvpkg , pg , VMKernel is good to follow. | |
| Recovery: Rename the port-groups from different class to use different names. | |

| | |
|---|---|
| Defect ID: DEFECT000548981 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: Configuration Fundamentals |
| Symptom: Access to the switch via management port (out of band) for all IPv6 protocols fails | |
| Condition: This issue is observed when a IPv6 ACL is configured on the management interface. For example "ipv6 permit any any" | |
| Workaround: Configure another matching entry in the ACL for permit based on protocol. Example "permit tcp any any" | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|---|
| Defect ID: DEFECT000550658 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: Momentary traffic loss is observed on NSX controller managed tunnels during VCS cluster formation. | |
| Condition: VCS has an overlay-gateway configuration with two or more RBridges attached and tunnel configurations are discovered by connecting to NSX controller. One of the RBridges attached to overlay-gateway is rebooted or upgraded via ISSU protocol. Traffic loss is observed on tunnels when such rebooted or upgraded RBridge rejoins VCS cluster. | |

| | |
|--|--|
| Defect ID: DEFECT000552520 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.1 | Technology: Static Routing (IPv4) |
| Symptom: Memory leak observed with repeated addition/deletion of VRFs using an automated script | |
| Condition: Adding and deleting VRFs repetitively | |
| Workaround: Remove routes before deleting the VRF | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000553066 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: Under certain condition, multiple "Invalid InterfaceId." log messages will be seen on the console. | |
| Condition: The symptom will be seen when a node joins the fabric and the fabric starts to rebuild. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000553426 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: "This command is not supported on this product." message displayed | |
| Condition: When trying to execute a dpod command on a VDX 8770 (principal node in VCS cluster) for a non-existent rbridge-id. | |

| | |
|---|--|
| Defect ID: DEFECT000554319 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: Switch does not generate a ColdStart trap on the VE interface configured in mgmt-vrf. | |
| Condition: When switch is configured with VE interface in mgmt-vrf, then we will observe this issue. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|---|
| Defect ID: DEFECT000554351 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: On the VDX8770 , executing certain Spanning Tree show commands may display the following warning message: "% Warning: Output Incomplete, VCS is in transient state" | |
| Condition: 1) Configure MSTP 2) Issue "show spanning-tree brief" or "show spanning-tree mst detail" CLI command | |
| Workaround: This is a cosmetic issue and these messages can be safely ignored. There is no impact to functionality. | |

| | |
|--|---|
| Defect ID: DEFECT000554472 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: Configuration Fundamentals |
| Symptom: When booting logical chassis with its default config, the principal Rbridge ends up with "system-monitor MM threshold" "down-threshold 0" while non-principal Rbridges end up with "system-monitor MM threshold" "down-threshold 2". | |
| Condition: This issue can be observed when using multiple VDX6740 switches in a logical chassis and booting from the default config. | |
| Recovery: Reconfigure the "system-monitor MM threshold" to make it sync. | |

| | |
|--|--|
| Defect ID: DEFECT000555059 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP walk should display Interfaces even when linecard is powered-off. | |
| Condition: When a linecard is powered off, the SNMP walk does not display and reference the interface. Fixed it as part of this defect. | |

| | |
|--|---|
| Defect ID: DEFECT000556025 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.1 | Technology: Configuration Fundamentals |
| Symptom: On port channel, Fabric Watch provides incorrect message SFP is absent when link is shut and SFP is not removed. | |
| Condition: On Chassis VDX switches, SFP info from Fabric Watch may mislead when port is shut. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000556553 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: The specific CLI command "no openflow" without specifying further parameters used to execute with default value instead of rejecting it. | |
| Condition: When "no openflow" was executed on CLI then it was getting accepted with first default value "enable". | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|---|
| Defect ID: DEFECT000558082 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.1 | Technology: RMON - Remote Network Monitoring |
| Symptom: when MAPS rules are triggered and MAPS is configured to generate e-mails, they are not generated | |
| Condition: MAPS has to be enabled with one of the default policies and e-mail action has to be enabled on the switch. Switch IP address needs to be configured as an IPv6 address. | |
| Workaround: Use IPv4 address for Switch IP if MAPs email action is desired | |

| | |
|--|--|
| Defect ID: DEFECT000558216 | Technical Severity: High |
| Reason Code: Design Limitation | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.1 | Technology: VRRPv2 - Virtual Router Redundancy Protocol Version 2 |
| Symptom: Attaching VE interface to another RBridge is taking more time. | |
| Condition: With more than 2K VE interfaces created, attaching a VE interface to another RBridge takes more 1 sec for each VE interface. | |

| | |
|--|--|
| Defect ID: DEFECT000558937 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS6.0.1 | Technology: AMPP - Automatic Migration of Port Profiles |
| Symptom: Sometimes, MAC addresses are shown twice in the output of 'show mac port-profile' command. | |
| Condition: 'show mac port-profile' command is issued | |
| Workaround: Re-issue to command to see a refreshed display | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000559023 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: The negation of command may be accepted without value if single value is present for the command | |
| Condition: "no" command will be accepted for single value command. For example - "no ip address" can remove existing IP address without specifying it on CLI. | |

| | |
|---|---|
| Defect ID: DEFECT000560092 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS6.0.1 | Technology: OpenStack Integration |
| Symptom: Unable to execute connector command due to config in error state. | |
| Condition: User did "copy default-config startup-config" on part of the cluster while the rest of the cluster is segmented. When cluster joined, it dynamically replayed back the connector config received from new principal at rejoin. This caused its connector configuration being partially saved to database on the local switch. | |
| Recovery: Recovery steps: <ol style="list-style-type: none"> 1. From principal switch, remove this secondary node from cluster by using "no vcs logical-chassis enable rbridge-id XX default-config" with XX being the secondary node RBridgeid. 2. On the secondary node, do "write erase" to clean up its config. 3. Restore the secondary switch vcs mode, vcsid and RBridgeId and the necessary ISL and see it rejoin the cluster. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000561651 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.0 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: Spelling of the word 'display' in 'rasman' command' help text should be corrected. | |
| Condition: Help text when 'rasman' command is executed. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000562427 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.2 | Technology: Logical Chassis |
| Symptom: When ISSU upgrade from 5.0.1b->5.0.2 with sflow collector config, does not allow to modify the existing configs. | |
| Condition: ISSU upgrade with sflow configs. | |
| Workaround: Remove the sflow config before the ISSU upgrade and the reconfig again. | |

| | |
|--|--|
| Defect ID: DEFECT000562672 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: When we try to associate 6th unique IPv4/IPv6 ACL to SNMP community/v3 user, we observe DB sync issue between front and backend. And also we observe unexpected reload of SNMP daemon when we try to associate 6th unique IPv4/IPv6 ACL to SNMP community/user. | |
| Condition: When we associate 6th unique IPv4/IPv6 ACL to SNMP community/user. | |

| | |
|--|--|
| Defect ID: DEFECT000563295 | Technical Severity: Medium |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.1 | Technology: ARP - Address Resolution Protocol |
| Symptom: With certain 3rd party devices, VDX learns ARP entries from a host belonging to a different subnet due to their GARP replies. This can potentially impact the traffic towards that source. | |
| Condition: When the device sends out GARP reply packet with source IP on different subnet than L3 interface IP. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000564701 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS6.0.1 | Technology: Syslog |
| Symptom: Syslog-server secure port option can not be used without using the use-vrf option. | |
| Condition: CLI usability issue where it shows use-vrf and secure options for syslog-server, where as secure option can be used along with use-vrf option only. Secure port can not be set alone for syslog server. | |

| | |
|---|---|
| Defect ID: DEFECT000565277 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: Configuration Fundamentals |
| Symptom: Under rare condition, when ACL is applied under management interface, it may not take effect and may display "Generic Error" when trying to remove the access-list. | |
| Condition: When applied under management interface | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|---|
| Defect ID: DEFECT000565590 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: Kernel panic results in switch reload after PVST configuration | |
| Condition: This memory corruption has only been observed in SWAT test with -ve test scenarios. | |
| Recovery: Node will automatically reboot. | |

| | |
|--|--|
| Defect ID: DEFECT000565913 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.1 | Technology: Software Installation & Upgrade |
| Symptom: Nodes of same cluster after reboot of all the nodes form cluster islands. | |
| Condition: This can happen when some nodes of a topology come up first and form their own cluster, and do not join the cluster formed by the rest of the nodes. | |

| | |
|--|--|
| Defect ID: DEFECT000567346 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: dot1qTpFdbTable of Q-BRIDGE-MIB does not give any output when queried using a SNMP Get/Get-Next or Get-Bulk request. | |
| Condition: This issue is specific to dot1qTpFdbTable of Q-BRIDGE-MIB. | |

| | |
|---|--|
| Defect ID: DEFECT000568362 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS5.0.2 | Technology: ACLs - Access Control Lists |
| Symptom: VDX running with Network OS5.0.2a experience unexpected reload due to pdm termination. | |
| Condition: Execution of the script (37 iterations) while config/unconfig the snmp community with ipv4 ACL and ipv6 ACL configured. | |

| | |
|--|--|
| Defect ID: DEFECT000568380 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: VLAN - Virtual LAN |
| Symptom: VDX running Network OS5.0.2x, [Pvlan]-Mac learning is not happening on community port when configured over port-channel. | |
| Condition: Mac learning. | |

| | |
|---|--|
| Defect ID: DEFECT000568854 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.1 | Technology: Static Routing (IPv4) |
| Symptom: Routing Information component (ribmgr) went through an unexpected restart | |
| Condition: Running soak tests with routing loop | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|--|
| Defect ID: DEFECT000569319 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: VLAN - Virtual LAN |
| Symptom: The MAC's gets learned on native VLAN for the shortest duration during RSTP root RBridge toggle. There will be loop during RSTP root RBridge toggle and the macs are getting classified to native VLAN causing the issue. The macs will be aged out after age time. | |
| Condition: The MAC's will be learned on native VLAN during RSTP toggle. There will be loop during RSTP root RBridge toggle and the macs are getting classified to native VLAN causing the issue. The MAC's will be aged out after age time. | |
| Workaround: Clearing FDB table will recover the mac tables. | |

| | |
|--|--|
| Defect ID: DEFECT000570086 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: NTP - Network Time Protocol |
| Symptom: Removal of NTP server is failing in the cluster. | |
| Condition: In rare conditions when multiple NTP servers are configured, removal of NTP server is failing. | |
| Workaround: Delete the NTP server second time if it failed on first time. | |

| | |
|---|--|
| Defect ID: DEFECT000570230 | Technical Severity: Medium |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: VLAN - Virtual LAN |
| Symptom: VDX running with Network OS5.x throws error messages of NOT A KNOWN Resource Id | |
| Condition: Execution of "switchport private-vlan trunk native vlan <ID>" | |

| | |
|---|---|
| Defect ID: DEFECT000570284 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS6.0.1 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: Root port of MSTP Instances will become Master port after adding a VLAN to an existing instances after MSTP-I convergence | |
| Condition: Addition of VLAN to an already converged MSTP Instances. | |
| Workaround: Shutdown/no shutdown of MSTP | |

| | |
|--|---|
| Defect ID: DEFECT000577094 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS5.0.2 | Technology: IPv4 Multicast Routing |
| Symptom: When a secondary IP is not configured, VDX sends option 24 in PIMV2 hello header which results in PIM adjacency failing to form on other switch. | |
| Condition: Happens with secondary IP address not configured. | |

| | |
|---|--|
| Defect ID: DEFECT000578640 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: Few millisecond packet loss observed while collecting supportsave when multiple data stream with high line rate is running. | |

