




NVIDIA

NVIDIA DATA PROCESSING UNIT ETHERNET SWITCH


Sebastian Kalcher

September 2021



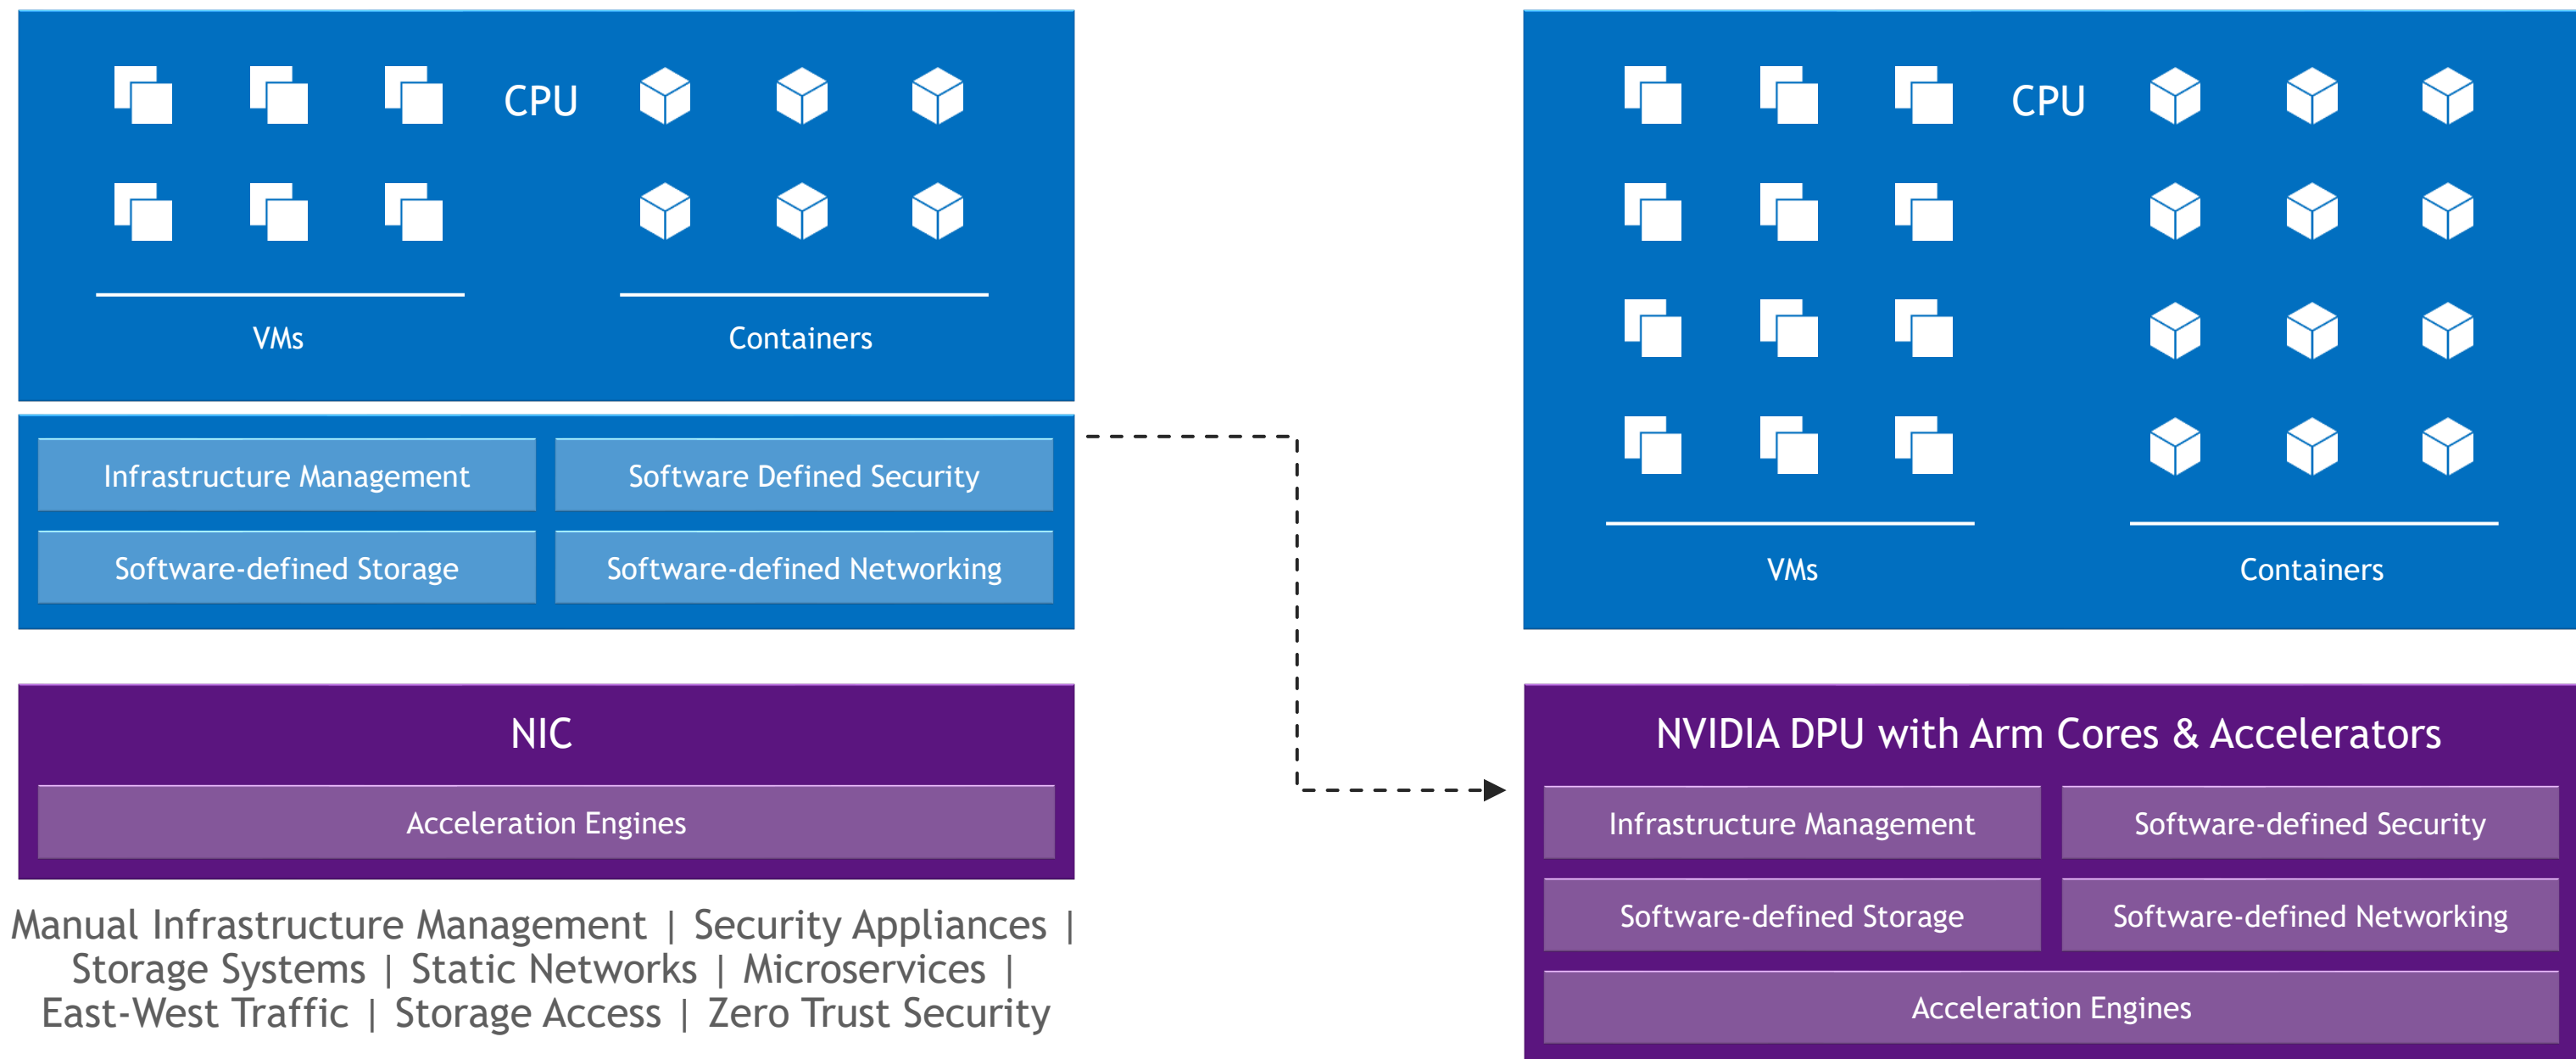
REINVENTING THE DATA CENTER

The Data Center is the New Unit of Computing



INTRODUCING THE DATA PROCESSING UNIT

Software-Defined, Hardware-Accelerated Data Center Infrastructure-on-a-Chip



NVIDIA BLUEFIELD-2 DATA PROCESSING UNIT

Data Center Infrastructure on a Chip

Up to 200Gb/s Ethernet and InfiniBand, PAM4/NRZ

ConnectX-6 Dx inside

8 Arm A72 CPUs subsystem - up to 2.75GHz

8MB L2 cache, 6MB L3 cache in 4 Tiles

Fully coherent low-latency interconnect


Integrated PCIe switch, 16x Gen4.0

PCIe root complex or end point modes

Single DDR4 channel

1GbE Out-of-Band management port

Accelerated security, storage, networking



NVIDIA DOCA

Enabling Broad DPU Partner Ecosystem

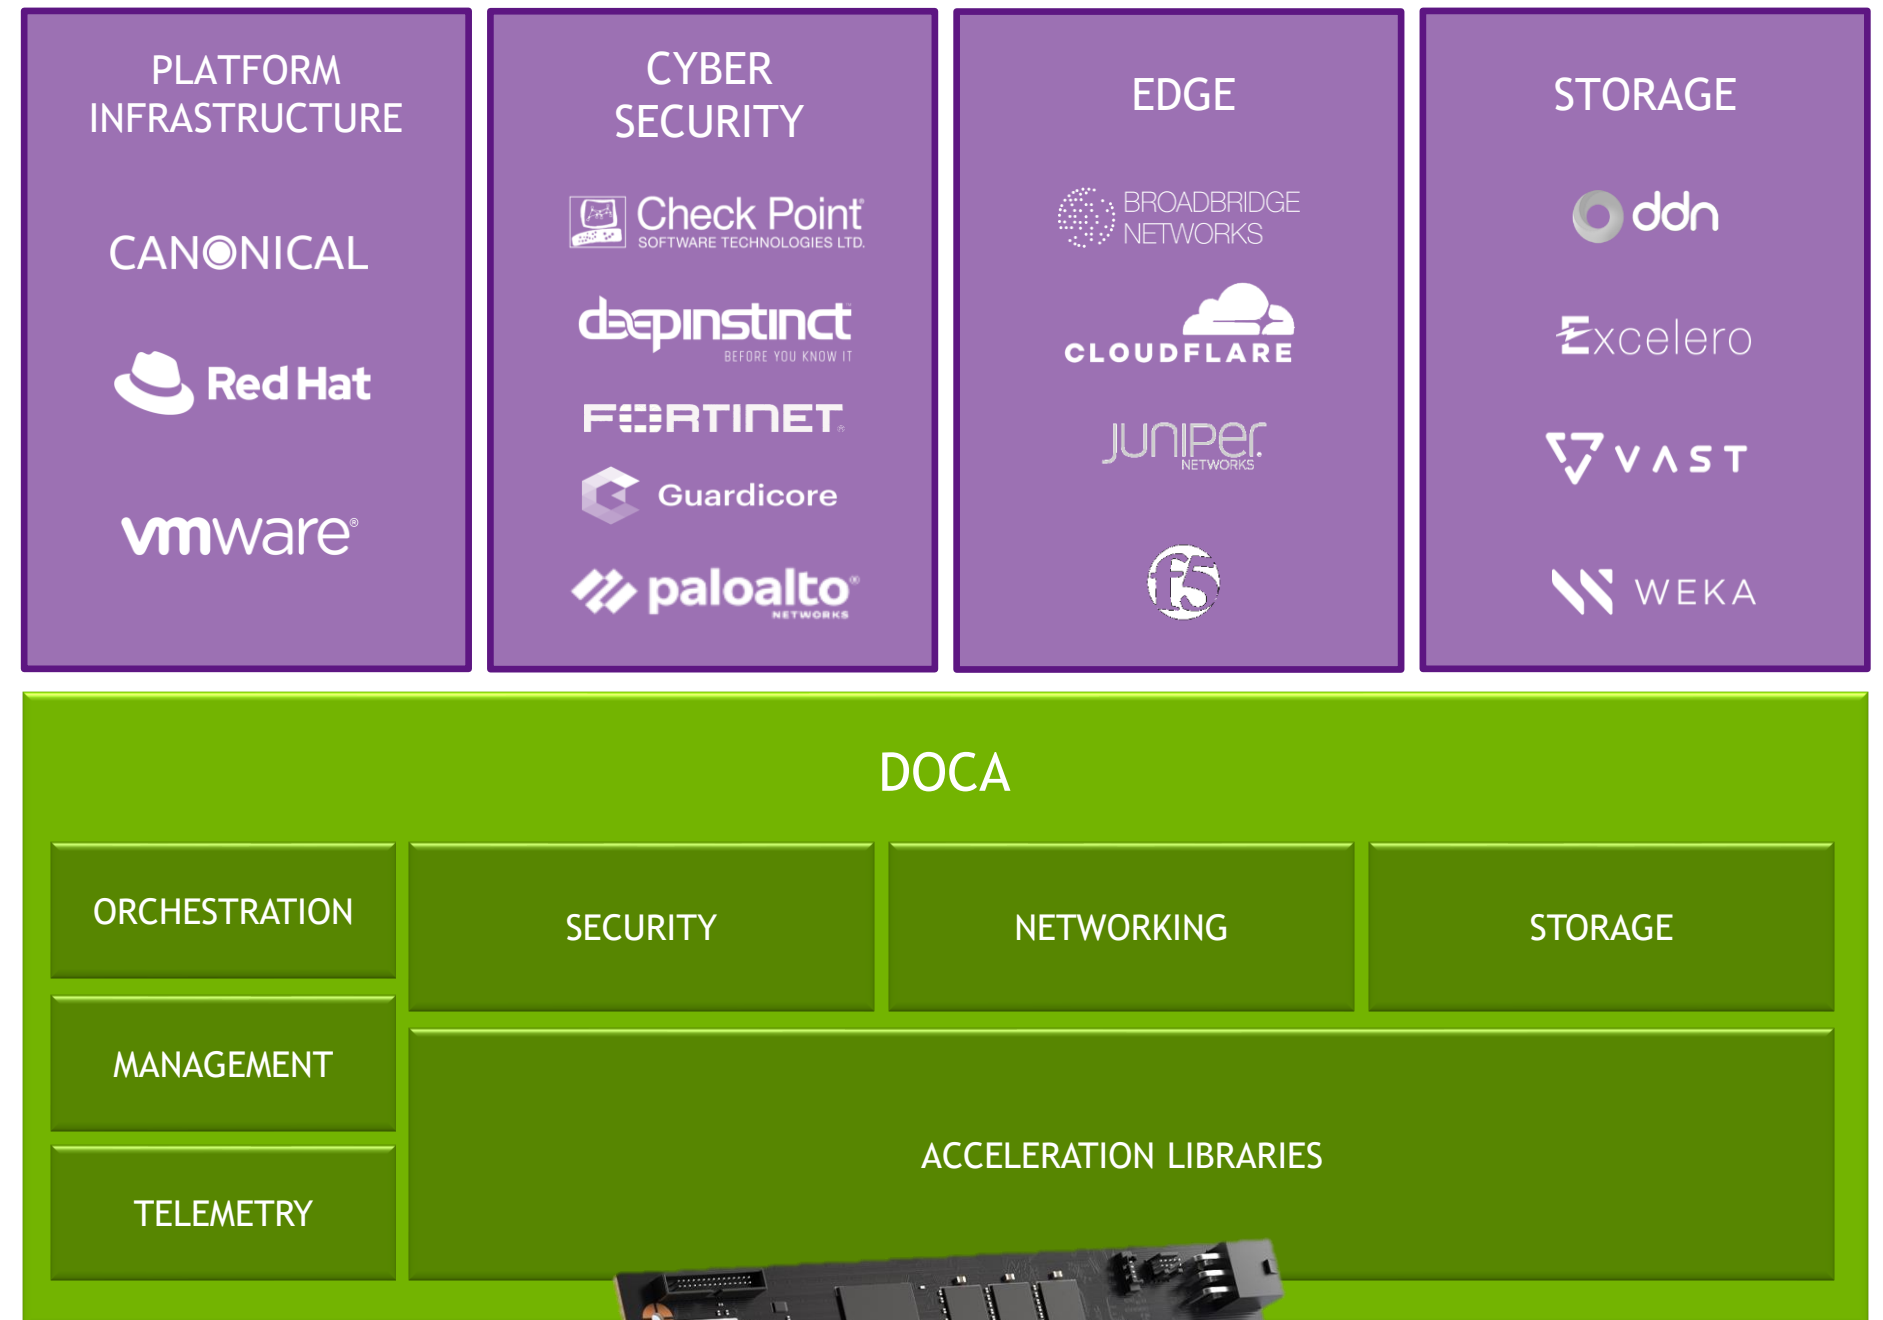
Software application framework for BlueField DPUs

