

Netvisor ONE Version 7.0.1 GSG for Dell Z9432F-ON Platforms



Table of Contents

Legal Notice	3
Dell Z9432F-ON Platform Introduction.....	4
ASIC Capabilities Supported on Netvisor ONE version 7.0.1	4
Transceivers, Optics, and CMIS support	5
Obtaining a Platform License.....	6
Dell Networking Firmware Updates	6
Installation and Configuration of Netvisor ONE	6
Installing Netvisor ONE Version 7.0.1 on Dell Z9432F-ON platforms.....	7
Fresh Install or Re-imaging with ONIE installer	7
Port Configuration Models	9
Command: port-breakout-config-modify	9
Command: port-fec-mode-modify	11
Guidelines or Limitations.....	13
Command: port-autoneg-mode-modify	13
Guidelines or Limitations.....	15
Command: port-xcvr-breakout-ability-show	16
Logical Port Numbering	18
Deployment Topology Recommendation.....	19
Functionality Differences on Z9432F-ON Platforms	20
Unsupported Features on Z9432F-ON Platforms.....	23
About Pluribus Networks.....	24

Legal Notice

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR PLURIBUS NETWORKS REPRESENTATIVE FOR A COPY.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE ARE PROVIDED “AS IS” WITH ALL FAULTS. PLURIBUS NETWORKS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL PLURIBUS NETWORKS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA, ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF PLURIBUS NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

©2022 Pluribus Networks, Inc. All rights reserved. Pluribus Networks, the Pluribus Networks logo, nvOS, Netvisor®, vManage, vRender, PluribusCare, FreedomCare™, Pluribus Cloud, and iTOR are registered trademarks or trademarks of Pluribus Networks, Inc., in the United States and other countries. All other trademarks, service marks, registered marks, registered service marks are the property of their respective owners. Pluribus Networks assumes no responsibility for any inaccuracies in this document. Pluribus Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Dell Z9432F-ON Platform Introduction

Netvisor ONE version 7.0.1 and later supports the Dell Z9432F-ON, which is a high-performance, high-density open networking 400GbE multi-rate aggregation switch with Trident 4 ASIC capabilities.

Figure 1 shows the port view of the Dell Z9432F-ON switch. For details on the switch capabilities, refer to the Dell Z9432F-ON datasheet.




Figure 1: Dell Z9432F-ON Platform Port View

ASIC Capabilities Supported on Netvisor ONE version 7.0.1

Netvisor ONE version 7.0.1 supports the following ASIC capabilities on Dell Z9432F-ON switch:

- Broadcom BCM56881 Trident4-X11 ASIC
- 32xQSFP56-DD ports backward compatible with QSFP28 (100GbE). That is, the switch can operate as 32x400GbE or 128x100GbE ports
- Each port has 8 serdes lanes that can operate at either 25Gbps or 50Gbps
- Netvisor ONE supports up to 4 logical ports per physical port or 128 logical ports per chassis
- Netvisor supports devices operating at 400GbE and 100GbE link speeds on the Z9432F-ON switch
 - 40GbE link speed is not supported on Netvisor ONE 7.0.1

Transceivers, Optics, and CMIS support

Netvisor ONE supports the following transceiver and cable types on a Dell Z9432F-ON switch. For additional details, refer to the [Transceiver Compatibility Matrix for Dell EMC Open Networking Switches](#).

Table 1: Transceivers and Cable Types

Device Type/Use Case	Description/Application	Dell Part Number
400G/QSFP56-DD - Straight Connection Options		
Passive Copper Cable (DAC) 0.5m, 1m, 2m reach	<ul style="list-style-type: none"> 400GbE QSFP56-DD to QSFP56-DD Passive Copper Direct Attach Cable 	<ul style="list-style-type: none"> DAC-Q56DD-400G-0.5M (W4JWV) DAC-Q56DD-400G-1M (22V6R) DAC-Q56DD-400G-2M (XR11M)
Active Copper Cable (ACC) 3-5m reach	<ul style="list-style-type: none"> 400GbE QSFP56-DD Active Copper Cable 	<ul style="list-style-type: none"> ACC-Q56DD-400G-3M (DH11M) ACC-Q56DD-400G-5M (86K95)
Optical Transceiver (FR4) 2km reach	<ul style="list-style-type: none"> 400GbE QSFP56-DD FR4 LC duplex 2km SMF 	<ul style="list-style-type: none"> Q56DD-400G-FR4 (78RHK)
4x100G/QSFP56-DD - Breakout Connection Options		
Active Copper Cable (ACC) 3-5m reach	<ul style="list-style-type: none"> 4x100GbE QSFP56-DD to 4xQSFP28 100GbE Active Copper Cable 	<ul style="list-style-type: none"> ACC-Q56DD-4Q28-100G-3M (0WDNV) ACC-Q56DD-4Q28-100G-5M (C0TP5)
Optical Transceiver (SR4.2) 70-100m reach	<ul style="list-style-type: none"> 4x100GbE QSFP56-DD SR4.2 BIDI* MPO12 70m/100m OM3/OM4 	<ul style="list-style-type: none"> Q56DD-400G-SR4.2 (0059F) <p>*peer device requires Dell 100GbE Bidi transceiver (Dell p/n 0X9CT) to interconnect with 4x100G SR4.2 using MPO12 to 4xLC MMF breakout cables (CBL-MPO12-4LC-OM4-xx).</p>

