多雲環境、容器支援、高效網路: 重新定義下一代超融合

Allen Tseng
產品技術經理
2018.07

傳統基礎架構的挑戰

複雜

Hampers speed, security, service delivery, scale

應用程式架構的發展

Microservices and distributed applications on/off premises

多雲架構

Developers demand choice, business demands economics

如何克服？

Hyperconverged Infrastructure

Cloud Like Resource Model

Speed, economics, simplicity for any application

速度,效能和簡化

Common platform across data center

Predictable, reliable performance

價格和成本

Pay-as-you-grow Economics with Resource Scaling On-Demand

操作簡化

Cloud like simplified management and orchestration

Cisco Approaches

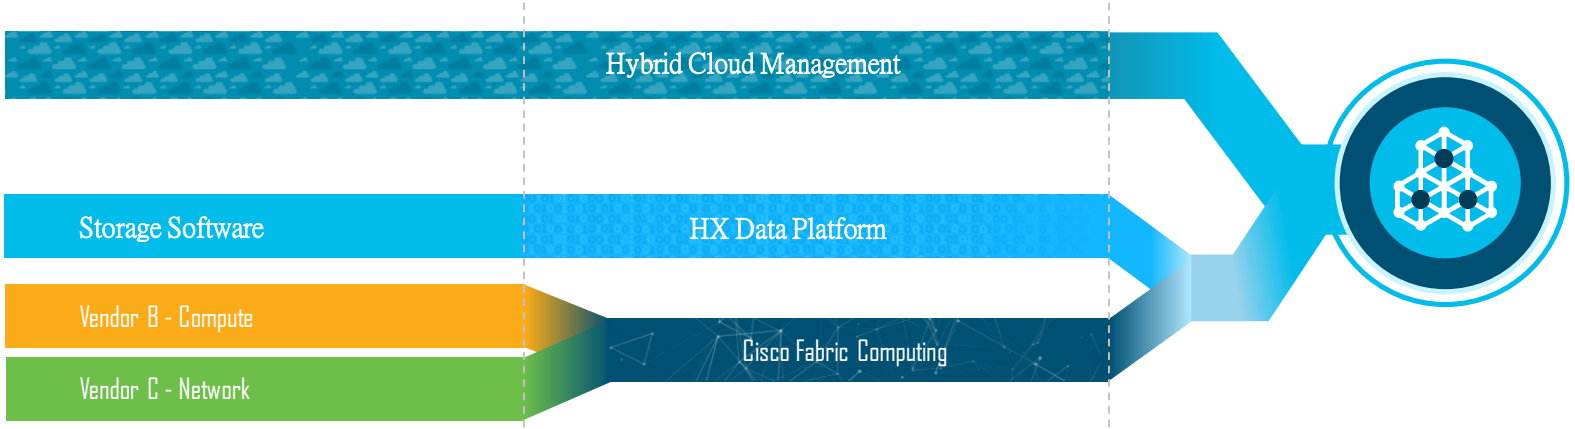

Traditional Infrastructure

Siloed Tiers of Compute, Storage and Networking from separate vendors

Siloed & Disaggregated Management

Cisco UCS Foundations

Converged Compute and Networking System

Cisco HyperFlex

Multicloud Era Hyperconverged Infrastructure

HyperFlex 架構差異

Hardware and Software Engineered Together

Pre-integrated UCS Rack and Blades + Cisco HX Data Platform

Single point of support backed by world class Cisco TAC Support

Complete Hyperconvergence with Integrated Network Fabric

Unified Network Infrastructure
Guaranteed QoS with low latency

Hyperflex

Simplified Policy Based Management

Simplified operations across compute, storage and network

Next Gen Distributed Data Platform

Enterprise Scale
Performance without hot spots
Extremely consistent IO

Hypervisor Choice ESXi and Hyper-V	Tetration ACI	UCS Intersight	CloudCenter AppDynamics CWOM	Security

HyperFlex Platform Evolution

- Data Distributed, Flash Friendly Foundation
- Turnkey Install, Easy Operations
- Tightly coupled with VMware

- Performance for Enterprise Apps
- Data protection, Native Replication
- Security, Encryption
- Edge/IoT

- Multi-Hypervisor
- Advanced Data Protection, Higher Scale, Capacity options
- Hybrid-cloud Ready, Multi-cloud Enabled
- Cisco Intersight Managed
- Enterprise Apps++

HYPERFLEX Multicloud Platform

NEW

Any App

Microsoft Hyper-V

Container FlexVolume Support

New Enterprise Application Validation

Any Cloud

Application Performance Monitoring

Application Placement

Cloud Mobility

IaaS and Private Cloud

Any Scale

HyperFlex Container Appliance
Cisco Google Hybrid Cloud

64 Node and Capacity Scaling w/ HA

Stretch Clusters

Cloud-Deployed w/ Intersight

Any Cloud

HyperFlex Flexible Multicloud Services

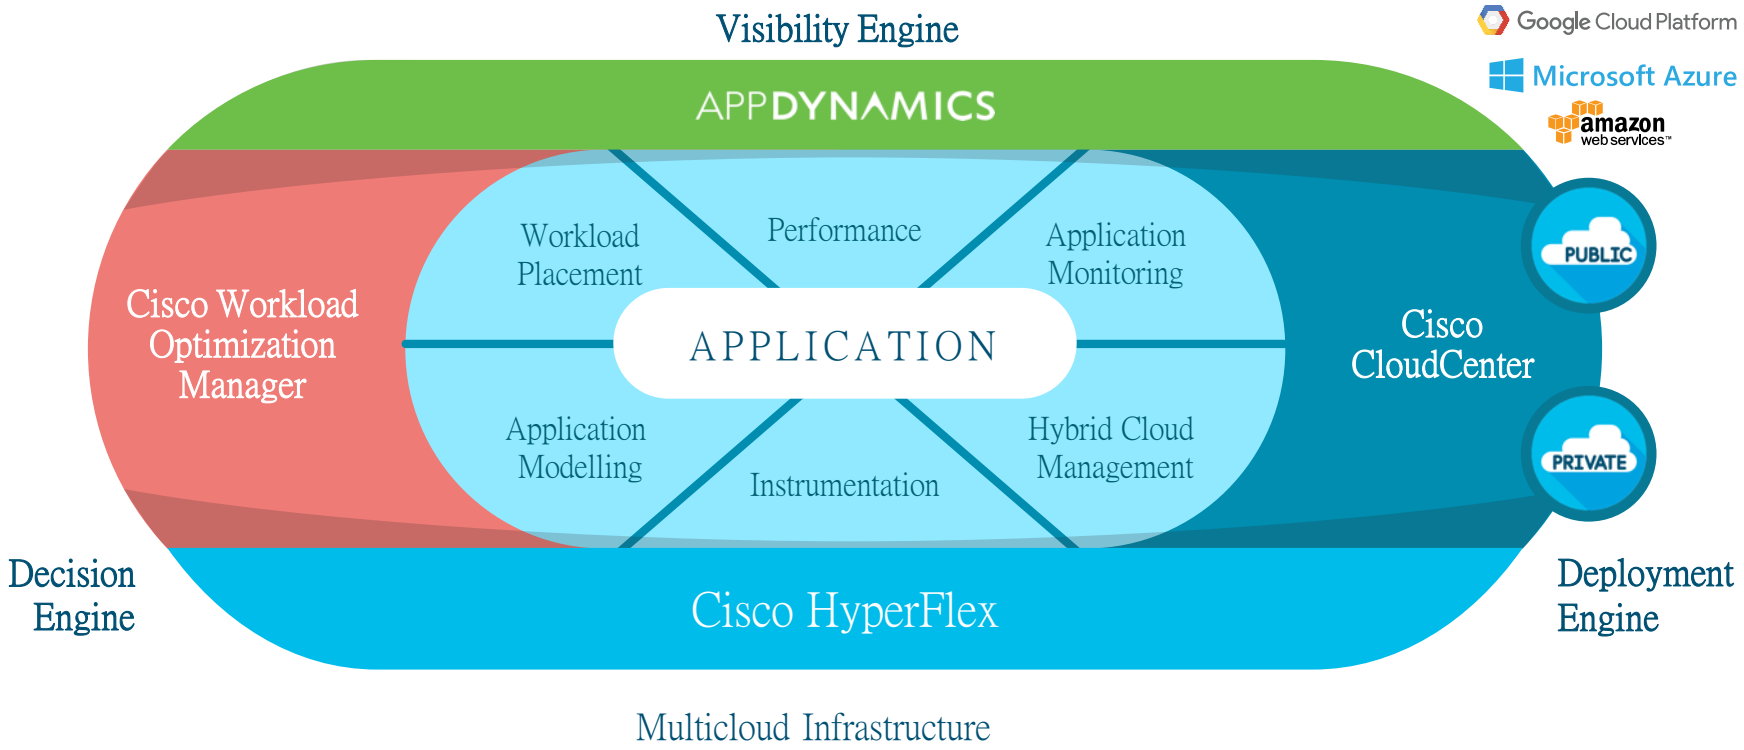

Multicloud Application Management
HyperFlex, CloudCenter & AppDynamics

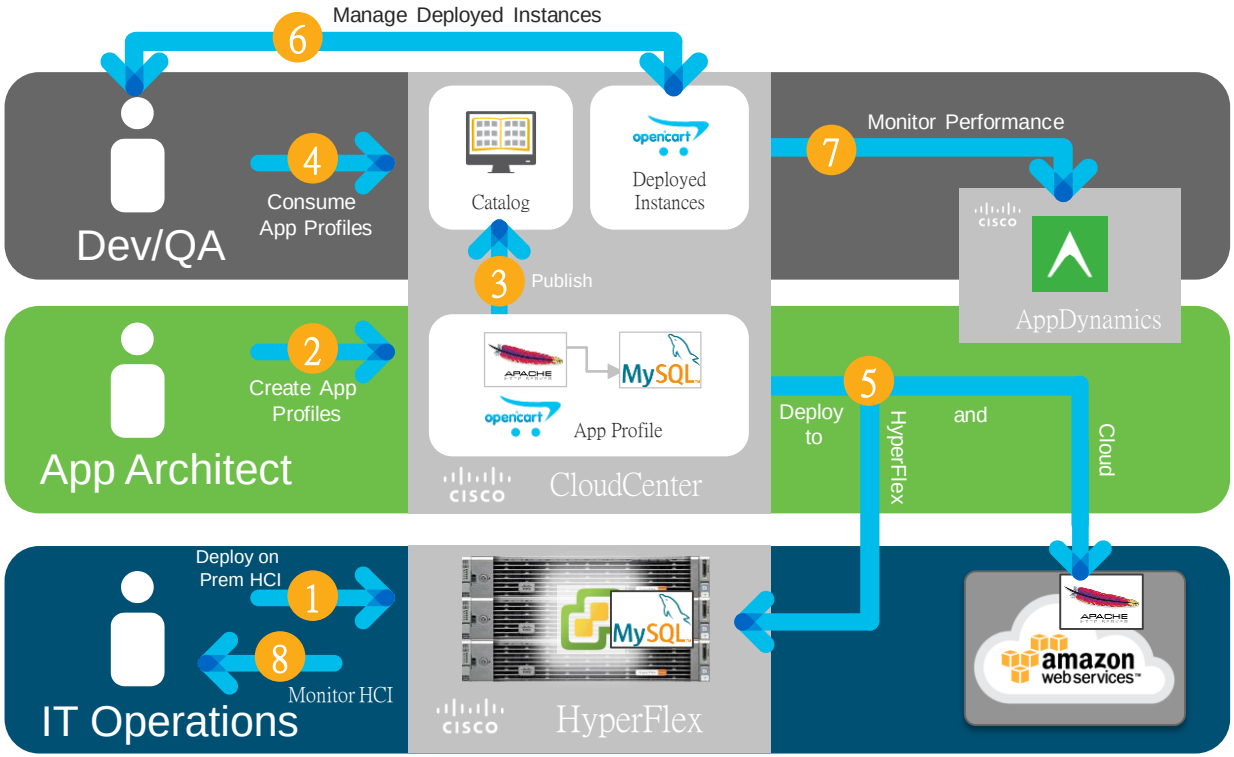

1. Deploy On Prem HCI
2. Create App Profiles
3. Publish App Profiles
4. Consume App Profiles
5. Deploy Applications
6. Manage Applications
7. Monitor Applications
8. Monitor HCI

