

How to Apply the Latest Advances in Hitachi Mainframe Storage

WALTER AMSLER, SENIOR DIRECTOR, OFFICE OF TECHNOLOGY AND PLANNING

JOHN HARKER, SENIOR PRODUCT MARKETING MANAGER, HITACHI DATA SYSTEMS



AGENDA: How to Apply the Latest Advances in Hitachi Mainframe Storage

- Hitachi mainframe storage
- Extended address volume (EAV)
- Dynamic volume expansion (DVE)
- Hitachi Dynamic Provisioning for Mainframe
- Space reclamation
- Interoperability with other program products
- Summary

Leverage IBM® compatibility with Hitachi value-add

Middleware and application level


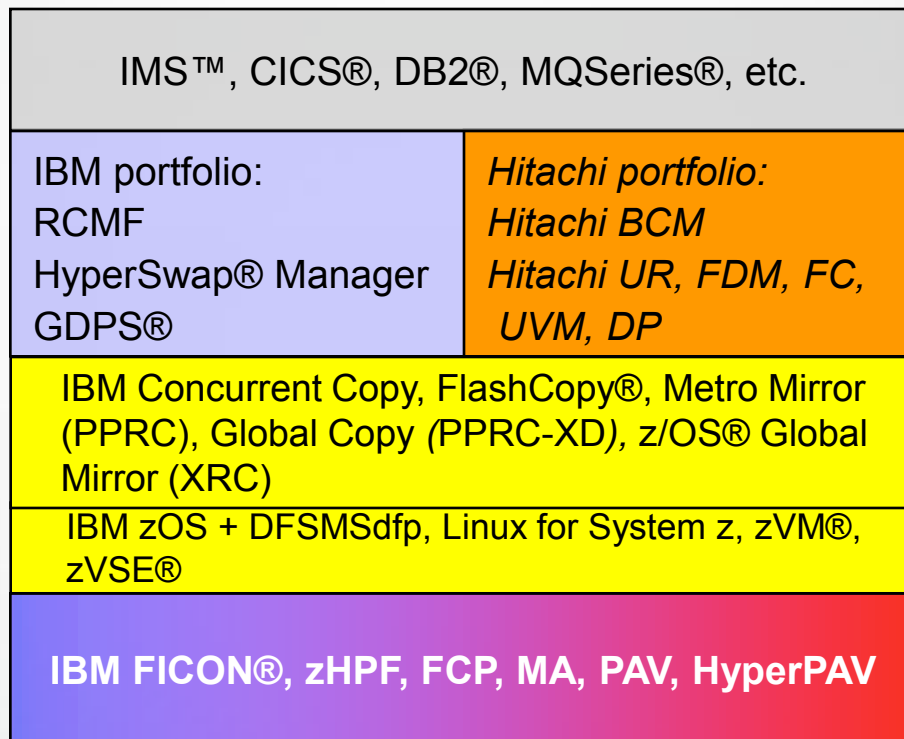
Solution level

Storage functions

Operating systems

Hardware interface

Storage subsystems



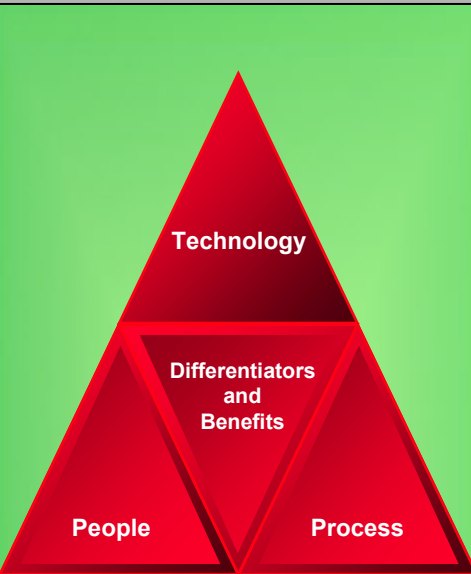
Hitachi strategy:

Development focus on Hitachi added-value solutions

IBM compatibility
(without host software)

IBM compatibility
(without host software)


IBM compatibility
(without host software)



