



June 13-15, 2023

DoubleTree by Hilton San Jose

SmartNICsSummit.com

The Datacenter need for DPUs

Paul Turner, VP Products VMware

Disclaimer




- This presentation may contain product features or functionality that are currently under development.
- This overview of new technology represents no commitment from VMware to deliver these features in any generally available product.
- Features are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind.
- Technical feasibility and market demand will affect final delivery.
- Pricing and packaging for any new features/functionality/technology discussed or presented, have not been determined.
- This information is confidential.

Broad Interest across every vertical

Segment	Use case
Federal	Sovereign Cloud. Securely consolidate multiple private clouds across data centers into a single datacenter
Financials	High Performance Workloads. Onboard network sensitive bare metal workloads onto VMware stack
Tier 2 CSPs	Server Fleet Management. Secure, multi-tenant, composable bare metal cloud with VMware
Healthcare	Secure Private Cloud. Run VMware's Private Cloud leveraging all security capabilities w/o CPU overhead
Retail	Secure Private Cloud. Running VMware's Private Cloud outside the datacenter at the Retail Edge
Telco	Emerging

Macro Trends

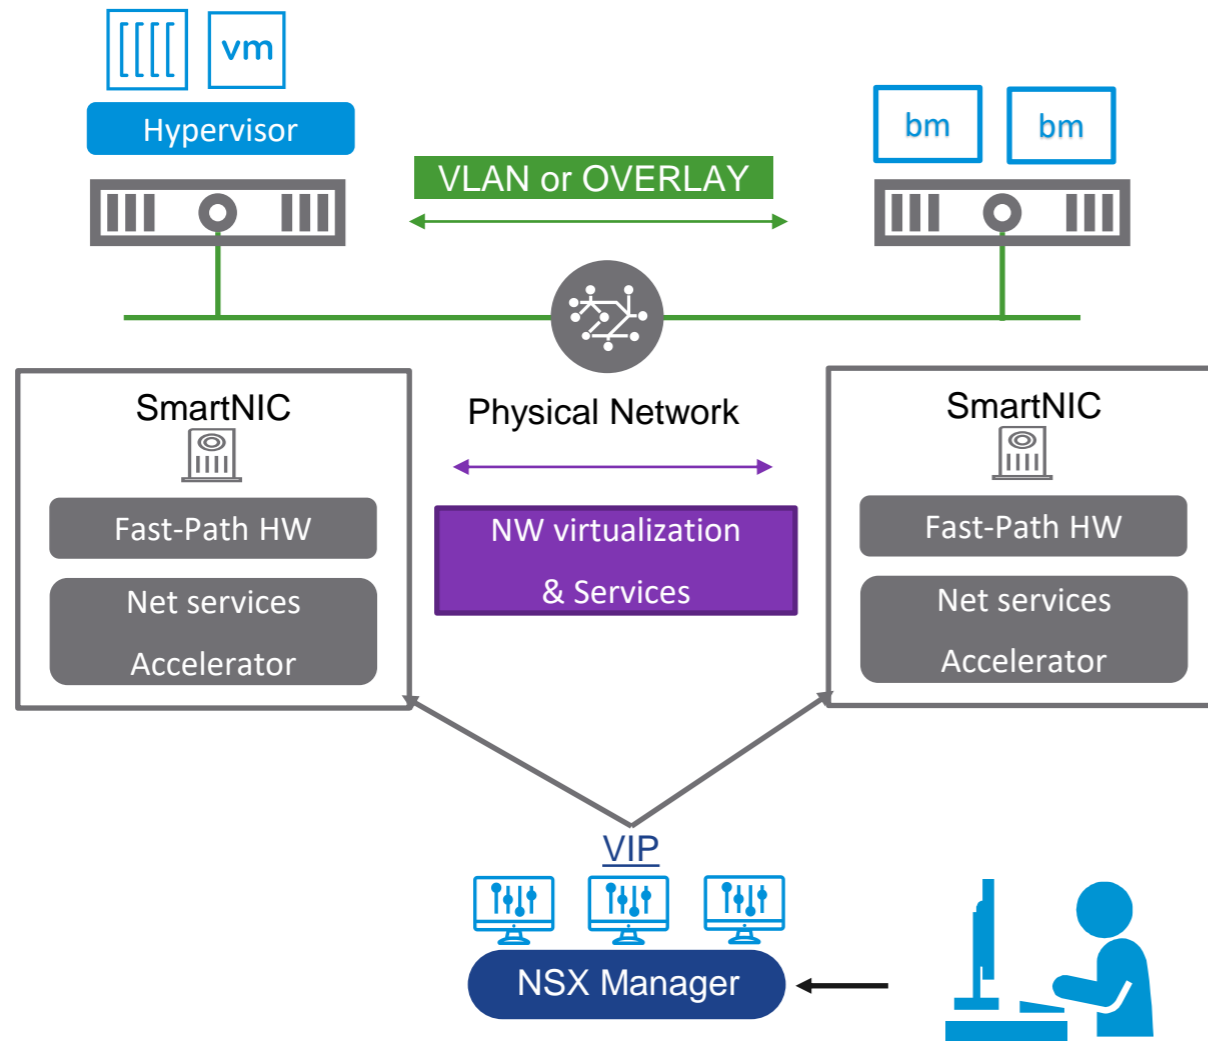
... and evolving customer needs

Macro Trends	Pain Points	Customer Needs
 <p>Rise of heterogenous infrastructure architectures based on accelerators</p>	<p>Inconsistent operating models across distributed clouds</p>	<p>A uniform operating model that unifies workload management across workloads</p>
 <p>Increasing demand for infrastructure services from higher volume of modern applications</p>	<p>Bigger portion of CPU capacity consumed by infrastructure services, leaving fewer cycles for workloads</p>	<p>Maximize utilization of compute resources for workloads without increasing infrastructure cost</p>
 <p>Expansion of the Enterprise perimeter to a distributed scale out architecture, across multi-clouds</p>	<p>Traditional CPU-centric security model exposes a single point of failure for workloads</p>	<p>Robust security model that provides better isolation between infrastructure and workload</p>