DOCA is for DPUs what CUDA is for GPUs

Protects developer investment for future DPUs

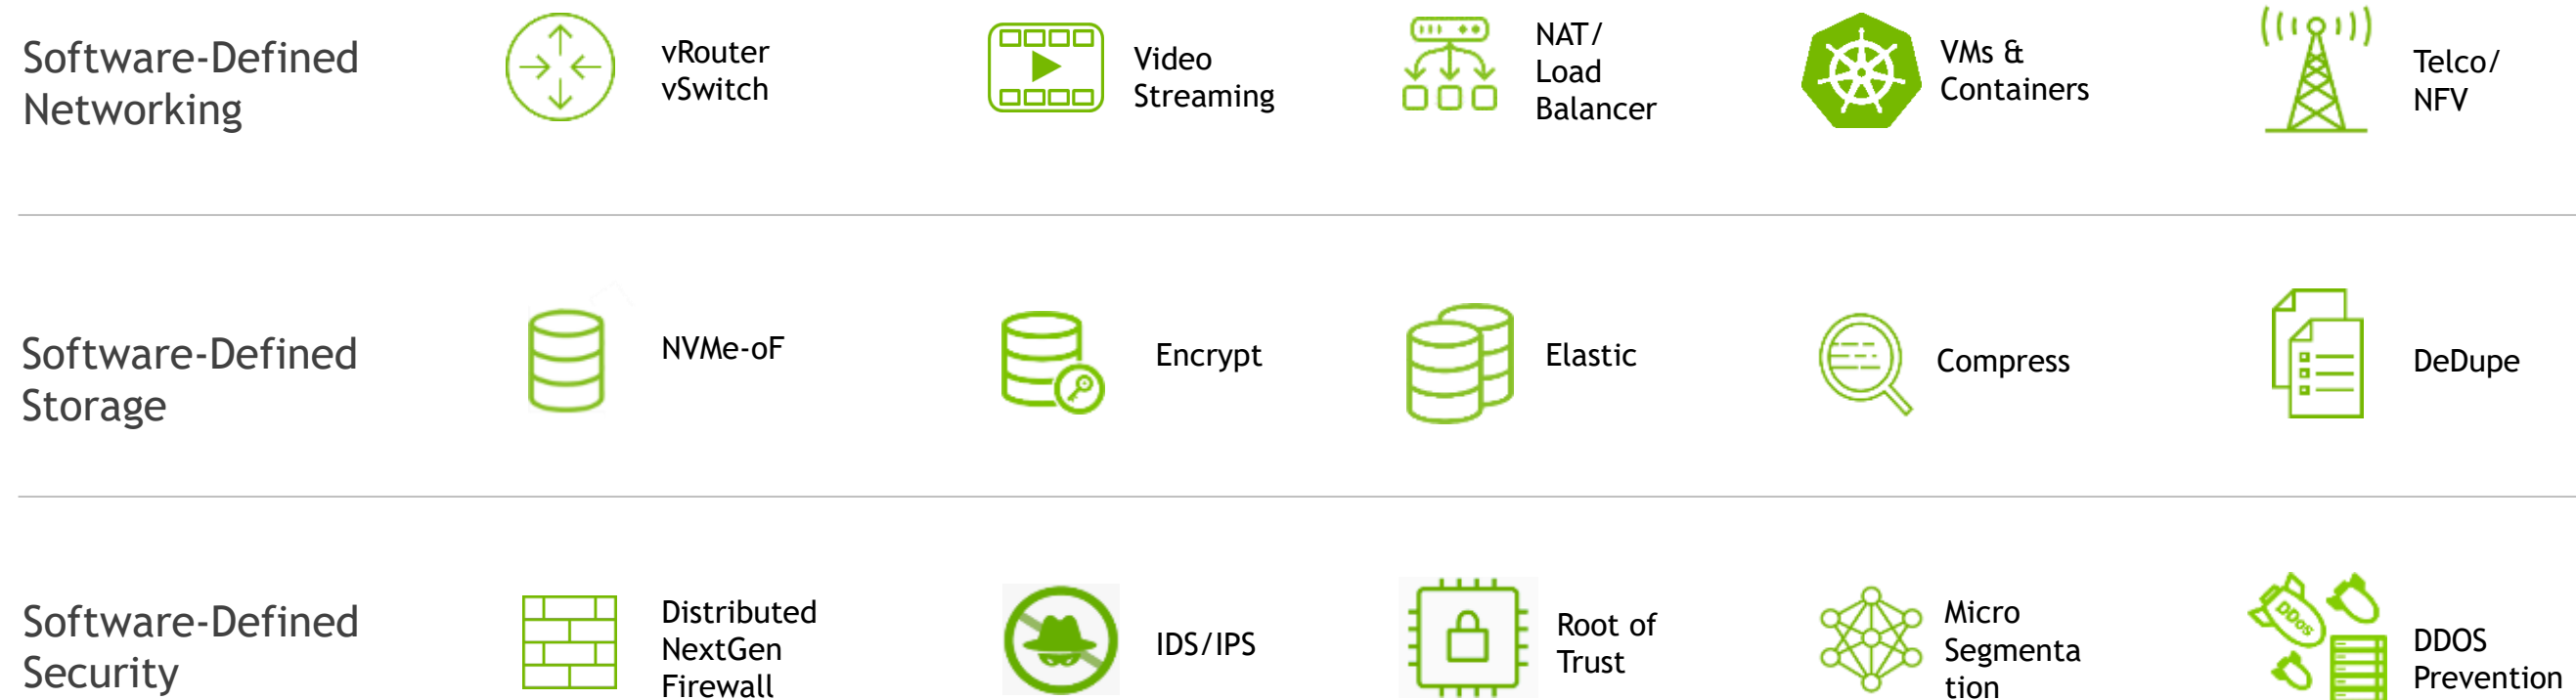
Certified reference applications, APIs & partner solutions

Rich partner ecosystem across industries and workloads



DOCA VISION

Software-Defined, Hardware-Accelerated Infrastructure



DOCA (Open and Programmable API Framework)
Easy, Flexible Programming of Infrastructure/ Acceleration and Security

NVIDIA Bluefield-2 DPU

BLUEFIELD REDEFINES BARE-METAL CLOUDS



Software-Defined Networking

Brings full-featured SDN capabilities to bare-metal clouds
Full control and orchestration through OVS/OVN
No installation of network driver in bare-metal host



Elastic Composable Storage

Virtualized storage flexibility w/ local storage performance
Dynamic allocate cloud storage & streamline data protection
Operating System is agnostic; no driver is required



