



## The science behind the report:

# Managing VLANs from advanced partitioned NICs: Dell PowerEdge MX760c servers with Broadcom 57504 Quad-Port 25GbE Mezzanine Adapter using Advanced NPAR

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report [Managing VLANs from advanced partitioned NICs: Dell PowerEdge MX760c servers with Broadcom 57504 Quad-Port 25GbE Mezzanine Adapter using Advanced NPAR](#).

We concluded our hands-on testing on May 4, 2023. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on May 4, 2023 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

The parts numbers for ordering the solution we tested from their respective vendors are as follows:

- Dell: 555-BGGB
- Broadcom: BCM957504-NGM250D

## System configuration information

Table 1: Configuration information for the Dell™ PowerEdge™ MX7000 server enclosure.

Server enclosure configuration information	Dell PowerEdge MX7000
Number of management modules	2
Management module firmware revision	2.10.00 (Build A00.116)
First type of I/O module	
Vendor and model number	Dell MX9116n Fabric Switching Engine
I/O module firmware revision	10.5.5.2.14 3.35.5.1-24
Number of modules	2
Occupied bay(s)	A1, A2
Second type of I/O module	
Vendor and model number	Dell MXG610s FC Switch 2SFP
I/O module firmware revision	8.1.0_Inx2
Number of modules	1
Occupied bay(s)	C1
Power supplies	
Vendor and model number	Dell
Number of power supplies	6
Wattage of each (W)	3,000
Cooling fans	
Vendor and model number	Dell
Number of fans	9

Table 2: Detailed configuration information for the servers under test.

Server configuration information	Dell PowerEdge MX760c
BIOS name and version	Dell 1.8.2
Non-default BIOS settings	Intel Turbo Boost enabled, Virtualization enabled
Operating system name and version/build number	DellEMC-VMware ESXi 7.0 Update 3 Build-20036589 (A06)
Date of last OS updates/patches applied	04/20/2023
Power management policy	Performance
Processor	
Number of processors	2
Vendor and model	Intel® Xeon® Gold 6430
Core count (per processor)	32
Core frequency (GHz)	2.10

Server configuration information		Dell PowerEdge MX760c
Memory module(s)		
Total memory in system (GB)	512	
Number of memory modules	8	
Vendor and model	Hynix HMCG94AEBRA109N	
Size (GB)	64	
Type	PC5-38400 (DDR5-4800)	
Speed (MHz)	4,800	
Speed running in the server (MHz)	4,400	
Network adapter		
Vendor and model	Broadcom BCM957504-N425G Adv Quad 25Gb Ethernet	
Number and type of ports	4 x 25GbE	
Driver version	Dell FW 22.31.13.70	
Local storage		
Number of drives	2	
Drive vendor and model	Dell EC NVMe™ ISE 7400 RI M.2	
Drive size (GB)	480	
Drive information (speed, interface, type)	NVMe 1.4, PCIE, M.2 SSD	
Number of drives	4	
Drive vendor and model	Dell Ent NVMe P5600 MU U.2	
Drive size (GB)	3,200	
Drive information (speed, interface, type)	NVMe 1.3, PCIE, SSD	

## How we tested

Our goal was to show the Advanced feature on the Dell PowerEdge MX chassis and the functionality of Dell OpenManage Enterprise Modular (OME-M) to support it. For testing, we used a Dell PowerEdge MX7000 chassis with three servers, a Dell MX9116n Fabric Switching Engine, and Broadcom BCM957504-N425G Adv Quad 25Gb Ethernet NIC.

Table 3: Components and their firmware versions.

Component	Type or firmware version
Dell PowerEdge MX chassis	2.10.00
Dell PowerEdge servers	PowerEdge MX760c
iDRAC	IDRAC9 Version 6.10.05.00(Build 53)
Lifecycle Controller	6.10.05.00
Broadcom BCM957504-N425G Adv Quad 25Gb Ethernet NIC	22.31.13.70
Dell PowerEdge MX switch	Dell MX9116n Fabric Switching Engine

For the hypervisor, we used DellEMC-VMware ESXi 7.0 Update 3 Build-20036589 (A06).

