- Cisco Live Las Vegas 2017

NetApp

Accelerate Graphics Performance in VDI Environments with NetApp HCI

Chris (C-Rod) Rodriguez, Principal Solution Architect, NetApp

Nachiket Karmarkar, Product Manager, NVIDIA

NetApp's Data Fabric accelerates data-driven digital transformation by simplifying and integrating data management.

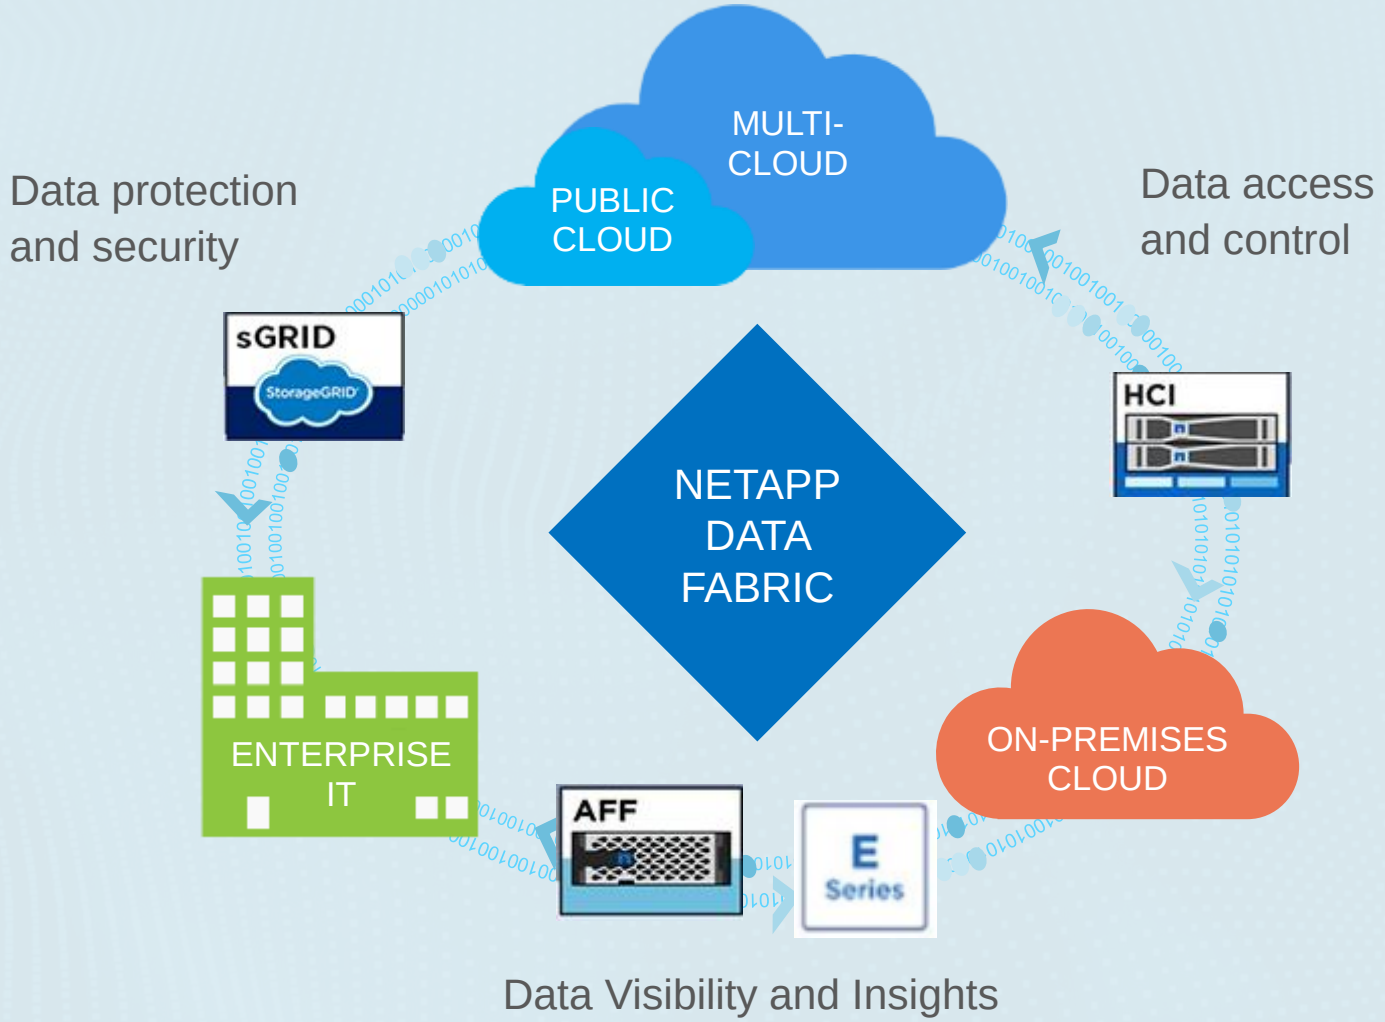

HCI | Solid Fire | AFF FAS | E-Series | Active IQ | NKS | Trident

NetApp

NetApp HCI Differentiators

What makes NetApp HCI great

1. Scale Compute Nodes and Storage Independently

2. Mix Different Models of Compute or Storage nodes

3. Add nodes or remove nodes during production

4. Meet SLA's/XLA's with Guaranteed Performance

5. Any VMware ISO automatically deployed or Manual deploy Linux / Bare Metal

6. Not locked into one Hypervisor (e.g. VSAN)

7. Support for all three major Cloud Vendors
   - AWS
   - Azure
   - Google

8. Easy On-Premise and Public Cloud support with NetApp Kubernetes Service

9. Pre-Validated Solutions for easy Partner Offerings

10. Integrated Solution with all NetApp Data Fabric products

Traditional Graphics Workspace Challenges

- High performance and resolution workstations