Network Acceleration and Offload

Offload Overlay and other network services - Free up compute resources

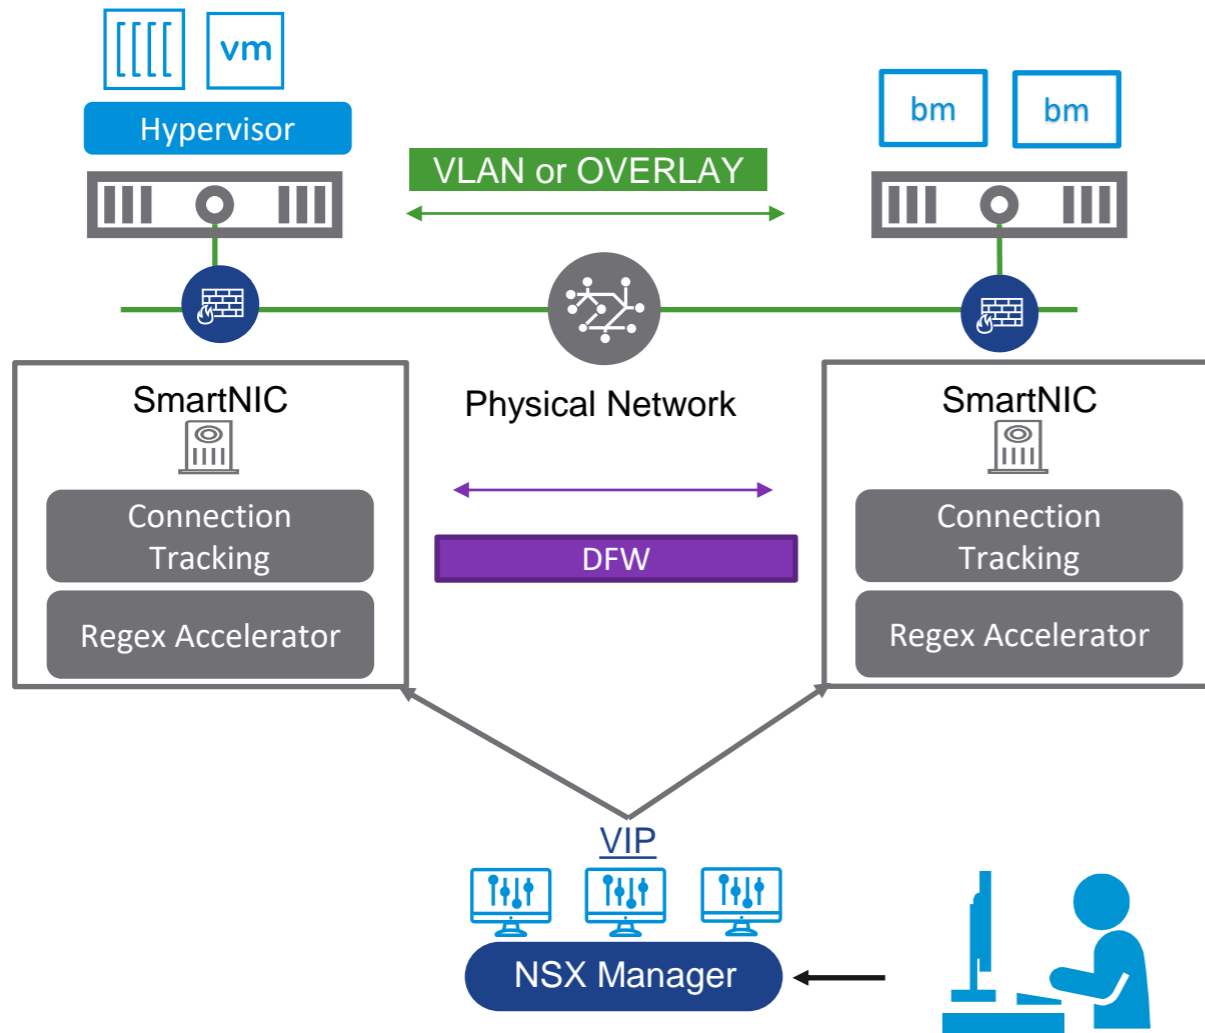


- Overlay Tunneling offload
 - VXLAN, Geneve, NVGRE
- Stateless offload
 - TCP/UDP/IP
 - LSO/LRO/Checksum/RSS/TSS/HDS
 - VLAN encaps/decaps
- VMDirectPath (UPTV2) with vMotion
 - SRIOV equivalent performance vMotion compatible.
- Advanced Network Services offload
 - LB, NAT etc.




Firewall and Advance Security – Airgapped

Leverage purpose-built hardware for DFW Services - Enforcement at NIC level.

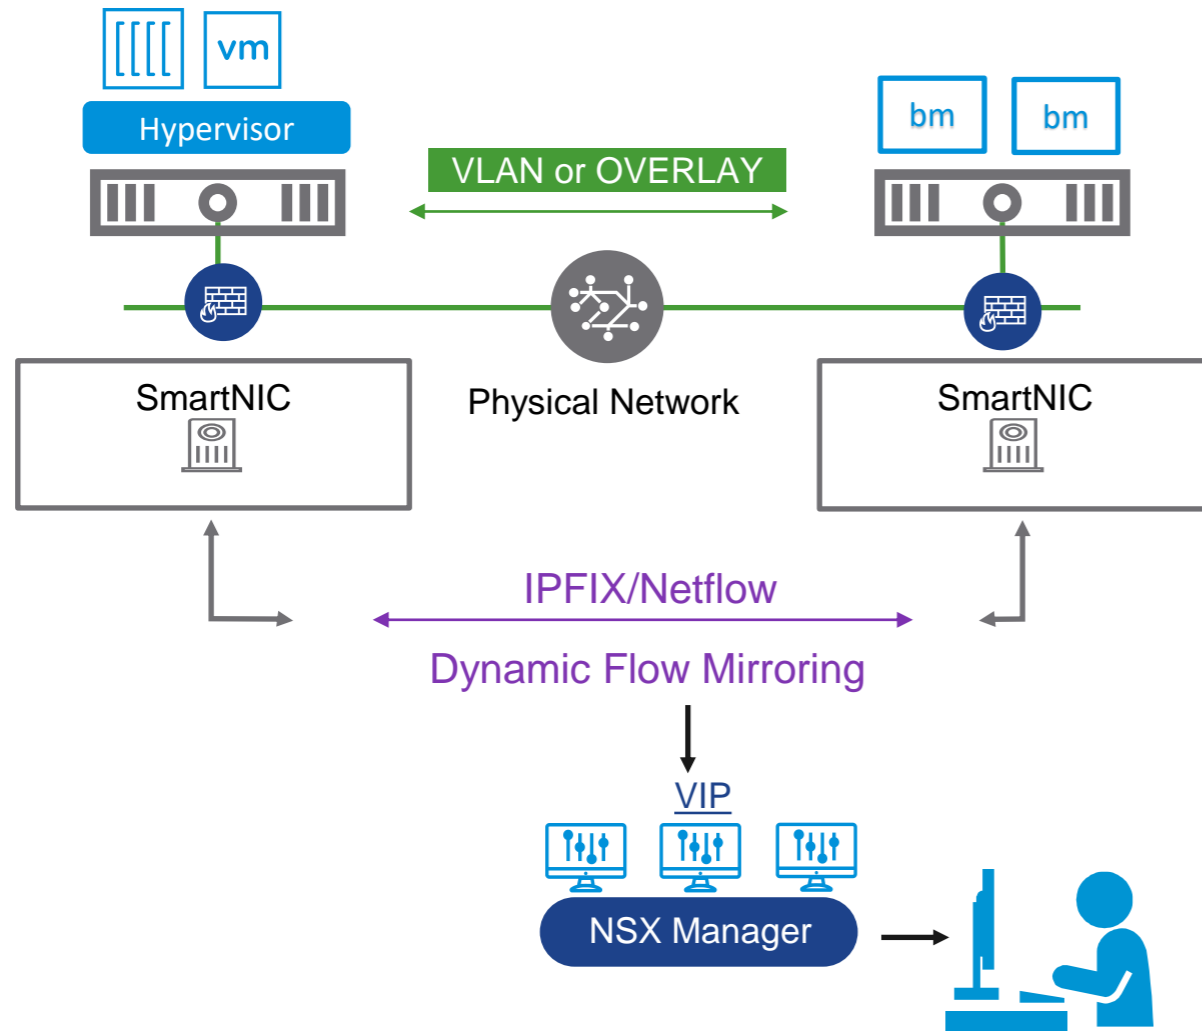


- L2 firewall, L4 firewall, L7 firewall.
- Spoofguard, Application layer gateway.
- High Perf, low-latency and low-jitter firewall.
- Intrusion detection and prevention.
- URL filtering, custom FQDN.
- Identify based firewall.
- Search and analyze IP traffic for malware detection and threat mitigation using Regex acceleration.
- Data in Motion Encryption.
- Time based rules.



