Solution operation and life cycle management

Initial setup and firmware updates for the DSE solution

Lenovo

vSphere on DPU

The concept of the Distributed Services Engine is vSphere on DPU. The solution allows VMware to offload and accelerate infrastructure service functions on to DPUs. For example, storage services, network services, and infrastructure management can all be offloaded from the CPU to the DPU, allowing more CPU resources to be used for your workload.

- Offload and accelerate infrastructure service functions on DPUs
- Simplify lifecycle management of DPUs
- Isolation of workload domain from the infrastructure domain

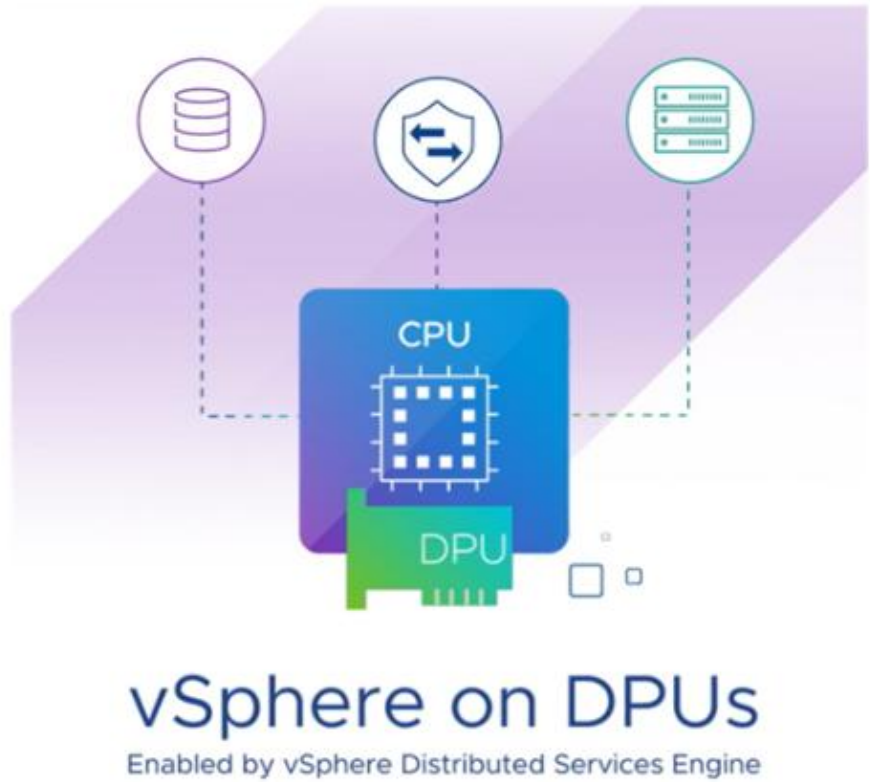

DPU-accelerated servers

The following figures show a comparison between traditional servers and DPU-accelerated servers.