- Limited collaboration features
- Data security
- Tied to the workplace
- Technology refreshes and patch management
- Expensive workstations and maintenance cost

NetApp

NetApp VDI Solution

Move the Intensive Graphic applications into a centralized computer room on a VDI desktop

Mitigate business challenges:

8K Resolution

- Network and power outages do not affect users' working session
- Users' resources expa workflow and multiwo
- Data stays within the regardless of where c accessed.
- Very low bandwidth requireme mean users can connect from the most remote locations.

4K Resolution

HD Resolution

HD Resolution

Local -or- Remote Access

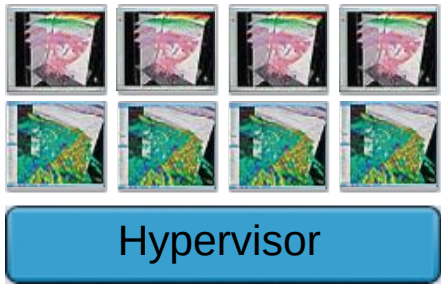

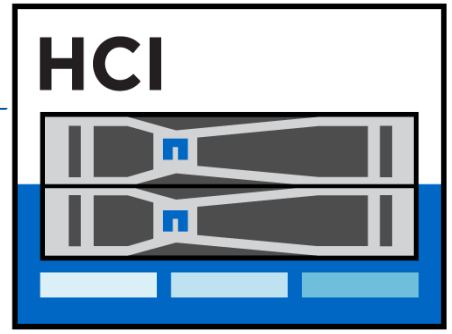

HCI

NetApp

NetApp H615C: 1U HCI Compute Node

Delivering platform flexibility to HCI deployments

- 1U form-factor compute node
- Intel Cascade Lake CPU generation
  - Higher CPU core counts and clock speed
  - Increased RAM capacity options
    - 384GB – 1.5TB
- Three NVIDIA Tesla T4 GPUs
  - Mid-range to High-end graphics
    - CAD, medical imaging, seismic exploration
  - AI Edge Processing
  - AI Inferencing