Observability and Advanced Visibility without penalty

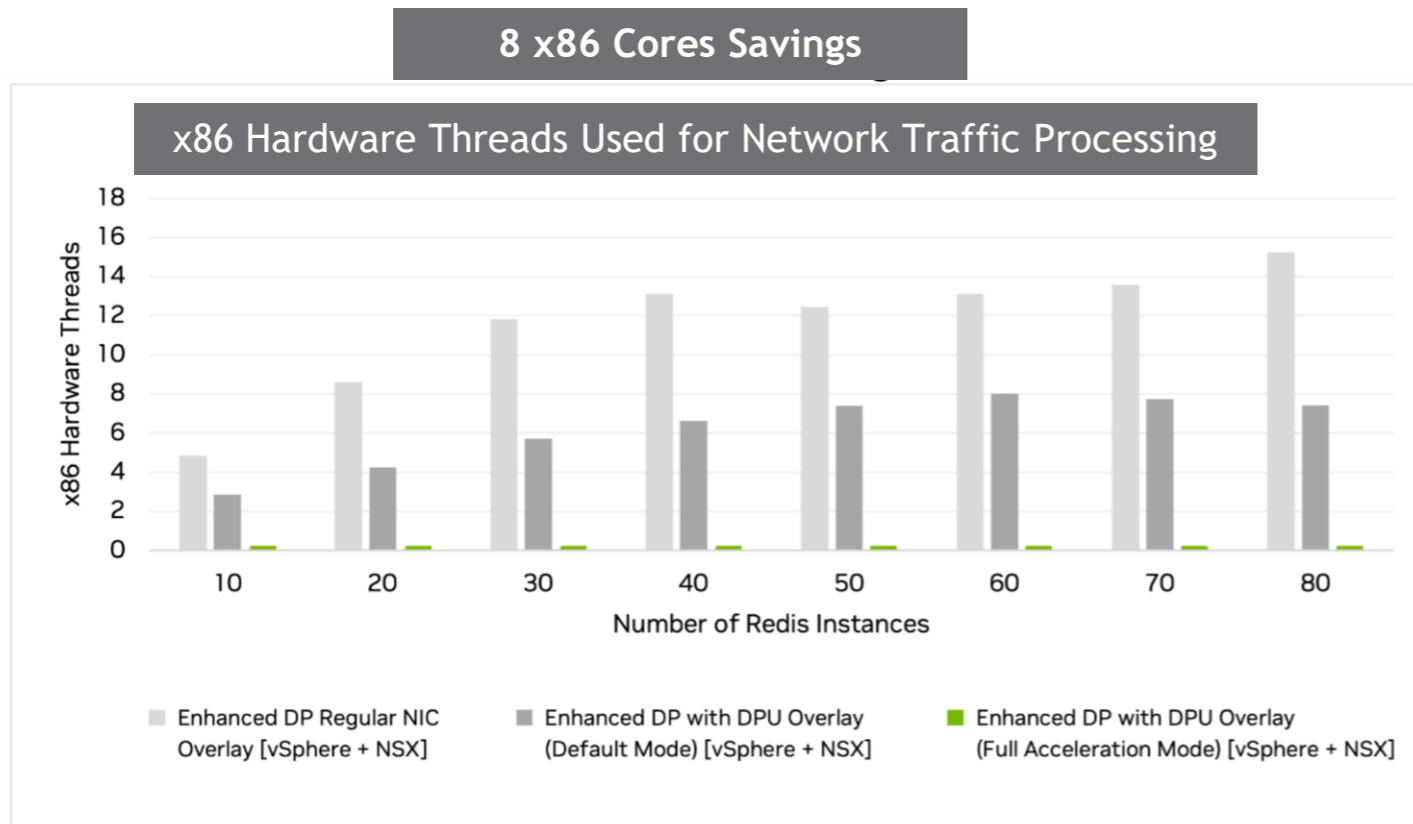
Observability for Troubleshooting and compliance - No need for SPAN or TAP network



- Troubleshooting
 - Topology view
 - Flow and packet level analysis
 - IPFIX support.
 - Packet capture, port mirroring (Offload).
- Logging & Compliance
 - Log analysis - Detailed syslog, Audit log.
 - Event co-relation
 - Audit and compliance check.
- Statistics, flows export.
- Operational thresholds



Accelerating Performance with DPU network offload

Redis In-Memory Key-Value Store testing



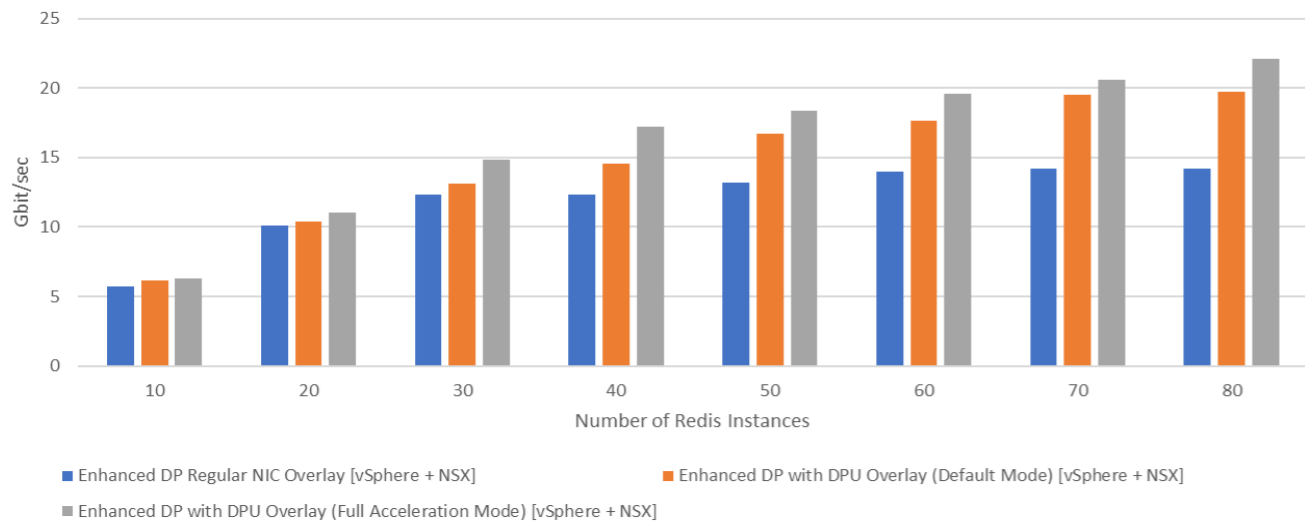
3 Dell PowerEdge R750s with one Intel® Xeon® Silver 4316 CPU running 40 physical cores (hyperthreading enabled) at 2.30GHz and 256 GB of main memory and one Nvidia Bluefield-2 DPU 25 GbE

<https://resources.nvidia.com/en-us-accelerated-networking-resource-library/nvidia-vmware-redis>