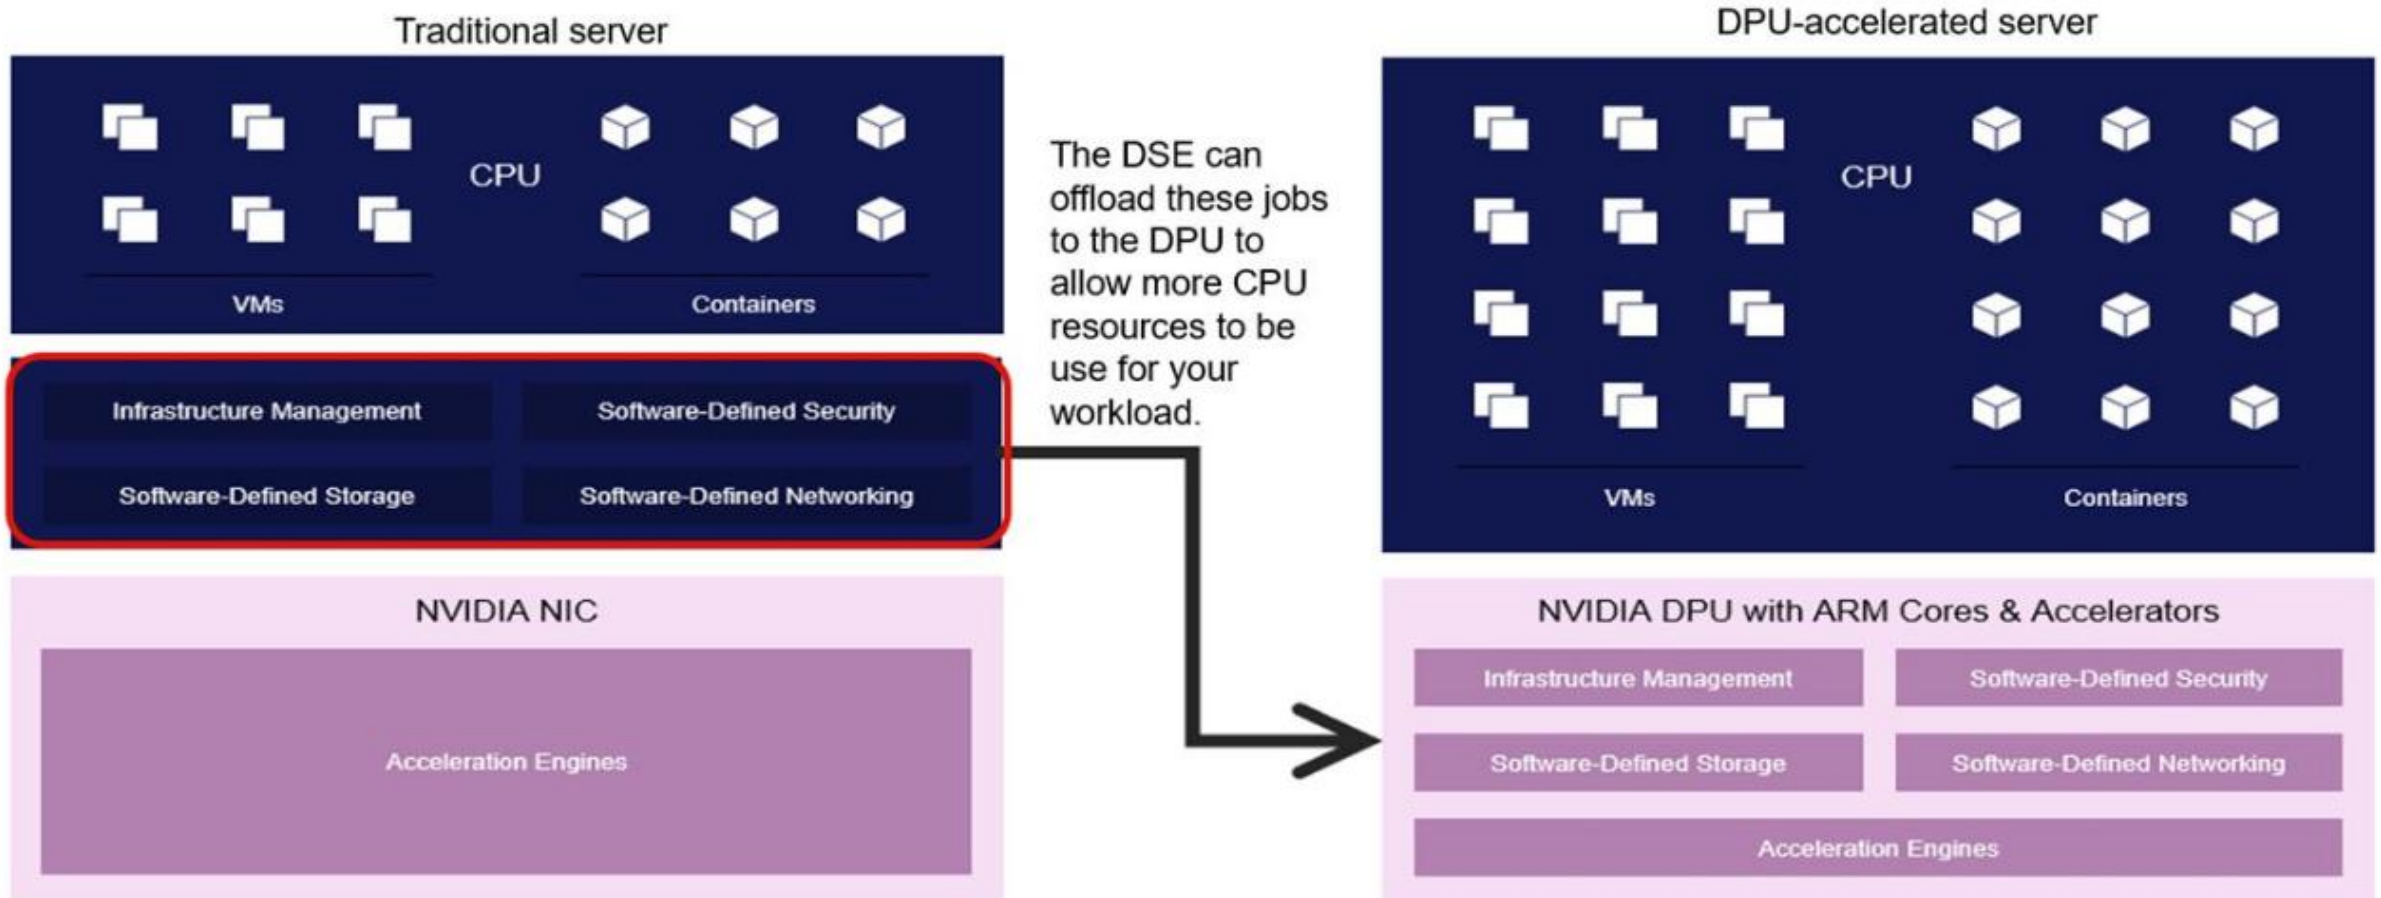

ThinkAgile VX DPU models

The ThinkAgile VX CTO models with the Distributed Services Engine have the following features:

- DPU settings applied at NVIDIA
- ThinkAglie VX node UEFI settings applied in the factory (also listed in the Best Recipe)
- ESXi 8.0 U1 preloaded in the factory for Certified Nodes (CN) and Integrated Systems (IS)
- The OS will be installed on both the M.2/7 mm drive and the DPU
- UEFI, XCC, and DPU firmware (NIC and ARM) and driver versions will be maintained in the existing Best Recipes Cycle releases
- All code can be updated via vLCM using the ThinkAgile VX Repository Pack – this will be included in LXCI June 2023 release.

UEFI settings for the DSE solution

To work with the DSE, apply the following UEFI settings:

- Set **System Boot Mode** to **UEFI Mode**
  - **F1➔ Boot Manager➔ Boot Mode**
- Enable **PCIe Error Recovery**
  - **F1➔ System Settings➔ Recovery and RAS➔ Advanced RAS**
- Enable **SRIOV** (needed for VMware NSX offloading)
  - **F1➔ System Settings➔ Device and IO Ports**
  - This item is only available when the correct adapters are installed in the system

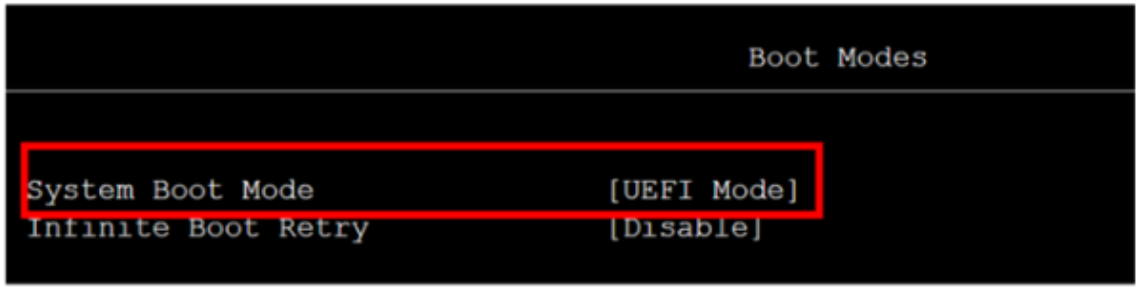

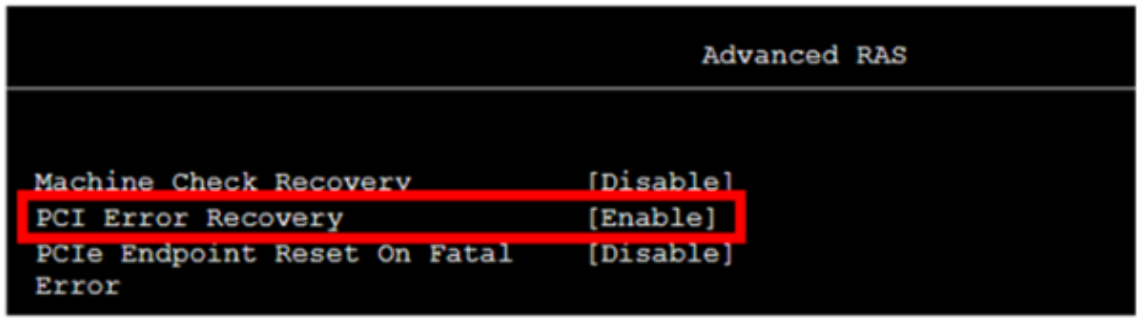

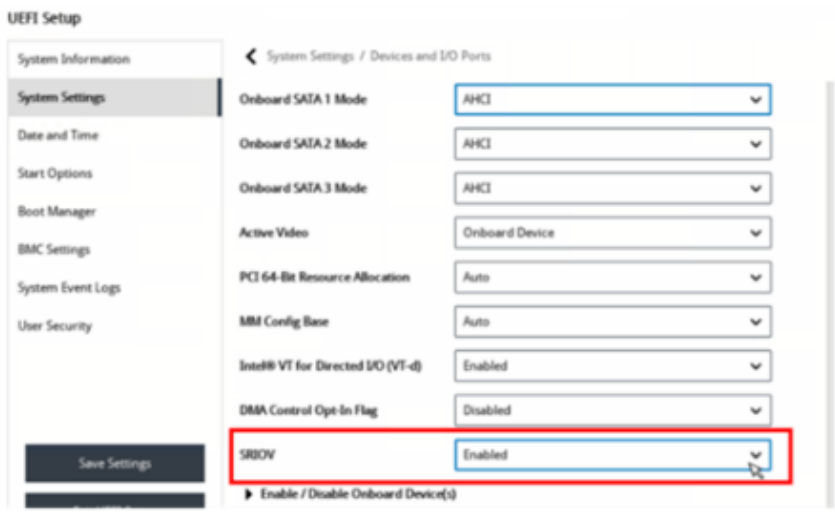

OCP Enablement Kit for vSphere DSE identification in XCC

If equipped, the OCP Enablement Kit for vSphere DSE will be identified as the **Rear_OCP_Card** in XCC. Users can also check the part number or FRU number for further confirmation.

