Hewlett Packard Enterprise

# GPU 加速型 VDI 環境的完美選擇

馬西聰
資訊系統技術經理
HPE Taiwan

# Current Customer Challenges

# Enterprise continues to go “virtual” as apps, mobility, & cloud take hold


VDI growth driven by

Mobility | BYOD | Security | Manageability

Hewlett Packard Enterprise | NVIDIA

# Graphics are Required for More Workloads

Today’s employees need access to a great user experience at all levels


**2x**

Number of applications that require graphics doubled since 2012

**56%**

Over half of enterprise users access at least one graphics accelerated app

**30%**

Windows 10 requires a 30% increase in CPU consumption, compared with Windows 7

Hewlett Packard Enterprise | NVIDIA

# Business user Graphical Requirements are rapidly Rising

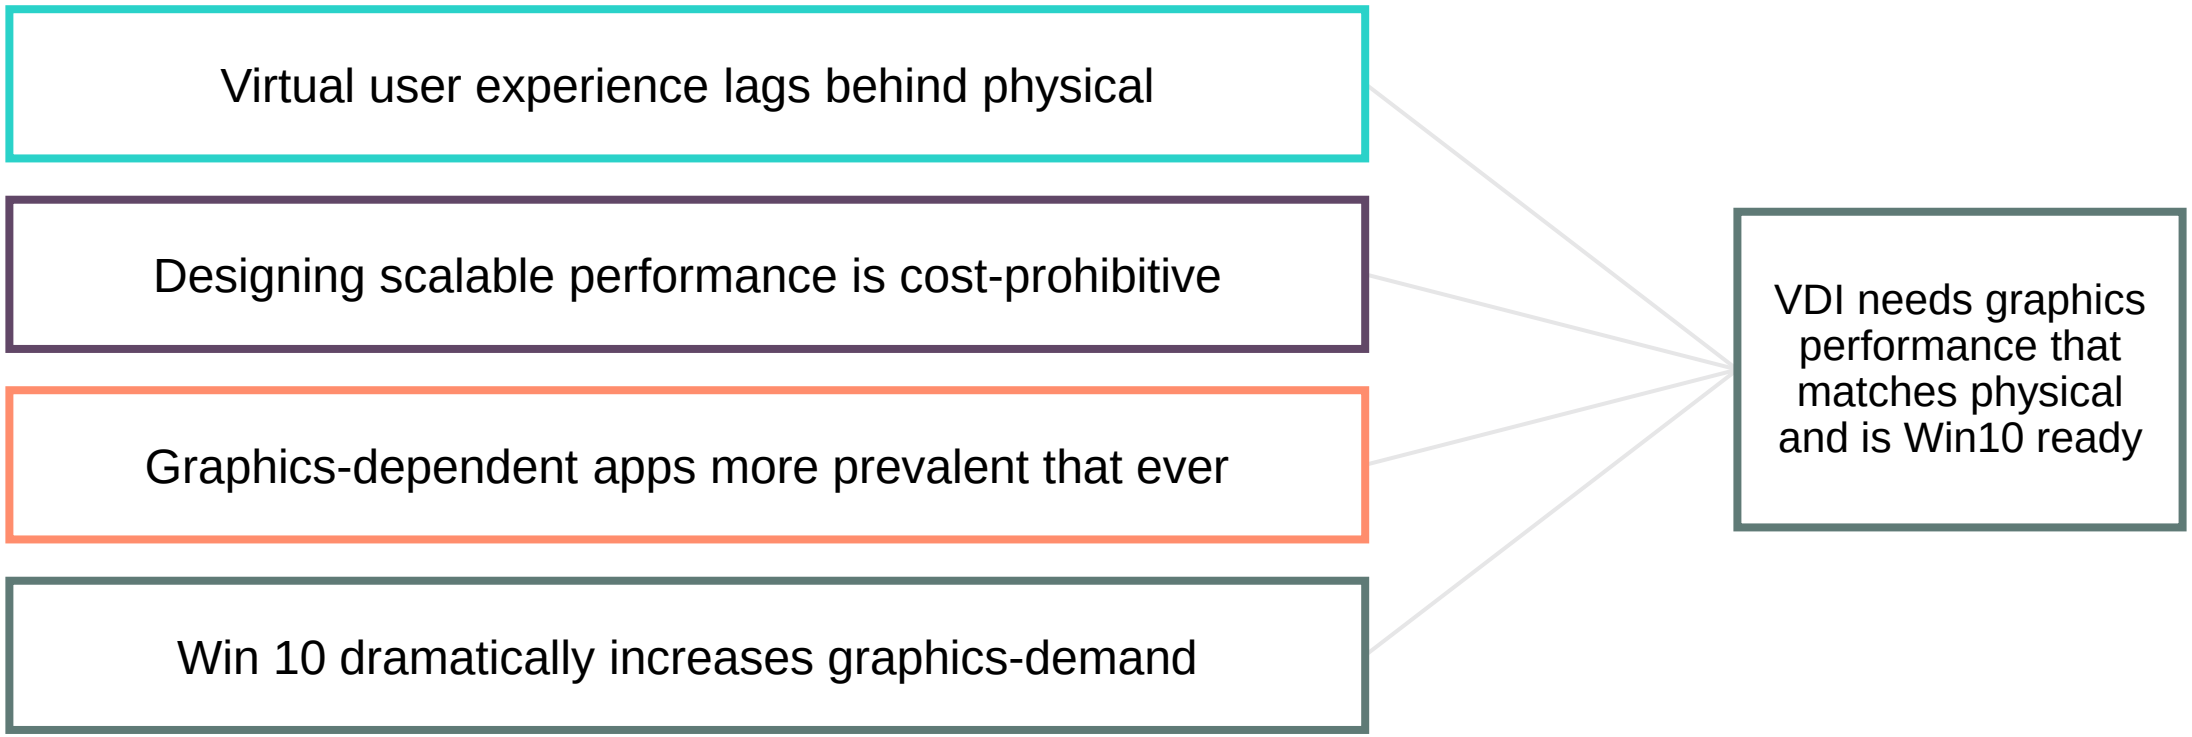


Hewlett Packard Enterprise | NVIDIA

# NVIDIA Virtual GPU Software

Hewlett Packard
Enterprise

# Virtualization for any workload

NVIDIA Virtual GPUs deliver the right level of performance for any use case

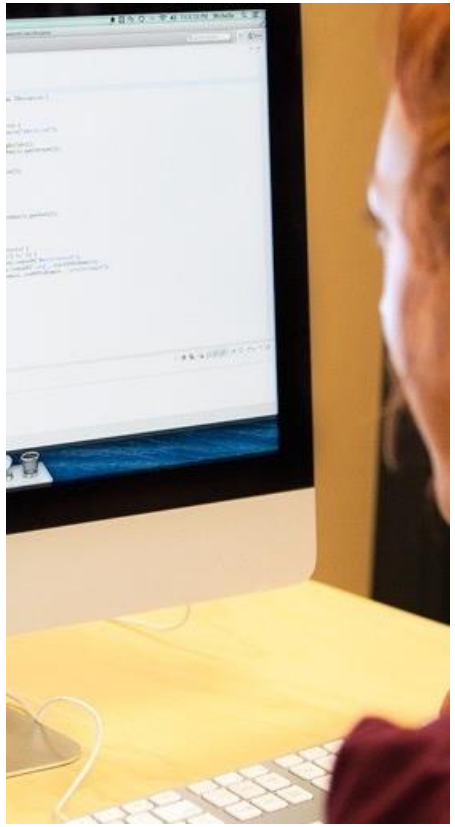



Hewlett Packard Enterprise | NVIDIA

# Performance from the Data center

NVIDIA Virtual GPU technology delivers graphics accelerated virtual desktops and applications