Obtaining a Platform License

To obtain your activation keys and activate your Dell Z9432F-ON platform, you must first register at <https://www.pluribusnetworks.com/get-started>.

Also, refer to the *Getting Started Guide* from the [Pluribus Technical Documentation](#) page.

Dell Networking Firmware Updates

To install Netvisor ONE on Dell Z9432F-ON platforms, you must have the correct firmware installed on the platform. For details, click on the download link under the Firmware category for the appropriate Dell Networking Firmware Updater BIOS version v3.51.0.D-9 or later:

<https://www.dell.com/support/home/en-in/product-support/product/networking-z9432f-on/drivers>

Installation and Configuration of Netvisor ONE

The installation and configuration documentation for Netvisor® ONE is located on the Pluribus Networks Web site:

<https://www.pluribusnetworks.com/support/technical-documentation/>

Installing Netvisor ONE Version 7.0.1 on Dell Z9432F-ON platforms

Fresh Install or Re-imaging with ONIE installer

1. Locate the ONIE image for Netvisor ONE Version 7.0.1 GA. This can be found on cloud-web.pluribusnetworks.com and then used in one of two ways:
 - a. Download the recommended ONIE installer image from cloud-web.pluribusnetworks.com. If the name is different, re-name the image as 'onie-installer' and copy it to the top level of a USB thumb drive, and attach it to the Netvisor ONE switch. Now, the switch automatically finds the onie-installer image and installs it.

Note: Ensure that the image name is *onie-installer*.

- b. Download the recommended *onie-installer* image from cloud-web.pluribusnetworks.com and copy it to a local Web server.
2. Connect to the serial console and reboot the switch.
 3. While the switch boots up, select the *onie* option from the grub menu.
 4. If you want to re-image the switch, select *onie-uninstall*.
 5. After uninstalling the onie-image, the switch reboots to the *onie-install* option.
 6. If you use an ONIE installer on a USB stick and you copied the file to the top-level directory, the switch automatically locates it and begins the installation.
OR
 7. If you want to use an ONIE installer kept on a local Web server in your network, you can start ONIE install using the command: `onie-nos-install <http-location-of-installer>`.

Note: On Dell Z9432F-ON platforms, you must first login with `root/<service_tag!>` credentials.

8. Wait for the ONIE installer to finish the installation and the switch automatically reboots.

Note: To install Netvisor ONE on the Z9432F-ON platforms, you must have installed the BIOS firmware version v3.51.0.D-9 or later, which is available for installation using Dell Networking Firmware Update 3.51.5.1-10 or later.

The ONIE installer checks the installed firmware version and if the firmware is not current, the ONIE install fails and displays the following message:

```
ERROR: Unsupported BIOS version 3.51.0.D-6. Please upgrade
firmware to Dell Networking Firmware Updater v3.51.5.1-10 or
greater. Aborting ONIE NOS install!!
```

If you get this message, upgrade the firmware using the Dell Firmware Updater version v3.51.5.1-10 or later, which includes the minimum BIOS firmware. This firmware updater is available on the [Dell website](#).

You can check the firmware version on the Z9432F-ON switch by:

To check BIOS version:

```
root@sag-colo-4:~# dmidecode -s bios-version
Version: 3.51.0.D-9
```

or

```
CLI (network-admin@z9432f) > switch-info-show
switch:                z9432f
model:                 Z9432F-ON
<snip>
bios-vendor:          American Megatrends Inc.
bios-version:         3.51.0.D-9
```


Port Configuration Models

The following port configurations are supported on Z9432F-ON platforms:

- Port speed and breakout configuration
- Port autoneg and FEC configuration modes
- Port cable type support including cable length
- Supported port configuration speeds
- Logical port numbering

For details on the port configuration of legacy switches, see the *Configuring Switch Ports* chapter of the *Configuration Guide*.