Other Hardware

Name	Manufacture ID	Serial Number	Manufacture Date	Part Number	FRU Number
IO_Board	LENOVO	L3HF29J00M1	2022-10-24	STA7A90421	03KM422
ROT_Board	LENOVO	L2HF2AD00SA	2022-10-17	STA7B28783	03KM033
Rear_OCP_Card	Lenovo	L1HF2C9000V	2023-01-04	SN37B41311	03GX844
PCIe_Riser_1	LENOVO	L1HF27D0019	2022-07-25	STA7A90440	03KL993
PCIe_Riser_2	LENOVO	L1HF27D0002	2022-09-19	STA7A90440	03KL993
SFF_Retimer_1	LENOVO	V1HZ2AP005Y	2022-11-07	SC57A26292	02JG693

BlueField-2 identification in XCC

The BlueField-2 DPU adapter will be identified as shown below. Users can also check the FRU number for further confirmation.

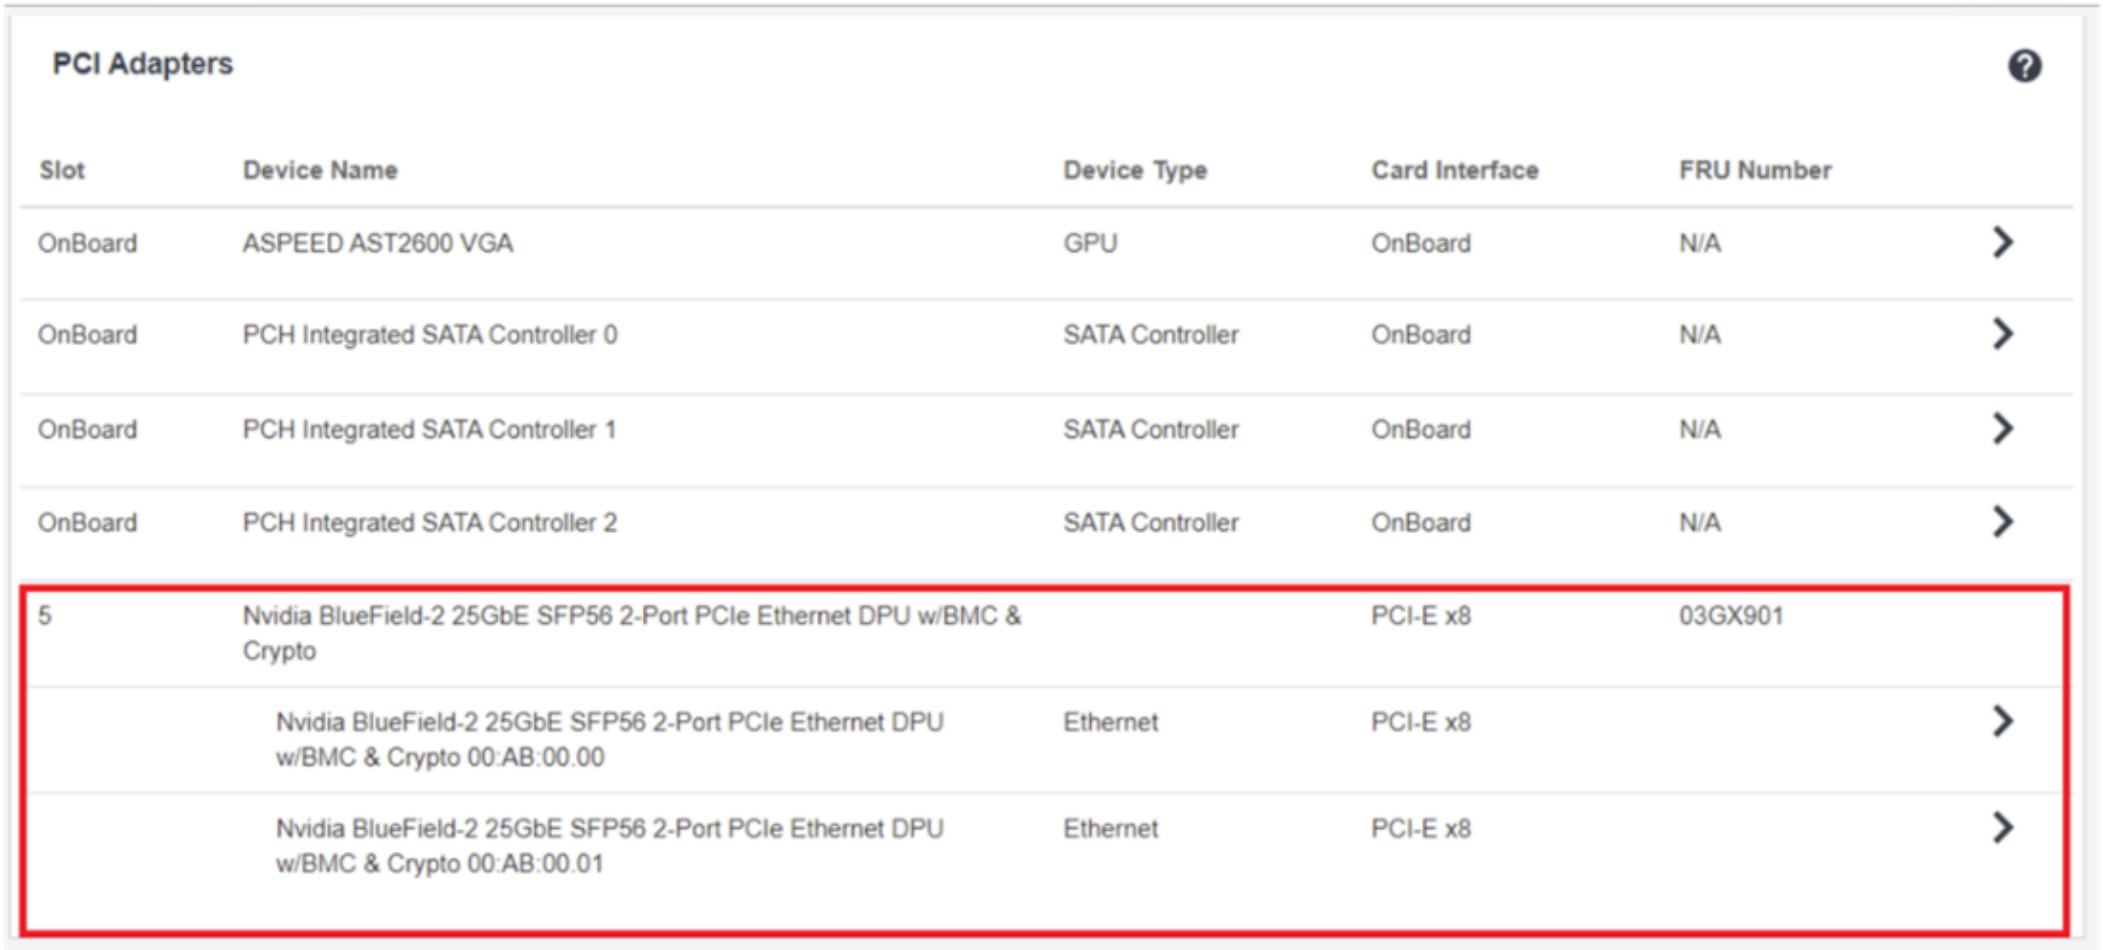

PCI Adapters

Slot	Device Name	Device Type	Card Interface	FRU Number
OnBoard	ASPEED AST2600 VGA	GPU	OnBoard	N/A
OnBoard	PCH Integrated SATA Controller 0	SATA Controller	OnBoard	N/A
OnBoard	PCH Integrated SATA Controller 1	SATA Controller	OnBoard	N/A
OnBoard	PCH Integrated SATA Controller 2	SATA Controller	OnBoard	N/A
5	Nvidia BlueField-2 25GbE SFP56 2-Port PCIe Ethernet DPU w/BMC & Crypto		PCI-E x8	03GX901
	Nvidia BlueField-2 25GbE SFP56 2-Port PCIe Ethernet DPU w/BMC & Crypto 00:AB:00.00	Ethernet	PCI-E x8	
	Nvidia BlueField-2 25GbE SFP56 2-Port PCIe Ethernet DPU w/BMC & Crypto 00:AB:00.01	Ethernet	PCI-E x8	

BlueField-2 identification in the ESXi installation

The BlueField-2 adapter will be identified as a Mellanox adapter in the ESXi installation (the adapter name will be amended in the future release). Customers will not generally see this screen, as the OS will be installed in the factory before the system is shipped.