HyperFlex 3.0 全面支持 Google Kubernetes 容器

Turnkey Kubernetes
- 簡介無縫的整合 K8S
- HyperFlex IaaS

企業級存儲
- 水準擴展HA高可用檔案系統
- 資料保護，高效彈性

企業級網路和安全
- 多租戶（Multi-tenant） 架構
- 微分區（Micro-segmentation）
- 安全性原則

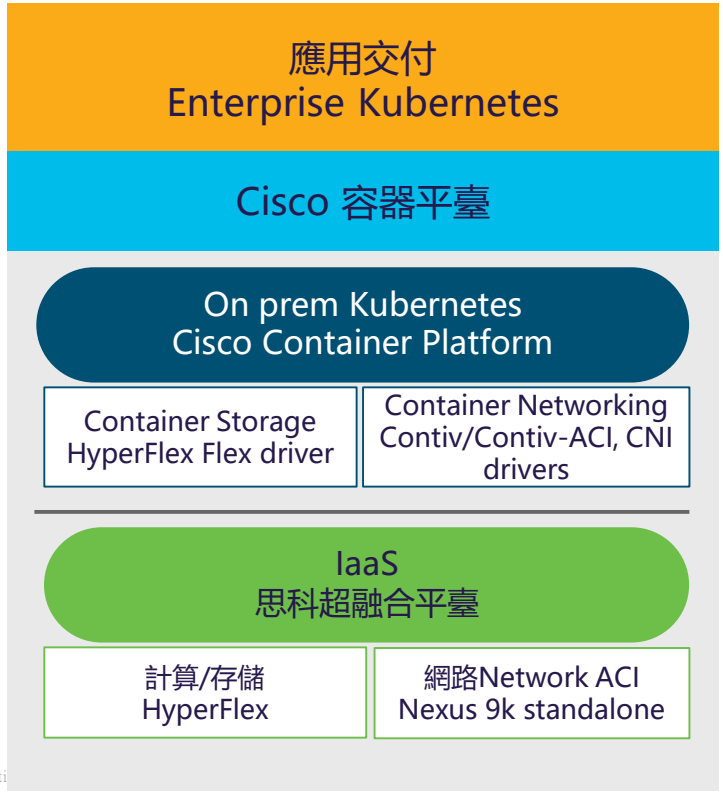

傳統和現代應用的通用平臺
- 虛擬機器和容器在同一平臺上共存

DevOps 就緒的 IT 架構
- 通過IT和安全性原則實現開發人員敏捷性
- 避免 IT 架構僵化

單一廠商支援
- 思科全球技術支援中心（TAC）全面支持
- 單一入口支援整體系統（entire stack）

HyperFlex 3.0 產品組成簡介

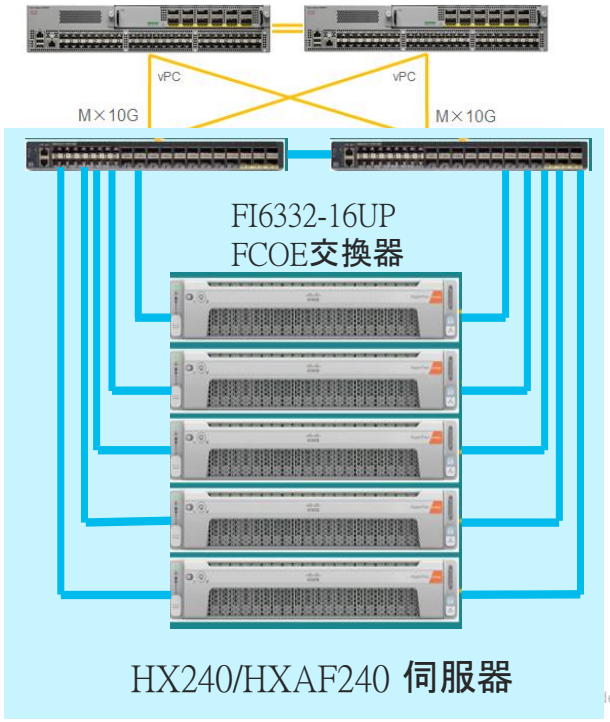

組件	數量	HX240	HXAF240（全快閃記憶體）
FCoE 交換機	2	FI6332-16up，40G 支持IP、FCoE、SAN 3.2微秒延遲，2.43Tbps頻寬	FI6332-16up，40G 支持IP、FCoE、SAN 3.2微秒延遲，2.43Tbps頻寬
UCS 超融合伺服器	3~64	2U機架伺服器 2顆CPU，24條記憶體，26個硬碟槽 CPU：Intel Scalable CPU 記憶體：128GB-3TB（DDR-2666） SSD：2*240GB+1*1.6TB 硬碟：6~23塊 10KRPM/7.2KRPM 硬碟 網口：2個40G FCoE埠 可虛擬出256個IP或SAN埠	2U機架伺服器 2顆CPU，24條記憶體，26個硬碟槽 CPU：Intel Scalable CPU 記憶體：128GB-3TB（DDR-2666） SSD：2*240GB+1*400GB 硬碟：6~23塊 SSD 網口：2個40G FCoE埠 可虛擬出256個IP或SAN埠
HXDP 分散式存儲軟體	3~64	支持 VMware、Hyper-V 虛擬化 支持 Google Kubernetes 容器 支持2/3副本資料保護 支援記憶體級線上去重壓縮 支援連接外置 SAN 存儲 支援全虛擬化網路 橫向獨立擴展計算、存儲、網路資源 支援遠端容災 支持 Stretch Cluster 雙活	支持 VMware、Hyper-V 虛擬化 支持 Google Kubernetes 容器 支持2/3副本資料保護 支援記憶體級線上去重壓縮 支援連接外置 SAN 存儲 支援全虛擬化網路 橫向獨立擴展計算、存儲、網路資源 支援遠端容災 支持 Stretch Cluster 雙活

HyperFlex 全球成功案例，2800+

Financial Services

Government


Education

Software / Services

Medical/ Pharmaceutical

Miscellaneous

HyperFlex價值

部署簡單！

使用簡單！

維護簡單！

From Cabling to Your Data Center Organization – UCS Simplifies

What does your data center organization look like?

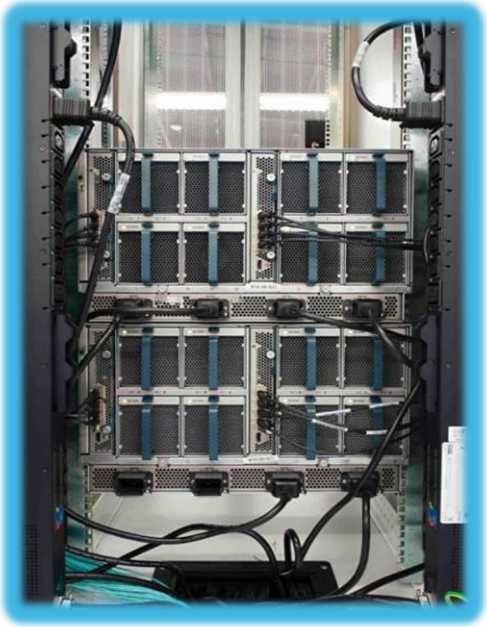

超融合架構的融合網路

1. Network Integrated HCI
2. Policy Based Management
3. End-to-End HCI Automation

HX220C
High Performance in Small Footprint
(Databases, VDI, VSI)

HX240C
High Density with Performance
(Databases, VDI, VSI, Test/Dev)

HX220C & EDGE
Smallest Footprint
(VDI, VSI, ROBO)

HX240C
Capactiy-Heavy
(VDI, VSI, Test/Dev)

40 GBPS / 10 GBPS
UCS FABRIC NETWORKING
90+ Nodes, Single Fabric,
Single Management Domain

HYBRID CLUSTER ALL FLASH CLUSTER

Cisco HyperFlex **All Flash M5** Configurations

HXAF220c M5 Edge Cluster	**HXAF220c M5** Cluster	**HXAF240c M5** Cluster	**HXAF M5 + Compute Node** Clusters
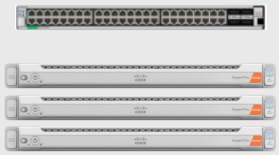	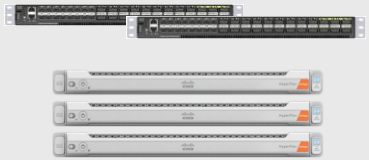	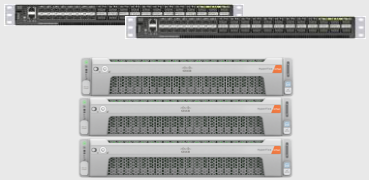	 HX M5 Cluster + HX Compatible M4 & M5 Compute Nodes
*3.6TiB – 38.5TiB	**4.8TiB – 274.2TiB	**4.8TiB – 788.3TiB	NOTE: Consult Release Notes for Compute Node Support Details
Smallest Footprint 3 Node Cluster (VSI, ROBO)	Smallest Footprint 3–32 Node Cluster (VSI, VDI, DB, ROBO)	Capacity-Heavy 3–32 Node Cluster (VDI, VSI & DB Workloads)	Compute-Heavy Hybrid (Compute Bound Apps/VDI)
Per-Node 1 x Cache SSD 3-8 x 960GB or 3.84TB Capacity HDDs	Per-Node 1 x Cache SSD 6-8 x 960GB or 3.84TB Capacity HDDs SED Options Orderable	Per-Node 1 x Cache SSD 6-23 x 960GB or 3.84TB Capacity HDDs Up to 2 x GPUs SED Options Orderable	3-32 HXAF220 or HXAF240 Node Cluster + Up to 32 Compute Nodes Blade or Rack Local Disk, SD Card or SAN Boot

*Edge does not support RF3. Usable capacity w/ RF2 before compression and depuplication

**Usable capacity w/ RF3 before compression and depuplication

Use HX sizer for details: https://hyperflexsizer.cloudapps.cisco.com/

Cisco HyperFlex **HYBRID M5** Configurations

HX220c M5 Edge Cluster	**HX220c M5** Cluster	**HX240c M5** Cluster	**HX240c M5 LFF** Cluster	**HX M5 + Compute Node** Clusters
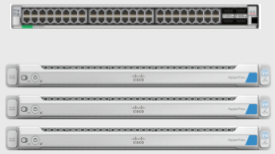	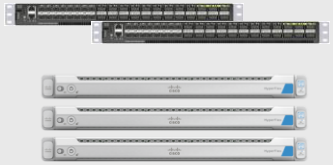			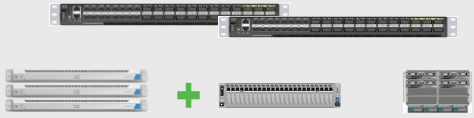 HX M5 Cluster + HX Compatible M3,M4 & M5 Compute Nodes
*4.5TiB – 12.0TiB	**6.0TiB – 128.5TiB	**6.0TiB – 369.5TiB	**30.1TiB – 214.2TiB	NOTE: Consult Release Notes for Compute Node Support Details
Smallest Footprint 3 Node Cluster (VSI, ROBO)	Smallest Footprint 3–32 Node Cluster (VDI, ROBO)	Capacity-Heavy 3–32 Node Cluster (VDI & VSI Workloads)	Capacity-Heavy 3–8 Node Cluster (high capacity Workloads)	Compute-Heavy Hybrid (Compute Bound Apps/VDI)
Per-Node 1 x Cache SSD 3-8 x 1.2TB Capacity HDDs	Per-Node 1 x Cache SSD 6-8 x 1.2TB or 1.8TB Capacity HDDs SED Options Available	Per-Node 1 x Cache SSD 6-23 x 1.2TB or 1.8TB Capacity HDDs Support up to 2 GPUs SED Options Available	Per-Node 1 x Cache SSD 6-12 x 6TB or 8TB Capacity HDDs Support up to 2 GPUs	3-32 HX220 or HX240 Node Cluster + Up to 32 Compute Nodes Blade or Rack Local Disk, SD Card or SAN Boot

*Edge does not support RF3. Usable capacity w/ RF2 before compression and depuplication

**Usable capacity w/ RF3 before compression and depuplication

Use HX sizer for details: https://hyperflexsizer.cloudapps.cisco.com/

快速且簡單的安裝

STEP 0

VMware and HX Data Platform Software Installed on Servers at Factory/Reseller Before Shipping to Customer