All devices have graphics

Virtual machines also need a GPU

Hewlett Packard Enterprise | NVIDIA

# How it Works

NVIDIA virtual GPU products deliver a GPU Experience to every Virtual Desktop

## CPU Only VDI

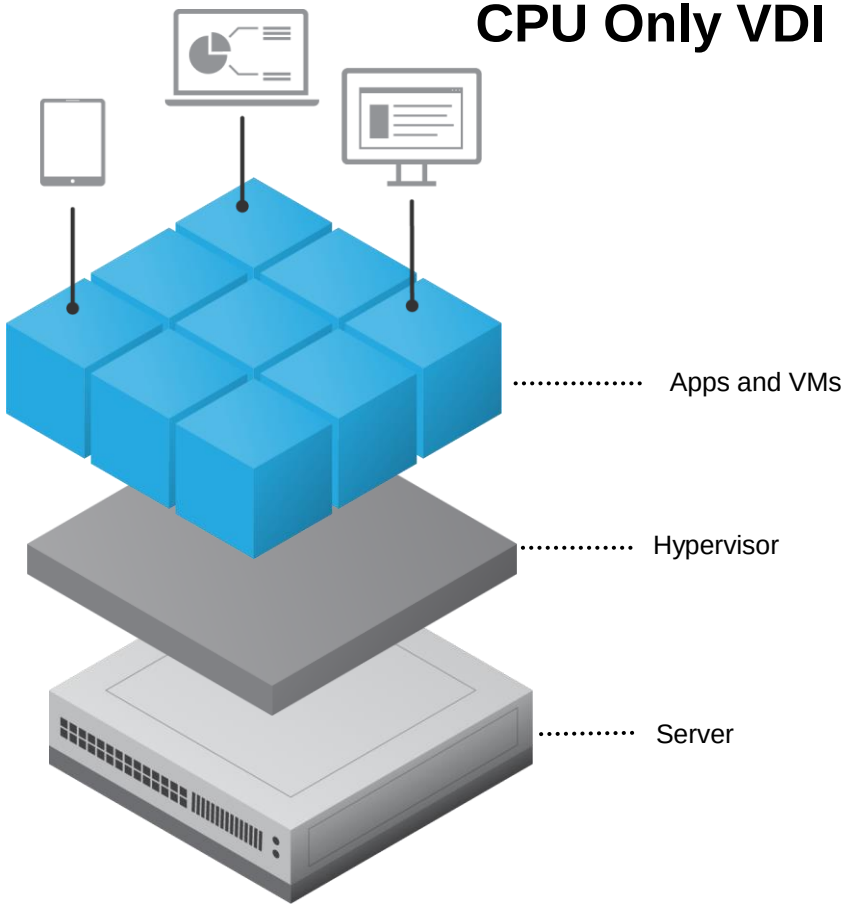


## With NVIDIA Virtual GPU

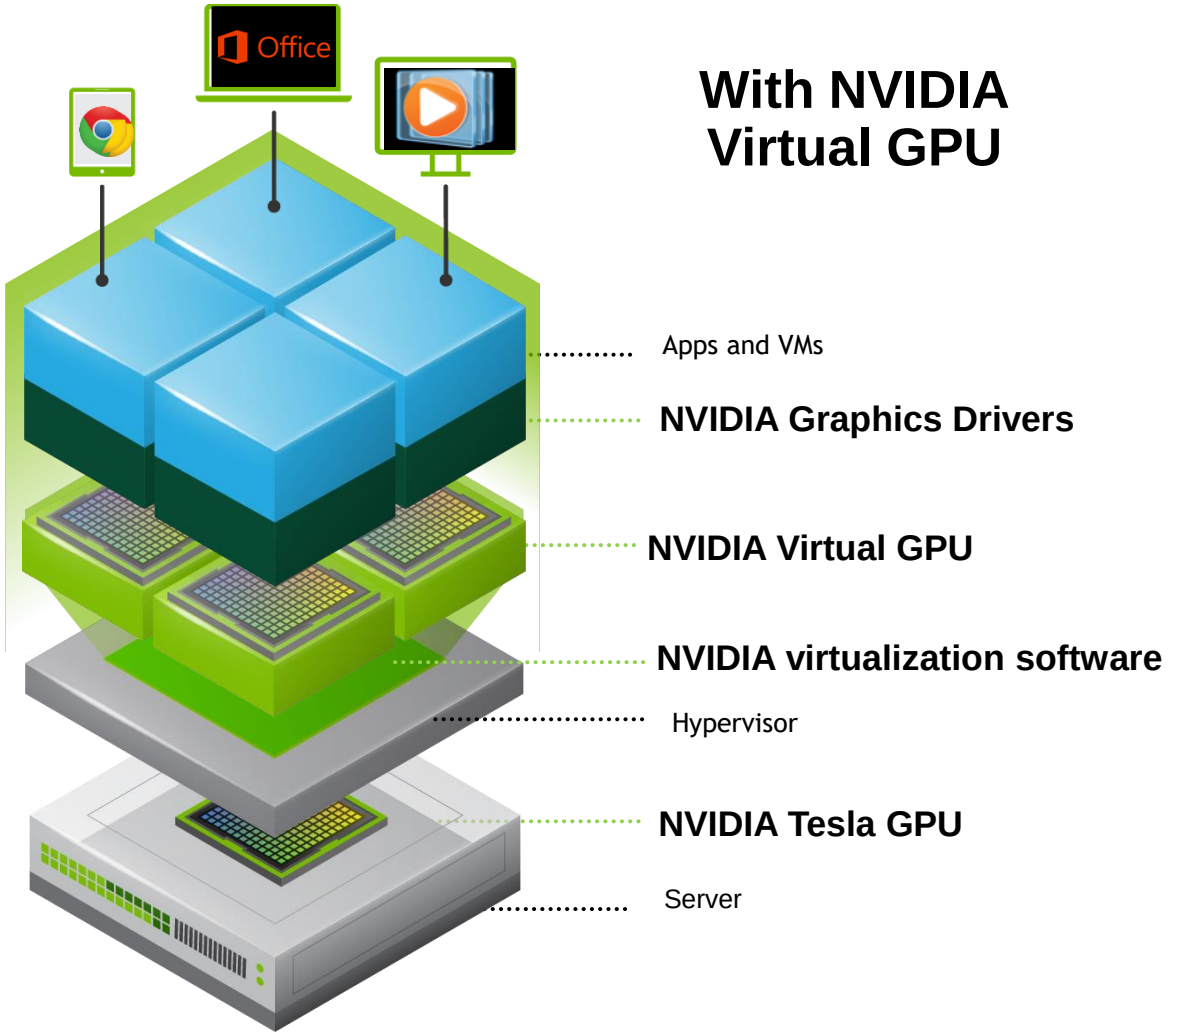


Hewlett Packard Enterprise | NVIDIA

# Accelerate productivity

## For Every User, Any App

| Knowledge Worker | Creative & Technical Professional |
| --- | --- |
| Providing business users the highest level of experience for all their apps on any device | Uncompromised experience for professional graphics users allowing them to design on the go |

