What's New in vSphere 8?

VMwareAI/ML

Table of contents

What's New in vSphere 8?

For what's new in vSphere 8 Update 1 see the following article.

What's New in vSphere 8 Update 1?

VMware vSphere 8 is the enterprise workload platform that brings the benefits of cloud to on-premises workloads. It supercharges performance with DPU and GPU based acceleration, enhances operational efficiency through the VMware Cloud Console, seamlessly integrates with add-on hybrid cloud services, and accelerates innovation with an enterprise-ready integrated Kubernetes runtime that runs containers alongside VMs.

The following is a highlight of what's new in vSphere 8 and not an exhaustive list of new features and capabilities. For more on vSphere 8, including deeper dives into certain topics, see core.vmware.com/vmware-vsphere-8.

VMware vSphere 8

Video: vSphere 8 What's New?

---

vSphere Distributed Services Engine

Introducing **vSphere Distributed Services Engine**, formerly Project Monterey. vSphere Distributed Services Engine unlocks the power of Data Processing Units (DPUs) for hardware accelerated data processing to improve infrastructure performance, boost infrastructure security and simplify DPU lifecycle management. vSphere 8 makes using DPUs easy for workloads to take advantage of these benefits.

Emergence of the Data Processing Unit

Data Processing Units (DPU) exist today and live in the hardware layer, similar to a PCIe device like a NIC or GPU. Today networking, storage and host management services run in the instance of ESXi virtualizing the x86 compute layer.

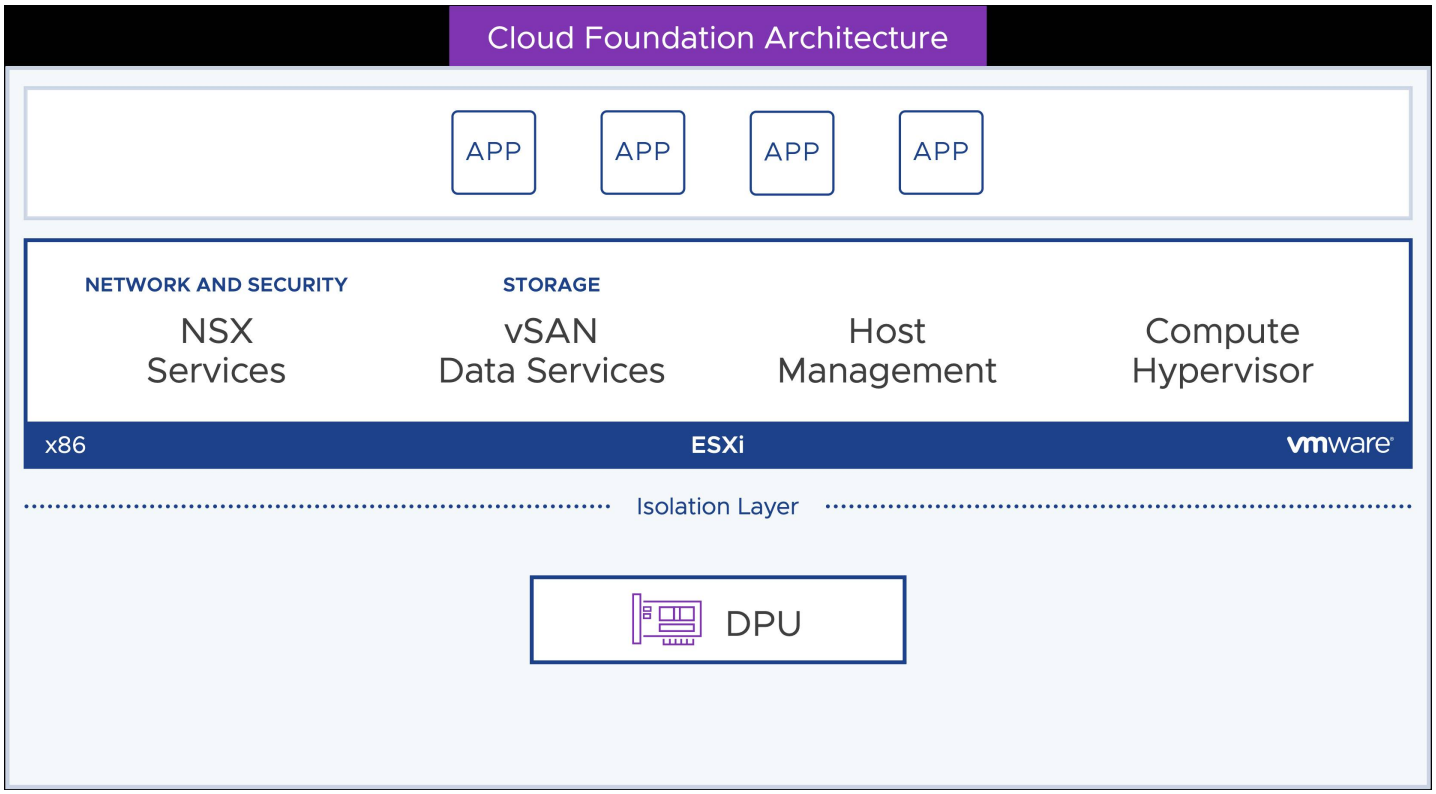

For more on Data Processing Units:

The Rise of DPUs in the Infrastructure

Meet vSphere Distributed Services Engine

In vSphere 8, an additional instance of **ESXi is installed directly on the Data Processing Unit**. This allows ESXi services to be offloaded to the DPU for increased performance.

In vSphere 8 GA. we support greenfield installations with support for network offloading with NSX. vSphere Distributed Services Engine is life-cycle managed using vSphere Lifecycle Manager. When remediating a host that contains a DPU ESXi installation, the DPU ESXi version is always remediated with the parent host and kept in version lock-step.

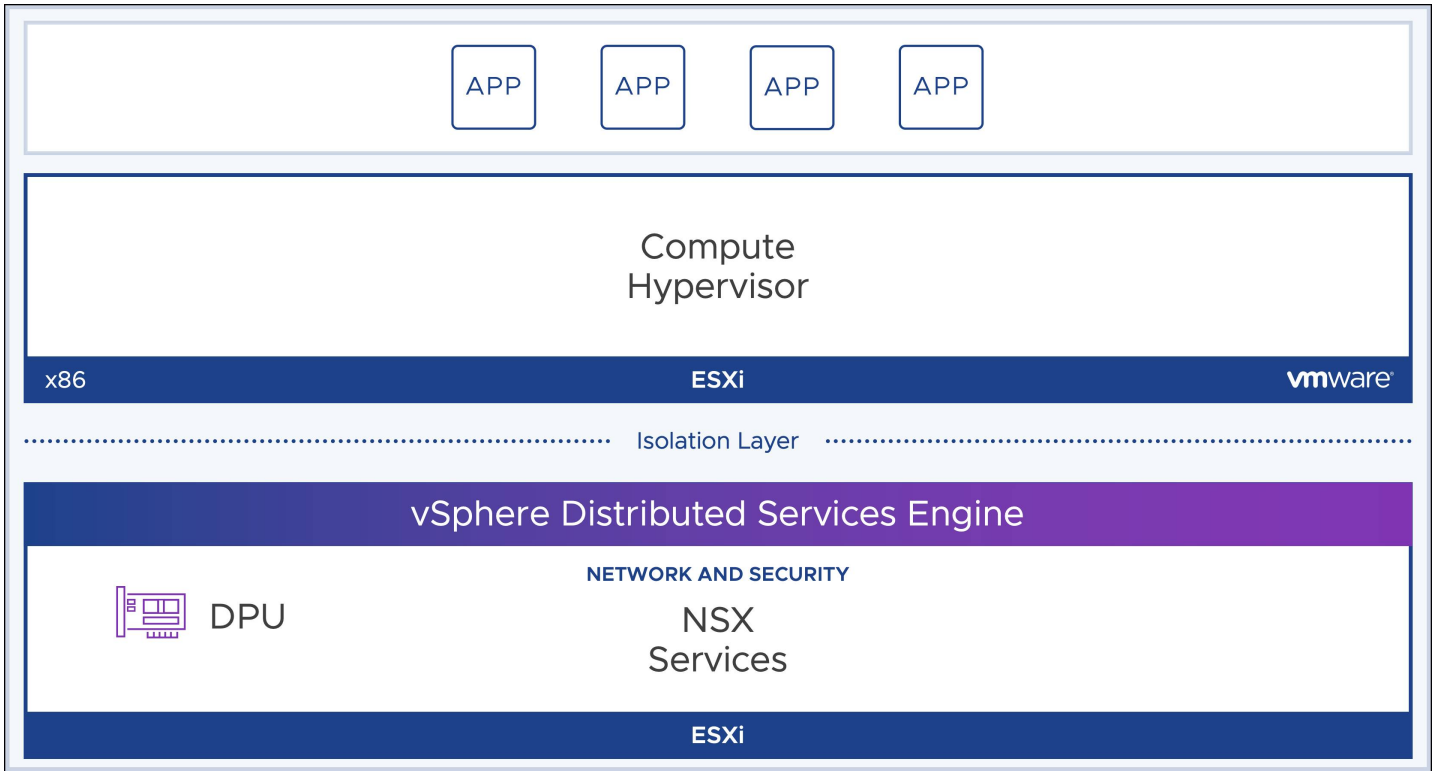

Simple Configuration for Network Offloads

Using a **vSphere Distributed Switch version 8.0 and NSX**, network services are offloaded to the DPU, allowing for increased network performance with no x86 CPU overhead, enhanced visibility for the network traffic and the security, isolation and protection you would expect from NSX.

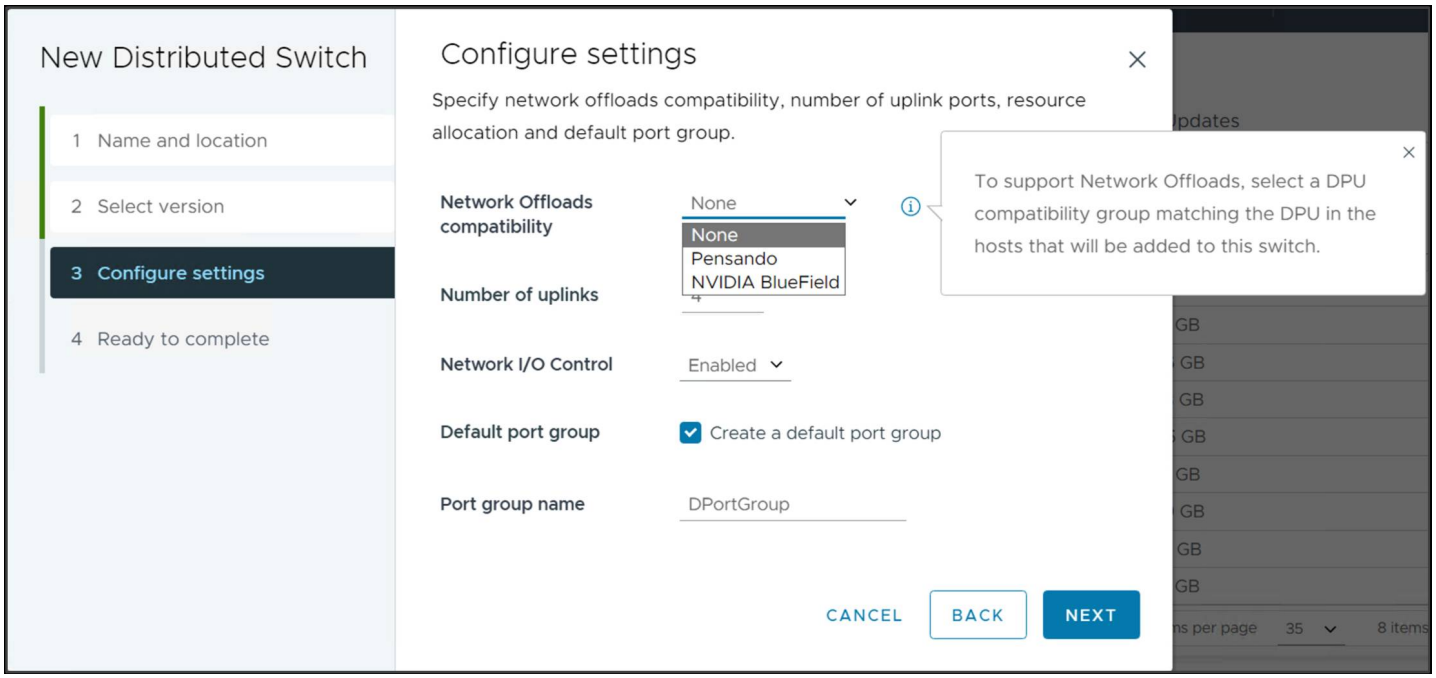

vSphere with Tanzu

Tanzu Kubernetes Grid on vSphere 8 consolidates the Tanzu Kubernetes offerings into a **single unified Kubernetes runtime** from VMware.

Workload Availability Zones are used to isolate workloads across vSphere clusters. supervisor clusters and Tanzu Kubernetes clusters can be deployed across zones to increase the availability of the clusters by ensuring that nodes are not sharing the same vSphere clusters.

ClusterClass is a way to declaratively specify your cluster's configuration through the open-source ClusterAPI project.

PhotonOS and Ubuntu base images can be customized and saved to the content library for use in Tanzu Kubernetes clusters.

