Rev. A01
October 2021

Dell EMC Networking OS10 Enterprise Edition 10.4.3.8 Release Notes

This document describes the new features, restrictions, and fixed and known issues in the OS10 Enterprise Edition, Release 10.4.3.8.

This release note document is applicable to all switches present in the Supported hardware section.

For documentation about the Dell EMC open network install environment (ONIE)-enabled hardware switches, see www.dellemc.com/networking.

Document revision history

Table 1. Revision History

Revision	Date	Description
A01	2021–10	10.4.3.8 Release—Added a documentation correction related to port breakout to the Documentation Correction section.
A00	2021–10	10.4.3.8 Release—Added AR-40573, AR-40686, and AR-40704 to the Fixed Issues in 10.4.3.8 section.

Supported hardware

The OS10 Enterprise Edition release 10.4.3.8 is supported on the following Dell EMC switches:

- S3048-ON
- S4048-ON, S4048T-ON
- S4112F-ON, S4112T-ON
- S4128F-ON, S4128T-ON
- S4148F-ON, S4148FE-ON, S4148T-ON, S4148U-ON
- S4248FB-ON, S4248FBL-ON
- S5148F-ON
- S5232F-ON, S5248F-ON, S5296F-ON
- S5212F–ON, S5224F–ON
- S6010-ON
- Z9100-ON
- Z9264F-ON

Restrictions

This section describes limitations and restrictions in this release.

Supported switches: Only a Dell EMC ONIE–based switch supports the OS10 Enterprise Edition 10.4.3.8 software. For a list of supported Dell EMC switches, see Supported hardware.

Third-party software: Dell EMC does not support third-party software and drivers, community projects, code development, or implementation and development of security rules and policies.

Flow-based ERPM: Flow-based ERPM is not applicable for packets whose destination IP prefix is not present in the routing table, or there is no default route configured. This restriction is only applicable to the S5148F-ON switch.

ECMP Static Routes: When you configure static route leaking, all the Equal-cost multi-path (ECMP) static routes from the source do not leak to the destination VRF instance. Only a single ECMP route, normally the best ECMP route, leaks to the destination VRF instance.

Multicast limitations: OS10 does not support the following:

- Fast leave support with a prefix list
- Static multicast group configuration
- Simple Network Management Protocol (SNMP) MIB for Internet Group Management Protocol (IGMP) or Protocol Independent Multicast (PIM)
- Dynamic Rendezvous Point (RP) discovery using Bootstrap Router (BSR)
- VLT in Active-Active mode
- Debug commands from CLI

> **NOTE:** Layer 3 (L3) PIM and IGMP multicast is not supported on the S5148F-ON and S3048-ON switches. IGMP and Multicast Listener Discovery (MLD) snooping is supported on all switches.

New in 10.4.3.8

None

Documentation Corrections

This section describes the errors identified in the current release of the OS10 Enterprise Edition User Guide.

- The OS10 Enterprise Edition User Guide, Release 10.4.3.0 states that for breaking out a 25G port to 10G ports you must use the `interface breakout` CLI command, this information is incorrect. To perform a breakout of an interface, use the `port-group 1/1/x` CLI command and then specify the breakout required.
- The OS10 Enterprise Edition User Guide, Release 10.4.3.0 announces support for 50G breakout interfaces on the S5148F-ON switch, this information is incorrect.
- The OS10 Enterprise Edition User Guide, Release 10.4.3.0 provides a partial list of supported platforms for OpenFlow. In addition to the list in the User Guide, the following platforms also support the OpenFlow protocol:
  - S5232F-ON
  - S5248F-ON
  - S5296F-ON

Known software behavior

802.1X

- 802.1X becomes fully functional only when the feature is enabled globally. If you do not enable 802.1X globally but enable only at the interface level, the system displays the "Dot1x Not Enabled" message.
- Dot1x multi-auth mode is not supported.

BGP

- By default, routes learned on multiple paths to EBGP peers are advertised to IBGP peers with the next-hop local IP address. This behavior allows for local repair of atomic failure of any external peers.
- Fast external failover is enabled by default. To disable or re-enable fast external failover, use the `[no] fast-external-fallover` command. For the `fast-external-fallover` command to take effect on an established BGP session, you must reset the session using the `clear ip bgp {* | `*`peer-ipv4-address`*` | `*`peer-ipv6-address`*`}` command.
- Enabling the BGP add-paths globally for all BGP neighbors is not supported (the `add-path` command in ROUTER-BGPv4-AF or ROUTER-BGPv6-AF mode). To enable the BGP add-path for one neighbor, use the `add-path` command in ROUTER-BGP-NEIGHBOR-AF mode.
- When you redistribute OSPFv3 routes to BGP, including External Type-2 routes, the multi-exit discriminator (MED) attribute is set to the OSPF route metric plus one instead of the OSPF route metric value.
- When you configure the `bgp bestpath router-id ignore` command, for non-best paths, the `show ip bgp` output displays `Inactive reason: Router ID.`
- Do not configure the IP address of the router as a BGP neighbor. This action causes the address being accepted as an invalid neighbor address.

Control-plane ACL

- The control-plane ACL feature is not supported on the S5148F-ON.
- IPv6 ACL is not supported on the S4200-ON series.
- MAC ACL in Control Plane mode is not supported for the management port.

Converged data center services

- When you do not enable PFC on some of the port channel members between the FIP snooping bridge (FSB) and NPIV proxy gateway (NPG), FCoE sessions are not established. You must enable all the members of a port channel with PFC, for the FCoE sessions to establish.

DHCP

- DHCP automatic address allocation—before you configure a DHCP address pool, you must configure a DHCP server interface with an IP address in the range used in the DHCP address pool. If you configure the DHCP address pool first and then configure a DHCP server interface, to enable automatic DHCP address allocation, you must restart the DHCP service using the `disable` and `no disable` commands. Select one of the choices for successfully configuring a DHCP address pool:
  - ○ Configure manual binding for a host/hardware MAC address in the IP address range used for the DHCP pool.
  - ○ Configure a network statement with a valid IP address range.
- DHCP client on management interface—DHCP client is enabled by default on the management interface. The management interface automatically tries to obtain an IP address from a DHCP server. To manually configure an IP address on the management port, disable the DHCP client using the `no ip address dhcp` command in Interface mode.

  ```
  OS10(conf-if-ma-1/1/1)# no ip address dhcp
  ```

- The DHCP server does not start unless at least one interface matches one of the configured network pools. An interface matches a network pool when you include the IP address in the subnet defined for that network pool. For example, an interface with IP address 10.1.1.1/24 matches a pool configured with network 10.1.1.0/24.