STEP 1

Rack Up Servers, Power on, and Add to vCenter

Drag and Drop Configuration (JSON) File

Create, Cluster, and Datastore

60 MINUTES OR LESS

Start Provisioning VMs

UCS 硬體架構 with HyperFlex
全自動部署網路和運算節點

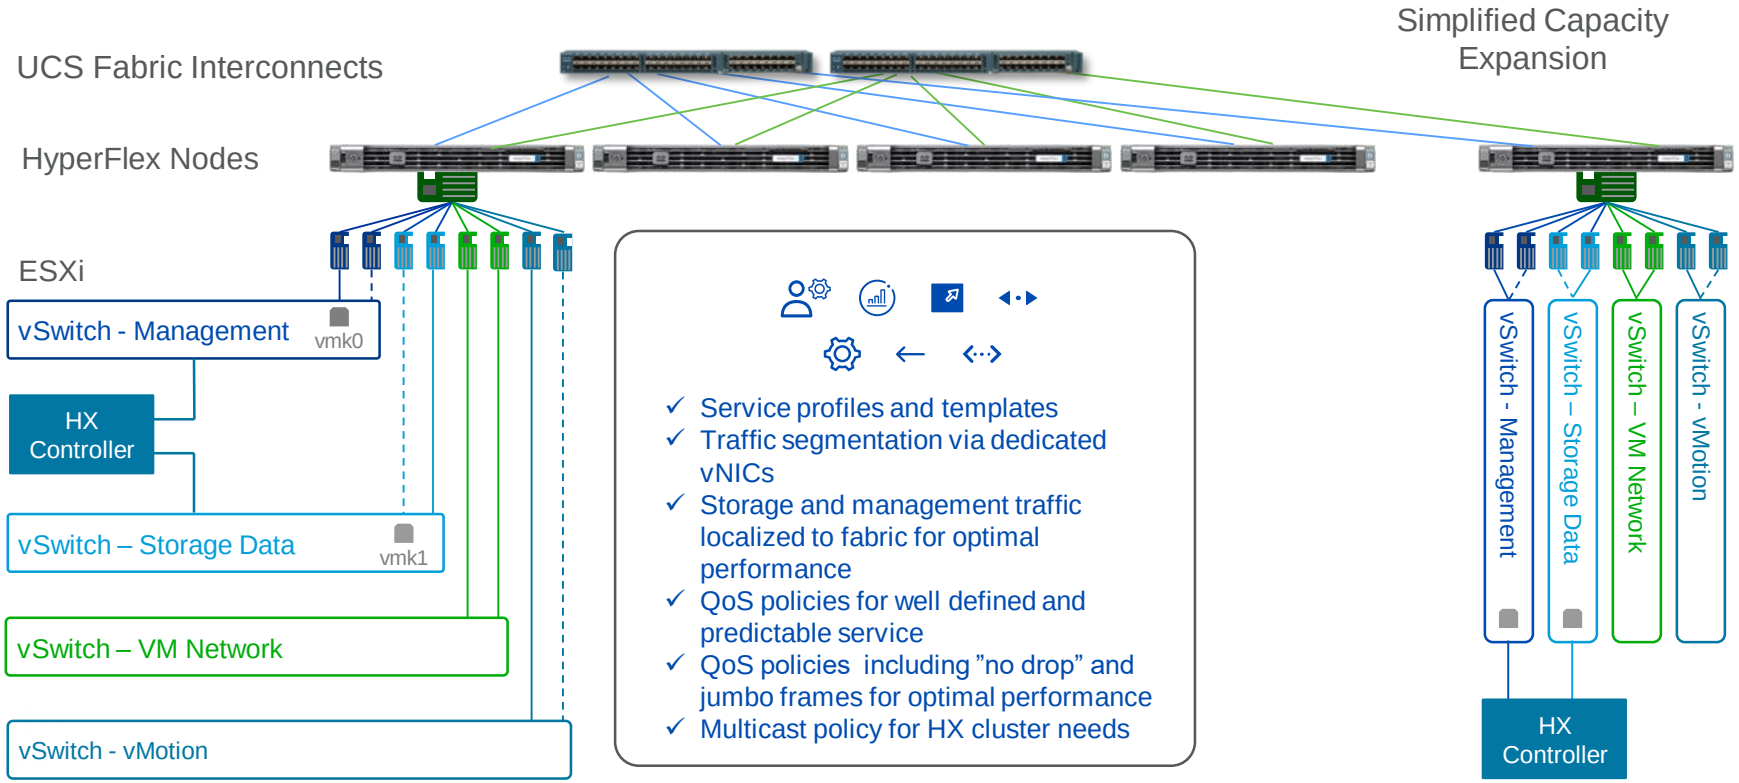

新一代資料平台

可訂製式系統，日誌結構檔案系統配合快閃記憶體快取，以重複資料刪除和壓縮作為基礎要素

可獨立擴充並橫向擴充的架構

可預期的擴充，隨成長付費

具備企業級儲存特徵

基於指標技術的快照

快速複製

重複資料刪除和壓縮

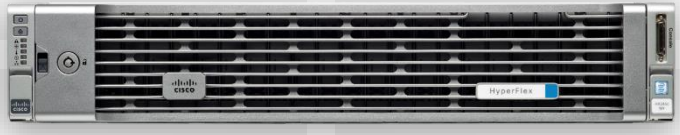

可用性高/自愈性強

單鍵滾動式更新無縫升級

可選購Call Home和7x24小時現場支援服務

企業資料保護

與vCenter整合

強大的的報告和分析

單一管理介面

動態分散資料

建構於傳統檔案系統之上，先寫入本地儲存然後再複製，造成性能瓶頸

- HX資料平台把資料同時分散儲存到各個節點，利用所有SSD的快取層達成快速寫入
- 儲存空間利用更平均， 虛擬機器遷移後無需遷移資料

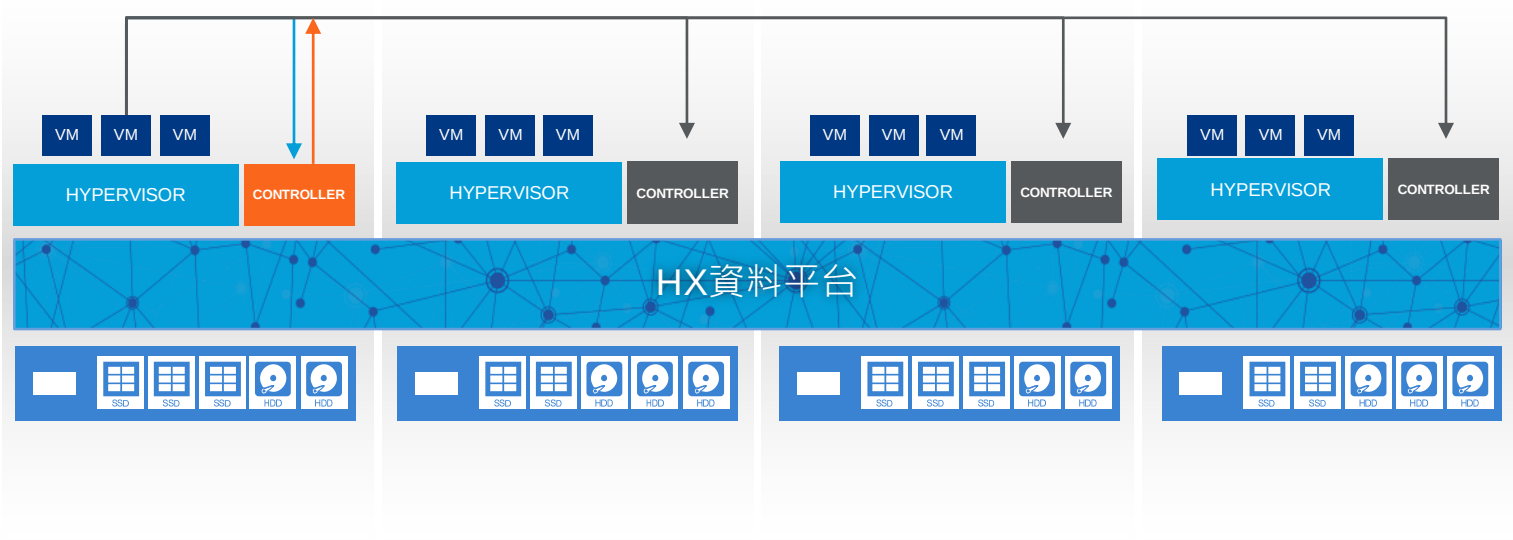

橫向擴充及分散式檔案系統

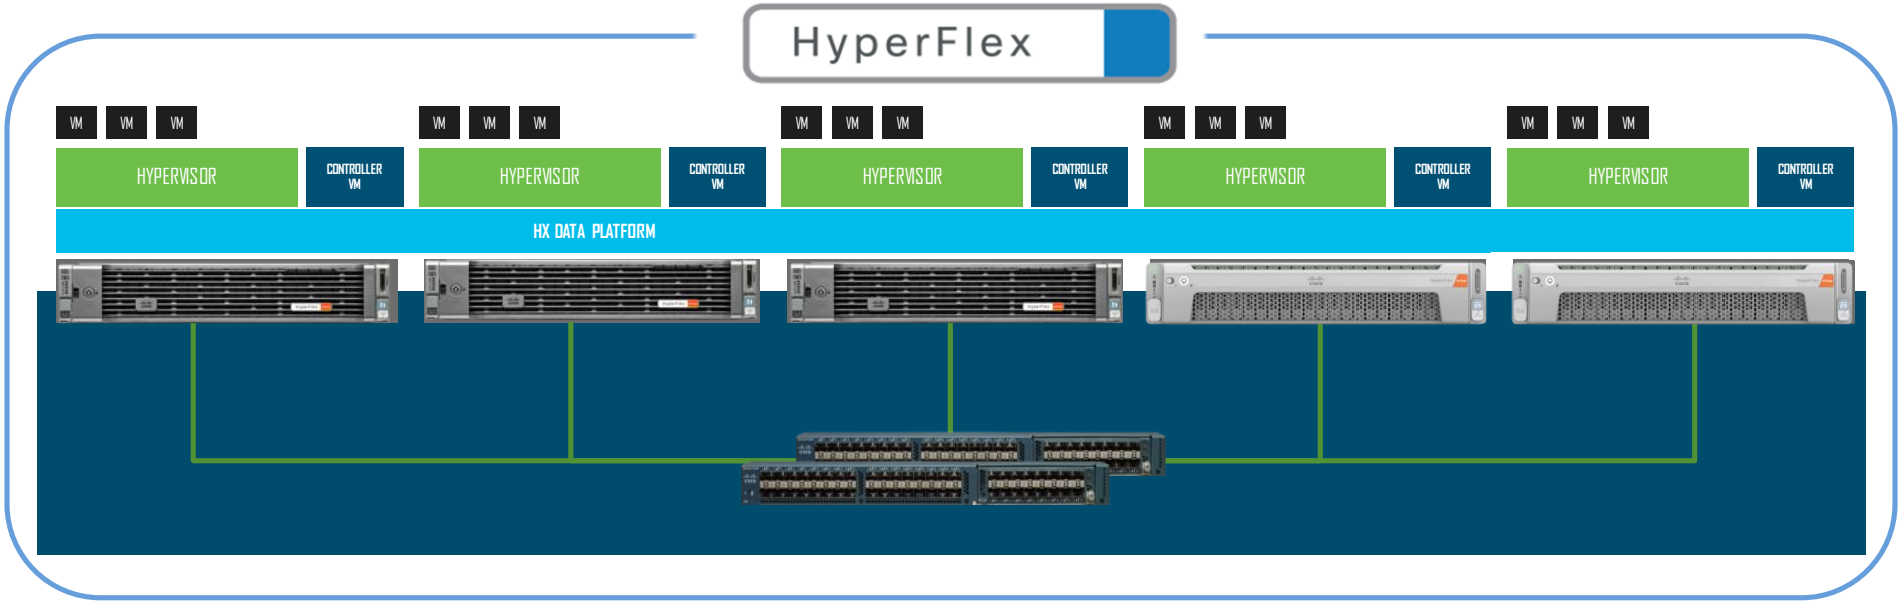

運算節點和儲存容量可獨立擴充

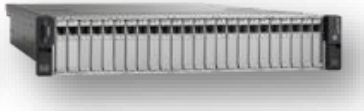

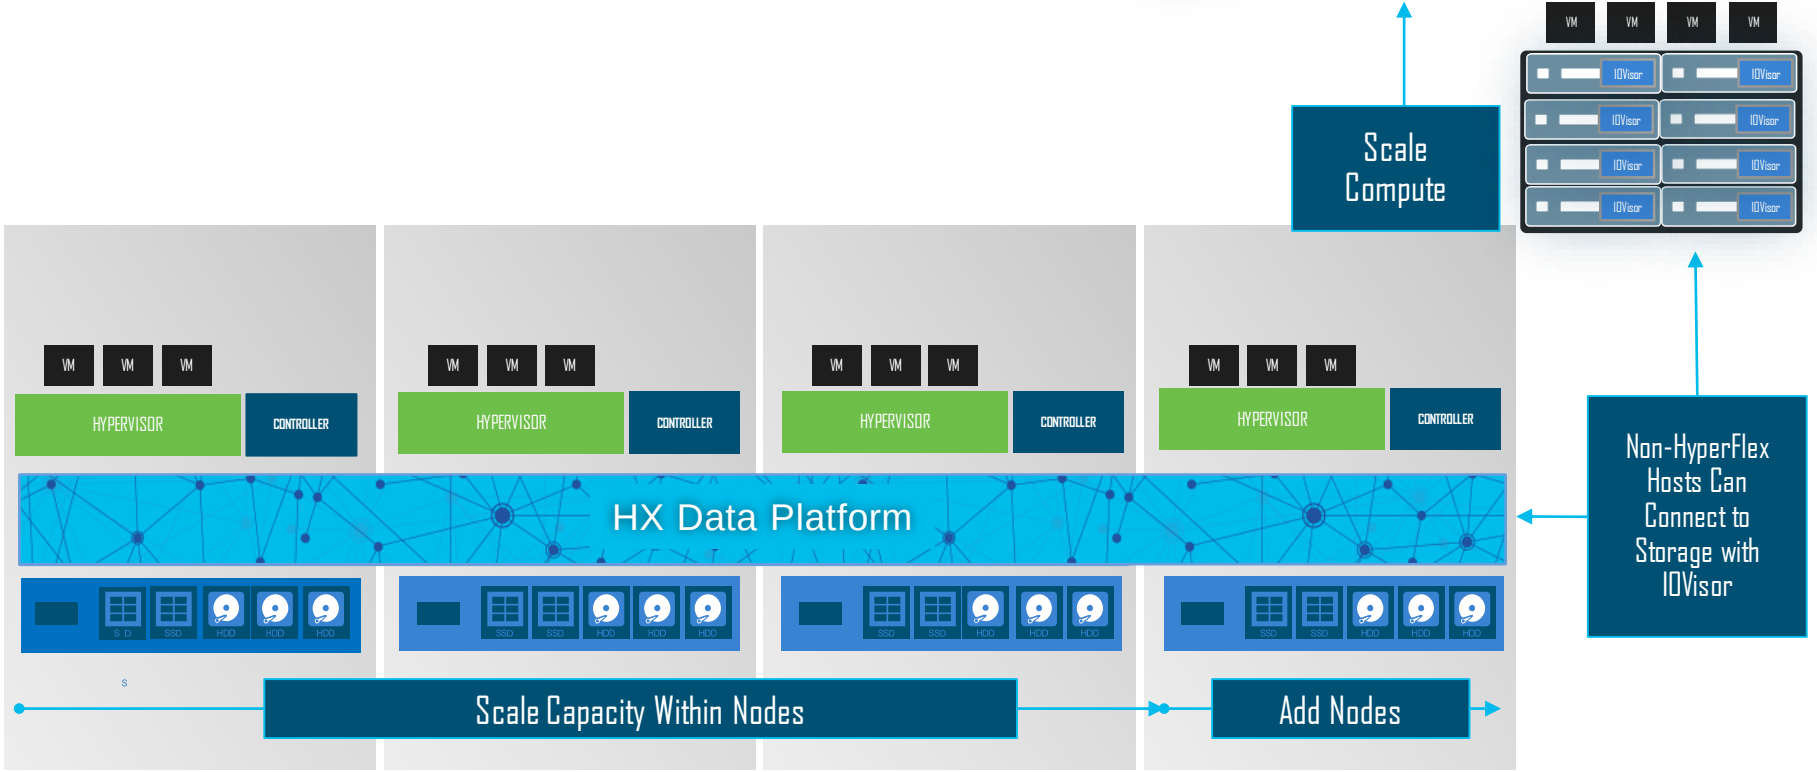

HX Cluster進階整合外接儲存

- HX Converged node supports mounting external storage arrays via (NFS, iSCSI, FC) Example uses:
  - Storage array must be on Cisco HCL
  - VM migration and data transfer
  - Co-existence with Existing SAN environment
  - Present RDM from FC SAN to VMs
  - Use VMware storage vMotion for migrations over Ethernet
  - Use for backup and other application support
- HX Datastore is not exportable outside the HX cluster