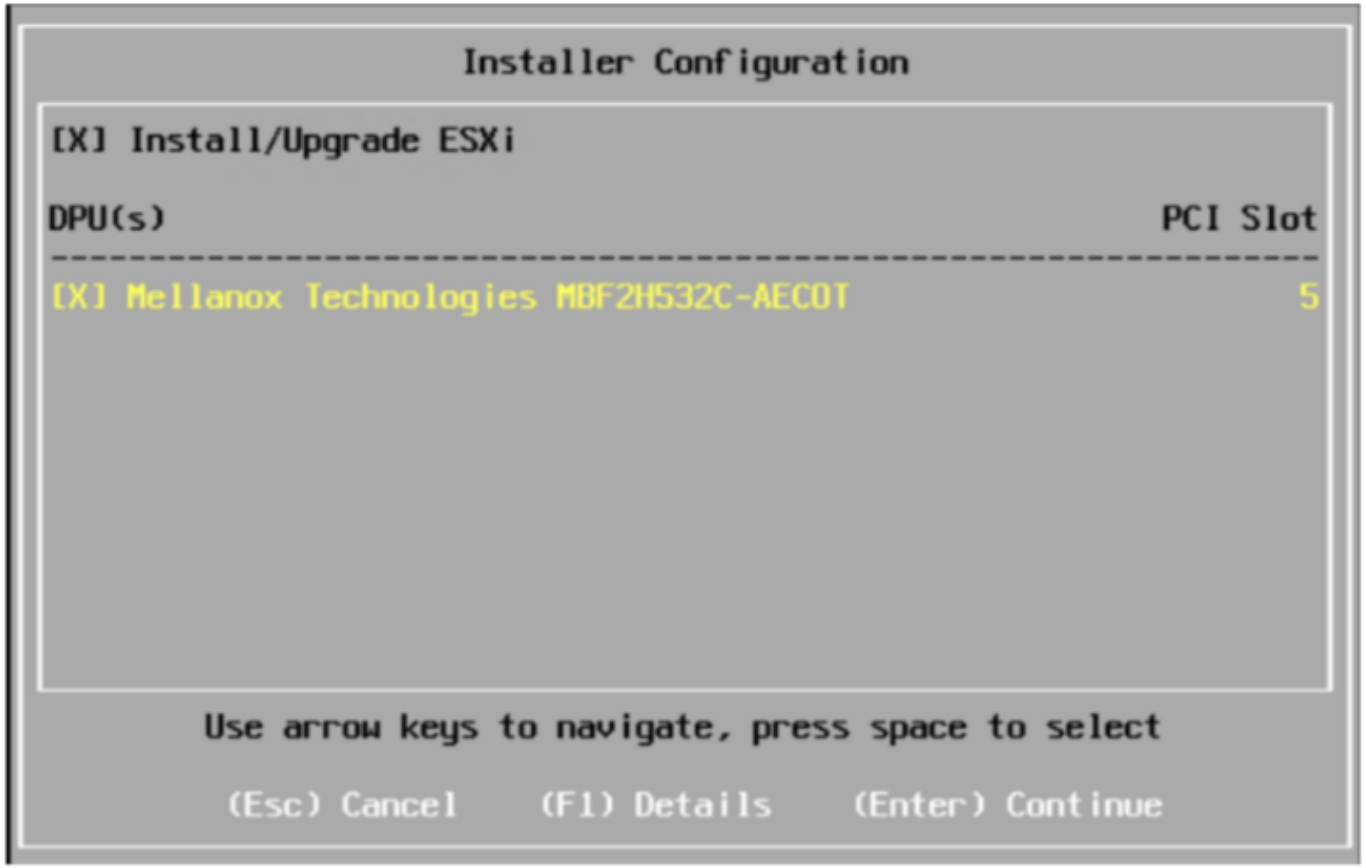

BlueField-2 firmware identification

The BlueField-2 firmware version can be identified in the XCC **Firmware Update** section.