Zero Trust Security

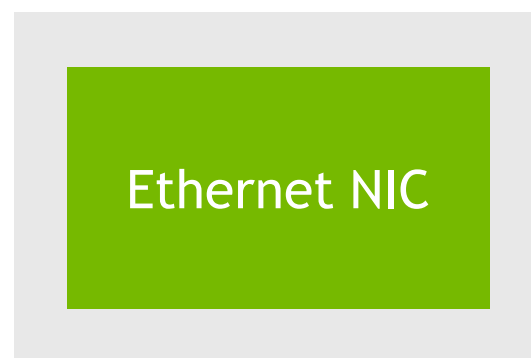
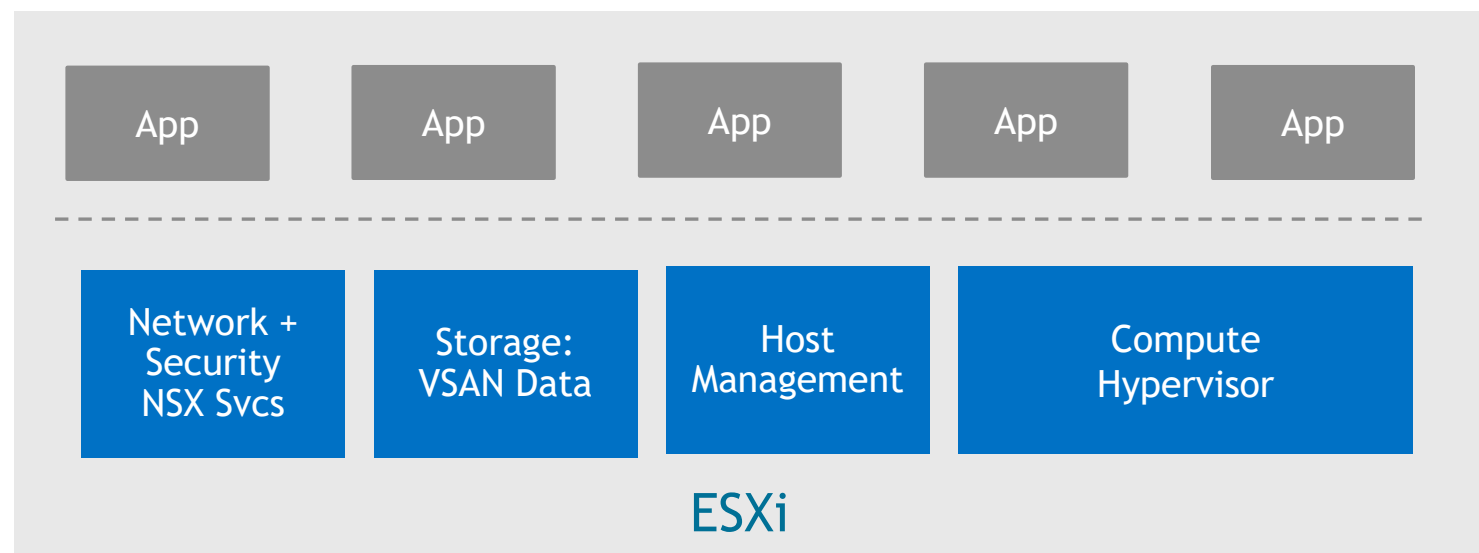
Isolation of security control from the tenant's workload
Enhancing data-center security
No local agents on servers or any performance impact



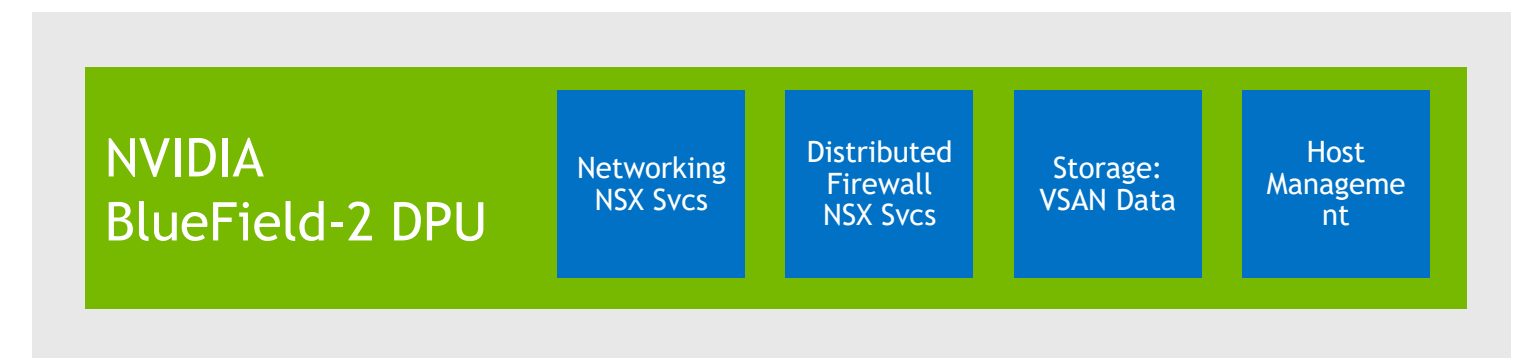

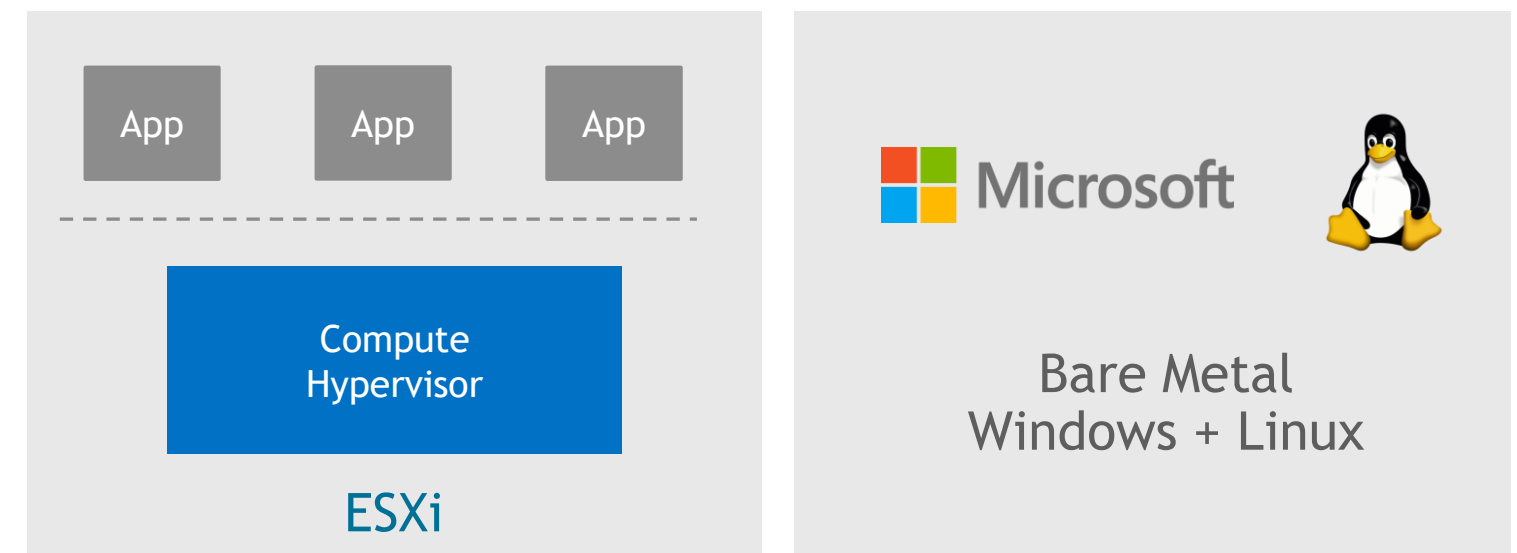
INTRODUCING PROJECT MONTEREY

Run Modern Workloads Efficiently Over New Composable, Disaggregated Infrastructure

Today's Environment



Project Monterey



NVIDIA & VMWARE ENABLE HYBRID CLOUD ARCHITECTURE

Run Modern Workloads Efficiently Over New Composable, Disaggregated Infrastructure



Peak Performance

Maximize network bandwidth while saving x86 cycles for top application performance




Unified, Consistent Operations

Support bare metal servers, simplify lifecycle management, and reduce TCO



Zero-Trust Security Model

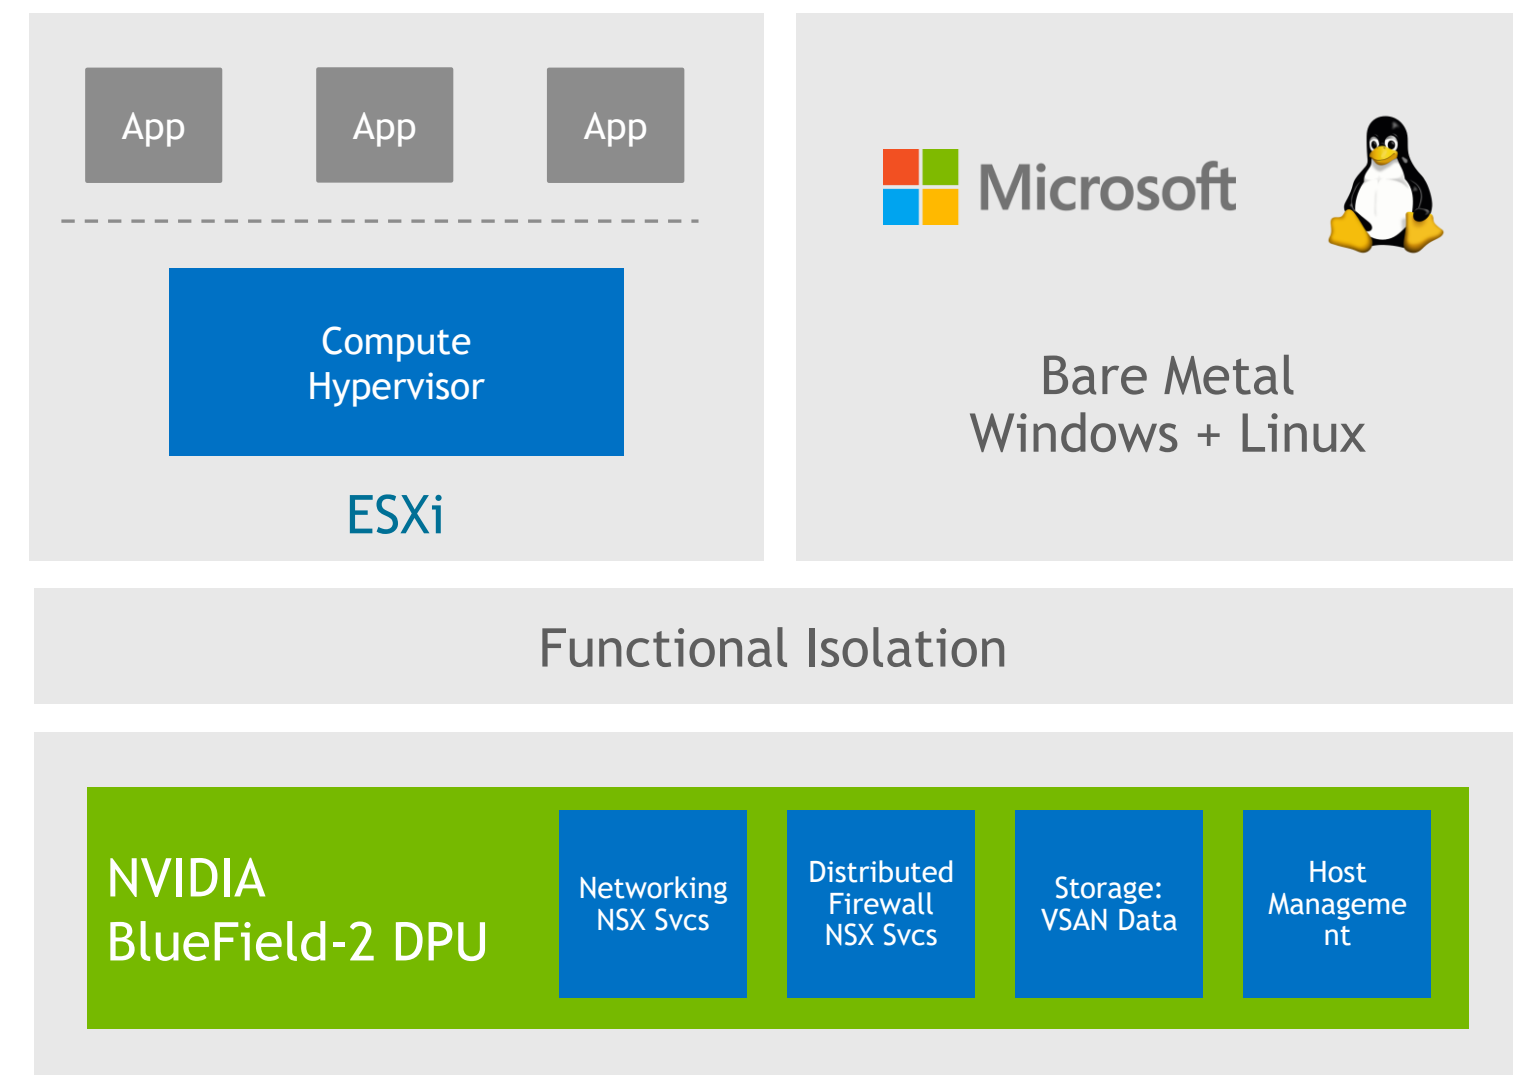
Ensure comprehensive application security without compromising performance