Pinniped Integration comes to the supervisor clusters and Tanzu Kubernetes clusters. Pinniped supports LDAP and OIDC federated authentication. You can define identity providers that can be used to authenticate users to the supervisor clusters and Tanzu Kubernetes clusters.

Improving Workload Resilience in Modern Apps

Workload Availability Zones allows Supervisor Clusters and Tanzu Kubernetes Clusters to span across vSphere Clusters for increased availability. vSphere Namespaces span the Workload Availability Zones to support Tanzu Kubernetes Clusters being deployed for increased availability across zones.

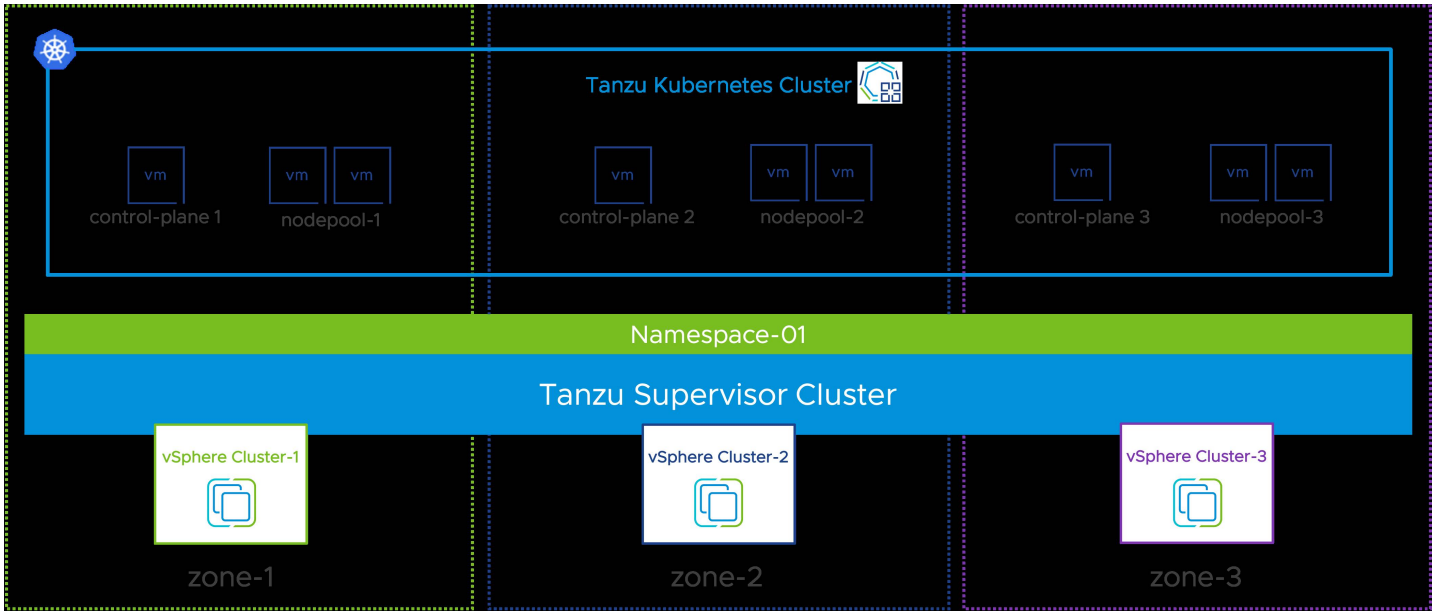

Three Workload Availability Zones are required for availability. During Workload Management activation, you have a choice to deploy across Workload Availability Zones or deploy to the single cluster option. At vSphere 8 GA, a Workload Availability Zone has a 1:1 relationship with a vSphere cluster.

Define It Once, Use It Many Times

ClusterClass provides a declarative way to define a Tanzu Kubernetes cluster configuration as well as the default installed packages. The platform team can decide the infrastructure packages that must be installed at cluster creation. This might include the networking, storage or cloud providers, as well as the authentication mechanism and metrics collection. The cluster specification references the ClusterClass.

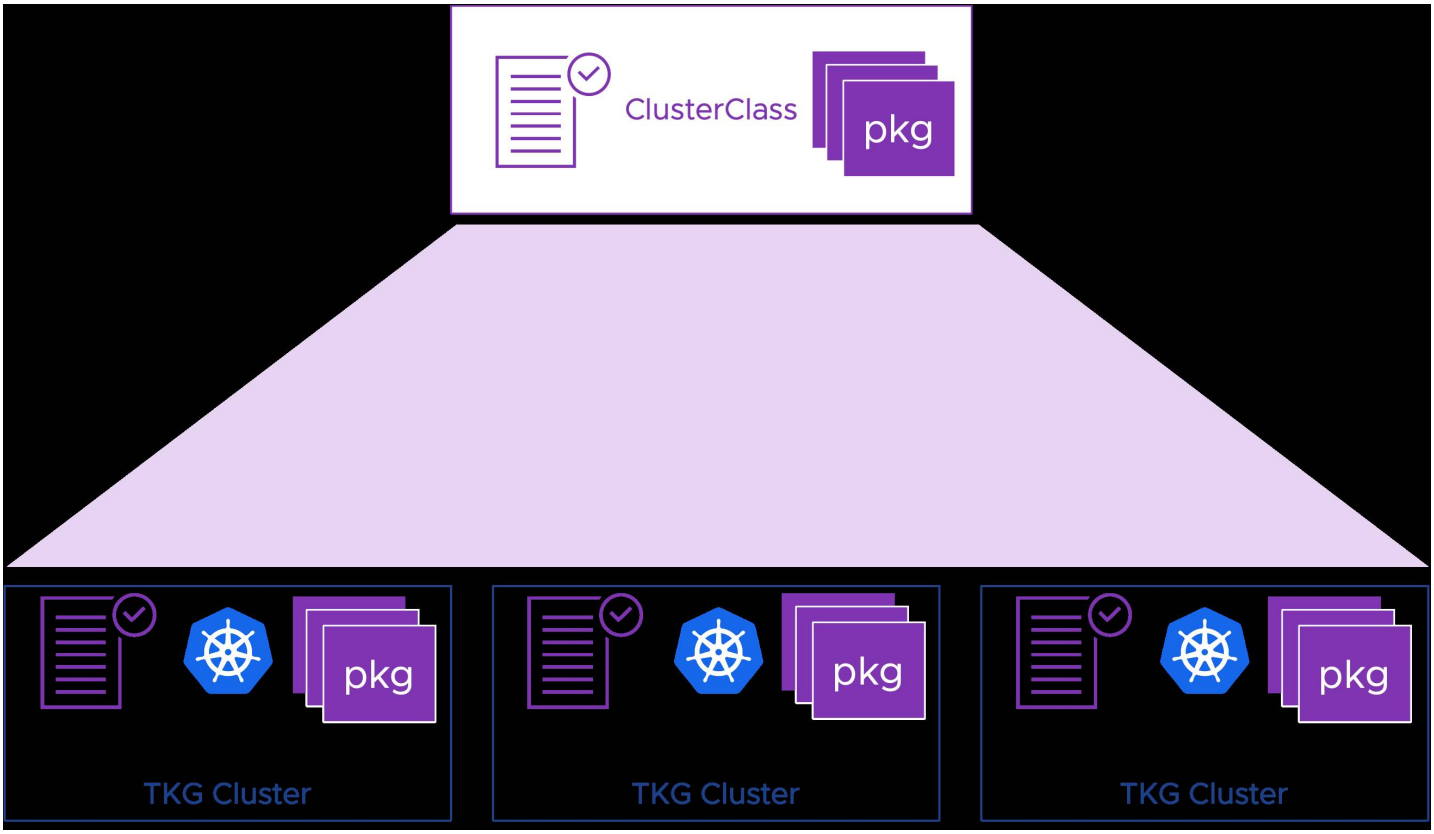

ClusterClass is an open-source specification that is part of the ClusterAPI project. ClusterAPI defines a declarative way to lifecycle manage Kubernetes clusters through an existing management Kubernetes cluster. In vSphere with Tanzu, that management cluster is the supervisor cluster.

Flexible Package Management

After cluster deployment, Developers or DevOps users can optionally add additional packages from the **Tanzu Standard Package Repository**. These packages might include Contour for ingress to the cluster, certificate management, logging, observability with Prometheus or Grafana or external DNS. These are managed as add-ons through the Tanzu CLI interface.

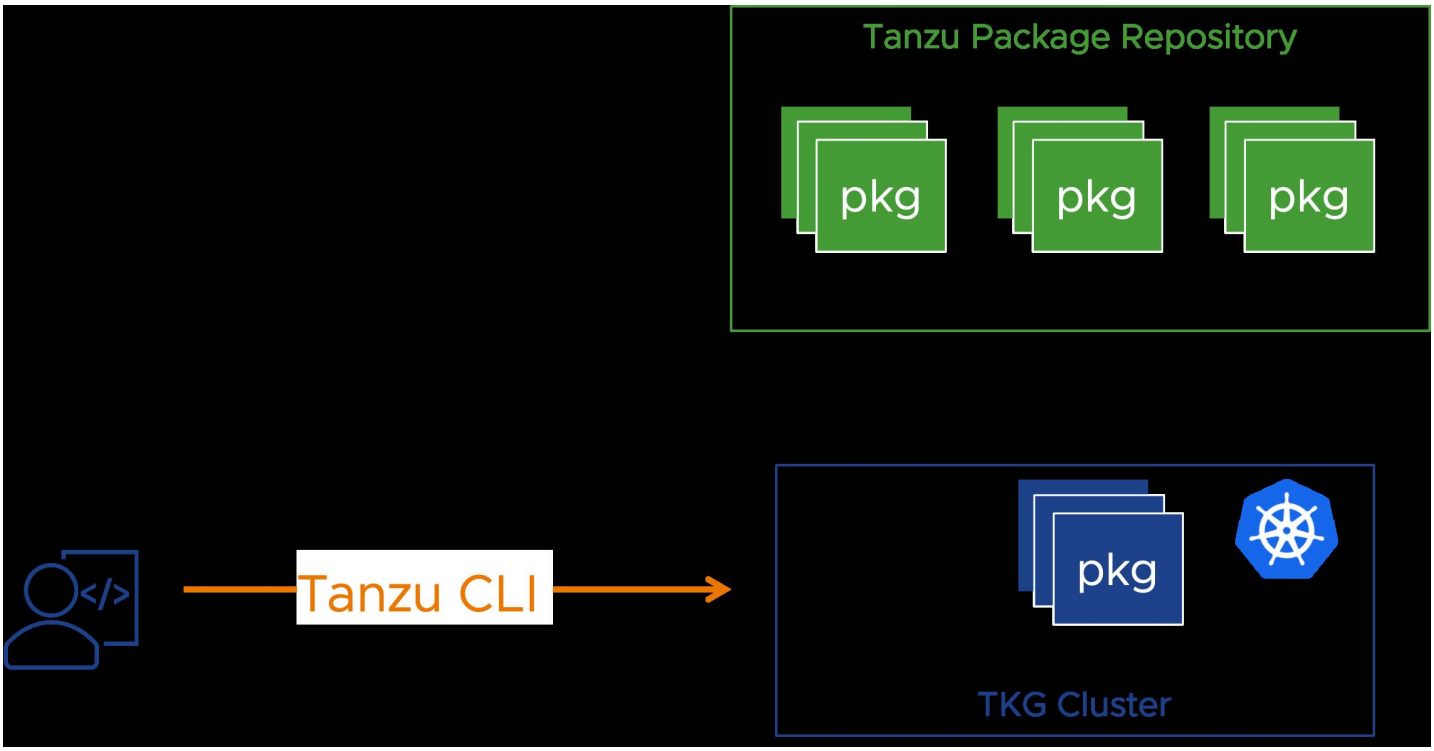

Bring Your Own Identity Provider

In vSphere 7, authentication is performed through integration with vCenter Single Sign-On. You can continue to use vCenter Single Sign-On in vSphere 8, but you also have an alternative. Using **Pinniped integration**, the supervisor cluster and Tanzu Kubernetes clusters can have direct access OIDC to an Identity Provider (IDP) without relying on vCenter Single Sign-On. Pinniped pods are automatically deployed in the supervisor cluster and Tanzu Kubernetes clusters to facilitate the integration.

