


NVIDIA IS REINVENTING THE DATA CENTER

The Data Center is the New Unit of Computing



CHALLENGES FOR HIGH FREQUENCY TRADERS

Increasing Market Data rates and bursts

- Market data feeds peaking above 10Gbps
- 25G and 40G handoffs available already

No innovation above 10G

- Existing Low Latency solutions capped at 10G

Obscured Network

- Lack of visibility for congestion
- Lack of real time buffer utilization

**Merchant
Network
Silicon is not
aligned with
HFT Market**

LOWEST LATENCY SWITCHES FOR HFT

ToR & Edge

ToR & Spine

Lowest Latency + Features

Spectrum



SN2010: 18x1/10/25G + 4x100G



SN2410: 48x1/10/25G + 8x100G



SN2100: 16x100G



SN2700: 32x100G

- ✓ 300ns Latency *
- ✓ PTP Boundary Clock
- ✓ Egress SPAN Timestamping

Low Latency + Advanced Features

Spectrum-2

Spectrum-3



SN3420: 48x1/10/25G + 12x100G

AI Switches



SN3700: 32x200G



SN4600: 64x200G



SN3700C: 32x100G



SN4600C: 64x100G



SN4800: 128x100G, 64x200, 32x400G

- ✓ NAT in Hardware
- ✓ 400 or 500ns Latency *
- ✓ PTP Boundary Clock
- ✓ Stratum3 Oscillator
- ✓ In-Band and SPAN Timestamping

* Not STAC Benchmarks

3 GENERATIONS OF ASICS

Lowest Latency, High Performance, HFT Feature Set




	Spectrum	Spectrum-2	Spectrum-3
Buffer Architecture	Fully Shared	Fully Shared	Fully Shared
Buffer Size	16MB	42MB	64MB
Latency	300ns *	400ns *	500ns *
Low Jitter	Yes	Yes	Yes
Forwarding Table	256k	512k	512k
Max. Multicast Scale	13k	100k+	100k+
PTP Boundary Clock	Yes	Yes	Yes
Stratum3 Oscillator	No	Yes	Yes
SPAN Timestamping	Yes	Yes	Yes
In-Band Timestamping	No	Yes	Yes
NAT in Hardware	No	Yes	Yes

* Not STAC Benchmarks

NETWORK PACKET BROKERING

cVu® 32100/32100E Series

- Data Center
- HPC Cluster
- HFT / FinServ



- ▶ Up to 100G real-time intelligence
- ▶ Non-blocking design with zero packet-loss
- ▶ Low latency monitoring with millisecond analytics
- ▶ Nanosecond time-stamping with PTP synchronization
- ▶ Header stripping for protocols
- ▶ Open monitoring architecture with RESTful
- ▶ Flexible aggregation through TAP/SPAN/RSPAN
- ▶ Centrally manageable through cPacket cClear®
- ▶ Integrated w/ cPacket cStor® Packet Capture & cProbe® flow devices



*“The NVIDIA Spectrum-2 switch provides the **lowest latency available** with a **non-blocking design** and **zero packet-loss** at speeds up to 100GbE. This performance is mandatory to provide **crucial network analytics and monitoring** in the **most demanding trading environments**.”*


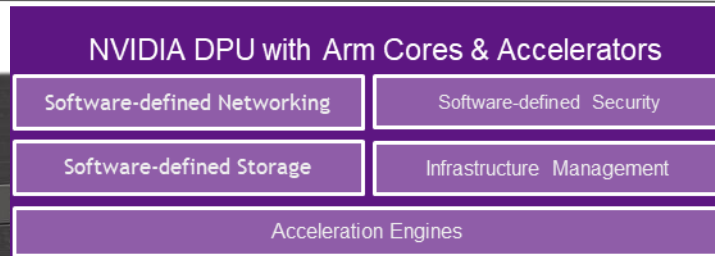

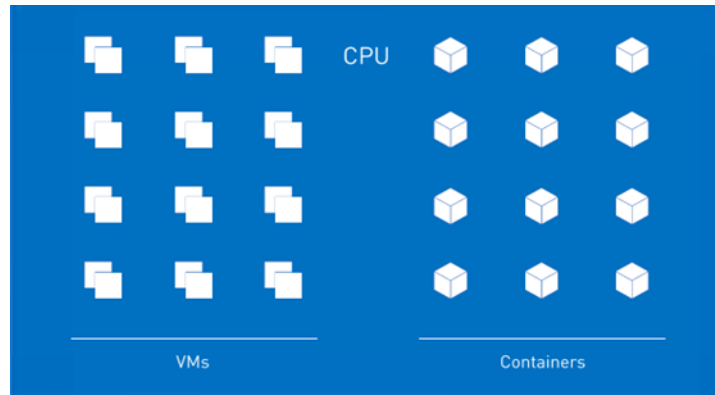
- Daryl Stanbery, Vice President, cPacket