Application-Driven Infrastructure

Dynamically assemble right-sized infrastructure at runtime based on unique application needs

Project Monterey



BLUEFIELD ENABLES CLOUD-NATIVE SUPERCOMPUTING

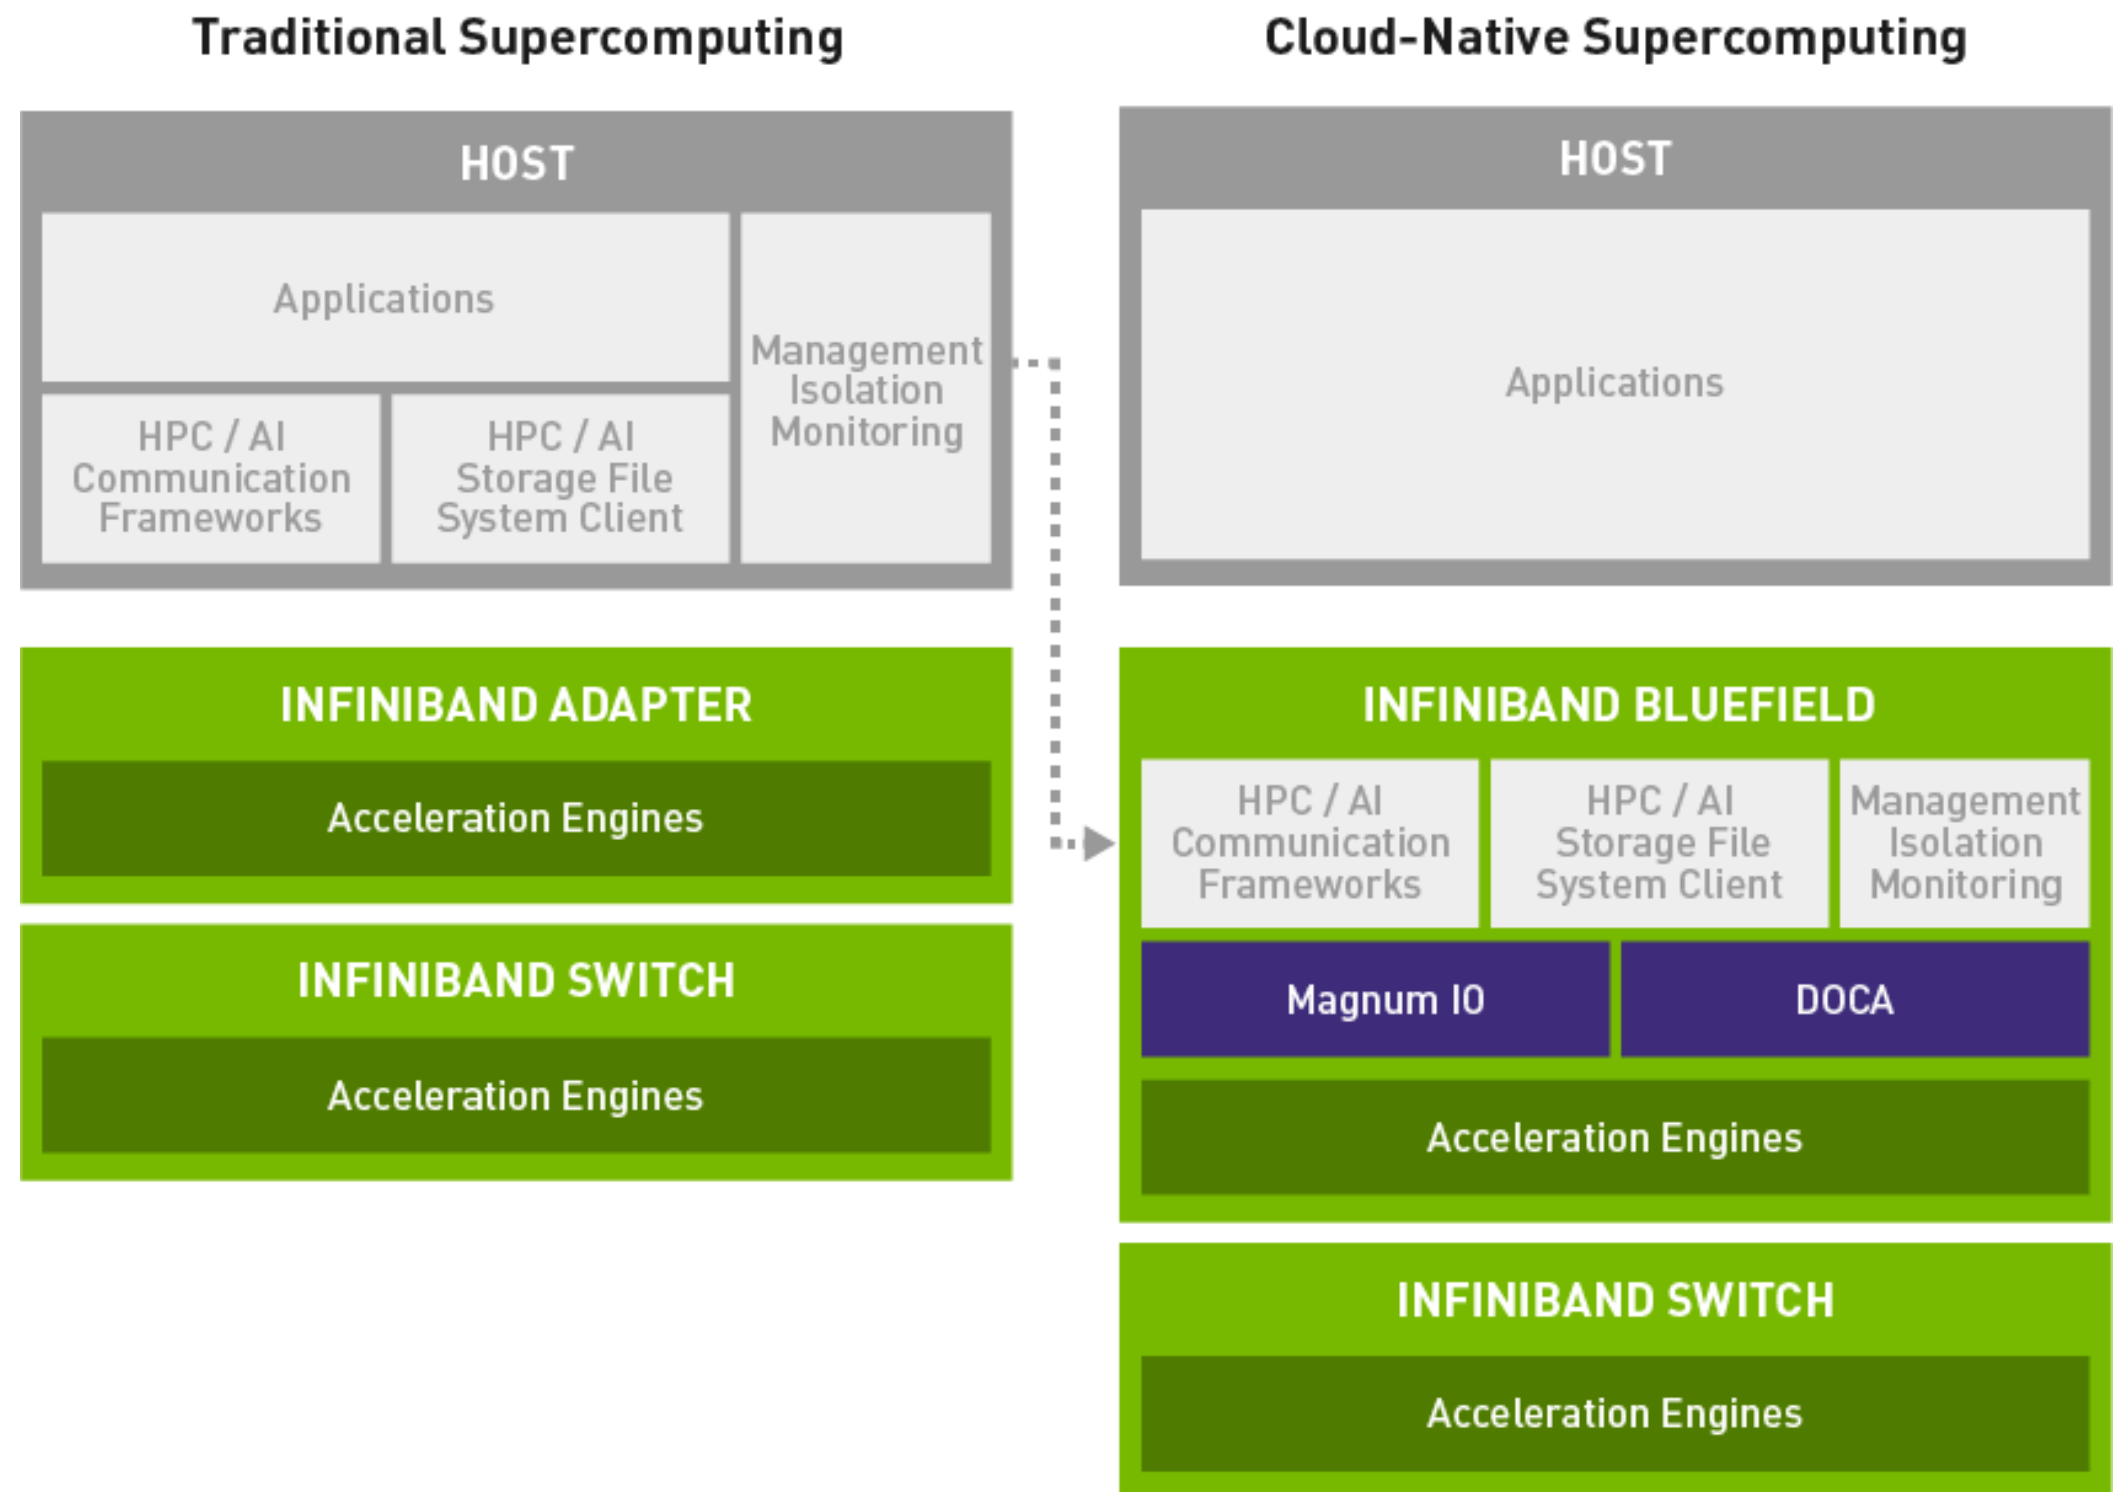
Multi-Tenancy with Zero-Trust Security

Collective offload with UCC accelerator


Smart MPI progression

User-defined algorithms

1.4X higher application performance



BLUEFIELD SUPERCHARGES CYBER SECURITY



Platform Security

Protect compute platforms from hardware-level attacks by enabling secure and measured boot of the OS with hardware root-of-trust