Note: Each bezel port on a Z9432F-ON platform supports: eight serial links at a line rate of 10.3125 Gbps to 56.25 Gbps (PAM4 56.25 Gbps or NRZ 28.125 Gbps).

The Z9432F-ON platform supports the following new CLI commands. For details on the legacy port configuration and commands, see the *Configuring Switch Ports* chapter of the *Netvisor ONE Configuration Guide*.

The Z9432F-ON specific commands in Netvisor ONE version 7.0.1 are:

- `port-breakout-config-modify` - configures geometry of a port
- `port-fec-mode-modify` - configures the FEC mode of a port
- `port-autoneg-mode-modify` - configures auto negotiation on a port
- `port-autoneg-status-show` - displays the auto negotiation mode on a port
- `port-xcvr-breakout-ability-show` - displays available xcvr breakout modes

Additionally, only for Z9432F-ON platforms, the existing `port-fec-status-show` command also displays the FEC mode of a port.

Command: `port-breakout-config-modify`

Starting with Netvisor ONE version 7.0.1, a new `port-breakout-config-modify` command identifies the port geometry on Z9432F-ON platforms and can be used to configure the port for the required speed.

Note: Additionally, the existing `port-config-modify port <port> speed <speed>` command should also be executed to set the port's speed configuration.

To configure port breakout, use the command:

```
CLI (network-admin@z9432f) > port-breakout-config-modify
```

<code>port-breakout-config-modify</code>	Configure the port breakout mode.
<code>port port-list</code>	The port number or list of ports
<code>breakout sys-default 400g-1x 100g-4x 100g-2x</code>	Specify the port breakout mode: System-default: each bezel port is 400G/PAM4 for the first two entries, then 100G/PAM4 for third entry, and 100G/NRZ for fourth entry.

To view the configuration details, use the command:

```
CLI (network-admin@z9432f) > port-breakout-config-show
```

<code>port-breakout-config-show</code>	Displays the port breakout configuration.
<code>port port-list</code>	The port number or list of ports
<code>breakout sys-default 400g-1x 100g-4x 100g-2x</code>	Specify the port breakout mode: System-default: each bezel port is 400G/PAM4 for the first two entries, then 100G/PAM4 for third entry, and 100G/NRZ for fourth entry.
<code>status sys-default 400g-1x 100g-4x 100g-2x</code>	Displays the breakout mode status.
<code>present no-present</code>	Indicates if a transceiver is present or not
<code>type unknown sfp qsfp28 qsfp56-dd sfp-dd</code>	Indicates the transceiver type.
<code>match unknown match not-match</code>	Indicates whether configured mode is a match for the installed transceiver.

Below is a sample configuration:

To view the port details on the switch:

```
CLI (network-admin@z9432f) > port-show port 5
switch port bezel-port config
-----
z9432f 5 2 fd,jumbo, fec, 400g
```

To configure the port breakout for port 5:

```
CLI (network-admin@z9432f) > port-breakout-config-modify port 5
breakout 100g-2x
Warning: Changing port breakout configure will cause port flap
and require reconfigure of all ports within the group.
Please confirm y/n (Default: n):y
```

Then, use the existing port-config-modify command to set the port speed:

```
CLI (network-admin@z9432f) > port-config-modify port 5 speed 100g
```

To verify the port configuration details, use the port-breakout-config-show command:

```
CLI (network-admin@z9432f) > port-breakout-config-show port 5

switch port breakout present type match
-----
z9432f 5      100g-2x  yes      qsf28 match
```

To view port details:

```
CLI (network-admin@z9432f) > port-show port 5

switch port bezel-port config
-----
z9432f 5      2          fd, 100g, jumbo, fec
```

Command: port-fec-mode-modify

The Z9432F-ON platforms provide different FEC modes for different port speeds and can be configured using the command:

```
CLI (network-admin@z9432f) > port-fec-mode-modify
```

port-fec-mode-modify	Associate a FEC mode with ports. Use this command to disable FEC or specify a user requested FEC mode.
port port-list	The port number or list of ports
mode sys-default none base-r rsfec rs544 rs544-2xn	Specify the port FEC mode: sys-default: depends on the configured speed. See default values in the table below. none: fec mode is disabled base-r: fec is using base-r fec rsfec: fec is using rsfec rs544: fec is using rs544 fec rs544-2xn: fec is using rs544-2xn Note: The FEC modes vary for different port speeds. See table below.

