

CXS1569BCN

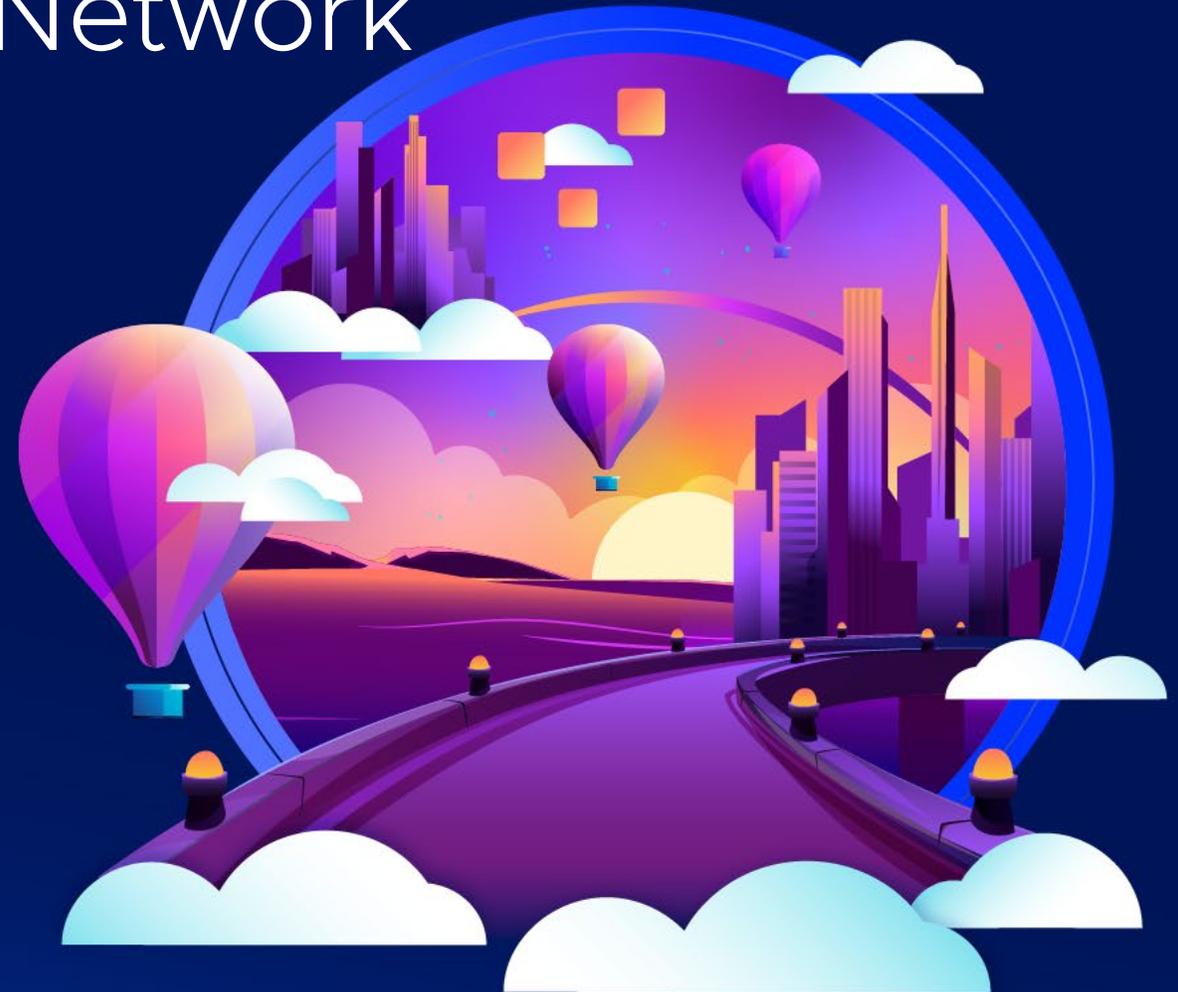
vmware® EXPLORE

# Accelerate Your Virtual Network Environment with NSX Powered by SmartNICs

Teddy Alerte

Staff Learning Solutions Architect

#vmwareexplore #CXS1569BCN



# Presenter



**Teddy Alerte**

Staff Learning Solutions Architect



Blog: [www.infrazxp.com](http://www.infrazxp.com)

vExpert: <https://vexpert.vmware.com/directory/4664>

# Agenda

Traditional Infrastructure Challenges

Explore the Future with SmartNICs

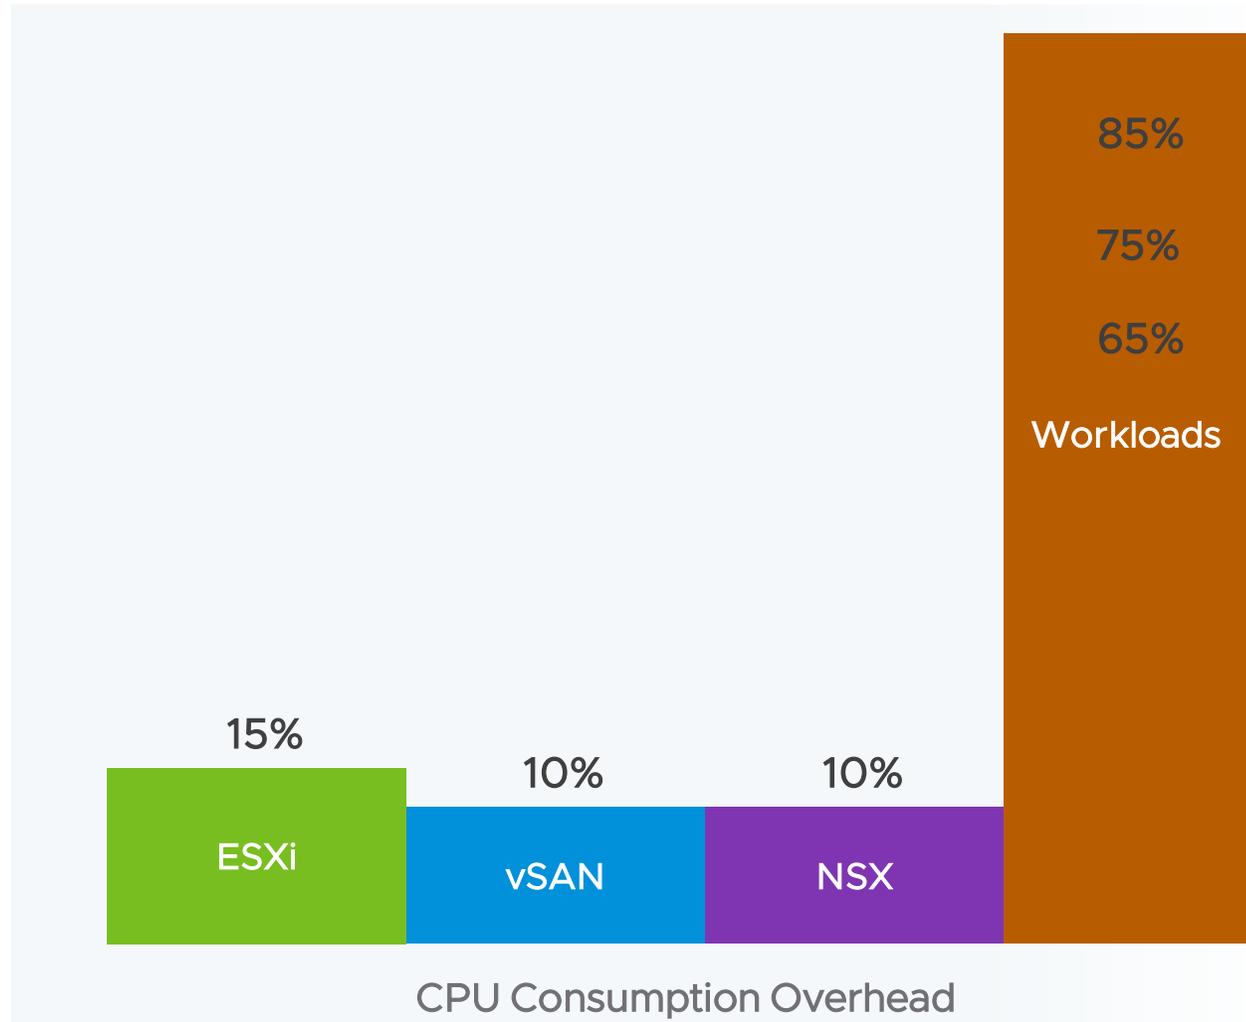
Use Cases and NSX Supported Features with SmartNICs