Crypto Acceleration

Secure data center communications and web application delivery with hardware-accelerated, inline IPsec and TLS data encryption



Security Policy Enforcement

Enforce distributed security policy in hardware by enabling advanced switching and packet processing (ASAP2) technology



BLUEFIELD SUPERCHARGES CYBER SECURITY



Deep Packet Inspection

Detect malicious activity by performing deep packet inspection on incoming traffic with built-in pattern matching acceleration



Isolated Security Control Plane

Deploy security applications on trusted environment, isolated from the host system for protecting against lateral movement and reducing the attack surface



BLUEFIELD ACCELERATES 5G TRANSFORMATION



Uncompromised SDN Acceleration

Advanced switching and packet processing (ASAP2) technology enables best of both worlds, software-defined networking with hardware-accelerated performance



Time-Triggered Transport Technology for Telco (5T) for 5G

Accurate timing synchronization across front-haul and mid-haul networks, enabling software-defined, GPU-accelerated RAN systems



Cloud-Native Telecoms

As telecom applications adopt cloud-native, micro-services architectures, NVIDIA BlueField DPUs bring high-performance solutions for Kubernetes environments



ANNOUNCING NVIDIA BLUEFIELD-3 DPU

First 400Gb/s Data Processing Unit

22B transistors


400Gb/s Ethernet & InfiniBand
Connectivity

400Gb/s Crypto Acceleration

16 ARM CPU Cores

300 Equivalent x86 Cores

18M IOP/s Elastic Block Storage



BLUEFIELD-3 USE CASES

Unprecedented Innovation for Modern Data Centers




Cloud Computing

Bare-Metal | Virtualized | Containerized
Private | Public | Hybrid Cloud




Cyber Security

Distributed Security | NGFW |
Micro-segmentation



HPC & AI

Cloud-Native Supercomputing |
Accelerated DLRM



Telco & Edge

Telco Cloud | CloudRAN |
Edge Compute



Data Storage

HCI | Elastic Block Storage |
Instance Storage



Media Streaming

Visual High Quality |
8K Video | CDN

INDUSTRY-LEADING PERFORMANCE

Massive Advancements, Built for Cloud Scale

	BlueField-2	BlueField-3
Bandwidth	200Gb/s	400Gb/s
DPDK Max msg Rate	215Mpps	250Mpps
RDMA max msg rate	215Mpps	330Mpps
Compute	SPECINT: 70	SPECINT: 350
VirtIO Acceleration	40Mpps (*)	80Mpps (*)
IPsec Acceleration	100Gb/s	400Gb/s
TLS Acceleration	200Gb/s	400Gb/s
MACsec Acceleration	X	400Gb/s bi-dir
NVMe SNAP	5.4M IOPs @4K	12M-18MIOPS @4K