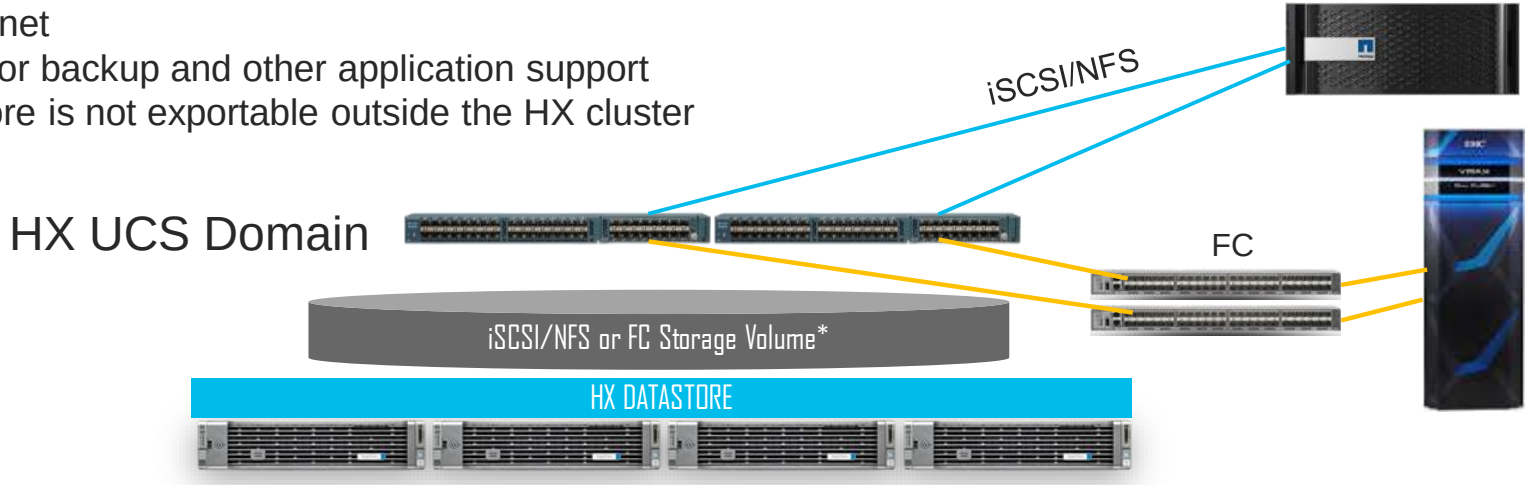

業界唯一的整合架構

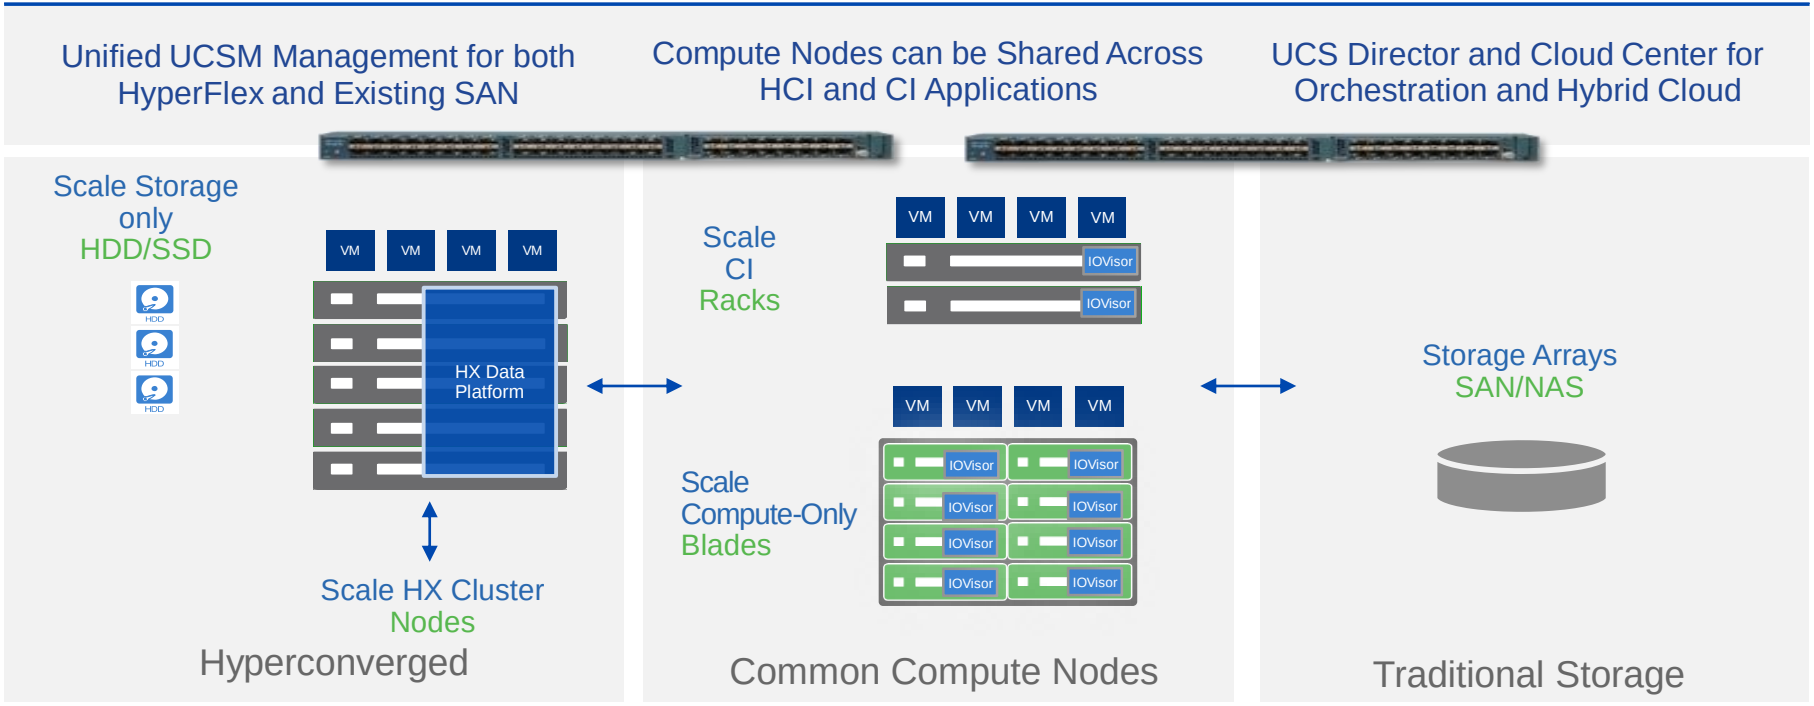

計算和儲存容量可獨立擴充

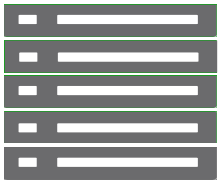

HX240C

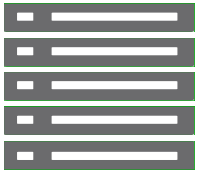

UCS C240

UCS B200

NVIDIA GPU CARD

SAN STORAGE ARRAY

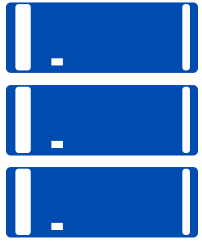

UCS S3260

HCI Compute and Capacity

Multipurpose Rack Compute

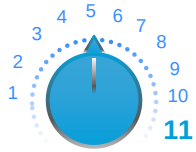

Multipurpose Blade Compute

Graphics Acceleration

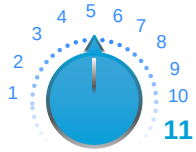

High Performance Storage

High Capacity Storage Servers

HyperFlex 資料保護
靈活滿足你的業務需求

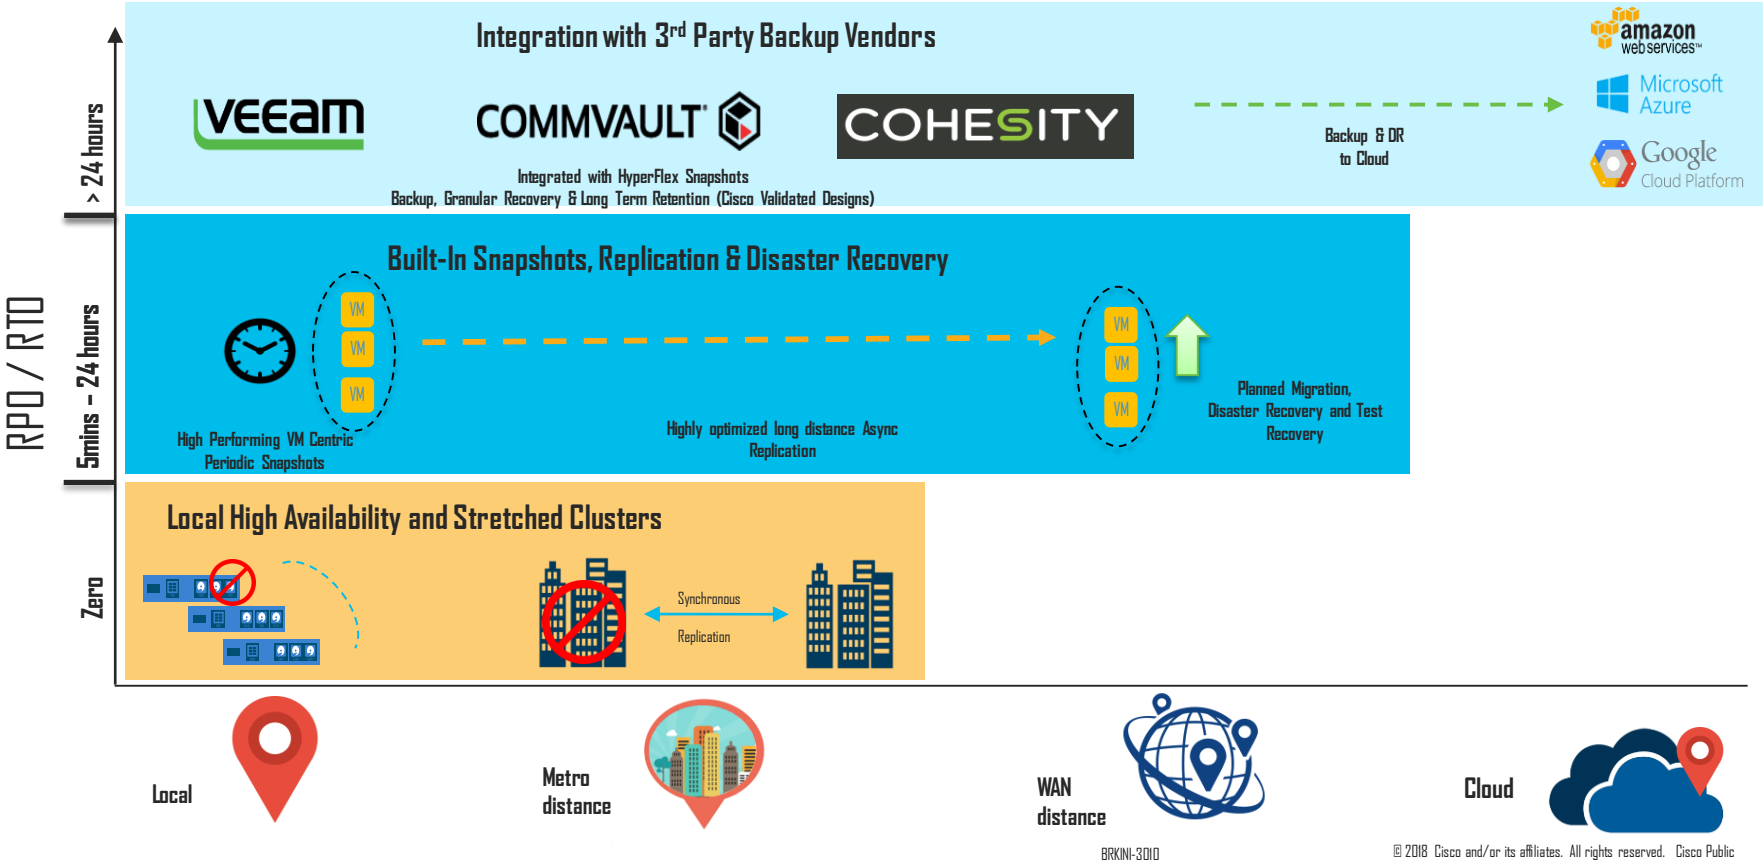

高效的資料服務

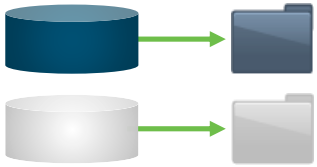

Pointer-Based Snapshots

- Space Efficient
- Fast creations and deletions
- VM-Level or VM Folder Level
- VAAI integrated
- vCenter Snapshot Manager
- Schedules, Retention period

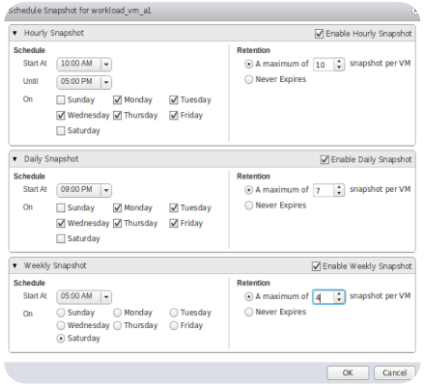

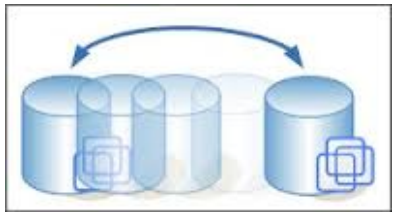

Ready Clones

- Batch creation GUI
- Apply unique names
- Use customization spec to apply IP

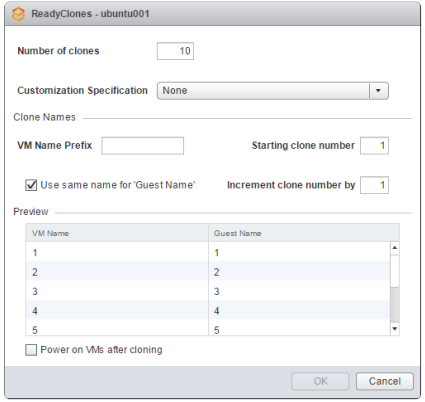

HyperFlex Data Protection
Built-In Replication for Business Continuity

Reliable

- TCP based reliable transmission
- Protected from network corruptions
- Robust fault handling
- Crash consistent VM snapshots

Performant

- Scale-out Replication streams
- Large IO transfers
- Network QoS
- Minimal impact on Primary IOs

Optimized

- incremental snapshots
- Compressed data on wire
- Intelligent pattern detection
- Different Primary and Target side configurations

HyperFlex Data Protection
Built-In Replication : Simple to Setup

1

Configure Network

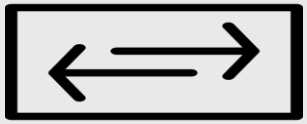

- Create a Replication VLAN
- Input network parameters
- Set Replication bandwidth
- One time setup

2

Pair Sites

- Assign a name
- Provide remote site credentials
- Map Datastores

3

Set Protection Group

- Create Protection Groups
- Provide a replication interval
- Select Crash Consistency