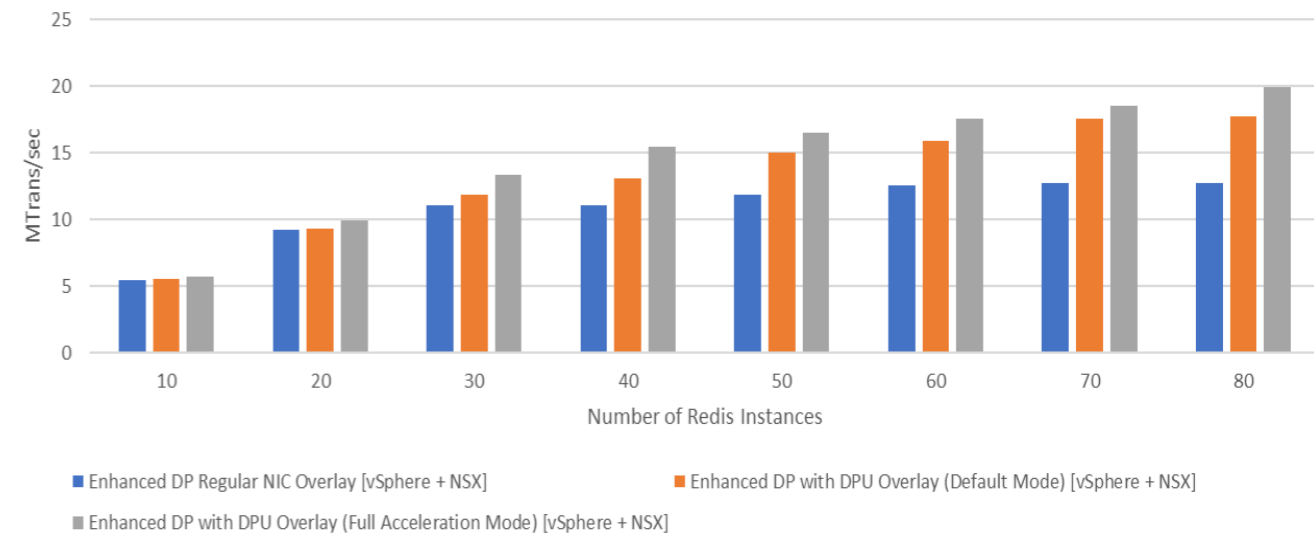


Bandwidth, Transaction rate and latency (Nvidia Bluefield 2 25G)