* Total packet rate for the sum of Tx and Rx



BROAD PARTNER ECOSYSTEM

Hybrid Cloud Compatibility | No Fork-Lift Upgrades | No Vendor Lock-In

CANONICAL



vmware®



FORTINET®

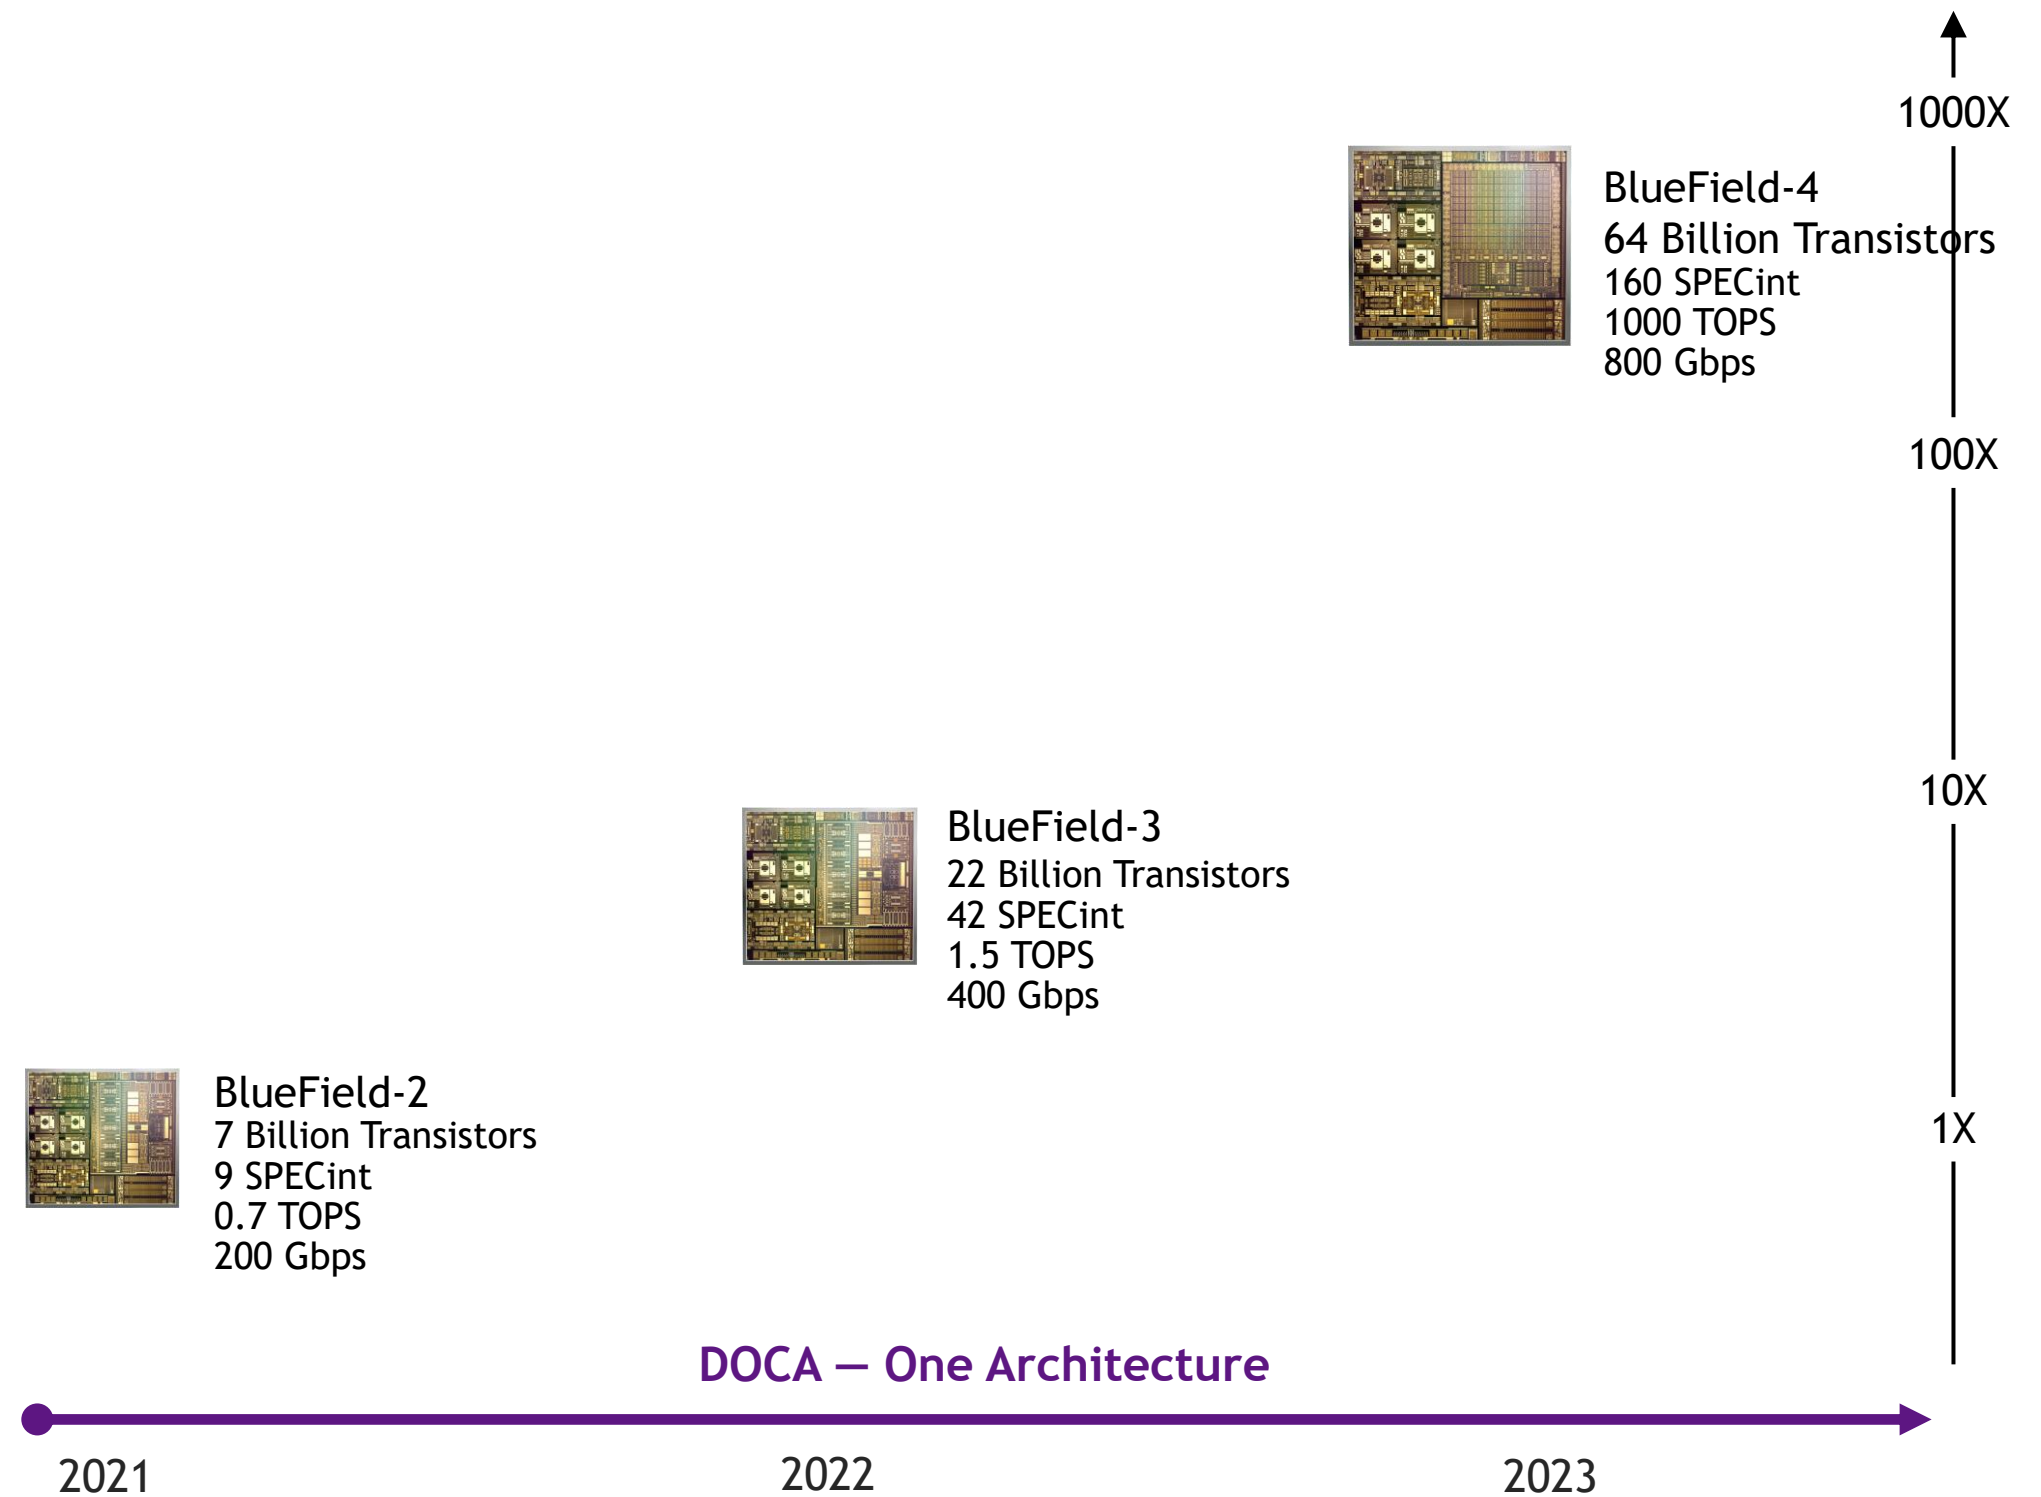


JUNIPER
NETWORKS



NVIDIA DPU ROADMAP

Programmable Data Center
Infrastructure-on-a-Chip





ETHERNET SWITCHING

NVIDIA NETWORKING

Cloud-Scale Efficiency

Operational Efficiency



Built for Automation

Hardware Efficiency



Best in Class Performance

Visibility



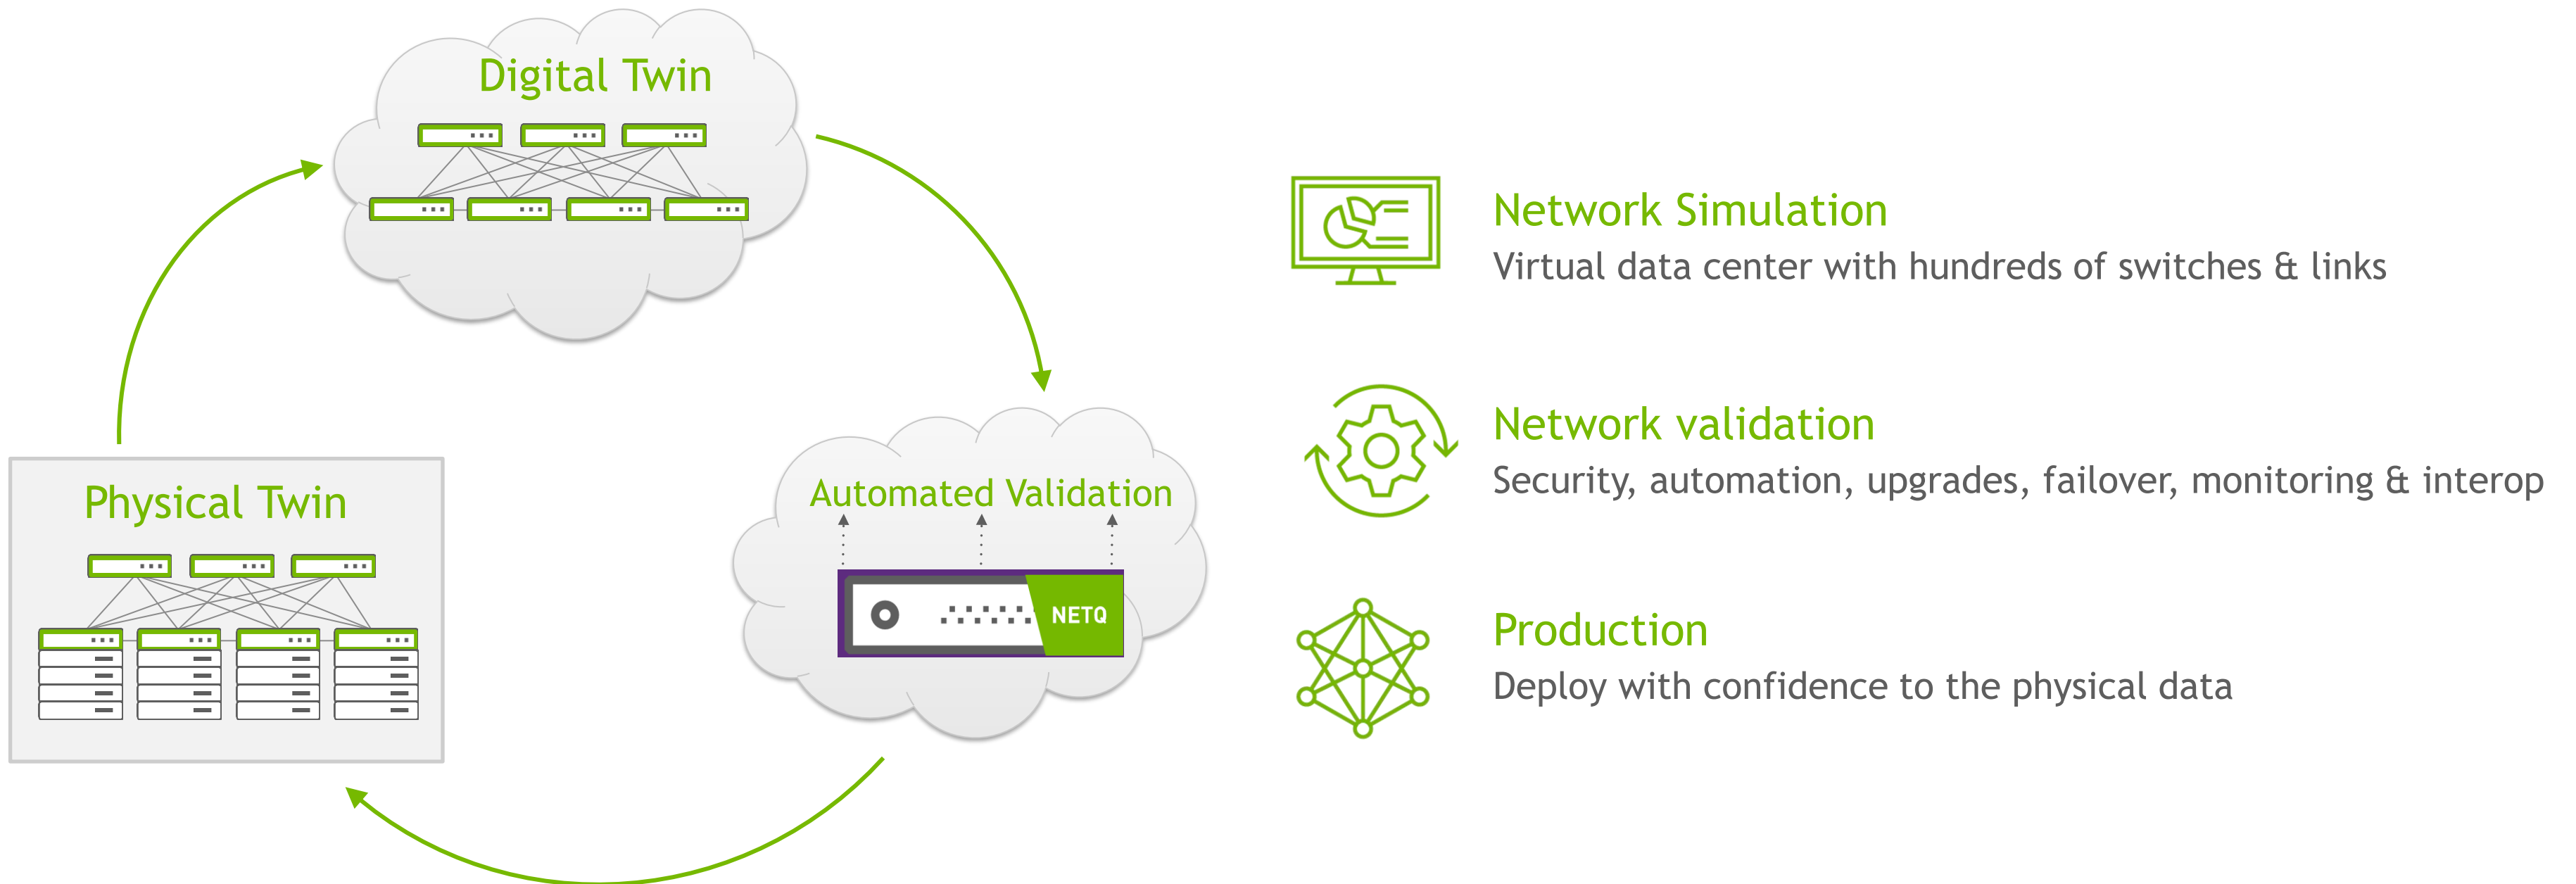
Advanced Telemetry

36% Less IT Time Keeping the Lights On | 2x More Time Focused on Innovation

Source: IDC, Business Value Survey, March 2021

95% FASTER SWITCH DEPLOYMENTS


NVIDIA NetDevOps with CI/CD



GETTING STARTED

Take a Virtual Spin Today

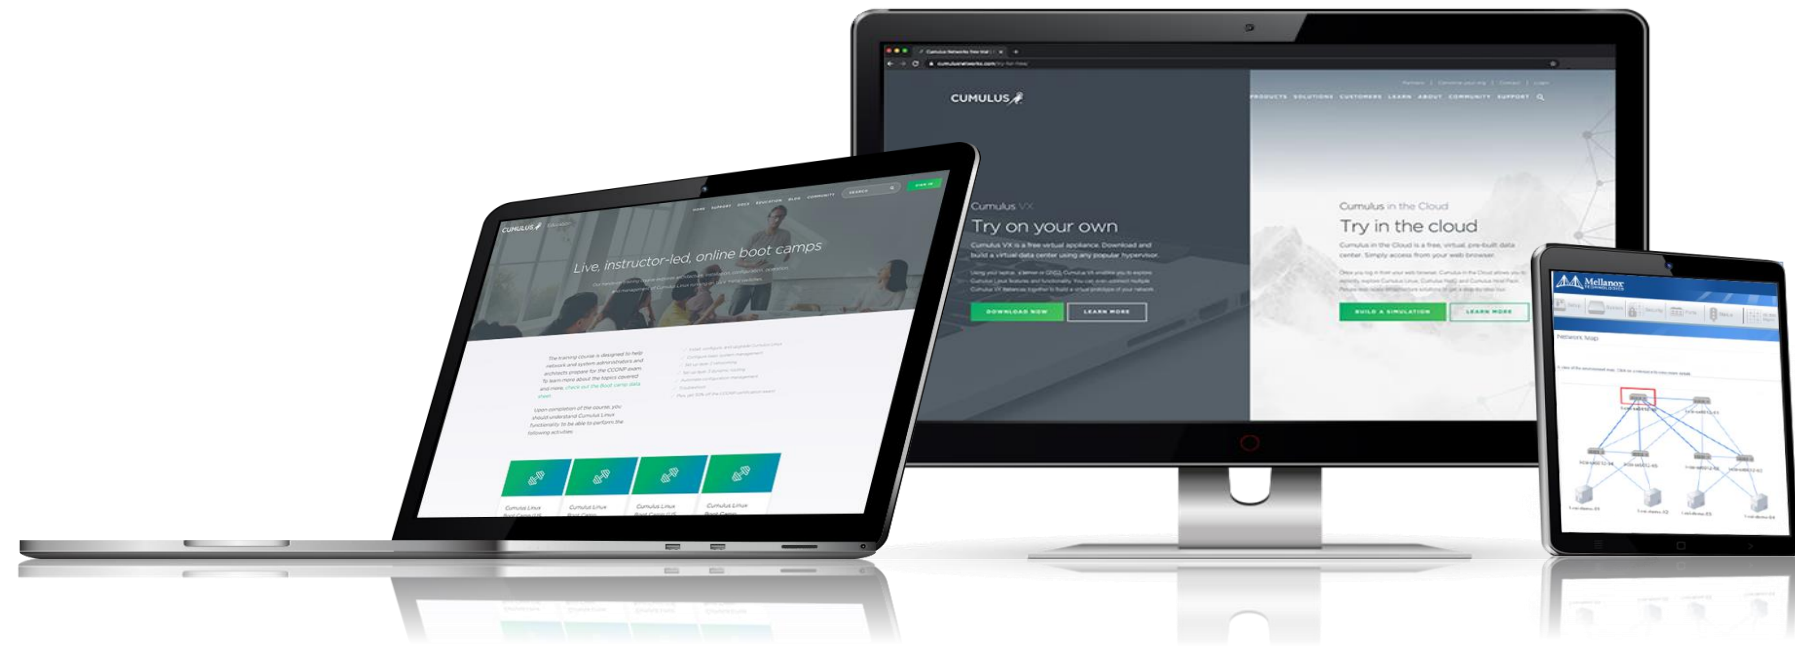
Existing Network



air.nvidia.com

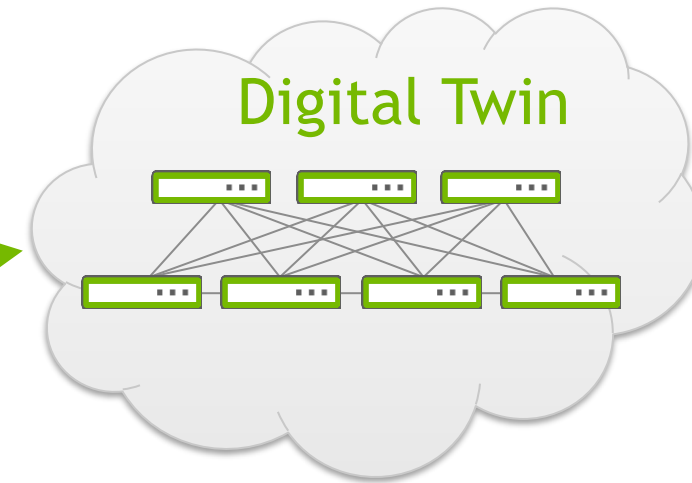
Digital Clone:

- Setup in minutes
- 100's of switches
- 1000's of servers
- Exact topology

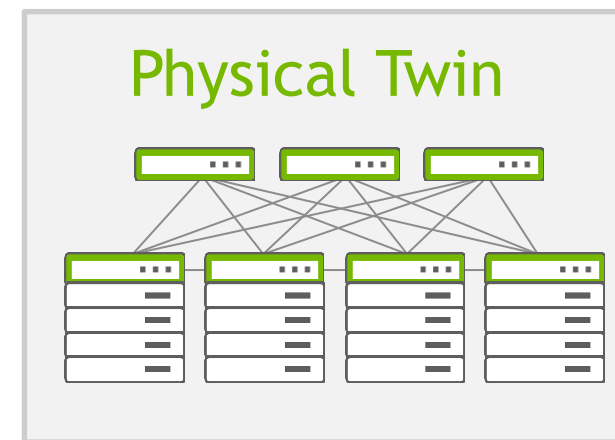


SA-assisted or Self-service
via browser to NVIDIA Air

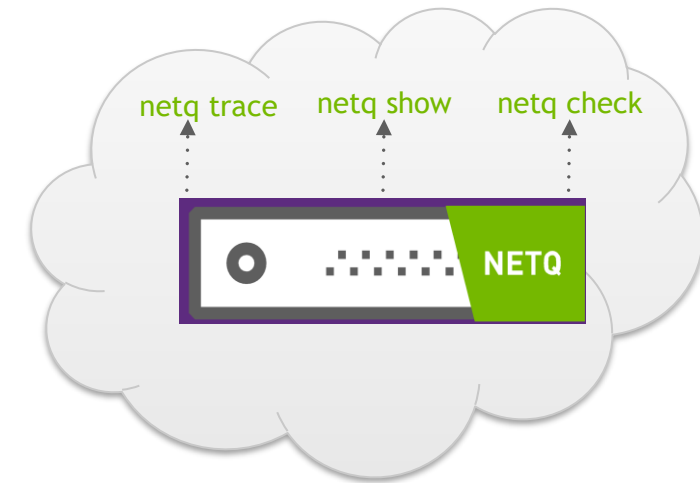
Digital Twin



Physical Twin



netq trace netq show netq check



NETWORK OPERATING SYSTEMS

Enabling Choice of Network Operating Experience

	NVIDIA Cumulus Linux	NVIDIA Mellanox Onyx	SONiC
	Commercial	Commercial	Free & Open Source
Linux	✓✓✓	✓	✓✓✓
Turnkey Automation	✓✓✓	✓	✓
EVPN/VXLAN Functionality	✓✓✓	✗	✗
Maturity	✓✓✓	✓✓	✓
ROCE functionality	✓✓✓	✓✓✓	✓
Industry Standard CLI	✓	✓✓✓	✗

NVIDIA NETWORKING

Cloud-Scale Efficiency

Operational Efficiency



Built for Automation

Hardware Efficiency



Best in Class Performance

Visibility

















Advanced Telemetry

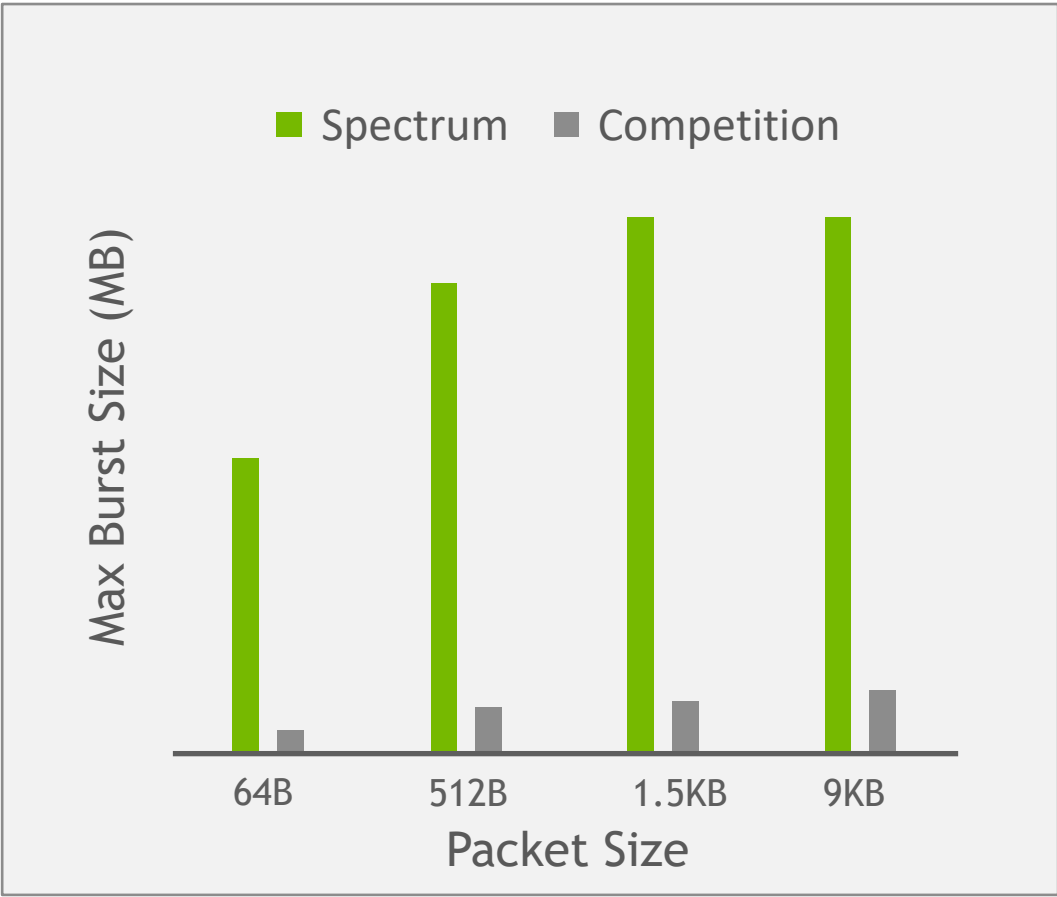
4 Generations of NVIDIA Switching Silicon Purpose-built for the Data Center

BROAD ETHERNET SWITCH PORTFOLIO

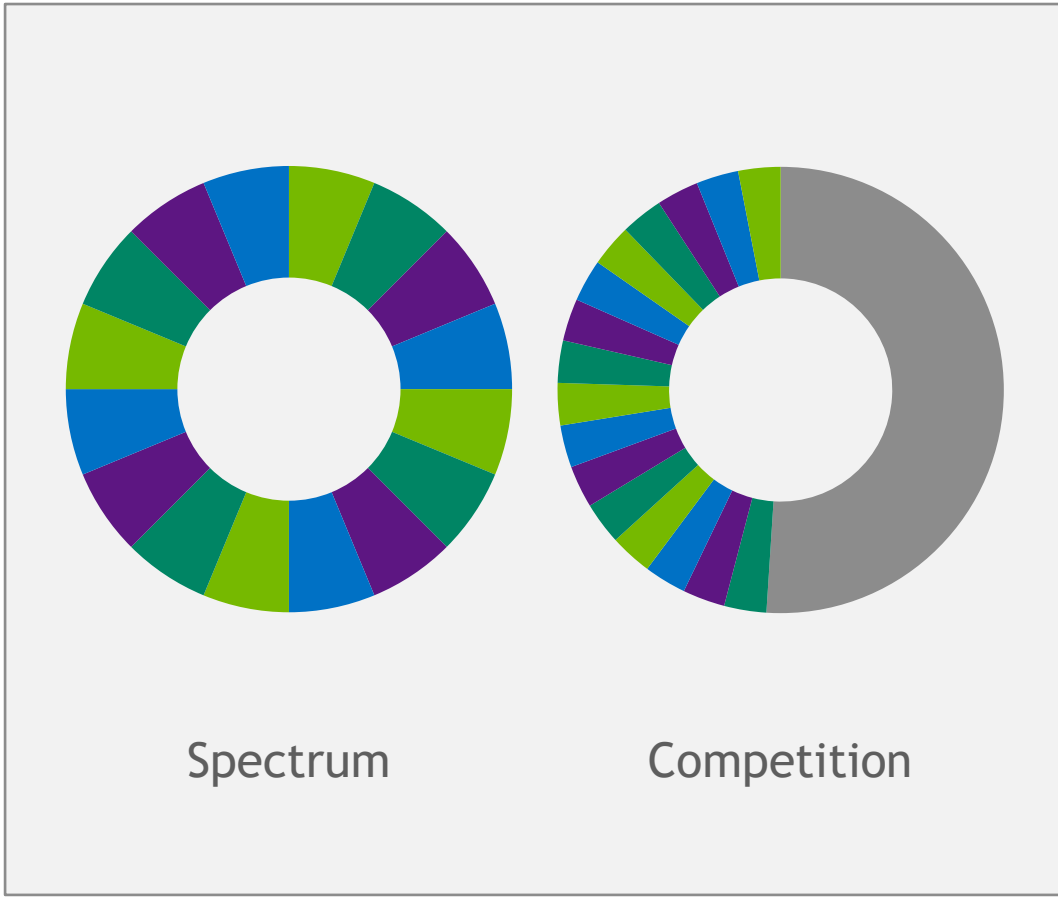
A Switch for Every Workload

Edge Switches	Leaf/Spine Switches	AI Switches	
 SN2100: 16x100G	 SN2410: 48x25G + 8x100G	 SN3700-C: 32x100G	 SN3700-V: 32x200G
 SN2010: 18x25G + 4x100G	 SN3420: 48x25G + 12x100G	 SN4600-C: 64x100G	 SN4600-V: 64x200G -Q3*
 SN2201: 48x 1G + 4x100G - Q1'22*	 SN4410: 48x100G + 8x400G - Q3*	 SN4700: 32x400G	
	 SN4800: 128x100G - Q4*	 SN5400: 64x400G - 2022*	
		 SN5600: 64x800G or 128x400G - 2022*	