HyperFlex Data Protection
Built-In 1-Click Disaster Recovery

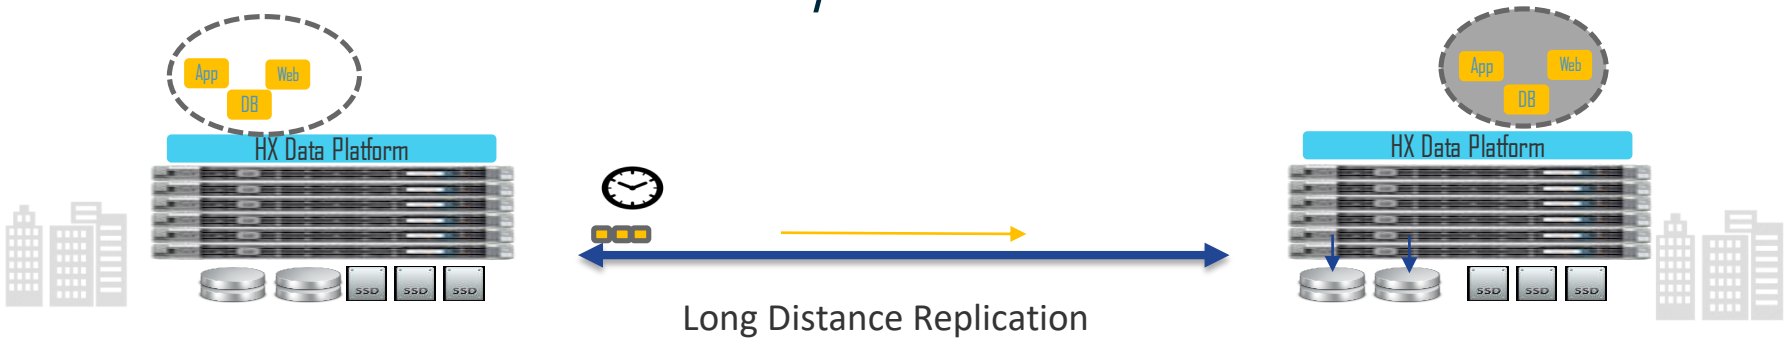

Test Recovery

- DR Readiness
- Customize DR Test parameters

Planned Migration

- Move VMs across Data Centers / Clusters
- Re-Protect after Migration

Unplanned Failover

- Recovery VMs after Disaster
- Re-Protect after Recovery

HyperFlex Stretched Cluster
Cloud Scale Data Platform

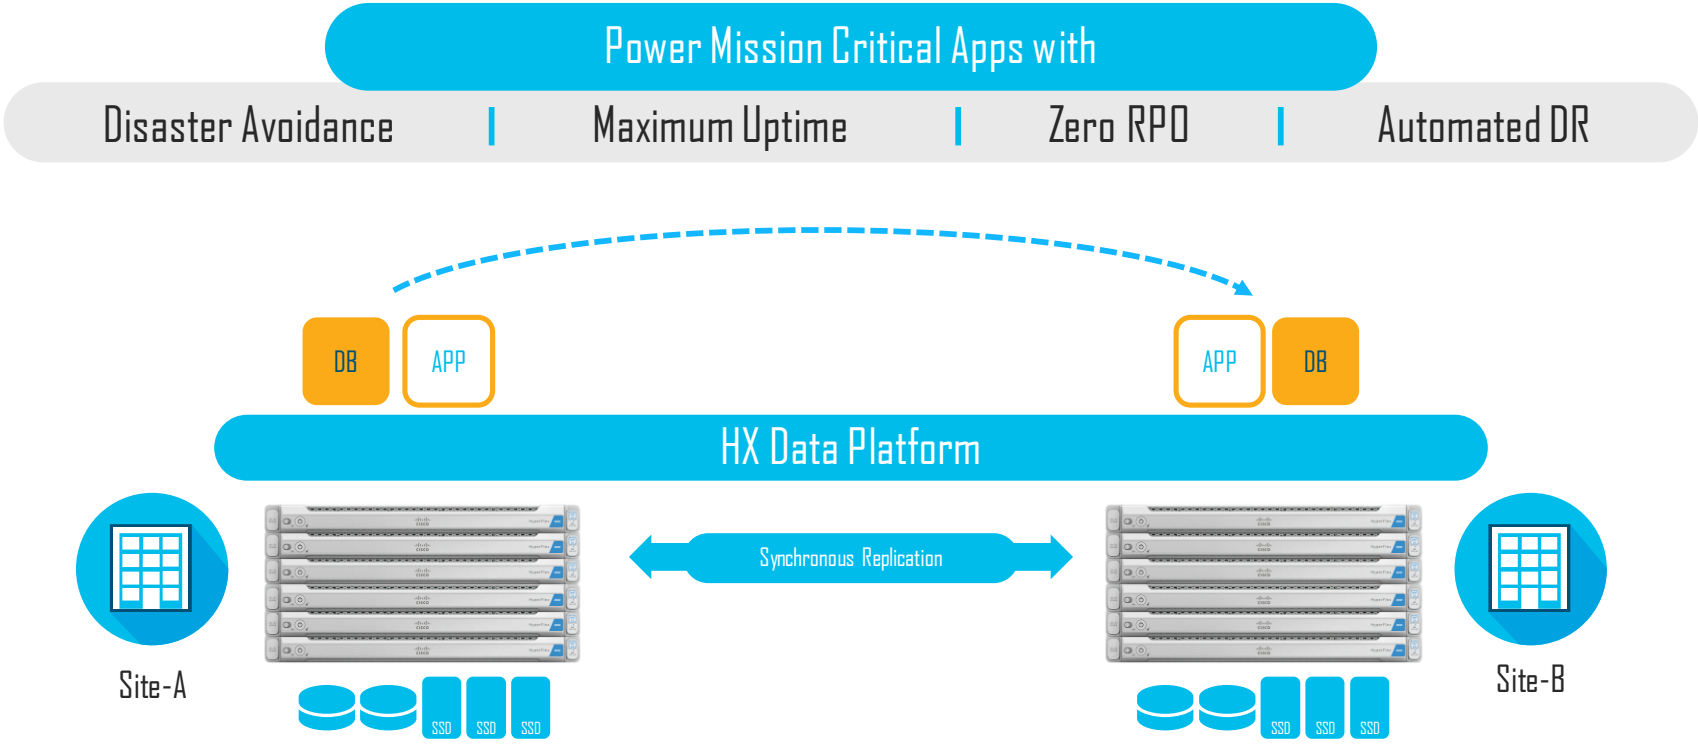

HyperFlex Stretched Cluster

ZERO RPO ! NEAR ZERO RTO!

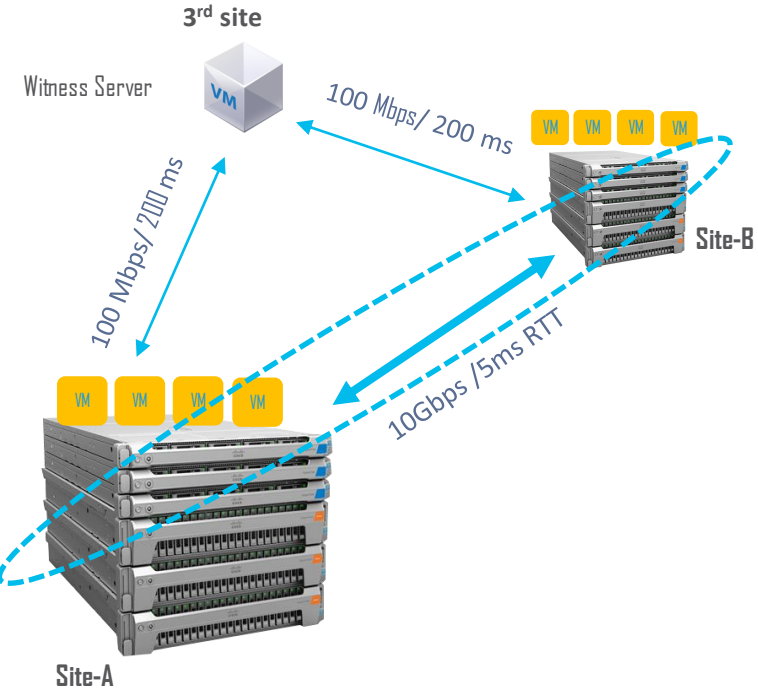

Configuration Support

- ✓ Single Stretched Cluster across 2 sites
- ✓ Symmetric Configuration
- ✓ Site to host a "Witness Server" (small VM)
- ✓ 8 HX nodes on each site

Management

- ✓ Cross site Cluster creation
- ✓ Non disruptive online rolling upgrade
- ✓ Site awareness in HXConnect
- ✓ Site specific Alarm and Events on a single Dashboard

IO Path

- ✓ Active-Active sites – VMs Active on each site
- ✓ VM Read IOs served locally
- ✓ VM Write IOs Sync-Writes across sites