To enable Advanced NPAR on our Broadcom NICs, we first updated firmware on our bare-metal servers. This allowed us to access the necessary settings within the Lifecycle Controller to enable Advanced NPAR on the NICs themselves. We also had to ensure that we disabled NIC+RDMA. Then, we built a profile to deploy Advanced NPAR to multiple servers and test advanced NPAR functionality.

To test the NPAR functionality we used two modes: SmartFabric mode and Full Switch mode. Each mode is mutually exclusive, so as we operated in one mode, the other was disabled. In the sections below, we detail the steps for enabling Advanced NPAR and how we configured Full Switch mode. Advanced NPAR functionality is the same for both SmartFabric and Full Switch modes.

To enable Advanced NPAR with Dell OME-M, servers must have advanced NPAR already enabled, and a configured profile template. Once a server has the necessary settings, admins can easily replicate Advanced NPAR settings through OME-M. In SmartFabric mode, admins may need to create an uplink and the tag VLANs appropriately for the network; however, most networking setups already have an existing uplink.

To create a server template, we first selected our original bare-metal server from the Template menu, and used it as a reference. Once we created the template, we edited it under networking, assigning an untagged VLAN for each NPAR port and specifying the correct VLAN. At this point in the process, Advanced NPAR works, though it may not be correctly configured. Additionally, OME-M propagates changes to the configuration almost instantly, with no significant wait times. Once the profile is configured correctly, OME-M can apply it to any number of servers, handling the application, configuration, and restart of the desired machines.

## Advanced NPAR with Dell OME Modular

### Configuring the bare-metal server

1. Navigate to OME Modular.
2. Log in with admin credentials.
3. From the homepage, mouse over Devices, and select Compute.
4. Click the name of the sled that will be used as the Template.
5. Click Launch Virtual Console.
6. From the Virtual Console, click Boot, select Lifecycle Controller, and click Yes.
7. Click Power, click Reset System (Warm Boot), and click Yes.

### Enabling Advanced NPAR

1. From the Lifecycle Controller landing page, click System Setup.
2. Click Advanced Hardware Configuration.
3. Click Device Settings.
4. Click the desired NIC Port for NPAR (in our case, Port 1).
5. Click Device Level Configuration.
6. From the Virtualization Mode drop-down list, select Advanced NPAR.
7. If NPAREP is disabled, click Enable.
8. Click Back.

## Configuring NIC partitioning

1. Click NIC Partitioning Configuration.
2. Click Partition 1 Configuration.
3. Ensure NIC+RDMA is set to Disabled.
4. Click Back.
5. Click Finish.
6. Click Finish.
7. Click Finish.
8. Click Exit, and click Yes.

## Creating an uplink

1. Return to OME Modular.
2. Mouse over Devices, and select Fabric.
3. Click the Fabric being modified.
4. Click Add Uplink.
5. Assign a Name, and add a description of the uplink.
6. In the Uplink Type: drop-down list, select Ethernet - No Spanning Tree.
7. Click Next.
8. Select a switch port to add ports from.
9. Select the appropriate switch ports.
10. Select the VLANs that will be used in Advanced NPAR.
11. Leave VLAN 1 untagged.
12. Click Finish.
13. Create and Deploy Server Template.
14. Mouse over Configuration, and click Templates.
15. Click Create Template, and click From Reference Device.
16. Give the Template a name and a short description.
17. Click Next.
18. Click Select Device.
19. Click Server.
20. Click the Server that was previously set up to be the template.
21. Click Finish.
22. Click Finish.

## Configuring the template

1. Click the newly created template.
2. Click Edit, and click Edit Template.
3. Click Edit Components.
4. Click Networking.
5. Scroll down and click the pencil icon next to the name of the displayed NIC.
6. On each port, set the untagged VLAN to the correct untagged VLAN.
7. On each Port, click Select VLANs, select the correct VLANs, and click the double arrow to activate them.
8. Click Finish.
9. Click Finish.
10. Click Next.
11. Click Finish.

Note: Servers receiving the profile must be up to date with patches to enable Advanced NPAR.

## Deploying the template to other servers

1. From the Template page, click on the template created in the previous steps.
2. Click Deploy Template.
3. Click Yes.
4. Click Deploy to Devices, and click Select Sheds.
5. Click Server.
6. Click the Server(s) that are eligible for Advanced NPAR and will receive the profile.
7. Click Finish.