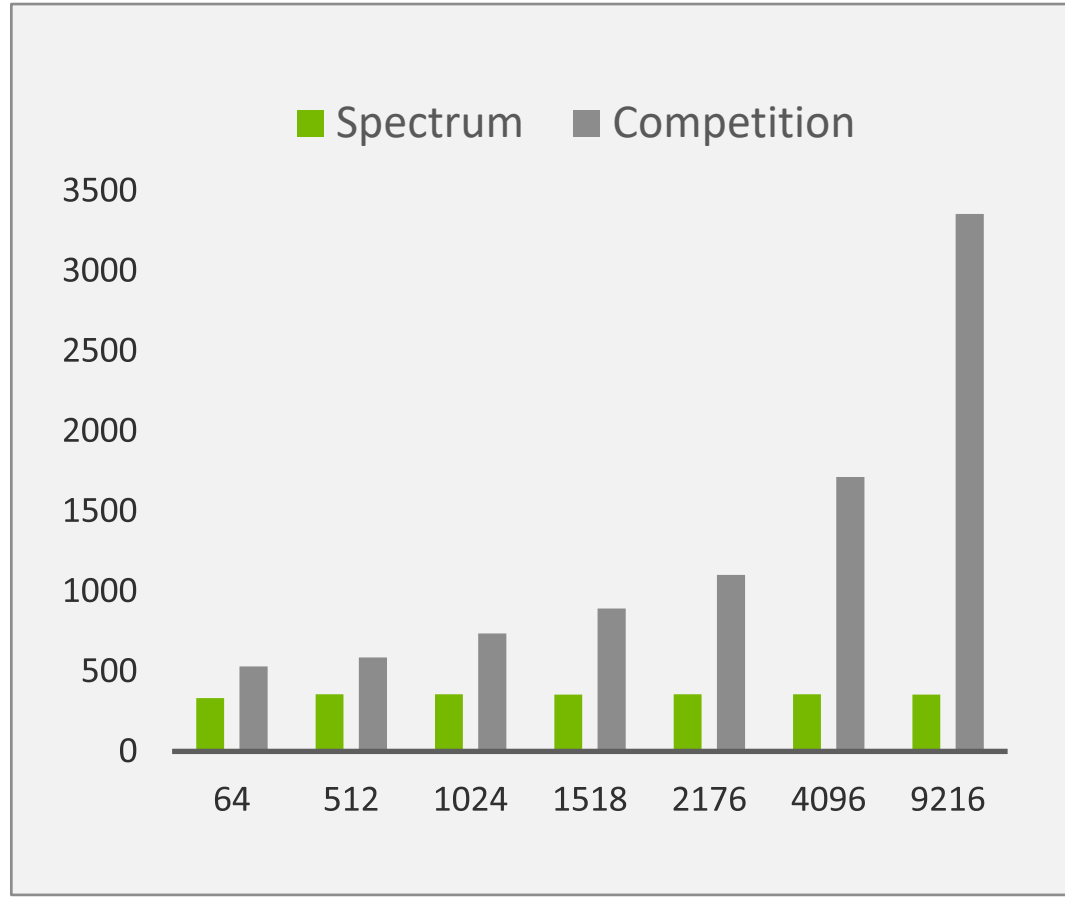
INDUSTRY LEADING SWITCH ASIC PERFORMANCE



Superior Burst Absorption
Microburst Absorption Capability



Fair Bandwidth Sharing
Bandwidth Sharing Across Ports




Lowest Latency
Packet Latency (ns)

NVIDIA SPECTRUM

Superior Buffer Architecture


Fair Bandwidth Distribution



Fully Shared Buffers

- Maximum burst absorption
- Fair
- Predictable

4 Times Larger Effective Buffer Size



Unfair & Unpredictable Bandwidth sharing

Split Buffers

- Unpredictable
- Fairness issues
- Burst size based on ingress port or egress port

Merchant Silicon Split-buffers

NVIDIA NETWORKING

Cloud-Scale Efficiency

Operational Efficiency



Built for Automation

Hardware Efficiency



Best in Class Performance


Visibility




Advanced Telemetry

ACCELERATING TIME TO ROOT CAUSE

Legacy Network Monitoring

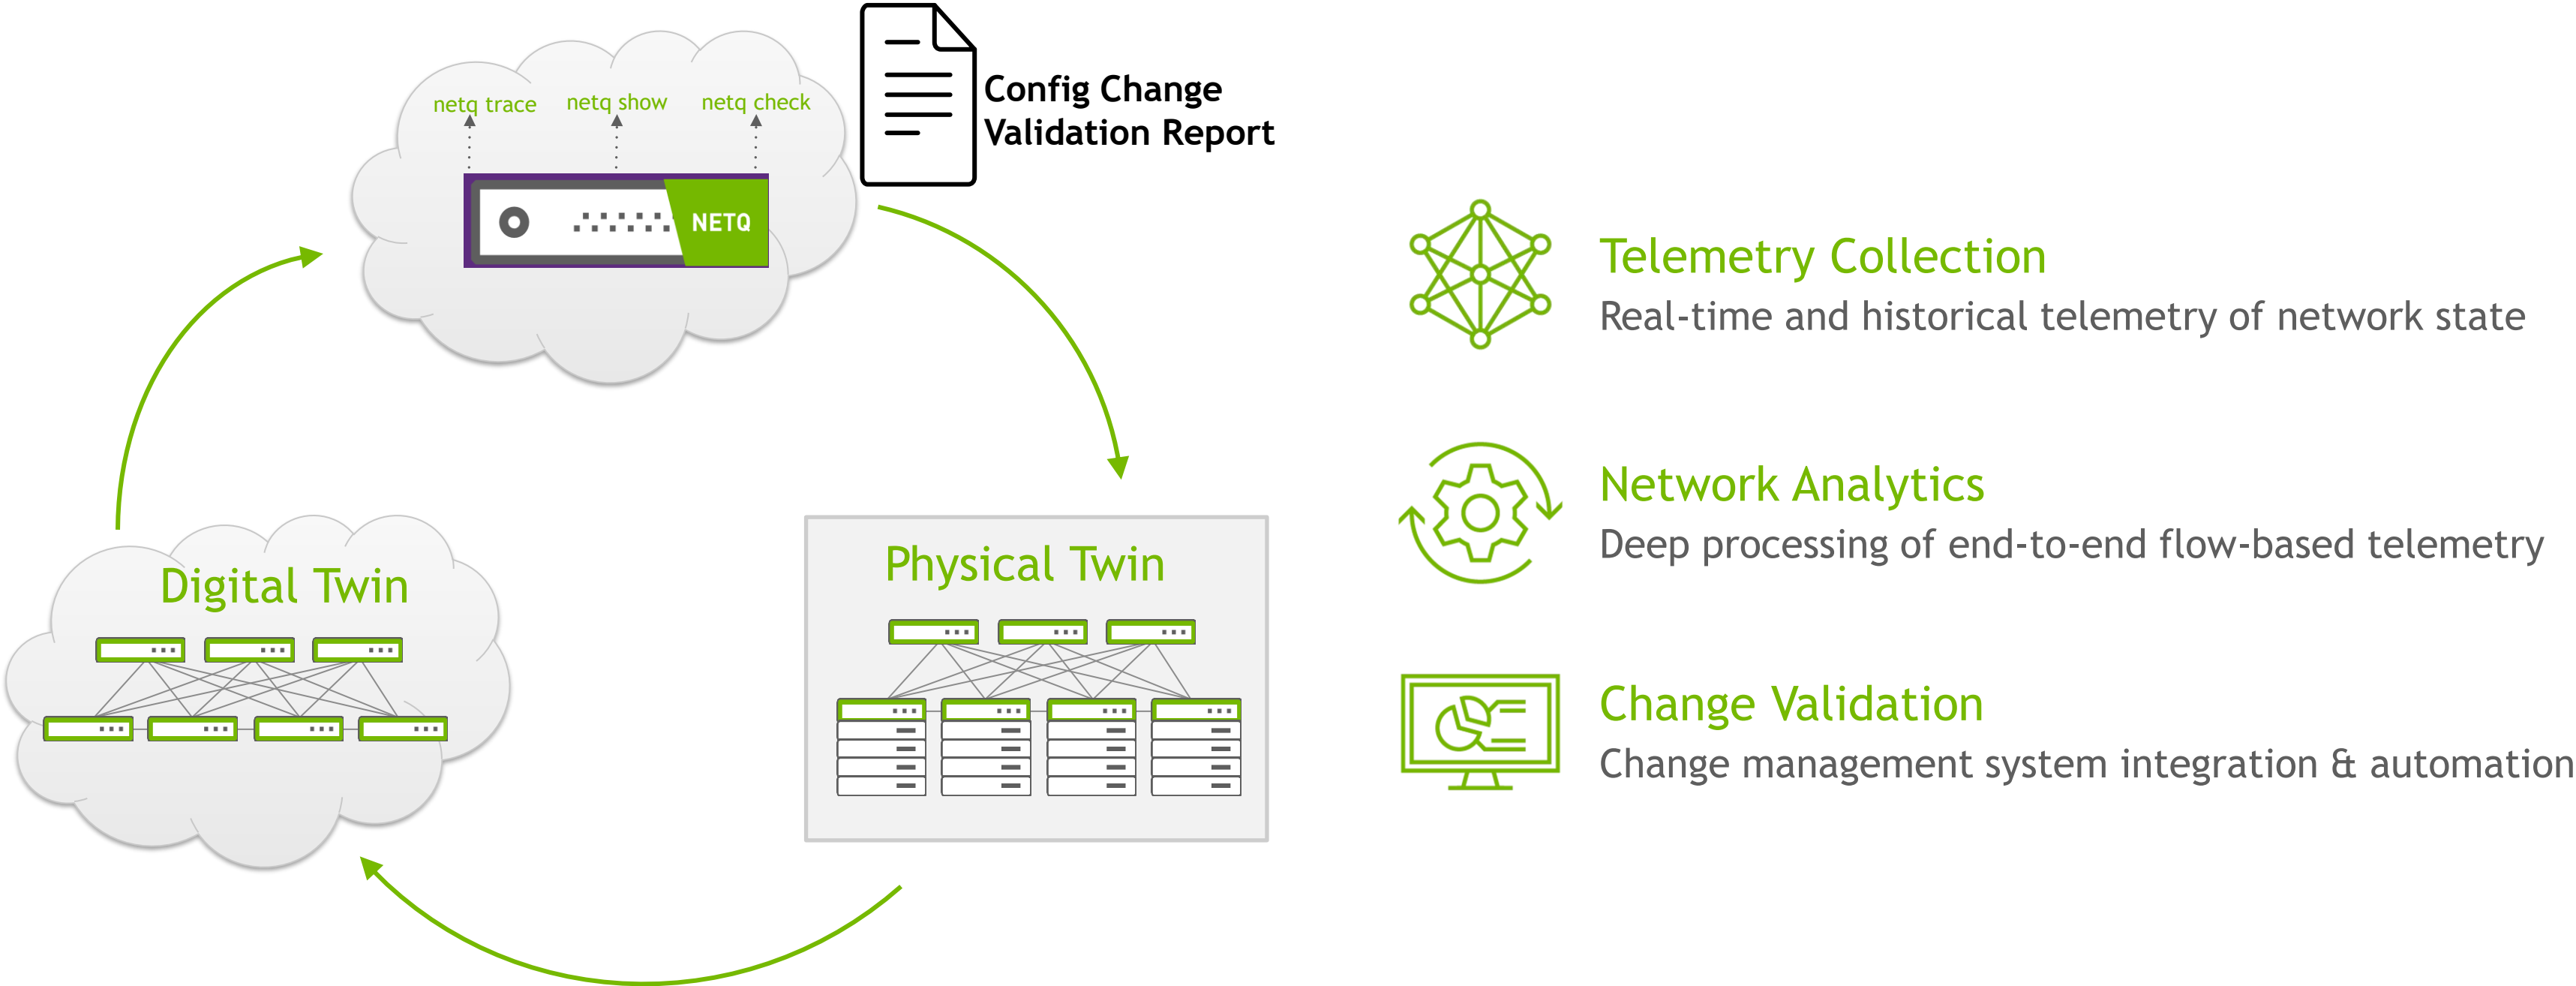


What Just Happened



NVIDIA NETQ

Automated Change Management using NetDevOps



CLOUD-SCALE NETWORKING STACK

SPECTRUM ASIC

Best-in-class performance with the lowest latency, highest packet rate, and best buffer architecture



HARDWARE ACCELERATED TELEMETRY

WJH advanced telemetry offers the deepest visibility to instantly identify root causes



BUILT FOR AUTOMATION

NetDevOps Operational model
Free Production Ready Automation



TURNKEY FABRIC VALIDATION

Complete lifecycle management and network fabric operations tooling with NVIDIA NetQ



FREE RANGE ROUTING

Best open-source routing stack with large, active development community driving unparalleled feature velocity



ROBUST YET SIMPLE EVPN

Cutting edge EVPN functionality with smart defaults
- 10 config lines vs “industry standard” 90+ lines