Note: By default, the FEC mode is enabled on Z9432F-ON platforms (on) for all speeds except 10G as opposed to being disabled on all legacy switches.

The below table provides the supported FEC mode configurations for different speeds (when *auto negotiation* is off for all below speeds):

Table 2: Port Breakout Vs Supported FEC modes

Port Breakout	Allowed FEC Configurations
400g-1x	rs544-2xn (sys-default)
100g-4x (50g Lane)	none rsfec rs544 (sys-default)
100g-2x (25g Lane)	none rsfec (sys-default) rs544

To view the configuration details, use the command:

```
CLI (network-admin@z9432f) > port-fec-mode-show port x
```

Below is a sample configuration:

```
CLI (network-admin@z9432f) > port-fec-mode-show port 5
```

```
switch port mode
-----
z9432f 1-130 sys-default
```

```
CLI (network-admin@z9432f) > port-fec-mode-modify port 5 mode
sys-default
none
base-r
rsfec
rs544
rs544-2xn
```

```
CLI (network-admin@z9432f) > port-fec-mode-modify port 5 mode
rs544-2xn
```

```
CLI (network-admin@z9432f) > port-fec-mode-show port 5
switch port mode
-----
z9432f 5 rs544-2xn
```

```
CLI (network-admin@z9432f) > port-fec-status-show port 5
```

```
switch port configured_mode negotiated_mode
-----
z9432f 5 RS544-2xn RS544-2xn
```

Guidelines or Limitations

On Z9432F-ON platforms, do not attempt to use the `fec/no-fec` keyword in the `port-config-modify` command (used in the case of legacy platforms) as the usage of this keyword is blocked. Instead, use the `port-fec-mode-modify` command. For example:

```
CLI (network-admin@z9432f) > port-config-modify port 5 fec
port-config-modify: Failed to set fec of port 5. Please use port-
fec-mode-modify to set the desired mode.
```

```
CLI (network-admin@z9432f) > Port-config-modify port 1 no-fec
```

This command does not display any error, but no functional change to the FEC state is done. This is because the legacy platforms require FEC mode to be part of the port configuration, but is always in *disable* state. Hence, the `no-fec` mode does not display any error but initiates a silent success. However, this command is not applicable on Z9432F-ON platforms.

Command: port-autoneg-mode-modify

The Z9432F-ON platforms support configuration of multiple types of autoneg and any particular flavor of *autoneg* is not unique to the port speed. Hence the existing `port-config-modify` command is not sufficient to configure autoneg on Z9432F-ON platforms.

Netvisor ONE provides different auto negotiation modes depending on the port's speed on Z9432F-ON platforms. To configure autoneg, use the command:

```
CLI (network-admin@z9432f) > port-autoneg-mode-modify port x mode
Y
```

<code>port-autoneg-mode-modify</code>	Configure autoneg on ports.
<code>port port-list</code>	The port number or list of ports
<code>mode sys-default none c173 c173- msa</code>	Specify the port autoneg mode: <code>sys-default</code> : depends on the configured speed. See default values below. <code>none</code> : auto negotiation mode is disabled <code>c173</code> : using clause 73 <code>c173-msa</code> : use clause 73 multi-source agreement Note: The FEC modes vary for different port speeds. See table below.

Note: On Z9432F-ON platforms, the auto negotiation is enabled (on) for DAC cables running at all speeds. See *default values*.

To view the configuration details, use the command:

```
CLI (network-admin@z9432f) > port-autoneg-mode-show
```

The below table provides the supported FEC mode and autoneg configurations for different speeds (when *auto negotiation* is on in all below speeds):

Table 3: Supported FEC and Augoneg modes for various Port Speed

Port Speed	FEC configurations	Autoneg
400g-1x	rS544-2xn	CL73-MSA
100g-4x (50g Lane)	none rsfec rs544	CL73
100g-2x (25g Lane)	none rsfec rs544	CL73

The system default value for auto negotiation mode depends on the configured speed and DAC cable:

- 400G-1x: cl73-msa
- 100G-4x: cl73
- 100G-2x: cl73

Note: When changing the configuration of port breakout mode, it is advantageous to revert the port autoneg and fec settings to `sys-default` to avoid an error due to the `autoneg` and/or `fec` setting not being valid for the new port breakout mode.