NVIDIA BLUEFIELD-2 DPU (DATA PROCESSING UNIT)

Data Center Infrastructure-on-a-Chip

X86 CPU

Bluefield-2 DPU



NVIDIA DOCA

Enabling Broad DPU Partner Ecosystem

Software application framework for BlueField DPUs

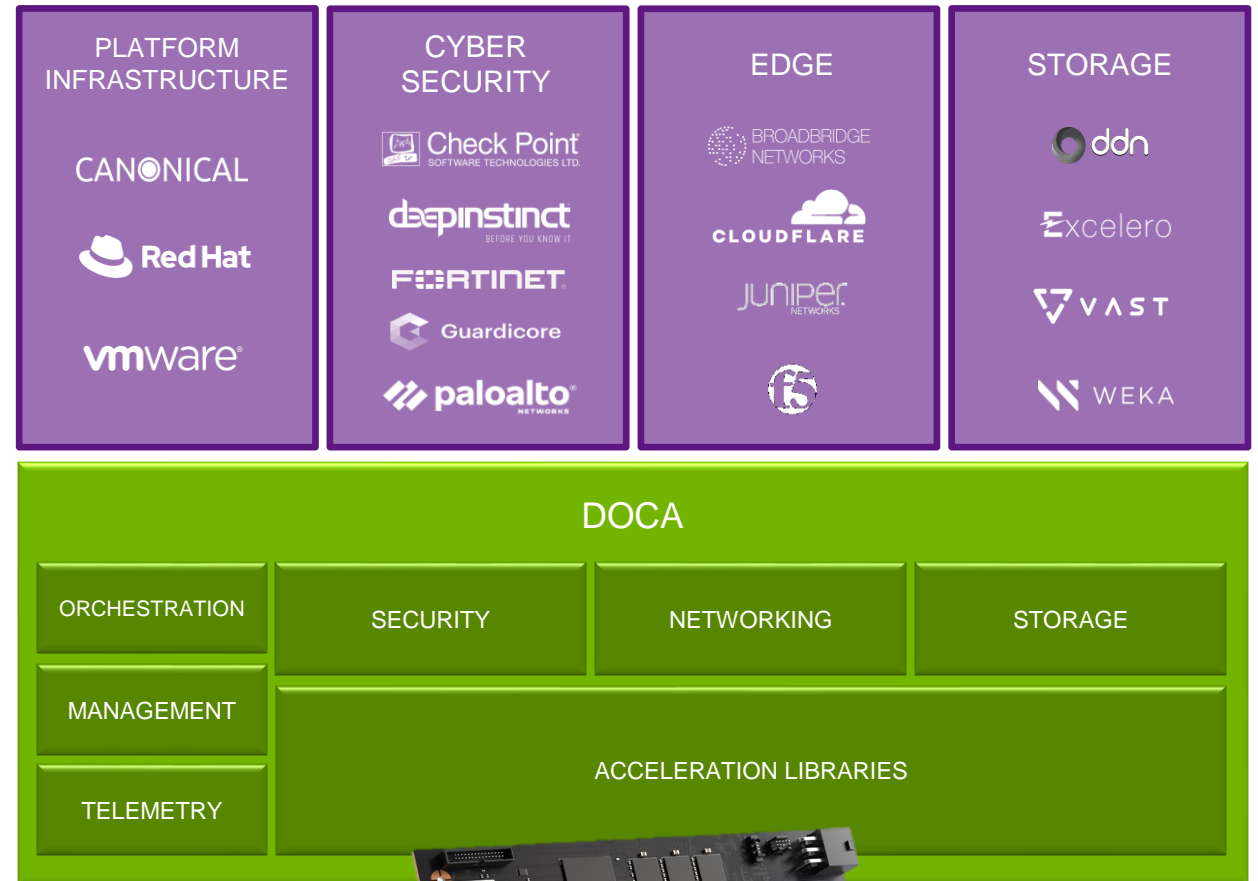
DOCA is for DPUs what CUDA is for GPUs

Protects developer investment for future DPUs

Certified reference applications, APIs & partner solutions

Rich partner ecosystem across industries and workloads

Sign up for DOCA access at <https://developer.nvidia.com/networking/doca>



PRECISION TIMING IN DPU & CONNECTX

Ultra-precise Network Time Stamps and Synchronization

Transaction Time Stamping


Precise transaction records and event logs at every location

Distributed Applications

Synchronize apps, storage and monitoring across multiple sites

Streaming Data

Optimize capture/broadcast of market data



View GTC Session w/ Facebook at:

https://gtc21.event.nvidia.com/media/Time%20Synchronization%20in%20Distributed%20Data%20Centers%20%5BS31889%5D/1_qytjmdhx

* Not STAC Benchmarks



ONLINE CONFERENCE

View Sessions at:

www.NVIDIA.com/GTC