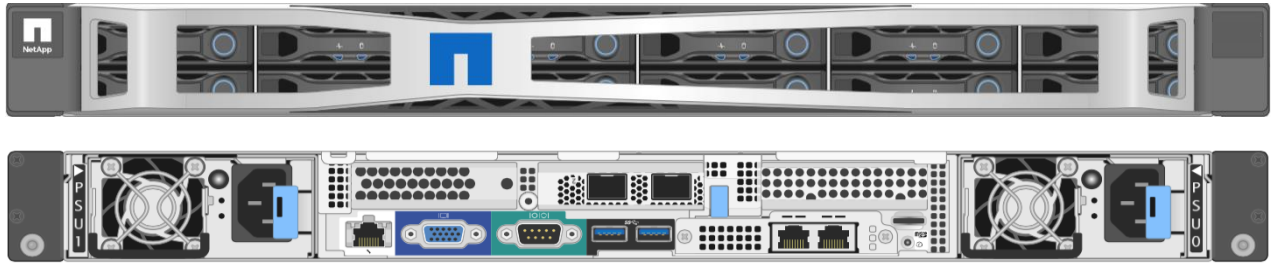

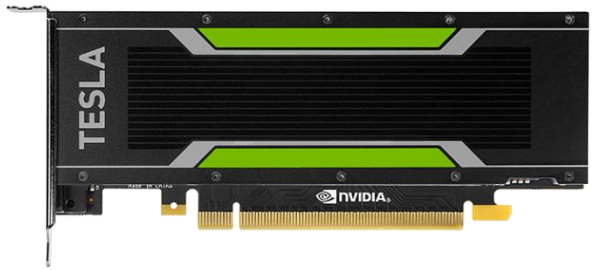

Use Cases

- **Oil&Gas**: Follow the sun engineering and collaboration for P&E like Schlumberger Petrel; Geosystemic applications (H615C coming)
- **Manufacturing**: CAD/CAE, PLM like CATIA, SolidWorks, Siemens PLM NX
- **Healthcare**: MRI, PET and CT image analysis
- **Media**: Show customers animations, bring remote artistic expertise.
- **ISP's**: 3D Desktop as a Service. Video rendering, manufacturing, telemedicine, collaboration,…