## Configuring SmartFabric mode

### Creating a SmartFabric

1. Navigate to OME-M.
2. Click Devices, and click Fabric.
3. Click Add Fabric.
4. Give the Fabric a name and a description, and click Next.
5. On Design, select Fabric Switching Engines in same chassis.
6. Under Chassis-X, select the correct Chassis.
7. Under Switch-A, pick the correct switch.
8. Under Switch-B, pick the other correct switch.
9. Click Next.
10. Click Finish.

### Configuring SmartFabric

1. Click the name of the fabric created in the previous steps.
2. Click Add Uplink.
3. Give the Uplink a name and Description.
4. From the Uplink Type drop down, select Ethernet - No Spanning Tree.
5. Click Next.
6. On Define, select the correct switch ports for the uplink to function on.
7. Select the desired VLANs under Tagged Networks.
8. Set the untagged network to the correct VLAN for untagged.
9. Click Finish.

## Configuring Full Switch mode

The configuration process for Full Switch mode is similar to a normal switch, as both are configured via the command-line interface.

```
# All commands run on both switches except where noted

# Configure hostname and set spanning tree mode to RSTP

MX9116N-A1# configure terminal
(switch-1) MX9116N-A1(config)# hostname MX9116N-A1
(switch-2) MX9116N-A2(config)# hostname MX9116N-A2
MX9116N-A1(config)# spanning-tree mode rstp
MX9116N-A1(config)# ip name-server 172.16.0.10

MX9116N-A1(config)# no multicast snooping flood-restrict
MX9116N-A1(config)# scale-profile vlan

# Configure ISL on ports 37-40

MX9116N-A1(config)# interface range ethernet1/1/37-1/1/40
MX9116N-A1(conf-range-eth1/1/37-1/1/40)# description ISL
MX9116N-A1(conf-range-eth1/1/37-1/1/40)# no switchport
MX9116N-A1(conf-range-eth1/1/37-1/1/40)# negotiation on
MX9116N-A1(conf-range-eth1/1/37-1/1/40)# exit

MX9116N-A1(config)# vlt-domain 255
(switch-1) MX9116N-A1(conf-vlt-255)# backup destination 172.16.1.12
(switch-2) MX9116N-A2(conf-vlt-255)# backup destination 172.16.1.13
MX9116N-A1(conf-vlt-255)# discovery-interface ethernet1/1/37-1/1/40
MX9116N-A1(conf-vlt-255)# exit

# Configure VLANs (in this example: 221, 1000, 2000, and 3000)

MX9116N-A1(config)# interface vlan221
MX9116N-A1(conf-if-vl-221)# no shutdown
```

```

MX9116N-A1(conf-if-vl-221)# no ipv6 mld snooping
MX9116N-A1(conf-if-vl-221)# no ip igmp snooping
MX9116N-A1(conf-if-vl-221)# exit

MX9116N-A1(config)# interface vlan1000
MX9116N-A1(conf-if-vl-1000)# no shutdown
MX9116N-A1(conf-if-vl-1000)# no ipv6 mld snooping
MX9116N-A1(conf-if-vl-1000)# no ip igmp snooping
MX9116N-A1(conf-if-vl-1000)# exit

MX9116N-A1(config)# interface vlan2000
MX9116N-A1(conf-if-vl-2000)# description vMotion
MX9116N-A1(conf-if-vl-2000)# no shutdown
MX9116N-A1(conf-if-vl-2000)# no ipv6 mld snooping
MX9116N-A1(conf-if-vl-2000)# no ip igmp snooping
MX9116N-A1(conf-if-vl-2000)# exit

MX9116N-A1(config)# interface vlan3000
MX9116N-A1(conf-if-vl-3000)# description vSAN
MX9116N-A1(conf-if-vl-3000)# no shutdown
MX9116N-A1(conf-if-vl-3000)# no ipv6 mld snooping
MX9116N-A1(conf-if-vl-3000)# no ip igmp snooping
MX9116N-A1(conf-if-vl-3000)# exit

# Configure uplink ports to ToR switches
# Set ports to 40g mode
MX9116N-A1(config)# port-group 1/1/13
MX9116N-A1(conf-pg-1/1/13)# mode Eth 40g-1x
MX9116N-A1(conf-pg-1/1/13)# exit
MX9116N-A1(config)# port-group 1/1/14
MX9116N-A1(conf-pg-1/1/14)# mode Eth 40g-1x
MX9116N-A1(conf-pg-1/1/14)# exit

# Configure port-channel
MX9116N-A1(config)# interface port-channel2
MX9116N-A1(conf-if-po-2)# no shutdown
MX9116N-A1(conf-if-po-2)# switchport mode trunk
MX9116N-A1(conf-if-po-2)# switchport access vlan 1
MX9116N-A1(conf-if-po-2)# switchport trunk allowed vlan 221,1000,2000,3000
MX9116N-A1(conf-if-po-2)# vlt-port-channel 2
MX9116N-A1(conf-if-po-2)# exit

# Assign interfaces to port-channel
MX9116N-A1(config)# interface range ethernet 1/1/41:1-1/1/42:1
MX9116N-A1(conf-range-eth1/1/41:1-1/1/42:1)# channel-group 2 mode active
MX9116N-A1(conf-range-eth1/1/41:1-1/1/42:1)# negotiation off
MX9116N-A1(conf-range-eth1/1/41:1-1/1/42:1)# exit

# Configure interfaces for Advanced NPAR mode:

MX9116N-A1(config)# interface ethernet1/1/1
MX9116N-A1(conf-if-eth1/1/1)# NPAR partition-4
MX9116N-A1(conf-if-eth1/1/1)# switchport mode trunk
MX9116N-A1(conf-if-eth1/1/1)# spanning-tree disable
MX9116N-A1(conf-if-eth1/1/1)# exit

MX9116N-A1(config)# interface ethernet1/1/2
MX9116N-A1(conf-if-eth1/1/2)# NPAR partition-4
MX9116N-A1(conf-if-eth1/1/2)# switchport mode trunk
MX9116N-A1(conf-if-eth1/1/2)# spanning-tree disable
MX9116N-A1(conf-if-eth1/1/2)# show configuration
MX9116N-A1(conf-if-eth1/1/2)# exit

MX9116N-A1(config)# interface ethernet1/1/5
MX9116N-A1(conf-if-eth1/1/5)# NPAR partition-4
MX9116N-A1(conf-if-eth1/1/5)# switchport mode trunk
MX9116N-A1(conf-if-eth1/1/5)# spanning-tree disable
MX9116N-A1(conf-if-eth1/1/5)# exit

MX9116N-A1(config)# interface ethernet1/1/6