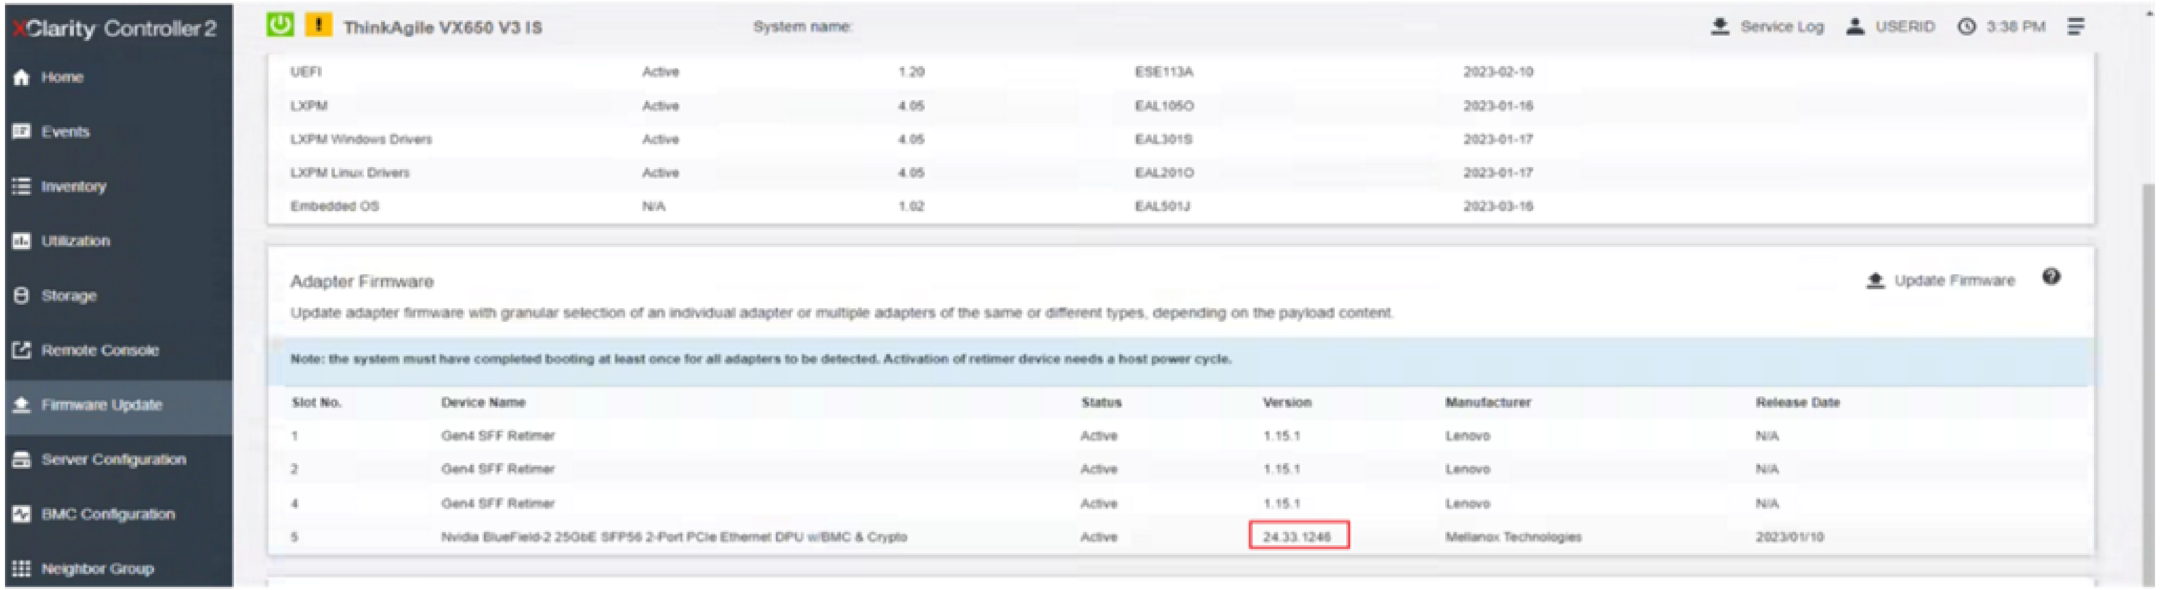

The OneCli `scan` command can also be used to check BlueField-2 firmware information. Click here to learn more about the command.

```
neCli.exe update scan --bmc <user>:<pw>@<ip>
==========================================================================================
No. |                    Updatable Unit                                   |   Slot    | Installed Version |
------------------------------------------------------------------------------------------
------------------------------------------------------------------------------------------
10  | Nvidia BlueField-2 25GbE SFP56 2-Port PCIe Ethernet DPU w/BMC & Crypto(BlueField UEFI) | PCI_Slot 2 | 18.2.0.12580      |
```

BlueField-2 settings in the ARM UEFI

This screenshot from the ARM UEFI shows the BlueField-2 settings applied by NVIDIA. This information is just for reference, and servicers or customers will not need to do anything with it.

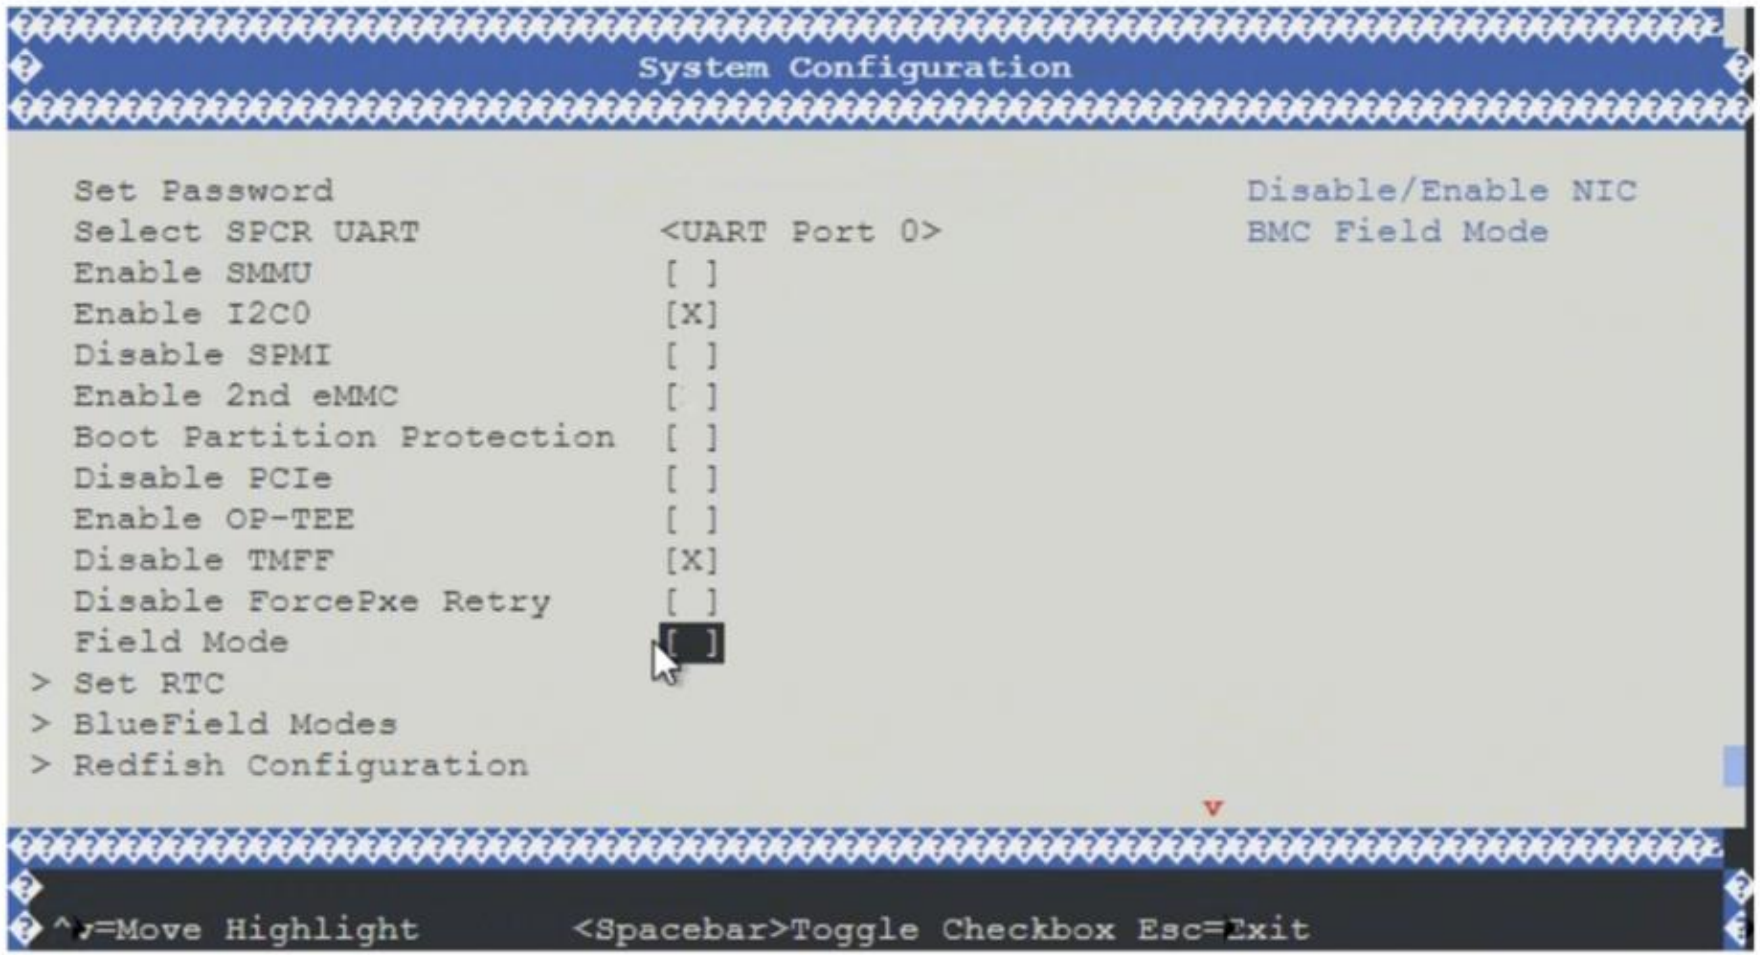

ThinkAgile VX DSE solution firmware updates with vLCM

To update the firmware of the ThinkAgile VX with DSE solution:

- Check ThinkAgile VX Best Recipe for the latest supported and certified code stack
- Only certified UEFI, XCC, and BlueField-2 code can be used on DPU models – for more information, refer to the Checking compatibility with the VMware Compatibility Guide page in this course
- All hosts in the cluster must have homogenous hardware configurations (the same DPU)
- Only one DPU per host

Note: For more information abut vLCM, refer to the following document: Working with the vSphere Lifecycle Manager function.