  In addition, if you attempt to enable (start) the DHCP server with an incorrect configuration, you must re-enable the DHCP server after you enter the correct configuration.

  Consider the following example, and assume that no interface matches either one of the network pools, `netdhcp1` and `netdhcp2`:

  ```
  OS10# show running-configuration ip dhcp
  !
  ip dhcp server
  no disable
  !
  ```

```
pool netdhcp1
lease infinite
network 35.1.1.0/24
!
pool netdhcp2
network 40.1.1.0/24

OS10# show ip interface brief
Interface Name            IP-Address          OK     Method       Status
Protocol
======================================================================================
===
Ethernet 1/1/1            unassigned          YES    unset        up
up
Ethernet 1/1/2            unassigned          YES    unset        up
up
…

Ethernet 1/1/32           unassigned          NO     unset        up          down
…
```

To resolve this issue, you must:

1. Configure a matching interface for pool `netdhcp2`—40.1.1.1/24 matches 40.1.1.0/24.

   ```
   OS10(config)# interface ethernet 1/1/2
   OS10(conf-if-eth1/1/2)# no shutdown
   OS10(conf-if-eth1/1/2)# no switchport
   OS10(conf-if-eth1/1/2)# ip address 40.1.1.1/24
   ```

2. Run the `show ip interface brief` command to verify if an IP address is assigned to ethernet 1/1/2 port.

   ```

   OS10# show ip interface brief
   Interface Name            IP-Address          OK     Method       Status
   Protocol
   ======================================================================================
   =====
   Ethernet 1/1/1            unassigned          YES    unset        up
   up
   Ethernet 1/1/2            40.1.1.1/24         YES    manual       up
   up
   …
   ```

3. Re-enable the DHCP server because it failed to start initially.

   ```
   OS10# configure terminal
   OS10(config)# ip dhcp server
   OS10(config-dhcp)# disable
   OS10(config-dhcp)# no disable
   OS10(config-dhcp)#
   ```

4. Verify the status of the DHCP server. You must log in with administrator privileges.

   ```
   OS10# system "systemctl is-active isc-dhcp-server"
   active
   ```

   If the DHCP server fails to start, you will see an output similar to the following:

   ```
   OS10# system "systemctl is-failed isc-dhcp-server"
   failed
   ```

- A DHCP client does not get an IP address via DHCP relay unless the DHCP server listens on all interfaces that provide connectivity to that DHCP relay agent.

  This problem happens when you connect a DHCP relay to the DHCP server through multiple paths or interfaces. For example, with ECMP. Also, this issue is seen only when you configure the OS10 switch as the DHCP server.

Consider the following, where the incoming interface IP range (100.1.1.1 to 100.1.*x*.*x*) is not configured as a DHCP server pool in the DHCP server. In this example, only the pool corresponding to the desired client IP address is configured.

```
OS10# configure terminal
OS10(config)# ip dhcp server
OS10(config-dhcp)# no disable
OS10(config-dhcp)# pool client
OS10(config-dhcp-client)# network 30.1.1.0/24
```

However, there are multiple paths between the DHCP server and the relay agent, as shown.

```
OS10# show ip route 30.1.1.0/24
Routing entry for 30.1.1.0/24
Known via bgp, type external
Distance 20, Metric 0
Last update 01:25:11
Routing descriptors Blocks:
    via 100.1.1.2
    via 100.1.1.3
    via 100.1.2.2
    via 100.1.2.3
    via 100.1.3.2
    via 100.1.3.3
    via 100.1.4.2
    via 100.1.4.3
    via 100.1.5.2
    via 100.1.5.3
    via 100.1.6.2
    via 100.1.6.3
    ...
```

The DHCP server listens only on interfaces that match one of the configured pools. An interface matches a DHCP pool when the subnet configured for that pool includes the IP address of the interface. For instance, an interface with the address 100.1.1.1/16 matches the pool with subnet 100.1.0.0/16. In this example, the paths between the relay agent and the DHCP server (100.1.1.1 to 100.1.*x*.*x*) are not configured and hence the DHCP DISCOVER packet does not reach the DHCP server.

To resolve this issue, configure the DHCP server to listen on all interfaces that communicate with the DHCP relay via ECMP.

In addition to the pool with the IP address range desired for the actual client, you must configure a pool with a matching IP range for each of the ECMP interfaces.

```
OS10# configure terminal
OS10(config)# ip dhcp server
OS10(config-dhcp)# no disable
OS10(config-dhcp)# pool client
OS10(config-dhcp-client)# network 30.1.1.0/24
OS10(config-dhcp-client)# pool ecmprange
OS10(config-dhcp-ecmprange)# network 100.1.0.0/16
```

The IP addresses of the interfaces corresponding to the ECMP paths are listed below. These addresses are part of the `ecmprange` pool configured in the example above. The DHCP server listens on all these interfaces and receives the DHCP DISCOVER packet.

```
OS10# show ip route 30.1.1.0/24
Routing entry for 30.1.1.0/24
Known via bgp, type external
Distance 20, Metric 0
Last update 01:25:11
Routing descriptors Blocks:
    via 100.1.1.2
    via 100.1.1.3
    via 100.1.2.2
    via 100.1.2.3
    via 100.1.3.2
    via 100.1.3.3
    via 100.1.4.2
    via 100.1.4.3
    via 100.1.5.2
```

```
     via 100.1.5.3
     via 100.1.6.2
     via 100.1.6.3
     ...
```

Fibre Channel

- On an NPG port, load balancing may not be efficient when all end devices send unicast solicitation. For the load balancing to be efficient, Dell EMC recommends using multicast solicitation.
- After you change the FC Map on FIP snooping enabled active VLAN sessions, use the `shut` and `no shut` commands to re-establish the FCoE sessions.
- Restrict the number of members in an FC zone to 255.
- For the default-zone settings to work properly, ensure that the maximum number of logged-in FC and FCoE nodes is less than 120.
- FCoE-generated ACLs take precedence over user-configured ACLs. A user-configured ingress ACL entry cannot deny FCoE and FIP snooping frames.
- After you remove the vfabric configuration from an interface, to configure the MTU to default value, configure the non-default MTU and then configure the default MTU.
- In a FIP snooping bridge, FIP and FCoE frames ingressing on a PFC mismatch interface are dropped.
- PFC mismatch on a port-channel member port drops FIP and FCoE frames ingressing on that member port, but the learned Enode/Session/FCF information associated with the port-channel is retained. This results in FCoE show commands displaying misleading information. To resolve this issue, check and correct the PFC configuration on both the ends.
- When you configure a port-channel as VLT port-channel, the port-channel goes down operationally and comes up in the local device. The physical interfaces are operationally up. This leads to the switch removing the FCoE sessions. The remote server is not aware of the port-channels being up and down, so the server maintains the FCoE sessions. As these sessions are not available in the switch, the FCoE frames are dropped in the switch. To resolve this, manually flap the port-channel.