Below is a sample configuration:

```
CLI (network-admin@z9432f) > port-autoneg-status-show port 81
```

```
switch port configured_mode negotiated_mode
-----
z9432f 81 none none
```

```
CLI (network-admin@z9432f) > port-autoneg-mode-show port 81
```

```
switch port mode
-----
z9432f 1-130 sys-default
```

```
CLI (network-admin@z9432f) > port-autoneg-mode-modify port 81
```

```
mode
sys-default
none
cl73
cl73-msa
```

```
CLI (network-admin@z9432f) > port-autoneg-mode-modify port 81
mode cl73
```

```
CLI (network-admin@z9432f) > port-autoneg-mode-show port 81
switch port mode
-----
z9432f 81    cl73
```

```
CLI (network-admin@z9432f) > port-autoneg-status-show port 81
switch port configured_mode negotiated_mode
-----
z9432f 81    cl73                none
```

In the above example, even though the *auto-negotiation* mode is configured, the link has not negotiated with this mode, possibly due to the port being disabled on either peer, or due to a mismatch in *autoneg* configuration between peers.

Guidelines or Limitations

The *autoneg/no-autoneg* keywords in *port-config-modify* are blocked for Z9432F-ON platforms.

On Z9432F-ON platforms, do not attempt to use the *autoneg/no-autoneg* keyword in the *port-config-modify* command (used in the case of legacy platforms) as the usage of this keyword is blocked. Instead, use the *port-autoneg-mode-modify* command. For example:

```
CLI (network-admin@z9432f) > port-config-modify port 81 autoneg
port-config-modify: Failed to set autoneg of port 81. Please use
port-autoneg-mode-modify to set the desired mode.
```

```
CLI (network-admin@z9432f) > Port-config-modify port 1 no-autoneg
```

This command does not display any error, but no functional change to the *autoneg* mode is done. This is because the legacy platforms require *autoneg* mode to be part of the port configuration, but is always in *disable* state. Hence, the *no-autoneg* mode does not display any error but initiates a silent success. However, this command is not applicable on Z9432F-ON platforms.

Command: port-xcvr-breakout-ability-show

To view the vendor-programmed transceiver application ID information for CMIS QSFP56-DD modules on Z9432F-ON platforms, use the `port-xcvr-breakout-ability-show <port-list>` command.

Note: In Netvisor ONE version 7.0.1, the `port-xcvr-breakout-ability-show <port-list>` command is available only on Z9432F-ON platforms with QSFP-DD connector ports.

```
CLI (network-admin@z9432f) > port-xcvr-breakout-ability-show
```

<code>port-xcvr-breakout-ability-show</code>	Displays available xcvr breakout modes.
<code>port port-list</code>	The port number or list of ports.
<code>port(s) port-list</code>	The list of logical breakout port(s).
<code>bezel-port bezel-port-string</code>	The associated bezel port.
<code>vendor-name vendor-name-string</code>	The transceiver vendor name.
<code>part-number part-number-string</code>	The transceiver part number.
<code>breakout sys-default 400g-1x 100g-4x 100g-2x</code>	The port breakout mode.
<code>avail Yes No</code>	Breakout is available via Netvisor ONE.
<code>advert advert-string</code>	CMIS application descriptor record number.
<code>app-record-hex app-record-hex-string</code>	Raw hex application record.
<code>host-ifc host-ifc-string</code>	Host Interface.
<code>media-type media-type-string</code>	Media Type.
<code>media-ifc media-ifc-string</code>	Media Interface.
<code>status status-string</code>	Sys-default and current breakout flags

Below is a sample configuration output for the `port-xcvr-breakout-ability-show` command:

- On port 9, AOC with 4 app records:

```
CLI (network-admin@z9432f) > port-xcvr-breakout-ability-show port
9
switch port(s) bezel-port vendor-name part-number breakout avail advert host-ifc media-type status
-----
9          3          DELL EMC    9GMDY      400g-1x   Yes   app-01 400G-GAUI8 active-cable sys-default
9,11       3          DELL EMC    9GMDY      100g-2x   Yes   app-02 CAUI4      active-cable
9-12       3          DELL EMC    9GMDY      100g-4x   Yes   app-03 100G-GAUI2 active-cable current
9-12       3          DELL EMC    9GMDY      50g-4x    No    app-04 50GAUI    active-cable
```



```
CLI (network-admin@z9432f) > switch-local
```

```
CLI (network-admin@z9432f*) > port-xcvr-breakout-ability-show
port 9
```

port(s)	bezel-port	vendor-name	part-number	breakout	avail	advert	host-ifc	media-type	status
9	3	DELL EMC	9GMDY	400g-1x	Yes	app-01	400G-GAUI8	active-cable	sys-default
9,11	3	DELL EMC	9GMDY	100g-2x	Yes	app-02	CAUI4	active-cable	
9-12	3	DELL EMC	9GMDY	100g-4x	Yes	app-03	100G-GAUI2	active-cable	current
9-12	3	DELL EMC	9GMDY	50g-4x	No	app-04	50GAUI	active-cable	