Office, Ai, Ps, Id

AUTODESK AUTOCAD, SOLIDWORKS, Petrel, CATIA, AUTODESK REVIT, ArcGIS, PLM Software SIEMENS NX, AUTODESK MAYA, 3DEXCITE

Scalability ⟷ Performance

Hewlett Packard Enterprise | NVIDIA

# NVIDIA Virtual GPUs for all “App-Driven” Use Cases

Office Ai Ps Id AUTODESK AUTOCAD SOLIDWORKS PLM Software SIEMENS NX 3DEXCITE

AUTODESK REVIT CATIA ArcGIS AUTODESK MAYA Petrel

## Business PC Users

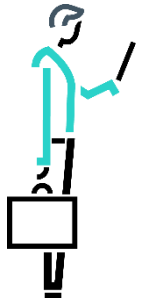

**NVIDIA GRID vPC/vApps Software**

High-quality user experience for office workers, that scales cost-effectively

## Workstation users

**NVIDIA Quadro Virtual Data Center Workstation Software**

Secure, mobile, high-performance experience for professional graphics users

# NVIDIA Tesla Gpus

## Recommended for Virtualization

| | V100 | P100 | P40 | P4 | M60 | M10 | M6 | P6 |
|---|---|---|---|---|---|---|---|---|
| **GPUs / Board (Architecture)** | 1 (Volta) | 1 (Pascal) | 1 (Pascal) | 1 (Pascal) | 2 (Maxwell) | 4 (Maxwell) | 1 (Maxwell) | 1 (Pascal) |
| **CUDA Cores** | 5,120 | 3,584 | 3,840 | 2,560 | 4,096 (2,048 per GPU) | 2,560 (640 per GPU) | 1,536 | 2,048 |
| **Memory Size** | 32 GB/16 GB HBM2 | 16 GB HBM2 | 24 GB GDDR5 | 8 GB GDDR5 | 16 GB GDDR5 (8 GB per GPU) | 32 GB GDDR5 (8 GB per GPU) | 8 GB GDDR5 | 16 GB GDDR5 |
| **H.264 1080p30 streams** | 36 | 36 | 24 | 24 | 36 | 28 | 16 | 24 |
| **vGPU Profiles** | ***1 GB, 2 GB, 4 GB, 8 GB, 16 GB, 32 GB*** | ***1 GB, 2 GB, 4 GB, 8 GB, 16 GB*** | 1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB | 1 GB, 2 GB, 4 GB, 8 GB | 0.5 GB, 1 GB, 2 GB, 4 GB, 8 GB | 0.5 GB, 1 GB, 2 GB, 4 GB, 8 GB | 0.5 GB, 1 GB, 2 GB, 4 GB, 8 GB | ***1 GB, 2 GB, 4 GB, 8 GB, 16 GB*** |
| **Form Factor** | PCIe 3.0 Dual Slot (rack servers) | PCIe 3.0 Dual Slot (rack servers) | PCIe 3.0 Dual Slot (rack servers) | PCIe 3.0 Single Slot (rack servers) | PCIe 3.0 Dual Slot (rack servers) | PCIe 3.0 Dual Slot (rack servers) | MXM (blade servers) | MXM (blade servers) |
| **Power** | 250W | 250W | 250W | 75W | 300W (225W opt) | 225W | 100W (75W opt) | 90W |
| **Thermal** | passive | passive | passive | passive | active/passive | passive | bare board | bare board |

***Performance***
Optimized

***Density***
Optimized

***Blade***
Optimized

Hewlett Packard Enterprise | NVIDIA

# NVIDIA Virtual GPUs and HPE Servers

# NVIDIA Virtual GPUs + HPE

Support, Updates, and Maintenance Subscription (SUMS)

Office, Ai, Ps, Id, Allscripts, Epic, Cerner, McKESSON

AUTODESK AUTOCAD, SOLIDWORKS, BIOVIA, Petrel, teraRecon, ArcGIS, CATIA, AUTODESK REVIT, AUTODESK MAYA, 3DEXCITE, PLM Software SIEMENS NX, Carestream, varian, FUJIFILM MEDICAL SYSTEMS, SECTRA

NVIDIA GRID vApps/vPC Software

NVIDIA Quadro vDWS Software

Hewlett Packard Enterprise

Servers

Synergy

Apollo

HPE ProLiant

HPE Edgeline Converged IoT

HPE Mission Critical Systems

HPE Qualified NVIDIA Tesla GPUs

V100

P100

P40

P4

M60

M10

M6

P6

An NVIDIA virtual GPU software license enables every virtual machine to get the benefits of a GPU just like a physical desktop has.