- Consolidation, high availability and security
 - Hitachi Universal Storage Platform V and Hitachi Virtual Storage Platform designed for zero downtime operation
 - Highest performance and scalability
 - Hardware-based encryption for data at rest
- Lower TCO, increase ROA
 - Hitachi storage virtualization to allow use of lower-cost external storage
 - Hitachi Content Platform for long-term archiving on disk instead of tape
- Backup
 - In-system replication technology such as Hitachi ShadowImage® and IBM® FlashCopy® provide disk-to-disk backup to minimize backup window
- Disaster Recovery
 - Hitachi Universal Replicator for 3-data-center disaster recovery
 - Cascading and/or multi-target topologies
 - Basic HyperSwap®
 - No impact on host resources

- **Management simplification**
- **Reduced cost**
- **Improved productivity**
- **Improved RTO and RPO**

- Leverage Hitachi storage virtualization
 - Attach external low-cost storage via Hitachi Universal Volume Manager (needs FC I/F)
 - Use case 1: HSM ML2 on disk
 - Use case 2: Mainframe archive on disk
 - Use case 3: Lower cost for PiT and disaster recovery



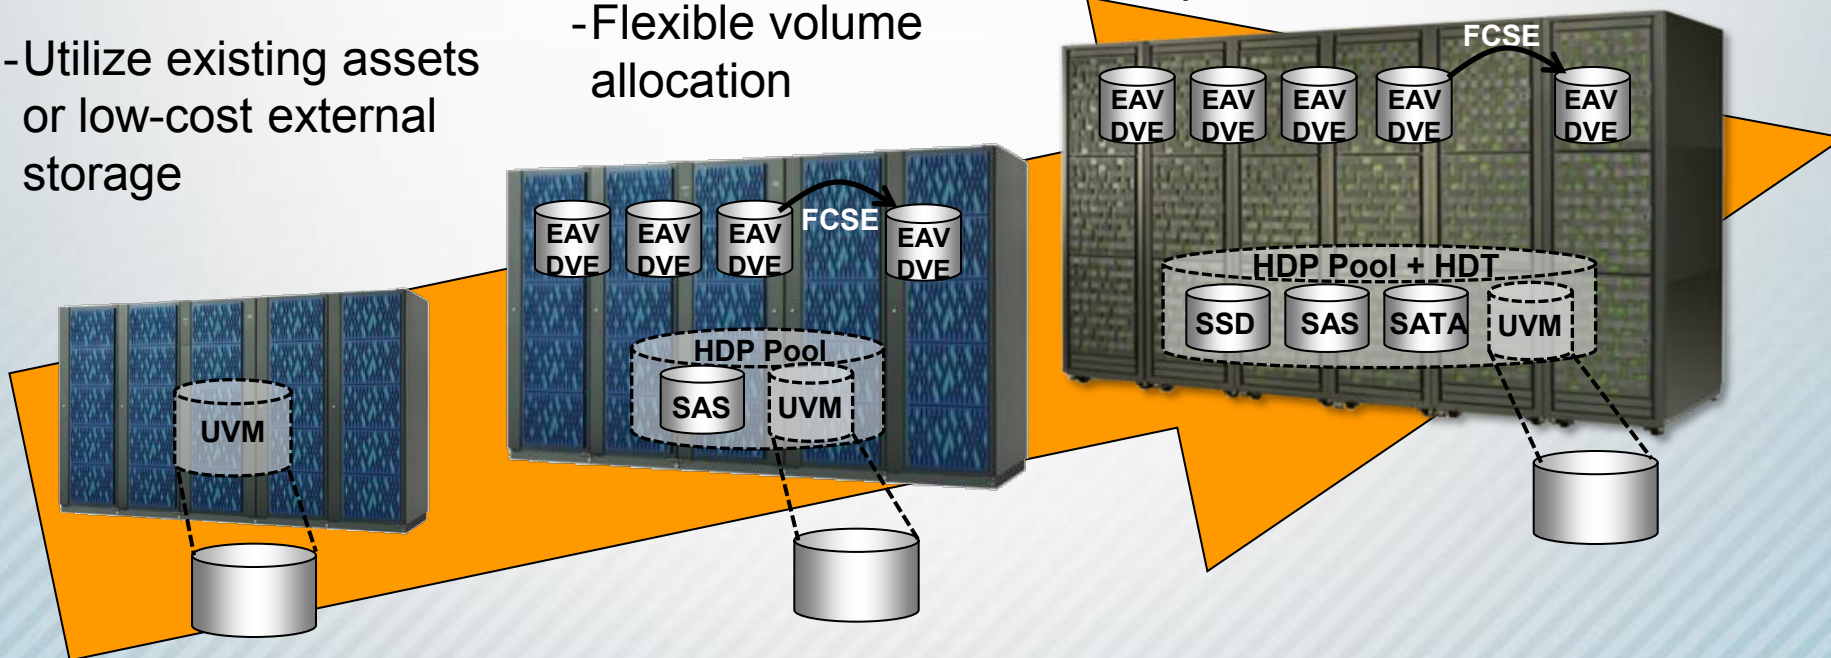
Benefits of Storage virtualization;
Eliminate tape and improve “nearline” storage