IGMP snooping

- If you configure and unconfigure a static connection to a multicast router on an interface using the `ip igmp snooping mrouter interface` command in VLAN mode, the router port still appears in the `show ip igmp snooping mrouter vlan` output. To remove the VLAN port from the show output, configure the VLAN port again using the `ip igmp snooping mrouter interface` command, and then unconfigure it using the `no` version of the command.

Interfaces

- To avoid loops in an L2 network with statically configured port channels, keep the `no switchport` configuration on an interface after you remove its port-channel configuration using the `no interface port-channel` command.
- On the S4112F, S4128F, S4148F and S4148FE switches, 10GBASE-T transceiver operation at 100M or 1G speed is not supported.

IPv6

- IPv6 processing is supported according to the OS10 interface type. The following interface-specific IPv6 settings apply:
  - Physical port and port-channel (LAG) interfaces are in L2 mode by default. IPv6 capability and forwarding are disabled in L2 mode. To enable IPv6 forwarding, set the interface in L3 mode using the `no switchport` and `commit` commands.
  - VLAN and Loopback interfaces come up in L3 mode with IPv6 capability and forwarding enabled by default.
  - On the management interface, IPv6 is enabled by default. IPv6 forwarding is disabled so that the interface operates in Host mode without routing traffic.
  - IPv6 stateless auto-configuration is disabled by default, except on the management interface. To enable auto-configuration, use the `ipv6 address autoconfig` command in Interface mode. Autoconfiguration acquires a global IPv6 address using the network prefix in Router Advertisements. When IPv6 auto-configuration is enabled, IPv6 forwarding is disabled on the interface.

    To disable auto-configuration, use the `no ipv6 address autoconfig` command. IPv6 forwarding remains enabled.

iSCSI optimization

- Do not change the description of a Dell SC series storage device; for example, `Storage Center 65849 SC9000 Version 07.02.01.138`. If you change the description, the SC storage device is not detected by the iSCSI auto-configuration.
- iSCSI auto-configuration on OS10 switch ports is not supported with Compellent storage arrays that use QLE4062 network adapters. To manually configure iSCSI, use the `iscsi profile-storage` *`storage-device-name`* command.
- Starting from release 10.4.1.1, when you perform a fresh installation of OS10, iSCSI autoconfig is enabled and flowcontrol receive is set to on. However, when you upgrade from an earlier release to release 10.4.1.1 or later, the existing iSCSI configuration is retained and the flowcontrol receive could be set to on or off, depending on the iSCSI configuration before the upgrade.
- When you re-configure the iSCSI TCP ports and IP addresses of target storage devices at the same time using the `iscsi target port` command, iSCSI optimization may fail on interfaces connected to the devices. To successfully enable iSSCI optimization, enter only the new TCP port number(s) in the command and do not specify an IP address
- On an S3048-ON switch, iSCSI auto-configuration is disabled, by default. You must manually enable iSCSI optimization on the switch using the `iscsi enable` command.

LLDP

- The Linux LLDPD doesn't recognize IPv6 management address TLVs. In order for OS10 LLDP to work seamlessly with Linux LLDPD, disable IPv6 management address TLV on the interface connected to the Linux device.

Management VRF

- Before you assign the management port to the management VRF instance, you must remove all configured settings on the management port, including the IP address. **Perform this action from the console**. Removing the IP address disconnects all existing SSH and Telnet sessions on the switch.

  The following example shows removing IP address, configuring management VRF, and then adding IP address:

  ```
  OS10(config)# interface mgmt 1/1/1
  OS10(conf-if-ma-1/1/1)# show configuration
  !
  interface mgmt1/1/1
   no shutdown
   ip address 10.16.208.125/16
   ipv6 address autoconfig
  OS10(conf-if-ma-1/1/1)# no ip address
  OS10(conf-if-ma-1/1/1)# no ipv6 address autoconfig
  OS10(conf-if-ma-1/1/1)# exit
  OS10(config)# ip vrf management
  OS10(conf-vrf)# interface management
  OS10(conf-vrf)# exit
  OS10(config)# interface mgmt1/1/1
  OS10(conf-if-ma-1/1/1)# no shutdown
  OS10(conf-if-ma-1/1/1)# ip address 10.16.208.125/16
  OS10(conf-if-ma-1/1/1)# ipv6 address autoconfig
  ```

OpenFlow

- The ONOS controller does not encode the DSCP flow entry values that are matched according to the Openflow 1.0 specification. Hence when you install a flow entry in OpenFlow 1.0, that matches the IP DSCP, the ONOS controller sets an incorrect flow-entry encoding value for IP DSCP.

Passwords

- When you enter a password in an OS10 command, either at a password prompt or in the command syntax, you can enter only alphanumeric and certain special characters — $ - _ . + ! * ' () — unencoded. You cannot enter any other special characters in the password. Use URL encoding instead.

  For example, in the `image download` command, the password `a@b` is not accepted: `image download ftp://`*`username`*`:a@b@10.11.63.122/`*`filename`*. You must enter the password as `image download ftp://`*`username`*`:a%40b@10.11.63.122/`*`filename`*. The URL encoding for `@` is `%40`. For information about other characters that require URL encoding, go to URL Encoding.

RADIUS Authentication

- RADIUS Server Source Interface and RADIUS Server VRF configuration work individually, but do not work together.