Closed Without Code Changes for Network OS 7.1.0

Condition: Collecting supportsave when multiple stream of very high rate data is running.

| | |
|---|--------------------------------------|
| Defect ID: DEFECT000579288 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.1 | Technology: High Availability |
| Symptom: A Functional switch hardware faces an abrupt reboot | |
| Condition: Multiple Chassis Disable/Enable, repeated reboots of switch | |

| | |
|--|---|
| Defect ID: DEFECT000580044 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.1 | Technology: Configuration Fundamentals |
| Symptom: After installing a new POD license and reserving the port using DPOD commands, show interface output indicates "no 10-G Port Upgrade License" erroneously. | |
| Condition: When the user has not yet bounced the link (ie, "no shut") and recovered back to an active state after adding the license, the "No 10G Port Upgrade license" detail will remain displayed. | |

| | |
|---|---|
| Defect ID: DEFECT000581124 | Technical Severity: High |
| Reason Code: Design Limitation | Probability: Low |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.0 | Technology: OpenStack Integration |
| Symptom: 40G Interface is either protocol down (or) administratively down after "no shut" operation. | |
| Condition: Breakout Config operation performed on 40G Interface connected to a 40G Interface | |

| | |
|---|--|
| Defect ID: DEFECT000581852 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP query ifDescr output will be partial. This is specific to VDX6940-144S platform. | |
| Condition: SNMP walk for ifDescr & Ifname table. | |

| | |
|---|--|
| Defect ID: DEFECT000583151 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: panic will be observed when SS is initiated on a chassis with SFM(with HW issue which still has to be identified) | |
| Condition: issue will occur when SFM (with HW issue) is used | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|---|
| Defect ID: DEFECT000583324 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: FCoE - Fibre Channel over Ethernet |
| Symptom: ISL fails to come up due to Trunking Error | |
| Condition: When port is enabled between VDX8770 LC48x10G and VDX 6940 4x10G breakout interfaces | |
| Recovery: Issue shut followed by no shut on the port. | |

| | |
|---|---|
| Defect ID: DEFECT000583859 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: L2 Loop can occur in VDX running with Network OS5.0.2x | |
| Condition: "tunnel tagged-ieee-bpdu" is not working after removing "bpdu-drop" from the interface. | |

| | |
|--|--|
| Defect ID: DEFECT000585015 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: AMPP - Automatic Migration of Port Profiles |
| Symptom: traffic may flood for the non-profiled macs even if the global knob is disabled. | |
| Condition: 'no allow non-profiled-macs' is configured. | |
| Workaround: configure and remove 'allow non-profiled-macs' again. | |

| | |
|--|---|
| Defect ID: DEFECT000585352 | Technical Severity: Medium |
| Reason Code: Design Limitation | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: The VxLAN traffic drops when the underlay VLAG interface links go through the state change. It is further observed that VTEP is not learning the MAC addresses of hosts located across the VXLAN tunnel. | |
| Condition: The problem occurs when VLAG is VXLAN underlay network AND Loopback IP is configured as VTEP IP. | |
| Workaround: VLAG as undelay and Loopback IP as VTEP IP is not supported. Hence, please use VRRP(-E) IP as VTEP IP when underlay network is comprised of VLAG(s). | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000585934 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Newly joined VRRP-E session preempts current Master in the group. | |
| Condition: Configure VRRP-E sessions over GVLAN VE with preempt mode disabled. Perform chassis disable on Master router. After other router becomes Master, perform chassis enable on earlier Master router. This router, which was Master initially, again becomes Master. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|---|
| Defect ID: DEFECT000586338 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: IN VDX 6940-144S, link flap occurs on 40 G ISL ports if breakout configuration mis-matched with any adjacent ports. | |
| Condition: In VDX 6940-144S, a 40 G Port with breakout QSFP is not configured as breakout may cause adjacent 40 G port to flap, whereas its peer port is configured as 40G breakout. | |
| Workaround: In VDX 6940-144S, configure 40 G port as breakout if the peer port is configured as 40G breakout. After that, the link flap on the port will stop. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000586556 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: Logical Chassis |
| Symptom: Switch may panic after enabling maintenance-mode on M4 because of memory corruption. | |
| Condition: This can happen randomly and inconsistently | |
| Recovery: None. Switch will reboot and recover automatically. | |

| | |
|---|--|
| Defect ID: DEFECT000587120 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: Low |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.0.0 | Technology: ACLs - Access Control Lists |
| Symptom: When a permit rule is added for a specific IP(range), it may not block the other unspecified IP(ranges) | |
| Condition: This might occur for rules having source IP and gateway values. | |
| Workaround: Adding an additional rule to deny all other IP will block all other unspecified IP(ranges) | |

| | |
|---|---|
| Defect ID: DEFECT000587135 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.2 | Technology: TRILL - Transparent Interconnection of Lots of Links |
| Symptom: Unexpected reload on standby MM in rare scenario. | |
| Condition: While changing VCS. | |

| | |
|---|---|
| Defect ID: DEFECT000587566 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: Some VLANs that the user is expecting to be provisioned in a BGP EVPN instance get unprovisioned. | |
| Condition: The issue is seen when the user uses some common VNIs between BGP EVPN instance and some physical ports and then the user removed the last port from all the VLANs. | |
| Workaround: The user should avoid configuring common VNIs between BGP EVPN and physical ports. | |
| Recovery: Deleting those VLANs and configuring them back again, will fix the issue. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|--|
| Defect ID: DEFECT000587880 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: IPv6 Addressing |
| Symptom: IPv6 DHCP relay SOLICIT packets are not getting intercepted after coldboot upgrade. | |
| Condition: Running DHCP relay in a IP Fabric EVPN environment sometimes. | |
| Workaround: Delete and re-configure the same L3 interface where relay config is present. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000588078 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Virtual Fabrics |
| Symptom: The switch may get rebooted when performing copy scp command | |
| Condition: This can occur after powercycle of the switch. | |

| | |
|---|---|
| Defect ID: DEFECT000588346 | Technical Severity: High |
| Reason Code: Design Limitation | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: All tunnels go operationally down when VRF associated with overlay-gateway is deleted. | |
| Condition: The "ip interface loopback X" or "ip interface ve X vrrp-extended-group Y" is configured for the overlay-gateway and corresponding ve/loopback interface is bound to a non-default vrf. | |
| Recovery: Configure the IP address for the Ve/loopback interface associated with overlay-gateway and then configure "no activate" & "activate" for overlay-gateway. | |

| | |
|---|---|
| Defect ID: DEFECT000588355 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: FCoE - Fibre Channel over Ethernet |
| Symptom: Dcmd may not go through the graceful shutdown process when switch is rebooted. | |
| Condition: Under some rare condition, Dcmd cannot shut down its database and thus fails to go through the graceful shutdown process. | |
| Workaround: None, the system will be recovered automatically | |

| | |
|---|---|
| Defect ID: DEFECT000588682 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | Technology: TRILL - Transparent Interconnection of Lots of Links |
| Symptom: In a two node cluster after shut of the ISL port on VDX6740 side the other node in the 2 node VCS goes down; this issue is rarely seen. | |
| Condition: 2 nodes must be part of same VCS. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000589040 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.0.0 | Technology: ACLs - Access Control Lists |
| Symptom: When a permit/deny rule is added for a specific IP(range), it may not block the other unspecified IP(ranges) | |
| Condition: This might occur for rules having source IP and gateway values. | |
| Workaround: Adding an additional deny rule with same source IP and current switch IP as destination for IP protocol will block all other unspecified IP(ranges) | |

| | |
|---|--|
| Defect ID: DEFECT000589259 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.0.0 | Technology: ACLs - Access Control Lists |
| Symptom: Applying a policy having rule with DSCP option will fail internally. Hence no IP ACL rule will be enforced. | |
| Condition: Applying a policy having rule with DSCP option will fail with internally with a script error. | |
| Workaround: Remove/Avoid adding the rule having DSCP value/option | |

| | |
|---|---|
| Defect ID: DEFECT000589297 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: xSTP - Spanning Tree Protocols |
| Symptom: VDX running Network OS5.0.2 displays issue in STP output, root bridge is wrongly displayed. | |
| Condition: Execution of "spanning-tree ieee-bpdu limit-vlan-flood" command | |

| | |
|---|--|
| Defect ID: DEFECT000590114 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.1 | Technology: AMPP - Automatic Migration of Port Profiles |
| Symptom: If user configures two AMPP port-profiles, one of them configured with access VLAN x and other configured with trunk VLAN x, Then ,In that case , It will not be shown as conflicting in "show port-profile name <pp1-name> name <pp2-name> validate" command output. | |
| Condition: When user creates 2 port-profiles, one port-profile with access VLAN x and other port-profile with trunk VLAN x and executes "show port-profile name <pp1-name> name <pp2-name> validate" CLI. | |

| | |
|---|--------------------------------------|
| Defect ID: DEFECT000590771 | Technical Severity: Medium |
| Reason Code: Already Fixed in Release | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.1 | Technology: High Availability |
| Symptom: The standby MM on 8770-8 went to faulty state. This caused switch HA failover not to work. Also seen on primary MM was DCMD daemon termination, causing primary MM to reboot. | |
| Condition: The primary MM was booted with wrong Model ID, resulting in communication failure with secondary MM and database corruption. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|--|
| Defect ID: DEFECT000591172 | Technical Severity: Medium |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.1 | Technology: IP Addressing |
| Symptom: It is seen that configuration of global VE interface is missing in output of " show running-config". Configuration of same global VE was already present in protocol daemon, hence configuration of global VE again is not allowed. | |
| Condition: In a rare scenario during global VE configuration. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000594276 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.1 | Technology: Logical Chassis |
| Symptom: Under a high scale of VCS nodes, the configuration applied for a range of interfaces across the VDX nodes may cause principal node to encounter an unexpected reload. | |
| Condition: When issuing a configuration command under an interface range in a large cluster (32+ nodes) | |
| Workaround: Avoid using interface range option in large clusters (32+ nodes) & instead configure the interfaces individually. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000595226 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.1 | Technology: Syslog |
| Symptom: IPv4 and IPv6 syslog servers were not working when configured together as default/non-default VRF. | |
| Condition: Defect exists in 7.0.0 also. | |