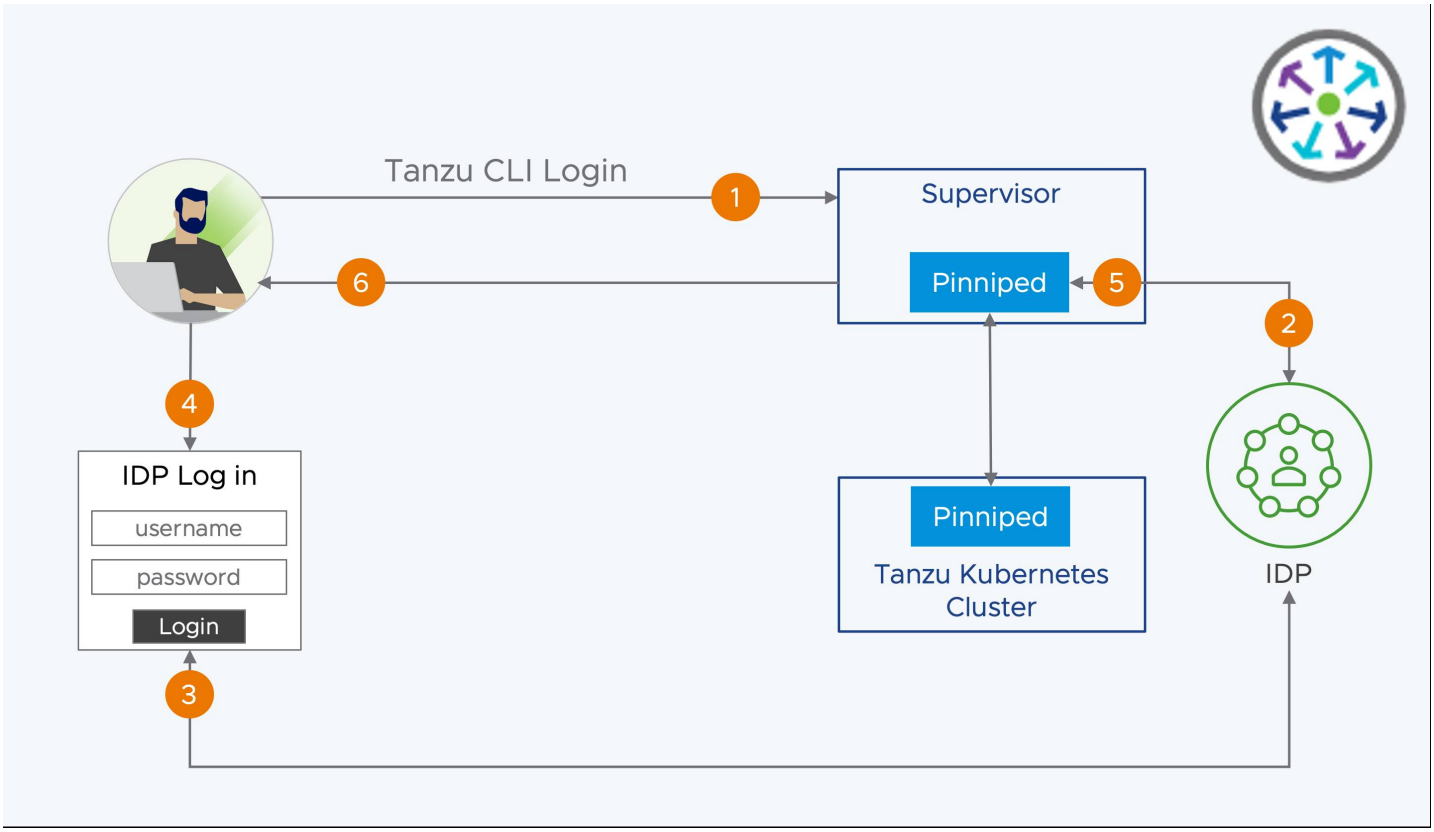

1. DevOps user uses Tanzu CLI login to authenticate to the Supervisor and/or TKC
2. Pinniped integration federates to an IDP
3. IDP returns a login link or window
4. DevOps user provides IDP credentials
5. Successful authentication to the IDP is returned to Pinniped
6. Tanzu CLI builds the kubeconfig file needed to access the Supervisor and/or TKC

Lifecycle Management

vSphere 8 introduces DPU support for vSphere Lifecycle Manager to automatically remediate the ESXi installation on a DPU in lock-step with the host ESXi version. Staging of update/upgrade payloads, parallel remediation and standalone host support combine to bring vLCM up to feature parity with Update Manager. Standalone hosts can be managed using vSphere Lifecycle Manager via API. VMware Compatibility Guide details what vLCM features a Hardware Support Manager can support.

A technical preview of vSphere Configuration Profiles introduces the next generation of cluster configuration management as a future replacement of Host Profiles.

Deprecation Awareness

Baseline lifecycle management, previously known as vSphere Update Manager, **is deprecated in vSphere 8.** This means that baseline lifecycle management is still supported in vSphere 8, but that vSphere 8 will be the last release that supports baseline lifecycle management.

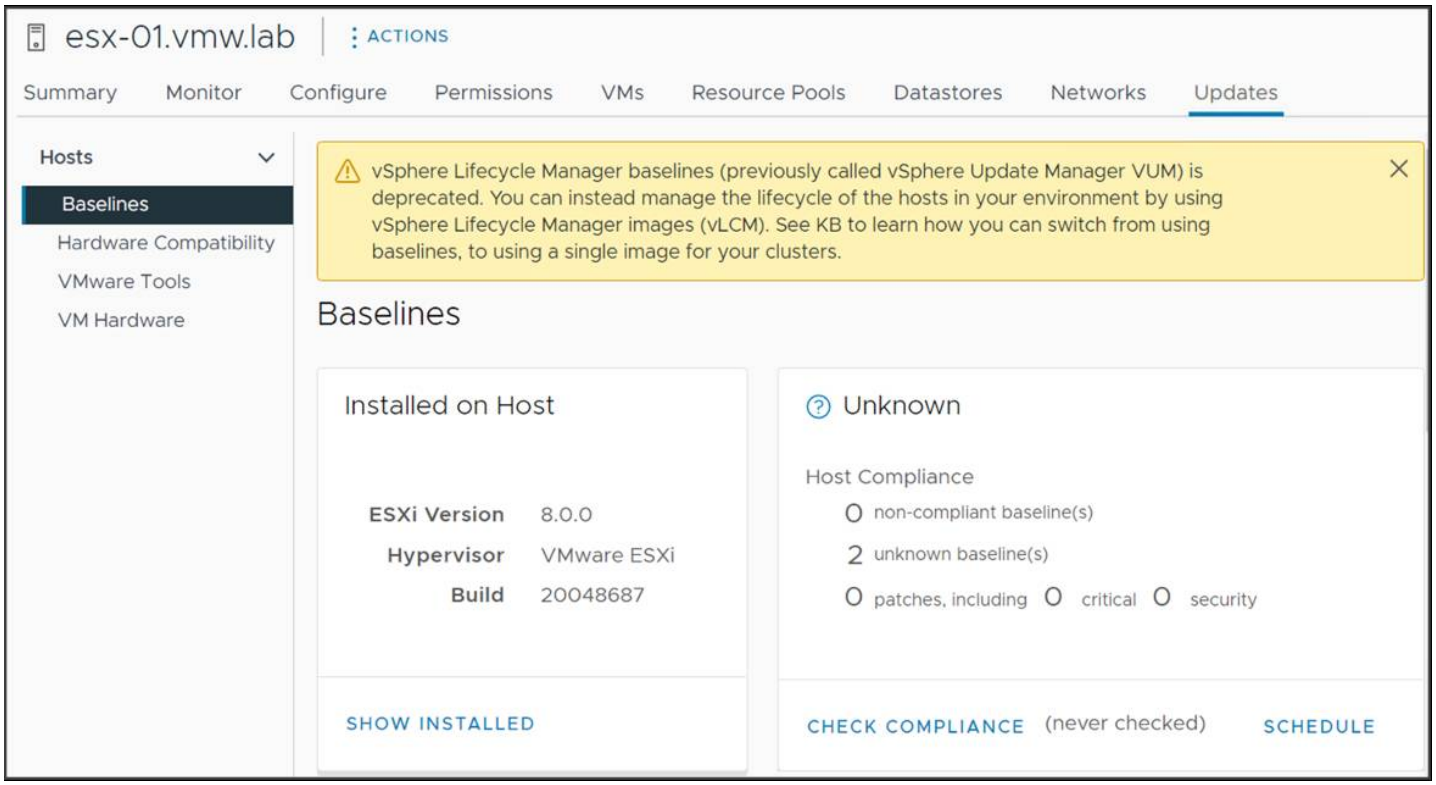

Stage Cluster Image Updates to Speed Up Remediation

vSphere Lifecycle Manager can stage update payloads to the hosts in advance of remediation. Staging can be performed without maintenance mode. This reduces the time needed for the hosts to spend in maintenance mode. Firmware payloads can also be staged with integration from a supported Hardware Support Manager.

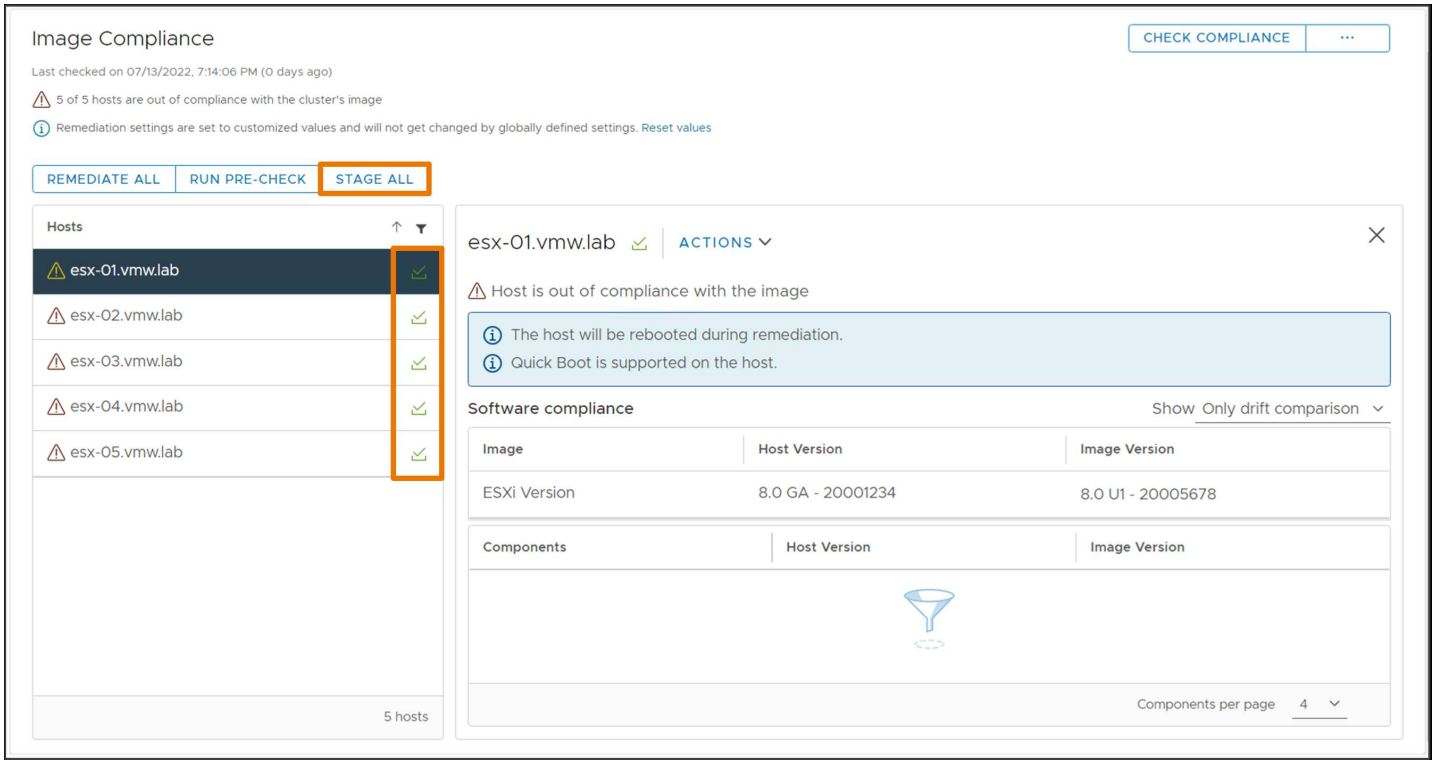

Quicker Cluster Remediation

vSphere Lifecycle Manager can **remediate multiple hosts in parallel**, dramatically reducing the overall time needed to remediate an entire cluster. Hosts placed into maintenance mode can be remediated in parallel. A vSphere administrator can choose to remediate all hosts in maintenance mode or define the number of parallel remediations to perform at a given time. Hosts not placed into maintenance mode are not remediated during this lifecycle operation.

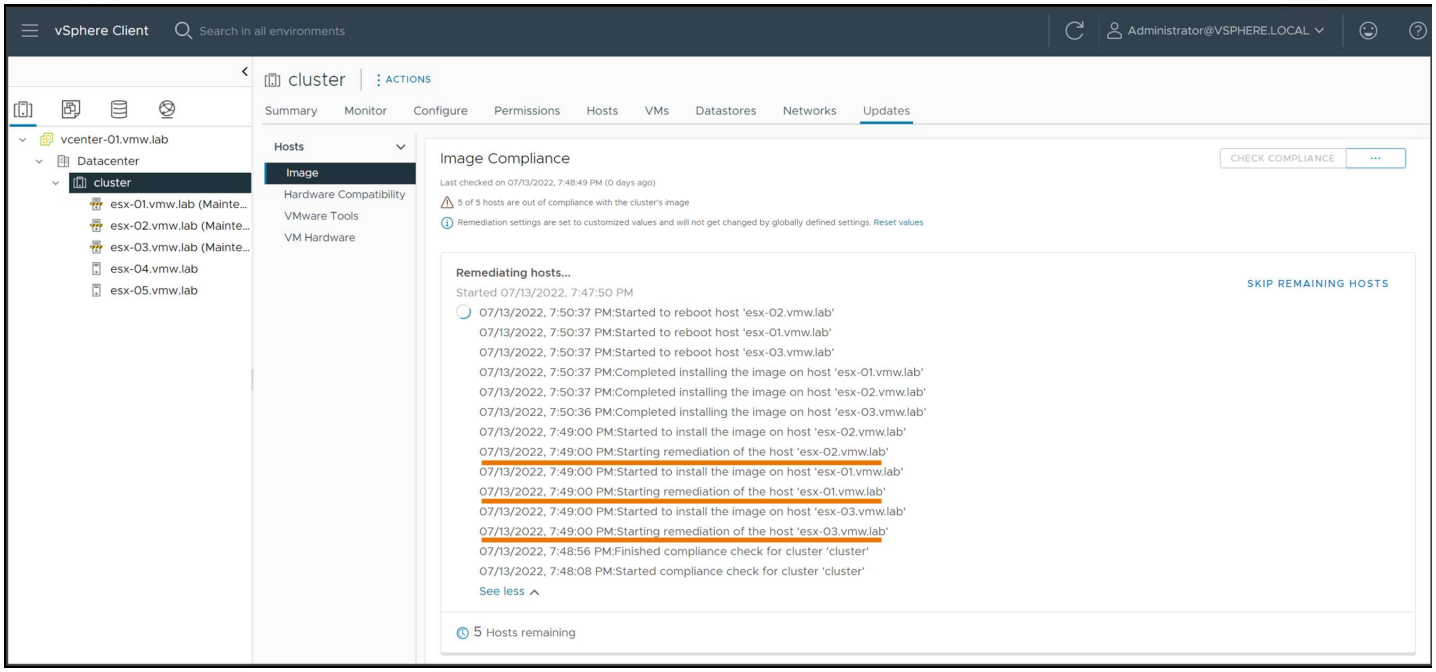

Configuration Management at Scale

vSphere 8 introduces a technical preview of **vSphere Configuration Profiles**, the next generation of cluster configuration management.