- ✓ 2x copies on each site

HA Operations

- ✓ Recover from a Site failure
- ✓ Recover from a Local failure
  - ✓ Failover of VM
  - ✓ vMotion of VM
  - ✓ Split Brain handling

資料最佳化

Log-Structured File System Yields More Efficient Data Optimization

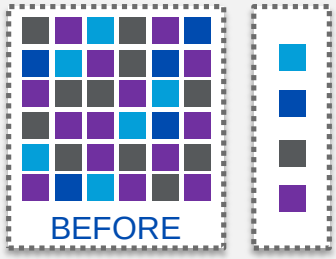

Inline Deduplication

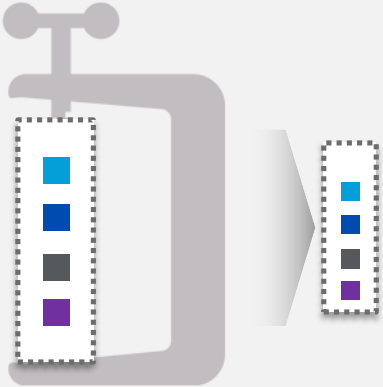

Inline Compression

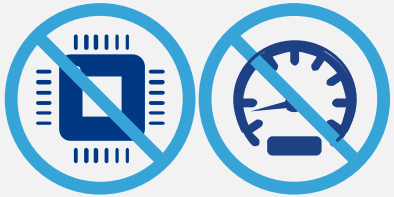

No Special Hardware
No Performance Impact

Upgrade Features

One-click, orchestrated ESXi hypervisor upgrades

Integrate with HXDP and UCS Server Firmware for truly one-click, full stack upgrades

*vCenter must be upgraded following VMware standard processes

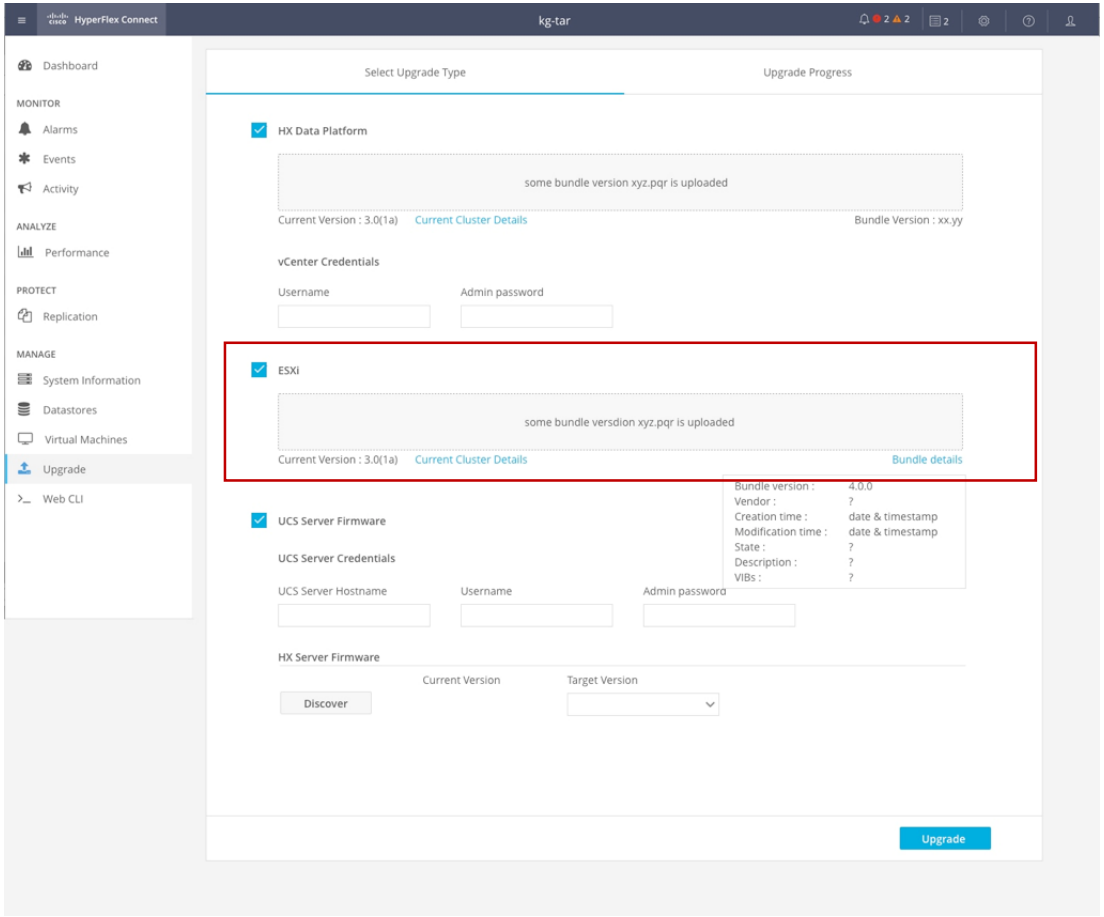

Enterprise Strategy Group | **Getting to the bigger truth.**™

ESG Lab Validation Preview Cisco HyperFlex

Consistent, High-performance Hyperconverged Infrastructure

Tony Palmer and Kerry Dolan, Senior Lab Analysts

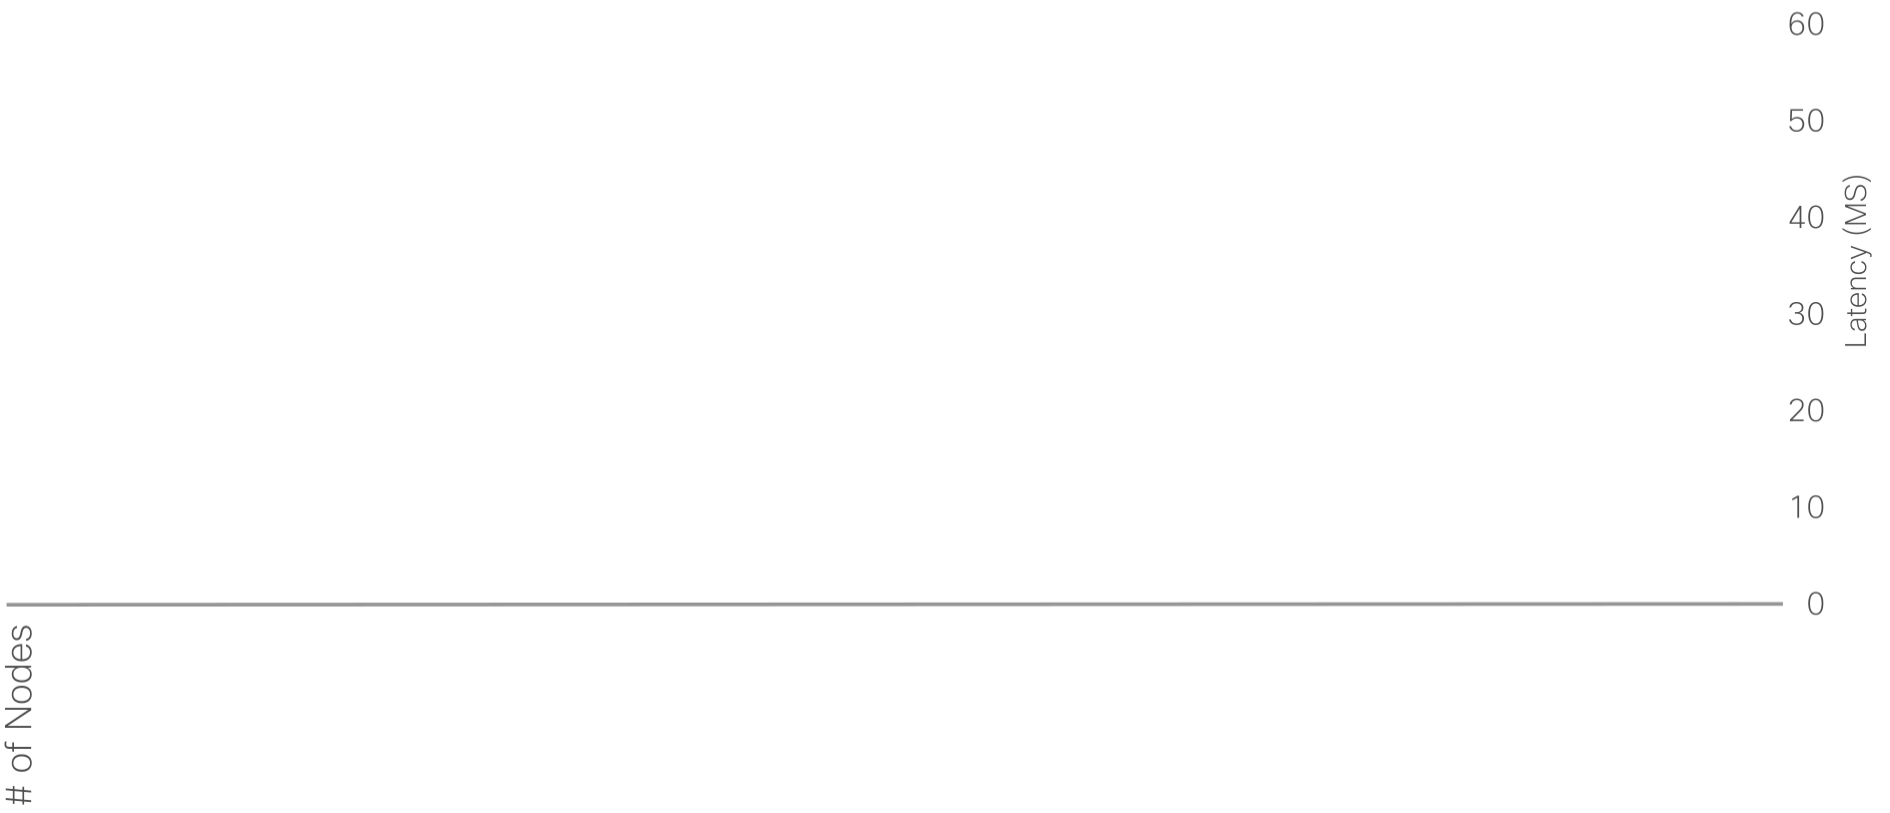
60
50
40
30
20
10
0
Latency (MS)
# of Nodes

