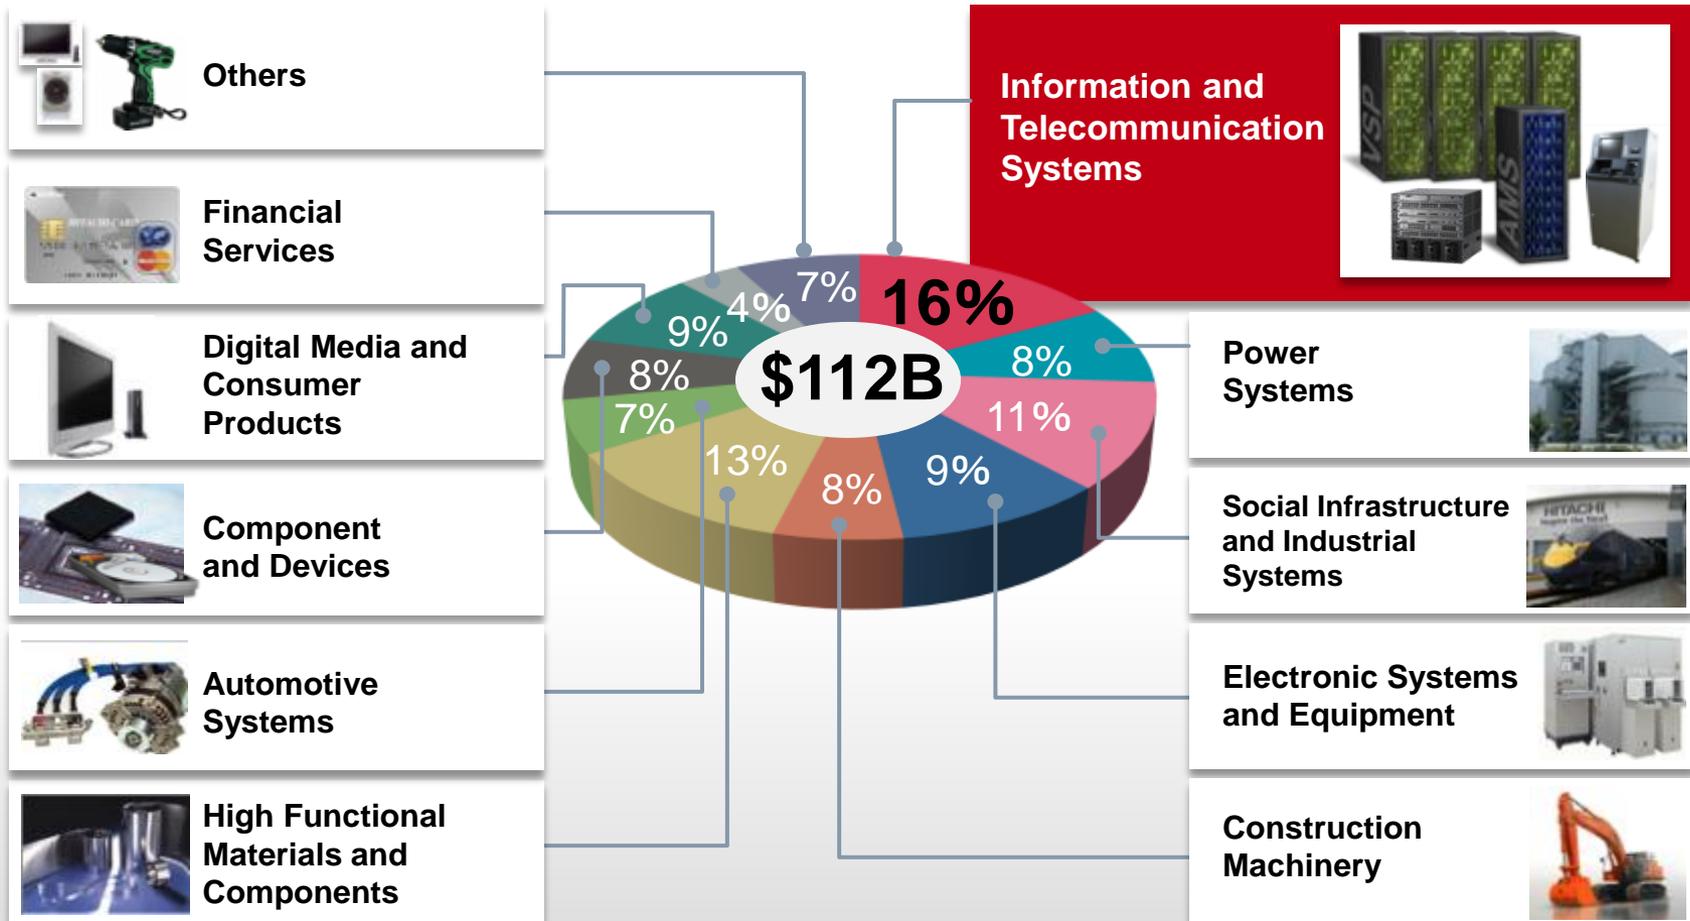


THE STORAGE INFRASTRUCTURE DESIGN WITH HITACHI STORAGE SYSTEM ON HYPER-V CLOUD DEPLOYMENT

STEPHEN SO



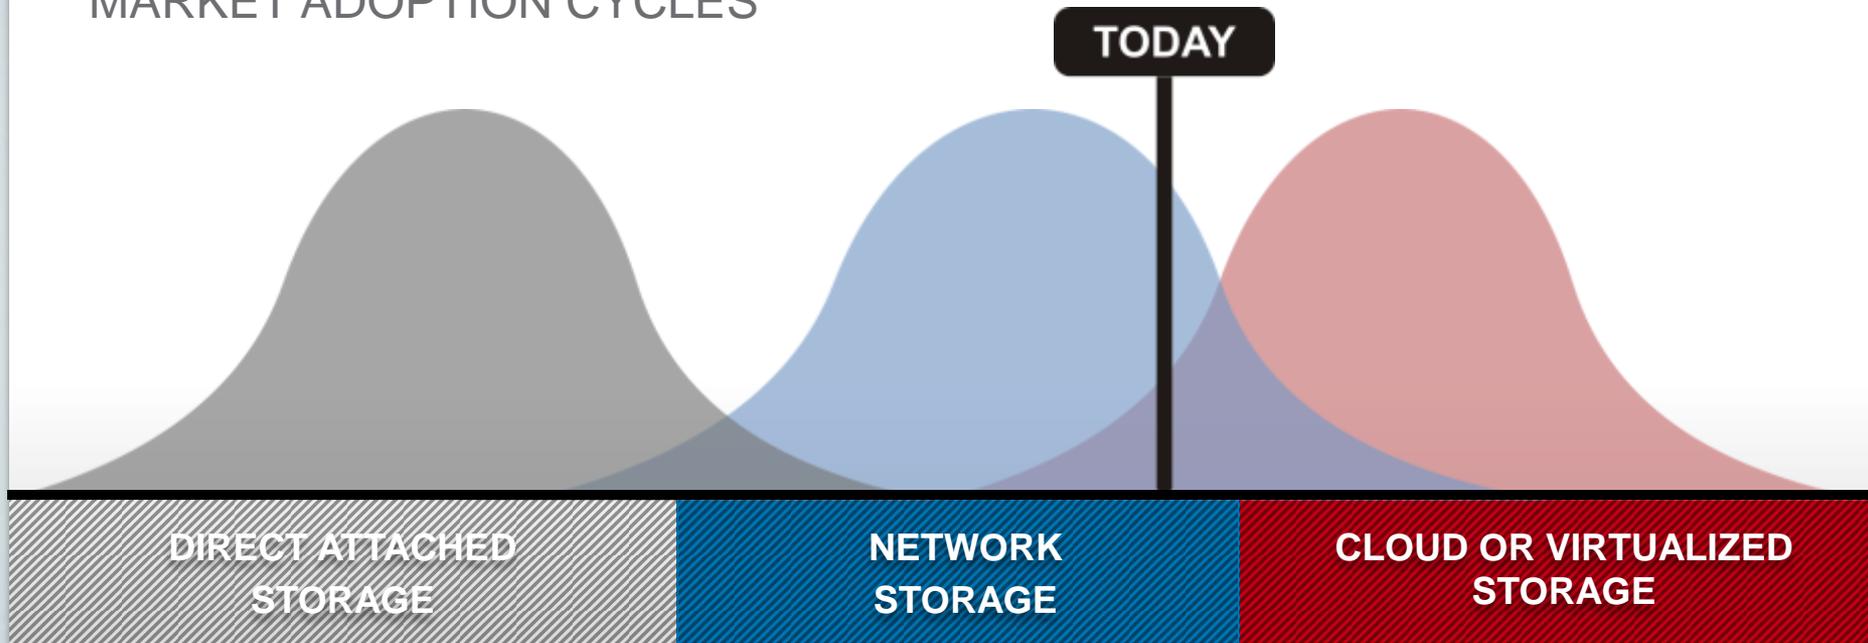
11 INDUSTRY SEGMENTS OF HITACHI GROUP GROUP REVENUE BY INDUSTRY SEGMENT



PARADIGM SHIFTS

STORAGE 3.0

MARKET ADOPTION CYCLES



- **Networked Storage** (NAS and SAN) consolidates and virtualizes disk (capacity) for improved provisioning flexibility and efficiency
- **Cloud-based or virtualized storage** consolidates and virtualizes disks and controllers to scale capacity and performance
 - It enhances resiliency and reduces both CAPEX and OPEX for service providers and end users

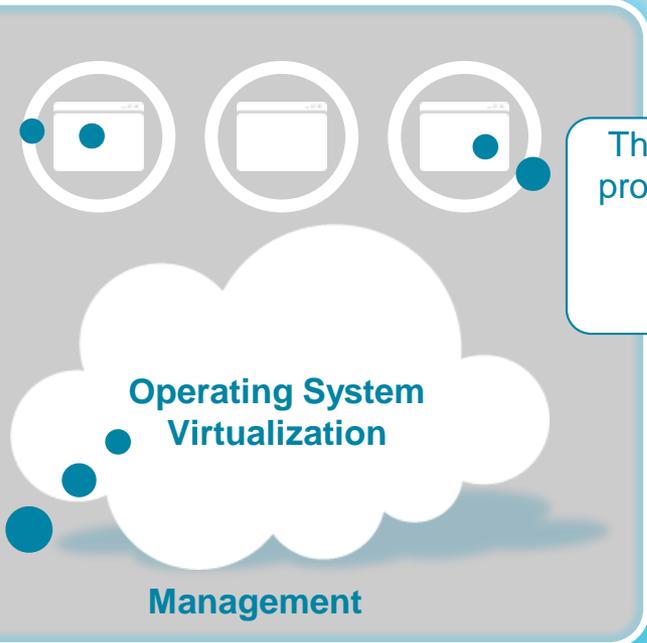
Benefits of a Private Cloud (IaaS) Beyond Large-Scale Virtualization



The Applications use as much or as little of the resource pool as needed
Dynamic



The Applications can be provisioned by end-users based on business requirements
Self-service

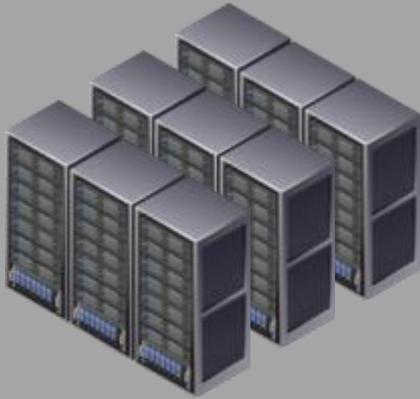


The Applications can be turned on, off, moved & scaled
Elasticity



The ability to scale is now based on the application's needs, *not* limited by the underlying resources; Server, VM, or even the datacenter

Physical, Legacy Environment



Transition to Virtual Environment



- Reduce capital and operational expenditures
- Increase flexibility to be more responsive to the business
- Improve service levels

From Managing Infrastructure Technology

To Providing Infrastructure Services

Silos of independent physical devices

Virtualized

Manage physical devices independently

Automated

Closed, inflexible monolithic environment

Cloud-ready

Costly, legacy systems

Sustainable

THE PROGRESSION OF IT SCALE UP...SCALE OUT...SCALE DEEP

HITACHI
Inspire the Next

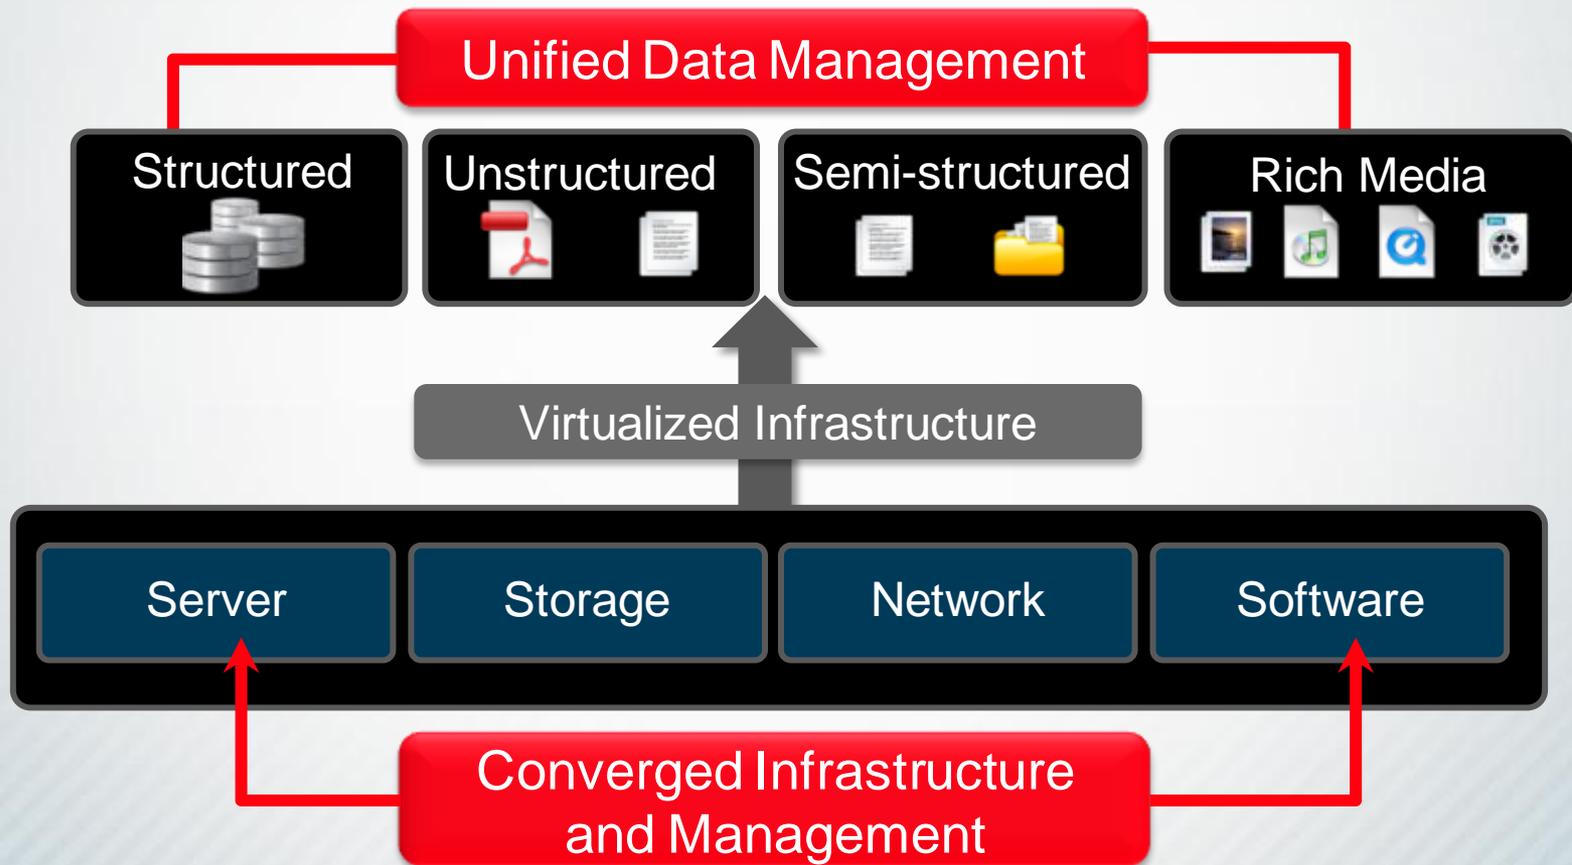
Servers

Network

Virtualization



A COMMON VIRTUALIZED PLATFORM FOR ALL DATA, CONTENT AND INFORMATION



HITACHI & MICROSOFT: EXPANDING OUR PARTNERSHIP



HITACHI
Inspire the Next

Hitachi Investments on the Microsoft Campus

- Over 1 PB Hitachi Storage in key Microsoft Labs
- Microsoft Partner Solutions Center
- Microsoft Enterprise Engineering Center
- Receiving strong support from Hitachi Ltd and joint program development

Microsoft Solution Centers:

- Enterprise Engineering Center - EEC
- Microsoft Partner Solutions Center - MPSC
- Microsoft Technology Centers - MTC's

Hitachi Server and Storage Solutions

- Hitachi Blades
 - BDS 2000's
- Hitachi Storage
 - USP V
 - AMS 2500
- POC Engagements
 - Hyper-V Cloud Fast Track
 - Exchange Workloads
 - SharePoint Workloads
 - SQL Workloads

Microsoft Qualification Programs

- Windows Logo Program
- Windows Hardware Qualification List
- Failover Cluster Configuration Program
- Exchange Solutions Reviewed Program
- SQL Always On Technologies
- System Center Alliance Program

Microsoft Gold Certified Partner

- Access to 24x7 solutions support
- 4 Business Hours for break-fix queries
- Microsoft Technical Account Management (TAM)
- Exchange, SQL and SharePoint TAP Programs



Microsoft Campus - Redmond, WA
Beijing, China - Tokyo, Japan
Irvine, CA - France - UK