Remote port mirroring

- When you configure a port as a source interface, and add the same port to the remote VLAN used for monitoring traffic, the configuration fails and the system does not display an error message. Dell EMC recommends adding the ports to the destination remote VLAN first and then configuring the source interface. In this case, when you configure the destination port as source, the system displays an error message.
- When you configure remote port mirroring and overwrite the transport VLAN by re-entering the `destination remote-vlan` *`vlan-id`* command with a different VLAN ID, an error message displays. The new remote VLAN configuration is not accepted. You must first remove the configured remote VLAN using the `no` version of the command, and then re-enter the command with the new remote VLAN ID.
- VLAN statistics on the remote port mirroring (RPM) VLAN interface are not incremented on the following switches: S4048-ON, S4048T-ON, S4100-ON, S6010-ON, and Z9264F-ON. For these switches, the `show interface vlan` *`rpm_vlanid`* command does not display statistics for the mirrored traffic.

sFlow

- Each time you enable sFlow at the interface level, the sFlow agent restarts, causing a delay in polling counter statistics. As a result, the sFlow counters for all sFlow-enabled ports may not be accurate.
- sFlow counter statistics that are individually reported for the port members of a port-channel data source are accurate. However, the counter statistics reported for the port channel may not be accurate. To calculate the correct counters for a port-channel data source, add together the counter statistics of the individual port members.
- When you enable sFlow in Per-Interface mode, the counter statistics of sFlow-enabled ports reset to zero when you add a new member port or remove an existing member port from any sFlow-enabled port-channel group. To avoid this impact, enable sFlow globally by using the `sflow enable all-interfaces` command.

  ```
  OS10(config)# sflow enable all-interfaces
  ```
- Dell EMC recommends deploying sFlow in Global mode, which is the native mode for sFlow.

SupportAssist

- SupportAssist requires that you configure a name-server and a default route using the `server url` *`server-url`* command.
- The `proxy server ip` command does not support an IPv6 address to reach the SupportAssist server.
- Automated email notification at the time of a hardware fault alert, automatic case creation, automatic part dispatch, or reports are not supported.

System administration

- **Config Partition Disk Utilization**

When the config partition has low disk space, you will see a syslog message as below:

```
SYS_STAT_LOW_DISK_SPACE: Warning! Configuration directory has 0.0% free. Please
delete unnecessary files from home directory.
```

When you see such errors, please delete unwanted files from the home directory or you may encounter degraded system performance.

- **Admin User**

  You can delete the default **admin** username, as long as there is a local user with sysadmin role present. The default admin user sees a message in MOTD, unless the user password is changed or the user is deleted.

- **Linux Admin User**

  Password of the linuxadmin user can be modified via a CLI. The linuxadmin user can also be enabled or disabled via another CLI.

System monitoring

- Logging is enabled by default on a terminal emulator connected to the console serial port. However, in an SSH or Telnet terminal session, logging is disabled by default. To enable logging on a remote terminal in an SSH or Telnet session, use the `terminal monitor` command in EXEC mode. To disable logging in a remote or directly connected terminal, use the `no terminal monitor` command.

VLANs

- The valid VLAN ID range displays as `1-4093`. VLAN IDs 4094 and 4095 are reserved for internal use.

VLT

- To check mismatch of MAC address table entries between VLT peers, use the `show vlt mac-inconsistency` command. To identify mismatches in VLT configuration on peer switches, use the `show vlt` *`domain-name`* `mismatch` command.

```
OS10# show vlt-mac-inconsistency
Checking Vlan 228 .. Found 7 inconsistencies .. Progress 100%
VLAN 128
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 1
----------
MAC 00:a0:c9:00:00:18 is missing from Node(s) 2
MAC 00:a0:c9:00:00:20 is missing from Node(s) 2
VLAN 131
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 132
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 135
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 137
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2

Run "show vlt d1 mismatch ..." commands to identify configuration issues
```

VRRP

- Priority 255 is not supported.

VXLAN

- In a static VXLAN, overlay routing is supported on:
  - S4100-ON Series
  - S4200-ON Series
  - S5200-ON Series
  - S4048T-ON
  - S6010-ON

Known software behavior — S4200 series

The following limitations are applicable to S4248FB-ON and S4248FBL-ON switches:

Ingress ACL

- The following applications require ACL tables: VLT, iSCSI, L2 ACL, L3 v4 ACL, L3 v6 ACL, PBR v4, PBR v6, QoS L2, QoS L3, FCoE. In ingress ACL, you can create ACL tables only for three applications at a time.
- In IPv6 ACL and PBR ACL, l4–destination-port, l4–source-port, flow label, and TCP flags are not supported.
- IP fragment supports only 2 options: non-fragment and head/non-head.

Egress ACL

- You can create either Layer 2 ACL or Layer 3 ACL. You cannot create both the tables at a time.
- In egress L3 IPv4 ACL, the fragment, TCP flags, and DSCP fields are not supported.
- In egress ACLs, L2 user table is utilized only for switched packets and L3 user table is utilized only for routed packets.

Counters

- In the `show interface vlan` command output, the VLAN octet counters are not displayed accurately.
- If a packet hits two ACL tables, the counter with higher priority statistics gets incremented and the other actions are merged and applied.

Layer 2

- The default MAC aging time is set as 550 seconds. This is the maximum value that can be configured.

Layer 3

- Though it is possible to configure more VRIDs in VRRP, the S4200–ON Series switches support only up to 16 VRIDs. This number decreases when VLT peer routing is enabled.

Load balancing

- The command `load-balancing` does not work with the `tcp-udp-selection` parameter.

QoS — CoPP

- Shaping does not support traffic less than 468 kbps. Configure the shaping rates in multiples of 468.
- In System Flow ACL, ARP request and ARP response packets share the same CPU queue.
- CPU queues support shaping instead of rate limiting.

- Port shaping, storm control rate shaping, and CoPP rates are converted to kbps internally, even when configured in pps.

QoS — Egress Scheduler

- If PFC is provisioned, the control packets injected by CPU shares queue-7 while egressing on a front panel interface. If queues other than queue-7 are provisioned as strict priority, it is recommended to provision queue-7 as strict priority too, to reduce latency or loss of control packets.

QoS — PFC and class maps

- Provisioning PFC is not supported when deep buffer mode is enabled.
- Configure the traffic class ID to queue mapping policy on egress interfaces.
- You cannot enable PFC on all the physical interfaces, when you have split the ports to multiple breakout interfaces. For more information, see the 'PFC configuration notes' section in the *Dell EMC Networking OS10 Enterprise Edition User Guide*.
- When you add or remove the PFC configuration on an interface, the interface gets flapped. Stop the traffic before applying or modifying the PFC configuration.

QoS — LLFC

- Provisioning LLFC is not supported when deep buffer mode is enabled.
- Stop the traffic before applying or modifying the LLFC configuration.

RADIUS Authentication

- RADIUS Server Source Interface and RADIUS Server VRF configuration work individually, but do not work together.

Resilient Hashing

- Resilient hashing is not supported on the S4200–ON Series switches.

sFlow

- Do not enable sFlow on per-port basis.

VXLAN

- On the S4248-ON, IPv6 overlay routing is not supported with static VXLAN. IPv6 overlay routing is supported with BGP EVPN.
- S4248-ON does not copy DSCP from the inner header to outer header when switching traffic to the same virtual network in remote VTEPs.
- S4248-ON does not copy dot1P from the inner header to outer header when routing traffic from one virtual network to another.