DPU Architecture in NSX

NSX Configuration with DPUs

# Traditional Infrastructure Challenges

# Traditional Infrastructure Challenges



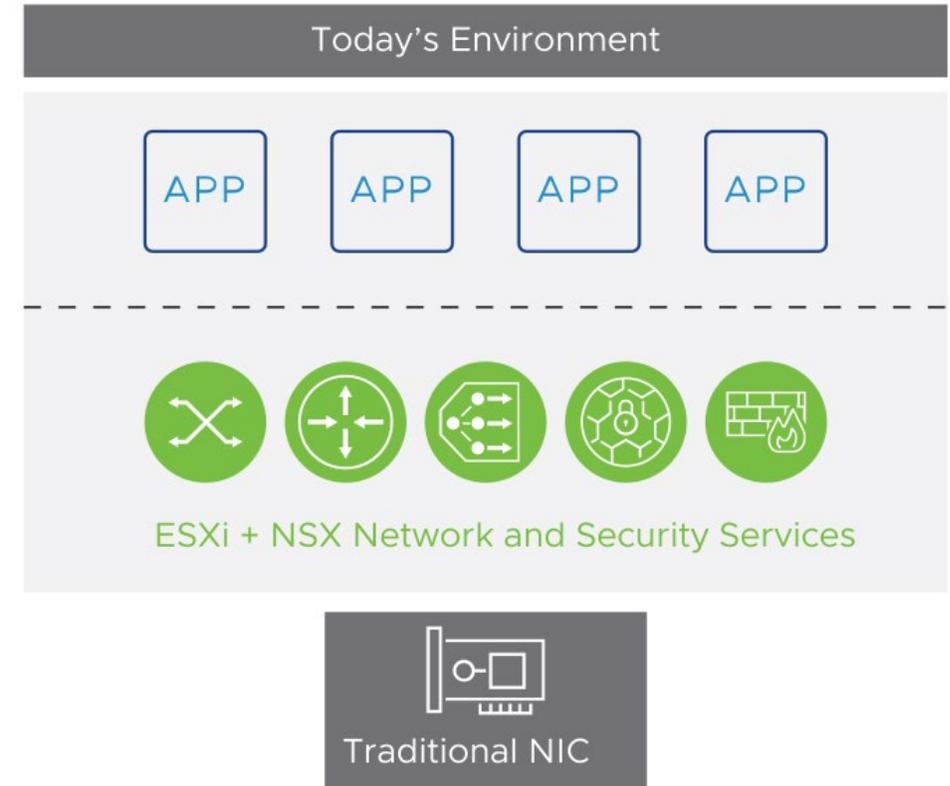
Note: Please note that these values are approximative.

# Traditional Infrastructure Challenges

## Today's challenges

Here are some of the challenges faced by many businesses in their environment today:

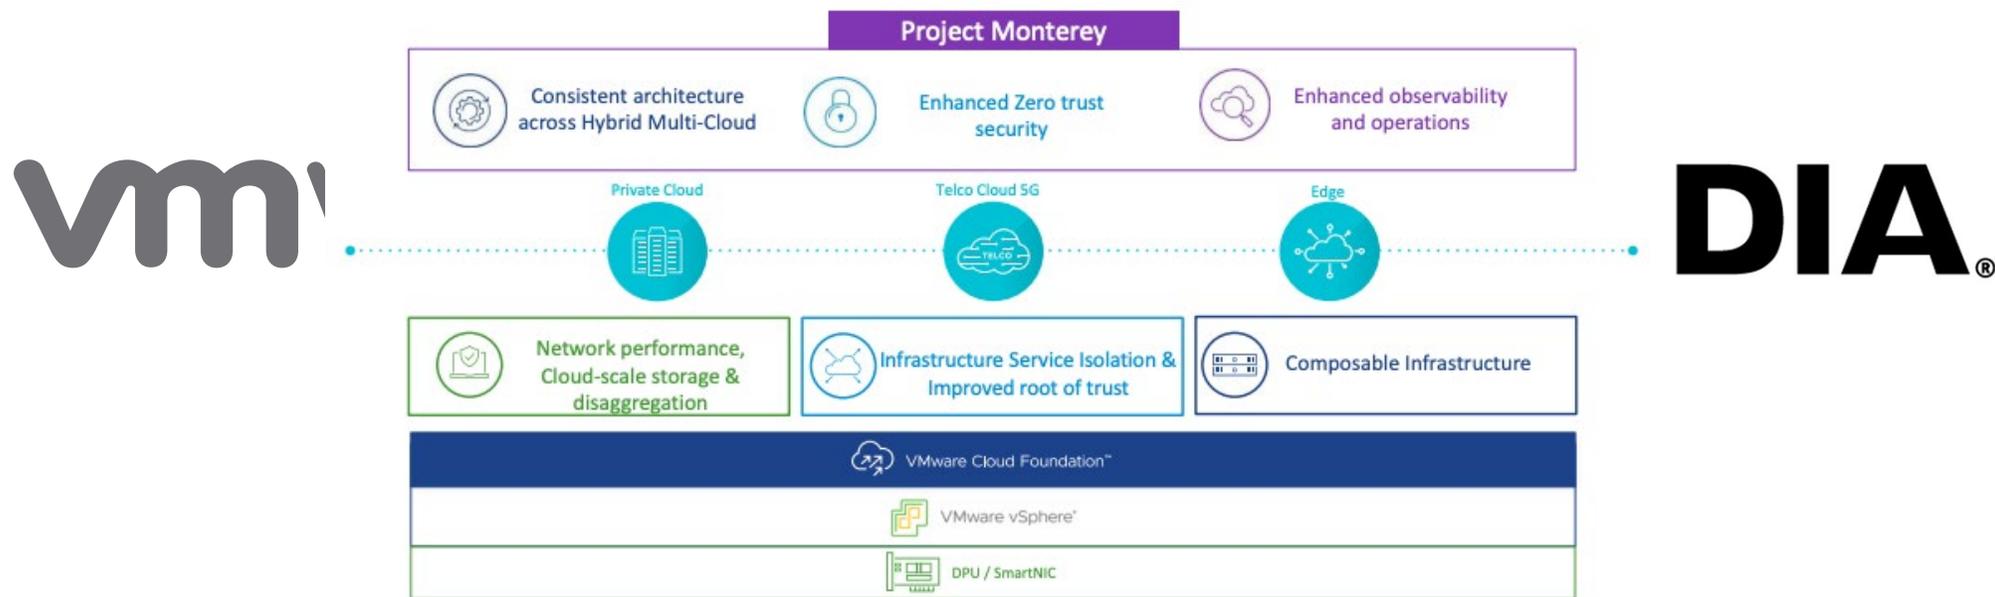
- New applications requirements
- Hypervisors sharing computing resources between application and infrastructure (compute, storage, network...)
- Relying on traditional NIC for networking and security services affects the performance of applications
- Inadequate isolation between provider and tenant in multitenant environments based on bare-metal clouds



# Traditional Infrastructure Challenges

At VMworld 2020, we have announced a new partnership with Nvidia and introduced Project Monterey.

Our goal is to help you face all these challenges, re-invent the virtual infrastructure, and accelerate and better secure your workloads.



# EXPLORE > THE FUTURE WITH SMARTNICs



# Explore the Future with SmartNICs

## What is a SmartNIC

A SmartNIC or DPU (Data Processor Unit) is a high-performance network adapter that incorporates advanced processing capabilities to offload computing resources from a bare-metal server.

This includes also programmable feature optimizing and accelerating network communication within data centers.