| | |
|---|--|
| Defect ID: DEFECT000595610 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Traffic Management |
| Reported In Release: Network OS4.1.3 | Technology: Rate Limiting and Shaping |
| Symptom: Customer could not enable the multicast rate-limit CLI as it gives error indicating it is not supported for VDX6746 platform. | |
| Condition: multicast rate-limit CLI throws error. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000597786 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.1 | Technology: Logical Chassis |
| Symptom: After ISSU: IP Fabric VXLAN tunnels can start flapping continuously. | |
| Condition: When there exist BGP-EVPN based IP Fabric VXLAN tunnels in the system and an ISSU is performed. | |
| Workaround: Disabling the "vtep-discovery" defined under "address-family l2vpn evpn" submode of "router bgp" mode, before ISSU. or Deletion of the "overlay-gateway" config, before ISSU. | |
| Recovery: Deletion and recreation of the "overlay-gateway" config. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|--|
| Defect ID: DEFECT000598965 | Technical Severity: Medium |
| Reason Code: Design Limitation | Probability: Medium |
| 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: Local configuration related to global configuration may not restore on "config snapshot restore". | |
| Condition: Customer using snapshot feature may see issues when running "attached rbridge-id add 1" missing from running-config after "vcs config snapshot restore rbridge-id <rb-id> snapshot-id <snapshot-id>". | |
| Workaround: Customer should configure the missing configurations again. | |

| | |
|---|--|
| Defect ID: DEFECT000599203 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: The SNMP IPV4 traps may not be received through in-band interface. | |
| Condition: The SNMP traps may not be received through in-band interface after upgrade from 6.0.1 to 7.0.1. | |
| Workaround: Configure source-interface in the SNMP host / v3host recipients. | |

| | |
|--|--|
| Defect ID: DEFECT000599993 | Technical Severity: High |
| Reason Code: Not Reproducible | 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: Performing back-to-back reload of VDX may result in some of the OSPF sessions getting stuck in Exchange state | |
| Condition: This is observed with a high number of L3 routes of the magnitude ~20K and performing back-to-back reload of VDX | |
| Recovery: Clear the sessions that are in stuck in exchange state | |

| | |
|--|--|
| Defect ID: DEFECT000600197 | Technical Severity: Medium |
| Reason Code: Design Limitation | 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: "show running-config overlay-gateway <name> vlan <vlan number>" throws error "% No entries found" even when VLAN is present. | |
| Condition: This happens only when a filter is specified after "overlay-gateway <name>". Otherwise command works fine when no filter is specified. | |
| Workaround: Instead of using the filter, use the " include <string>" for filters like following: "show running-config overlay-gateway <name> include "vlan <vlan-number>" | |

| | |
|---|--|
| Defect ID: DEFECT000600482 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.2 | Technology: ARP - Address Resolution Protocol |
| Symptom: Inter-VLAN traffic that is routed on VDX is failing for specific end hosts. | |
| Condition: Frequent ARP addition/deletion, topology changes, high CPU bound ARP/L3 traffic | |
| Recovery: 'clear arp no-refresh' command | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|---|
| Defect ID: DEFECT000601318 | Technical Severity: High |
| Reason Code: Will Not Fix | Probability: Medium |
| 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: Blade may become faulty with error code 97 during firmware upgrade with the coldboot option. | |
| Condition: That can happen due to a rare MI/ISC issue during the blade initialization process. | |
| Workaround: None is required | |
| Recovery: The blade will be reset and will be recovered automatically. | |

| | |
|---|---|
| Defect ID: DEFECT000601398 | Technical Severity: High |
| Reason Code: Not Reproducible | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.0.1 | Technology: OpenStack Integration |
| Symptom: VDX 8770 6x100 GbE port, when shut, no shut, may cause link partner VDX 6940-144S 100 GbE port to go administratively down, and stay administratively down. | |
| Condition: VDX 8770 6x100 GbE port, when shut, no shut, may cause link partner VDX 6940-144S 100 GbE port to go administratively down, and stay administratively down. | |
| Workaround: Perform no shut on VDX 6940-144S 100 GbE port | |
| Recovery: Perform no shut on VDX 6940-144S 100 GbE port | |

| | |
|---|--------------------------------------|
| Defect ID: DEFECT000602153 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.2 | Technology: Inband Management |
| Symptom: No response for SNMP query from loopback IP address. | |
| Condition: If ingress and egress interfaces are different. | |
| Workaround: configure route to take ingress and egress path from same interface. | |
| Recovery: Use the same interface for ingress and egress side. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000602861 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.1 | Technology: Logical Chassis |
| Symptom: High disk usage that ended up out of space. | |
| Condition: Postgres log file(Dcmd.Linux.powerpc.pg_ctl.log) unconditionally growing | |
| Recovery: Delete Dcmd.Linux.powerpc.pg_ctl.log file. | |

| | |
|---|---|
| Defect ID: DEFECT000604917 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.2 | Technology: TRILL - Transparent Interconnection of Lots of Links |
| Symptom: VDX throws Application communication failure error. | |
| Condition: Execution of VCS ID change CLI. | |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|------------------------------------|
| Defect ID: DEFECT000613594 | Technical Severity: Low |
| Reason Code: Will Not Fix | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS5.0.2 | Technology: Logical Chassis |
| Symptom: show commands couldn't be accepted due to "application communication failure". | |
| Condition: Deletion of snmp-community config after ISSU upgrade from Network OS 502a to Network OS 502b1, can cause the issue of show command. | |
| Workaround: Remove the snmp community config before the upgrades and apply it back | |

| | |
|---|---|
| Defect ID: DEFECT000615168 | Technical Severity: Medium |
| Reason Code: Will Not Fix | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: CLI - Command Line Interface |
| Symptom: User may not be able to login to VDX switch after ISSU upgrade from Network OS7.0.0 to Network OS7.0.1 with no activation option. Existing telnet sessions will not be impacted. Cold boot upgrades are not impacted | |
| Condition: ISSU upgrade from Network OS7.0.0 to Network OS7.0.1 with no activation option | |
| Workaround: Perform ISSU firmware Install with auto activation. (Don't install firmware on a single node VDX running Network OS7.0.0 to Network OS7.0.1 with "noactivate" option. In case "logical-chassis" keyword used, don't install firmware on single or multiple nodes in cluster running Network OS7.0.0 to Network OS7.0.1 without "auto-activate" option.) Existing open telnet sessions will not be impacted. Thus if you really want to perform ISSU firmware install without activation, then keep principal node telnet/ssh/console session open with infinite terminal time out using "terminal timeout 0". This will allow to perform "firmware activate" to recover from the impacted state | |
| Recovery: The user may login using another node in the VCS that is not yet upgraded and carry out principal switchover to make that VDX a Principal switch. Once done, execute "firmware activate". Alternatively, add a new switch to VCS cluster, make it principal and run command "firmware activate" which would recover all switches in VCS cluster | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000625386 | Technical Severity: High |
| Reason Code: Already Fixed in Release | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.2 | Technology: Logical Chassis |
| Symptom: Unable to login to the device and customer should do netinstall to bring up the device | |
| Condition: After upgrading the firmware form version nos6.0.2 to nos7.0.0, user unable to login the switch. | |
| Workaround: Upgrade the firmware version from nos6.0.2 to nos6.0.2c before upgrading it to nos7.0.0 or higher version. | |
| Recovery: Netinstall is required to recover the switch state. | |

Known Issues for Network OS v7.2.0c

This section lists open software defects with Critical, High, and Medium Technical Severity as of June 5th, 2019 in Network OS v7.2.0c.

- NONE

Known Issues for Network OS v7.2.0b

This section lists open software defects with Critical, High, and Medium Technical Severity as of December 7, 2018 in Network OS v7.2.0b.

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

Known Issues for Network OS v7.2.0a

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

- NONE

Known Issues for Network OS v7.2.0

This section lists open software defects with Critical, High, and Medium Technical Severity as of July 10, 2017 in Network OS v7.2.0.

| | |
|--|--|
| Defect ID: DEFECT000517329 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS5.0.0 | Technology: BGP4 - IPv4 Border Gateway Protocol |
| Symptom: 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. | |

| | |
|--|--|
| Defect ID: DEFECT000577800 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS7.0.0 | Technology: MAPS - Monitoring and Alerting Policy Suite |
| Symptom: device connectivity config should be consistent on all the links in the port-channel | |

Closed Without Code Changes for Network OS 7.1.0

| |
|--|
| Condition: port-channel members configured as different type NAS, iSCSI |
| Workaround: Configure all members to be in same type. |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000581284 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme 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. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000584685 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS6.0.1 | Technology: VMWare |
| Symptom: Support save collected on the switch would not include vCenter specific outputs. | |
| Condition: Support save collected on the switch would not include vCenter specific outputs. | |
| Workaround: The same data can be obtained by dumping the sqllite database included as part of the supportsave. | |
| Recovery: There is no loss of functionality and hence no recovery. Same data is available in the sqllite database | |

| | |
|--|---|
| Defect ID: DEFECT000588886 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |

| | |
|---|---|
| Defect ID: DEFECT000596415 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.1.0 | Technology: ICMP - Internet Control Message Protocol |
| Symptom: 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. | |

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

| | |
|------------------------------------|------------------------------|
| Defect ID: DEFECT000596930 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |

Closed Without Code Changes for Network OS 7.1.0

| | |
|--|---|
| Reported In Release: NOS6.0.2 | 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. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000616434 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.1.0 | Technology: Logical Chassis |
| Symptom: In Cluster, during firmware upgrade, Principal Node may experience an unexpected reload. | |
| Condition: Principal and secondary nodes in the cluster are running different firmware versions. One the node is rebooted as a result of firmware upgrade. On the other node, at the same time user issued "vcs vesid <id> rbridge-id <id> " command. This sequence of events may cause this issue. | |

| | |
|--|--|
| Defect ID: DEFECT000616985 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |

| | |
|---|---|
| Defect ID: DEFECT000617251 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS6.0.2 | 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. | |
| Recovery: Recovery is to do one of the following - shut / no shut of the interface - un-config/ config of OSPF BFD. | |

| | |
|---|--|
| Defect ID: DEFECT000617700 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS7.0.1 | 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. | |
| Recovery: This is a cosmetic issue; no functional impact. | |

| |
|-----------------------------------|
| Defect ID: DEFECT000618254 |
|-----------------------------------|