THE PROBLEM: DATA CENTER COMPLEXITY IS OUT OF CONTROL

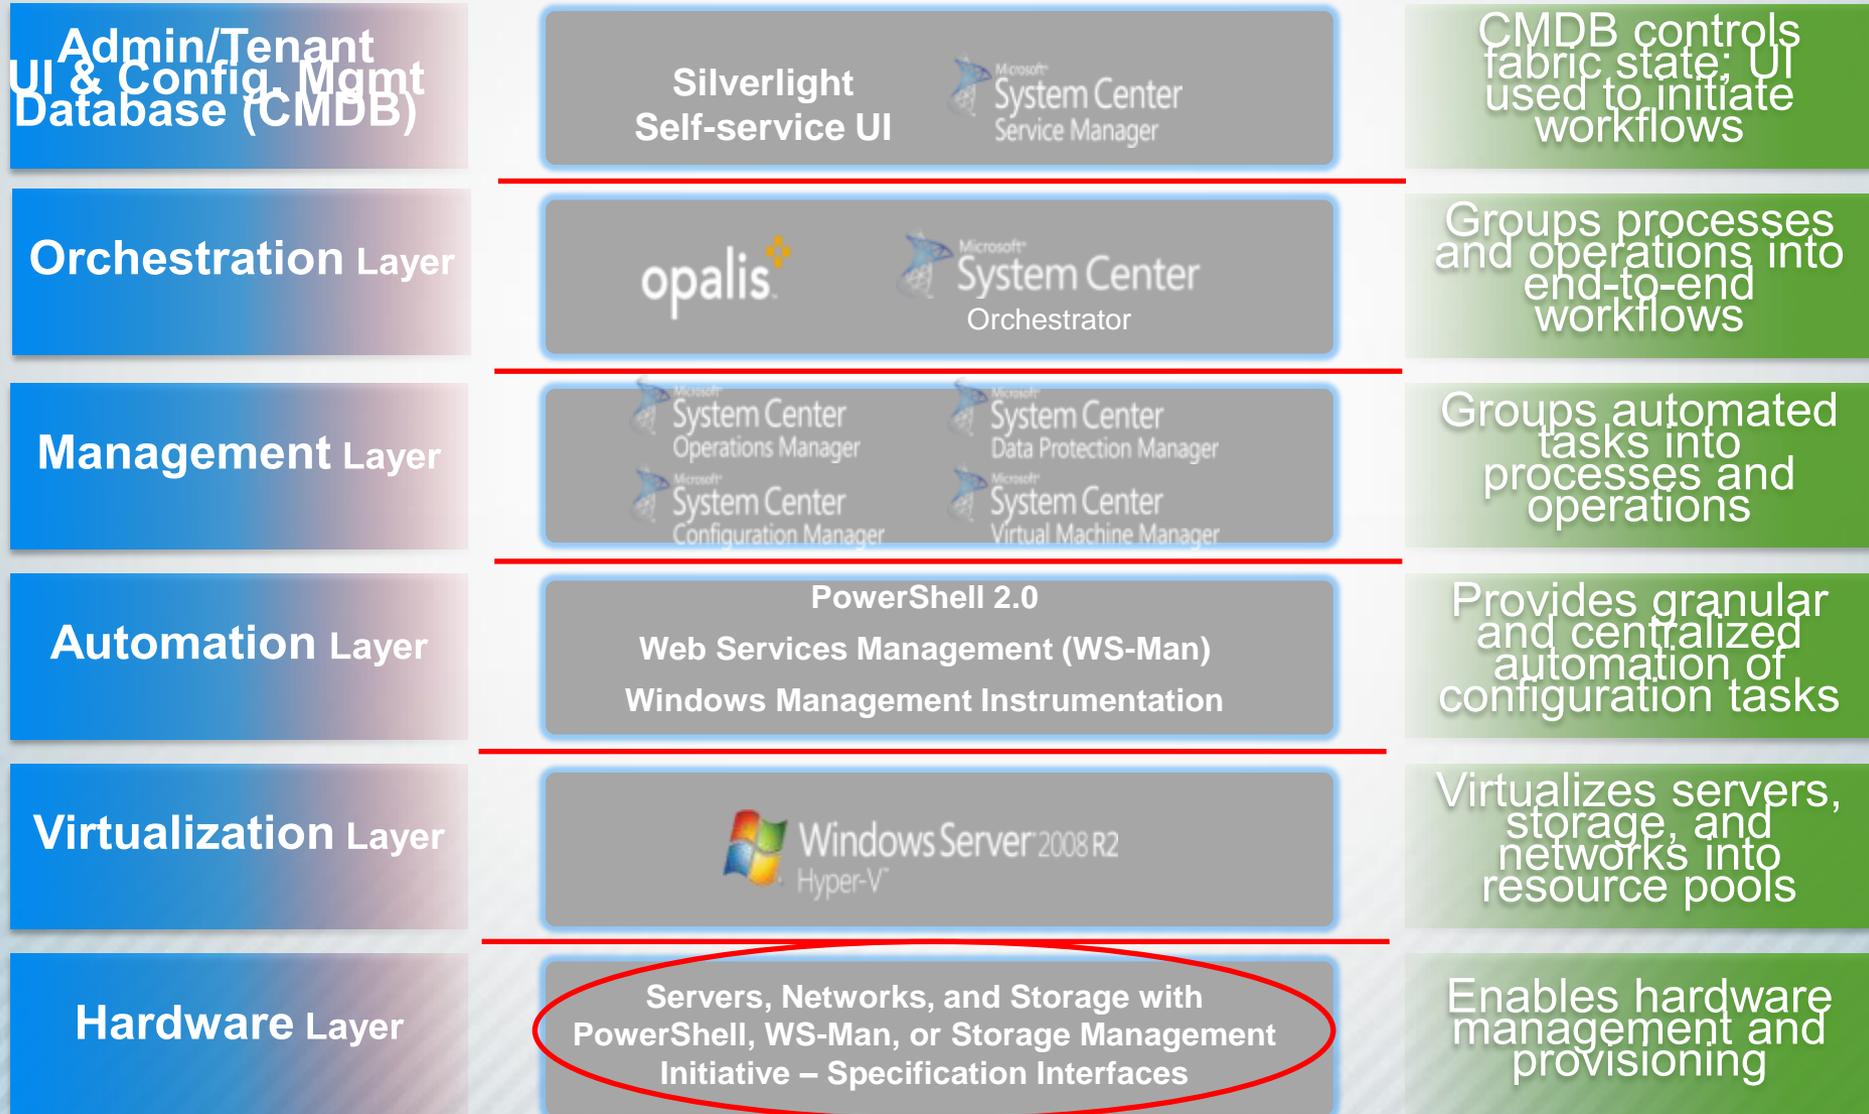
- Technology sprawl
- Resource constraints
- Multiple vendor management
- Skill gaps
- Faster deployment
- Skyrocketing operations costs
- Difficulty aligning value delivery with core business



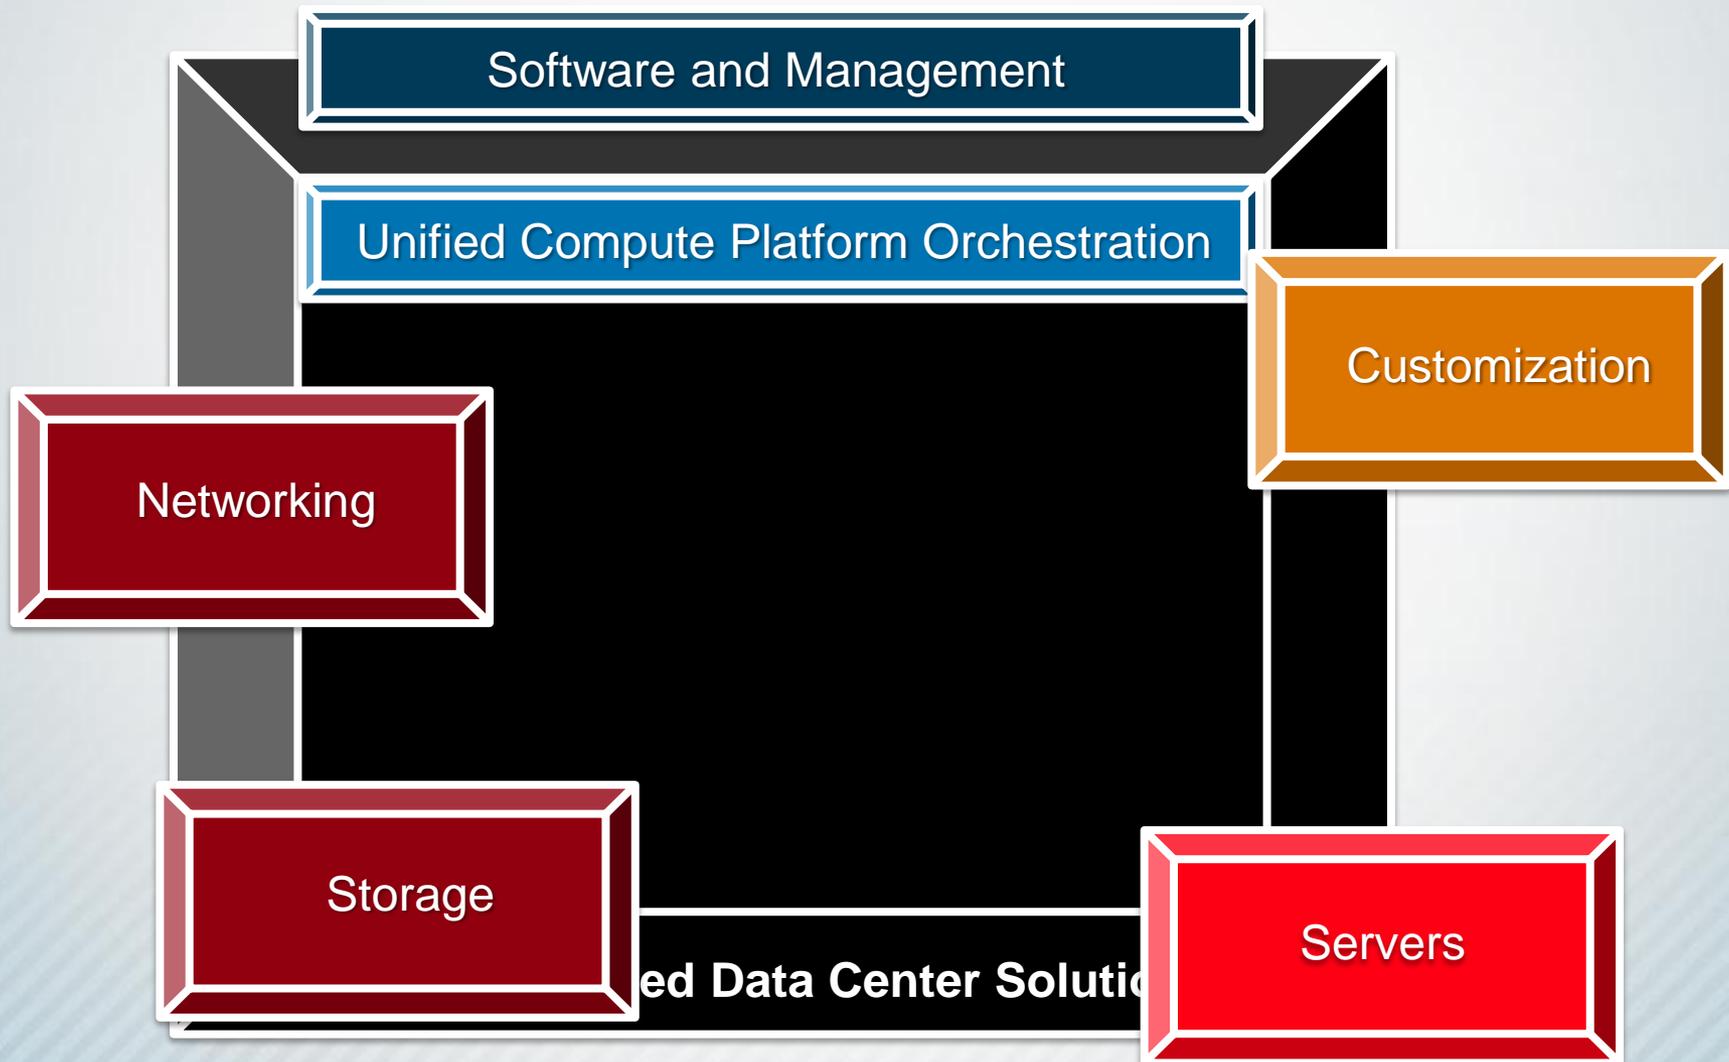
MICROSOFT HYPER-V CLOUD FAST TRACK

INTRODUCTION

IAAS ARCHITECTURE BLUEPRINT

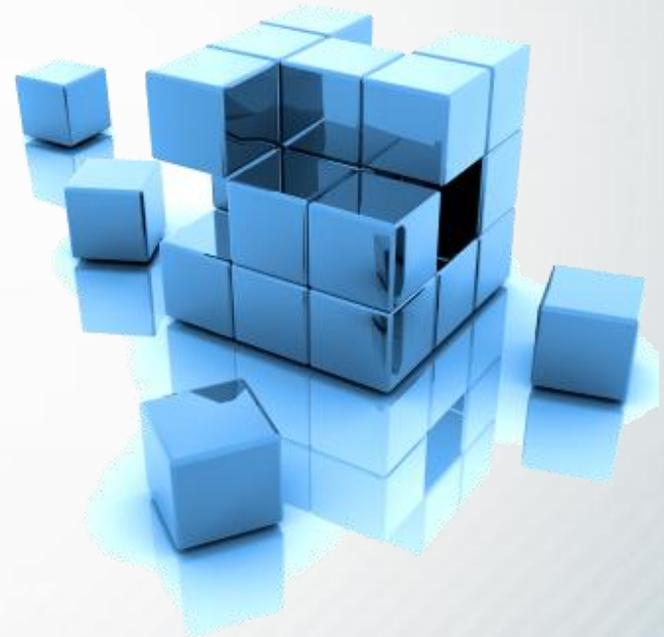


CONVERGED DATA CENTER SOLUTIONS



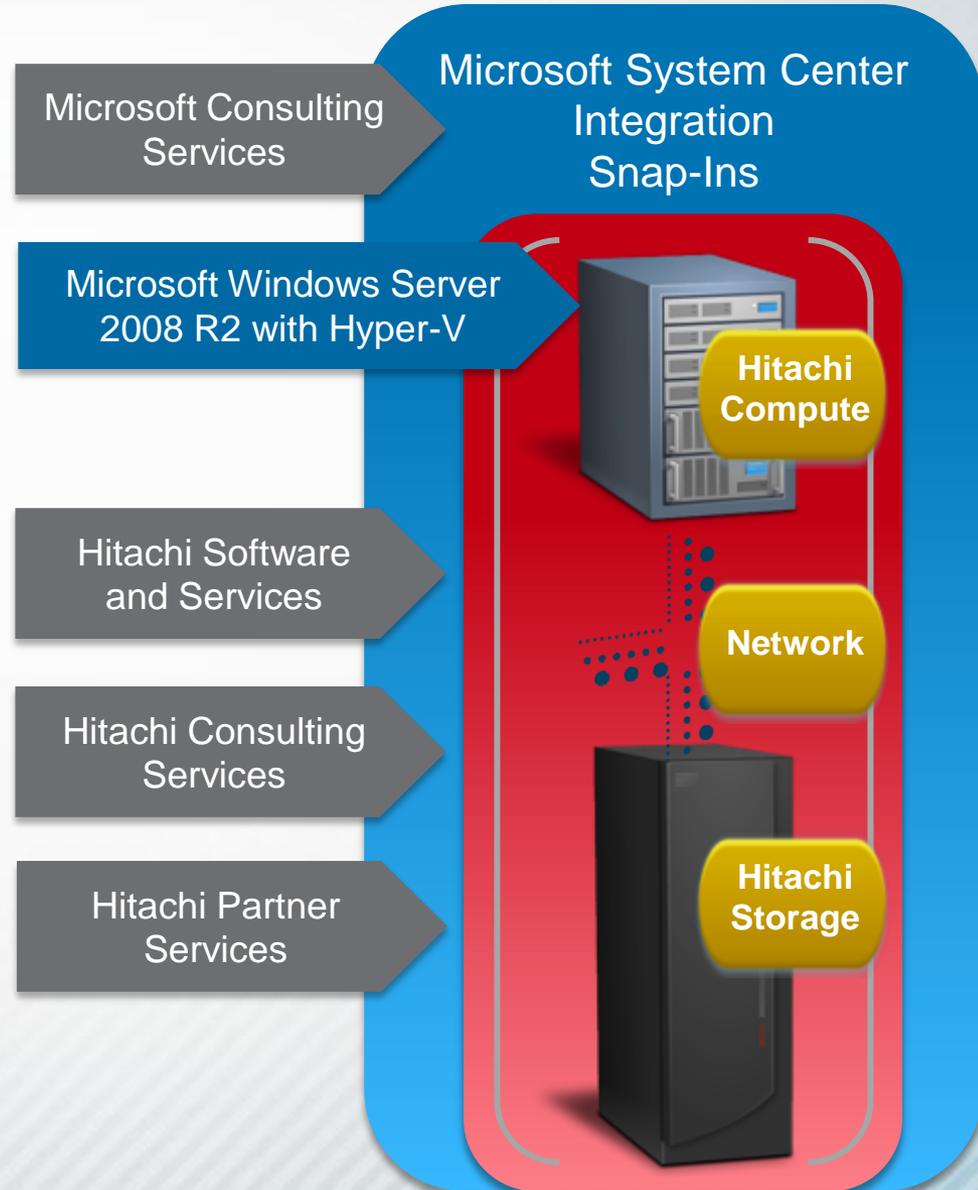
BUILDING BLOCK FOR PRIVATE CLOUD INFRASTRUCTURE

- **Predictable, repeatable, reliable**
 - Pre-validated reference architectures, guides and services
 - High availability Infrastructure
- **Exceptional performance**
 - High density computing and throughput
 - Optimized I/O balancing and concurrency
- **Fast time to value**
 - Quick solution deployment
- **Advanced flexibility**
 - Extreme adaptability to changes in workload
 - Various infrastructure choices



PARTNERSHIP

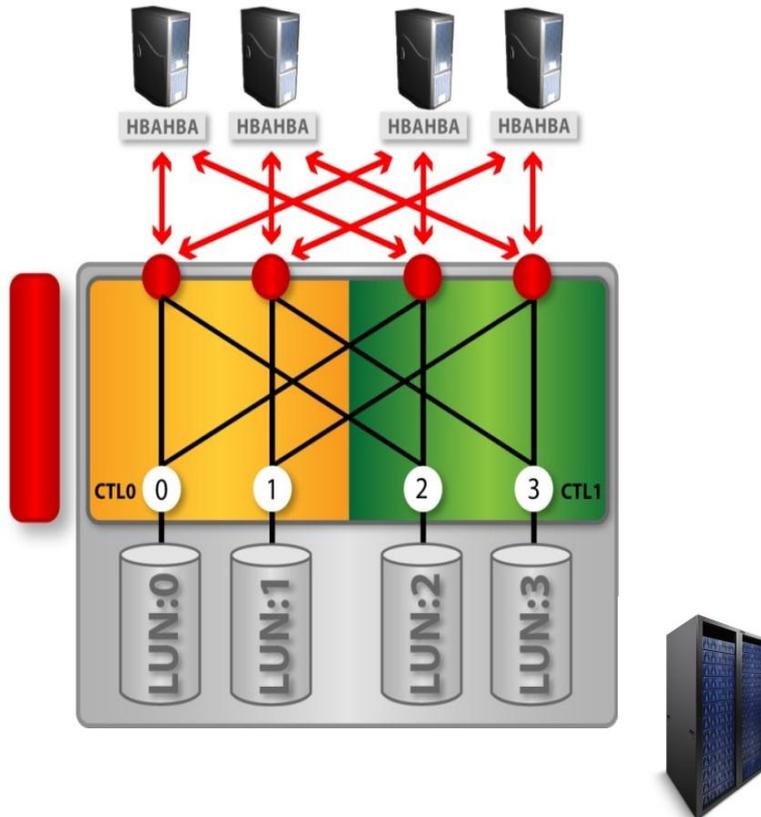
- Microsoft provides
 - Guidance, requirements
 - Validation of solution
- Hitachi Data Systems and Hitachi, Ltd. provide
 - Reference architectures
 - Bill of materials
 - Microsoft integrations
 - Deployment guides
 - Operations guides
 - Software and services