NetApp

NetApp 3D HCI Solution Simplifies Graphics

3D HCI removes the expensive workstation from the desk

Mobility / Security

Performance

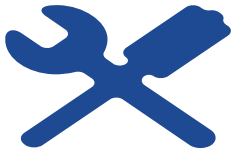

SLA

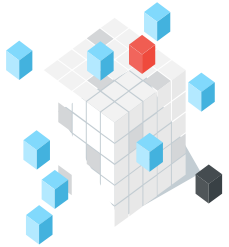

TCO

NetApp

User experience is the most important factor for VDI and DaaS

Mobility

1 in 3 workers are remote or multilocation today

Security

Loss of IP and data from cyberattacks, natural disaster, human error

Cost

Achieve highest user density while maintaining performance

Pro 3D Apps

Deliver a full Quadro experience from data center or cloud

Modern Apps

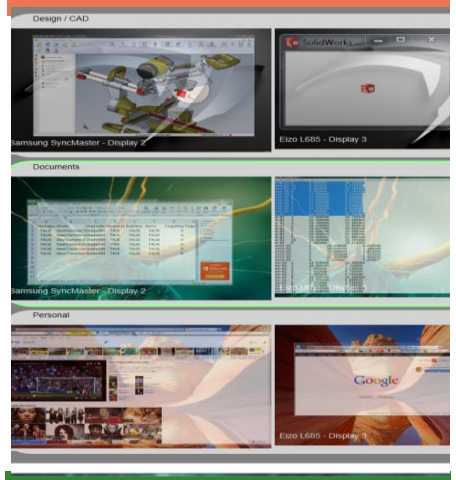

82% of global IP traffic will be video by 2021

Sources: Gartner, Lakeside, Cisco Visual Networking Forecast

GPU VIRTUALIZATION FOR ANY WORKLOAD

NVIDIA delivers GPU virtualization for both graphics and compute workloads