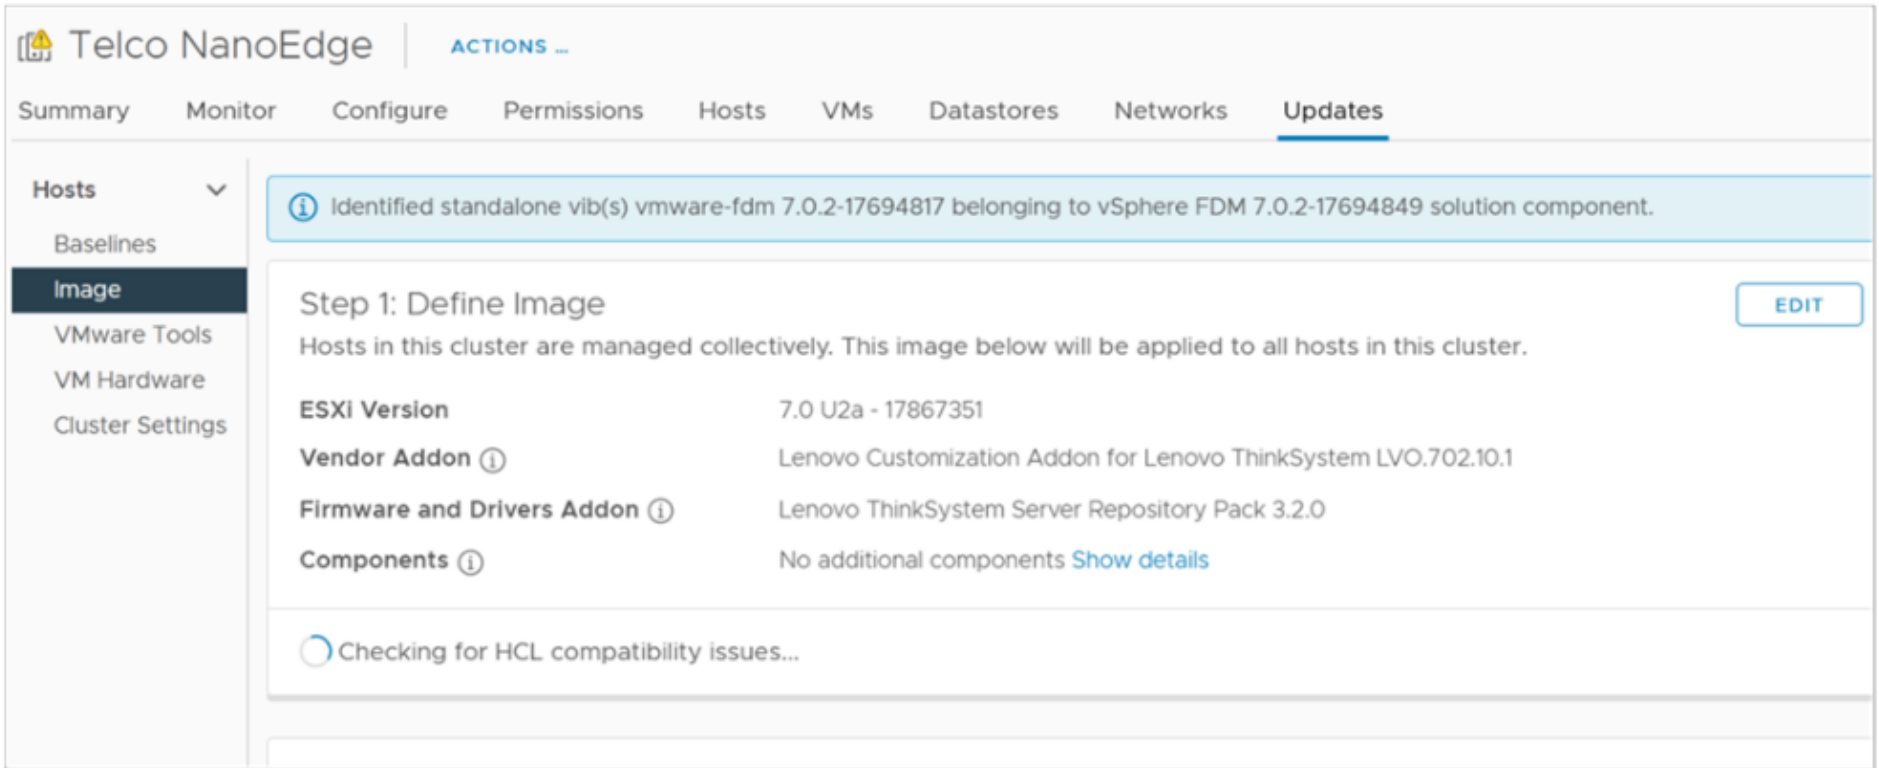

Manual ARM firmware update process

After installing the BMU (Bare Metal Update) OS, the firmware can be seen in XCC and updated with the new firmware update package.