THE HITACHI ADVANTAGE – THE CONTROLLER

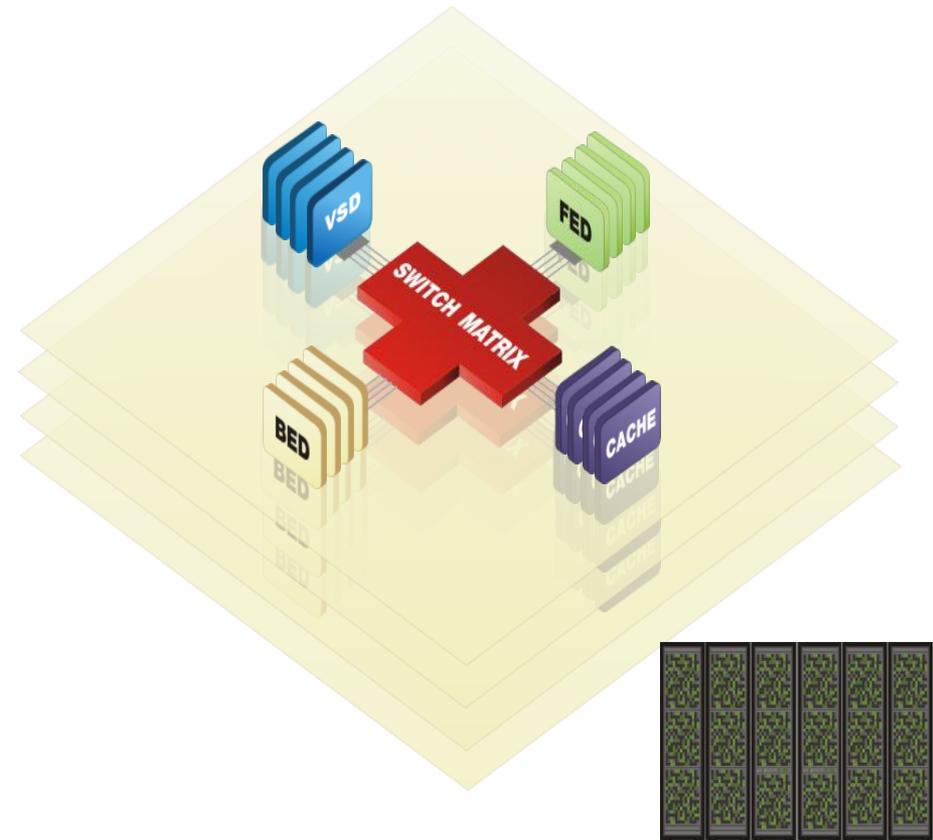
AMS 2000 Industry Differentiator

Hitachi Symmetric Active/Active Controller: Virtually eliminates controller bottlenecks found in traditional asymmetric controller designs



VSP – Virtual Storage Director

Increase performance, capacity, and connectivity tightly coupled through a global cache to meet increasing server demands



HITACHI VIRTUALIZATION

A KEY ELEMENT OF CLOUD STORAGE SERVICES

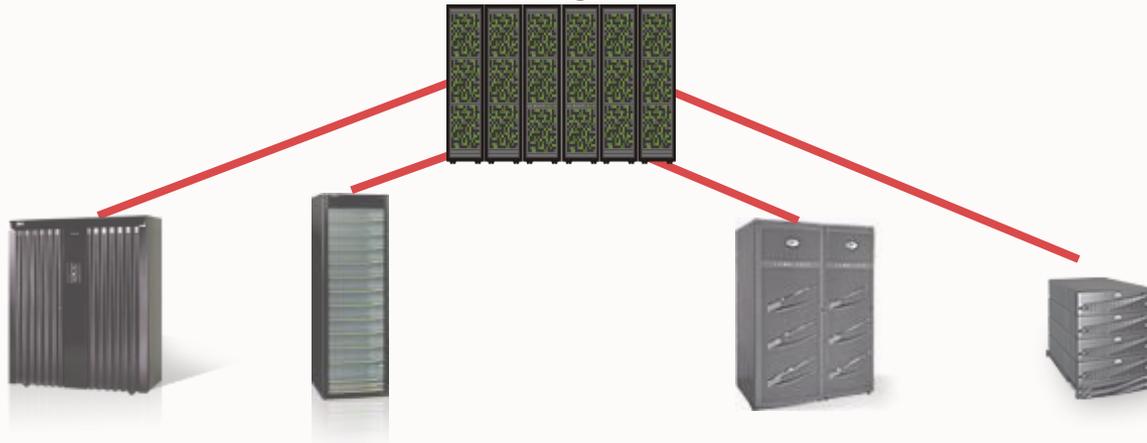
HITACHI
Inspire the Next

All connected storage inherits...



Virtual Storage Platform

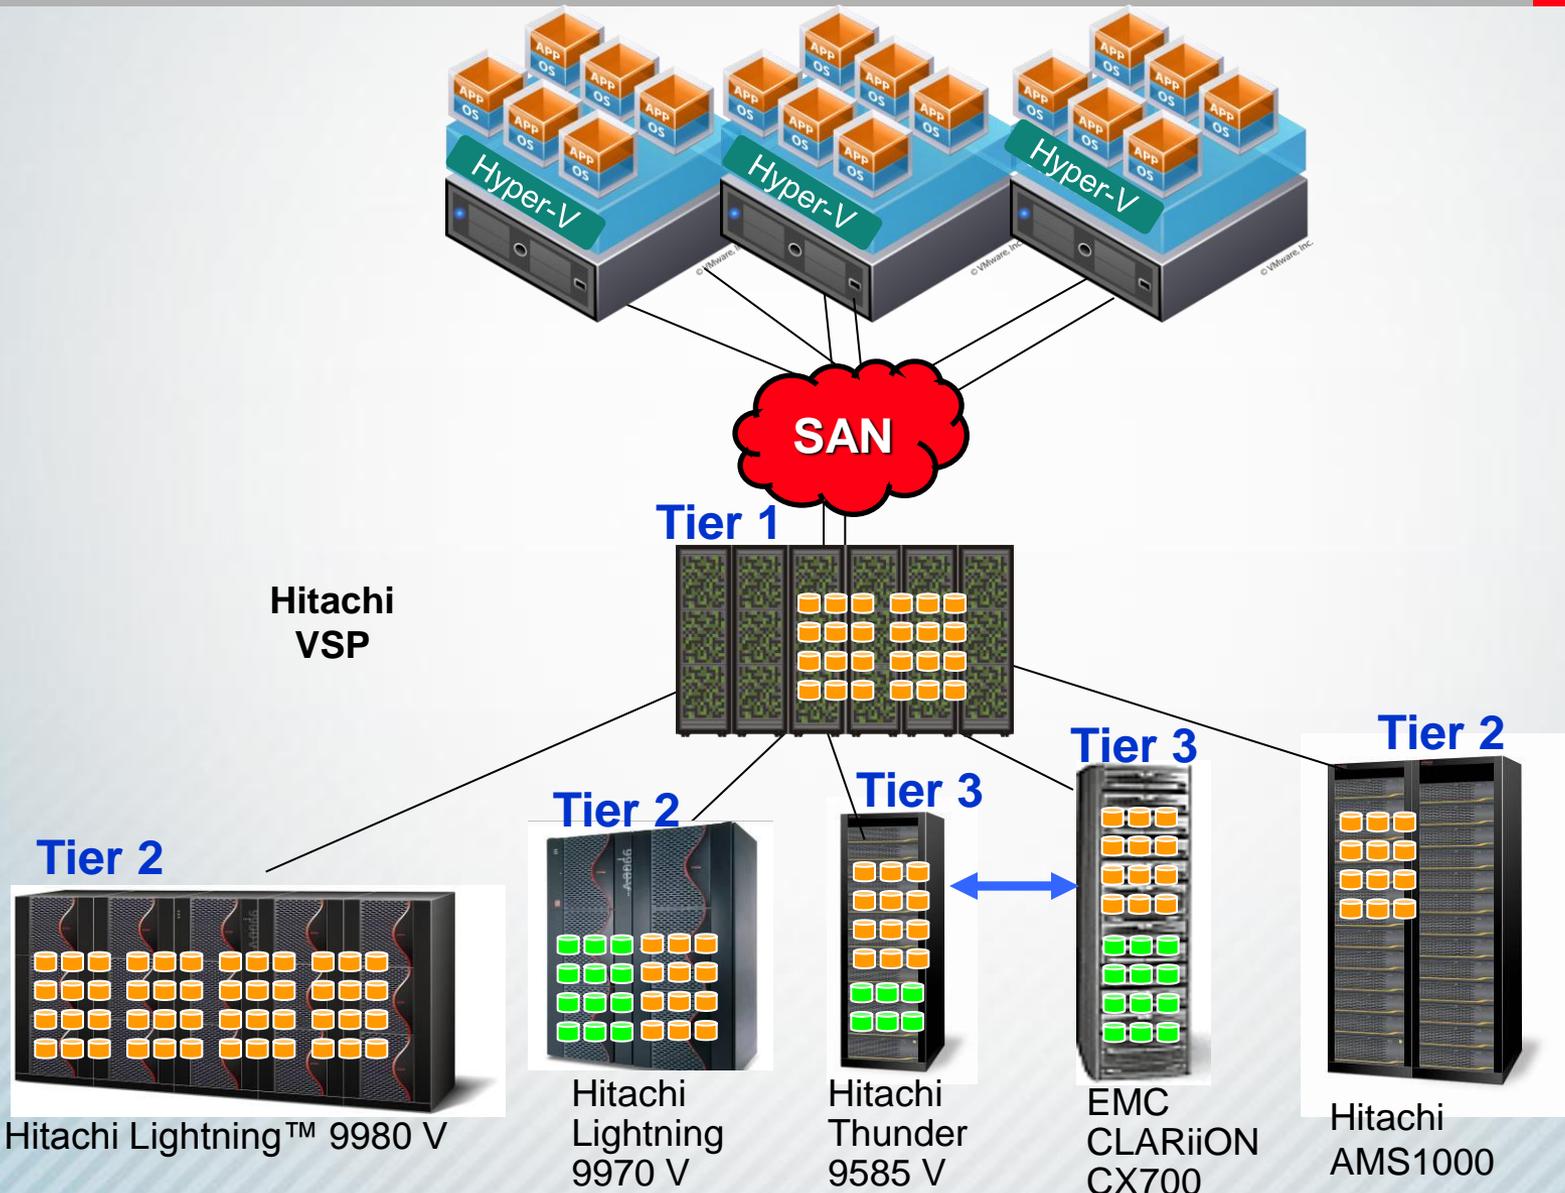
Inherited
Functionality



What does that mean to customers?

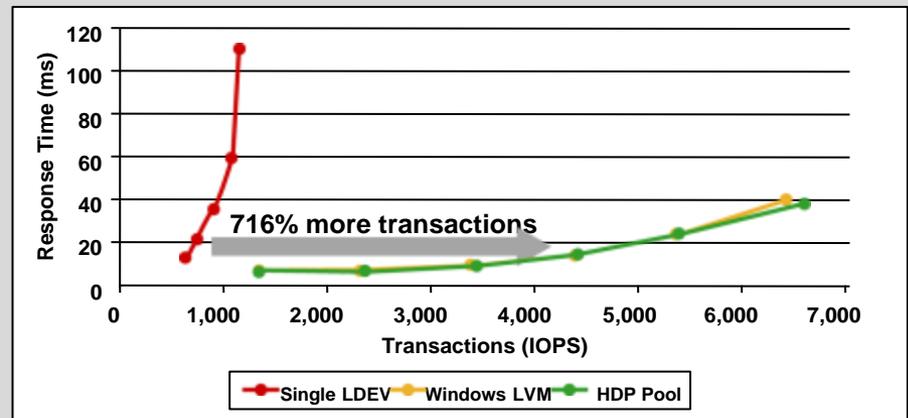
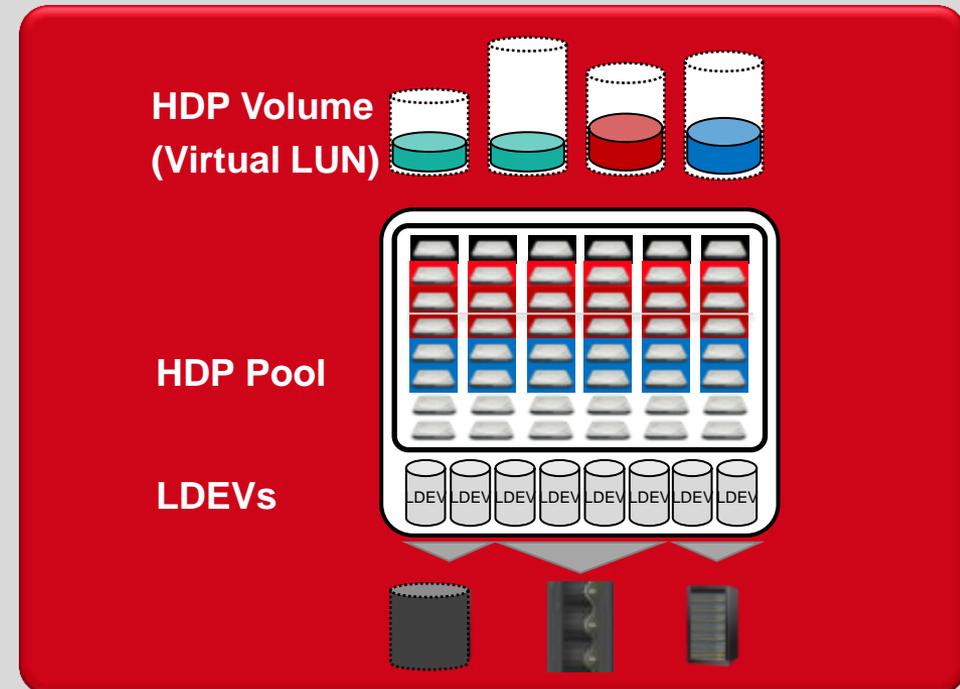
**A consistent and technologically advanced storage services strategy
across heterogeneous platforms**

CONSOLIDATION AND VIRTUALIZED STORAGE NON-DISRUPTIVE DATA MIGRATION



HITACHI DYNAMIC PROVISIONING PAY-PER-USE UTILITY MODEL

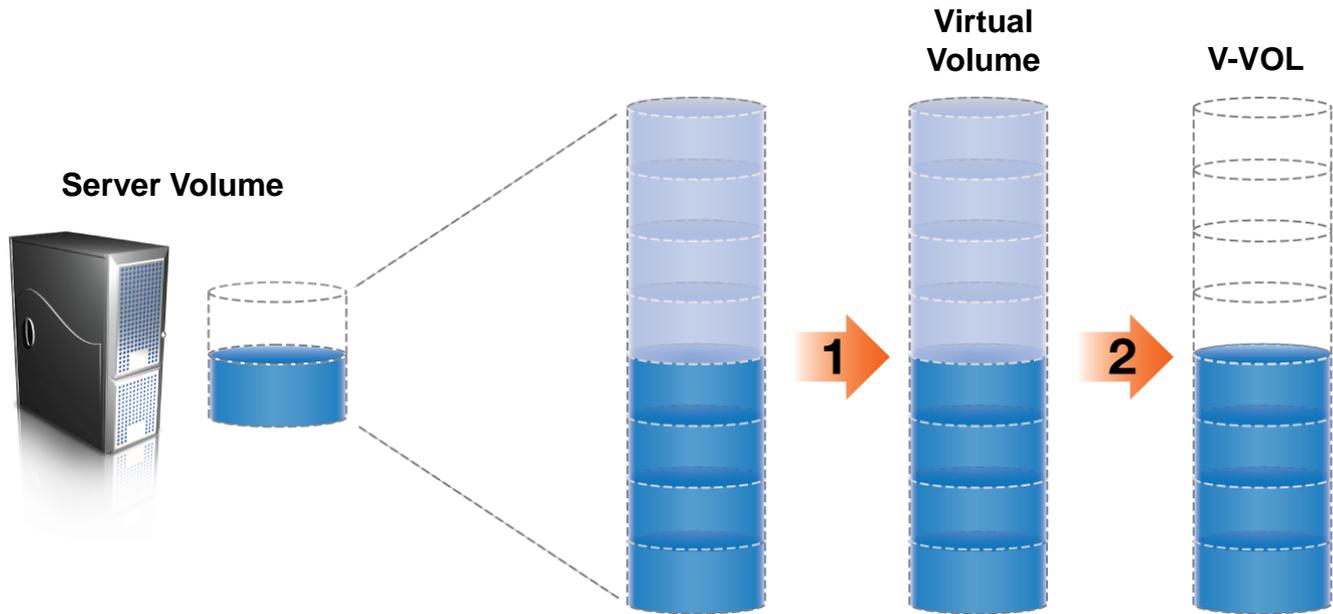
- Virtualize devices into a pool of capacity and allocate by pages
- Dynamically provision new servers in seconds
- Eliminate allocated but unused waste by allocating only the pages that are used
- Extend Dynamic Provisioning to external virtualized storage
- Convert fat volumes into thin volumes by moving them into the pool
- Optimize storage performance by spreading the I/O across more arms