Known software behavior — S5200 series

The following limitations are applicable to the S5212F-ON, S5224F-ON, S5232F-ON, S5248F-ON and S5296F-ON switches:

Ingress ACL

- When you configure QoS service-policy on an S5200-ON switch that is in a VLT setup with MAC and IP ACLs configured, an error appears. This issue occurs because of ACL group width limitation in the S5200-ON series switches. VLT, IP, MAC, and

QoS ACLs require double-width ACL table slice. The S5200-ON series switches support only three applications that require double-wide ACL table slice at a time. An error appears because the QoS application configuration requires a fourth ACL table slice.

Known hardware behavior

Fan LED

- On an OS10 S3048-ON switch with reverse airflow:
  - When all fans are operational, the Fan LED is solid amber.
  - When a fan fails, the Fan LED is blinking amber.

  On an OS10 S3048-ON switch with normal airflow:
  - When all fans are operational, the Fan LED is solid green.
  - When a fan fails, the Fan LED is blinking green.

Fixed issues in 10.4.3.8

AR-40573 Security updates for authentication related vulnerability in RESTCONF API.

AR-40686 The following CVEs have been addressed:

- CVE-2021-3711
- CVE-2021-3712
- CVE-2021-23839
- CVE-2021-23840
- CVE-2021-23841
- CVE-2020-1968
- CVE-2020-1971
- CVE-2019-1547
- CVE-2019-1551
- CVE-2019-1552
- CVE-2019-1559
- CVE-2019-1563
- CVE-2018-0732
- CVE-2018-0734
- CVE-2018-0737
- CVE-2018-0739
- CVE-2018-5407
- CVE-2017-3735
- CVE-2017-3736
- CVE-2017-3737
- CVE-2017-3738

The CVE database can be accessed here: https://cve.mitre.org/cve/search_cve_list.html.

AR-40704 Security updates for authentication related vulnerability in SmartFabric Services.

Fixed issues in 10.4.3.7

AR-25279 In certain scenarios, a software exception may occur when removing an SNMP user.

AR-34744 If a MAC is learned on a VLT port channel, it is not learned on an orphan port.

AR-37753 DHCP relay does not function when DHCP reply packets are incorrectly hashed in a VLT peer node.

NOTE: This fix is specific to the S5148F–ON switch.

AR-39052 In certain scenarios, if telemetry is enabled, the switch may reboot.

NOTE: This fix is specific to the S5148F–ON switch.

AR-39573 User configured cluster security-profile does not take effect after upgrading the operating system and this impacts VLT and DNV cluster convergence with peers that use the same security-profile.

AR-39599 Updated the old default X.509v3 certificate with a new X.509v3 default certificate.

AR-39777 VLT fails to converge after installing a custom SSL certificate.

Fixed issues in 10.4.3.6P3

AR-35300 On S4112-ON switches, the SFS master fails to sync with the backup switch after upgrade to 10.5.0.x.

Fixed issues in 10.4.3.6

AR-26136 Untagged data traffic may be truncated or corrupted if traffic egresses via ICL.

NOTE: This fix is specific to the S5148F-ON switch.

AR-26350 Random SFP28 25GBASE interfaces fail to come up after a peer device is rebooted.

NOTE: This fix is specific to the S5148F-ON switch.

AR-28419 Interfaces with 10G DAC fail to come up randomly when connected to an Intel NIC.

NOTE: This fix is specific to the S5148F-ON switch.

Workaround: Toggle the interface admin state on the switch using the `shutdown` command followed by the `no shutdown` command.

Fixed CVEs

The **Common Vulnerabilities and Exposures** (**CVE**) system provides a reference-method for publicly known information-security vulnerabilities and exposures.

The following CVEs are addressed in this release:

- CVE-2019-3846
- CVE-2019-5489
- CVE-2019-9500
- CVE-2019-9503
- CVE-2019-10126
- CVE-2019-11477
- CVE-2019-11478
- CVE-2019-11479
- CVE-2019-11486
- CVE-2019-11599
- CVE-2019-11815
- CVE-2019-11833
- CVE-2019-11884

Fixed issues in 10.4.3.5

AR-25393 Ports may go into down state randomly.

ⓘ **NOTE:** This fix is specific to the S5148F-ON switch.

Workaround: Move the media to another port.

AR-26615 When using the `access-list` command with the `range port` command, the policy module sends multiple access-list rules with changed ports. Though these rules are seen as duplicate by the SM module, they are still inserted in the `sm_stub_access_tree`, which results in a loop in the SM module.

AR-26800 In certain scenarios, stack corruption occurs in SNMP module leading to a switch crash.

AR-27149 When executing the `show ip bgp l2vpn evpn summary` command, an incorrect comparison is made in the BGP module for afi or safi settings that causes a failure to add the peer to the output display buffer.

AR-27426 In certain scenarios, the port channel does not come up when the VLT port channel is created using LACP.

AR-27503 On a multicast snooping enabled VLAN (IGMP or MLD), if the IP multicast packet has the TTL value set to 1, then the packet may not be forwarded based on the snooping database.

ⓘ **NOTE:** This fix is applicable to all platforms except the S4200-ON Series and S5148F-ON switch.

Fixed issues in 10.4.3.4

AR-26646 RADIUS authentication may fail on a dot1x enabled client.

AR-26989 In a VLT environment, modifying switch port access vlan configuration may fail, and sometimes the configured switch port access vlan is not retained after reload.

Fixed issues in 10.4.3.3

AR-24678 In certain scenarios, after a VLT node is upgraded or reloaded, the VLT APP may fail to write a few vlan configurations to a persistence database. This may result in po1000 not being added to those vlans.

AR-25110 In a VLT environment, certain transit traffic with a particular byte offset value can incorrectly match a hardware entry and gets copied to the CPU.

Fixed issues in 10.4.3.2

AR-24509 Login banner does not display in the ssh session when upgrading from release 10.3.x to 10.4.2 or 10.4.3.

AR-25007 TCP port number 443 is not closed by default even if the REST service is not enabled.

AR-25565 Link LED’s might not work correctly for certain ports after upgrade to 10.4.3.0 release.

AR-25869 Unable to make any interface configuration changes if `delete config://startup.xml` command is used with a non default switch-profile configuration.

Workaround: Use `delete startup-configuration` command.

AR-25906 BGP neighbor commands are lost if primary partition is loaded with 10.4.3.1 and secondary partition is loaded with a release earlier than 10.4.3.0.

Workaround: Load the same image in both partitions.