Closed Without Code Changes for Network OS 7.1.0

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

| | |
|--|---|
| Defect ID: DEFECT000619146 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.0.1 | Technology: IP Fabric |
| Symptom: ISSU upgrade from 7.0.1 to NOS7.0.1a can cause some traffic loss if BFD is configured. | |
| Condition: ISSU upgrade to NOS7.0.1a when BFD is configured | |
| Workaround: BFD can be disabled during upgrade. | |

| | |
|---|---|
| Defect ID: DEFECT000621633 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.0.1 | Technology: OSPFv3 - IPv6 Open Shortest Path First |
| Symptom: Not able to change the IPv6 OSPF cost to 1 when auto-cost reference bandwidth is configured. | |
| Condition: The issue is observed for below sequence of steps: | |
| <ol style="list-style-type: none"> 1. Configure auto-cost for IPv6 OSPF using CLI: "auto-cost reference-bandwidth 100000" 2. Go to config-rbridge-Ve-<id> interface mode and configure OSPF cost using CLI: "ipv6 ospf cost 1" 3. Run show command to display interface OSPF parameters using CLI: "show ipv6 ospf in ve <id> rb <id>" <p>It is observed that cost field is not changed.</p> | |
| Workaround: Change the cost value to any non default-value and then back to default-value. | |

| | |
|--|--|
| Defect ID: DEFECT000621736 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS7.0.1 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: User may see the fewer audit logs instead of the exact count provided to the command "show logging audit-log reverse count". | |
| Condition: This issue might not happen always and the issue is seen in some of the nodes which are in cluster. | |

| | |
|--|--|
| Defect ID: DEFECT000622356 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.1.0 | Technology: IPv6 Addressing |
| Symptom: The running-configuration of port-channel interfaces have IPv6 nd cache expire time configured to 240 on an upgrade from 7.0.1 to 7.1.0, even though the user hasn't explicitly configured it. | |
| Condition: This issue is seen by the user whenever he upgrades the firmware from 7.0.1 to 7.1.0 | |

| | |
|--|---|
| Defect ID: DEFECT000623446 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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 | |

Closed Without Code Changes for Network OS 7.1.0

| |
|------------------------|
| in the BGP EVPN table. |
|------------------------|

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

| | |
|---|---|
| Defect ID: DEFECT000624566 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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 | |

| | |
|--|---|
| Defect ID: DEFECT000625402 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.1.0 | Technology: OSPF - IPv4 Open Shortest Path First |
| Symptom: OSPF authentication key configured on interface is getting lost after config-replay from backup configuration. | |
| Condition: The issue is observed if configurations are done in below order for RBridge sub-mode: | |
| <ol style="list-style-type: none"> 1. Configure router OSPF and create area 2. Create VE interface and go to its sub-mode 3. Configure authentication-key using CLI: "ip ospf authentication-key 2 <password>" 4. Save configuration using copy command 5. Delete VE interface 6. Run config-replay command using CLI: "copy flash://<file_name> running-config" <p>Admin can verify that configured authentication-key is lost by using command show running-config for the Ve interface.</p> | |
| Workaround: Configure the OSPF authentication key on the interface using CLI after config-replay is done. | |

| | |
|---|------------------------------|
| Defect ID: DEFECT000625616 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.1.0 | Technology: Metro VCS |
| Symptom: 10G ISL using tunable ZR optics (57-1000266-01) does not form between VDX6740 and VDX6940-144s after performing a single "shut/no shut" | |
| Condition: Performing "shut/no shut" on 10G ISL using tunable ZR optics (57-1000266-01) between VDX6740 and VDX6940-144s | |

| | |
|---|--|
| Defect ID: DEFECT000625831 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme 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. | |

Closed Without Code Changes for Network OS 7.1.0

Workaround: VxLAN tunnel is terminated on edge ports that are non-TRILL Ports.

| | |
|--|------------------------------------|
| Defect ID: DEFECT000625956 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme 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. | |

| | |
|---|--|
| Defect ID: DEFECT000626331 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.0.1 | Technology: VLAN - Virtual LAN |
| Symptom: User configured Vlan names are not displayed after reload of cluster in "show vlan br". It changes to default Vlan name | |
| Condition: Execution of "show vlan brief" CLI after reload. | |

| | |
|---|---|
| Defect ID: DEFECT000627564 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS6.0.2 | 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 | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000629684 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS5.0.2 | Technology: Management GUI |
| Symptom: Unexpected reload of standby management module in VDX8770. | |
| Condition: Reloading of standby management module without any user intervention. | |

| | |
|---|---|
| Defect ID: DEFECT000629838 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.0.1 | Technology: IP Fabric |
| Symptom: Traceroute for leaked route in vrf is not supported by Linux. | |
| Condition: Running traceroute for leaked route in vrf | |
| Workaround: Reachability can be performed by ping | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000631332 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS4.0.0 | Technology: Syslog |
| Symptom: Some Internal RAS log [Ex: BL-5282] are important and good to monitor those, but we don;t display internal RAS log on Console and we also don;t redirect them to syslog server. | |
| Condition: RAS log monitoring through Console or syslog. | |

| | |
|--------------------------------------|------------------------------------|
| Defect ID: DEFECT000631934 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS5.0.2 | Technology: Logical Chassis |

Closed Without Code Changes for Network OS 7.1.0

| |
|---|
| 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. |

| | |
|---|---|
| Defect ID: DEFECT00063313 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS5.0.2 | Technology: FCoE - Fibre Channel over Ethernet |
| Symptom: Changing fcoe advertisement interval not working sometimes | |
| Condition: 1. Change advertisement interval to a higher value than default 2. Reboot the switch 3. Keep alives from switch are still going out every 8 seconds | |

| | |
|--|--|
| Defect ID: DEFECT000634086 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.1.0 | Technology: IPv6 Addressing |
| Symptom: 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). | |

| | |
|--|---|
| Defect ID: DEFECT000634260 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS7.1.0 | Technology: Security Vulnerability |
| Symptom: Switch allows Non-admin user to execute certain show commands even though its denied by RBAC Rule. | |
| Condition: when user is moved from Admin role to Non-admin role after Ha Failover. | |

| | |
|--|---|
| Defect ID: DEFECT000634629 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS7.1.0 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: "BSR-candidate interface" and "RP-candidate interface" configuration is lost during configuration replay from external server. | |
| Condition: Configuration replay from external serve | |
| Recovery: Reconfigure after configuration replay | |

| | |
|--|---|
| Defect ID: DEFECT000634672 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS6.0.2 | 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 | |

| | |
|---|---|
| Defect ID: DEFECT000634913 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.0.1 | Technology: OSPF - IPv4 Open Shortest Path First |
| Symptom: If distribute-list is configured to filter out local connected routes and same external prefix is | |

Closed Without Code Changes for Network OS 7.1.0

| |
|---|
| advertised by multiple ASBRs to which there is there is no intra/inter area connectivity, then prefix learnt via one ASBR will be present in route table. |
| Condition: Distribute list is configured to filter out local connected routes and same external prefix received from multiple ASBRs. |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000636297 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.2.0 | Technology: Logical Chassis |
| Symptom: STP (PVST) protocol on vLAG interface would not be converged properly when bulk vlans are creating. Customer would see the Loop in network for new vlans created and customer also seen Portchannel (vLAG) stuck in DESGINATED role and LISTEN state forever. | |
| Condition: when user creating vlans in bulk say "vlan 2-128" with PVST protocol and vLAG is configued in VCS setip. | |
| Workaround: Do "shut" and "no shut" of vLAG interface to solve the issue. | |

| | |
|--|--|
| Defect ID: DEFECT000637037 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |

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

| | |
|---|---|
| Defect ID: DEFECT000638862 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: 100mb speed cli configured on physical interface in VDX6940-144S platform is not removed when downgraded to 7.1.0 | |
| Condition: 100mb speed configuration on VDX6940-144S platform | |
| Recovery: Remove manually after downgrade | |

| | |
|---|--|
| Defect ID: DEFECT000639033 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: NOS7.1.0 | Technology: RAS - Reliability, Availability, and Serviceability |
| Symptom: saving support save may fail some times | |
| Condition: VCS fabric with large numbers of nodes and support save is triggered for all the nodes. | |

| | |
|-----------------------------------|-------------------------|
| Defect ID: DEFECT000639398 | |
| Technical Severity: High | Probability: Low |

Closed Without Code Changes for Network OS 7.1.0

| | |
|---|--|
| Product: Extreme 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. | |

| | |
|--|---|
| Defect ID: DEFECT000639680 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.2.0 | Technology: IP Fabric |
| Symptom: When a user tries IPV6 prefix list config under rbridge range config mode and the config happens only on the principal node, but not on other nodes. | |
| Condition: IPV6 prefix list config does not happen on nodes other than principal node while the nodes being included in the rbridge range mode. | |
| Workaround: User can go to the specific node and do the same config. | |
| Recovery: User can go to the specific node and do the same config. | |

| | |
|--|--|
| Defect ID: DEFECT000639723 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS7.2.0 | Technology: ACLs - Access Control Lists |
| Symptom: Observing "Internal Error" error message, while enforcing acl configuration on management interface. | |
| Condition: Configured ACL names were similar to other enforced ACL name in "case-insensitive" scenario. | |
| Workaround: Can create ACL names which is not similar to existing ACL names by not only differentiating between capital and lower-case letters. | |

| | |
|--|---|
| Defect ID: DEFECT000639793 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |
| Recovery: User can go to the specific rbridge and can unconfigure the specific config. | |

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

| | |
|------------------------------------|--|
| Defect ID: DEFECT000640199 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |

Closed Without Code Changes for Network OS 7.1.0

| Reported In Release: NOS7.2.0 | Technology: VLAN - Virtual LAN | | | | | | | | | | | | |
|---|--|--------------------------|--------------------------|--------------------------|----------|---|-------------|-----------------------|--|---|-------------|--------------------------|--|
| Symptom: More than 1 protected vlan mapped to the same internal vlan, when the mod value results in to same when (protected vlan % 1024) is performed to pick a new internal isolated vlan. | | | | | | | | | | | | | |
| Condition: Above problem occurs since, available reserved internal vlans are only 1K and the protected-vlans can be anything in 7K range, since, input is 7K and the available o/p is only 1K, this results in to collision. | | | | | | | | | | | | | |
| Workaround: Release note so that only following set of vlans can be used for Protected Vlans to avoid same internal vlan mapping. | | | | | | | | | | | | | |
| | <table border="1"> <thead> <tr> <th>SI No</th> <th>Vlan Range (dot1q/gvlan)</th> <th>Internal Ivid Allocation</th> <th>Comments</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>3K - 3.5K-1</td> <td>(7168+0) – (7168+511)</td> <td></td> </tr> <tr> <td>2</td> <td>6.5K - 7K-1</td> <td>(7168+512) – (7168+1023)</td> <td></td> </tr> </tbody> </table> | SI No | Vlan Range (dot1q/gvlan) | Internal Ivid Allocation | Comments | 1 | 3K - 3.5K-1 | (7168+0) – (7168+511) | | 2 | 6.5K - 7K-1 | (7168+512) – (7168+1023) | |
| SI No | Vlan Range (dot1q/gvlan) | Internal Ivid Allocation | Comments | | | | | | | | | | |
| 1 | 3K - 3.5K-1 | (7168+0) – (7168+511) | | | | | | | | | | | |
| 2 | 6.5K - 7K-1 | (7168+512) – (7168+1023) | | | | | | | | | | | |
| Recovery: | | | | | | | | | | | | | |

| | |
|---|--|
| Defect ID: DEFECT000640460 | |
| Technical Severity: High | Probability: High |
| Product: Extreme 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. | |

| | |
|---|--|
| Defect ID: DEFECT000641475 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS7.0.1 | 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. | |