# Explore the Future with SmartNICs

Supported DPU Vendors with NSX



# Use Cases and NSX Supported Features with SmartNICs

# Use Cases and NSX Supported Features with SmartNICs

## Use Cases with SmartNICs or DPUs in NSX

DPU-based acceleration has the following use cases:

Applications with high network bandwidth demand and low latency



Security services offloading for a better performance



Enhanced observability requirements



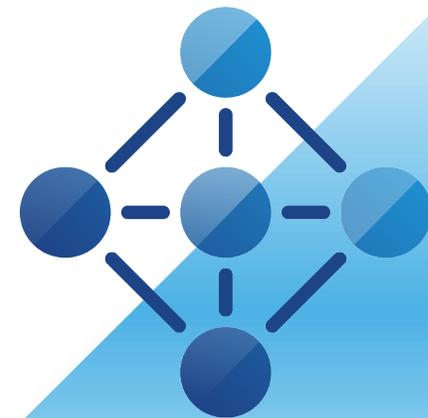
# Use Cases and NSX Supported Features with SmartNICs

## NSX Supported Features with SmartNICs

NSX 4.0.1 supports the following features with SmartNICs/DPUs:

- **Networking**
  - Overlay and VLAN based segments
  - Distributed IPv4 and IPv6 routing
  - NIC teaming across the SmartNIC / DPU ports
- **Security (Tech Preview)**
  - Distributed IDS/IPS
  - Distributed Firewall
- **Visibility and Operations**
  - Traceflow
  - IPFIX
  - Packet Capture
  - Port Mirroring
  - Statistics
- **Supported DPU Vendors**
  - NVIDIA Bluefield-2 (25Gb models only) – (UPT - Tech Preview)
  - AMD Pensando (25Gb and 100Gb models)

# DPU Architecture in NSX

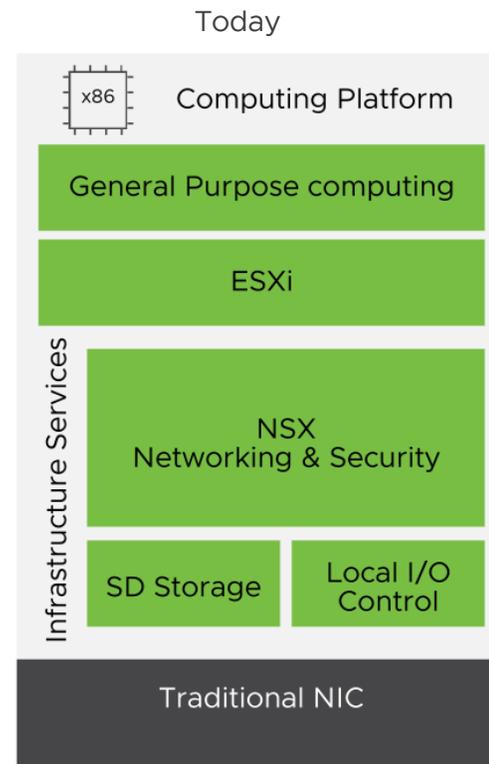


# DPU Architecture in NSX

## Architecture Changes with DPUs

DPU based architecture includes the following changes:

- ESXi, NSX, Storage and other and Infrastructure services through standard NICs using CPU cycle of the host.
- ESXi and NSX instances run directly in the DPU's processor.

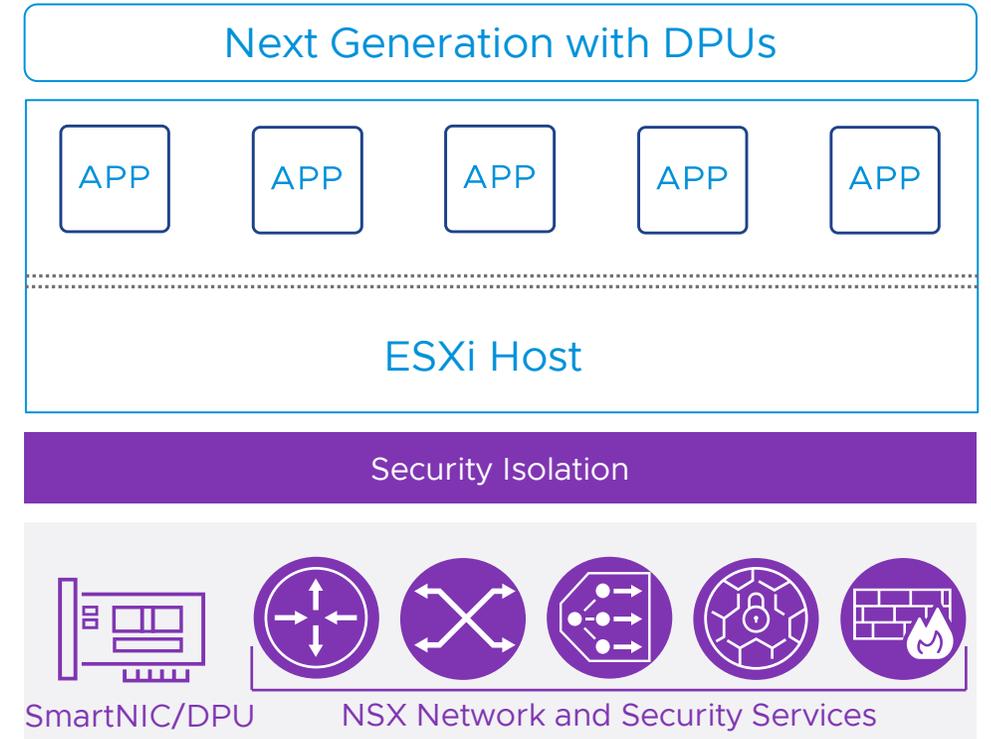


# DPU Architecture in NSX

## Next-Gen Infrastructure with SmartNICs

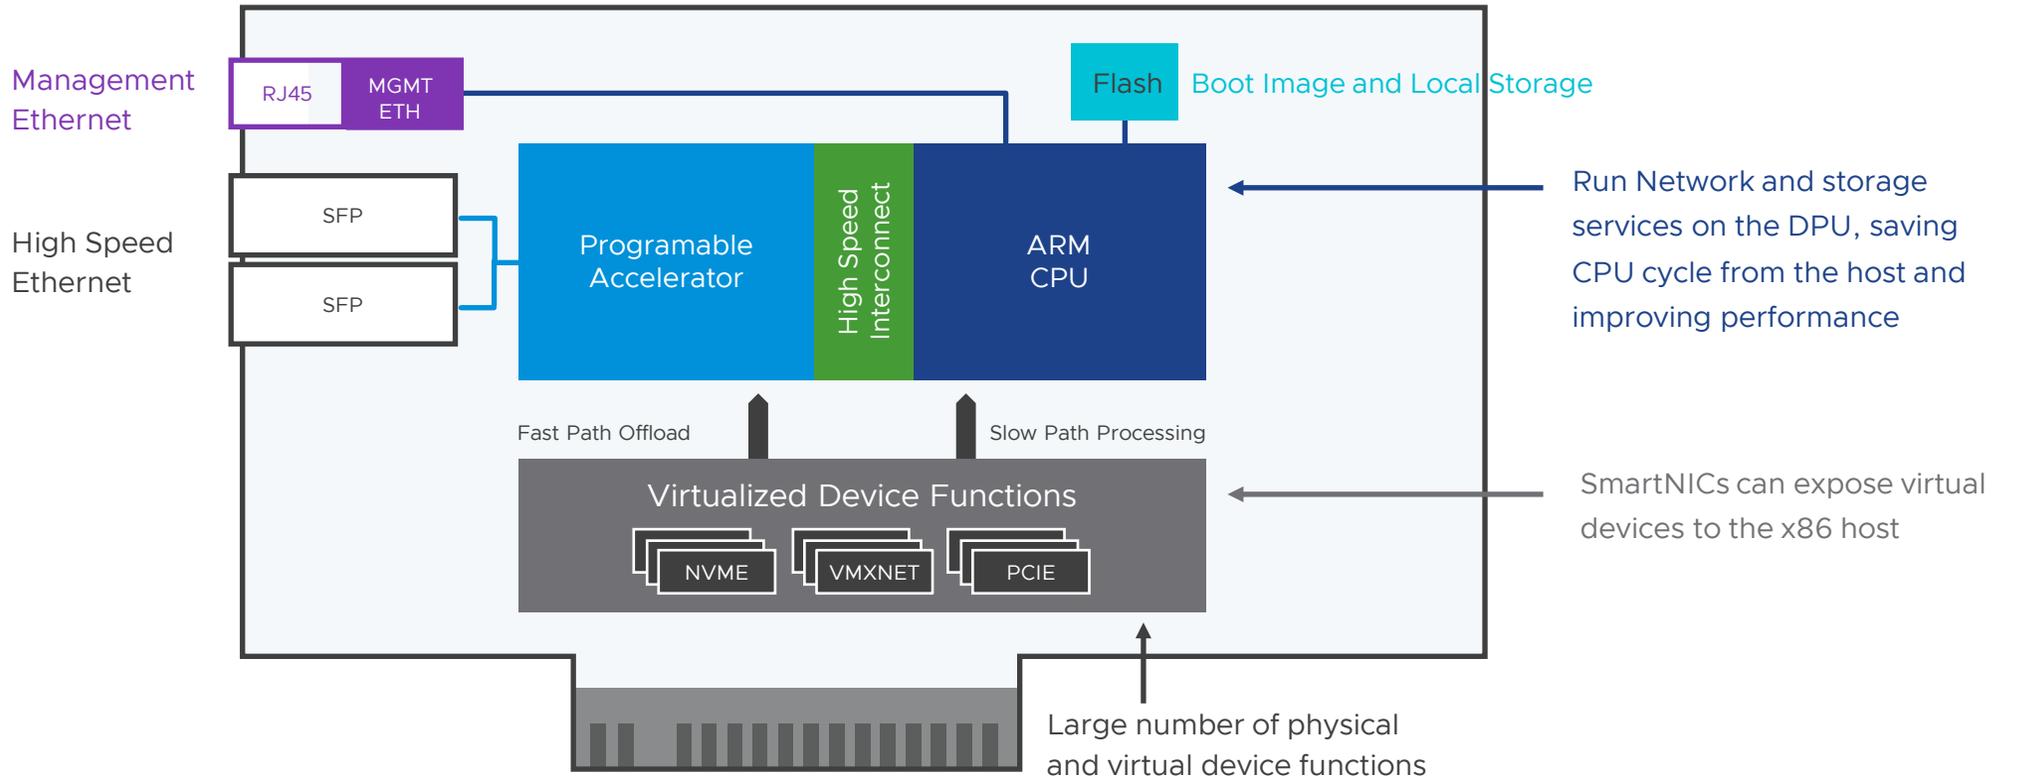
DPU offer the following advantages:

- Dedicated computing resources and hardware acceleration
- Full datapath offloading to achieve high throughput and low latency
- Security services enabled in the DPU without performance impact
- Enhanced observability and operations capabilities for monitoring, troubleshooting, logging, and compliance
- Isolation between tenant and provider both on ESXi and bare-metal platforms (available in future releases)



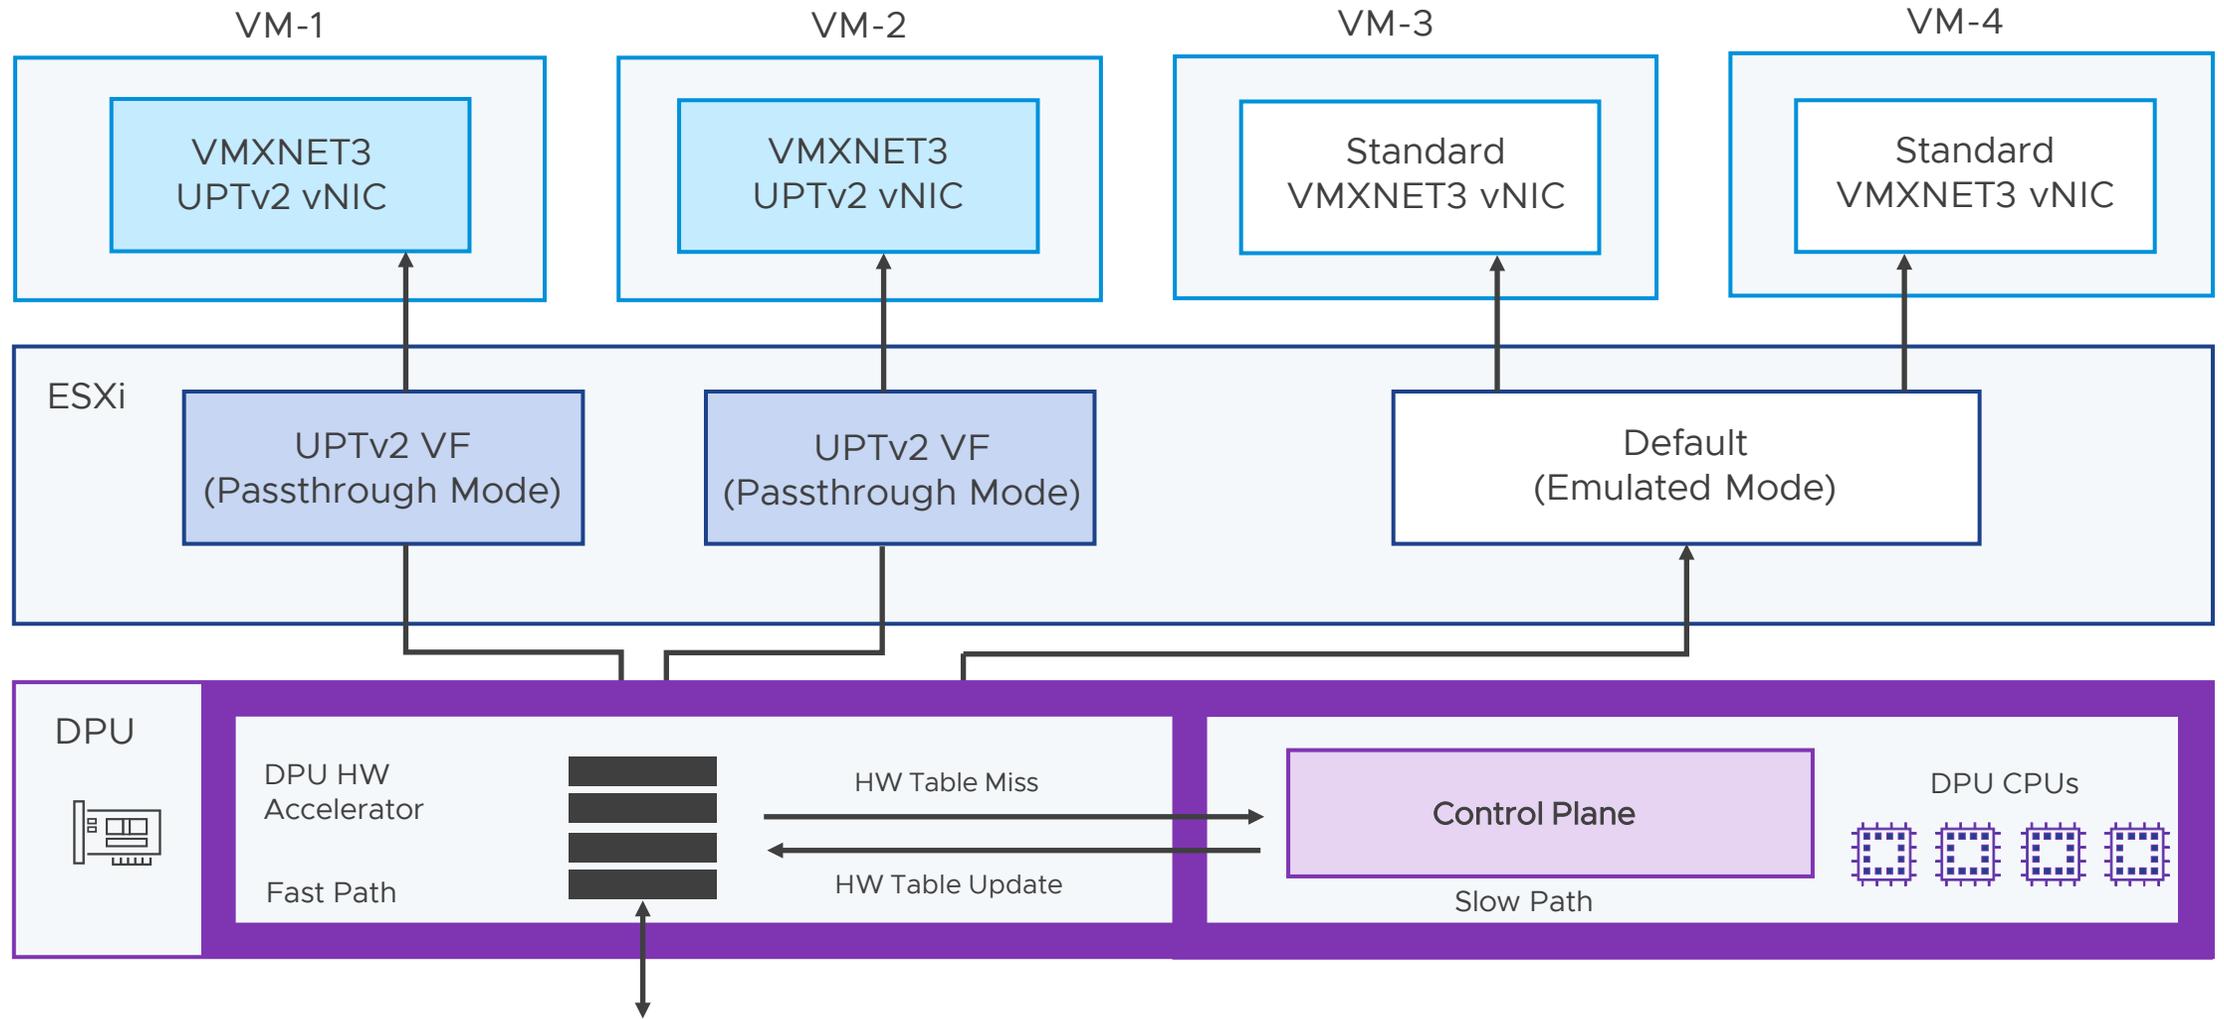
# DPU Architecture in NSX

## DPU Generic Hardware



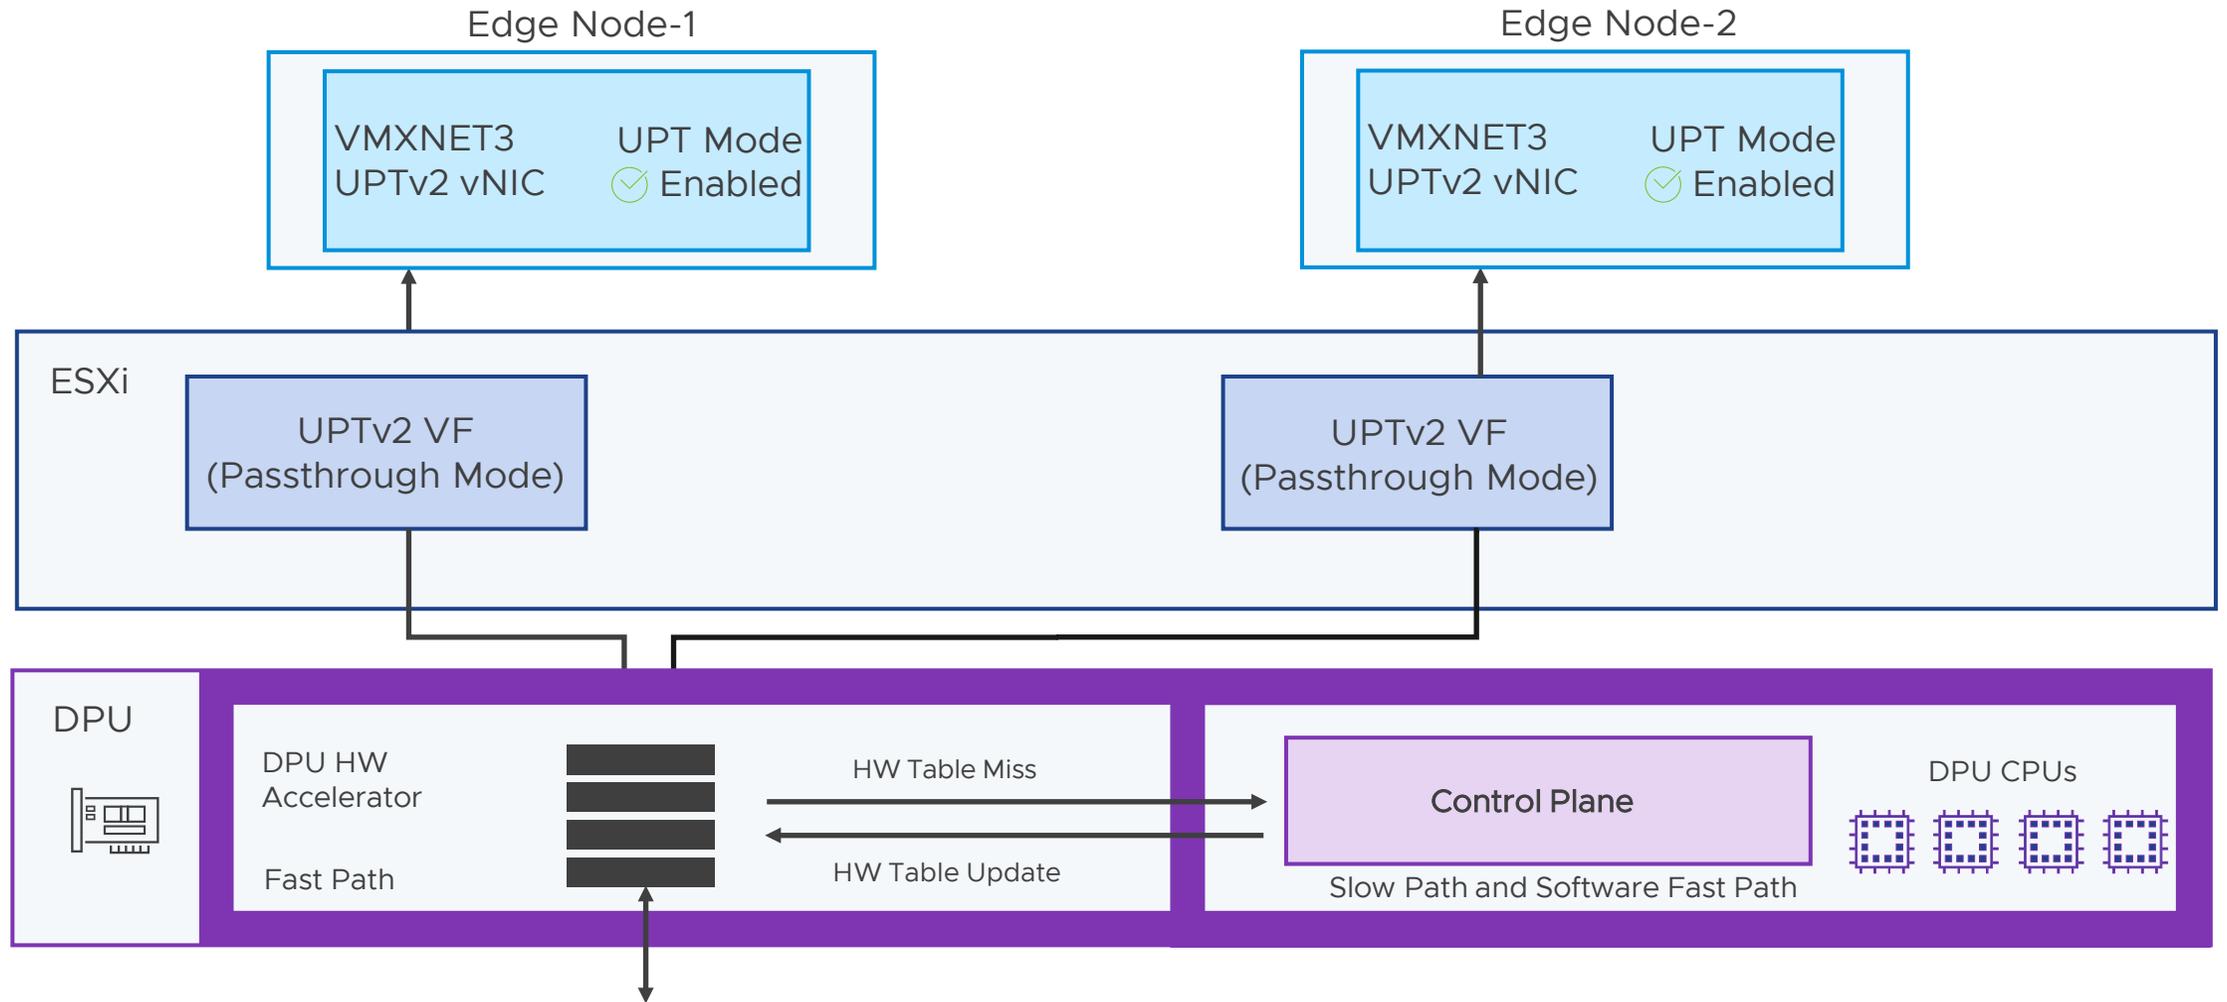
# DPU Architecture in NSX

## DPU Network Offloading with UPTv2



# DPU Architecture in NSX

## DPU Network Offloading for VM Edge Node



# DPU Architecture in NSX

## DPU Properties for NSX/vSphere

### UPTv2 (Passthrough)

- Near zero CPU Consumption
- Complete Guest Memory Reservation
- Requires VMware Tools for ESXi 8.x minimum
- High Performance, Low Latency Data Path

### Default

- Higher CPU footprint than UPTv2
- No Guest Memory Reservation
- No driver dependencies on GuestOS
- Performance accelerated by HW in DPU