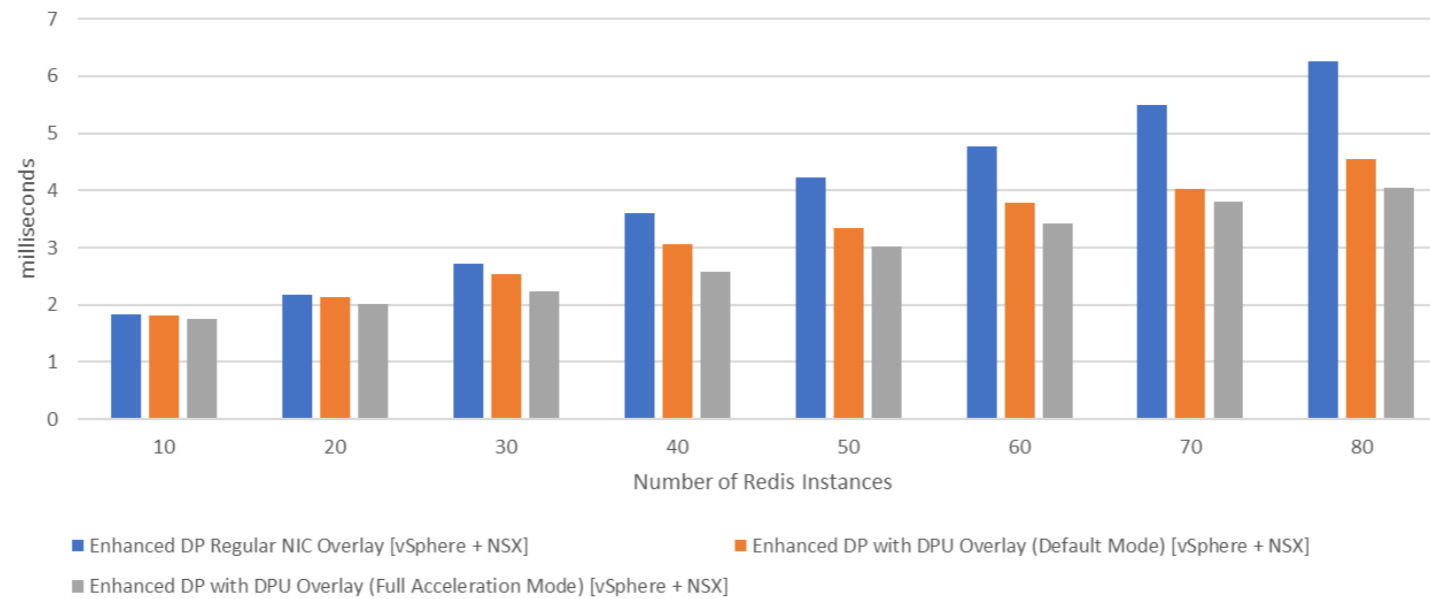
Redis Bandwidth



Millions of Redis Transactions/sec

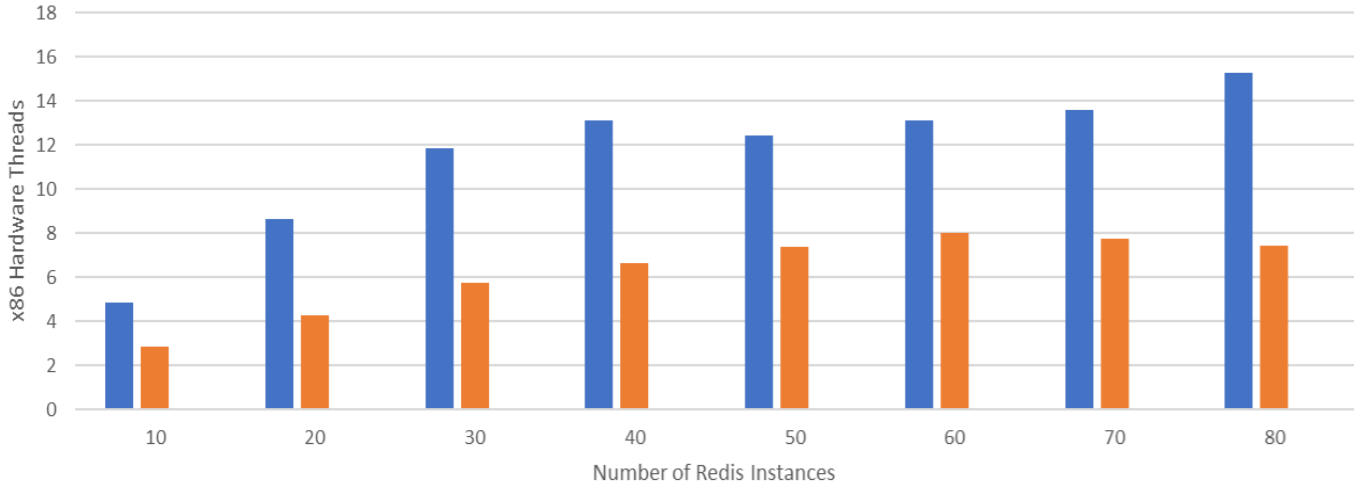


Average Redis Latency



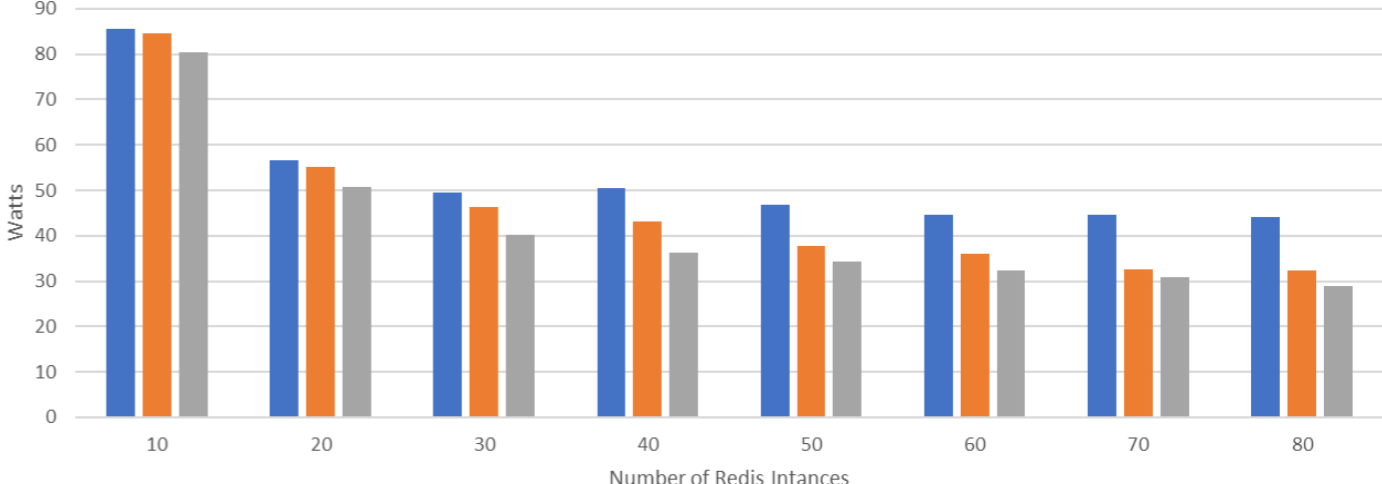
Core and Power Savings

x86 CPUs Used For Network Traffic Processing



■ Enhanced DP Regular NIC Overlay [vSphere + NSX]
 ■ Enhanced DP with DPU Overlay (Default Mode) [vSphere + NSX]
 ■ Enhanced DP with DPU Overlay (Full Acceleration Mode) [vSphere + NSX]

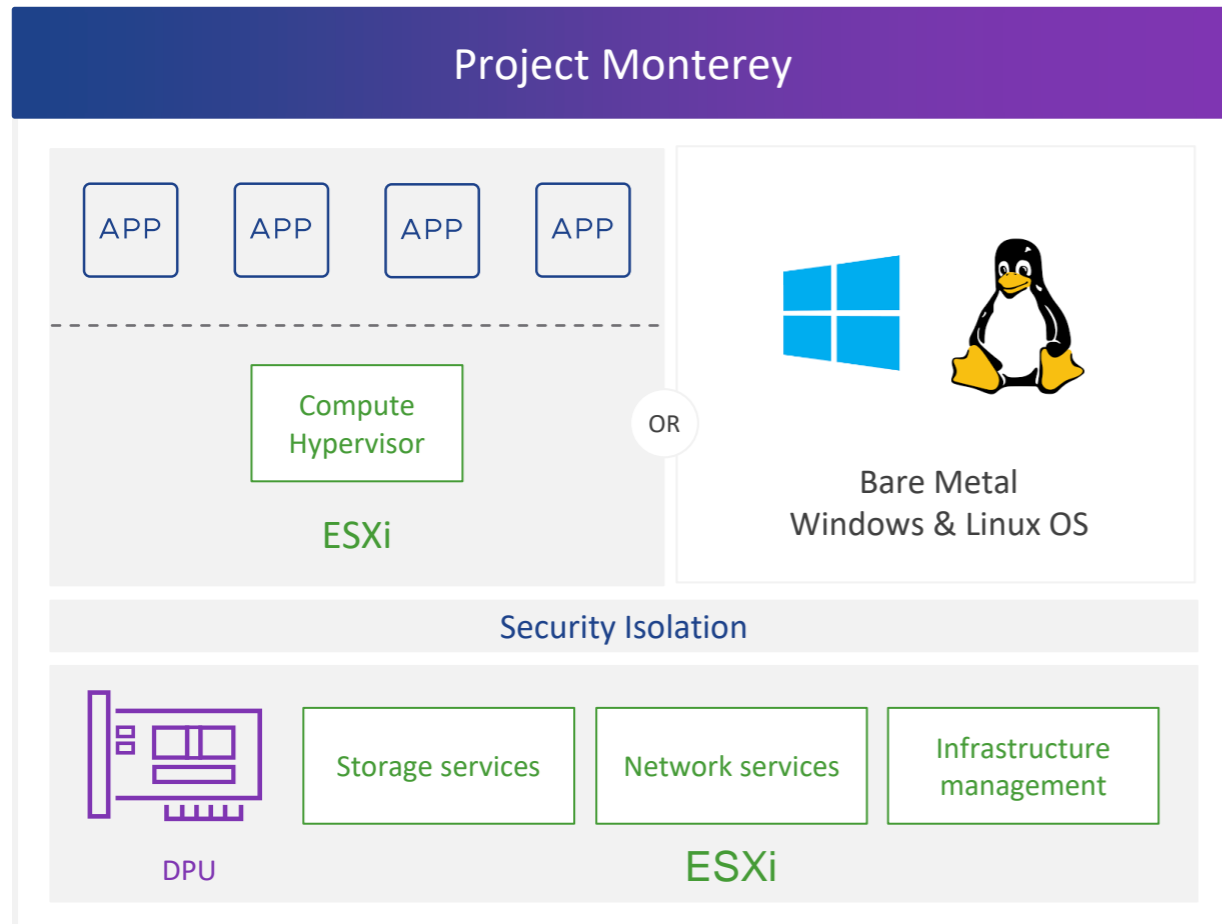
Watts per million transactions/sec



■ Enhanced DP Regular NIC Overlay [vSphere + NSX]
 ■ Enhanced DP with DPU Overlay (Default Mode) [vSphere + NSX]
 ■ Enhanced DP with DPU Overlay (Full Acceleration Mode) [vSphere + NSX]

vSphere on DPUs (Project Monterey)

Securing and accelerating modern and traditional workloads on composable hardware