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

| | |
|--|---|
| Defect ID: DEFECT000641514 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS7.2.0 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: Multicast Source route would not get learnt on PIM router acting as Rendezvous Point (RP). May result in traffic loss for affect routes. | |
| Condition: Issue can be seen when multiple RP are present in network, and Priority value for non elected RP is updated such that it becomes newly elected RP. | |
| Recovery: Clearing affected routes from FHR router may recover the forwarding states | |

| | |
|--|---|
| Defect ID: DEFECT000641722 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.2.0 | Technology: VCS Fabric |
| Symptom: When upgrading an TOR switch with the coldboot option, when the standby GOS is booting up, the | |

Closed Without Code Changes for Network OS 7.1.0

| |
|---|
| customers may see the following message "Error happens on service instance chassis 0: command load failed or timed out (Critical)". |
| Condition: This is due to a corner case that causes delay in standby GOS booting up. |
| Recovery: The switch will be rebooted and recovered automatically. |

| | |
|--|--|
| Defect ID: DEFECT000641952 | |
| Technical Severity: High | Probability: High |
| Product: Extreme 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. | |

| | |
|---|--|
| Defect ID: DEFECT000642029 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.2.0 | Technology: VLAN - Virtual LAN |
| Symptom: When STP is enabled , traffic received on protected port is not egressing from uplink ROOT PORT | |
| Condition: Traffic not egressing from uplink port. | |
| Recovery: | |

| | |
|--|--|
| Defect ID: DEFECT000642115 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.1.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP Traps from principal node is received with source IP as VCS or chassis IP address instead of management IP address. | |
| Condition: SNMP traps enabled in chassis and VCS or chassis IP configured. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000642475 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.2.0 | Technology: Logical Chassis |
| Symptom: spanning tree Root port state moves to discard state when "ha failover" command is issued. also user would notice the traffic not forwarded on root port interface. | |
| Condition: when a break-out port of a switch in VCS cluster connected to root-bridge with rapid spanning tree protocol(RSTP) configured and followed by issuing "ha failover" . This issue would occur only with break-out port connected to root-bridge. | |
| Workaround: User recommended to use "shutdown" and "no shutdown" command on break-out to resolve the issue. | |

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

Closed Without Code Changes for Network OS 7.1.0

| |
|------------------------------|
| 'SR' SFP+ and the link is up |
|------------------------------|

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

| | |
|---|---|
| Defect ID: DEFECT000644067 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: NOS7.2.0 | Technology: IP Fabric |
| Symptom: PIM neighbor-ship between L3 interfaces over extended VLANs between leaf switches in IP fabric, may be lost or timed-out. | |
| Condition: Issue is only seen when IP Address of the PIM enabled L3 interface over VLAN is changed on one of the leaf acting as PIM neighbor. | |
| Workaround: By not modifying the IP address of the interface participating in PIM neighborhood between leaves, issue can be avoided. | |
| Recovery: Disabling and Enabled the interface admin state can recover the failed state. Reloading the router can also recover the failed state. | |

| | |
|--|--|
| Defect ID: DEFECT000644145 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: NOS7.2.0 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP MIB counters (a) ifOutUcastPkts (b) ifHCOutUcastPkts (c) ifHCInUcastPkts showing incorrect values | |
| Condition: When user sends Multicast/Broadcast L2 traffic, SNMP MIB counters (a) ifOutUcastPkts (b) ifHCOutUcastPkts (c) ifHCInUcastPkts showing incorrect values | |
| Workaround: User can use CLI to get the accurate values for (a) ifOutUcastPkts (b) ifHCOutUcastPkts (c) ifHCInUcastPkts | |

| | |
|---|--|
| Defect ID: DEFECT000644224 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |
| Recovery: | |

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

| | |
|---|--|
| Defect ID: DEFECT000644252 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme 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" | |

| | |
|--|--|
| Defect ID: DEFECT000644324 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.2.0 | Technology: VLAN - Virtual LAN |
| Symptom: Interface throws error as "Interface not in service" in a scaled configuration. | |
| Condition: Interface throws error as "Interface not in service" while trying to enable protected configuration on it with scaled configuration. | |
| Recovery: | |

| | |
|--|--|
| Defect ID: DEFECT000644331 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.2.0 | Technology: VLAN - Virtual LAN |
| Symptom: LC of VDX8770-8 goes to faulty sometimes when 1K VLANs are present on few ports and no VLANs on few ports of a line card. | |
| Condition: LC of VDX8770-8 goes to faulty while trying to un-configure and configure protected port configuration on all ports of a line card using range command with 1K VLANs on few ports and no VLANs at all on few ports and tried to execute 'no protected enable' and try to configure VLANs on an ISL port and again tried to execute 'no protected enable' multiple times. | |
| Recovery: | |

| | |
|--|---|
| Defect ID: DEFECT000644590 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS7.2.0 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: Multicast Source registration between FHR and RP may fail, and may result in traffic outage. | |
| Condition: Issue is seen only when the intermediate router between FHR and RP, is reloaded/rebooted. | |
| Recovery: Clearing affected multicast routes from FHR router may recover the failed state. | |

| | |
|--|--|
| Defect ID: DEFECT000644612 | |
| Technical Severity: High | Probability: High |
| Product: Extreme 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. | |
| <p>STEP 1. Configure Vp as primary, Vi as isolated, Vc as community vlan</p> <p>STEP 2. Associate Vp to Vi & Vc on primary vlan Vp.</p> <p>STEP 3. Configure A1A as trunk promiscuous port, A2A as trunk isolated, A3A as trunk community, A4A as trunk PVLAN port.</p> <p>STEP 4. Try enabling IGMP snooping on secondary vlans Vi & Vc.</p> <p>STEP 5. Enable PVST/RPVST globally.</p> <p>STEP 6. Try configuring bridge priority for vlan Vi & Vc</p> <p>STEP 7. Now disable spanning-tree globally on all nodes in cluster.</p> <p>STEP 8. Try creating ve interface corresponding to secondary VLANs Vi & Vc.</p> | |

Recovery:

| | |
|---|--|
| Defect ID: DEFECT000644663 | |
| Technical Severity: High | Probability: High |
| Product: Extreme 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. | |
| Recovery: | |

| | |
|--|--|
| Defect ID: DEFECT000644727 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme 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 DHCP configuration in a sequential order with ACL configuration. | |

| | |
|---|--|
| Defect ID: DEFECT000644836 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: NOS7.2.0 | Technology: VLAN - Virtual LAN |
| Symptom: Un-tagged traffic will be learnt on protected port, when native-vlan is enabled. | |
| Condition: Un-tagged traffic will be learnt on protected port, when native-vlan is enabled at interface level using 'no switchport trunk tag native-vlan'. | |
| Workaround: Apply protected port configuration on interface and enable native-vlan globally. | |
| Recovery: | |

| | |
|--|--|
| Defect ID: DEFECT000644854 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: NOS7.2.0 | Technology: ACLs - Access Control Lists |
| Symptom: Observing "Internal Error" message on ACL configuration. | |
| Condition: Configuring IPV6 ACL on management interface. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000645034 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000645061 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme 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. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000645175 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: NOS7.2.0 | Technology: Logical Chassis |
| Symptom: Sometimes DCMDdaemon terminates and system reboots when CLI command "show ip igmp groups" is issued in scaled scenario. | |
| Condition: a. Scaled up environment with large number of IGMP groups b. Firmware upgrade was in progress simultaneously. | |
| Workaround: Do not run cli " show ip igmp groups" while upgrading firmware with large number of IGMP groups. | |
| Recovery: Remove IGMP configurations after reboot. | |

| | |
|--|---|
| Defect ID: DEFECT000645336 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: NOS7.2.0 | Technology: PIM - Protocol-Independent Multicast |
| Symptom: Symptoms of this issue would include reload or reboot of the router, after a significant memory consumption. Reload of the router may affect the traffic forwarding. | |
| Condition: When PIM protocol is scaled upto near maximum supported number of routes, Null-Register packet periodic exchange between FHR and RP, is causing the memory leak on RP node and other nodes on RP tree. | |

| | |
|--|--|
| Defect ID: DEFECT000645359 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VPN |
| Reported In Release: NOS7.2.0 | Technology: EVPN - Ethernet VPN |
| Symptom: Multicast Source registration between FHR and RP may fail, and may result in traffic outage. | |
| Condition: Issue is seen only when the intermediate router between FHR and RP, is reloaded/rebooted. | |
| Recovery: Clearing affected multicast routes from FHR router may recover the failed state. | |

| | |
|--|---|
| Defect ID: DEFECT000645515 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: NOS7.2.0 | Technology: OSPFv3 - IPv6 Open Shortest Path First |
| Symptom: system may undergo unexpected reload after executing the command 'no ipv6 ospf cost' | |
| Condition: execution of the command 'no ipv6 ospf cost' | |

| | |
|--|---|
| Defect ID: DEFECT000645882 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme 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. | |

| | |
|--|--|
| Defect ID: DEFECT000646181 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme 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. | |

Closed Without Code Changes for Network OS 7.1.0

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.

| | |
|--|---|
| Defect ID: DEFECT000646314 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme 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 | |