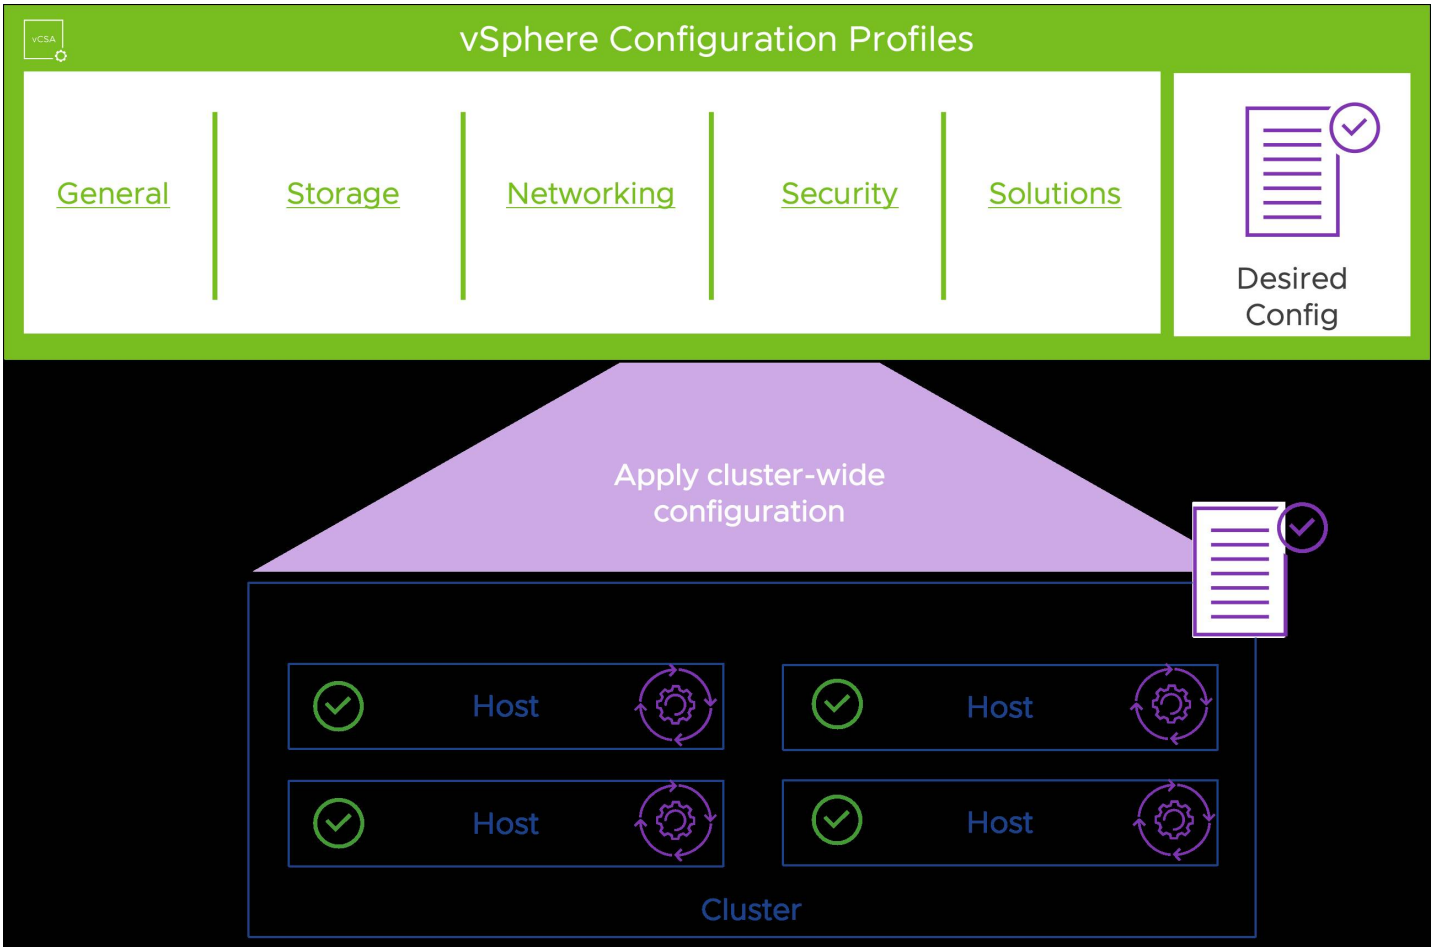

Desired configuration is defined at the cluster object and applied to all hosts in the cluster. All hosts in the cluster have a consistent configuration applied. Configuration drift is monitored for and notified about. A vSphere Administrator can remediate the configuration drift.

vSphere Configuration Profiles is still under active development and future releases of vSphere 8 will expand and enhance its support. Host Profiles continue to be supported in vSphere 8.

Enhanced Recovery of vCenter

vCenter reconciles cluster state after a restore from backup. ESXi hosts in a cluster contain a distributed key-value store of cluster state.

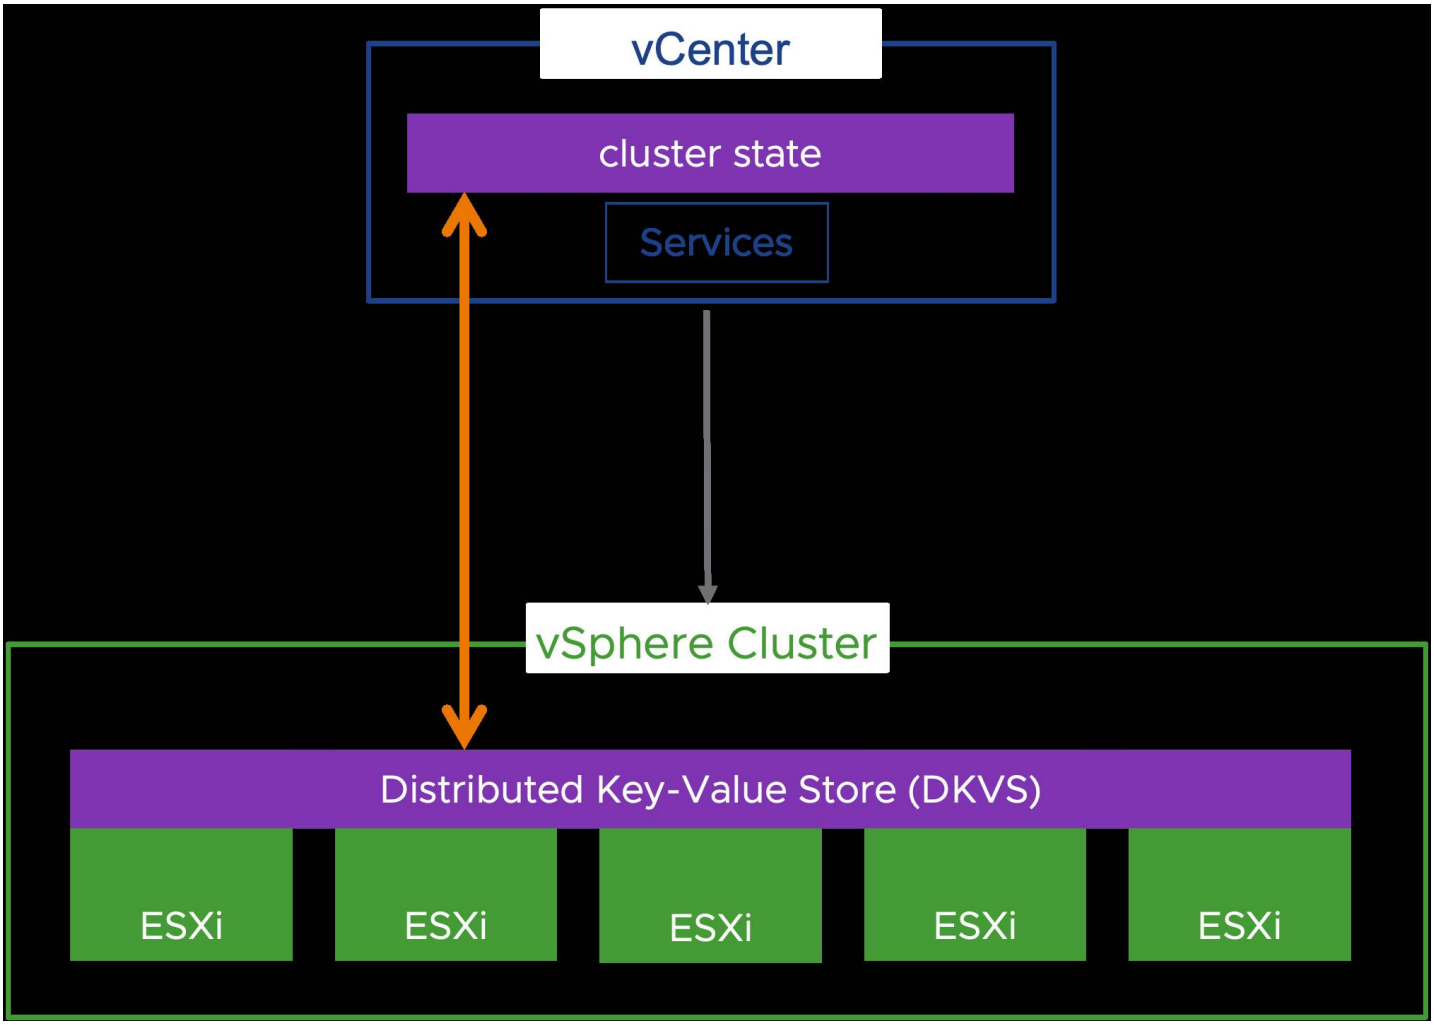

The distributed key-value store is the source of truth for the state of the cluster. If vCenter is restored from a backup, it will reconcile the cluster state and configuration with the distributed key-value store. In vSphere 8 GA, host-cluster membership is reconciled with additional configuration and state planned for support in future releases.

AI & ML

Unified Management for AI/ML Hardware Accelerators

Device Groups makes Virtual Machines consuming complementary hardware devices simpler in vSphere 8. NIC and GPU devices are supported in vSphere 8 GA. Compatible vendor device drivers are required and subject to vendor release. NVIDIA® will be the first partner supporting Device Groups with upcoming compatible drivers.