Increase Utilization

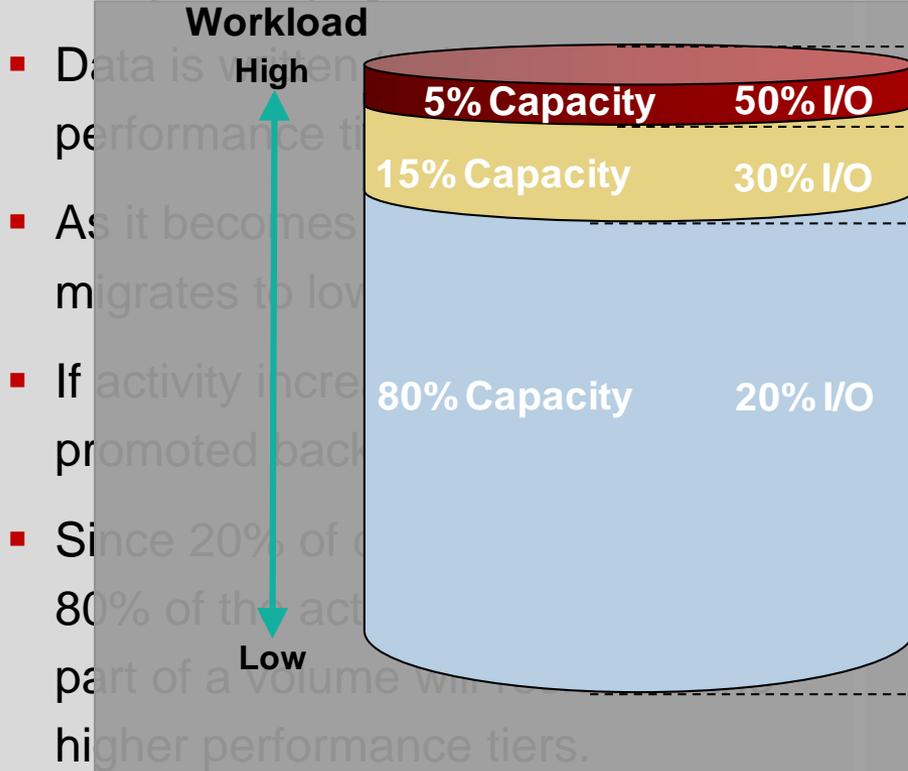
Capacity Reclamation/Zero Page Reclaim for Cloud

Step

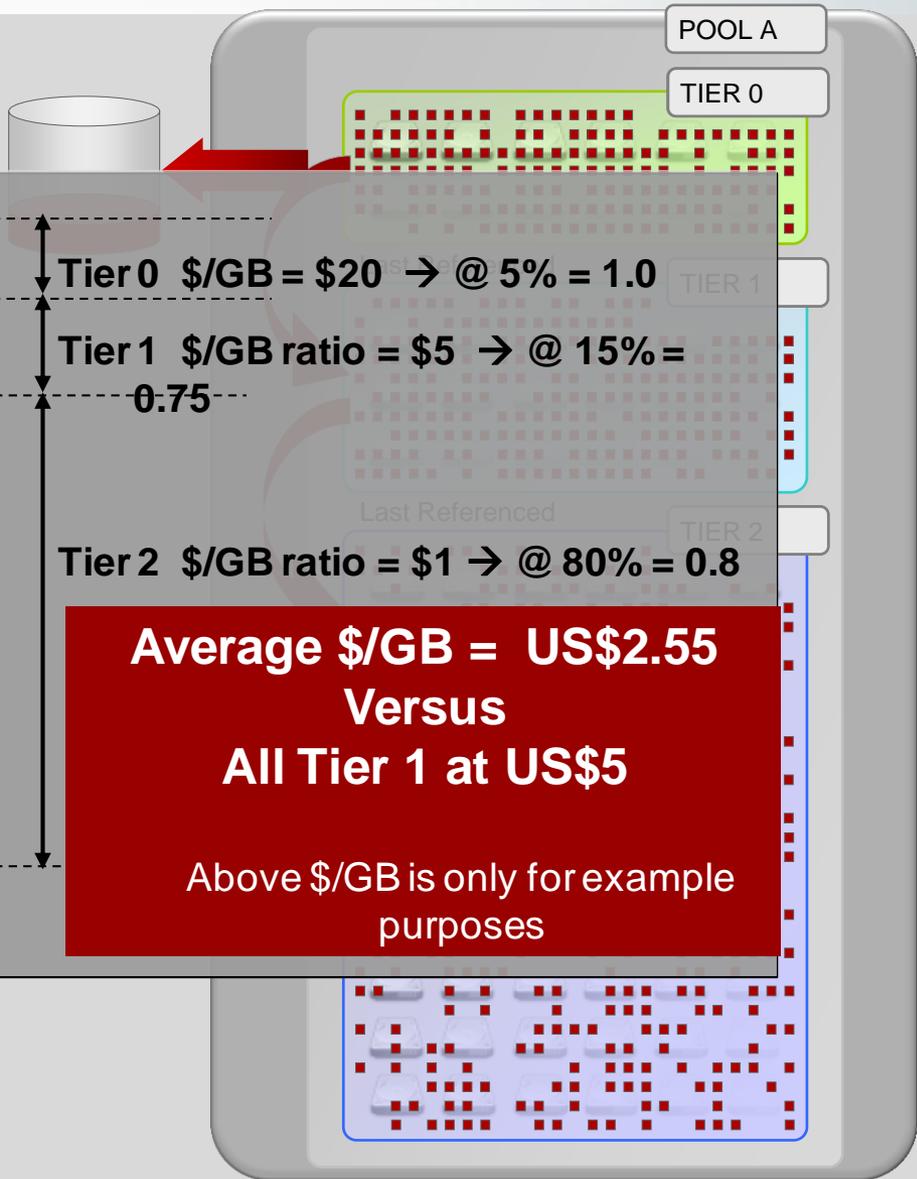


PAGE LEVEL TIERING

- Different tiers of storage are now in one pool of pages



- Data is written to higher performance tiers
- As it becomes less active, it migrates to lower performance tiers
- If activity increases, it is promoted back to higher performance tiers
- Since 20% of capacity is used for 80% of the activity, the average cost per GB is lower than if all data were in the higher performance tiers.
- An extension of Dynamic Provisioning with wide striping

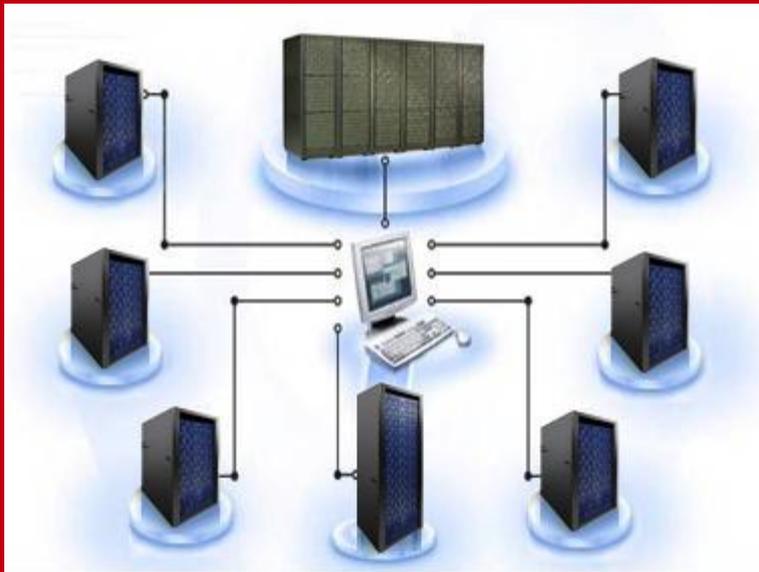


SIMPLIFY MANAGEMENT

COMMON MANAGEMENT FOR CLOUD STORAGE INFRASTRUCTURE

HITACHI
Inspire the Next

**Single management
framework with the breadth
to manage storage, servers
and the IT infrastructure**



Solution Capabilities

- Unified management for block, file, and content across all Hitachi storage
- Common management of virtualized heterogeneous storage assets
- End-to-end visibility and correlation of applications, virtual machines, virtual servers, and logical storage devices for traditional and hypervisor environments
- Roadmap to comprehensive IT management system starting with orchestration layer

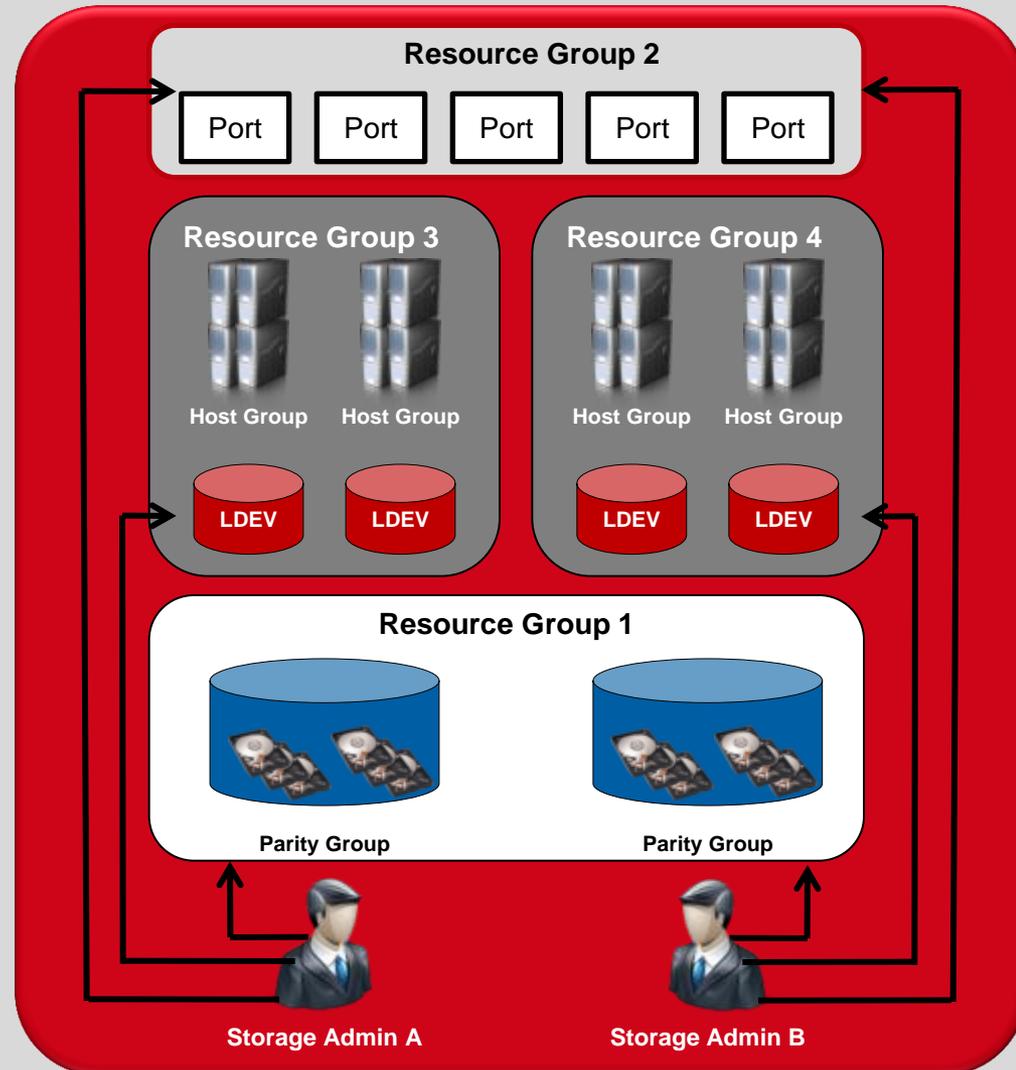
Business Value

- Properly monitor and plan resource needs for virtual server farms
- Consolidate multivendor storage management
- Integrated, open solution with orchestrated management lowers complexity and costs

Ensure QoS

Storage Logical Partitioning for multi-tenancy

- Ensure safe multitenancy and quality of service through partitioning of cache and port address space
- Allows greater administration granularity
- Flexible partitioning
 - Multitenancy
 - Cloud

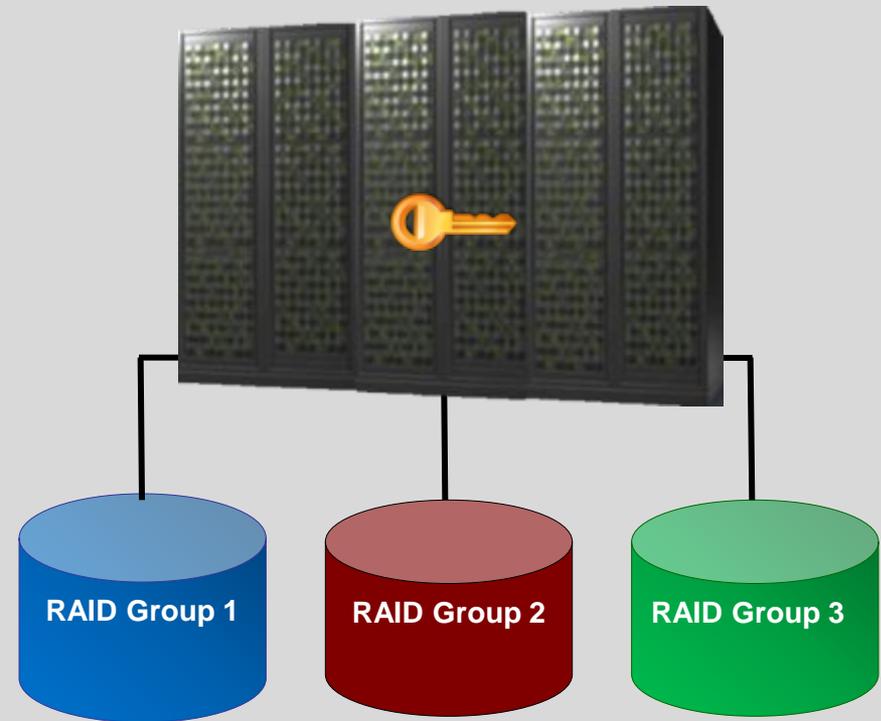


Data Security/Privacy

Encryption for multi-tenancy

HITACHI
Inspire the Next

- **Controller-based encryption for at-rest data**
- **Encryption is performed in hardware with no degradation in throughput**
- **Green encryption**
- **New encryption mode of operations**
 - XTS-AES 256 bit encryption
- **Expanded key support**
 - 32 keys per array
 - Encryption as access control

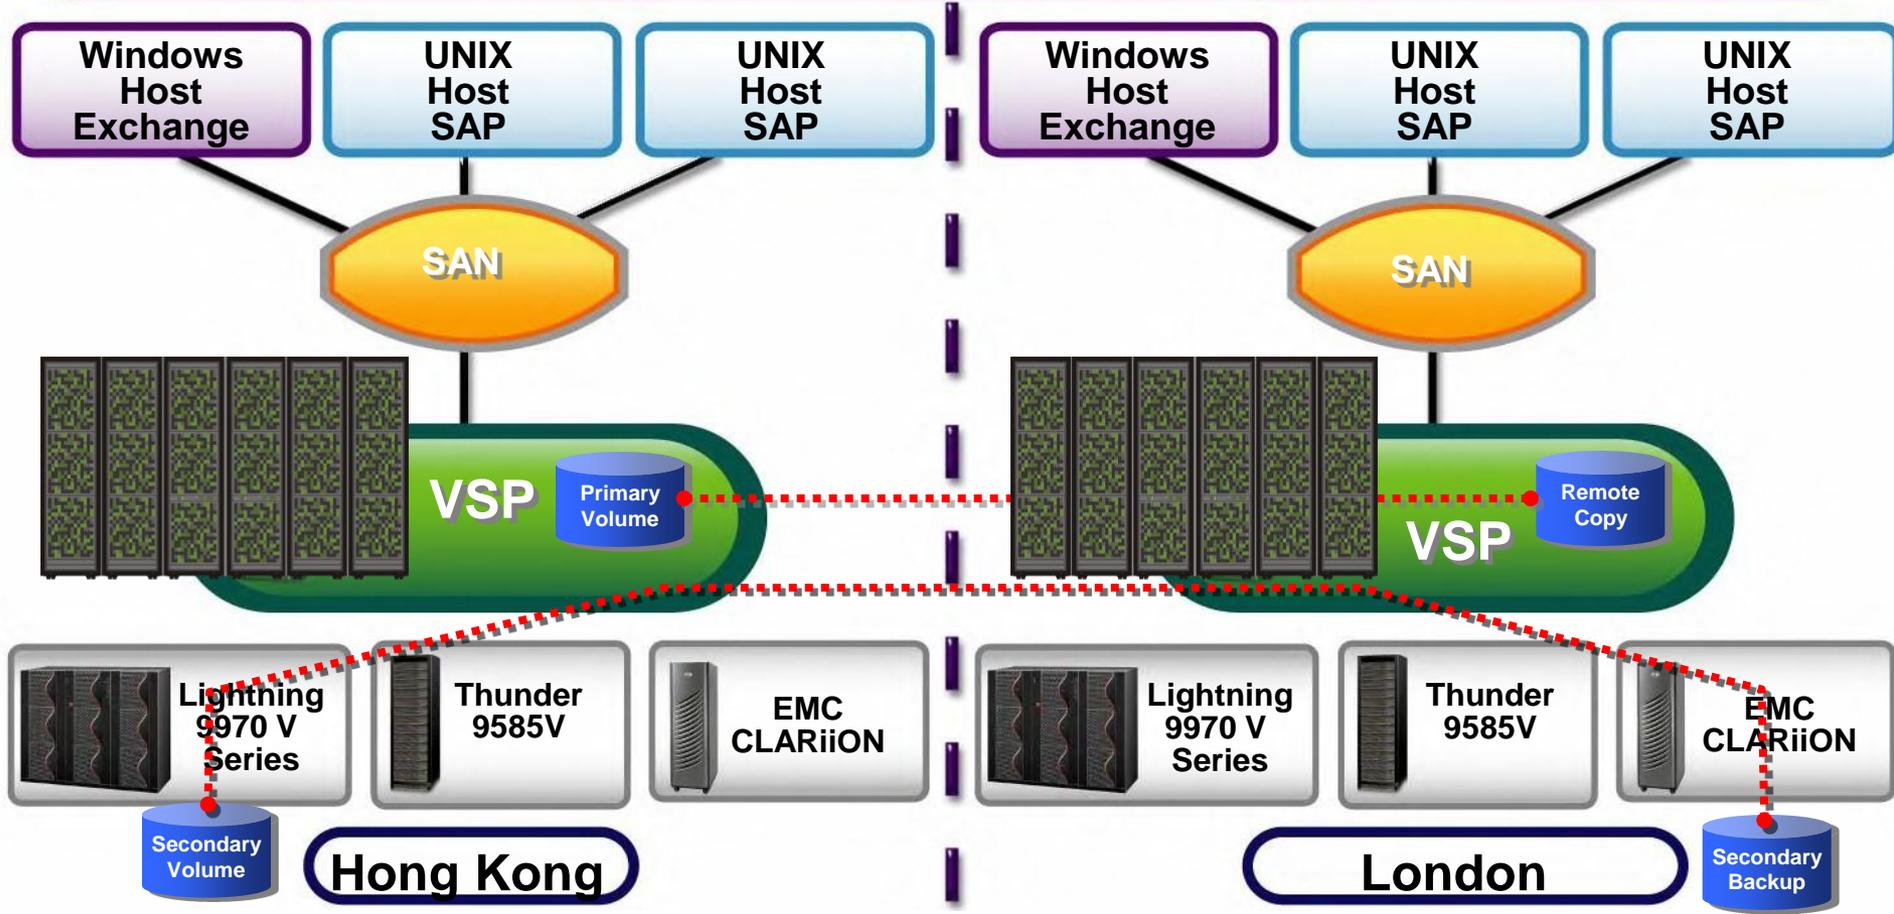


BUSINESS CONTINUITY SOLUTION

COMMON REMOTE REPLICATION SOLUTION FOR CLOUD INFRASTRUCTURES



Remote Replication with Universal Replicator Software



CLOUD-ENABLING ENGINES

HITACHI



- Isolation and security
 - Logical partitioning
- Application throughput
 - Hybrid I/O
- Workload performance
 - SMP Compute
- Reliability and low risk
 - N+ 1