Known Issues for Network OS v7.1.0

This section lists open software defects with Critical, High, and Medium Technical Severity as of November 22, 2016 in Network OS v7.1.0

| | |
|---|--|
| Defect ID: DEFECT000472972 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS3.0.1 | Technology: ARP - Address Resolution Protocol |
| Symptom: ARP Packet capture gets enabled for all VE interfaces even when the user has enabled it on a single VE interface. | |
| Condition: This issue is seen while enabling ARP PCAP on a single VE interface. | |

| | |
|---|--|
| Defect ID: DEFECT000510114 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS4.1.2 | Technology: VLAN - Virtual LAN |
| Symptom: In VDX 6740, when we have different load balancing scheme configured on the port channel, we see unexpected results with respect to load balance. | |
| Condition: If we have different load balancing schemes applied on VDX 6740, the latest configured value will take effect on the switch. | |
| Workaround: Use the same LB scheme for all PO in VDX 6740. | |
| Recovery: Re-configure the same LB scheme wherever required. | |

| | |
|--|--|
| Defect ID: DEFECT000517329 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.0 | 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. | |

| | |
|---|---|
| Defect ID: DEFECT000541449 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: Peer group configuration is not accepting update-source of IPv6 address | |
| Condition: Peer group configuration with update-source of IPv6 address | |

| | |
|--|--|
| Defect ID: DEFECT000543579 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.1 | Technology: IP Addressing |
| Symptom: Switch may reload due to low on memory when DHCP relay address and DHCP gateway CLIs are repeatedly used to configure and unconfigure. | |
| Condition: Repeated configure/un-configure of DHCP relay address and DHCP gateway CLI may lead to unexpected switch reload due to increased memory consumption. | |

| | |
|--|--|
| Defect ID: DEFECT000546702 | |
| Technical Severity: Low | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS5.0.1 | Technology: ACLs - Access Control Lists |
| Symptom: When user tries to login with wrong credentials using default-vrf, debug messages are seen on console. | |
| Condition: When user tries to login with wrong credentials using default-vrf, debug messages are seen on console. | |

| | |
|--|--|
| Defect ID: DEFECT000550982 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.1 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: Switch management port does not generate a ColdStart trap if a Management port is configured to acquire the IP address via DHCP. | |
| Condition: when switch is configured to acquire IP address via DHCP, then we will observe this issue. | |
| Workaround: If IP is configured statically, the issue will not happen. | |

| | |
|--|--|
| Defect ID: DEFECT000562214 | |
| 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: Source MAC may not get learnt on port channel configured with primary VLAN | |
| Condition: When Secondary VLANs, which are associated with other Primary VLAN are deleted | |

| | |
|---|---|
| Defect ID: DEFECT000568674 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS5.0.2 | Technology: OpenStack Integration |
| Symptom: Customer would see references to IPv4 address when running the "ping ipv6" command. This is a cosmetic issue and won't affect the ping functionality. | |
| Condition: That would happen when running the "ping ipv6" command. | |

| | |
|--|---|
| Defect ID: DEFECT000577571 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS5.0.2 | Technology: IPv4 Multicast Routing |
| Symptom: Configuration: L3 PIM protocol enabled in a scaled topology with 760 sources and are learnt on an interfaces. In addition VRRPE is also enabled. Symptom: When the specific interface is disabled and enabled back, high CPU utilization is seen for PIM, MCASTSS daemons on the system. In addition, learning of new forwarding entries is delayed by 5 minutes. | |
| Condition: The PIM protocol is busy after the interface is re-enabled. The protocol is busy in processing the route updates within the system. | |
| Workaround: Do not disable the interface. | |
| Recovery: The system is stable 5 minutes after the interface is enabled. | |

Known Issues for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000577800 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.0 | 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. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000580922 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS5.0.2 | Technology: sFlow |
| Symptom: sFlow samples goes out of switch with SRC-IP as management IP instead of Inband IP configured. | |
| Condition: When 2 sFlow collectors are configured with same IP and different VRFs | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000581284 | |
| Technical Severity: Low | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.0 | 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. | |

| | |
|---|---|
| Defect ID: DEFECT000584172 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: CLI - Command Line Interface |
| Symptom: When using special characters in password with the 'certutil import ssh' command, error message are thrown and it fails to configure. | |
| Condition: Special characters in password can cause the issue. | |
| Workaround: Please use back slash (\) when use special character in password. | |

| | |
|---|---|
| Defect ID: DEFECT000584534 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: A transport-vlan has been configured with multiple ctags on a port, however only a single ctag is seen in the command, "show vlan brief" | |
| Condition: This issue is seen when the following configuration is done in the following order. <ol style="list-style-type: none"> 1. configuring a vni of a vlan in the evpn-instance. 2. configuring the corresponding vlan as a transport-vlan (TVlan). <p>To avoid the issue, do the configuration in the order #2 & #1.</p> | |

| | |
|---|---|
| Defect ID: DEFECT000584634 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS6.0.2 | Technology: VCS Fabric |
| Symptom: 40G port will notice frequent online and offline events if one side is configured as breakout and the other side is not | |
| Condition: Failure to issue breakout on a QSFP 40G port, which is supposed to work in 4X10G mode. | |

| | |
|--|---|
| Defect ID: DEFECT000585008 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: When config apply error happens, user doesn't know which line of the config had the issue. | |
| Condition: Upon config replay on VDX devices. | |

| | |
|---|---|
| Defect ID: DEFECT000586790 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.0 | Technology: BGP4+ - IPv6 Border Gateway Protocol |
| Symptom: Using RBridge range configuration command, even after BGP VRF instance is deleted, configuration under BGP VRF instance is allowed and can cause BGP daemon termination, HA failover, and or reboot of the switch. | |
| Condition: Using RBridge range configuration command, BGP VRF instance is removed using "no address-family vrf .." command. And immediately after, without exiting from the configuration mode, another command applicable under (obsolete) BGP VRF instance configuration mode is issued. | |
| Workaround: After removing the BGP VRF instance while using RBridge range command, exit the obsolete configuration mode using "top" , "end", or "exit" commands. | |

| | |
|--|---|
| Defect ID: DEFECT000588886 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS7.0.0 | 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. | |

| | |
|---|--|
| Defect ID: DEFECT000589210 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: SNMP - Simple Network Management Protocol |
| Symptom: SNMP traps may not be received for SNMP-v1/v2/v3 hosts configured with IPv6 address. | |
| Condition: This issue is observed when IPv6 address is configured as trap recipient and VCS virtual IPv6 address is removed in the switch. | |
| Workaround: VCS virtual IPv6 should be configured to receive IPv6 traps. | |
| Recovery: VCS virtual IPv6 should be configured to receive IPv6 traps. | |

Known Issues for Network OS 7.1.0

| | |
|---|--|
| Defect ID: DEFECT000591398 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.2 | Technology: IPv6 Addressing |
| Symptom: IPv6 ping timeout option did not work properly. | |
| Condition: Execution of IPv6 ping. | |

| | |
|---|--|
| Defect ID: DEFECT000592597 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.0 | Technology: Software Installation & Upgrade |
| Symptom: Allowing for N+2 version upgrade with default config. | |
| Condition: It's a RFE to allow upgrade of N+2 version. | |

| | |
|---|---|
| Defect ID: DEFECT000592879 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: Configuration Fundamentals |
| Symptom: After LC power on/off in VDX8770, uplink interfaces from the LC are missing on show track summary output. | |
| Condition: When Link State Tracking (LST) configuration is present on a linecard, after slot power off/on the uplink configuration will be lost. | |
| Workaround: Uplinks need to be reconfigured again after slot power on. | |

| | |
|--|--|
| Defect ID: DEFECT000593537 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.2 | Technology: IP Addressing |
| Symptom: Host ARP is learnt even when host IP subnet does not match to VE IP subnet. | |
| Condition: Host is connected to a VLAN where the Ve IP subnet is different than the host IP subnet. | |
| Workaround: Disable proxy ARP on VE | |

| | |
|---|--|
| Defect ID: DEFECT000594793 | |
| Technical Severity: Medium | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: Software Installation & Upgrade |
| Symptom: System may display: "qman_recovery_exit_local: DEBUG: the FQID 516 has dest_wq as chaqman_recovery_exit_local: DEBUG: the WQ lengths for pool channel of portal 1 on cpu1 are: 0:0:0:0:0:0:0" | |
| Condition: This bug appears when partitions are switched with heavy traffic. | |
| Recovery: Reboot the system. | |

Known Issues for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000595199 | |
| 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: Chassis disable may fail, when same is tried with scale configuration. | |
| Condition: When the scale configuration is present and chassis enable did not complete, subsequent chassis disable command may fail due to processing of time consuming events. | |

| | |
|---|---|
| Defect ID: DEFECT000596415 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.1.0 | 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. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000596658 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.0.1 | Technology: Logical Chassis |
| Symptom: Traffic getting dropped indefinitely after reload. | |
| Condition: Due to /32 route functionality the packets are getting trapped twice (on local and remote leaf). | |

| | |
|--|--|
| Defect ID: DEFECT000596774 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS5.0.2 | Technology: IP Addressing |
| Symptom: Switch reloads with termination of ribmgr daemon | |
| Condition: Static route is leaked to multiple VRF's | |
| Workaround: Do not configure a static route more than once with the next-hop belonging to different VRF's | |

| | |
|---|---|
| Defect ID: DEFECT000596775 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.2 | Technology: CLI - Command Line Interface |
| Symptom: When the user configures IPv6 RA interval with the default value 600, the running-config shows the default RA value without suppressing it. | |
| Condition: The issue is seen by the user every time the RA interval is configured with the default value. | |

Known Issues for Network OS 7.1.0

| | |
|--|---|
| Defect ID: DEFECT000596930 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.2 | 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. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000598248 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.1 | Technology: Port Mirroring |
| Symptom: Span on tunnel is not working after ha failover. | |
| Condition: Span on tunnel after ha failover | |
| Recovery: Unconfigure and configure back monitor session will resolve the issue. | |

| | |
|---|---|
| Defect ID: DEFECT000600230 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.1 | Technology: IP Fabric |
| Symptom: "show running-config rbridge-id evpn-instance <vni -name> vni add <vni-range>" throws an error message. | |
| Condition: Customer doing show running configuration with VNI range in EVPN instance. | |
| Workaround: Use the following command: "show running-config rbridge-id evpn-instance vni add" . | |

| | |
|--|---|
| Defect ID: DEFECT000600233 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Traffic Management |
| Reported In Release: Network OS7.1.0 | Technology: QoS - Quality of Service |
| Symptom: A red profile and a non-conflicting policy doesn't co-exists | |
| Condition: red profile and non-conflicting policy are tried to apply to same interface. | |

| | |
|--|--|
| Defect ID: DEFECT000600385 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VLAN - Virtual LAN |
| Symptom: Duplicate ARP entries are observed. | |
| Condition: This can happen after an ISSU upgrade and a new IP address is allocated via DHCP for a connected host. | |
| Workaround: Execute "clear arp ip <IP address>" for the old IP address of host. | |
| Recovery: Execute "clear arp ip <IP address>" for the old IP address of host. | |

| | |
|--|---|
| Defect ID: DEFECT000601293 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: COS Priority tag frames egressed as Untagged frames | |
| Condition: Over VxLAN tunnel COS Priority tag frames are egressed as untagged frames. | |

| | |
|---|---|
| Defect ID: DEFECT000602319 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS5.0.2 | Technology: LAG - Link Aggregation Group |
| Symptom: BPDU packets creates loop in network and the network can become unstable. | |
| Condition: Network OS5.x with Protected vLAG BPDU received on Backup PO were flooded out from Active PO. | |