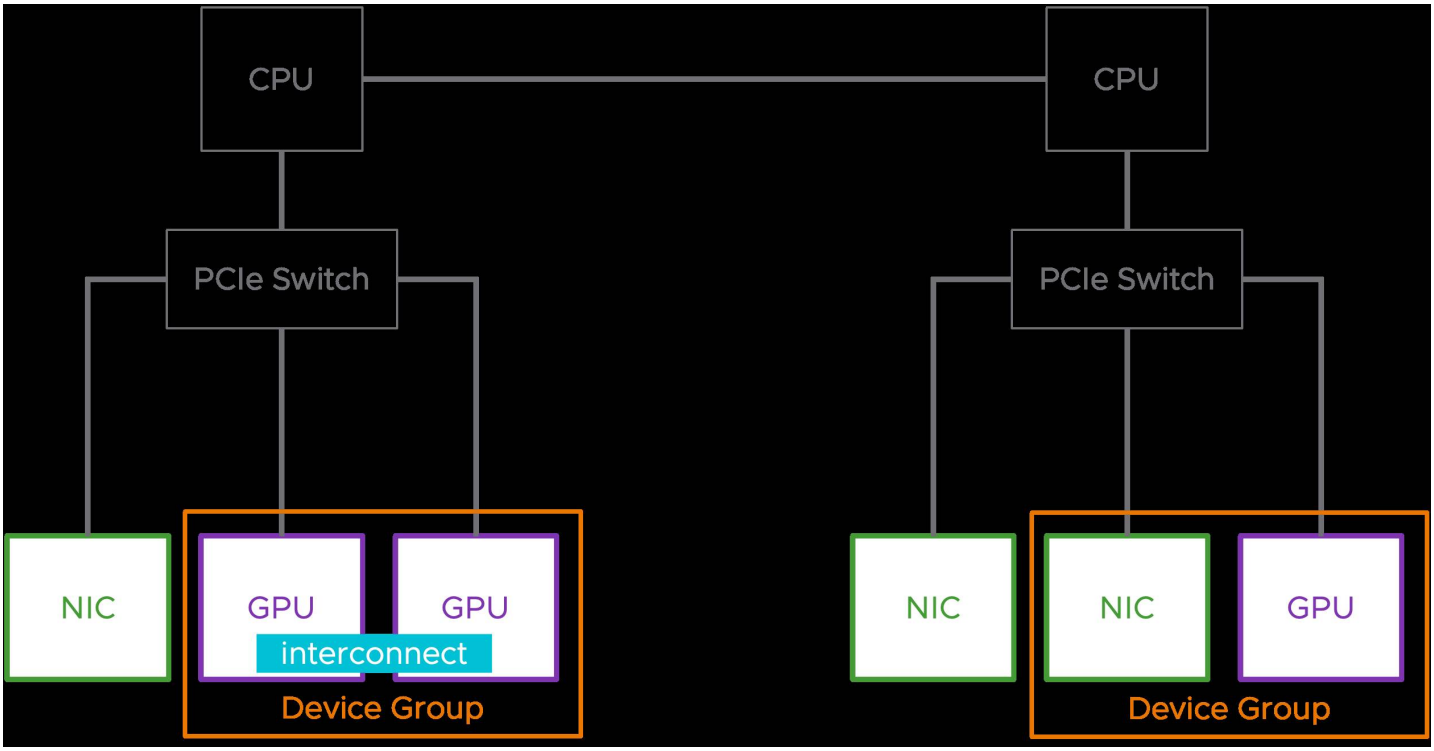

Device groups can be composed of two or more hardware devices that share a common PCIe switch or devices that share a direct interconnect between each other. Device groups are discovered at the hardware layer and presented to vSphere as a single unit that represents the group.

Simplified Hardware Consumption with Device Groups

Device Groups are added to virtual machines using the existing Add New PCI Device workflows. vSphere DRS and vSphere HA are aware of device groups and will place VMs appropriately to satisfy the device group.

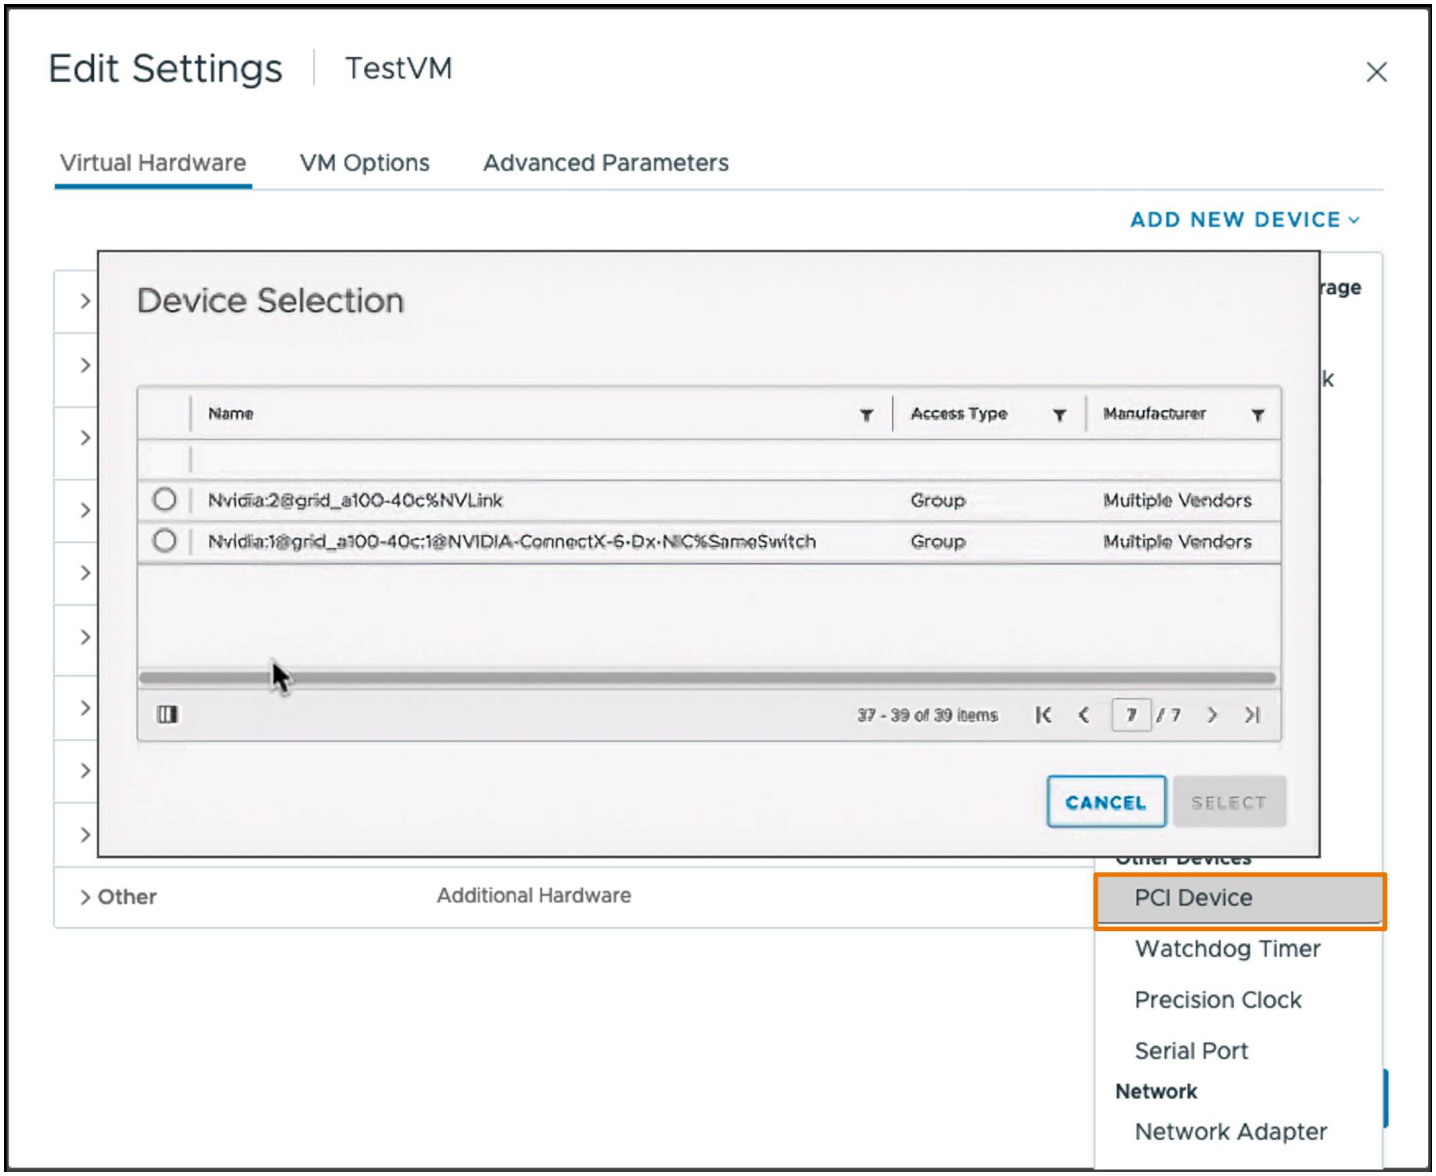

Next-Generation of Virtual Hardware Devices

Enhanced VMDirectPath I/O (also known as Device Virtualization Extensions or DVX) builds on Dynamic DirectPath I/O and introduces a new framework and API for vendors to create hardware-backed virtual devices. Enhanced VMDirectPath I/O allows greater support for virtualization features such as live migration using vSphere vMotion, suspending and resuming a virtual machine and support for disk and memory snapshots.

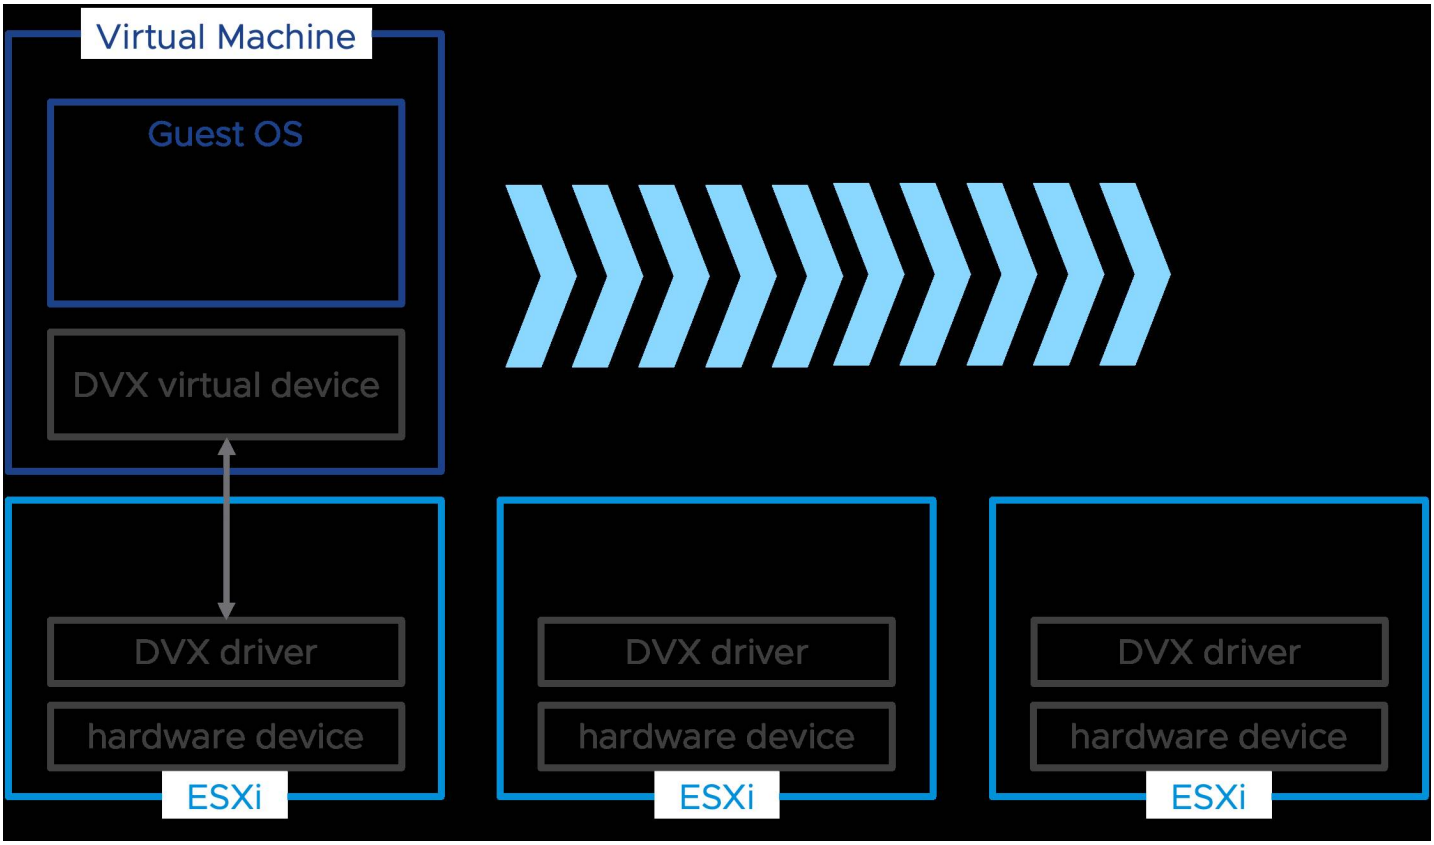

A compatible driver must be installed on the ESXi hosts and accompanied with a corresponding guest OS driver for the virtual device equivalent. A virtual machine consuming an Enhanced VMDirectPath I/O virtual device can be migrated using vSphere vMotion to another host that supports that same virtual device.

Guest OS & Workloads

Virtual Hardware Version 20

Virtual Hardware version 20 is the latest virtual hardware version. The theme for hardware version 20 brings new virtual hardware innovations, enhances guest services for applications, and increases performance and scale for certain workloads.