AR-25908 Unable to delete the snmp-server host configuration if the switch is upgraded from 10.4.2.2 to 10.4.3.1.

AR-26384 A software exception occurs when the user releases the IP address of the management interface acquired through DHCP, and then configures the NTP server.

Fixed issues in 10.4.3.1

AR-24986 10G copper SFP links do not come up after upgrading the OS to 10.4.3.0.

AR-25001 In certain scenarios, the system boots up with the default configuration after upgrade to 10.4.3.0.

AR-25068 SSH session exits when the user configures description with an empty string under interface range context.

Workaround: Configure the description field with any string instead of an empty string.

AR-25085 1G copper SFP links do not come up after upgrading the OS to 10.4.3.0.

NOTE: This fix is specific to the S5148F-ON switch.

Fixed issues in 10.4.3.0

AR-21718 IGMP snooping CLI commands should not be available in the default-vlan.

AR-22017 User data traffic not load balanced, leading to over utilization of links.

AR-22021 Unable to remove route-map from BGP redistribute configuration.

Workaround: Use different route-map for each of the address-family/source protocol.

AR-22545 xSTP process crashes if the default vlan-id is different from the vlan-id for the STP BPDUs.

Workaround: Set the default vlan-id to 1.

AR-22937 In a VLT environment, when repeated spanning-tree topology change notifications are received, memory accumulation might occur in an internal process.

AR-23084 In rare scenarios, the system may get stuck in the loading state when REST API calls are continuously sent to the system.

AR-23090 In a VLT environment, OSPF neighbor ship with a switch running OS9, might transition to the INIT state if an intermediate node is reloaded.

Workaround: Unconfigure ip dhcp snooping.

AR-23383 The switch hangs when the `show vlt` command is executed.

AR-23524 In a VLT environment, switches running port configurations involving 10G and 100G setup may experience receive or transmit failure of control packets.

NOTE: This fix is specific to the S4100–ON Series switches.

AR-23528 Following hang, VLT peer fails to add all vlans to VLTi.

AR-23751 The `show ip route` command displays the gateway of the last resort based on the default route.

AR-23881 Unable to configure valid IP addresses that contain number 255 in any octet postion as the RADIUS or TACACS server host.

AR-23901 The system reports unnecessary logs when RADIUS or TACACS authentication is enabled.

AR-24081 The TACACS source interface is not sourced from IPv4 or IPv6 loopback address after reload.

AR-24471 VRRP VIP is not reachable via ICMP.

AR-24555 Unable to add an interface to a VRF when the configuration `ipv6 address autoconfig` is not present in the interface configuration.

AR-24573 Ping to a virtual IP fails, if the same VRRP group is configured in all the 4 nodes in an EVLT setup.

Workaround: Configure the `no vrrp mode active-active` command in the affected interfaces.

AR-24675 System may crash when the `license install` command is executed with a "/" at the end.

Workaround: Install the license with the exact file path.

Known issues in 10.4.3.0

AR-22782

All ECMP static route paths may not get leaked in static route leaking. Only first best route path is leaked to a destination VRF. If the leaked best route path goes down then next best path will get installed in the destination VRF.

AR-22793

Batch command operation for certain commands may fail when used from an existing running configuration.

Workaround: Modify the configuration list in batch commands to make sure it aligns with the inter-dependency of the configuration. Example: `OSPF dead-interval` and `Hello-interval` should be corrected to the following order in the batch file, if they are copied from the running config :

```
Router# interface vlan1041
description Interface-1
no shutdown
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
```

AR-23065

On the S5232F/S5248F and S5296F, for ICMP / IGMP PDUs, control-plane, ACL may not take effect.

Workaround: To deny ingress of corresponding traffic streams, an equivalent ingress ACL can be applied on the required ports.

AR-23341

When a 25G DAC 5m cable is plugged in, the port may not come OPER UP with `Auto Negotiation off` on both the ends for the S5200–ON Series and S5100–ON Series switches.

Workaround: Configuring `Auto negotiation on` at both the ends will bring up the port.

AR-23723

Ports with 2*50G DAC cables may be OPER DOWN for S4100–ON Series switches if connected to a non-S4100–ON Series switch at the remote end.

Toggle `Auto neg` on the S4100–ON Series ports with 2*50G DAC cables.

AR-23899

On the Z9264F-ON switch running 10.4.3.0, the `show system` command output on a working AC PSU incorrectly displays as `Type:DC`.

AR-24149

When VRF and IP config are removed from a PIM enabled interface and added back without removing PIM config, PIM interface does not come up. This Issue can occur when dn_sm tries to set PIM config on an interface before the VRF name association to the interface.

Workaround: When removing IP and VRF config, PIM config should be removed and then configs should be re-applied. If not, once the issue comes up, removing and re-applying the PIM config will resolve the issue.

AR-24215

Even after Virtual-network-interface admin-down, Link-Local-IPv6 neighbors learnt on those interfaces may still be shown in `show ipv6 neighbors` output, but those link-local-IPv6 neighbor entries are correctly removed from the hardware table.

AR-24266

On the S4200–ON Series switch, when L3 vxlan traffic is received by a network port, the source MAC will not be learnt, as MAC learning of native-SA in routing over overlay (ROO) packets is not supported on this switch.

In IPv4 environment, the MAC address is learnt from ARP packets. As ARP packets are L2 switched, but in IPv6 environment, the MAC addresses are learnt from ICMPv6 NS, NA packets. As these are also routed packets the MAC address will not get learnt here. So for IPv6 environment, static L3 vxlan is not supported.

Workaround:

1. The MAC address can be added statically.
2. Issue is only with static L3 vxlan. In case of BGP EVPN, the MAC address is configured from the control plane hence this will not impact BGP EVPN.

3. By sending IPv6 switched traffic, the MAC address can be learnt.
4. If the environment is a dual stack IPv4/IPv6, then MAC address will be learnt from IPv4 ARP packets.

AR-24408

When the last receiver in a broadcast domain (VLAN) has joined a Group G and it sends a leave for that group, the group gets removed immediately from the Querier while it gets removed after 2 minutes (membership expiry) from the non-querier.

AR-24676

Management VRF users are unable to install a license using tftp, scp, ftp or http.