| | |
|--|--|
| Defect ID: DEFECT000606036 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS6.0.2 | Technology: Software Installation & Upgrade |
| Symptom: System reload happen occasionally(very rare occurrence) at the time of firmware upgrade | |
| Condition: system reload could happen at the time of firmware upgrade in a switch having more no of user names and roles. | |

| | |
|--|-------------------------------------|
| Defect ID: DEFECT000607522 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.1.0 | Technology: Access Gateway |
| Symptom: FFDC level log may be seen when loading config. | |
| Condition: This issue can be seen when a port is being monitored while the configuration is changing. | |

| | |
|---|---|
| Defect ID: DEFECT000609410 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.1.0 | Technology: OpenStack Integration |
| Symptom: When using 10 GbE CAT-5 twisted-pair SFP cabled between VDX 8770 and 6940 switches, the link may exhibit CRC errors and/or not form ISL link if the SFP is removed, then inserted repeatedly in quick succession. | |
| Condition: When using 10 GbE CAT-5 twisted-pair SFP cabled between VDX 8770 and 6940 switches, the link may exhibit CRC errors and/or not form ISL link if the SFP is removed, then inserted repeatedly in quick succession. | |
| Workaround: When inserting 10 GbE CAT-5 twisted-pair SFP into either side of VDX 8770 or 6940, please wait 20 seconds before removing and re-inserting the module. 10 GbE CAT-5 twisted-pair SFP links, in general, require more time to stabilize the link than other 10 GbE media. Re-insert the module if needed to recover from CRCs or failed ISL linkup. | |
| Recovery: When inserting 10 GbE CAT-5 twisted-pair SFP into either side of VDX 8770 or 6940, please wait 20 seconds before removing and re-inserting the module. 10 GbE CAT-5 twisted-pair SFP links, in general, require more time to stabilize the link than other 10 GbE media. Re-insert the module if needed to recover from CRCs or failed ISL linkup. | |

Known Issues for Network OS 7.1.0

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000609676 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.1.0 | Technology: Management GUI |
| Symptom: Raslog for thermal shutdown cancellation may indicate an incorrect temperature. | |
| Condition: Extensive changes were made to the thermal policy in this release. | |
| Workaround: This is cosmetic only. | |

| | |
|---|---|
| Defect ID: DEFECT000610251 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS4.1.3 | Technology: ICMP - Internet Control Message Protocol |
| Symptom: Dcmd process termination can occur. | |
| Condition: A script which launches multiple simultaneous "copy running-config <file>" operations can trigger the Dcmd process to terminate. Manually invoking simultaneous operations will not hit the small time window achievable by a script. | |
| Workaround: Ensure that the script does not invoke multiple simultaneous "copy running-config <file>" operations to a given switch. | |

| | |
|---|--|
| Defect ID: DEFECT000611303 | |
| Technical Severity: Medium | Probability: Low |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS5.0.1 | Technology: AMPP - Automatic Migration of Port Profiles |
| Symptom: Unexpected reload. | |
| Condition: After configuring vCenter and enabling CDP on ESXi vSwitch, due to very mild memory leak. | |

| | |
|---|---|
| Defect ID: DEFECT000611625 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.1.0 | Technology: OpenStack Integration |
| Symptom: VDX 6740 4x10 GbE port may go offline after firmware upgrade to Network OS Release 7.1.0 due to unstable 4x10 GbE link. User may see FFDC excessive interrupts on the problem port resulting in the port going offline. | |
| Condition: VDX 6740 4x10 GbE port may go offline after firmware upgrade to Network OS Release 7.1.0 due to unstable 4x10 GbE link. User may see FFDC excessive interrupts on the problem port resulting in the port going offline. | |
| Workaround: Perform shut / no shut to bring the port back online. | |
| Recovery: Perform shut / no shut to bring the port back online. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000612521 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.1.0 | Technology: Logical Chassis |
| Symptom: Unexpected reload of switch | |
| Condition: Unexpected reload of switch while taking supportsave when ismd and ssmc core files present. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000612542 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.1 | Technology: sFlow |
| Symptom: When IPv6 address is not configured on management port, sFlow sampled packets are sent OOB management interface with inband Ve interface IP as source IP. | |
| Condition: IPv6 address is not configured on management port. | |

| | |
|--|---|
| Defect ID: DEFECT000612933 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: Audit log is not updated with the user login/logout information | |
| Condition: This applicable only when the user accesses the device through the REST interface. | |

| | |
|---|--|
| Defect ID: DEFECT000615424 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.1.0 | Technology: MAC Port-based Authentication |
| Symptom: Interface does not come up after admin shut and no shut operation | |
| Condition: Interface is enabled with MAC-authentication-bypass and host is not directly connected to switch. | |
| Workaround: Remove the mac authentication configuration | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000616434 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.1.0 | Technology: Logical Chassis |
| Symptom: In Cluster, during firmware upgrade, Principal Node may experience an unexpected reload. | |
| Condition: Principal and secondary nodes in the cluster are running different firmware versions. | |
| <p>One the node is rebooted as a result of firmware upgrade. On the other node, at the same time user issued "vcs vcsid <id> rbridge-id <id> " command.</p> <p>This sequence of events may cause this issue.</p> | |

| | |
|---|--|
| Defect ID: DEFECT000616966 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.0.1 | Technology: VLAN - Virtual LAN |
| Symptom: In extremely rare case, Kernel panic can be seen when VDX6940 is in idle state. | |
| Condition: This has been seen only once and has not been reproducible. The following config was running: | |
| <ul style="list-style-type: none"> - L2/L3 node with virtual-fabric turned on. - 200 VLANs with ipv4 vrrp-e & 100 with v6. - In total box had 200 VLANs. | |
| Recovery: Reboot the box. | |

Known Issues for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000616985 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.1.0 | 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. | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000617058 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.1.0 | Technology: Sysmon |
| Symptom: Sometimes hardware will reboot during firmware upgrade, Downgrade or HA failover | |
| Condition: So far seen only on one hardware. Mostly seen during firmware upgrade, Downgrade or HA failover | |
| Recovery: Will recover after reboot | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000617197 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.1.0 | Technology: Logical Chassis |
| Symptom: Warning logs with tag 'FFDC' is displayed during cluster wide reload. | |
| Condition: Cluster reload. | |

| | |
|---|---|
| Defect ID: DEFECT000617251 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.2 | 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: <ul style="list-style-type: none"> - shut / no shut of the interface - un-config/ config of OSPF BFD. | |
| Recovery: Recovery is to do one of the following <ul style="list-style-type: none"> - shut / no shut of the interface - un-config/ config of OSPF BFD. | |

| | |
|--|---|
| Defect ID: DEFECT000617284 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS6.0.2 | Technology: BFD - BiDirectional Forwarding Detection |
| Symptom: Unassociated BFD session for the IP address change operation may result in BFD session Down. | |
| Condition: BFD packet not reaching the system, resulting in BFD Session going Down. | |
| Workaround: BFD packet reception is affected on the system when the IP address on an unrelated interface is removed. Not definitive on the workaround for this problem. | |
| Recovery: BFD session recovers itself after going down and comes back to UP state. | |

Known Issues for Network OS 7.1.0

| | |
|---|--|
| Defect ID: DEFECT000617700 | |
| 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: "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. | |
| Recovery: This is a cosmetic issue; no functional impact. | |

| | |
|---|---|
| Defect ID: DEFECT000617830 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.1.0 | Technology: CLI - Command Line Interface |
| Symptom: VDX6940 takes longer to come up after a reload operation. | |
| Condition: This can happen when there are loads of configuration done on the switch. | |

| | |
|--|---|
| Defect ID: DEFECT000617887 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.0.1 | Technology: IP Fabric |
| Symptom: On upgrade from 6.x to 7.x, one of the python CLI libraries may not be carried forward with all the changes & might impact some of the python scripts. | |
| Condition: Under certain unknown condition, when upgrading from 6.x to 7.x. | |
| Recovery: Copy CLI.py file manually to restore the script function. | |
| <p>After restoration, test using below CLI & it should appear as below with "splitlines()" instead of "split()":</p> <pre>sw0:FID128:root> grep -A 2 get_output /etc/fabos/Dcmd/python/CLI.py def get_output(self): return (self.output.splitlines())</pre> | |

| | |
|---|--|
| Defect ID: DEFECT000618052 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Monitoring |
| Reported In Release: Network OS7.0.1 | Technology: Hardware Monitoring |
| Symptom: Incorrect units displayed in the o/p of the command, "show media" for Lumentum SFP. | |
| Condition: In the o/p of "show media" for Lumentum SFP, the units for Wavelength field should be displayed as "units 0.025nm" instead of "units nm". | |

| | |
|--|--|
| Defect ID: DEFECT000618553 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.1.0 | Technology: VLAN - Virtual LAN |
| Symptom: SPAN configuration is not successful with the error ""% Error: Destination port cannot have 802.1x configuration on it." | |
| Condition: Dot1x is configured and removed on an interface and now this interface is made as SPAN destination | |

| | |
|--|--|
| Defect ID: DEFECT000619578 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.1.0 | Technology: ACLs - Access Control Lists |
| Symptom: Unexpected reload. | |
| Condition: When below command is used for viewing the enforced IP ACL's: "show access-list interface Management <interface id> in " Example: "show access-list interface Management 1/0 in" | |
| Workaround: Use below commands to view the enforced policy ' show running config ip(v6) access-list' 'show running-config interface management' | |

| | |
|---|-------------------------------------|
| Defect ID: DEFECT000620205 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS5.0.2 | Technology: Management GUI |
| Symptom: VDX-6740T Interface doesn't linkup as 1G by default and it comes up as 100Mb. | |
| Condition: VDX-6740T Interface linkup as 100Mb by default when other device has SEMI-CROSS LINK. | |
| Workaround: Configuration of speed 1000 on both side can make link 1G. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000620577 | |
| Technical Severity: Low | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.2 | Technology: Logical Chassis |
| Symptom: The output of "show interface description" command for port channel is not displayed in sorted order as per port channel interface number. | |
| Condition: The issue is seen in case of multi node cluster. | |

| | |
|--|---|
| Defect ID: DEFECT000620878 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.1.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: If disk is full due to too many core files, firmware download may not be successful on Draco-T device. | |
| Condition: Draco-T device already has limited disk space due to too many core files. | |
| Workaround: Delete old core files to free up space and fwdnld will be fine. | |
| Recovery: Delete old core files to free up space and fwdnld will be fine. | |

| | |
|---|--|
| Defect ID: DEFECT000621191 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: Software Installation & Upgrade |
| Symptom: The standby GOS is unable to boot up during ISSU. | |
| Condition: It is due to a rare QMAN initialization issue during the GOS boot up process. | |
| Recovery: The switch will need to be rebooted for recovery. | |