# SYNERGY VS RACK SERVER ARCHITECTURE

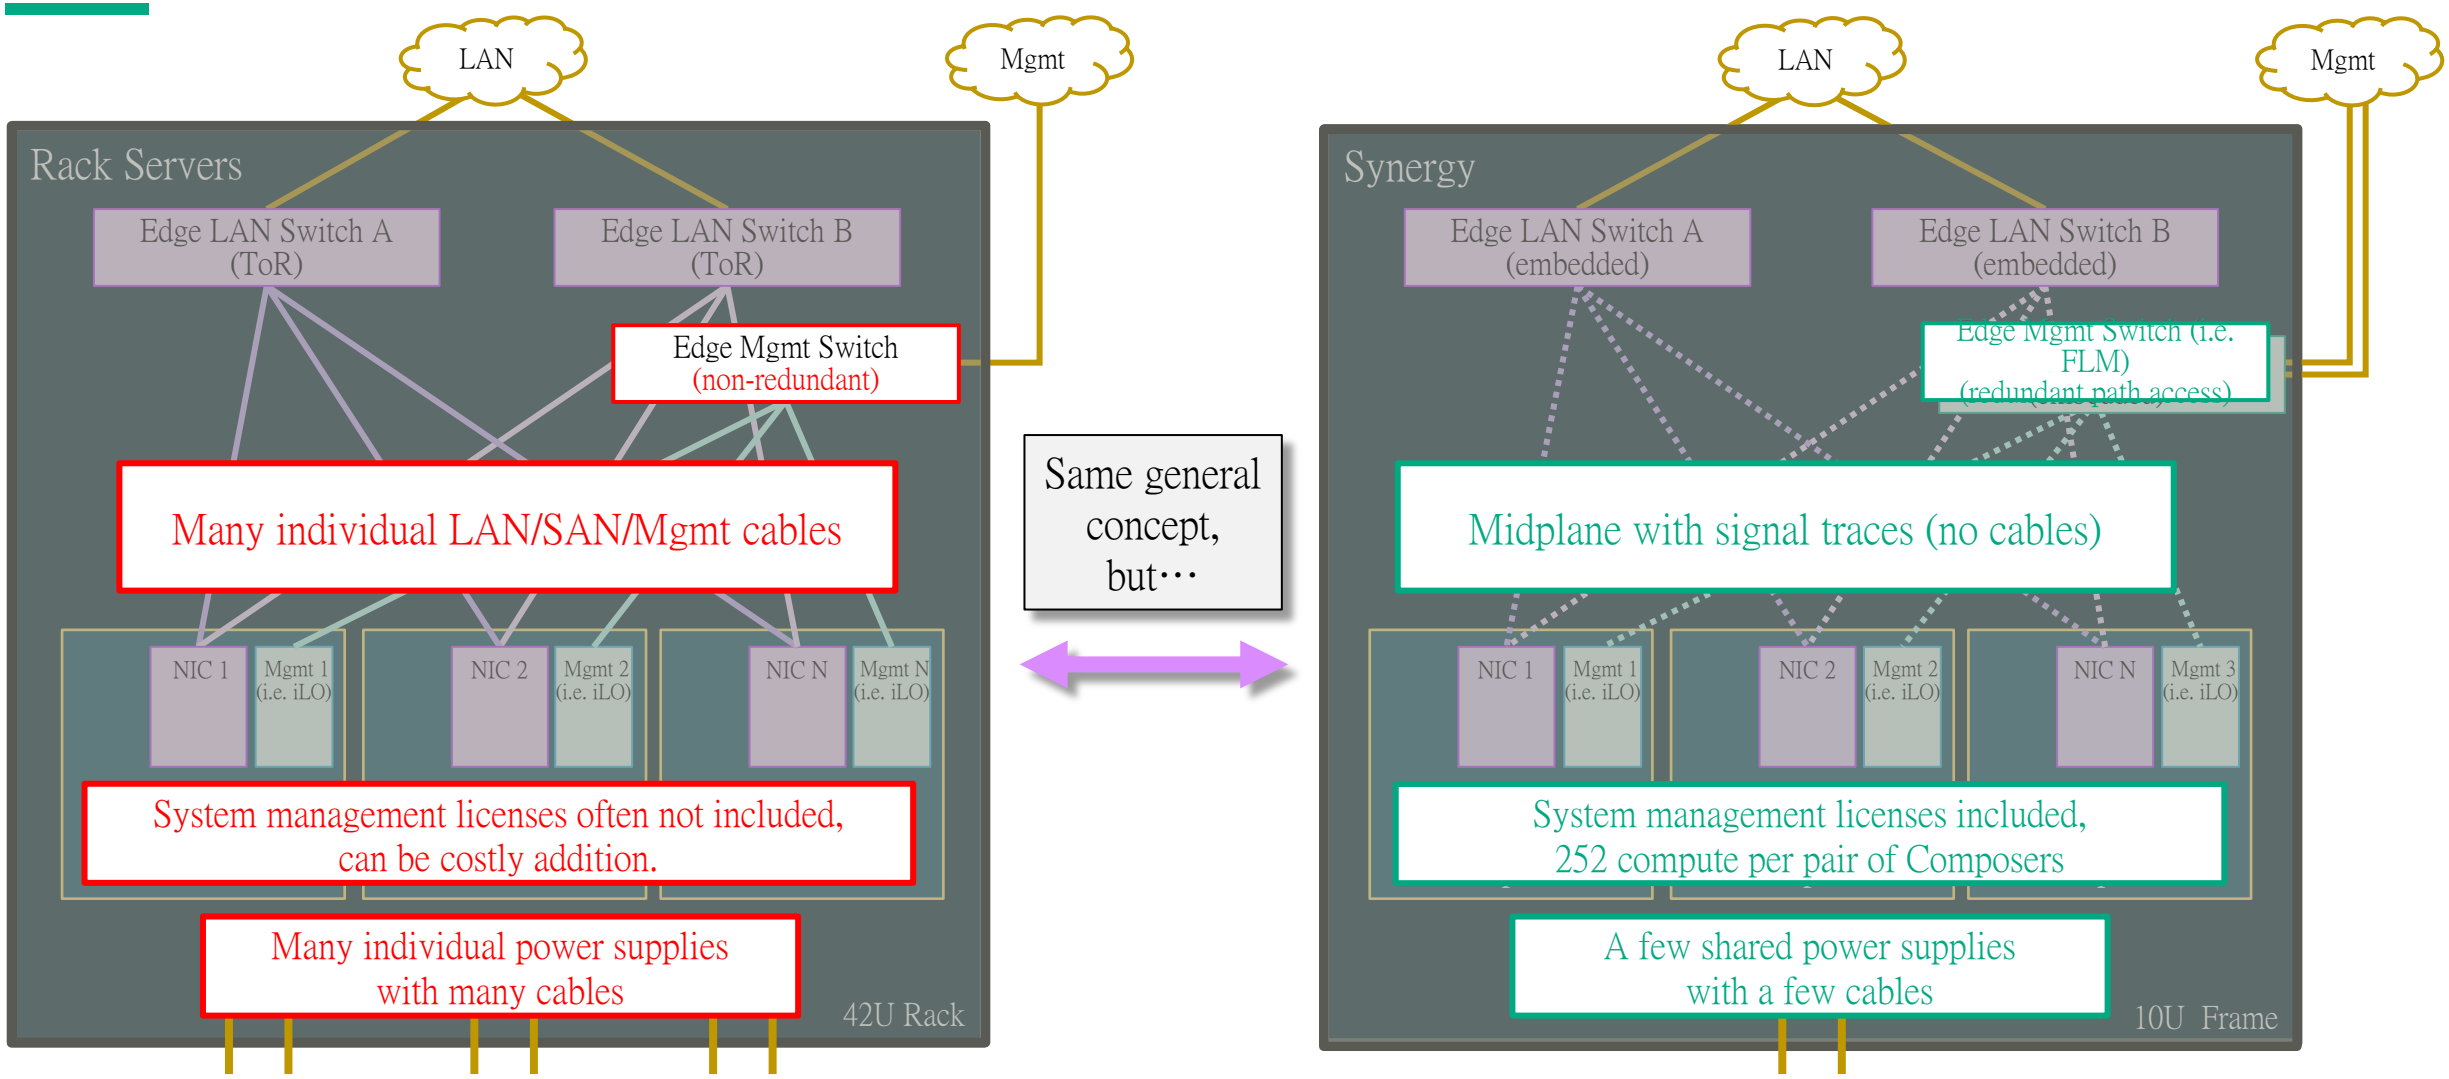

# SYNERGY CONSOLIDATION VS RACK SERVERS

## Rack Server (41U)

**ToR Switches**
- 2x LAN
- 2x SAN*
- 1x Mgmt

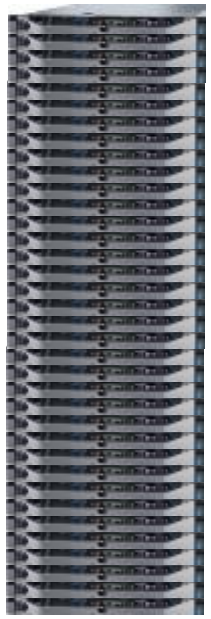

**36x 2P Servers**

| | | |
|---|---|---|
| 72 | Redundant LAN | 4 |
| 72 | Redundant SAN* | 0 |
| 82 | N+N Power | 18 |
| 36 | Management | 3 |
| **262** | **Cable Total** | **25** |

* Or an optional second LAN fabric at 72 vs 4 cables.