Workaround: Download the license file using tftp, scp, ftp or http to the device's local user home directory and install the license using the `license install localfs://home/admin/<filename>` CLI command. Follow these steps:

1. Copy the license file from the ftp server location to the home directory on the system.

```
OS10# copy ftp://admin:admin@10.11.95.101//home/admin/LADF/7B900Q2-
NOSEnterprise-License.XML home://7B900Q2-NOSEnterprise-License.XML
```

2. (optional) Use the `show copy-file status` command to check the status of the file copy.

```
OS10# show copy-file status
File Transfer State:    idle
---------------------------------------------------
State Detail:           idle
Task Start:             2019-02-15T00:46:35Z
Task End:               2019-02-15T00:46:36Z
Transfer Progress:      100 %
Transfer Bytes:         3795 bytes
File Size:              3795 bytes
Transfer Rate:          8 kbps
```

3. Verify that the license is present in the home directory of your system.

```
OS10# dir home

Directory contents for folder: home
Date (modified) Size (bytes) Name
--------------------- ------------ ------------------------
2019-02-15T00:47:25Z 3795 7B900Q2-NOSEnterprise-License.XML
```

4. Execute the `license install` command specifying the path to the home directory location where the license was downloaded in step 1.

```
OS10# license install localfs://home/admin/7B900Q2-NOSEnterprise-
License.XML
[ 5784.994389] EXT4-fs error (device dm-0):
ext4_has_uninit_itable:3039: comm CPS_API_instanc: Inode table for bg
0 marked as needing zeroing
License installation success.
```

AR-24867

If a multicast data source is connected to a router which is both FHR and RP, an interface flap on the multicast source connected interface would cause data traffic to stop flowing for ~3 minutes (till expiry timer for S,G expires). After ~3 minutes, multicast traffic automatically recovers.

Workaround: Avoid issuing `shut` or `no shut` command on the interface connected to a multicast source or avoid having the same router as both FHR and RP.

Installation

For complete installation and upgrade information using the ONIE installer, follow the instructions in the *Dell EMC Networking OS10 Enterprise Edition User Guide*. Before installing, download the 10.4.3.8 image and a license.

NOTE: After installing the NOS/DIAG-OS, if you boot into ONIE Install mode, ONIE assumes ownership of the switch—ONIE Install mode is sticky. In this situation, ONIE stays in Install mode until NOS/DIAG-OS is successfully installed again. If you want to boot into ONIE for any reason other than installation, use **Rescue** mode or **Update** mode.

Install OS10

ⓘ **NOTE:** Use the `image download` command to download the software image to the current active partition — it does not install the software on your device. The `image install` command installs the downloaded image to the standby partition. To verify OS10 installer image, see Verify OS10 installer image.

△ **CAUTION: Please do not use copy commands to download the image to the switch, as it may install the image on the config partition resulting in loss of disk space for critical system applications and functions.**

ⓘ **NOTE:** If the active partition contains any modified text files or custom packages installed, they would not be available in the standby partition. Backup the modified files and re-install the packages after downloading the image.

1. (Optional) Backup the current running configuration to the startup configuration in EXEC mode.

   ```
   OS10# copy running-configuration startup-configuration
   ```

2. Backup the startup configuration in EXEC mode.

   ```
   OS10# copy config://startup.xml config://<backup file name>
   ```

3. Download the new software image from the Dell Support Site, extract the *bin* files from the *tar* file, and save the file in EXEC mode.

   ```
   OS10# image download file-url
   ```

   For example:

   ```
   OS10# image download ftp://userid:passwd@hostip:/filepath
   ```

   ⓘ **NOTE:** Some Windows unzip applications insert extra carriage returns (CR) or line feeds (LF) when they extract the contents of a `.tar` file, which may corrupt the downloaded OS10 binary image. Turn off this option if you use a Windows-based tool to untar an OS10 binary file.

4. (Optional) View the current software download status in EXEC mode.

   ```
   OS10# show image status
   ```

5. Install the 10.4.3.8 software image in EXEC mode.

   ```
   OS10# image install image-url
   ```

   For example:

   ```
   OS10# image install image://filename.bin
   ```

6. (Optional) View the status of the current software install in EXEC mode.

   ```
   OS10# show image status
   ```

7. Change the next boot partition to the standby partition in EXEC mode.

   ```
   OS10# boot system standby
   ```

8. (Optional) Check whether the next boot partition has changed to standby in EXEC mode.

   ```
   OS10# show boot detail
   ```

9. Reload the new software image in EXEC mode.

   ```
   OS10# reload
   ```

10. After the installation is complete, use the `show version` command to check if the latest version of the software is running in the system.

The following example shows the 10.4.3.8 software installed and running:

```
OS10# show version
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2021 by Dell Inc. All Rights Reserved.
OS Version: 10.4.3.8
Build Version: 10.4.3.8
Build Time: 2021-05-28T20:45:46-0800
System Type: S5232F-ON
Architecture: x86_64
Up Time: 06:58:41
OS10#
```

Verify OS10 installer image

The OS10 installer image tar file (OS10_Enterprise_10.4.3.8.tar) contains the following files.

- The installer binary (PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin)
- The checksum (PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin.sha256)
- Detached binary signature (PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin.gpg)

This section describes how to verify OS10 installer image using any one of the following methods:

- Verify the signature of the installer image
- Verify the checksum of the installer image

Dell EMC recommends that you verify the signature of the installer image.

Verify the signature of the installer image

1. Download the Dell EMC official key.

```
# gpg --keyserver pool.sks-keyservers.net --recv-keys A9FCCB28
gpg: key D836AEA3: public key "DellEMC OS10 Networking Signing Key
<gpg.NW@dell.com>" imported
gpg: Total number processed: 1
gpg:               imported: 1  (RSA: 1)
```

2. Verify the signature. You must see a "Good signature" message similar to the following:

```
# gpg --verify PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin
 PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin
gpg: Signature made Fri 1 Oct 2021 12:13:32 PM PST using RSA key ID
A9FCCB28gpg: Signature made Fri 1 Oct 2021 12:13:32 PM PST using RSA key ID
A9FCCB28
gpg: Good signature from "DellEMC OS10 Networking Signing Key <gpg.NW@dell.com>"
gpg: WARNING: This key is not certified with a trusted signature!
gpg:          There is no indication that the signature belongs to the owner.
Primary key fingerprint: 8AFD 006B D6B7 363B 67F6  7608 58DF 862A D836 AEA3
     Subkey fingerprint: AD62 94D3 C8F6 803A 252A  F876 F1DB D597 A9FC CB28
```

If the .bin file is corrupted, a message similar to the following one appears.

```
gpg: Signature made Fri 1 Oct 2021 12:13:32 PM PST using RSA key ID
A9FCCB28
gpg: BAD signature from "DellEMC OS10 Networking Signing Key <gpg.NW@dell.com>"
```

If you do not see the "Good signature" message, stop the process and re-download the installer image or contact Technical Support.