```

```

MX9116N-A1 (conf-if-eth1/1/6)# NPAR partition-4
MX9116N-A1 (conf-if-eth1/1/6)# switchport mode trunk
MX9116N-A1 (conf-if-eth1/1/6)# spanning-tree disable
MX9116N-A1 (conf-if-eth1/1/6)# exit

# Assign VLANs to NPAR interfaces:

# (this example is only for switch-1)
MX9116N-A1 (config)# interface eth-NPAR1/1/1/1
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/1)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/1)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/1)# switchport access vlan 1000
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/1)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/1/2
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/2)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/2)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/2)# switchport access vlan 2000
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/2)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/1/3
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/3)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/3)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/3)# switchport access vlan 3000
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/3)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/1/4
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/4)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/4)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/4)# switchport access vlan 221
MX9116N-A1 (conf-if-eth-NPAR-1/1/1/4)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/5/1
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/1)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/1)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/1)# switchport access vlan 1000
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/1)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/5/2
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/2)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/2)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/2)# switchport access vlan 2000
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/2)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/5/3
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/3)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/3)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/3)# switchport access vlan 3000
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/3)# exit

MX9116N-A1 (config)# interface eth-NPAR1/1/5/4
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/4)# no shutdown
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/4)# switchport mode trunk
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/4)# switchport access vlan 221
MX9116N-A1 (conf-if-eth-NPAR-1/1/5/4)# exit

```

Read the report at <https://facts.pt/nVb3vx1>



This project was commissioned by Dell Technologies.



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners.

**DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:**

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.