**Hitachi Compute
Blade 2000**

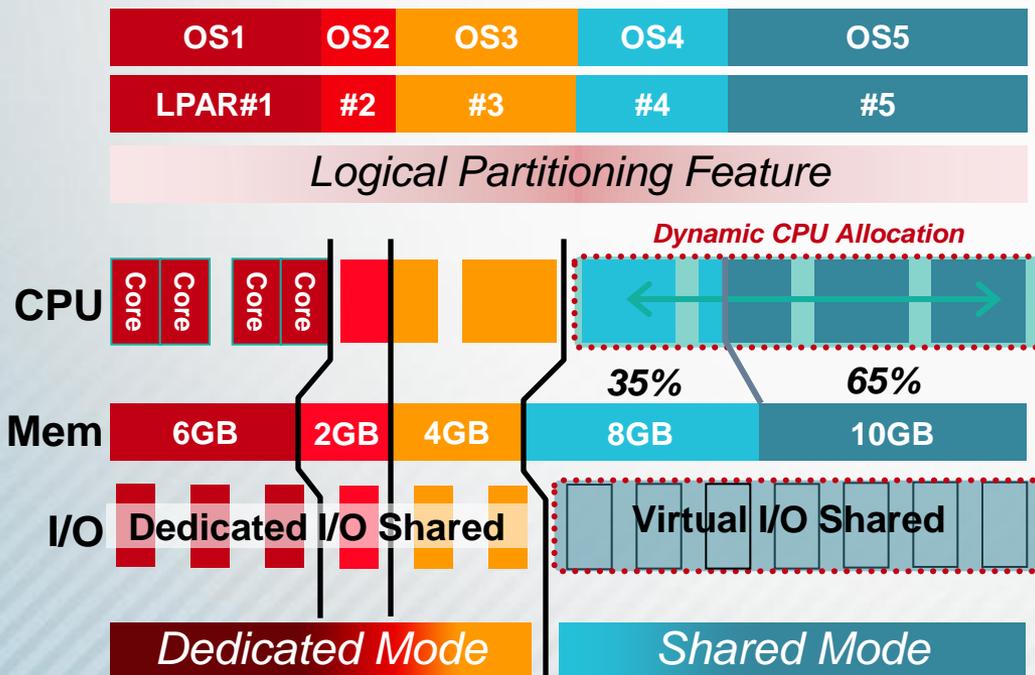
Virtualization Aware
Applications



Non-Virtualization Aware
Applications

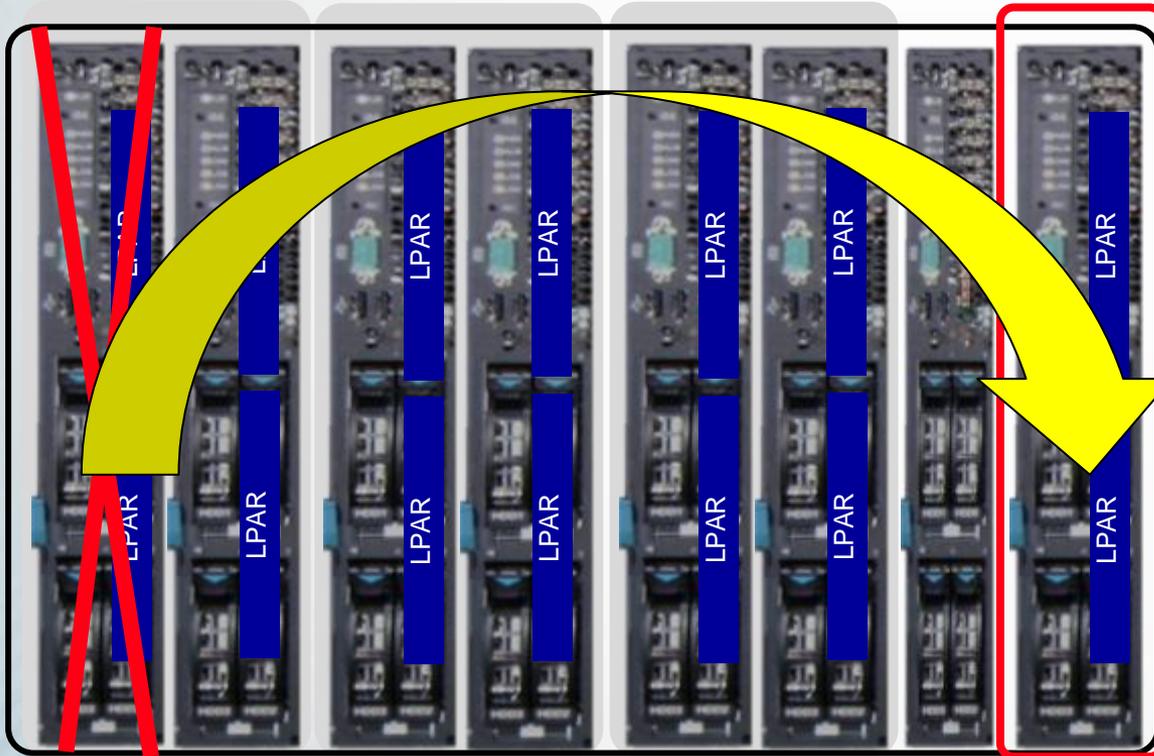
Dedicated and Shared Resources

- **Dedicated/Shared CPU Modes:** Exclusive CPU assignment for performance versus dynamic CPU load balancing
- **Dedicated/Shared I/O Resources:** Exclusive or shared I/O for SAN and network access



(N+1) BUSINESS ADVANTAGE INCREASE SERVICE AVAILABILITY

Hitachi Compute Blade 2000



N+1 Cold Standby

Minimize management resource

- When hardware failure occurs, the failover to the standby blade will be performed automatically and seamlessly.

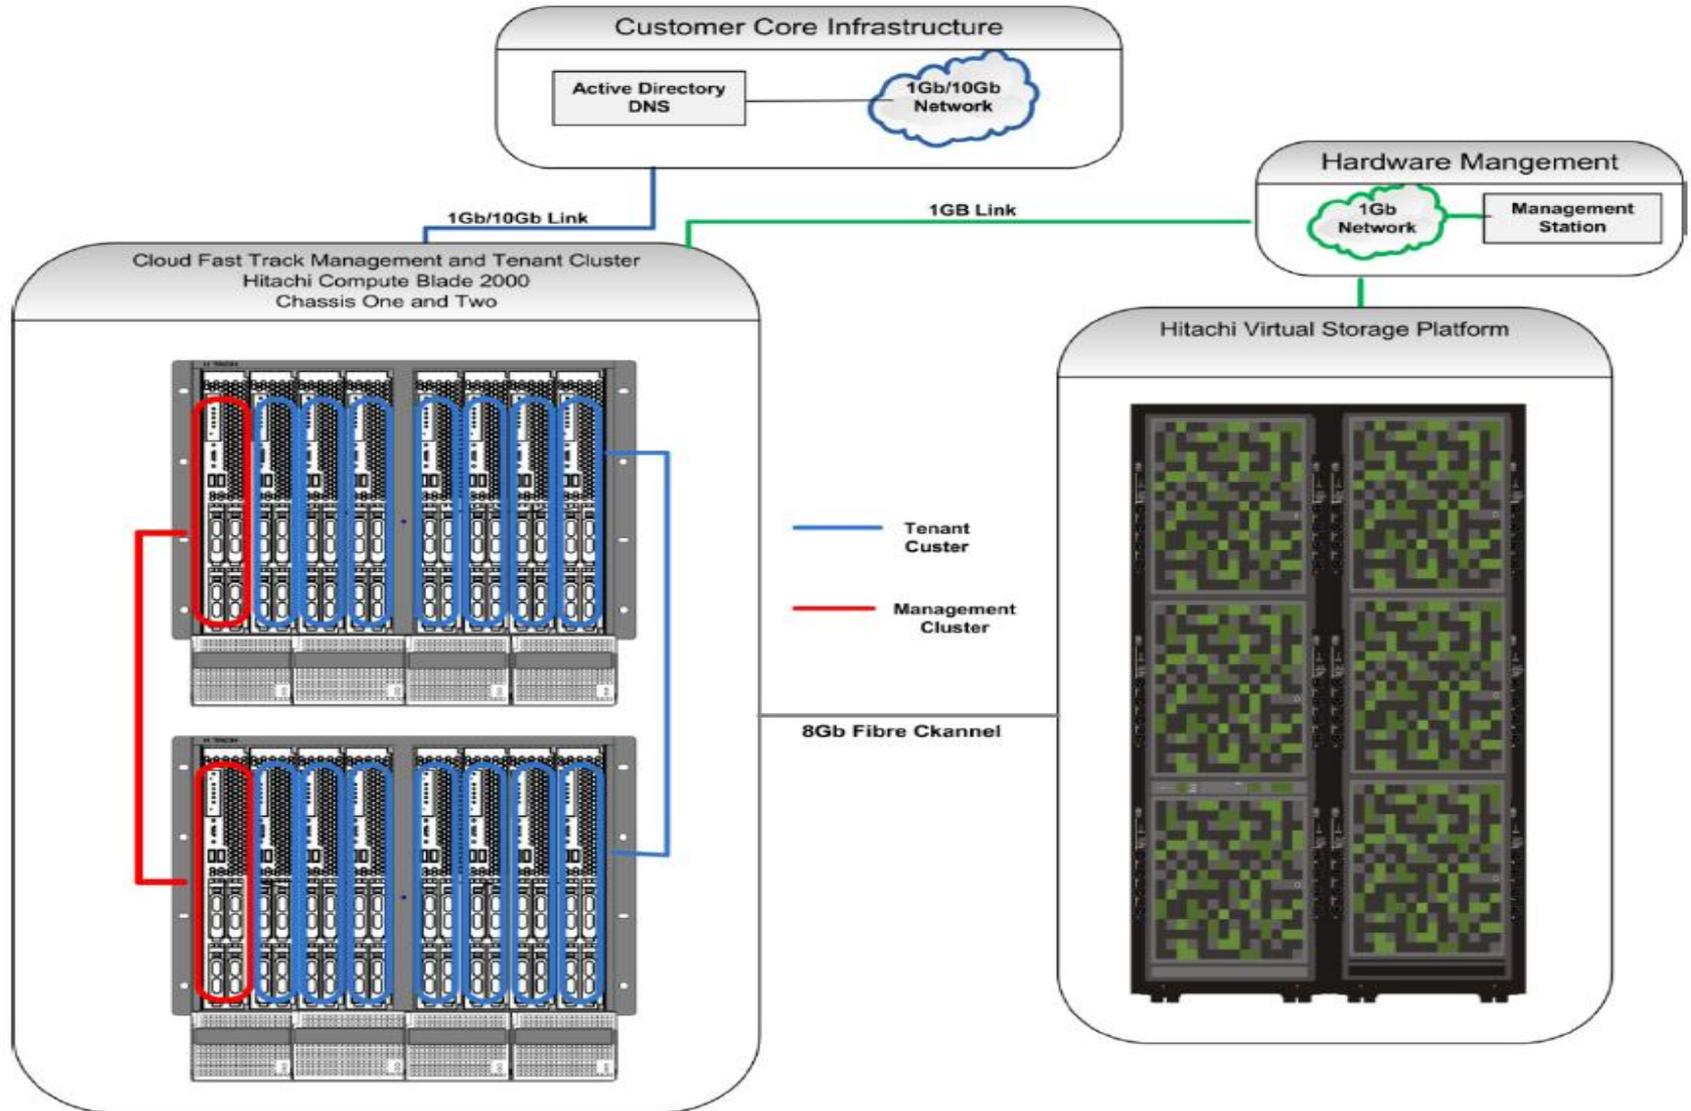
Minimize complexity

- Hardware-related Information will be transferred automatically over the blades during the failover.
- No need to change SAN or LAN configuration after the failover



**MICROSOFT
HYPER-V CLOUD
FAST TRACK**
REFERENCE ARCHITECTURE

HITACHI SOLUTION COMPONENTS FOR HYPER-V CLOUD



SECOND CONFIGURATION ARCHITECTURE

- Fault-tolerant
- Scalable

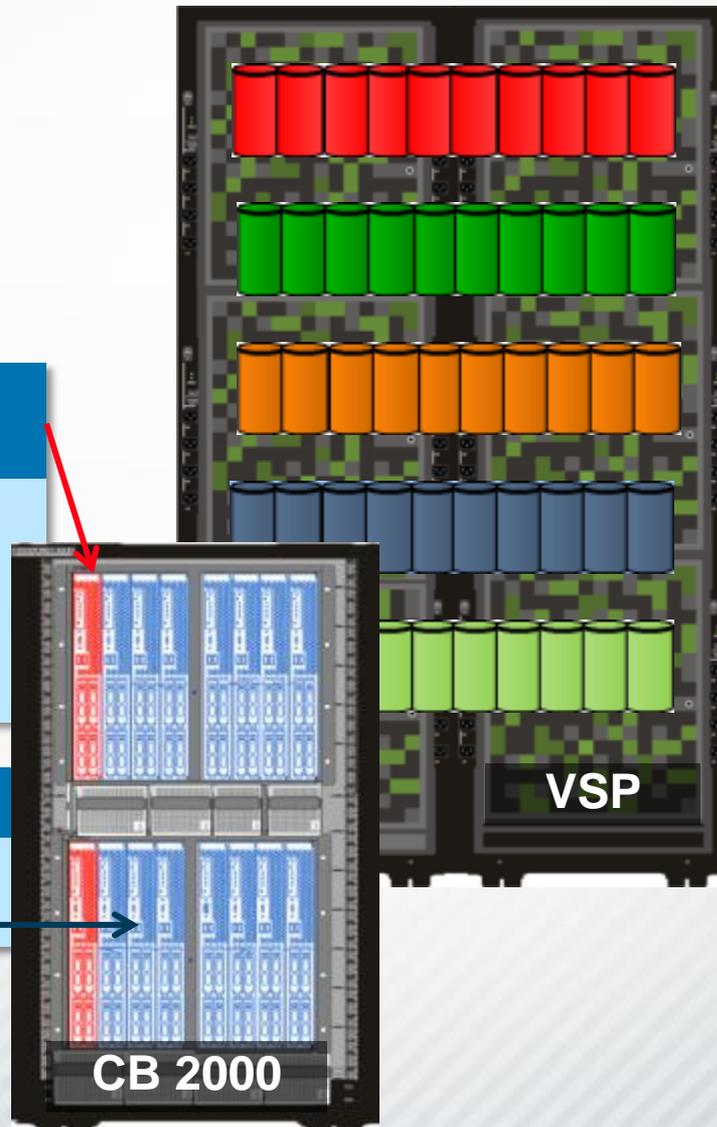
2-NODE MANAGEMENT CLUSTER

System Center

- Operations Manager
- Configuration Manager
- Virtual Machine Manager with Self-Service Portal

14-NODE TENANT

- Windows Server 2008 R2 Hyper-V Failover Clustering



-  Dynamic Provisioning Pool 1 - OS and VHDs
-  Dynamic Provisioning Pool 2 - Application Data
-  Dynamic Provisioning Pool 3 - Application Data
-  Dynamic Provisioning Pool 4 - Application Data
-  Dynamic Provisioning Pool 5 - Backup

HITACHI SOLUTION COMPONENTS FOR HYPER-V CLOUD

SOFTWARE	8 AND 16 SERVERS	OPTIONS AND UPGRADES
MICROSOFT MANAGEMENT	<ul style="list-style-type: none"> ▪ System Center Suite <ul style="list-style-type: none"> ▪ Virtual Machine Manager (SCVMM) ▪ Self Service Portal (SSP) ▪ Operation Manager (SCOM) ▪ Configuration Manager (SCCM) ▪ Deployment Toolkit 	<ul style="list-style-type: none"> ▪ System Center Opalis ▪ Volume Shadow Copy Service (VSS) ▪ Data Protection Manager
WINDOWS OS HYPER-V SQL SERVER	<ul style="list-style-type: none"> ▪ Windows Server 2008 R2 and Hyper-V ▪ SQL Server 2008 R2 ▪ Windows Deployment Services 2008 R2 	
HITACHI MANAGEMENT AND INTEGRATIONS	<ul style="list-style-type: none"> ▪ CB 2000 Server Conductor software ▪ Hitachi Command Suite ▪ Hitachi Dynamic Provisioning ▪ Hitachi Storage Navigator Modular 2 ▪ SCOM server management pack ▪ SCOM storage management pack (AMS only) ▪ VMM SSP Powershell integration (AMS only) 	<ul style="list-style-type: none"> ▪ VSS hardware provider

- **Virtual machine high availability**—With the Hitachi Compute Blade 2000 running Microsoft Hyper-V failover clustering, the virtual machines deployed in the failover cluster are made highly available. In case one of the blades in the cluster fails, the virtual machines residing on that blade automatically fail over to another blade in the cluster.
- **Virtual machine live migration**—The administrator can live migrate a virtual machine from one blade in the cluster to another blade. This can be used to balance workloads or to move the virtual machine before performing server maintenance.
- **Template based virtual machine provisioning**—Virtual machine templates allow administrators to deploy virtual machines rapidly.
- **Self-service virtual machine provisioning**—Administrators can delegate authority to other users or a group of business owners. This allows them to create virtual machines using a web interface, based on a set of predetermined templates.
- **Integration with System Center Operations Manager**—Hitachi provides monitoring packs for the Hitachi Compute Blade 2000. This enables the administrator to be notified of any alerts that require attention.