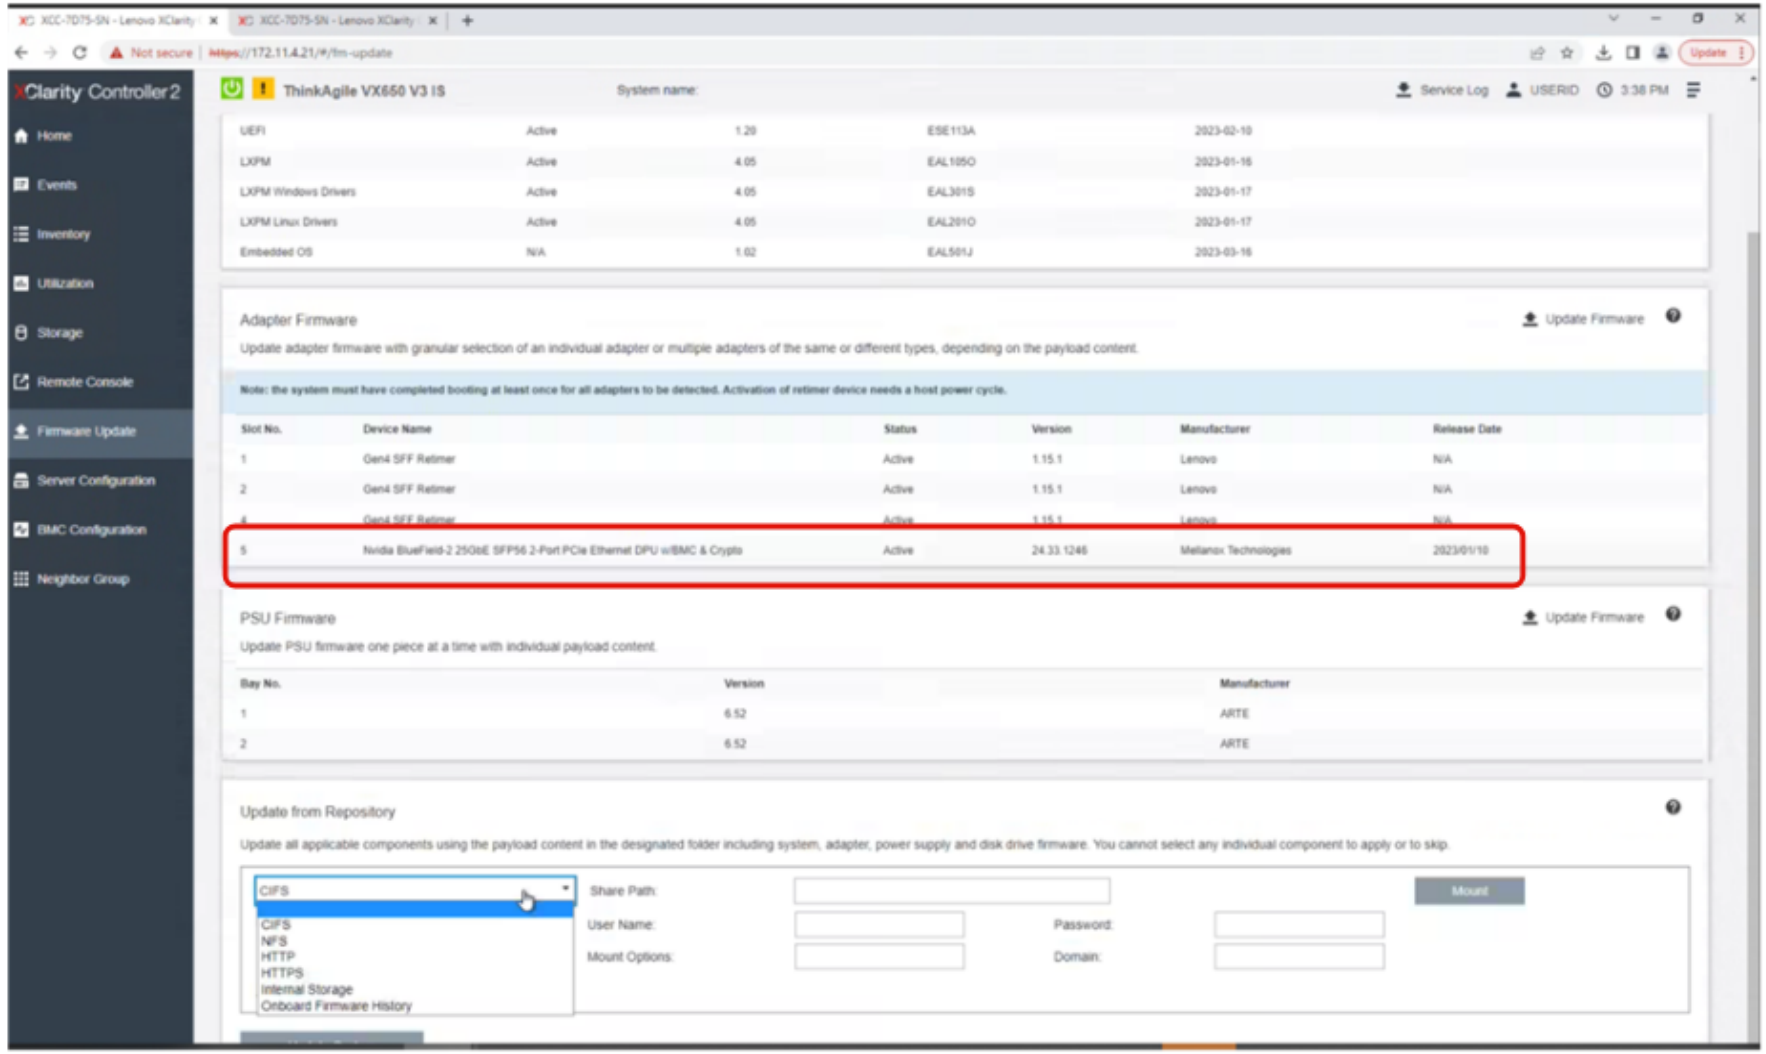

Enabling network offloads in vSphere

Work through the following steps to enable network offloads in vSphere:

1. Create a vSphere Distributed Switch
2. Add Hosts to a vSphere Distributed Switch
3. Configure NSX host transport node on DPU-enabled vSphere Lifecycle Manager cluster
4. View the Topology of Network Offloads Switch

Creating a new distributed switch with a DPU in vSphere

- Install ESXi on the boot device (M.2/7 mm) and DPU
- Create a distributed switch
  - In the vSphere Client, right-click a data center from the inventory tree
  - Datacenter➔ Distributed Switch➔ New Distributed Switch
- Create a name or accept the default generated name
- Select **Distributed Switch 8.0.0**
- Select **Type of Networking Offloads Capability**
- Select **Configure Settings➔ Network Offloads compatibility➔ NVIDIA BlueField**
- Select the number of uplinks (equal to the port number of the BlueField-2 adapter installed in the system)
- Enable **Network I/O Control**
- (Optional) Create port groups