## HPE Synergy (36U)

**3x SY Frames**

# 90% CABLE REDUCTION / 27% RACK SPACE SAVINGS

# SYNERGY CABLE REDUCTION

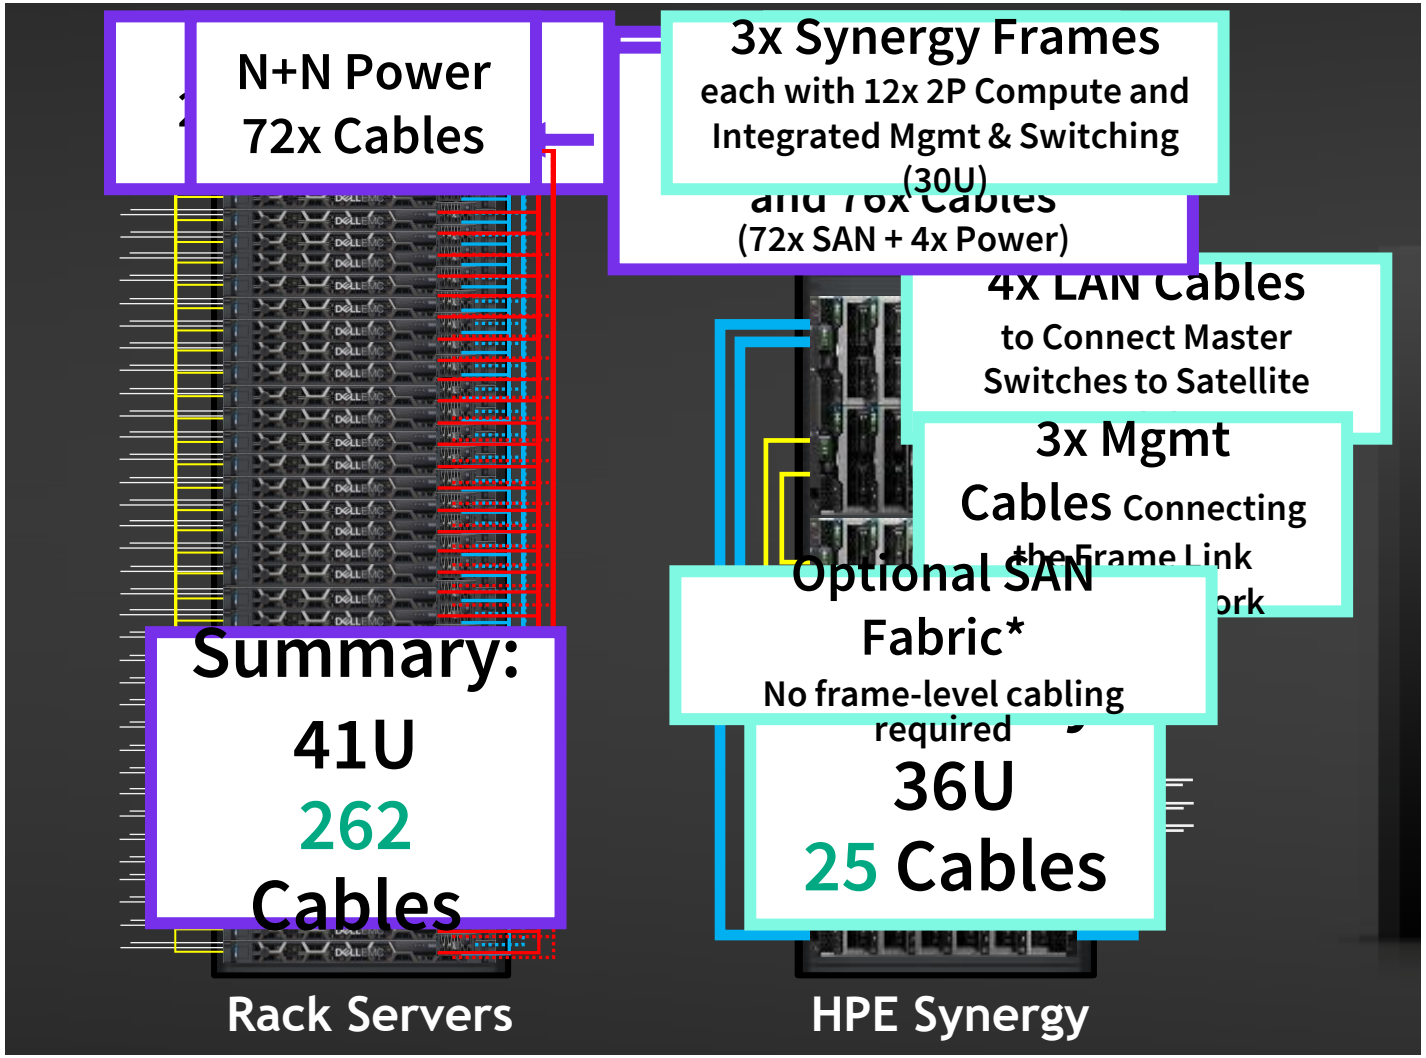


## Servers

- **Rack:** 36x individual 2P servers (36U)
- **Synergy:** 3x frames with integrated compute, management, and fabric (30U)

## Power

- **Rack:** N+N redundant, 72x cables
- **Synergy:** N+N redundant, 18x cables

## Management

- **Rack:** 1x ToR switch (1U) + 38x cables
- **Synergy:** 2x cables to enable management ring

## LAN

- **Rack:** 2x ToR switches (2U) + 76x cables
- **Synergy:** 4x cables for master-satellite fabric @ 25G (8 cables @ 50Gb)

## SAN Fabric (*or 2nd LAN)

- **Rack:** 2x ToR switches (2U) + 76x cables
- **Synergy:** SAN: No chassis-level cabling required
  LAN: 4-8x cables for master-sat fabric