How It Works

NVIDIA virtual GPU products deliver a GPU experience to every virtual desktop

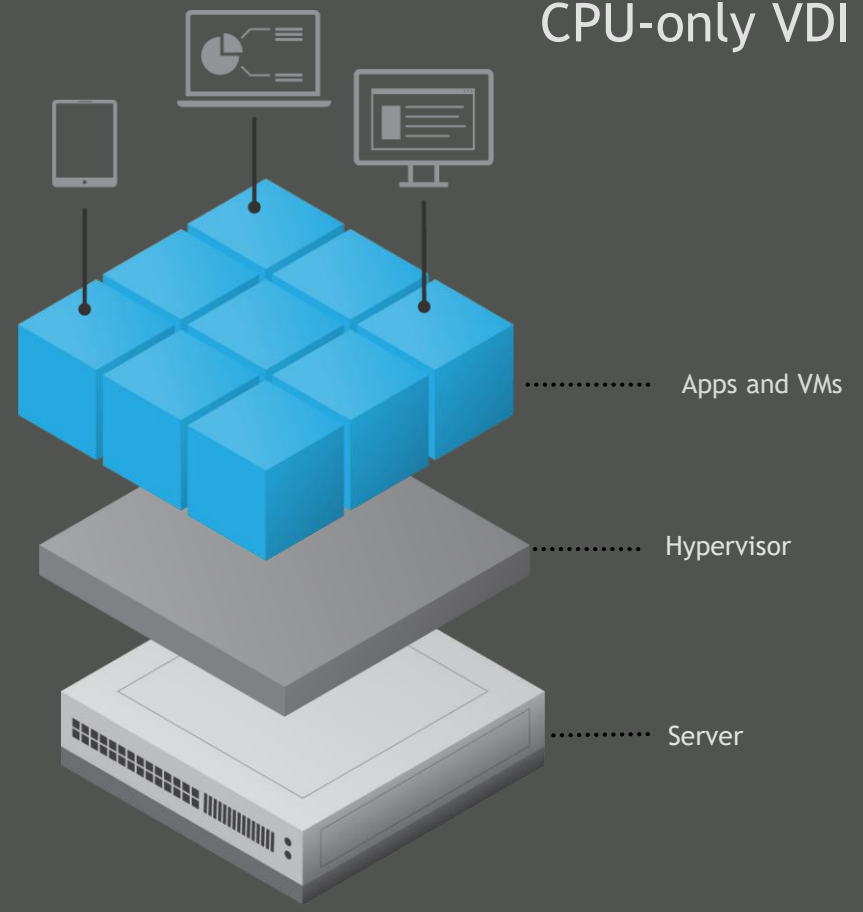

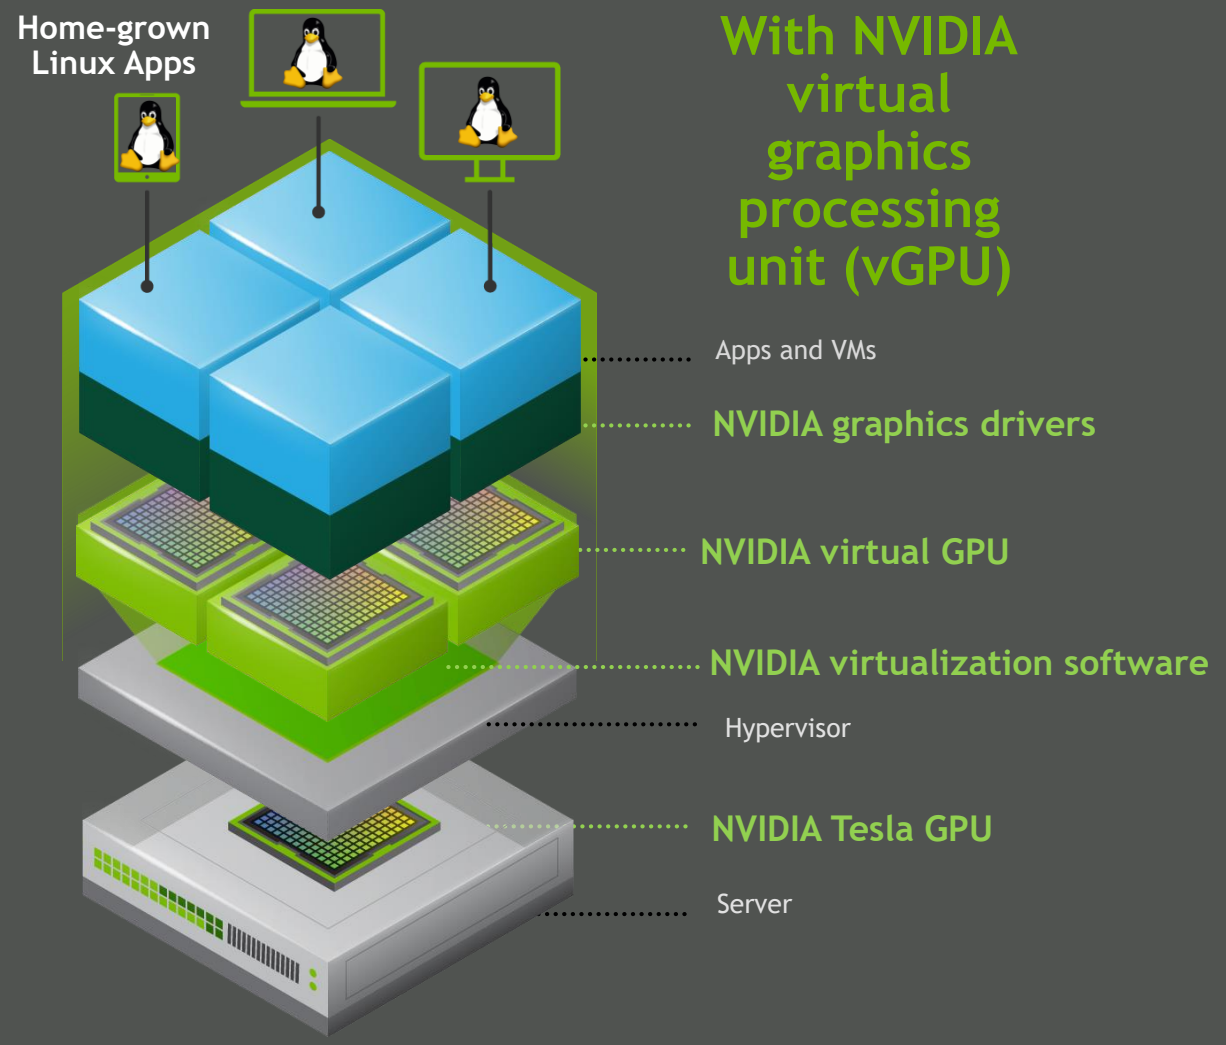

NVIDIA T4 Key Specifications

GPU Architecture	**NVIDIA Turing**
NVIDIA CUDA® Cores	**2,560**
NVIDIA Turing™ Tensor Cores	**320**
RT Cores	**40**
Giga Rays/second	**5**
Memory Size	**16 GB GDDR6**
Memory BW	**Up to 320 GB/s**
vGPU Profiles	**1 GB, 2 GB, 4 GB, 8 GB, 16 GB**
Form Factor	**PCIe 3.0 single slot (half height & length)**
Power	**70W**
Thermal	**Passive**

RTX Performance in a Quadro Virtual workstation

Support for up to 5 Giga Rays/Sec

Real-time Rendering

Simulation, modeling, design

Rendering, design

Real-Time Inference Performance

Quadro Virtual Workstation for deep learning inferencing workloads

Support for NVIDIA GPU Cloud (NGC)