- On port 17, an optical module with 2 app records:

```
CLI (network-admin@z9432f*) > port-xcvr-breakout-ability-show
port 17
```

port(s)	bezel-port	vendor-name	part-number	breakout	avail	advert	host-ifc	media-type	status
17	5	DELL EMC	TP2F0	400g-1x	Yes	app-01	400G-GAUI8	mmf	sys-default
17,19	5	DELL EMC	TP2F0	200g-2x	Yes	app-02	200G-GAUI4	mmf	current

- On port 41, a DAC with only one app record but with additional breakouts supported by Netvisor:

```
CLI (network-admin@z9432f*) > port-xcvr-breakout-ability-show
port 41
```

port(s)	bezel-port	vendor-name	part-number	breakout	avail	advert	host-ifc	media-type	status
41	11	DELL EMC	XR11M	400g-1x	Yes	app-01	400GBASE-CR8	passive-copper	sys-default
41,43	11	DELL EMC	XR11M	200g-2x	Yes	implicit			
41-44	11	DELL EMC	XR11M	100g-4x	Yes	implicit			current

- On port 89, a module that has 2 app records, only one of which is supported:

```
CLI (network-admin@z9432f*) > port-xcvr-breakout-ability-show
port 89
```

port(s)	bezel-port	vendor-name	part-number	breakout	avail	advert	host-ifc	media-type	status
89-92	23	DELL EMC	0059F	100g-4x	No	app-01	100G-GAUI2	mmf	
89-92	23	DELL EMC	0059F	100g-4x	Yes	app-02	100G-GAUI2	mmf	sys-default current

Logical Port Numbering

See **Table 4** below for details on the mapping of logical port with bezel port numbering on Z9432F-ON platforms:

Table 4: Mapping of Logical Port Number with the Bezel Port Number

Logical Port #	Bezel Port#	Logical Port #	Bezel Port#	Logical Port #	Bezel Port#	Logical Port #	Bezel Port#
1	1	17	5	33	9	49	13
2	1.2	18	5.2	34	9.2	50	13.2
3	1.3	19	5.3	35	9.3	51	13.3
4	1.4	20	5.4	36	9.4	52	13.4
5	2	21	6	37	10	53	14
6	2.2	22	6.2	38	10.2	54	14.2
7	2.3	23	6.3	39	10.3	55	14.3
8	2.4	24	6.4	40	10.4	56	14.4
9	3	25	7	41	11
10	3.2	26	7.2	42	11.2
11	3.3	27	7.3	43	11.3	125	32
12	3.4	28	7.4	44	11.4	126	32.2
13	4	29	8	45	12	127	32.3
14	4.2	30	8.2	46	12.2	128	32.4
15	4.3	31	8.3	47	12.3	129*	33
16	4.4	32	8.4	48	12.4	130*	34

***Note:** In the above table the ports 129 and 130 are 10G SFP ports.

Deployment Topology Recommendation

The following deployment scenarios are recommended in Netvisor ONE version 7.0.1:

With Netvisor ONE version 7.0.1, Dell Z9432F-ON switch can be deployed as a spine switch connecting via 100GbE links to non-Z9432F-ON leaf switches in a fabric. You can also configure non-Z9432F-ON switches as spine switches in the same fabric.

Note: However, in the fabric, you must **not** configure the software features that are unsupported on Z9432F-ON platforms in Netvisor version 7.0.1.

Note: In the case of an existing fabric with software features that are unsupported on Z9432F-ON switches, deploy the Z9432F-ON switch in a separate fabric and then connect this new fabric to the existing deployment.

Functionality Differences on Z9432F-ON Platforms

Feature functionality differences on the Z9432F-ON platforms in Netvisor ONE version 7.0.1 as compared to other platforms:

- vFlow/QoS (for more configuration details please refer to the *Configuring vFlow*, *Configuring QoS*, and *Configuring CPTP* chapters in the Configuration Guide).
 - A new vFlow table is introduced called `Ingress-Policer-Table-1-0`. The `vflow-table-show` (see below) and `vflow-table-profile-show` commands display the new table.
 - The `max-bw` vFlow/QoS action is supported in the new `Ingress-Policer-Table-1-0` as well as in the `Egress-Table-1-0`.
 - The `copy-to-port` vFlow action is not supported in the ingress tables (you can use a port mirror instead)
 - The `setvlan` action is supported in the ingress and egress tables (such as `System-L1-L4-Tun`, `IPv6-Table-1-0`, `Egress-Table-1-0`, `Egress-Table-v6-1-0`), but not in the `System-VCAP-table-1-0`.
 - The `set-metadata` action is supported in the ingress tables (`System-L1-L4-Tun` and `IPv6-Table-1-0`), but not in the `System-VCAP-table-1-0`.
 - The `vlan` qualifier is not supported in the `Egress-Table-1-0` and `Egress-Table-v6-1-0` (for it you can use the `System-VCAP-table-1-0` or the ingress tables)
 - The `set-dscp` action is not supported in the `Egress-Table-1-0` and `Egress-Table-v6-1-0` (use the ingress tables instead)
 - The `proto=ICMPv6` and `proto=IP` qualifiers are not supported in the `Egress-Table-v6-1-0`.
 - For CPTP classes a total of 48 queues are supported. Since one additional custom CPU class is reserved for internal use, only one user-defined CPTP class is supported with the `cpu-class-create` command.
 - The vFlow action `set-vlan-pri` is not supported in `System-VCAP-Table-1-0`. Only the `Ingress-Policer-Table-1-0` and `System-L1-L4-Tun-1-0` tables support it. For example:

```
CLI (network-admin@switch) > vflow-table-show
```

name	flow-max-per-group	flow-used	flow-tbl-slices	capability	flow-profile
Egress-Table-1-0	1023	0	3	match-metadata	system
Egress-Table-v6-1-0	1023	0	1	none	egress-v6
Ingress-Policer-Table-1-0	6143	0	2	none	system
IPv6-Table-1-0	3071	0	2	none	ipv6
System-L1-L4-Tun-1-0	3071	51	4	set-metadata	system
APP-vlan-L1-L4-1-0	12287	0	1	set-metadata	application
System-VCAP-table-1-0	511	0	3	none	system
VCAP-IPv6-table-1-0	255	0	1	none	ipv6-vcap

Note: The `flow-max-per-group` values vary and the above values may not be correct (PR 43479). See Release Notes for Netvisor ONE version 7.0.1 for details.

To configure and view the vFlow details:

```
CLI (network-admin@switch) > vflow-create name band scope local
in-port 129 bw-max 1M table-name Ingress-Policer-Table-1-0
precedence 12
```

```
CLI (network-admin@switch) > vflow-show
```

name	scope	type	in-port	bw-max	precedence	transparency	tracking	table-name
band	local	vflow	129	1M	12	disable	disable	Ingress-Policer-Table-1-0

- The below new header/packet offsets are supported in UDF objects:
 - 14-outer-tcp
 - 14-outer-udp
 - 13-outer-v4
 - 13-outer-v6
 - packet-start-untagged
 - packet-start-tagged
- Note that packet-start is no longer a valid UDF header value and therefore there is no default value for the header field in the udf-create command. The packet-start keyword has been replaced by two new parameters: packet-start-untagged and packet-start-tagged, but they are not equivalent to packet-start. They differ in the following ways:
 - The packet-start-untagged parameter must be used to match untagged traffic, whereas the packet-start-tagged parameter must be used to match tagged traffic.
 - The packet-start-untagged parameter corresponds to a base offset from the start of the first octet of the Ethertype field in the Ethernet header, whereas the packet-start-tagged parameter corresponds to a base offset from the start of the first octet of the (outer) VLAN tag in the Ethernet header. Both are different from packet-start which corresponds to a base offset from the start of the first octet of the Ethernet header.
 - The packet-start-untagged and packet-start-tagged parameters can no longer be used to match the source and destination MAC addresses of a packet.
- Switch ports support 4 dedicated queues for multicast, broadcast and unknown unicast (BUM, in short) traffic. The default weight of each queue is:

```
CLI (network-admin@switch) > port-cos-weight-show
cos8-weight: 32
cos9-weight: 64
cos10-weight: 64
cos11-weight: 127
```

- For BUM traffic CoS priorities from 0 to 4 are mapped to queue 8, CoS 5 is mapped to queue 9, CoS 6 is mapped to queue 10, and CoS 7 is mapped to queue 11.
- Queues 8-11 can be configured like other queues: weights and bandwidth can be set and verified with the port-cos-weight-modify/show and port-cos-bw-modify/show commands, while the queue stats can be displayed with the port-cos-stats-show command.