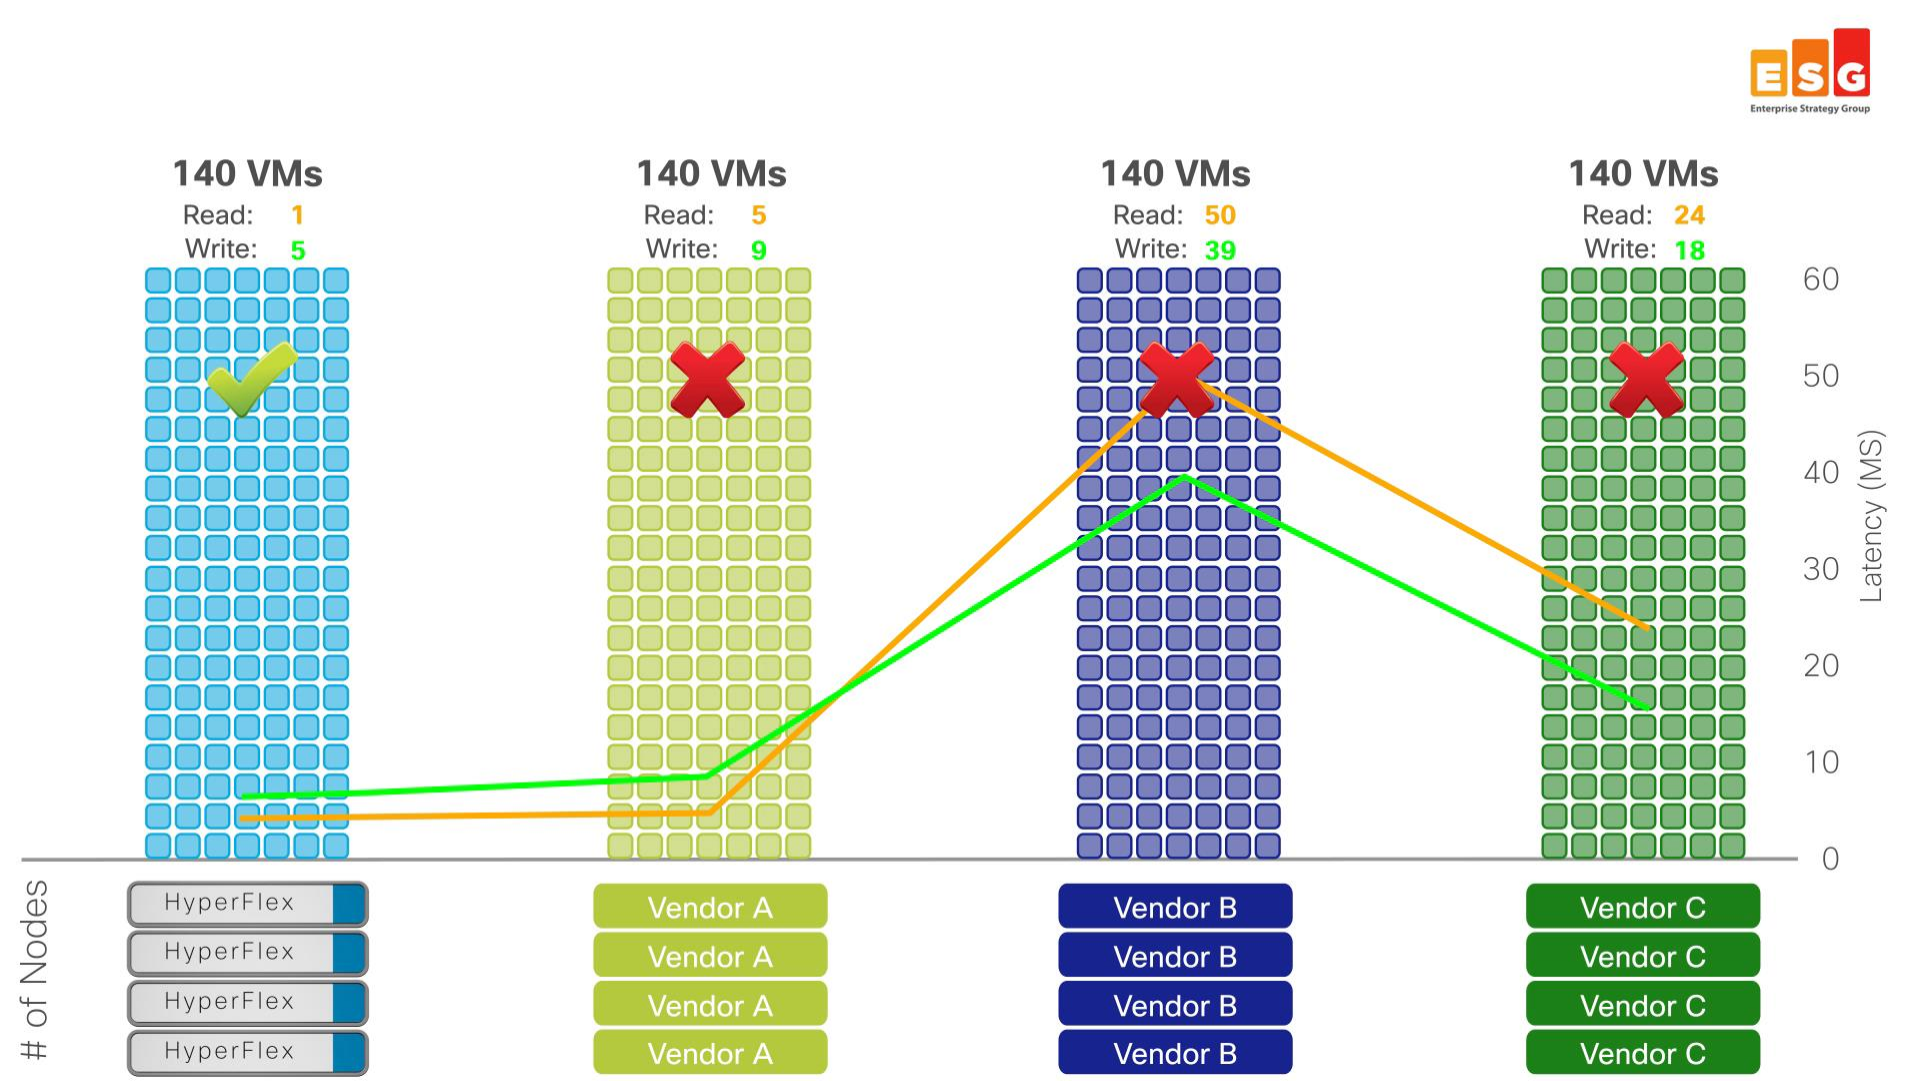
140 VMs
Read: 1
Write: 5
140 VMs
Read: 5
Write: 9
140 VMs
Read: 50
Write: 39
140 VMs
Read: 24
Write: 18
60
50
40
30
20
10
0
Latency (MS)
# of Nodes
HyperFlex
HyperFlex
HyperFlex
HyperFlex
Vendor A
Vendor A
Vendor A
Vendor A
Vendor B
Vendor B
Vendor B
Vendor B
Vendor C
Vendor C
Vendor C
Vendor C