REFERENCE ARCHITECTURE

HITACHI COMPUTE BLADE 2000 SERVER ARCHITECTURE

Chassis 1

Server Blade	Server Name	Role
Blade 0	Chassis One-CFT-Node0	Hyper-V Host Hyper-V Management
Blade 1	Chassis One-CFT-Node1	Hyper-V Host Tenant
Blade 2	Chassis One-CFT-Node2	Hyper-V Host Tenant
Blade 3	Chassis One-CFT-Node3	Hyper-V Host Tenant
Blade 4	Chassis One-CFT-Node4	Hyper-V Host Tenant
Blade 5	Chassis One-CFT-Node5	Hyper-V Host Tenant
Blade 6	Chassis One-CFT-Node6	Hyper-V Host Tenant
Blade 7	Chassis One-CFT-Node7	Hyper-V Host Tenant

Chassis 2

Server Blade	Server Name	Role
Blade 0	Chassis One-CFT-Node0	Hyper-V Host Hyper-V Management
Blade 1	Chassis One-CFT-Node1	Hyper-V Host Tenant VM's
Blade 2	Chassis One-CFT-Node2	Hyper-V Host Tenant VM's
Blade 3	Chassis One-CFT-Node3	Hyper-V Host Tenant VM's
Blade 4	Chassis One-CFT-Node4	Hyper-V Host Tenant VM's
Blade 5	Chassis One-CFT-Node5	Hyper-V Host Tenant VM's
Blade 6	Chassis One-CFT-Node6	Hyper-V Host Tenant VM's
Blade 7	Chassis One-CFT-Node7	Hyper-V Host Tenant VM's

Each blade runs Microsoft Windows 2008 R2, Datacenter Edition, with two 4-core Xeon X5640 2.66GHz processors and 72GB of RAM.

This Hitachi Compute Blade 2000 configuration can host a total of 32 virtual machines per server blade. This means there can be a total of 448 virtual machines in a fourteen node Hyper-V failover cluster, as described in this reference architecture.

Each server blades Microsoft Hyper-V host OS and paging files are located on two local 146GB SAS drives, configured as RAID-1 for high performance and availability.

To satisfy the requirement to support up to 448 virtual machines this reference architecture uses the following:

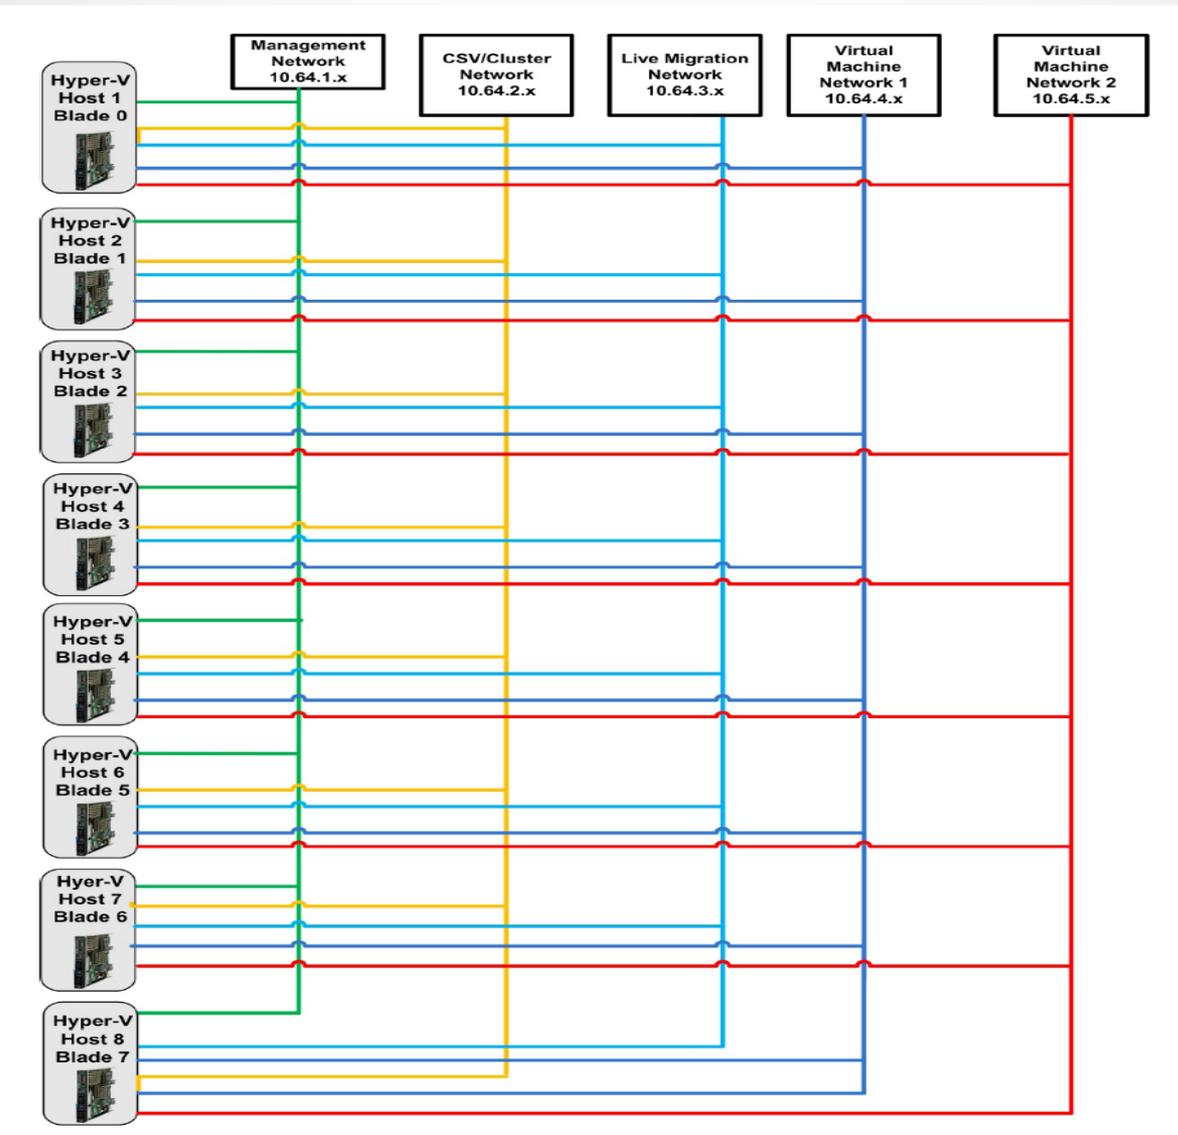
- A RAID-6 (6D+2P) configuration of 600GB 15K RPM SAS drives and 146GB 15K RPM SAS drives to host the CSVs.
- A RAID-6 (6D+2P) pool consisting of 2TB 7.2K RPM SATA drives to support backup of the CSV volumes.

The four pools created with Hitachi Dynamic Provisioning for this solution were created from 59 RAID-6 (6D+2P) groups on the Hitachi Virtual Storage Platform.

<i>Dynamic Provisioning Pool</i>	<i>Number of RAID Groups</i>	<i>Number of Drives</i>	<i>Drive Size</i>	<i>Drive Speed</i>	<i>Usable Pool Capacity (TB)</i>
0	7	56	146GB	15K RPM	22.5
1	8	64	600GB	10K RPM	6.4
2	15	120	600GB	10K RPM	48.3
3	13	104	600GB	10K RPM	41.8
4	16	128	2TB	7.2K RPM	177.6

REFERENCE ARCHITECTURE

NETWORK ARCHITECTURE



REFERENCE ARCHITECTURE

MICROSOFT SYSTEMS CENTER VIRTUAL MACHINE MANAGER R2 (SCVMM)

This environment has these roles enabled within SCVMM:

- SCVMM Administrator
- Administrator Console
- Command Shell
- SCVMM Library
- SQL Server Database (remote)

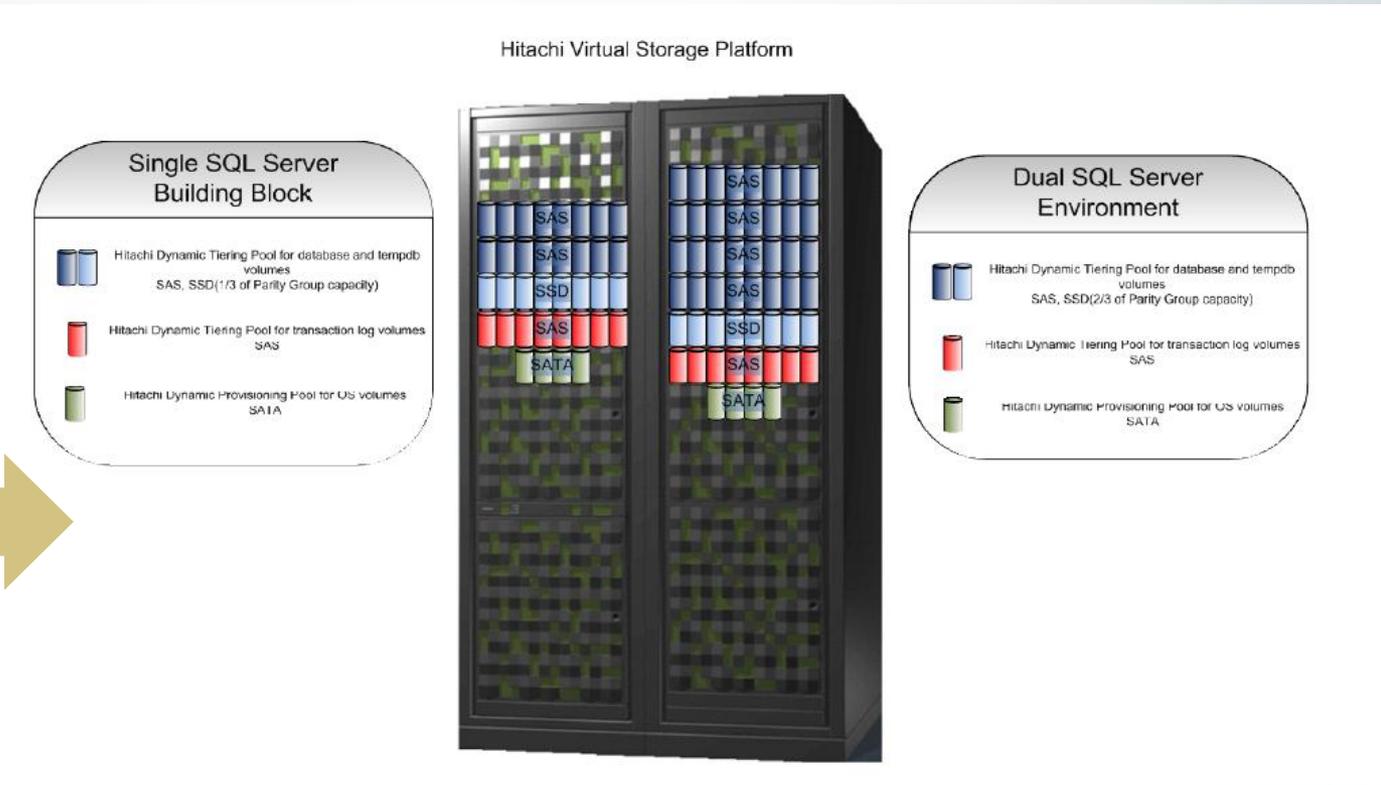
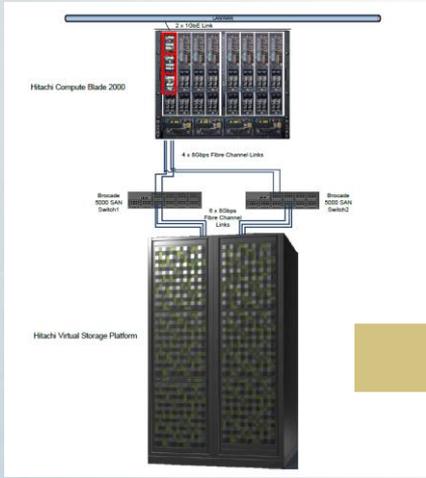
<i>Template</i>	<i>Specs</i>	<i>Network</i>	<i>OS</i>
Template 1 (Small)	1 vCPU, 2GB Memory, 50GB Disk	VLAN 2503 VLAN 2504	Microsoft Windows Server 2008 R2 SP1 Enterprise
Template 2 (Medium)	2 vCPU, 4GB Memory, 100GB Disk	VLAN 2503 VLAN 2504	Microsoft Windows Server 2008 R2 SP1 Enterprise
Template 3 (Large)	4 vCPU, 8GB Memory, 200GB Disk	VLAN 2503 VLAN 2504	Microsoft Windows Server 2008 R2 SP1 Enterprise

Virtual Machine Templates

MICROSOFT HYPER-V CLOUD FAST TRACK

MSSQL

DEPLOYING MICROSOFT SQL SERVER 2008 R2 WITH LOGICAL PARTITIONING ON THE HITACHI VIRTUAL STORAGE PLATFORM WITH HITACHI DYNAMIC TIERING



This solution's building block architecture achieves the following design goals for an OLTP workload:

- Optimize storage configuration on the Hitachi Virtual Storage Platform with **Hitachi Dynamic Tiering for best I/O throughput, database latency, and ease of management.**
- Maximize storage utilization while maintaining or improving the environment performance levels using Hitachi Dynamic Tiering.
- Deliver sustainable and acceptable levels of IOPS falling within the **1ms to 5ms response time** range for **transaction log file I/O** and **1ms to 20ms response time range for data file I/O.**
- Deliver at least 80 percent disk capacity utilization for the database volumes.

Several design factors were considered to arrive at the building block architecture, including:

- For a **very heavy user profile**, you might need **additional disks** available to support the **required IOPS**. As with any SQL server deployment, make sure to test your environment to ensure the proper number of disks are available to the database from an IOPS perspective.
- Determine the granularity of scale by the database size along with the SQL server logical partition provisioned hardware. Provision additional storage by adding RAID-5 groups for database, tempdb files, and the transaction logs. For the dynamic tiering pools, the specific drive type depends on the initial building block environment that meets both performance and capacity requirements for your environment.

- In a dynamically provisioned configuration, allocate dynamic tiering pool space by adding an LDEV from a RAID-5 group to the dynamic tiering pool. After that, create the required LUNs and assign them to a SQL server host.
- Using a 1 hour monitoring cycle, allow between 1 to 8 cycles to take place prior to establishing whether the correct ratio of tiers is devised for your environment. It may take between at least **4 monitoring cycles for your SQL server data to move to the appropriate tier in order to provide you with the best performance for your environment.**

Engineering Validation

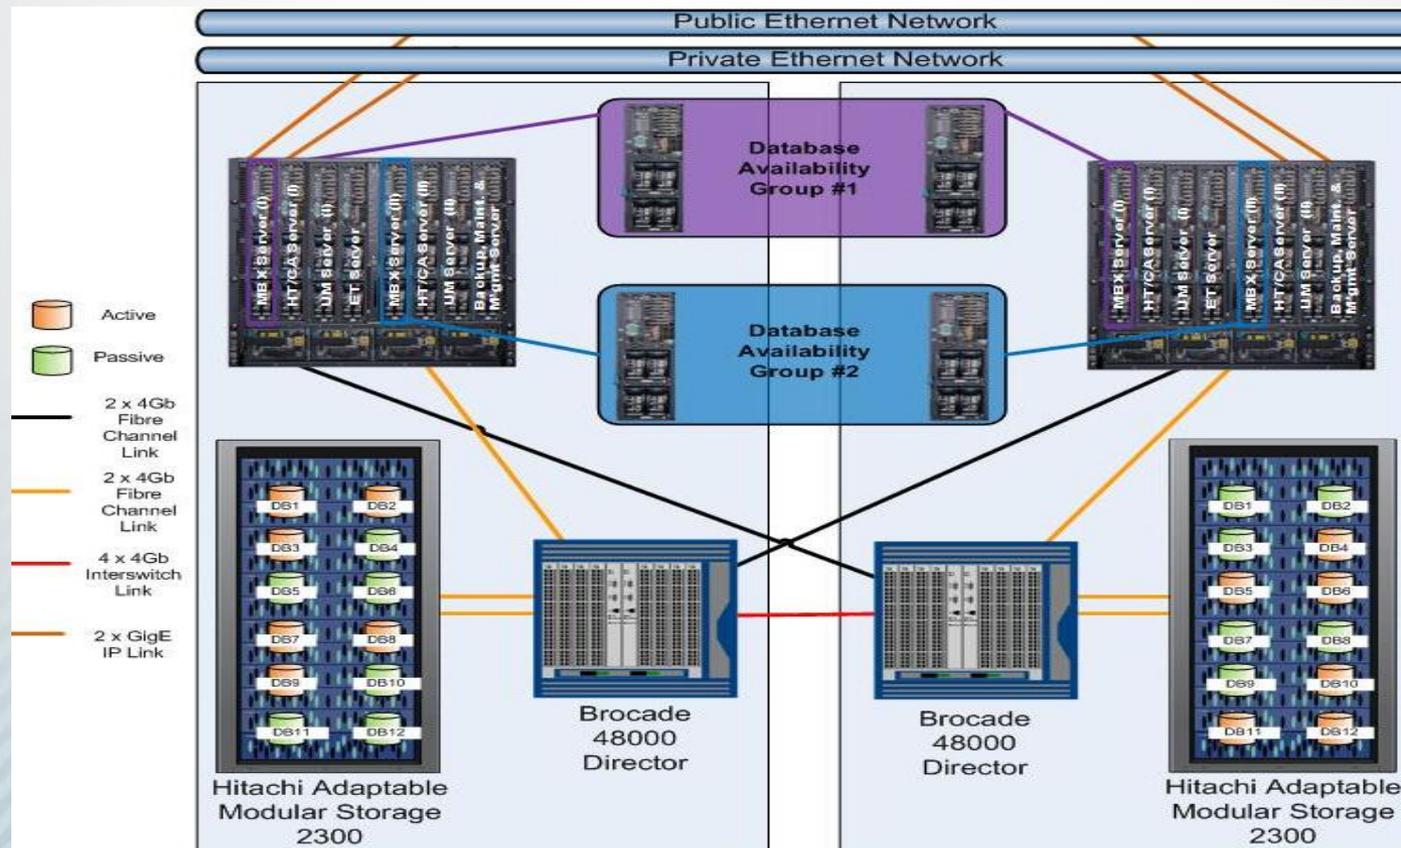
<i>Metric</i>	<i>Microsoft Success Criterion</i>	<i>Result</i>
Achieved Database IOPS	Varies	21,562 IOPS
Database Avg. Disk sec/Read	≤ 20ms	8.75ms
Database Avg. Disk sec/Write	≤ 20ms	2.5ms
Transaction Log Avg. Disk sec/Read	≤ 5ms	N/A
Transaction Log Avg. Disk sec/Write	≤ 5ms	1ms

MICROSOFT HYPER-V CLOUD FAST TRACK

EXCHANGE 2010

EXCHANGE CONVERGED SOLUTION ARCHITECTURE

- Optimized infrastructure for all Exchange 2010 server deployments
 - Design based on Microsoft ESRP configurations and best practices
 - Modular design mapping Exchange server roles to separate blades
- LPAR capabilities of logical partitioning and software partitioning of Microsoft Hyper-V



MICROSOFT EXCHANGE SOLUTION REVIEWED PROGRAM (ESRP)

- Hitachi Data Systems is an ESRP 3.0 participant
 - Resource for planning and designing your storage solution
- Extensive results published with large environment examples
 - Hitachi Dynamic Provisioning
 - SAS drive architecture
 - RAID 5

HDS

5000 Mailboxes or greater

- [Hitachi Adaptable Modular Storage 2300 Dynamically Provisioned 20,000 Mailbox Exchange 2010 Mailbox Resiliency Storage Solution \(Aug 31, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 5,800 mailbox Exchange 2010 Mailbox Resiliency Storage Solution \(Aug 3, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 9,440 mailbox Exchange 2010 Mailbox Resiliency Storage Solution \(Aug 3, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 28,000 mailbox Exchange 2010 Mailbox Resiliency Storage Solution \(Aug 3, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2500 Dynamically Provisioned 100,800 mailbox Exchange 2010 Mailbox Resiliency Storage Solution \(Aug 3, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2300 Dynamically Provisioned 34,300 User Exchange 2010 Resiliency Storage Solution \(May 15, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2500 Dynamically Provisioned 68,800 User Exchange 2010 Resiliency Storage Solution \(May 15, 2010\)](#)
- [Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 17,200 User Exchange 2010 Mailbox Resiliency Storage Solution \(May 4, 2010\)](#)

<http://technet.microsoft.com/en-us/exchange/ff182054.aspx>

- **Hitachi Dynamic Provision (HDP)**
 - Reduce management cost and eases user management. Does not require separation of heavy users, average users and light users.
 - Avoids performance “Hot Spots” by striping across a set of disks.
 - Allows addition of storage into the pool as you grow.