Verify the checksum of the installer image

You can verify the checksum of the installer image using the following commands. The output of these commands is a 64-bit character string.

```
# sha256sum PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin
e5e985cdea20a7ccdaaaa9192beca46eafc7bfeda235dc1eb78eb63602ae8571
PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin
```

```
# cat PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin.sha256
4aac9169e0165636940831d530bdae5776d1f8c17886bc9a527b9020b32c3d5e
PKGS_OS10-Enterprise-10.4.3.8.244stretch-installer-x86_64.bin
```

The checksum value, similar to the following one appears.

```
4aac9169e0165636940831d530bdae5776d1f8c17886bc9a527b9020b32c3d5e
```

Ensure that the checksum value displayed in the outputs of the two commands match. If the checksums do not match, re-download the image or contact Technical Support.

Upgrade and Downgrade

To upgrade or downgrade, download the image and license of the required version.

Upgrade OS10

NOTE: The default X.509v3 certificate used for VLT peer and SFS cluster convergence expires on July 27, 2021. When the certificate expires, if there is a network event such as a reload or a flap, the VLT peers and SFS cluster members will cease to communicate. As a result, network traffic will be affected. See *Section 3, Preparing for an upgrade* in the Dell EMC SmartFabric OS10 Installation, Upgrade, and Downgrade Guide for details.

NOTE: When upgrading from 10.4.3.2, 10.4.3.3, 10.4.3.4, 10.4.3.5 or 10.4.3.6, 10.4.3.7 to 10.4.3.8, a backup of the desired running configuration should be taken prior to the upgrade. Steps to backup and restore the running configuration:

- Before upgrade, save the latest running configuration to a file using the `copy running-configuration config://backup` command.
- After upgrade to 10.4.3.8 or higher, copy the saved configuration to the running configuration using the `copy config://backup running-configuration` command, followed by the `write memory` command.

NOTE: During the image upgrade process in a VLT setup, when the VLT peers are running different software versions, no configuration changes should be done in any of the VLT peers. Ensure that both the nodes are upgraded to the same version before you make any configuration change.

NOTE: When you upgrade VLT peers from 10.4.0E(R2) or earlier to 10.4.0E(R3) or later, upgrade both the VLT nodes at the same time during the maintenance window, as there is a possibility of traffic impact during the upgrade.

1. If you have installed the image for 10.4.3.7, then the networking OS partition contains the following:

Partition A	10.4.3.7
Partition B	10.4.3.8
Configuration Partition	`Startup.xml` specific to 10.4.3.7 Other configuration files
License Partition	License related files

Boot details:

- Active: A
- Standby: B

2. Back up `Startup.xml` to an external device as `10.4.3.7-Startup.xml`.
3. Download the latest image, using the `image download` command.
4. Install the downloaded image, using the `image install` command.
5. For example, if you have installed the image for 10.4.3.7, then the networking OS partition contains the following:

Partition A	10.4.3.7
Partition B	10.4.3.8
Configuration Partition	`Startup.xml` specific to 10.4.3.7 Other configuration files
License Partition	License related files

   Boot details:

   - Active: A
   - Standby: B

6. To upgrade to 10.4.3.8, change the boot system to `standby` and reload the device.

```
OS10# boot system standby
```

7. The switch boots up with 10.4.3.8 successfully.
8. The networking OS partition contains the following:

Partition A	10.4.3.7
Partition B	10.4.3.8
Configuration Partition	`Startup.xml` specific to 10.4.3.7 Other configuration files
License Partition	License related files

   Boot details:

   - Active: B
   - Standby: A

Password-less login

In a 10.4.3.7 release, if you have enabled password-less login to an SSH server using a script and upgrade to 10.4.3.8, you can enter the `username` *`username`* `sshkey` *`filename`* command to re-enable SSH password-less login. In the command, *`filename`* is public-key from the `id_rsa.pub` file from a Linux switch. A script is not required.

The password-less access is applicable only on the Linux switch from which you have copied the public key.

Enter the entire filename string within double quotes ("").

Example:

```
OS10(config)# username sshtest sshkey “ssh-rsa…….@netlogin-eqx-04”
```

Downgrade OS10

CAUTION: The following procedure requires administrative access to the Linux server. Be careful when you run the following commands.

To downgrade to the previous OS10 version, execute the following steps:

1. Change the boot system to A, the partition that has the old image, using the `boot system standby` command.
2. Restore the original `Startup.xml`, that was saved as `10.4.3.7-Startup.xml`, from the external device to the configuration partition.
3. Delete the configuration database and reboot the device system using the following command:

```
OS10# system "sudo sh -c 'rm -vf /config/var/lib/redis/dump.rdb; sync; reboot -f'"
[sudo] password for admin: <<Enter the admin password>>
```

```
removed '/config/var/lib/redis/dump.rdb'
Failed to read reboot parameter file: No such file or directory
Rebooting.
```

4. The networking OS partition contains the following:

Partition A	10.4.3.7
Partition B	10.4.3.8
Configuration Partition	`Startup.xml` specific to 10.4.3.7 The configuration database is deleted
License Partition	License related files

5. After rebooting, the system:
   - Loads up with 10.4.3.7 image.
   - Loads `10.4.3.7-Startup.xml`.
   - Automatically creates the new configuration database.

Support resources

The Dell EMC Support site provides a range of documents and tools to assist you with effectively using Dell EMC devices. Through the support site you can obtain technical information regarding Dell EMC products, access software upgrades and patches, download available management software, and manage your open cases. The Dell EMC support site provides integrated, secure access to these services.

To access the Dell EMC Support site, go to www.dell.com/support/. Sign in with a previously created account or create a new account. To display information in your language, scroll down to the bottom of the page and select your country from the drop-down menu.

- To obtain product-specific information, enter the 7-character service tag or 11-digit express service code of your switch and click **Submit**.

  To view the service tag or express service code, pull out the luggage tag on the chassis or enter the `show chassis` command from the CLI.

- To submit service requests or to contact technical support by phone or chat, click **Contact Us**, then click **Technical Support**.

To access product documentation and resources that might be helpful to configure and troubleshoot the OS10 Networking operating system, see the Dell EMC Networking OS10 Info Hub.

To search for drivers and downloads, see www.dell.com/drivers/.

To participate in Dell EMC community blogs and forums, see www.dell.com/community.

Notes, cautions, and warnings

ⓘ **NOTE:** A NOTE indicates important information that helps you make better use of your product.

△ **CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.**

⚠ **WARNING: A WARNING indicates a potential for property damage, personal injury, or death.**