# SYNERGY OPEX SAVINGS VS TRADITIONAL RACK

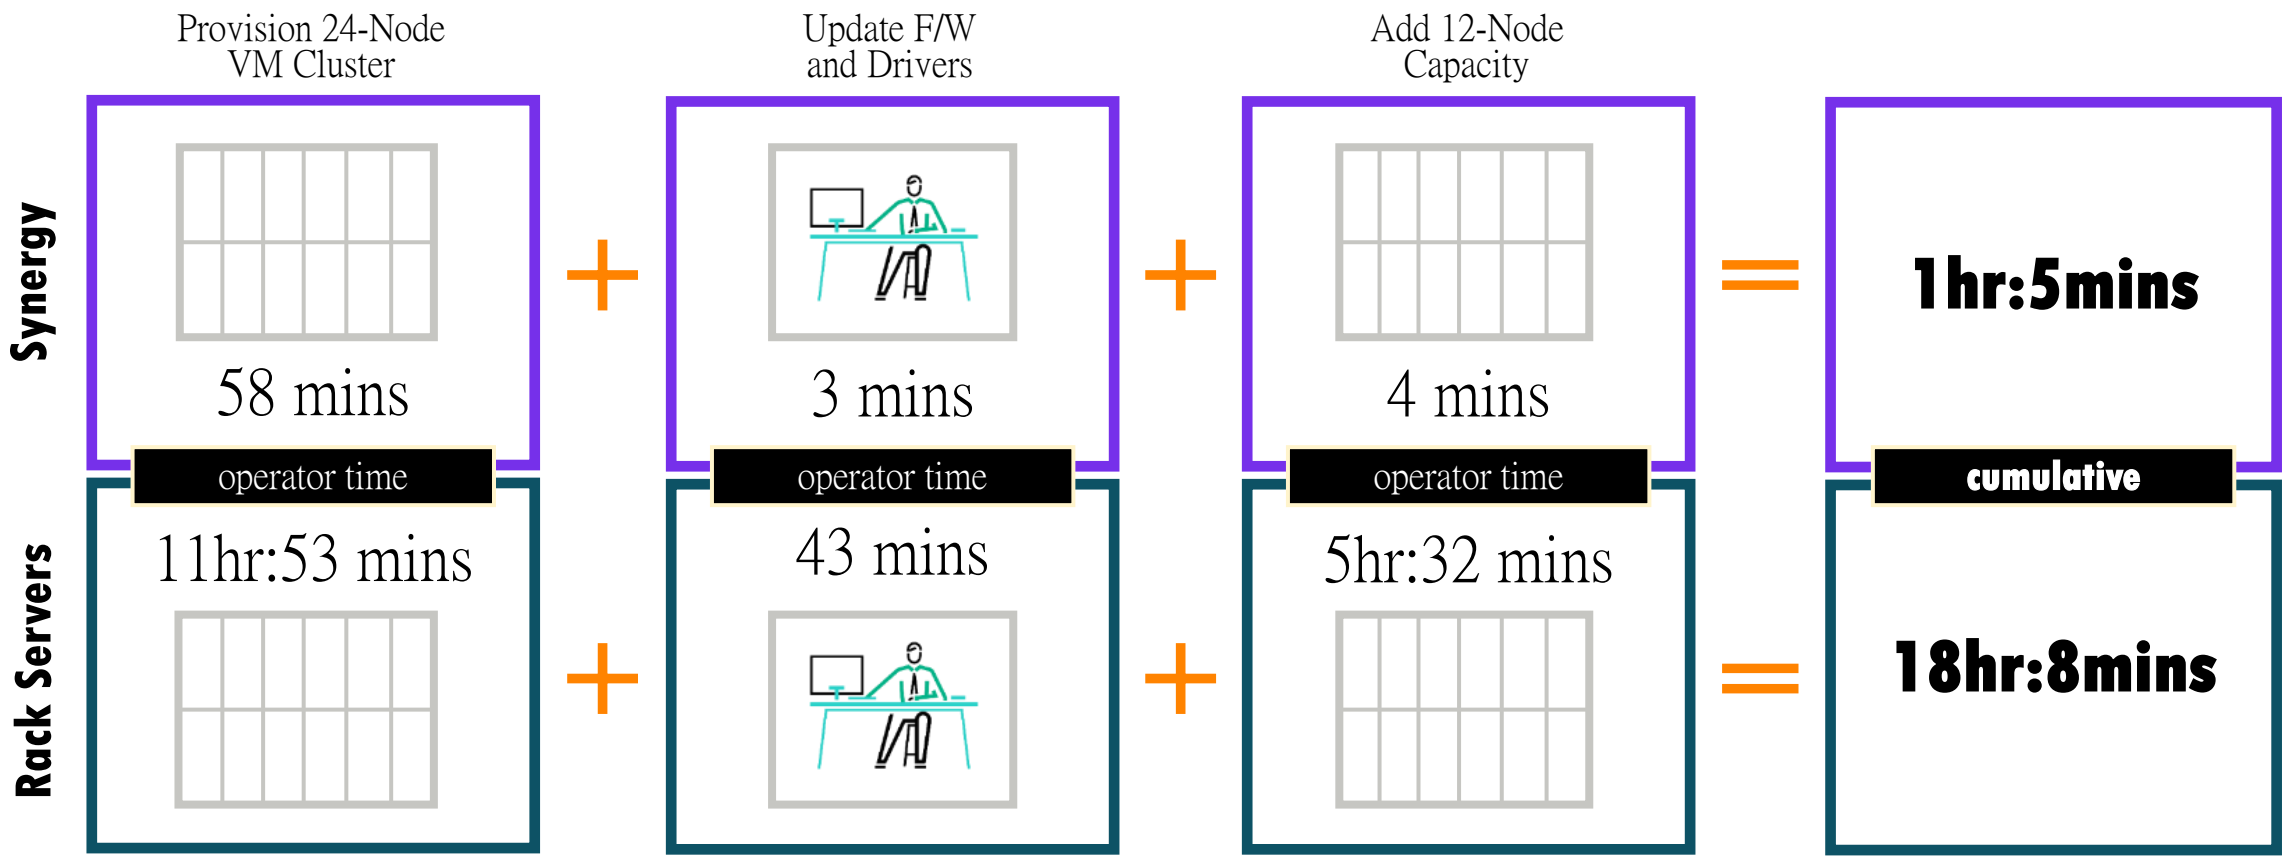


Synergy operator time savings (17x faster)

## WHAT USED TO TAKE 2 BUSINESS DAYS NOW TAKES AN HOUR

# ARCHITECTURAL ADVANTAGES TRANSLATE TO SAVINGS

## Synergy design = savings

Converged Infrastructure consolidating cabling, power, fans, and management