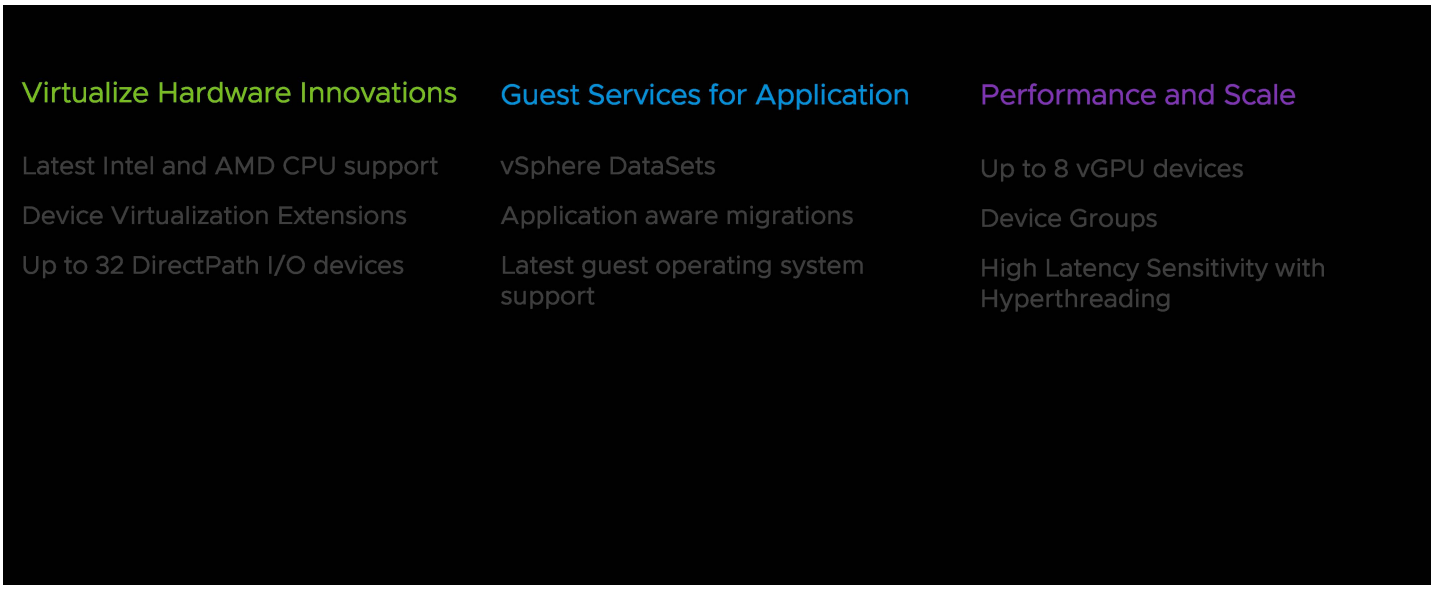

Deploy Windows 11 at Scale

Introducing **TPM Provision Policy**. Windows 11 requires vTPM devices to be present in virtual machines. Cloning a VM with a vTPM VM can introduce a security risk as TPM secrets are cloned.

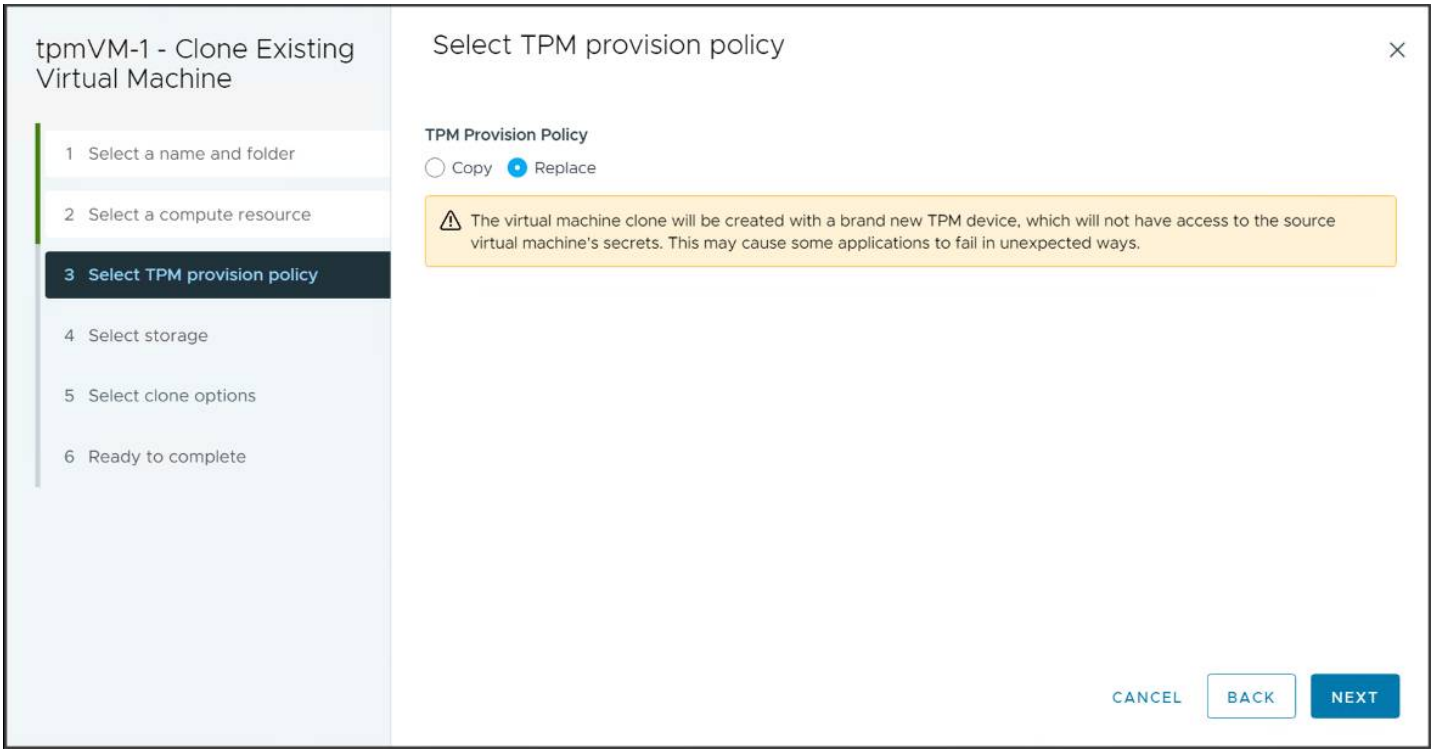

In vSphere 8 vTPM devices can be automatically replaced during clone or deployment operations. This allows best practices that each VM contain a unique TPM device be followed and improves vSphere support for Windows 11 deployment at scale. vSphere 8.0 also includes the **vpxd.clone.tpmProvisionPolicy** advanced setting to make the default clone behaviour for vTPMs to be replace.

Note: The TPM provision policy is used for any virtual machines that contain a virtual TPM device. It is not exclusively for Windows 11 VMs.

Reduce Outages by Preparing Applications for Migration

Certain applications cannot tolerate the stuns associated with vSphere vMotion. These applications can be written to be **migration aware** to improve their interoperability with vSphere vMotion. Applications can prepare for a migration event. This could be gracefully stopping services or performing a failover in the case of a clustered application. The application can delay the start of the migration up until the configured timeout but cannot decline or prevent the migration from occurring.

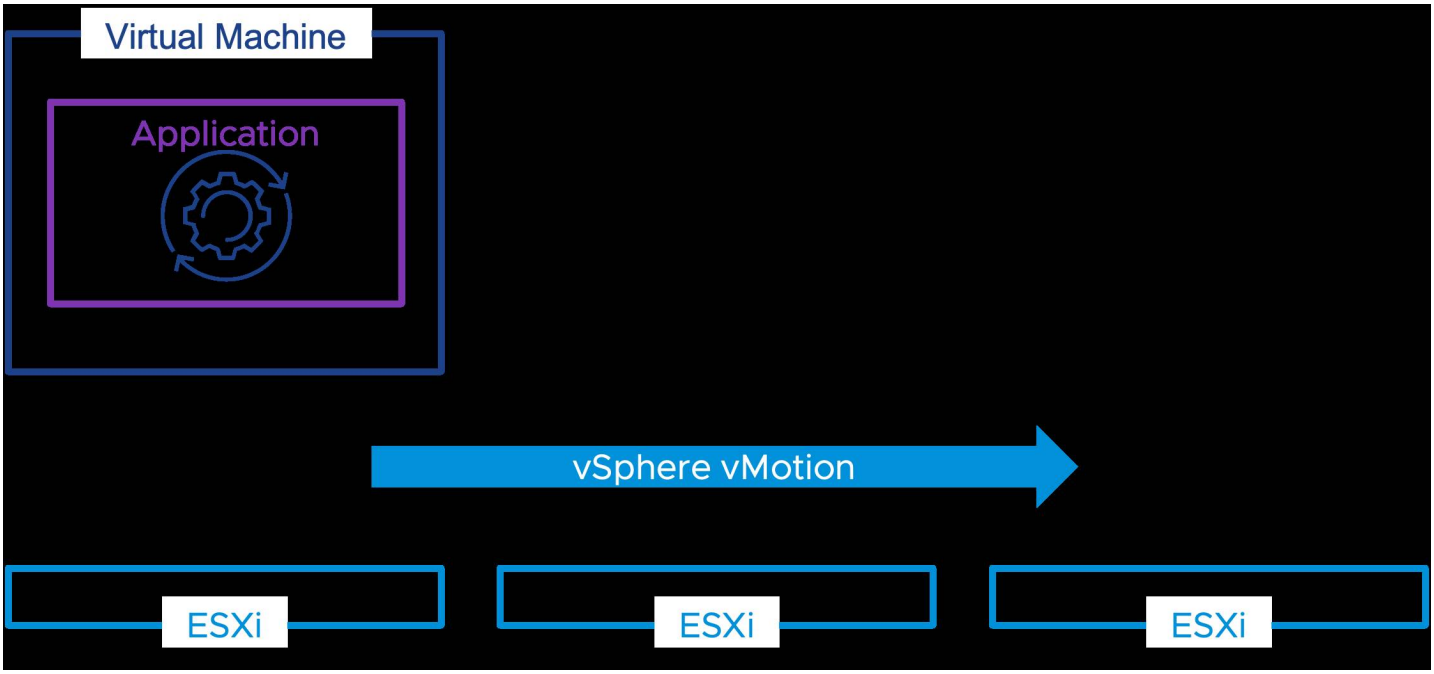

Use-cases include:

- Time-sensitive applications
- VoIP applications
- Clustered applications

Maximize Performance for Latency Sensitive Workloads

Emerging Telco workloads require increased support for latency sensitivity applications. **High Latency Sensitivity with Hyper-threading** is designed to support these workloads and deliver improved performance. A virtual machine's vCPU are scheduled on the same hyper-threaded physical CPU core.

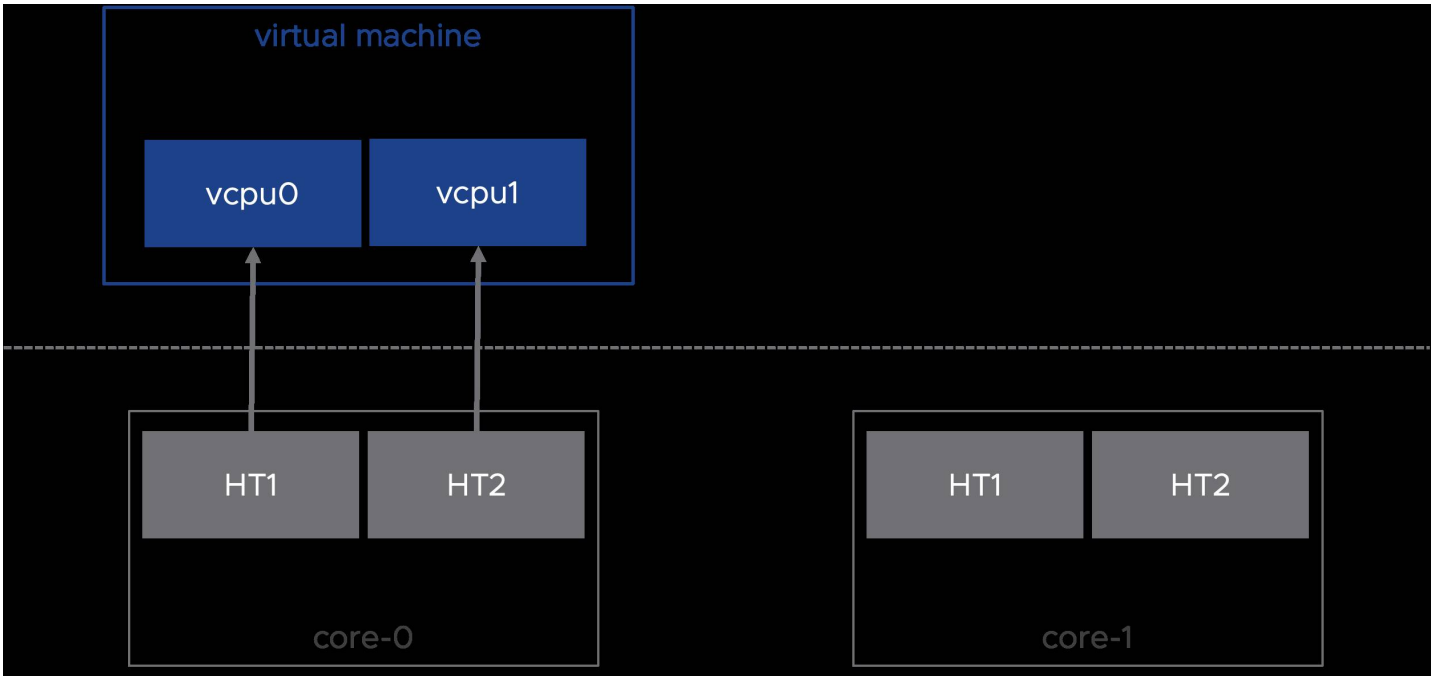

High Latency Sensitivity with Hyper-threading requires virtual machine hardware version 20 and is configurable in the Advanced settings for a virtual machine.