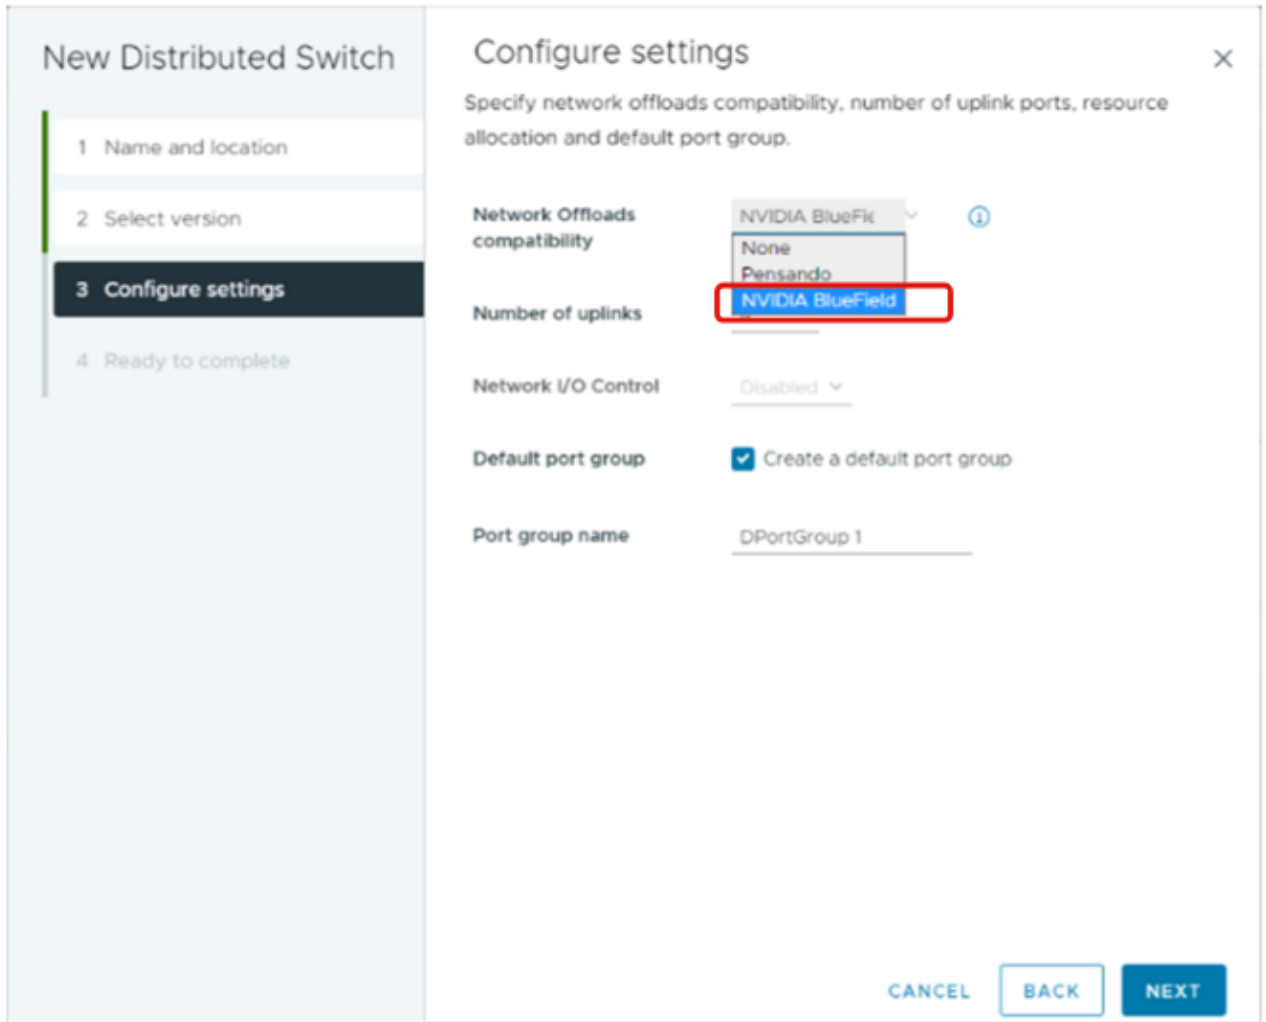

vSphere DPU identification

After completing the configuration, the DPU can be identified in vSphere.

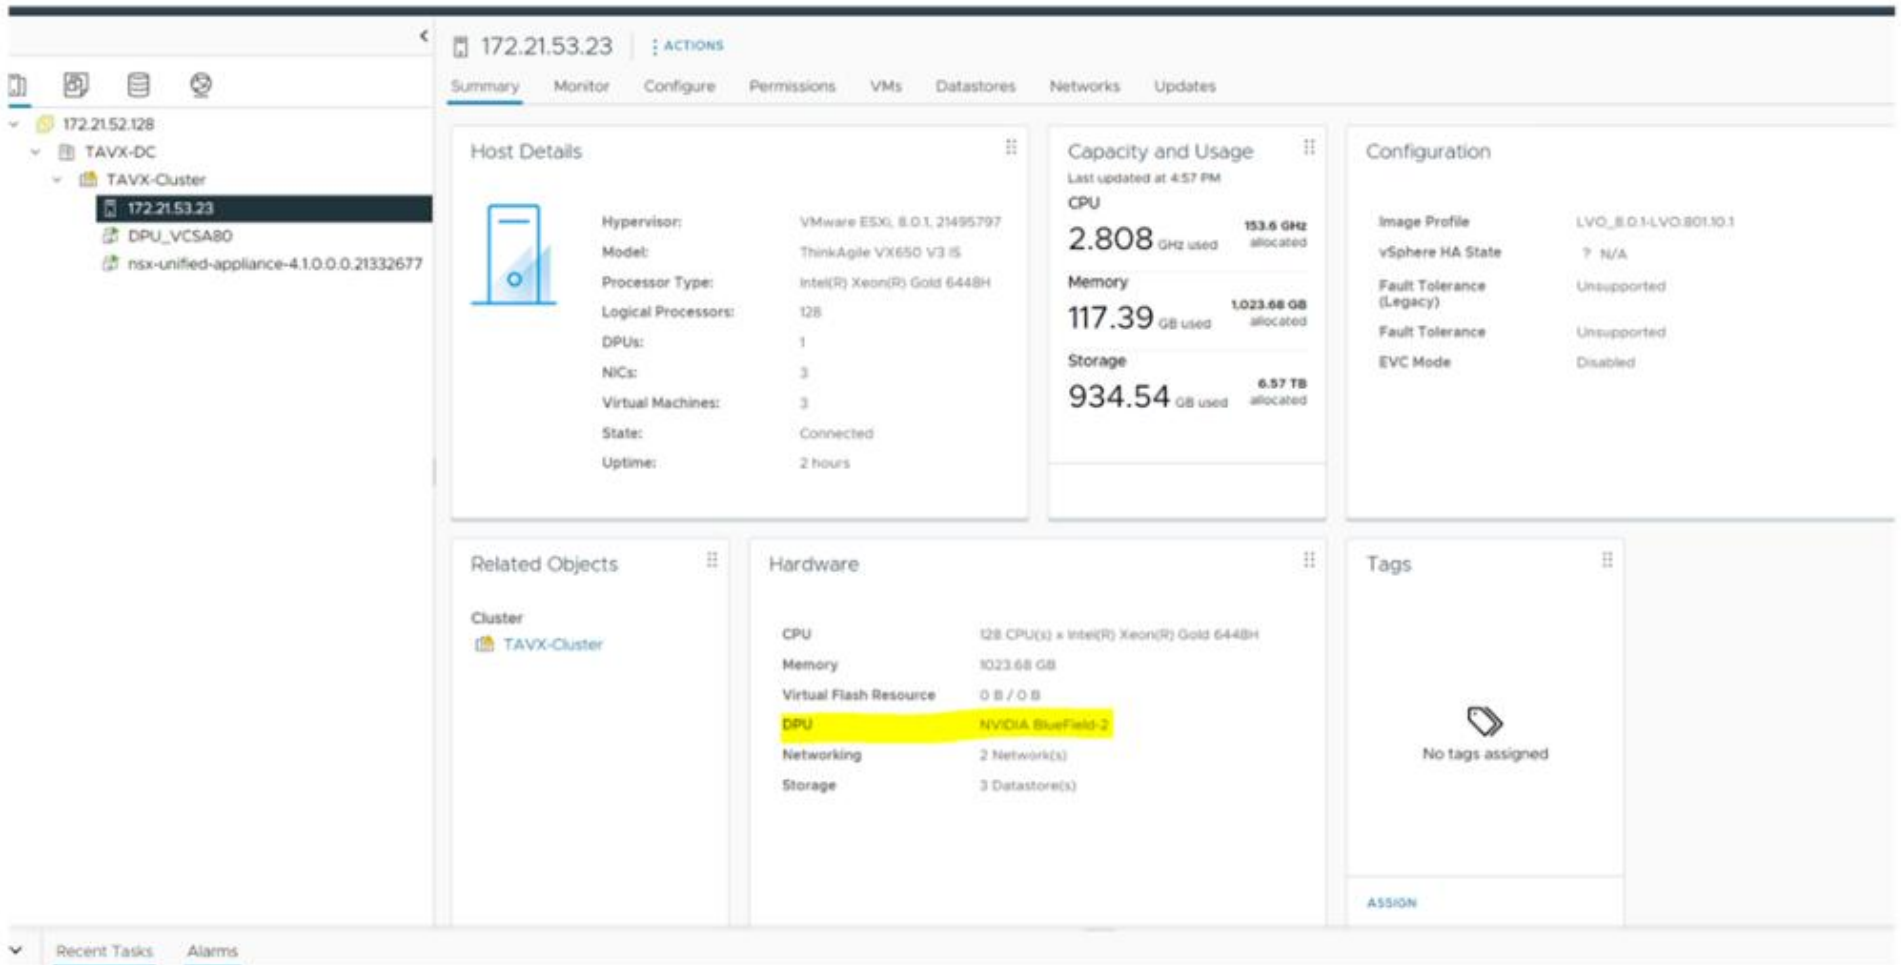

vSphere DPU identification - continued

After completing the configuration, the DPU can be identified in vSphere.