- Hitachi is bringing storage virtualization technology to the mainframe
- HDP for Mainframe provides FCSE, DVE and EAV compatibility
- HDP is technology foundation for next wave of mainframe virtualization

- Utilize existing assets or low-cost external storage

- Better space efficiency
- I/O load balancing
- Flexible volume allocation


- Automatic Hitachi Dynamic Tiering
- Effective use of high-speed SSD



EXTENDED ADDRESS VOLUME (EAV)



- z/OS® 1.10: IBM® introduced EAVs
 - Volumes
 - Larger than 65,520 cylinders and 54GB
 - Architectural limit raised to 268,434,453 cylinders
 - Current implementation supports 262,668 cylinders
 - Addresses the problem of z/OS 4-digit unit control blocks (UCBs)
 - Limit of 64K UCBs
 - Running out of addressable storage
 - 3390-A volumes
 - Maintain 3390 track format
 - HyperPAV facilitates multiple accesses to a volume
 - Addresses problem of queuing on a very large volume with a single UCB
 - Track-managed and cylinder-managed space



■ Track-managed space

- Area on an EAV located within the first 65,520 cylinders
- Intended as storage for “small” data sets
 - Extended Addressing Space (EAS)-ineligible data sets
- Comparable to the same space on non-EAVs
- Space is allocated in track or cylinder increments

■ Cylinder-managed space

- Area on an EAV located above the first 65,520 cylinders
- Intended as storage for “large” EAS-eligible data sets
- Space is allocated in multi-cylinder units of 21 cylinders.
 - System may round space requests up to a multiple of an MCU (i.e., 315 tracks)

DYNAMIC VOLUME EXPANSION (DVE)



DYNAMIC VOLUME EXPANSION (DVE) CONSIDERATIONS

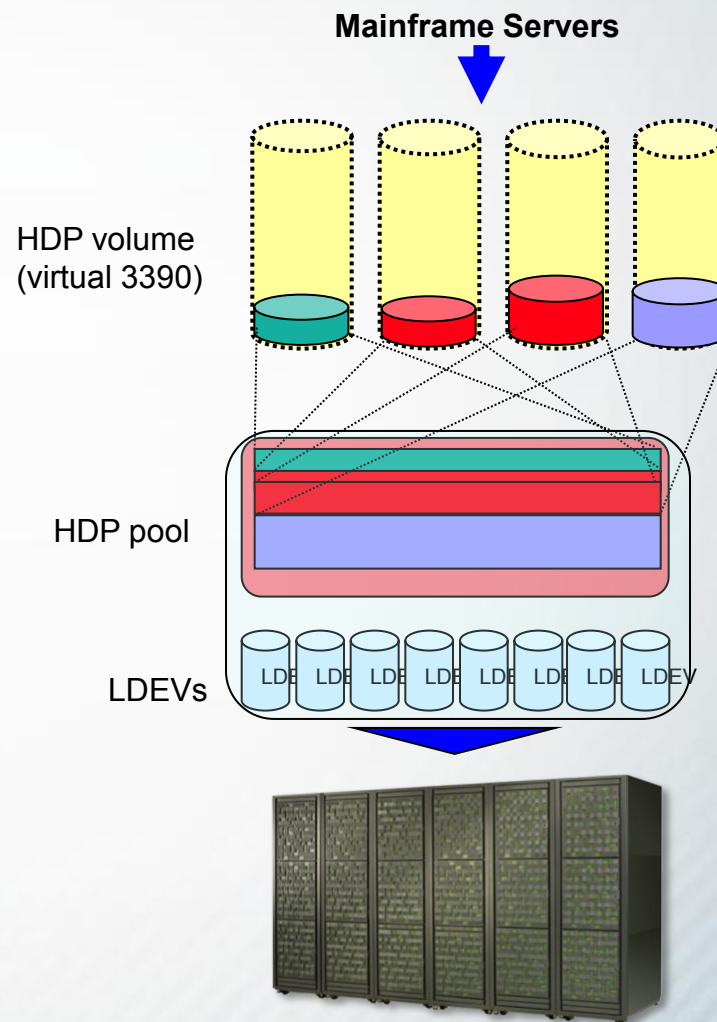
- Enables transparent capacity increase to meet application demands
 - Increases volume capacity for defined online volumes
 - Pre-z/OS® V1R10 up to 65,520 cylinders online
 - z/OS V1R10 online expands past 65,520 cylinders online (3390-A)
- Supports increased capacity with existing unit control blocks (UCBs)
 - Provides easier management
 - Reduces need to define new volumes requiring additional UCBs
 - Does not require updating IOCD or sysgen for new device addresses
- IBM® DFSMS device support code will detect the volume size change
 - Refreshes volume table of contents
 - Issues message IEA019I
 - Creates system management facility record type 22 with the old and new volume size
- DVE cannot be used with volumes involved in replication relationships*
 - Relationship must be broken before expanding volume