- The functionality of the port and vFlow commands has remained unchanged. The only difference is the addition of support for queues 8-11. For example, queue 8's bandwidth limit can be configured on port 1 like so:

```
CLI (network-admin@switch*) > port-cos-bw-modify port 1 cos 8
max-bw-limit 90
```

```
CLI (network-admin@switch*) > port-cos-bw-show cos 8,9,10,11
```

```
cos port                min-bw-guarantee max-bw-limit weight
-----
8 3,5,7,9,13,15,17,19,21,25,29,33-37 0%                100%          32
8 1 0%                90%           32
9 1,3,5,7,9,13,15,17,19,21,25 0%                100%          64
10 1,3,5,7,9,13,15,17,19,21,25 0%                100%          64
11 1,3,5,7,9,13,15,17,19,21,25 0%                100%         127
```

- Adding ports to vxlan-loopback trunk is not required (single pass only)
- Maximum vFlow table entries based on hardware constraints (note that the Z9432F-ON platforms support the flexible allocation of resources to each vFlow table so a fixed pool of resources may not be reserved to a flow table, in contrast to older platforms):

vFlow Tables	Max. Table entries
System-VCAP-table-1-0	511 allocatable entries
VCAP-IPv6-table-1-0	255 allocatable entries
Ingress-Policer-Table-1-0	6143 allocatable entries
System-L1-L4-Tun-1-0	3071 allocatable entries (of which 51 are reserved system entries and the remainder is shared with the other ingress tables)
IPv6-Table-1-0	2047 allocatable entries
APP-vlan-L1-L4-1-0	12287 allocatable entries
Egress-Table-1-0	1023 allocatable entries
Egress-Table-v6-1-0	1023 allocatable entries

Unsupported Features on Z9432F-ON Platforms

The following features are not supported on Z9432F-ON Platforms in Netvisor ONE version 7.0.1:

- Features not supported in 7.0.1 release – CLI commands may be present but are not supported:
 - L1 virtual wire
 - Symmetric hashing
 - sFlow
 - Private vNET/VLAN CLI
 - Open vSwitch
 - Policy Based Routing (PBR)
 - VXLAN end-points (that is, hardware VTEPs are not supported). However, transporting VXLAN overlay traffic as an underlay switch is supported.
 - IGMP Snooping with VXLAN
 - Unicast and Multicast Distributed VRFs
 - Bridge domains
 - QinQ
 - PDU Transparency (a.k.a. Layer 2 Protocol Tunneling)
 - TPID configuration
 - EVPN
 - vLE
 - vPG
 - vFlow statistics currently don't support byte and packet counts for policer drops
 - CPTP with cpu hog protection does not work for BFD protocol packets
 - The `port-buffer-stats-show` command is not supported
 - The `system-setting-modify manage-broadcast on setting` is not supported.

About Pluribus Networks

Pluribus Networks delivers an open, controllerless software-defined network fabric for modern data centers, multi-site data centers, and distributed cloud edge environments.

The Linux-based Netvisor® ONE operating system and the Adaptive Cloud Fabric™ have been purpose-built to deliver radically simplified networking and comprehensive visibility along with white box economics by leveraging hardware from our partners Dell EMC, Edgecore, Celestica, and Champion ONE, as well as Pluribus' own Freedom™ Series of switches.

The Adaptive Cloud Fabric provides a fully automated underlay and virtualized overlay with comprehensive visibility and brownfield interoperability and is optimized to deliver rich and highly secure per-tenant services across data center sites with simple operations having no single point of failure.

Further simplifying network operations is Pluribus UNUM™, an agile, multi-functional web management portal that provides a rich graphical user interface to manage the Adaptive Cloud Fabric. UNUM has two key modules - UNUM Fabric Manager for provisioning and management of the fabric and UNUM Insight Analytics to quickly examine billions of flows traversing the fabric to ensure quality and performance.

Pluribus is deployed in more than 275 customers worldwide, including the 4G and 5G mobile cores of more than 75 Tier 1 service providers delivering mission-critical traffic across the data center for hundreds of millions of connected devices. Pluribus is networking, simplified.

For additional information contact Pluribus Networks at info@pluribusnetworks.com, or visit www.pluribusnetworks.com.

Follow us on Twitter [@pluribusnet](https://twitter.com/pluribusnet) or LinkedIn at <https://www.linkedin.com/company/pluribus-networks/>.

Corporate Headquarters

Pluribus Networks, Inc.
5201 Great America Parkway, Suite 422
Santa Clara, CA 95054
1-855-438-8638 / +1-650-289-4717

India Office

Pluribus Networks India Private Limited
Indique Brigade Square, 4th Floor
21, Cambridge Road
Bangalore 560008

Document Version - July 2022