# DPU Architecture in NSX

## Switching Modes Compatibility

VDS Operational Modes	Compatibility
Standard VDS	Not supported for DPU Offload
Enhanced Data Path – Standard	Supported for DPU Offload and Acceleration
Enhanced Data Path – Performance	Supported for DPU Offload and Acceleration

# DPU Architecture in NSX

## VM Modes Compatibility

VM Operational Modes	Compatibility
Default Mode	DPU Based Mode, vNIC passthrough Disabled, Some ESXi CPU usage networking stack is offloaded to the DPU Available with ESXi 8.x (virtual hardware 20 minimum) and NSX 4.x
UPTv2 Mode	DPU Based Mode, Enabled vNIC x86 passthrough, vSphere HA and vMotion features preserved Networking stack is offloaded to the DPU Available with ESXi 8.x (virtual hardware 20 minimum) and NSX 4.x

# DPU Architecture in NSX

## NSX 4.1.x Enhancement with DPUs

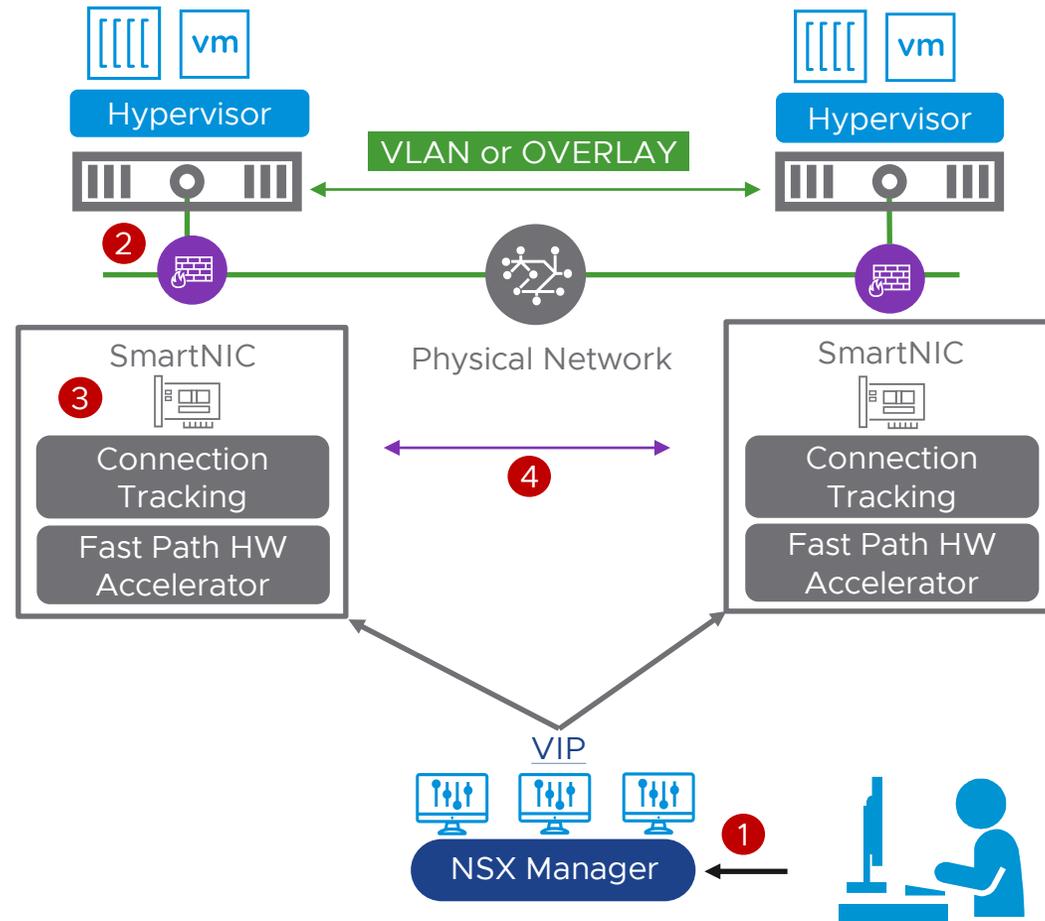
With the release of NSX 4.1.x, new enhancements are now available with DPUs:

- NVIDIA BlueField-2 100Gbps DPUs is supported.
- Support of NSX Distributed Firewall with DPU on production environments.

# DPU Architecture in NSX

## NSX Distributed Firewall on DPU

Starting NSX 4.1, NSX Distributed Firewall is supported with DPU on production environments.



# NSX Configuration with DPUs

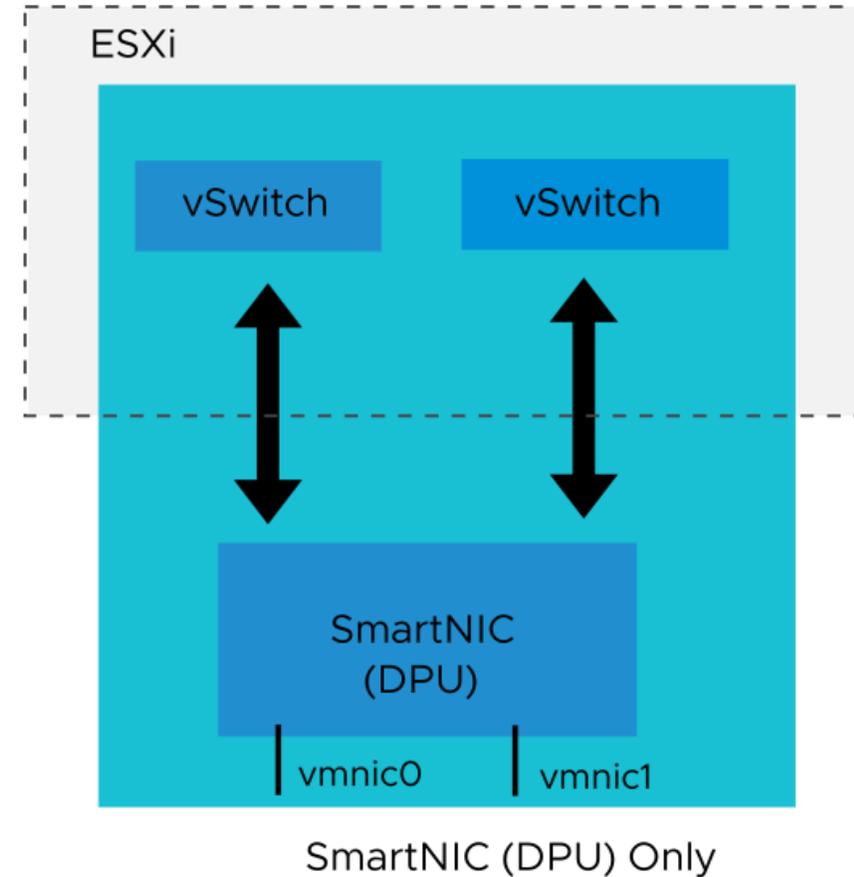


# NSX Configuration with DPUs

## Networking Configurations: SmartNIC Only

You can configure an ESXi with a SmartNIC with the following considerations:

- Only one SmartNIC is supported per host in NSX 4.1.x
- SmartNIC uplinks can only be attached to vSphere Distributed Switch (VDS) compatible with network offloads.
- A SmartNIC has two uplinks that can be associated with a maximum of two virtual switches.