Ideal for deep learning labs and classrooms

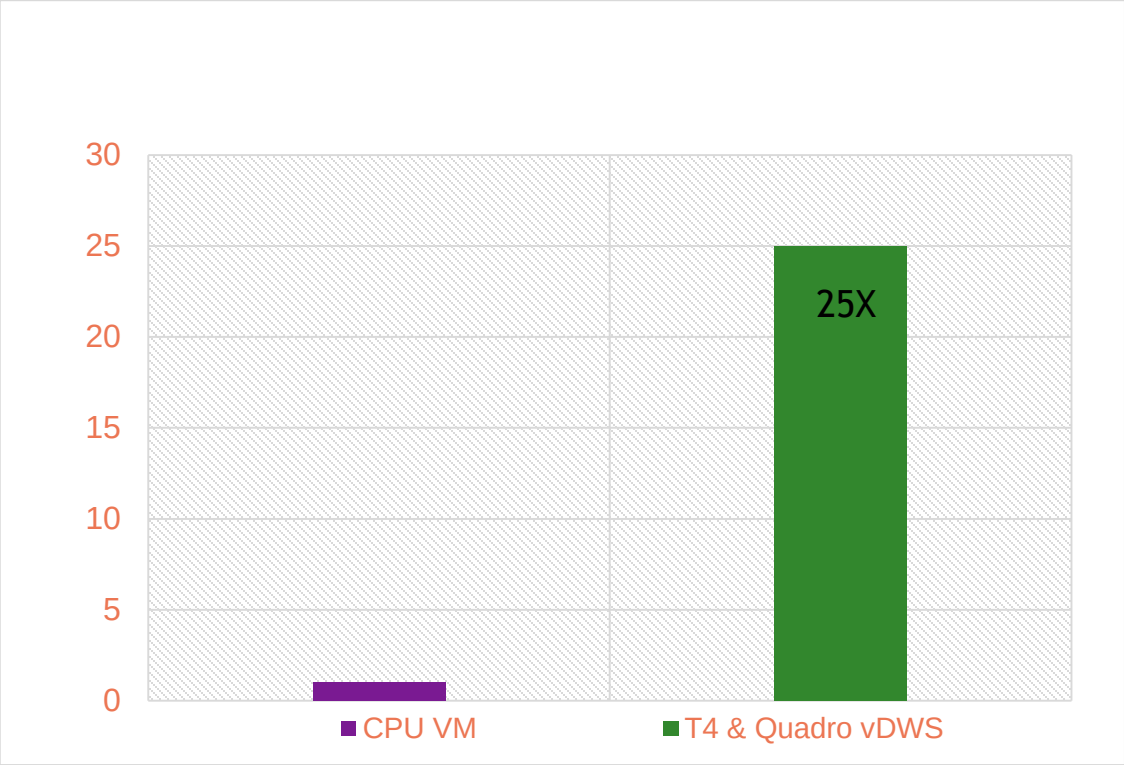

Speedup: 25x faster

ResNet-50 (7ms latency limit)

Tested on a server with Intel Xeon Gold 6154 (18C, 3.0 GHz), Quadro vDWS with T4-16Q, VMware ESXi 6.7, host/guest driver 410.87/412.10, VM config, Ubuntu 16.04, 8 vCPU, 32GB memory. 25X performance improvement over CPU VM.

Powerful, Versatile Platform for VDI

Powerful **virtual workstation** for the engineer, professional designer, and data scientist

Deep learning inferencing for virtual labs and classrooms

High density **virtual desktops** for the best user experience for Windows 10

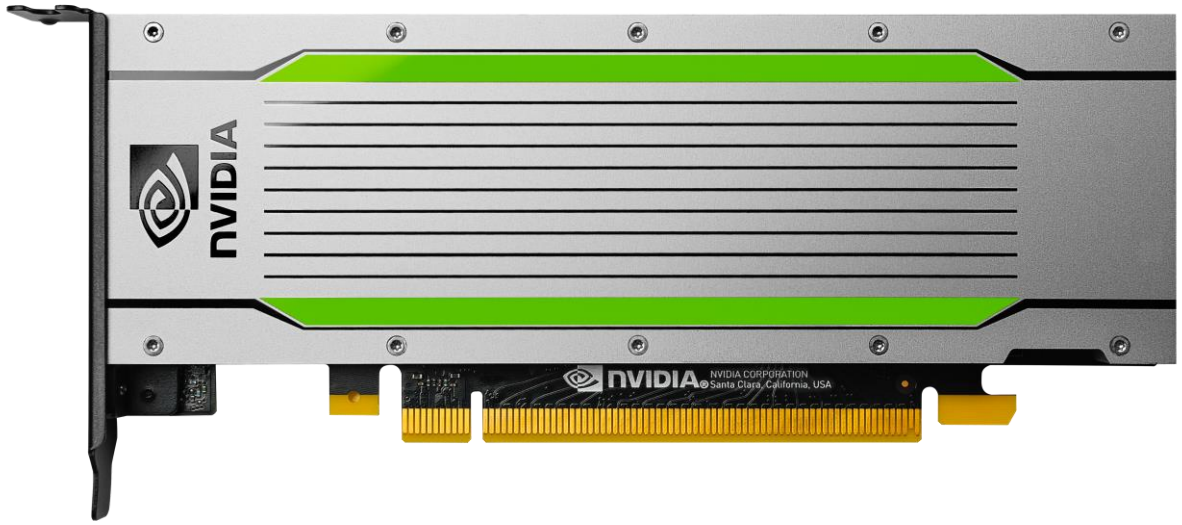

Questions?