Single, secure, operating model across workload types

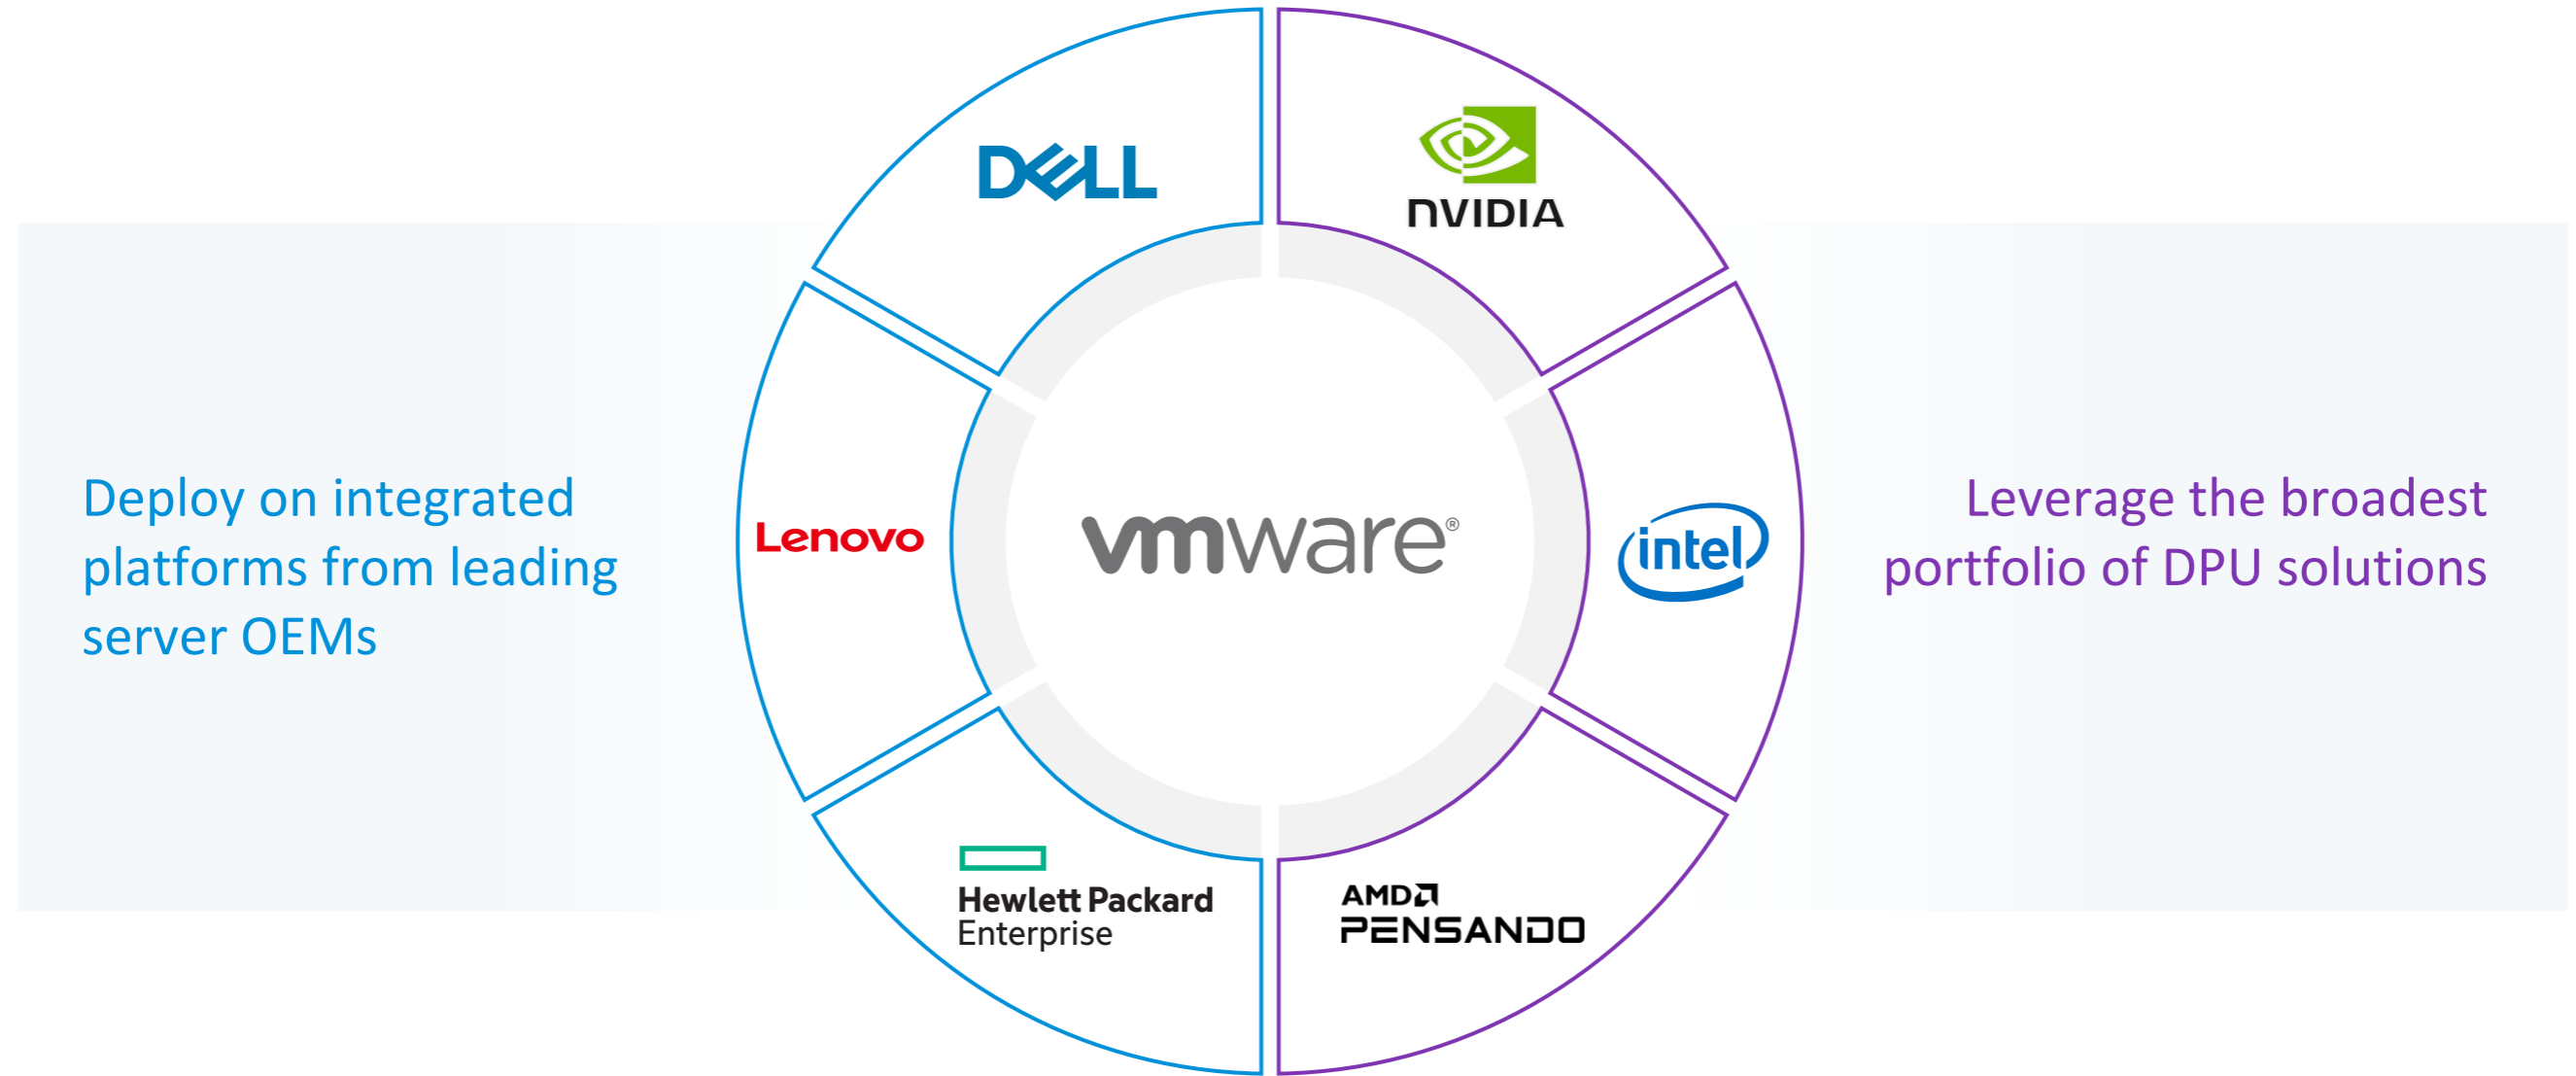


Isolation of workload domain from the infrastructure domain



Offload infrastructure service functions to DPUs

vSphere 8: Building a Flexible and Open Ecosystem of DPU Solutions



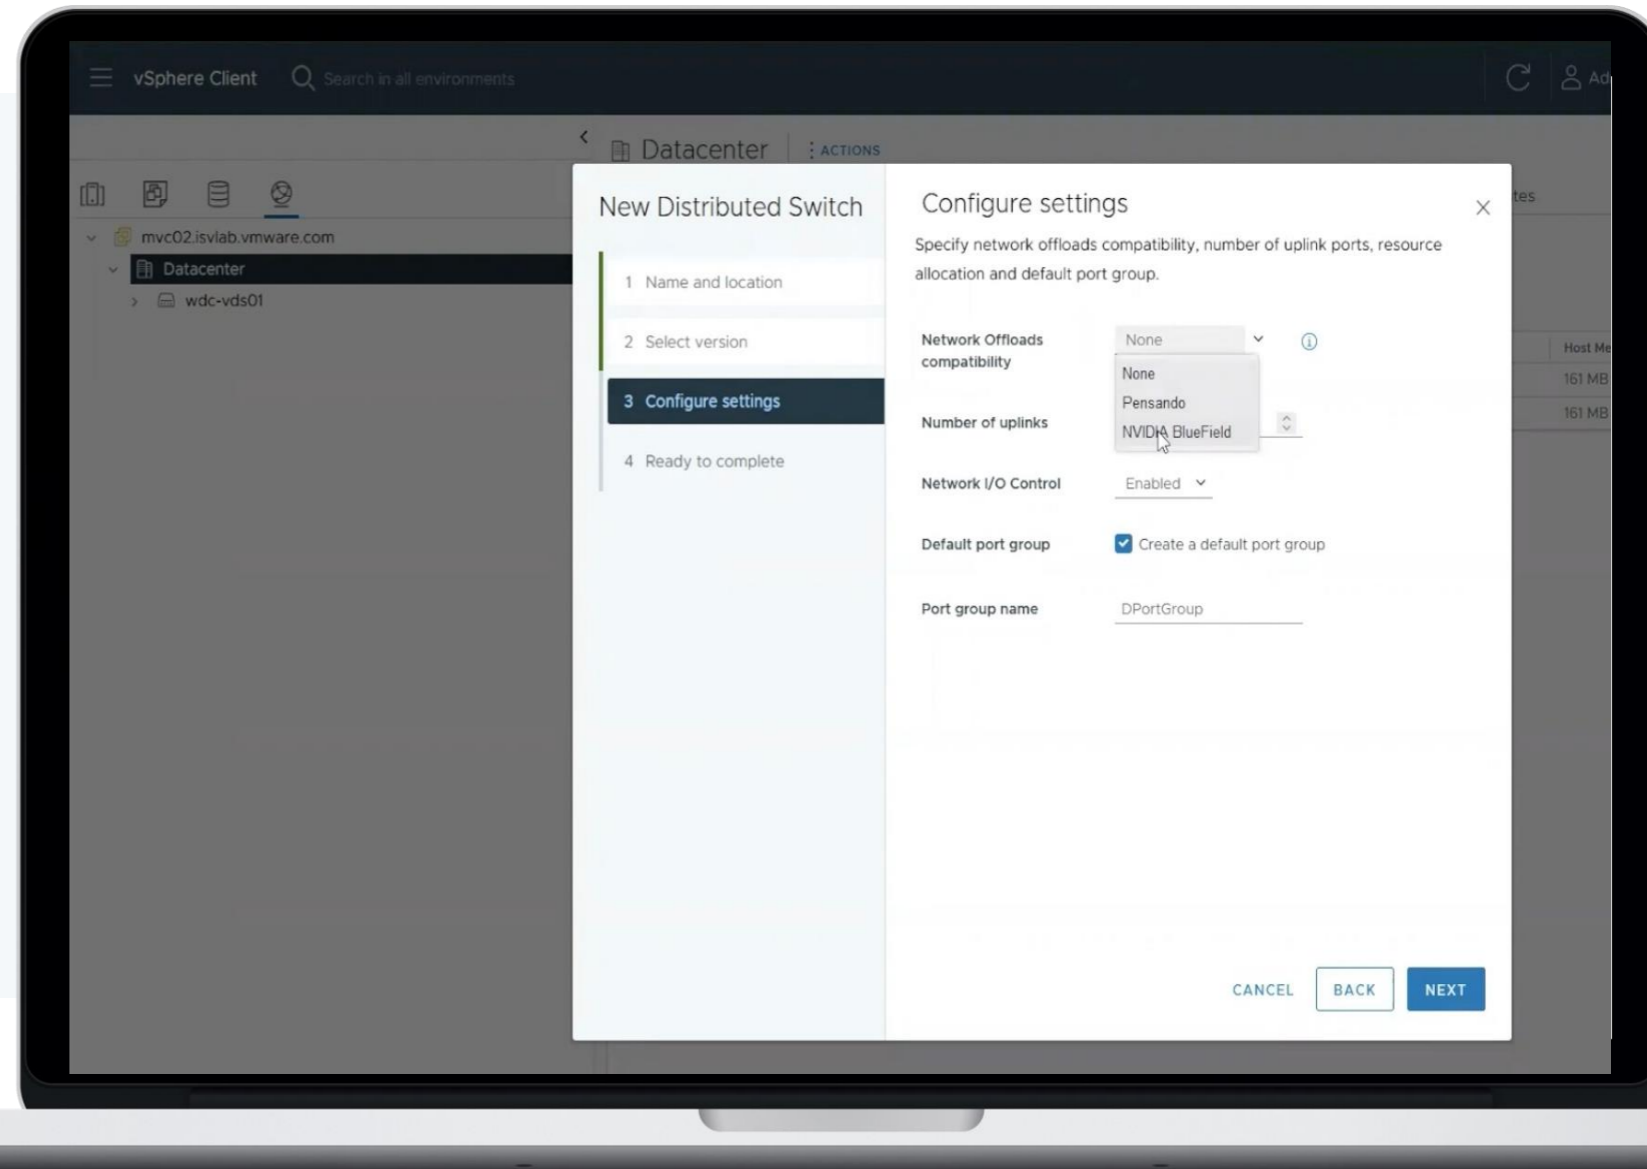
Meet the Throughput and Latency Needs of Modern Workloads

Accelerate infrastructure functions with vSphere Distributed Services Engine



Improve Infrastructure Performance

- Accelerate vSphere Distributed Switch, NSX Networking⁽¹⁾ & Observability⁽²⁾
- Achieve higher workload consolidation per host
- Get passthrough performance while preserving VM mobility through DRS and vMotion
- Scale your security operations⁽³⁾ at line rate without the need for software agents