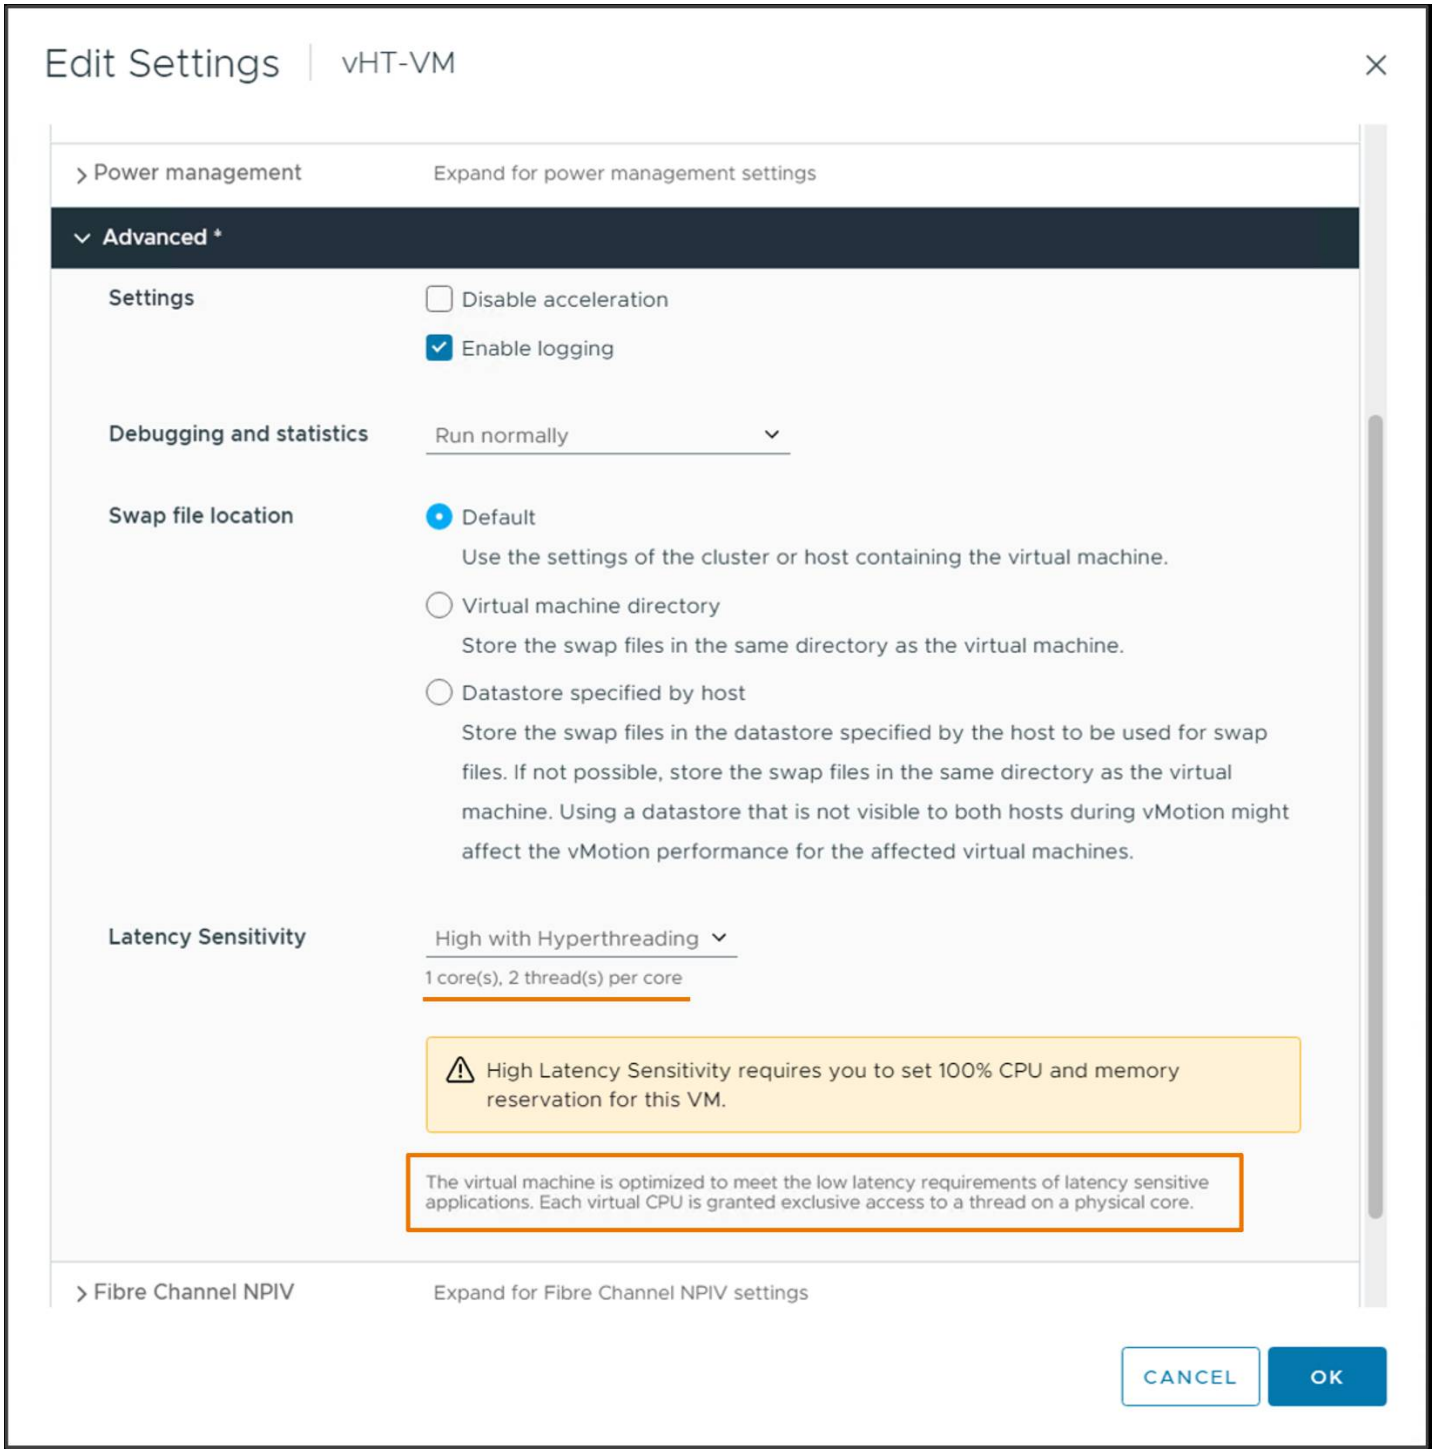

Simplified Virtual NUMA Configuration

vSphere 8 and hardware version 20 allows you to use the vSphere Client to configure the vNUMA topology for virtual machines.

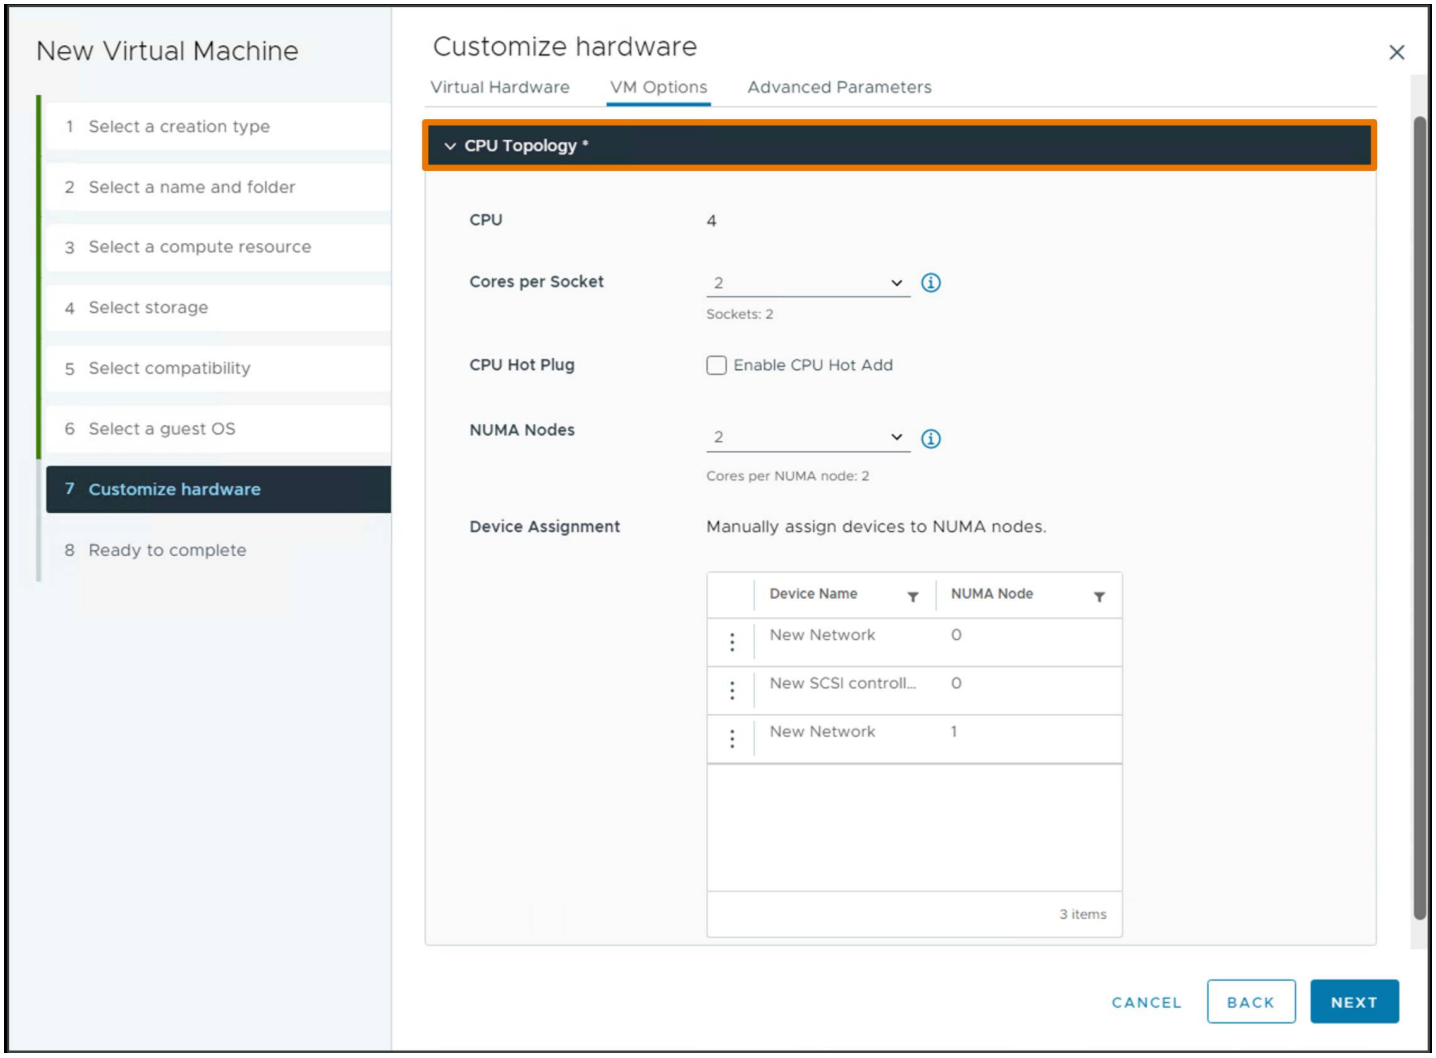

A new CPU topology tile is visible on the VM summary tab displaying the current topology.

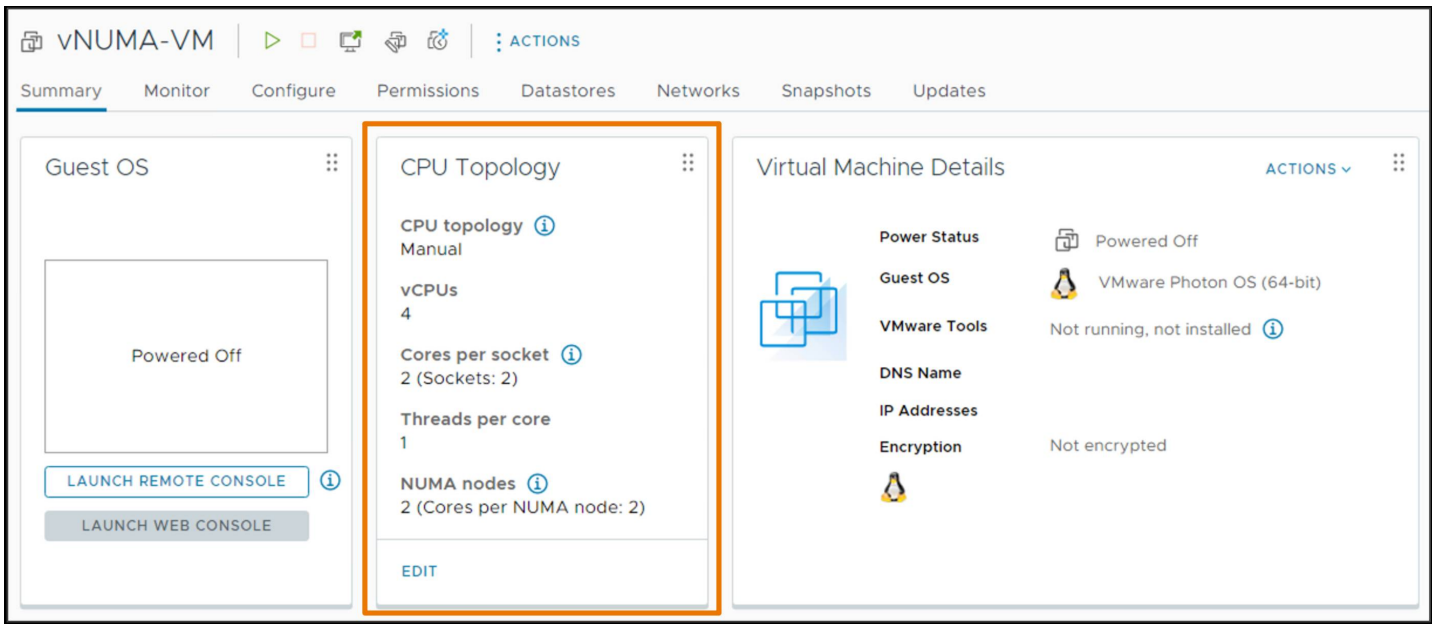

API Driven vSphere and Guest Data Sharing

vSphere DataSets provide an easy method distribute small, infrequently changing data between the vSphere management layer and a guest operating system running in a virtual machine with VMware Tools installed. Uses cases may include Guest deployment

status, Guest agent configuration or Guest inventory management. vSphere DataSets live with the VM object, and will move with the VM if migrated, even across vCenter Server instances.