- **Double Protection (RAID + Exchange 2010 DAG)**
 - RAID protection reduces the number of Database copies required within the same data center.
 - Hitachi Dynamic Sparing Technology – Avoid RAID rebuild before drive failure.
 - Ease of administration by REDUCING drive fail-over/recovery scenario.

- **System Efficiency (Green Data Center)**
 - Reduces network traffic by minimizing the # of database copies required within the same data center.
 - Preserves system resource (CPU, Memory) on Exchange server NOT IO operations.

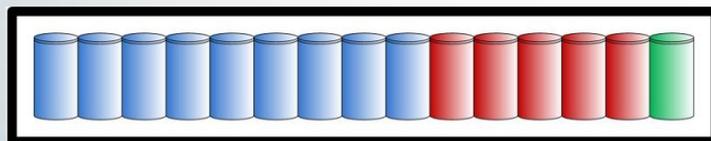
- **Virtualization**
 - Ability to scale, share, and repurpose storage when requirement changes.
 - Ability to virtualize behind Tier-1 VSP

Mailbox Count	2500
Number of Data Disks	9
Number of Log Disks	5
RAID type databases	RAID-5(8+1)
RAID type logs	RAID-5(4+1)*
Spares	1
Disk type	450GB, 15K SAS
Mailbox size	1 GB
IO Profile	0.12

AMS2100 can scale up to 20,000 users with this building block.

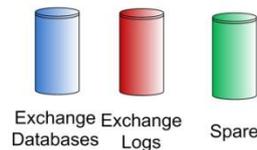
AMS2300 will scale to 40,000 users with this building block.

AMS2500 will scale 80,000 users with this building block.



RKA 1

Adaptable Modular Storage 2100
Exchange 2010 2500 Mailboxes
Serial Attached SCSI (SAS)
450 GB 15K Drives



*The pool for logs can handle 2 more sets of database disks.

Customer Use Case	AMS2100	AMS2300	AMS2500
Type of Drive	450GB SAS	450GB SAS	450GB SAS
RAID Configuration	RAID-10 (2+2)	RAID-10 (2+2)	RAID-10 (2+2)
Total # of Data Drives	108	216	432
Total # of Log Drives	12	24	48
Total # of Drives	120	240	480
Design Limitation - Capacity			
Total # of User	17200	34400	68800
Mailbox Size (GB)	1	1	1
IO Profile	0.12	0.12	0.12
# of Database Copies	2	2	2
Database Maintenance	On	On	On
Total Capacity Required (GB)	17200	34400	68800
IOPS Required	2064	4128	8256
Status	ESRP Under Review	ESRP Under Review	ESRP Under Review

High Performance, Linear Scalability

- When formatting a newly partitioned LUN, HDS recommends that for **database** and **log files** the ALU is set to **64K** and **4K** respectively.
- **Disk alignment is no longer** required when using Microsoft Windows Server 2008.
- Keep the Exchange workload **isolated from other applications**. Mixing another I/O intensive application whose workload differs from Exchange can cause the performance for both applications to degrade.
- Due to the difference in I/O patterns, isolate the Exchange database from the log groups. Create a **dedicated Dynamic Provisioning Pool for the databases and a separate Pool for the logs**.
- Hitachi Data Systems recommends RAID-5 or RAID-1+0 RAID groups for both the database pools and for the log pool. Use of **RAID-1+0** allows **more writes at a lower response time under heavier loads**. RAID-1+0 also has a **shorter RAID group rebuild time** on failure of a disk. Hitachi Data Systems does not recommend LU concatenation.
- Use **Multipathing** (Hitachi **Dynamic Link Manager** software or the **native MPIO** software with the Round-Robin algorithm), redundant switches and HBAs

MICROSOFT HYPER-V CLOUD FAST TRACK

LIVE MIGRATION

HSC for Hyper-V with Live Migration

Hitachi Storage Cluster (HSC): www.hds.com/go/HyperV

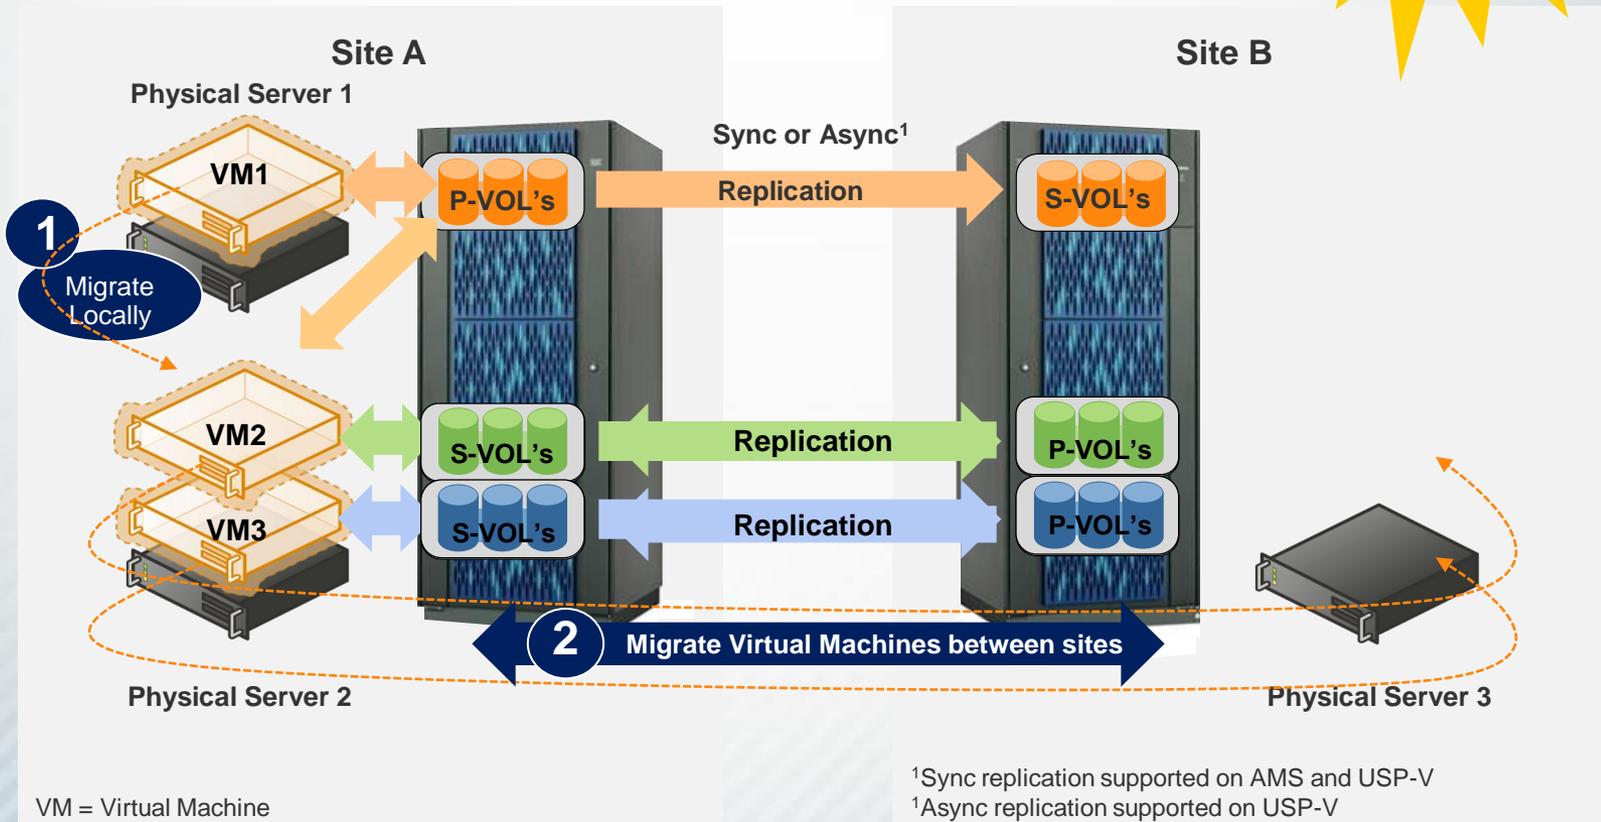
One Disaster Recovery solution for all Windows deployments:

1. Migrate Locally - across physical servers
2. Migrate Between Sites - completely transparent to workloads



Workloads:

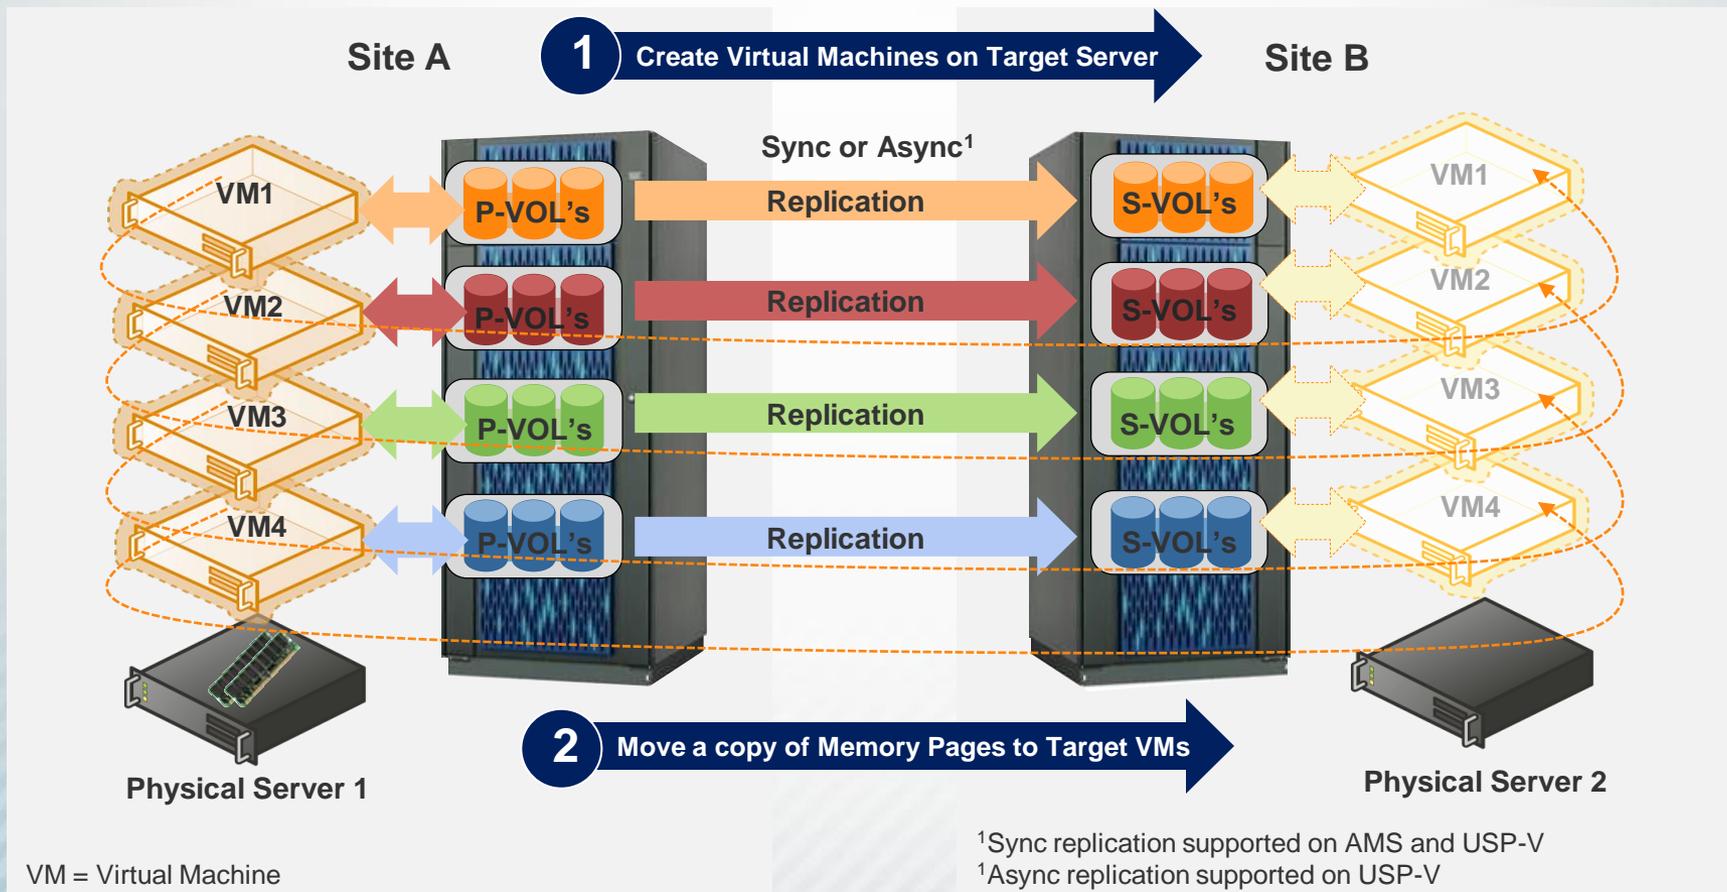
- Exchange
- SQL
- SharePoint
- Middleware
- Home Grown
- Legacy
- etc.



HSC for Hyper-V with Live Migration: Site A to Site B

Virtual Machine State & Memory Transfer

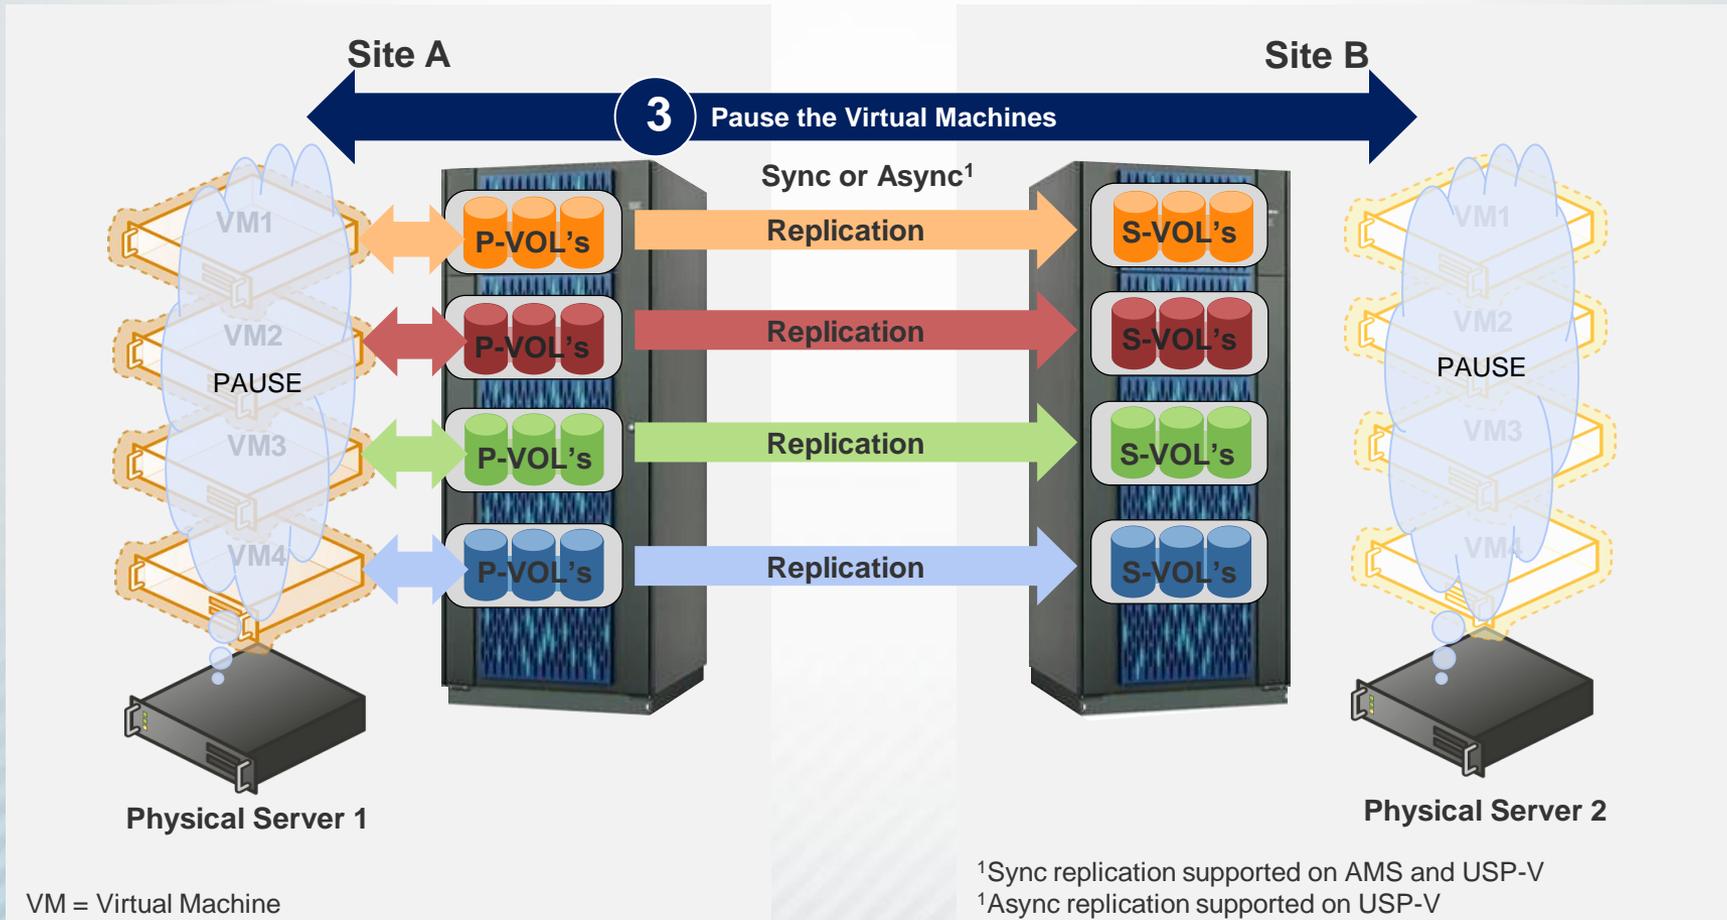
1. Create Virtual Machines on the target server
2. Move the Memory Pages from the source to the Target Server via Ethernet