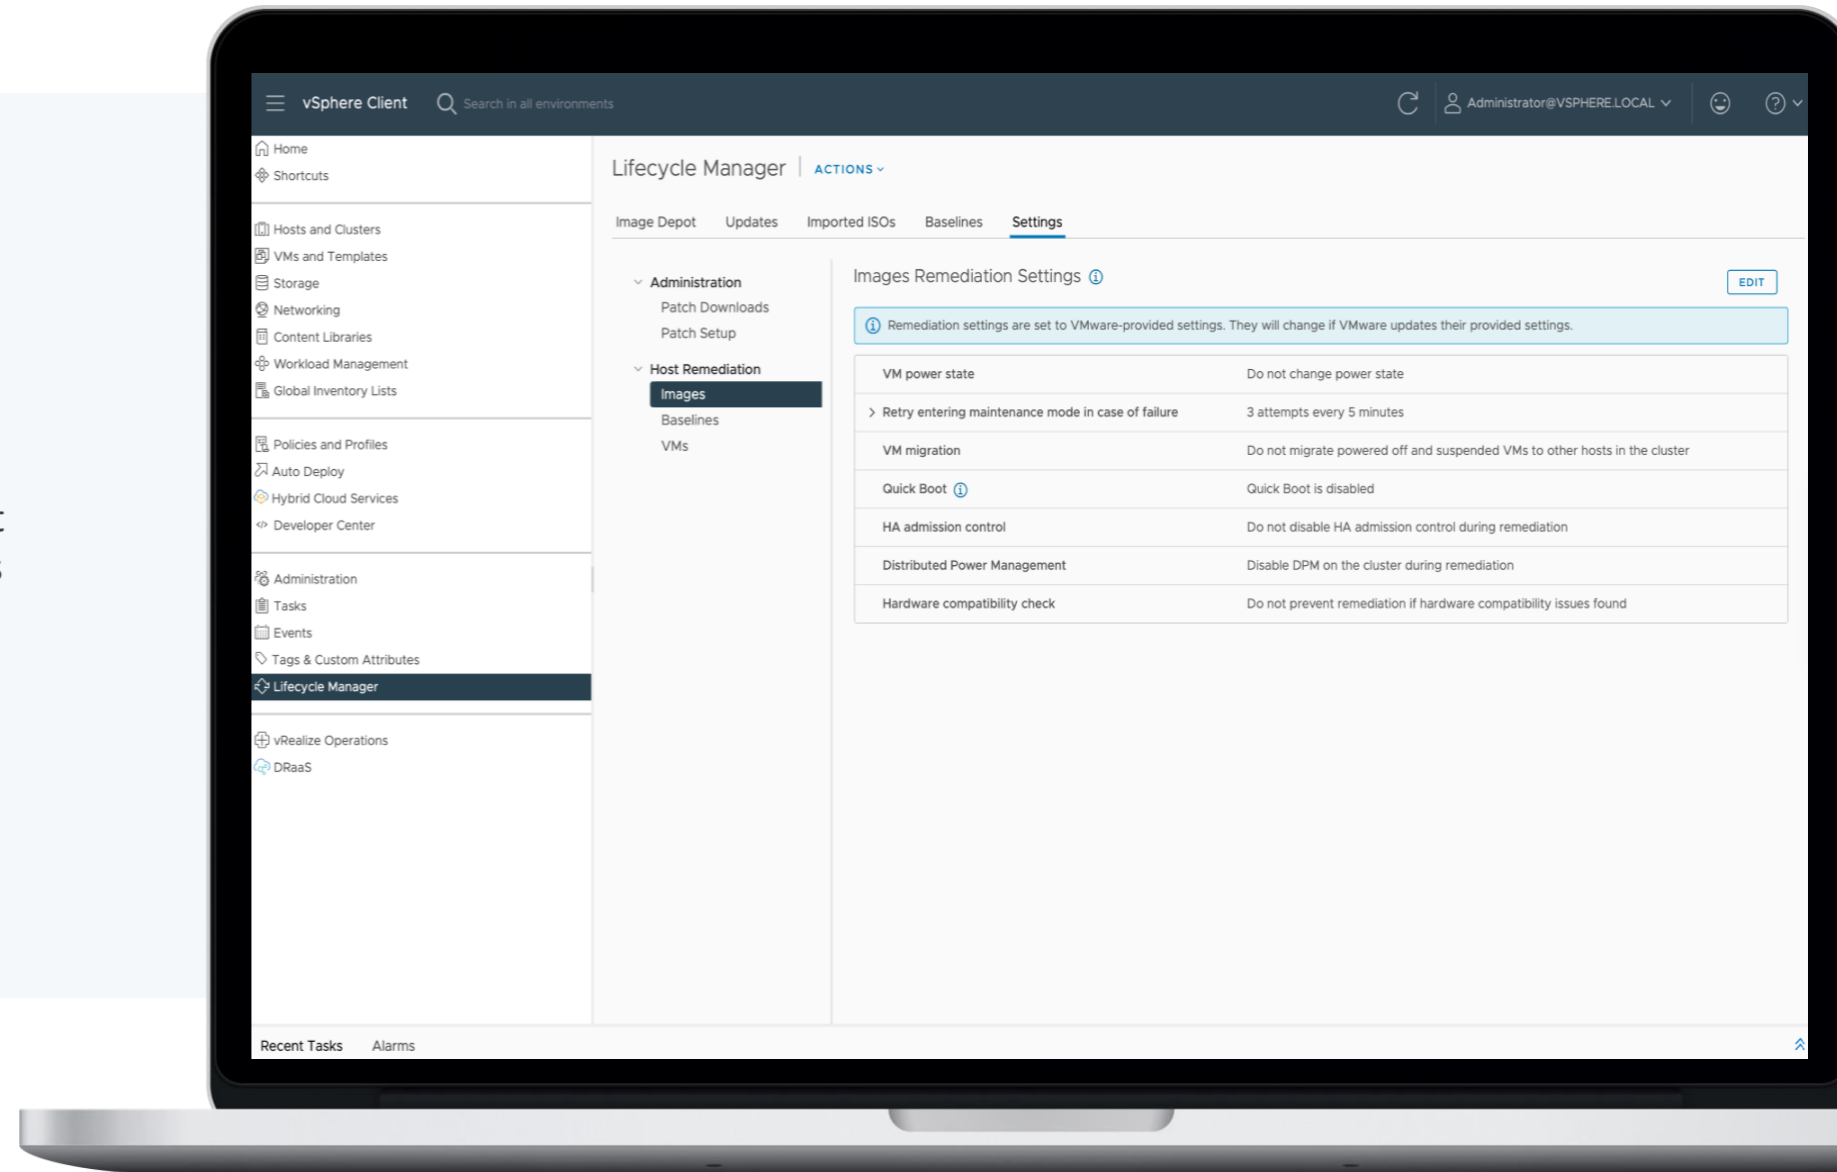
Reduce Operational Overhead

Manage DPU lifecycle with vSphere distributed services engine



Simplify DPU Lifecycle Management

- Reduce operational overhead of DPU lifecycle management with integrated vSphere workflows
- Use proven vCenter interfaces to get DPU alerts and performance metrics
- Proactively monitor, identify and mitigate infrastructure bottlenecks



Discussion Topic

What can we do to accelerate adoption?

What holds us back?



June 13-15, 2023
DoubleTree by Hilton San Jose
SmartNICsSummit.com



Thank You