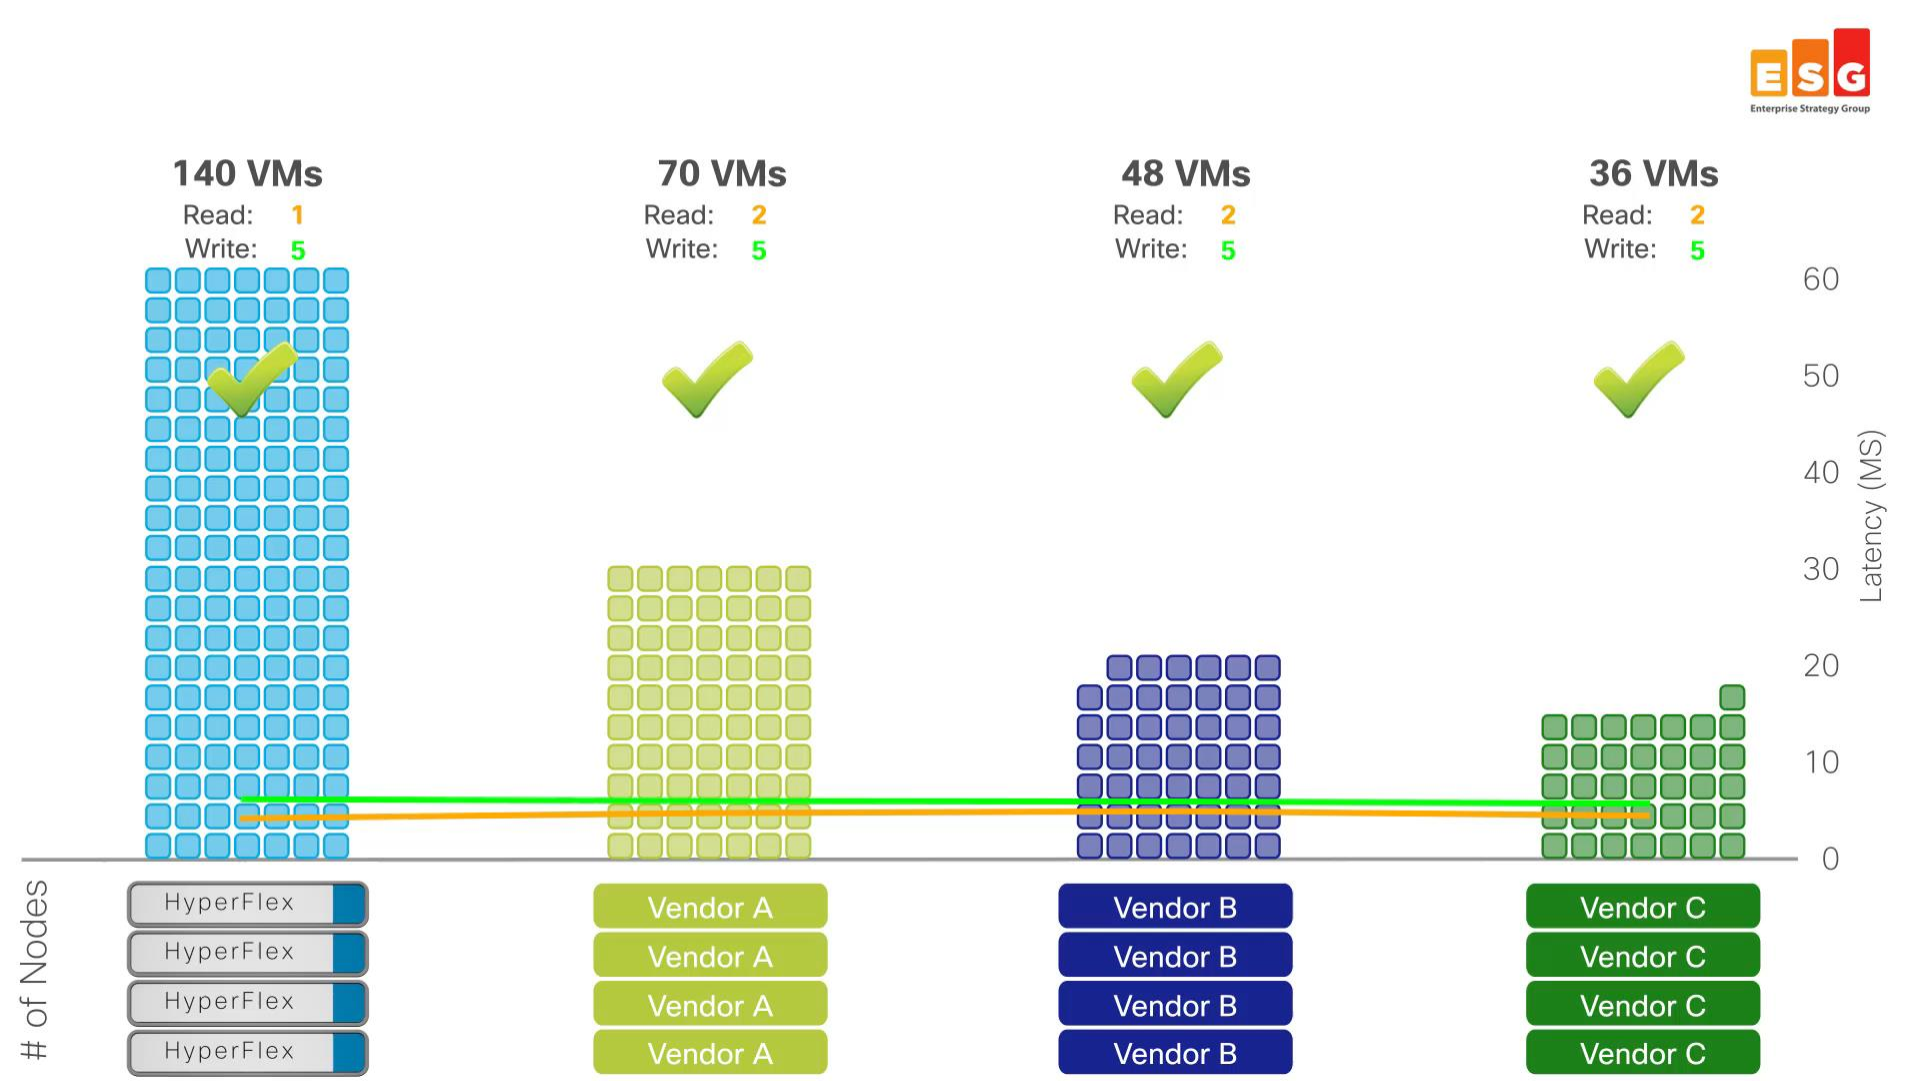
140 VMs
Read: 1
Write: 5
70 VMs
Read: 2
Write: 5
48 VMs
Read: 2
Write: 5
36 VMs
Read: 2
Write: 5
60
50
40
30
20
10
0
Latency (MS)
# of Nodes
HyperFlex
HyperFlex
HyperFlex
HyperFlex
Vendor A
Vendor A
Vendor A
Vendor A
Vendor B
Vendor B
Vendor B
Vendor B
Vendor C
Vendor C
Vendor C
Vendor C

HyperFlex All-Flash Performance Consistency

Each dot represents a Virtual Machine and its average IOPS over an hour of load

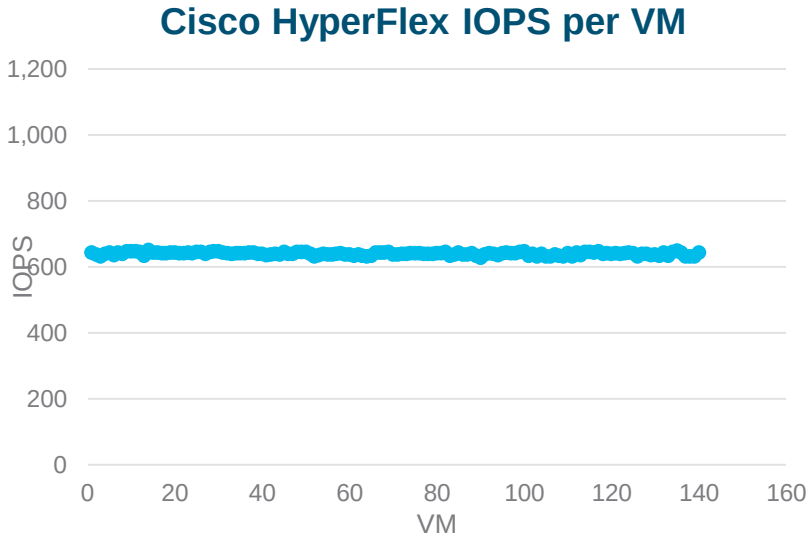

Consistent Performance for All VMs on the Cluster

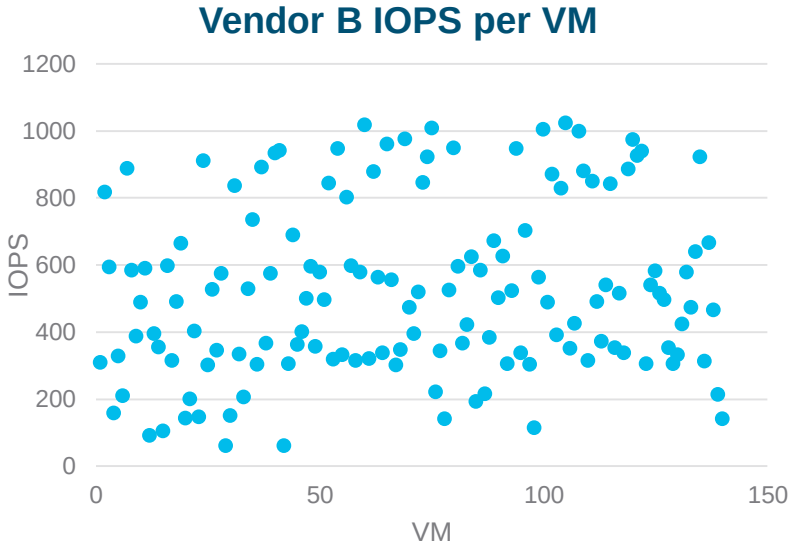

Wildly Inconsistent VM Performance Across the Cluster

HyperFlex Performance Testing Findings

3x Higher VM Density

3x Better TCO

3x Reduced Read/Write Latency

More Workloads on Hyperconverged

7:1 Reduced IOPS Variability

Predictable End-User Experience

Read the full ESG report:
cisco.com/go/hyperflex

HyperFlex Product Differentiation

Architected to Optimize Across Hardware, Software, Networking and Management.
Integrated Solution with Single Point of Support

High Performance & Scalable Data Platform

#1 Performing HCI platform

Consistent , Low latency performance

3X Lower TCO, 3X Higher VM Density, 64 node scale, linear scale out performance

Enterprise Class Data Services & Storage Optimization

Integrated Dedup & Compression w/ no performance penalty

Seamless integration of Converged & Hyperconverged

Investment protection of existing storage and compute investment

Independent Scaling of Compute & Capacity

Cost optimization through Compute-only node support

Deployment Automation & Simplicity

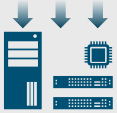

Out-of-the-box service profiles, install/upgrade automation, automated cluster scaling

Integrated High Performance Network Fabric

10G/40G VIC/Fabric
Factory installed, integrated networking, fabric QoS

Data Protection, High Availability & Resiliency

Native replication, backup/DR, Stretch Cluster, Availability Zones, Fault tolerant HA architecture

Cloud based centralized management

Monitoring, Telemetry, Analytics, Policy, Orchestration, Proactive TAC, HX Cluster management

Broad Range Of Supported Workloads	ROBO (Branch, IOT)	VSI (app/web)	VDI (Citrix, Horizon)	Collaboration (UC, HCS)	Databases (Oracle, SQL)	Mission Critical & ERP (SAP)	Analytics (Splunk)	Cloud-Native Apps (Docker, Kubernetes)


CISCO

HX Edge

Fully GA & redundant 10GbE Edge Network Option

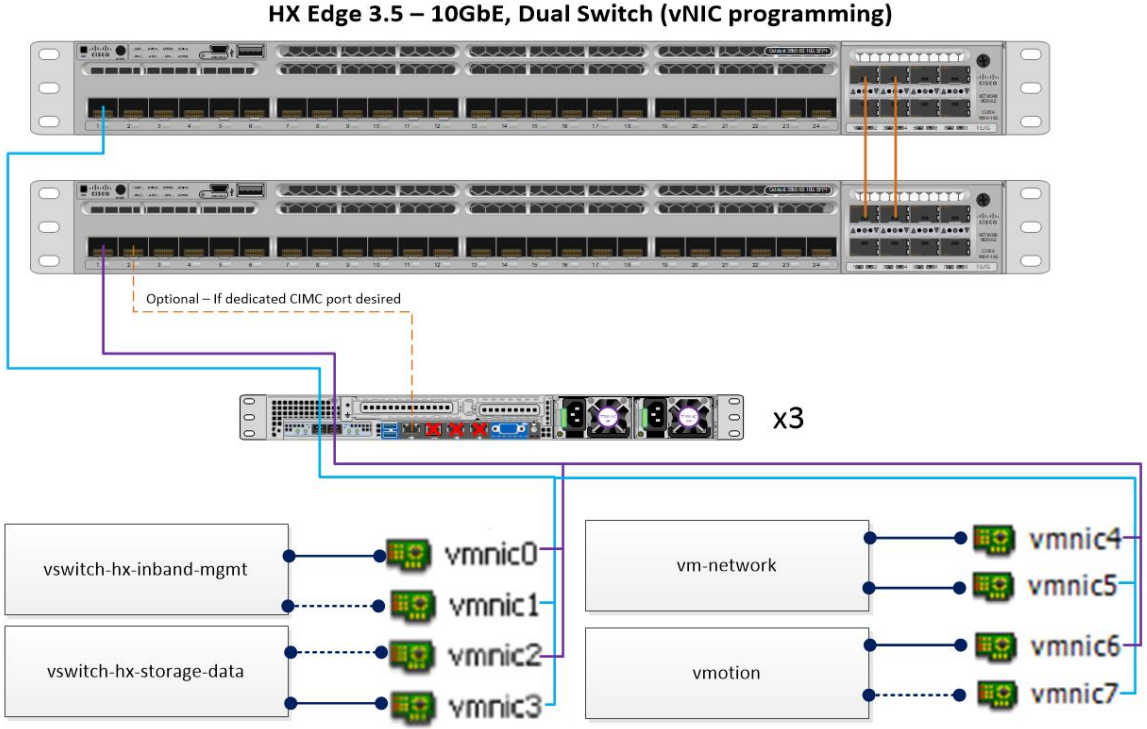

HyperFlex+AI =HX AIM

Use Case Conception → Feasibility Study → Model Design → Model Deployment → Model Maintenance

HX AIM

Quick Installation of AI Hardware and Software for Exploration, Experimentation and Lifecycle Management

cloudera HORTONWORKS hadoop APACHE Spark

Big Data Support for Data Preparation and Batch Training and Inference

NVIDIA V100

Offload and Accelerate Training at Scale with Performance Optimized Systems

NVIDIA P4

Distribute and Scale Inference from the Data Center to the Edge