HSC for Hyper-V with Live Migration: Site A to Site B

Final State Transfer & Virtual Machine Restore

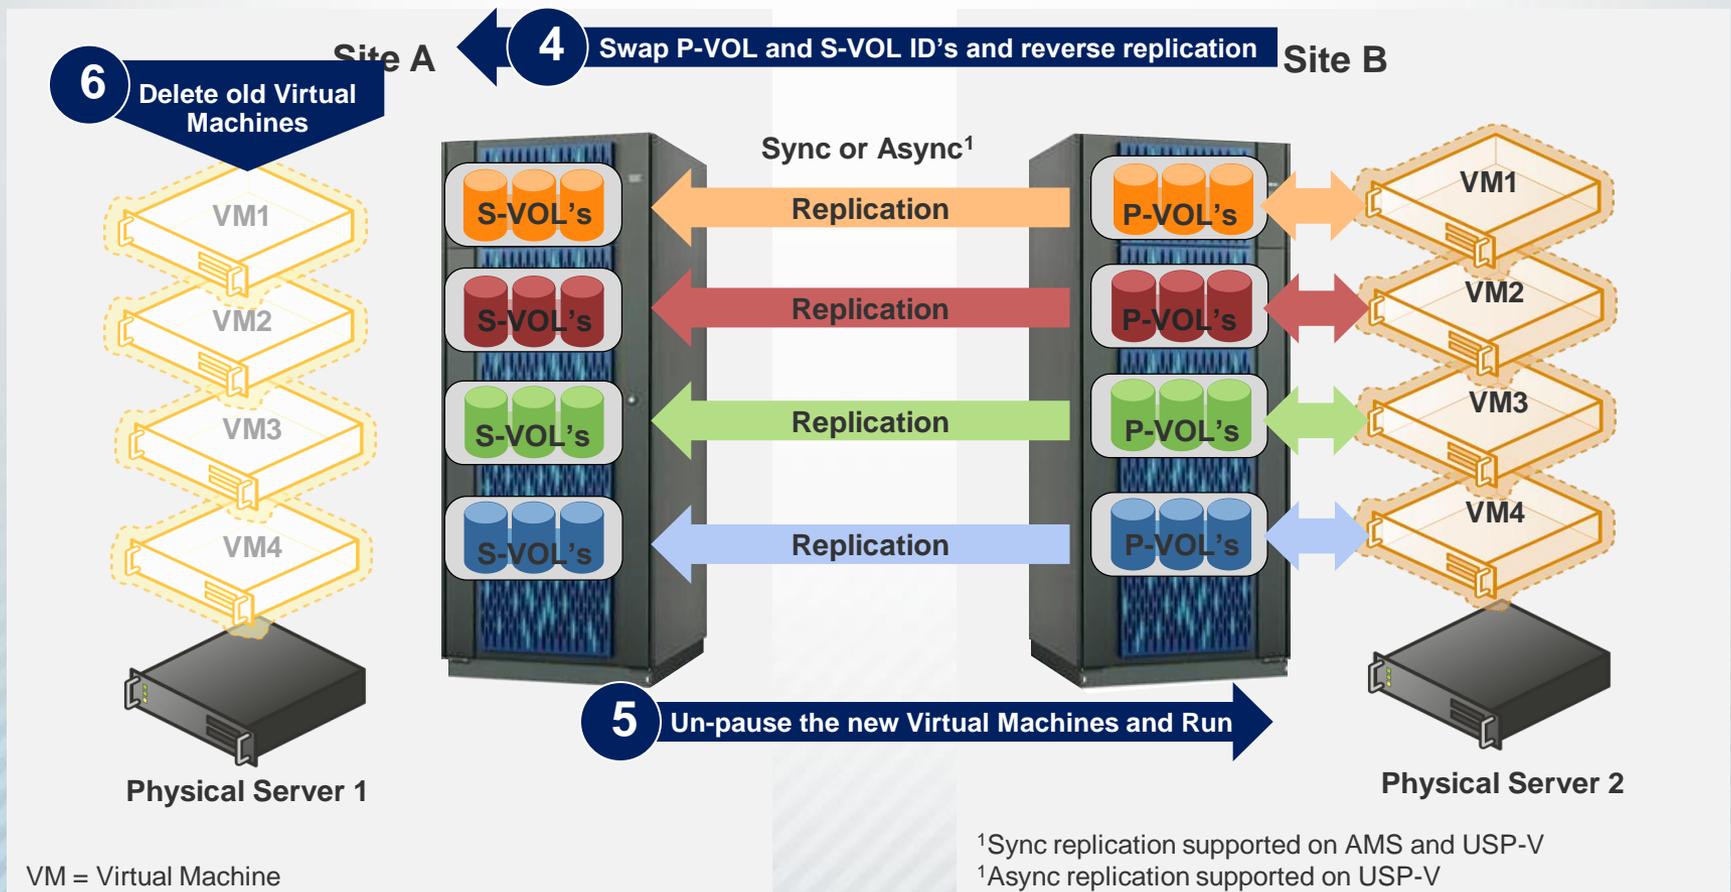
3. Pause the Virtual Machine(s)



HSC for Hyper-V with Live Migration: Site A to Site B

Reverse Replication, Swap ID's and Run

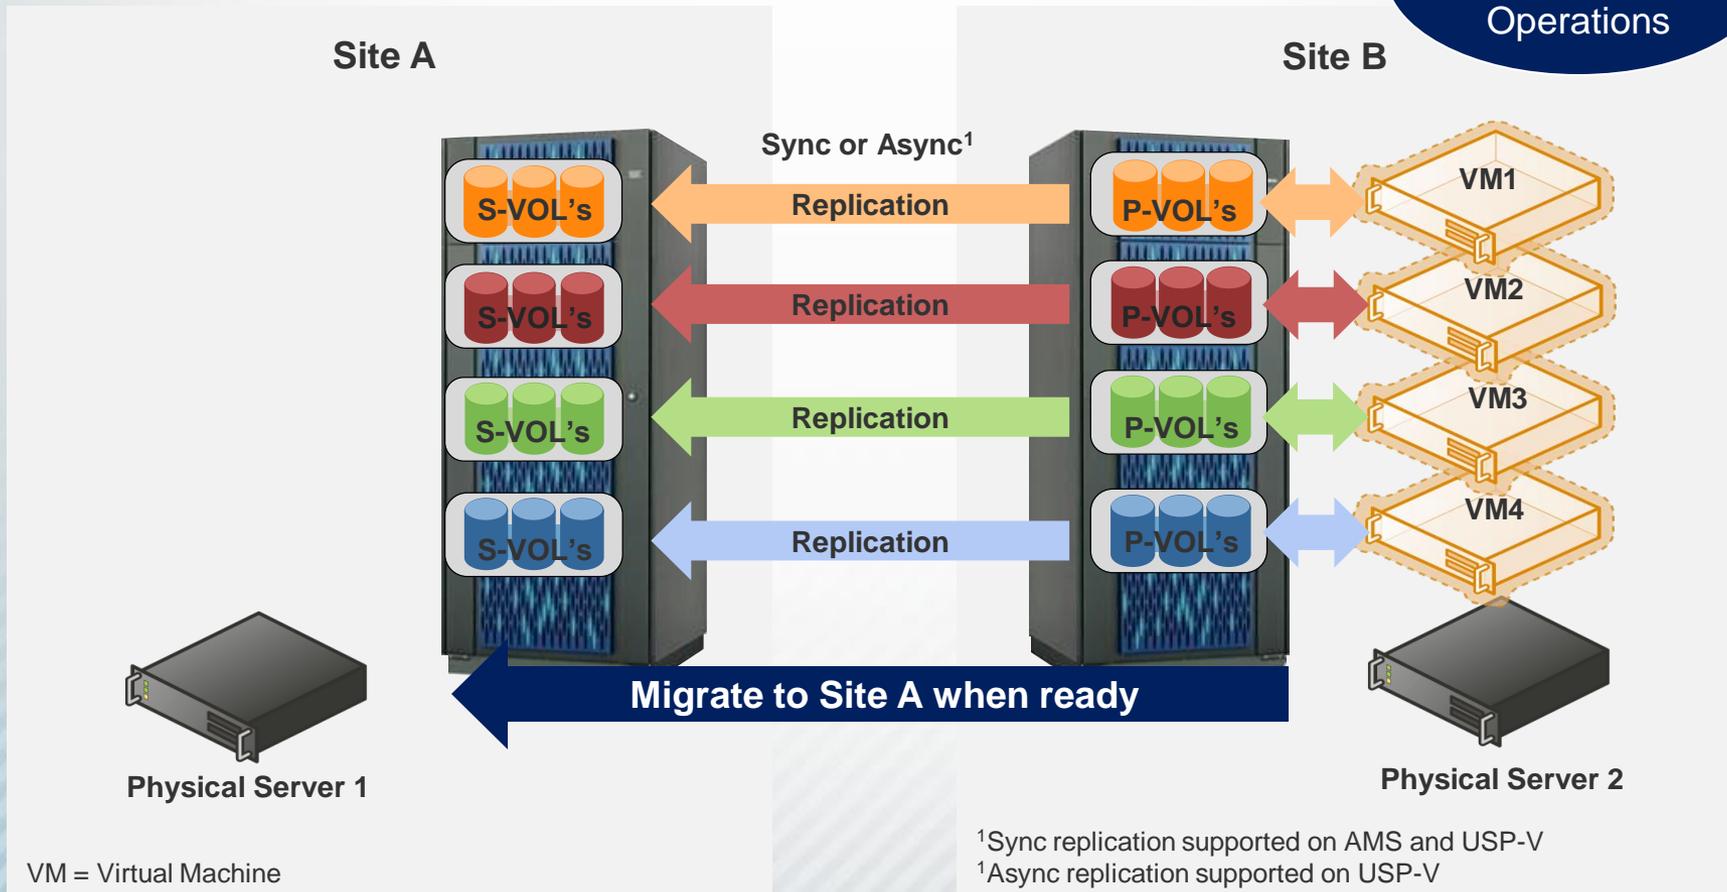
4. Swap P-VOL and S-VOL ID's and Reverse Replication
5. Un-pause the new Virtual Machine(s) and Run
6. (delete the old Virtual Machine(s))



HSC for Hyper-V with Live Migration: Site B to Site A

Benefits

- ✓ Fully Automated & Nearly Instant - completes within milliseconds
- ✓ Completely transparent to applications and workloads
- ✓ Managed from System Center Virtual Machine Manager



HSC for Hyper-V DR Failover & Recovery

Hitachi Storage Cluster (HSC): www.hds.com/go/HyperV

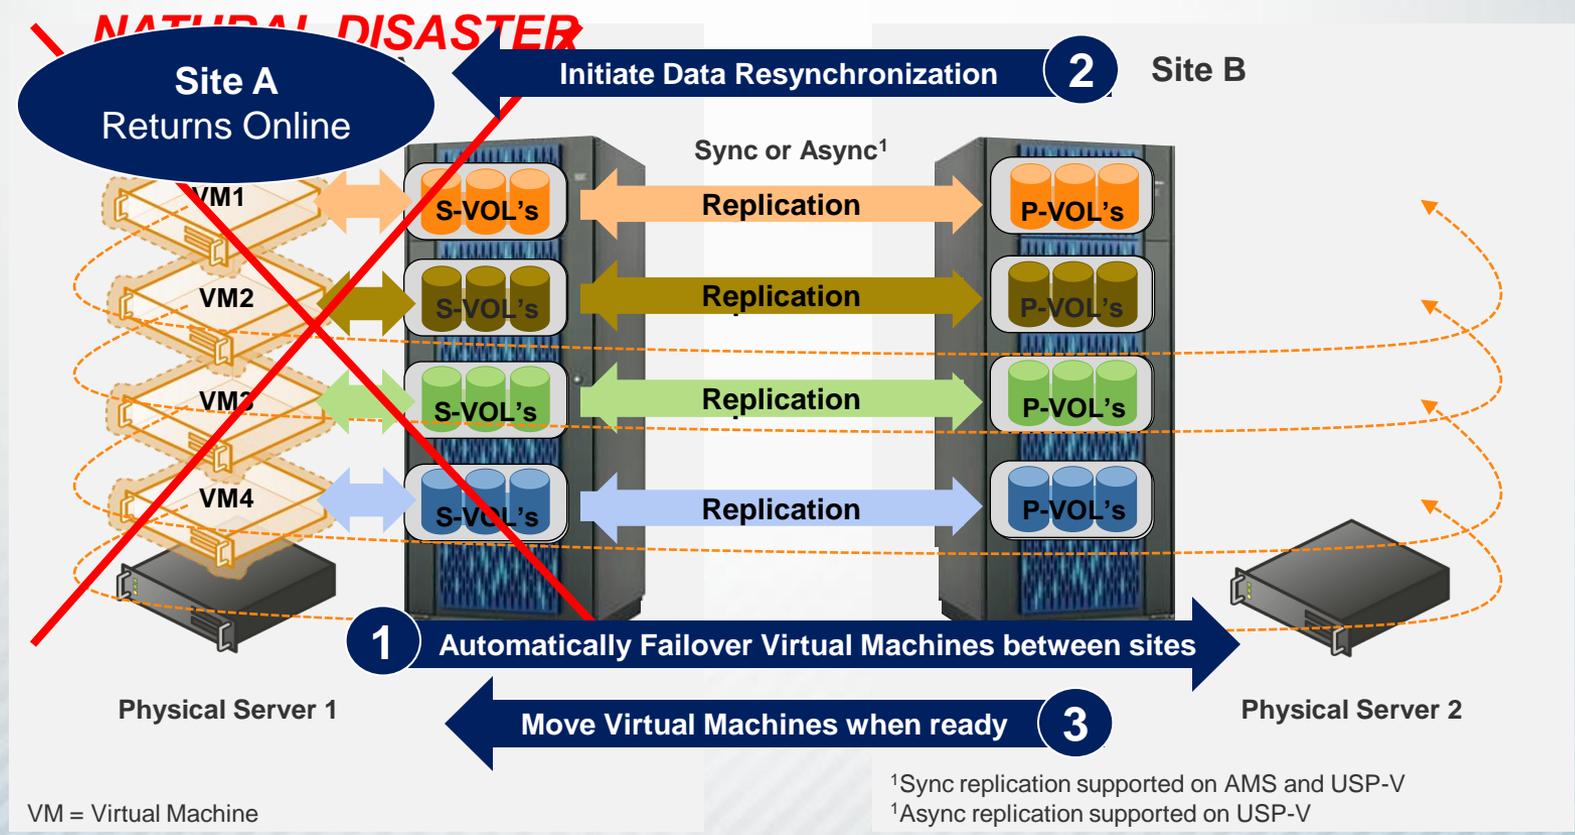
One Disaster Recovery solution for all Windows deployments:

1. Automated Failover between sites
2. Initiate Data Resynchronization upon primary site recovery
3. Move Virtual Machines back to Site A when ready

Site B
Continues All Operations

Workloads:

- Exchange
- SQL
- SharePoint
- Middleware
- Home Grown
- Legacy
- etc.



DATA ARCHIVING SOLUTION

SHAREPOINT



Hitachi Data Discovery for Microsoft SharePoint provides SharePoint users with a secure, Microsoft-oriented interface to essential productivity and business improvement features:

- **Reduce SharePoint SQL storage footprint**

Current and required SQL storage

- **Reduce SharePoint SQL backup requirements**

Move SQL data to an archive platform

- **Reclaim expensive SAN or Tier 1 capacity**

Tier data to less expensive storage media

- **Improve SharePoint SQL performance**

Query and retrieval time

Data movement transparent to user

- **Dramatically improve SharePoint scalability**

Grow and consolidate

Cloud enabled

- **Retention (For Hitachi Content Platform Only)**

- Advanced searches that include data that is not in SharePoint and is stored in Hitachi Content Platform
- Restoration of single files to SharePoint
- Setting retention, shredding and litigation holds to search results

- Decreases total cost of ownership of SharePoint, because less hardware and fewer SQL servers, licenses and Windows operating systems are required

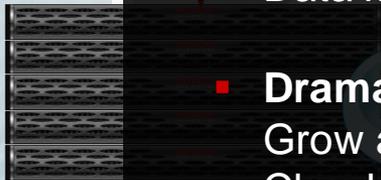
- Improves efficiency by automating data archiving and management processes

- Mitigates compliance and litigation risks by taking action on data based on policies

Encoded workflow

HCP or HNAS

Retrieval



- HNAS = Hitachi NAS Platform, powered by BlueArc®
- HCP = Hitachi Content Platform (Default Namespace Only)

GET READY – GET GOING



- Hitachi Data Systems Global Solution Services and Microsoft Professional Services
 - Planning
 - Implementation
 - Integration
 - Monitoring

Assess
and
Consult

Plan
and
Design

Install
and
Implement

Transition
and
Integrate

Manage
and
Optimize

Support

Education

**THE HITACHI PHILOSOPHY
HAS ALWAYS BEEN TO
ENHANCE SOCIETY
THROUGH TECHNOLOGY**

THANK YOU