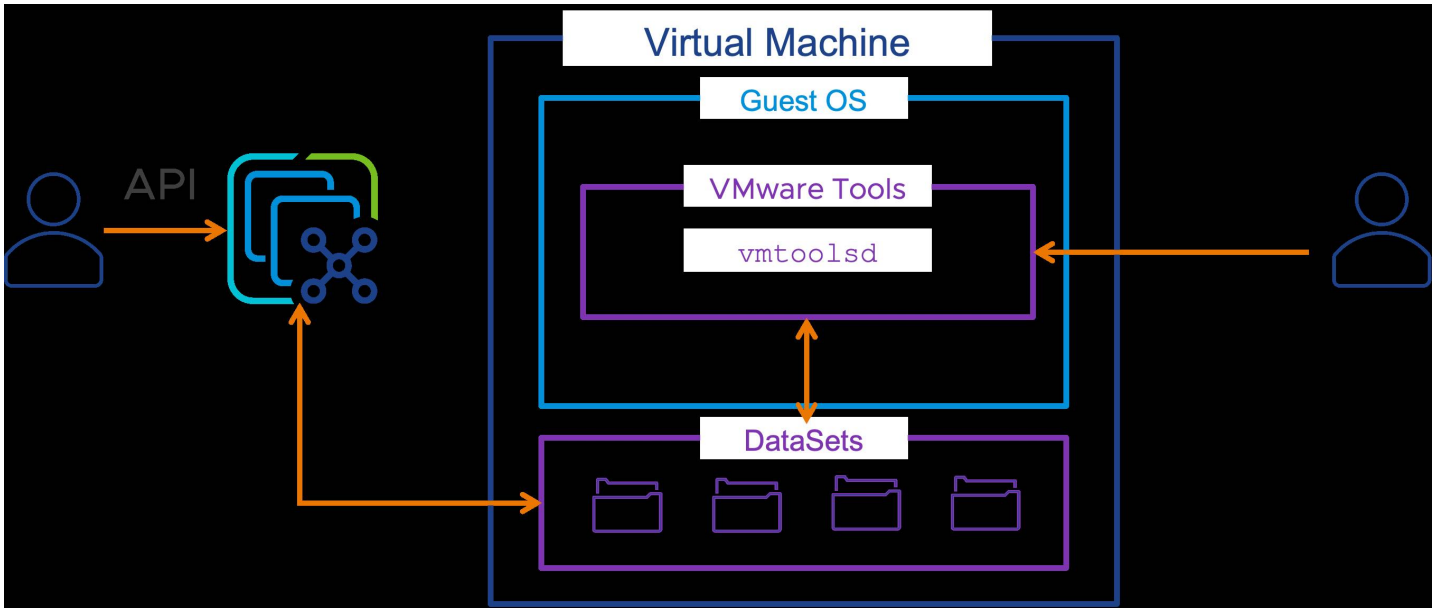

Resource Management

Enhanced DRS Performance

VMware has introduced a new feature with vSphere 7.0U3 called **vSphere Memory Monitoring and Remediation** (vMMR). vMMR helps bridge the need for monitoring by providing running statistics at both the VM (bandwidth) and Host levels (bandwidth, miss-rates). vMMR also provides default alerts and ability to configure custom alerts based on the workloads that run on VMs. vMMR collects data and provides visibility of performance statistics so you can determine if your application workload is regressed due to Memory Mode.

In vSphere 8, DRS performance can be significantly improved when PMEM is present by leveraging memory statistics, resulting in optimal placement decisions for VMs without affecting performance and resource consumption.

For more on Persistent Memory (Pmem):

Understanding Persistent Memory (PMem) in vSphere

Monitor Energy and Carbon Emissions

vSphere Green Metrics introduces new power consumption metrics for hosts and virtual machines. These metrics allow administrators monitor the energy consumption of the vSphere infrastructure and determine, based on the energy sources used to power the data center, the energy efficiency of the vSphere infrastructure.

The three new metrics track :

1. power.capacity.usageSystem: Power consumption of a host's system activities; how much power the host is using not attributed to VMs
2. power.capacity.usageSystem: Power consumption of a host's idle activity; how much power the host is using when it's not doing anything except being on
3. power.capacity.usageVm: Power consumption of a host due to VM workloads; how much power the host is using to run VM workloads

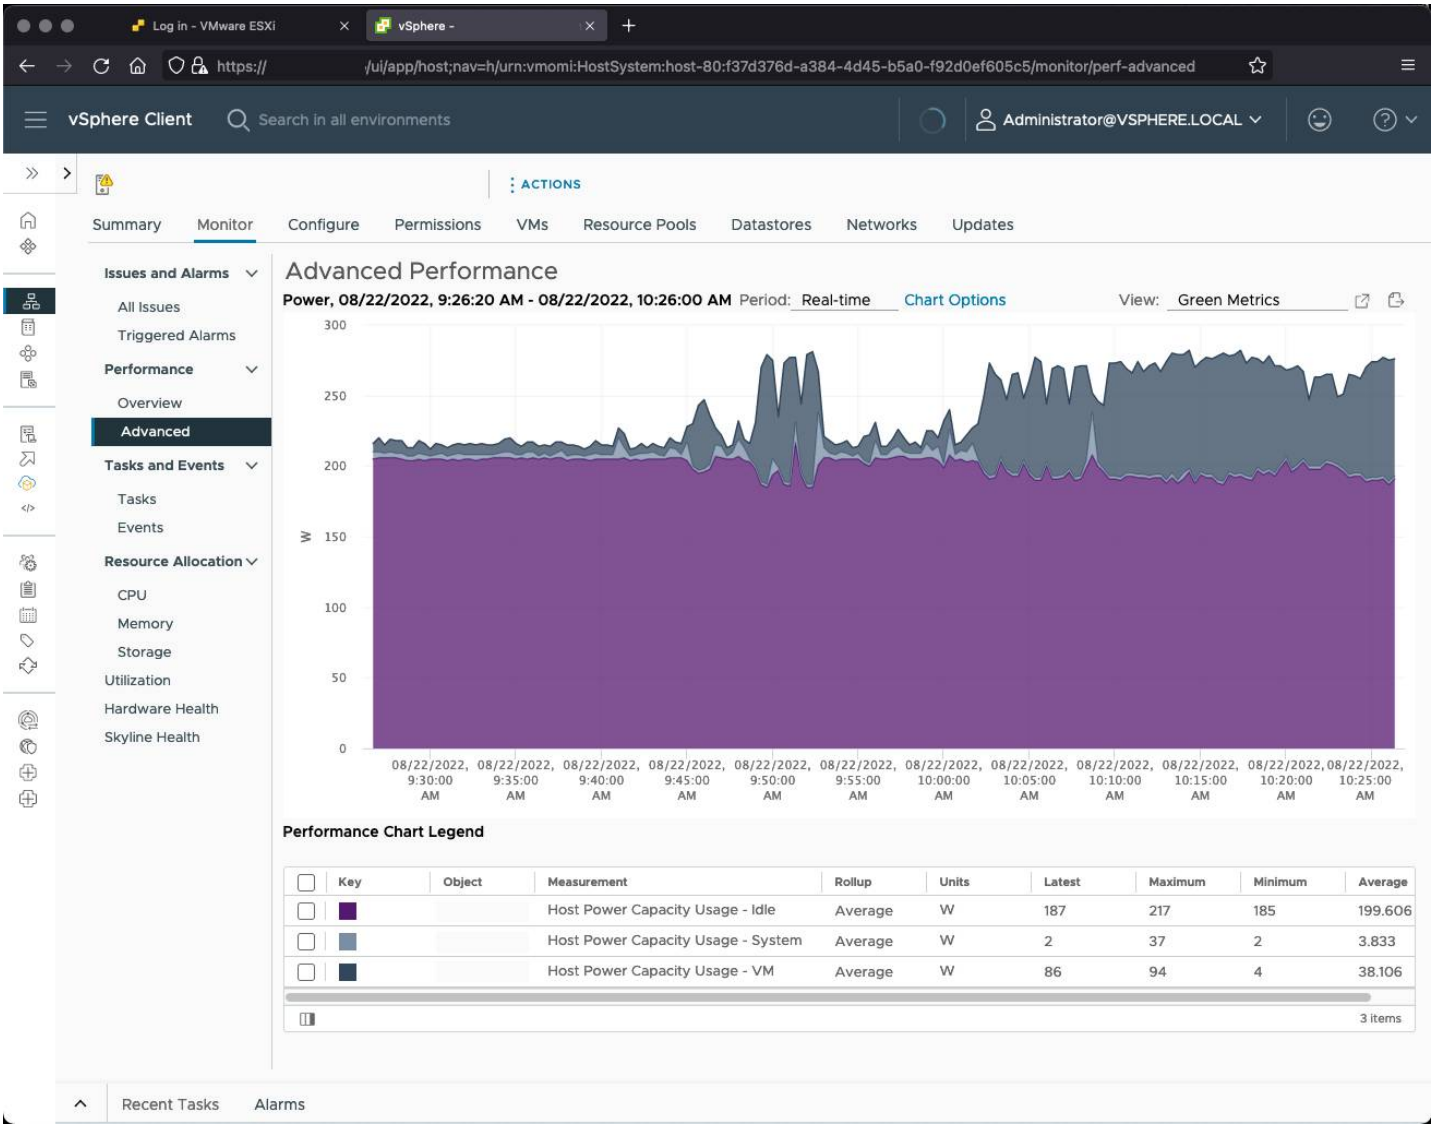
Log in - VMware ESXi
vSphere -
https://
/ui/app/host;nav=h/urn:vmomi:HostSystem:host-80:f37d376d-a384-4d45-b5a0-f92d0ef605c5/monitor/perf-advanced
vSphere Client
Search in all environments
Administrator@VSPHERE.LOCAL
ACTIONS
Summary
Monitor
Configure
Permissions
VMs
Resource Pools
Datastores
Networks
Updates
Issues and Alarms
All Issues
Triggered Alarms
Performance
Overview
Advanced
Tasks and Events
Tasks
Events
Resource Allocation
CPU
Memory
Storage
Utilization
Hardware Health
Skyline Health
Advanced Performance
Power, 08/22/2022, 9:26:20 AM - 08/22/2022, 10:26:00 AM
Period: Real-time
Chart Options
View: Green Metrics
300
250
200
150
100
50
0
W
08/22/2022, 9:30:00 AM
08/22/2022, 9:35:00 AM
08/22/2022, 9:40:00 AM
08/22/2022, 9:45:00 AM
08/22/2022, 9:50:00 AM
08/22/2022, 9:55:00 AM
08/22/2022, 10:00:00 AM
08/22/2022, 10:05:00 AM
08/22/2022, 10:10:00 AM
08/22/2022, 10:15:00 AM
08/22/2022, 10:20:00 AM
08/22/2022, 10:25:00 AM
Performance Chart Legend
Key
Object
Measurement
Rollup
Units
Latest
Maximum
Minimum
Average
Host Power Capacity Usage - Idle
Average
W
187
217
185
199.606
Host Power Capacity Usage - System
Average
W
2
37
2
3.833
Host Power Capacity Usage - VM
Average
W
86
94
4
38.106
3 items
Recent Tasks
Alarms

Security & Compliance

vSphere strives to be secure out of the box. In vSphere 8, further measures are taken to make vSphere secure by default.

Prevent execution of untrusted binaries: ESXi 8.0 will turn on the execInstalledOnly option by default. This prevents the execution of binaries that are not installed via a VIB.

TLS 1.2 only: vSphere 8 will not support TLS 1.0 and TLS 1.1. Both have previously been disabled by default in vSphere 7 and are now removed in vSphere 8.

Sandboxed Daemons: ESXi 8.0 daemons and processes run in their own sandboxed domain where only the minimum required permissions are available to the process.

Discontinuation of Trusted Platform Module (TPM) 1.2: ESXi 8.0 displays a warning during installation or upgrade if a TPM 1.2 device is present. The install or upgrade is not prevented.

```
This host has TPM1.2 hardware which is no longer supported.  For full use of vSphere features, use TPM 2.0.
```