Known Issues for Network OS 7.1.0

| | |
|---|--------------------------------------|
| Defect ID: DEFECT000621402 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: Inband Management |
| Symptom: Telnet access to VDX is blocked via default-VRF and user defined VRF. | |
| Condition: Firmware install with "no-activate" option. | |
| Workaround: Activate the firmware which was installed with "no-activate" option. | |

| | |
|---|---|
| Defect ID: DEFECT000621408 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.0.1 | Technology: Security Vulnerability |
| Symptom: Though telnet service on MGMT-VRF is shutdown, telnet access to VDX is allowed via MGMT-VRF. | |
| Condition: 1. Shutdown telnet service on MGMT-VRF 2. Firmware install with "no-activate" option 3. Recover the firmware using "firmware recover" | |
| Recovery: Activate the partially installed firmware. | |

| | |
|---|---|
| Defect ID: DEFECT000621633 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.0.1 | Technology: OSPFv3 - IPv6 Open Shortest Path First |
| Symptom: Not able to change the IPv6 OSPF cost to 1 when auto-cost reference bandwidth is configured. | |
| Condition: The issue is observed for below sequence of steps: <ol style="list-style-type: none"> 1. Configure auto-cost for IPv6 OSPF using CLI: "auto-cost reference-bandwidth 100000" 2. Go to config-rbridge-Ve-<id> interface mode and configure OSPF cost using CLI: "ipv6 ospf cost 1" 3. Run show command to display interface OSPF parameters using CLI: "show ipv6 ospf in ve <id> rb <id>" <p>It is observed that cost field is not changed.</p> | |
| Workaround: Change the cost value to any non default-value and then back to default-value. | |

| | |
|---|--|
| Defect ID: DEFECT000621696 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.0.1 | Technology: Software Installation & Upgrade |
| Symptom: Firmware download fails on standby GOS with error code 26. | |
| Condition: It can happen due to a rare network connectivity issue between the active and standby GOS partitions. | |
| Recovery: Firmware download will be aborted and filesystems will be recovered automatically. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000622093 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.1.0 | Technology: Logical Chassis |
| Symptom: Line card goes into fault state. one of the line card go to faulty state. | |
| Condition: After firmware download from 6.0.2c to 7.1.0 one of the line cards go to faulty state with reason code 119. What should the customer do now? | |

Known Issues for Network OS 7.1.0

| | |
|--|--|
| Defect ID: DEFECT000622356 | |
| 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: The running-configuration of port-channel interfaces have IPv6 nd cache expire time configured to 240 on an upgrade from 7.0.1 to 7.1.0, even though the user hasn't explicitly configured it. | |
| Condition: This issue is seen by the user whenever he upgrades the firmware from 7.0.1 to 7.1.0 | |

| | |
|---|---|
| Defect ID: DEFECT000622864 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.1.0 | Technology: Configuration Fundamentals |
| Symptom: Unexpected reload of switch while taking support save | |
| Condition: After diag test, chassis enable command will fail, switch may go through unexpected reload while taking supportsave | |
| Workaround: Reboot the switch after diag test before trying any other command | |
| Recovery: reboot the switch | |

| | |
|---|---|
| Defect ID: DEFECT000623446 | |
| 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: 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. | |

| | |
|--|---|
| Defect ID: DEFECT000623579 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.1.0 | Technology: VXLAN - Virtual Extensible LAN |
| Symptom: OVSDDB Physical_Switch table incorrectly shows "mac_table_exhaustion" fault status after reboot or ISSU or failover. | |
| Condition: Switch is rebooted or fails over after it learns around 16000 macs from NSX controller. | |
| Recovery: Remove the NSX controller configuration via "no nsx-controller <name>" command and configure it back. | |

| | |
|---|------------------------------------|
| Defect ID: DEFECT000624075 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS6.0.2 | Technology: Logical Chassis |
| Symptom: Disruptive firmware upgrade (coldboot) will fail. | |
| Condition: Number of SNMP communities associated with IPv4/IPv6 ACL configurations is greater than 20. | |
| Workaround: Limit the number of SNMP communities associated with IPv4/IPv6 ACL configurations to less than 20. | |

Known Issues for Network OS 7.1.0

| | |
|---|---|
| Defect ID: DEFECT000624561 | |
| 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: Rebooting host connected to VDX 6940 4x10 GbE breakout port may cause one or more 4x10 GbE ports to become unstable, which could result in port faulting with FFDC excessive interrupts on the port(s). | |
| Condition: Rebooting host connected to VDX 6940 4x10 GbE breakout port may cause one or more 4x10 GbE ports to become unstable, which could result in port faulting with FFDC excessive interrupts on the port(s). | |
| Workaround: Perform shut / no shut on the port to bring it back online. | |
| Recovery: Perform shut / no shut on the port to bring it back online. | |

| | |
|---|---|
| Defect ID: DEFECT000624566 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Data Center Fabric |
| Reported In Release: Network OS7.1.0 | Technology: IP Fabric |
| Symptom: Switch experience Out Of Memory (OOM) condition and reboots | |
| Condition: Using Scaled Configurations | |

| | |
|--|--|
| Defect ID: DEFECT000624714 | |
| Technical Severity: High | Probability: Medium |
| 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: Exceedingly rare (not reproduced so far) error during failover that will cause the system to become faulted. | |
| Condition: Exceedingly rare error during failover. | |
| Recovery: Reboot faulted system. | |

| | |
|---|---|
| Defect ID: DEFECT000624729 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.1.0 | Technology: OpenStack Integration |
| Symptom: VDX 6940 40 GbE port may go offline after upgrade from Network OS Release 7.0.0x to 7.1.0. | |
| Condition: VDX 6940 40 GbE port may go offline after firmware upgrade from Network OS Release 7.0.0x to 7.1.0. | |
| Workaround: Perform shut / no shut on the problem 40 GbE port to bring the port back online. | |
| Recovery: Perform shut / no shut on the problem 40 GbE port to bring the port back online. | |

| | |
|---|--|
| Defect ID: DEFECT000624805 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: Layer 3 Routing/Network Layer |
| Reported In Release: Network OS7.1.0 | Technology: IP Addressing |
| Symptom: Show command is not showing "ip icmp unreachable" under physical interface. | |
| Condition: After configuring the "ip icmp unreachable" under physical interface. | |
| Workaround: This is a cosmetic issue and can be ignored. | |

| | |
|--|-----------------------------------|
| Defect ID: DEFECT000624872 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.1.0 | Technology: Zoning |
| Symptom: In certain cases when joining RBridges together, if one of the RBridges has an empty zone configuration with a default zone mode set to All Access and the other RBridge being joined has an effective zone configuration but a default zone mode set to No Access, this will result in a zone conflict and the cluster will not form. | |
| Condition: This issue can be seen when joining RBridges together that have mismatched default zoning policies. | |
| Recovery: Once in this state, to recover, the customer will need to change the default zone policies so that they match across the RBridges that are being joined and then reattempt joining them together. | |

| | |
|---|--|
| Defect ID: DEFECT000624921 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Security |
| Reported In Release: Network OS7.1.0 | Technology: MAC Port-based Authentication |
| Symptom: After executing "no switchport" on a physical interface with "dot1x mac-auth-enable", dot1x mac-auth-bypass cannot be configured. | |
| Condition: This scenario occurs only if "dot1x mac-auth-enable" is configured while executing "no switchport" | |
| Workaround: Remove "dot1x mac-auth-enable" configuration before doing "no switchport" | |

| | |
|---|--|
| Defect ID: DEFECT000625263 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Layer 2 Switching |
| Reported In Release: Network OS7.1.0 | Technology: VLAN - Virtual LAN |
| Symptom: system may go for unexpected reload | |
| Condition: mac-authentication is enabled with more than 1500 source streams getting authenticated with layer 2 loop existing in the network. Here also loop detection is disabled. | |

| | |
|--|---|
| Defect ID: DEFECT000625402 | |
| 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 authentication key configured on interface is getting lost after config-replay from backup configuration. | |
| Condition: The issue is observed if configurations are done in below order for RBridge sub-mode: <ol style="list-style-type: none"> 1. Configure router OSPF and create area 2. Create VE interface and go to its sub-mode 3. Configure authentication-key using CLI: "ip ospf authentication-key 2 <password>" 4. Save configuration using copy command 5. Delete VE interface 6. Run config-replay command using CLI: "copy flash://<file_name> running-config" <p>Admin can verify that configured authentication-key is lost by using command show running-config for the Ve interface.</p> | |
| Workaround: Configure the OSPF authentication key on the interface using CLI after config-replay is done. | |

Known Issues for Network OS 7.1.0

| | |
|---|------------------------------|
| Defect ID: DEFECT000625616 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.1.0 | Technology: Metro VCS |
| Symptom: 10G ISL using tunable ZR optics (57-1000266-01) does not form between VDX6740 and VDX6940-144s after performing a single "shut/no shut" | |
| Condition: Performing "shut/no shut" on 10G ISL using tunable ZR optics (57-1000266-01) between VDX6740 and VDX6940-144s | |

| | |
|---|--|
| Defect ID: DEFECT000625670 | |
| Technical Severity: High | Probability: High |
| Product: Extreme Network OS | Technology Group: Management |
| Reported In Release: Network OS7.1.0 | Technology: Software Installation & Upgrade |
| Symptom: On rare occasions a SW error may be seen during HA synchronization. | |
| Condition: Rare occurrence that is not linked to a specific change. | |

| | |
|---|--|
| Defect ID: DEFECT000625831 | |
| Technical Severity: High | Probability: Low |
| Product: Extreme Network OS | Technology Group: IP Multicast |
| Reported In Release: Network OS7.1.0 | 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. | |

| | |
|--|------------------------------------|
| Defect ID: DEFECT000625956 | |
| Technical Severity: Medium | Probability: High |
| Product: Extreme Network OS | Technology Group: VCS |
| Reported In Release: Network OS7.1.0 | 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. | |

| | |
|--|---|
| Defect ID: DEFECT000626825 | |
| Technical Severity: High | Probability: Medium |
| Product: Extreme Network OS | Technology Group: Network Automation and Orchestration |
| Reported In Release: Network OS7.1.0 | Technology: OpenStack Integration |
| Symptom: VDX 6740 and 6940 may have linkup issues and/or CRCs after dynamically configuring 4x10GbE breakout mode. | |
| Condition: VDX 6740 and 6940 may have linkup issues and/or CRCs after dynamically configuring 4x10GbE breakout mode. Probability is increased if transceiver is changed from optical to copper or vice versa. | |
| Workaround: Reload switch to recover the links. | |
| Recovery: Reload switch to recover the links. | |

Known Issues for Network OS v7.0.0