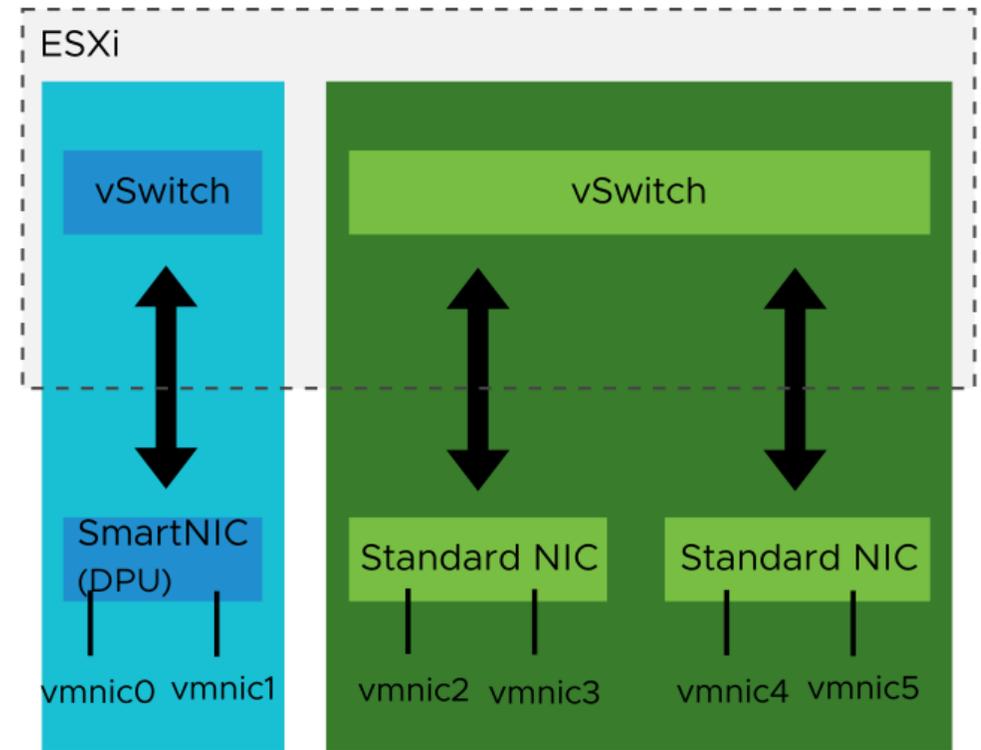


# NSX Configuration with DPUs

## Networking Configurations: SmartNIC with Standard NICs

You can configure ESXi with a SmartNIC and standard NICs with the following considerations:

- Uplinks from SmartNICs and standard NICs must not be part of the same VDS.
- Configure a dedicated Uplink profile for SmartNICs.



SmartNIC (DPU) with Standard NICs

# NSX Configuration with DPUs

## Installing NSX with DPUs (1)

Select **Physical adapters** by navigating to the host menu **Configure** > **Networking** to validate that the DPU adapters (SmartNIC) are detected.

The screenshot shows the VMware vSphere interface for host `w2-hs-dmz-f0610.isvlab.vmware.com`. The **Configure** tab is active, and the **Physical adapters** section is selected in the left sidebar. The main panel displays a table of network adapters with columns for Device, Actual Speed, Configured Speed, Switch, MAC Address, and DPU Backed. The 'DPU Backed' column shows 'Yes' for vmnic2 and vmnic3, and '--' for vmnic0, vmnic1, and vusb0.

Device	Actual Speed	Configured Speed	Switch	MAC Address	DPU Backed
vmnic0	25 Gbit/s	Auto negotiate	wdc-vds01	e4:3d:1a:13:3c:30	--
vmnic1	25 Gbit/s	Auto negotiate	wdc-vds01	e4:3d:1a:13:3c:31	--
vmnic2	25 Gbit/s	Auto negotiate	nsx-offload-dswitch	08:c0:eb:5d:ae:98	Yes
vmnic3	25 Gbit/s	Auto negotiate	nsx-offload-dswitch	08:c0:eb:5d:ae:99	Yes
vusb0	100 Mbit/s	100 Mbit/s	vSwitchBMC	b0:7b:25:d3:77:89	--

# NSX Configuration with DPUs

## Installing NSX with DPUs (2)

Create a vSphere 8.0 VDS version with the network offloads compatibility enabled.

### New Distributed Switch

- 1 Name and location
- 2 Select version
- 3 Configure settings**
- 4 Ready to complete

### Configure settings

Specify network offloads compatibility, number of uplink ports, resource allocation and default port group.

Network Offloads compatibility	None <input type="button" value="v"/> <input type="button" value="i"/>
Number of uplinks	4 <input type="button" value="+"/>
Network I/O Control	Enabled <input type="button" value="v"/>
Default port group	<input checked="" type="checkbox"/> Create a default port group
Port group name	DPortGroup

# NSX Configuration with DPUs

## Installing NSX with DPUs (3)

Manage DPU with a DPU installed on the VDS.

The screenshot displays the 'nsx-offload-dswitch - Add and Manage Hosts' configuration window. The left sidebar shows a progress list with '3 Manage physical adapters' selected. The main panel is titled 'Manage physical adapters' and includes instructions and a table of network adapters.

**Manage physical adapters**

Add or remove physical network adapters to this distributed switch.  
Adapters compatibility on this switch is set to NVIDIA BlueField.

Adapters on all hosts | Adapters per host

To associate a physical network adapter with an uplink, use "Assign uplink". This assignment would be applied to all the hosts that have the same physical network adapter available.

Physical network adapters	Compatible hosts	In use by switch	Assign uplink
>> vmnic0	None	1 host / 1 switch	None
>> vmnic1	None	1 host / 1 switch	None
>> vmnic2	1 host (All)	This switch	Uplink 1
>> vmnic3	1 host (All)	This switch	Uplink 2
>> vusb0	None	1 host / 1 switch	None

0 compatible physical network adapters

CANCEL BACK NEXT

1 hosts

CANCEL BACK NEXT

# NSX Configuration with DPUs

## Installing NSX with DPUs (4)

Select the cluster and click **Install NSX** to install NSX in a cluster with the DPU adapters.

The screenshot displays the VMware NSX System Overview page. The left sidebar contains navigation options: System Overview, Lifecycle Management, Quick Start (selected), Backup & Restore, Upgrade, Migrate, Configuration, Appliances, Hosts, Edges, Container Clusters, Switches, Transport Zones, and Compute Managers. The main content area is titled 'Prepare Clusters for Networking & Security' and features a 'BACK TO QUICK START' link and an 'INSTALL NSX' button. Below this is a table with columns for Cluster, Compute Managers, and Hosts. Cluster-01 is selected with a checkmark.

Cluster	Compute Managers	Hosts
<input checked="" type="checkbox"/> > Cluster-01	vCenter-01 (v8.0) 10.10.10.10	8
<input type="checkbox"/> > Cluster-02	vCenter-02 (v6.5) 10.10.10.11	8
<input type="checkbox"/> > Cluster-03	vCenter-03 (v6.5) 10.10.10.12	8
<input type="checkbox"/> > Cluster-04	vCenter-04 (v7.0) 10.10.10.13	8
<input type="checkbox"/> > Cluster-05	vCenter-05 (v7.0) 10.10.10.14	4

# NSX Configuration with DPUs

## Installing NSX with DPUs (5)

Verify that the DPU uplinks exist in the cluster by reviewing the NSX installation details.

Installation Details

NSX will be installed for the selected cluster with the following configuration.

Cluster01

8 Hosts  
ESXi 8.0

IP Assignment: DHCP

Customize Host Switch

System recommendations are prepopulated, you can modify anytime after installation.

VDS Switch: Select VDS switch

Transport Zone: TZ-01, TZ-VLAN

Uplink Profile: nsx-default-uplink-hostswitch-profile

Switch Teaming Policy Mapping

Uplinks	VDS Uplinks
Uplink-1 (active)	uplink1 (DPU)
Uplink-2 (standby)	uplink2 (DPU)

CANCEL INSTALL

# NSX Configuration with DPUs

## Installing NSX with DPUs (6)

Select **Host Transport Nodes** from the **Configuration** menu and verify that the host transport nodes are DPU supported.

The screenshot displays the NSX configuration interface. On the left, the 'Configuration' menu is open, with 'Nodes' selected. The main area shows 'Host Transport Nodes' managed by 'mvc04.isvlab.vmware.com'. A cluster named 'vLCM' is selected, containing three nodes: 'w2-hs-dmz-f0610.isvlab.vmware.com', 'w2-hs-dmz-f0611.isvlab.vmware.com', and 'w2-hs-dmz-f0612.isvlab.vmware.com'. The 'w2-hs-dmz-f0610.isvlab.vmware.com' node is selected, and the 'Physical Adapters' tab is active. A table lists the physical adapters for this node, with the 'DPU Backed' column highlighted in blue. The 'vmnic2' and 'vmnic3' adapters are marked as 'Yes' in the 'DPU Backed' column.

Interface Id	Admin Status	Link Status	MTU	Interface Details	DPU Backed
vmk0	Up	Up	1500	1	No
vmk1	Up	Up	1500	1	No
vmk2	Up	Up	9000	1	No
vmk3	Up	Up	9000	1	No
vmnic0	Up	Up	9000	1	No
vmnic1	Up	Up	9000	1	No
vmnic2	Up	Up	1600	1	Yes
vmnic3	Up	Up	1600	1	Yes
vusb0	Up	Up	1500	1	No

# NSX Configuration with DPUs

## Enabling UPT mode on Edge Nodes

You enable **Enable Uniform Passthrough** (UPT) mode on an NSX Edge during Edge node setting configuration.