Low cost Composer2 with HPE OneView and iLO Advanced licenses for 21 frames (252 compute)

Economical master-satellite fabric architecture

Integrated Frame Link Module management port consolidation

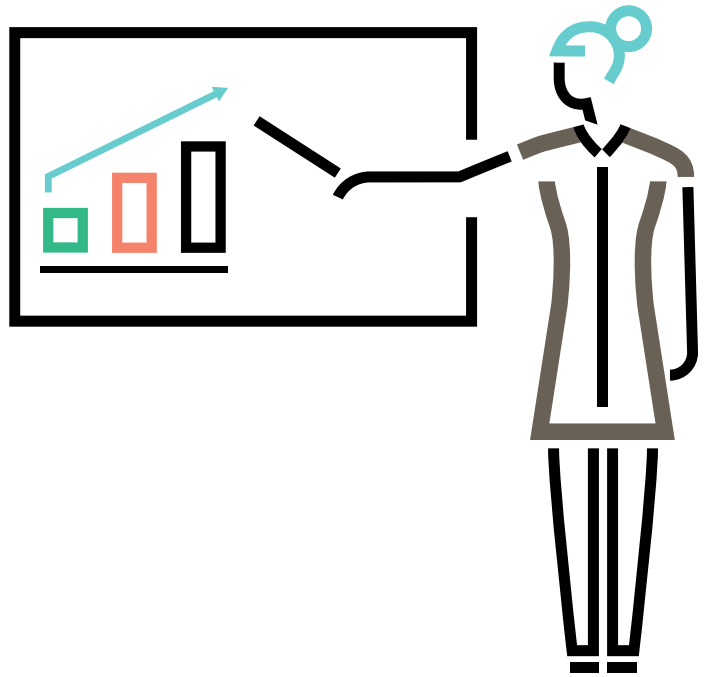

Hewlett Packard
Enterprise

# Thank you