The screenshot shows the 'Add Edge Node' configuration wizard in NSX, specifically the 'Configure Node Settings' step. The left sidebar lists the steps: 1 Name and Description, 2 Credentials, 3 Configure Deployment, 4 Configure Node Settings (highlighted), and 5 Configure NSX. The main configuration area includes the following settings:

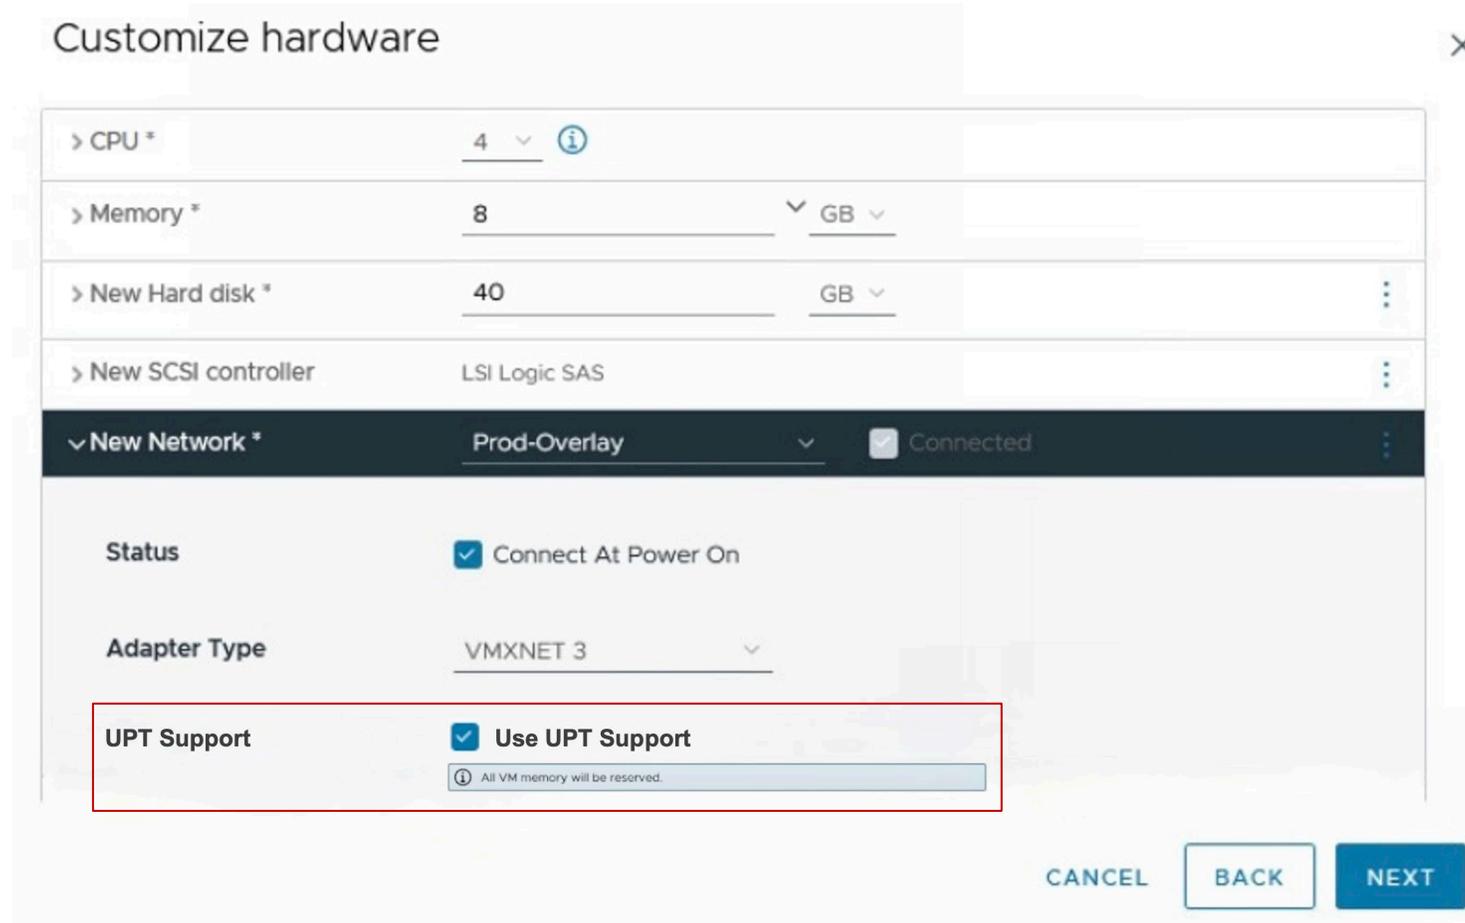
- Management IP Assignment:** Static (selected), DHCP (unselected). Management IP: 172.20.10.65/24. Default Gateway: 172.20.10.10.
- Management Interface:** pg-SA-Management (Distributed VI...)
- Search Domain Names:** vclass.local
- DNS Servers:** 172.20.10.10
- NTP Servers:** 172.20.10.100
- Enable UPT mode for datapath interface:** Yes (toggle is turned on).

At the bottom, there are buttons for CANCEL, PREVIOUS, and NEXT. A blue callout box points to the 'Enable UPT mode for datapath interface' toggle with the text 'Enable Uniform Passthrough Mode'.

# NSX Configuration with DPUs

## Virtual Machine Configuration

To configure a VM for DPU acceleration and offloading, select the **UTP Support** checkbox.



The screenshot shows the 'Customize hardware' dialog box in VMware vSphere. The dialog is titled 'Customize hardware' and has a close button (X) in the top right corner. It contains several configuration sections:

- CPU \***: 4 (with an info icon)
- Memory \***: 8 GB
- New Hard disk \***: 40 GB
- New SCSI controller**: LSI Logic SAS
- New Network \***: Prod-Overlay, Connected (checked)
- Status**:  Connect At Power On
- Adapter Type**: VMXNET 3
- UTP Support**:  Use UTP Support

A red box highlights the 'UTP Support' section. Below the 'Use UTP Support' checkbox, there is a small information icon and the text: 'All VM memory will be reserved.'

At the bottom right of the dialog, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

# NSX Configuration with DPUs

## DPUs Basic Troubleshooting

To verify that the traffic is offloaded to the DPU, use the `esxcfg-nics -l` command to review the vNICs associated to DPUs, access the DPU shell from the host, and then dump the flows table using the `nsxdp-cli ens flow-table dump` command.

```
[root@w2-hs-dmz-f0609:~] esxcfg-nics -l
[root@w2-hs-dmz-f0609:~] sshdpu
^
^
^The time and date of this login have been sent to the system logs.
^
^WARNING:
^  All commands run on the ESXi shell are logged and may be included in
^  support bundles. Do not provide passwords directly on the command line.
^  Most tools can prompt for secrets or accept them from standard input.
^
^VMware offers supported, powerful system administration tools. Please
^see www.vmware.com/go/sysadmintools for details.
^
^The ESXi Shell can be disabled by an administrative user. See the
^vSphere Security documentation for more information.
[root@localhost:~] nsxdp-cli ens flow-table dump
FT dstMAC          srcMAC          VLAN  srcPort  srcIP          dstIP          proto  VNI      srcPort/ty
pe dstPort/code  Actions hwHits          hwBytes
-----
L4 00:50:56:b6:1f:9c 00:50:56:b6:9a:50 0      2        192.168.10.11  192.168.10.10  1      69632   0
0      bmap:0x1400080 inval(s):103 cg:110 dp:0x6 len:704; GENEVE DECAP; DFW on dstPort; 0 0
L4 00:50:56:b6:9a:50 00:50:56:b6:1f:9c 0      6        192.168.10.10  192.168.10.11  1      0       8
0      bmap:0x20000c0 inval(s):136 cg:110 dp:0x2 len:814; DFW on srcPort; VNI: 69632; GENEVE ENCAP VNI: 69632; 0 0
```

# Want to Learn More?

VMware Learning courses relating to this topic

VMware NSX: Install, Configure, Manage [V4.0] - ILT/On Demand  
(5 DAYS)

VMware NSX: What's New [V4.0] - On Demand  
(3 DAYS)

VMware vSphere: What's New [V8] - ILT/Digital  
(2 DAYS)

[vmware.com/learning](https://www.vmware.com/learning)

Thank You