*IBM restriction on DVE volumes in replication relationships

HITACHI DYNAMIC PROVISIONING FOR MAINFRAME BASICS



- Optimize storage performance by spreading the I/O across all available disk drives
- Optimize storage capacity by only allocating capacity that is actually used
- Achieve FlashCopy® space efficiency for target volumes
- Leverage dynamic volume expansion to seamlessly grow physical and logical capacity



Simplifies capacity expansion: no IOGEN, no documentation updates, no changes to GDPS parameters, XRC sessions, FlashCopy configs, etc.

HITACHI DYNAMIC PROVISIONING FOR MAINFRAME STORAGE

Hitachi Storage Navigator Modular 2
Settings and reference of Hitachi Dynamic Provisioning are generally performed with Storage Navigator Modular 2.

STORAGE NAVIGATOR

SVP

V-VOL=DP-VOL (3390-A)

Volumes that are accessed by host and have a capacity recognized by host but do not have physical area

Can be created with a larger capacity than the installed physical Pool-VOL

V-VOL

HOST

PORT

PORT

PORT

3390-A
A

3390-A
B

3390-A
C

Mapping management information

Management area for managing mapping information of V-VOL, pool, and Pool-VOL and monitoring free capacity

SHARED MEMORY

Mapping management information backup area

POOL-VOL

A B

C

Pool-VOL (3390-V)


Volume in which data is stored

Necessary to have enough physical capacity for creating a Pool-VOL

Pool

Area created from 1 or more RAID Groups and in which Pool-VOLs are registered

RAID group



HITACHI DYNAMIC PROVISIONING FOR MAINFRAME STORAGE


DP POOL	PAGE SIZE	<ul style="list-style-type: none"> • 38MB 	<ul style="list-style-type: none"> • 672 3390 tracks
	CAPACITY	<ul style="list-style-type: none"> • Min: 8GB • Max: 696TB (internal) • Max: 928TB (external) 	<ul style="list-style-type: none"> • 1024 x 3390-V • 1024 x 3390-V
	VOLUME EMULATION TYPE	<ul style="list-style-type: none"> • 3390-V 	<ul style="list-style-type: none"> • Internal: 44 to 837,760 cyls • External: 44 to 1,117,760 cyls
	MINIMUM POOL SIZE	<ul style="list-style-type: none"> • 9,676 cylinders 	<ul style="list-style-type: none"> • 8GB
DP-VOLS	EMULATION TYPE	<ul style="list-style-type: none"> • 3390-A 	<ul style="list-style-type: none"> • EAV volumes
	CAPACITY	<ul style="list-style-type: none"> • Minimum size 870KB • Maximum size 218GB 	<ul style="list-style-type: none"> • 1 cylinder • 262,668 cylinders
	3390-A VOLS > 65,520 CYLINDERS	<ul style="list-style-type: none"> • Rounded up to multiple of 1,113 cylinders 	<ul style="list-style-type: none"> • 1,113 cyls = cylinder group of 25 mainframe HDP pages
	RECOMMENDED SIZE	<ul style="list-style-type: none"> • Multiple of 1,113 cylinders 	
	USER 3390-A DEFINITION	<ul style="list-style-type: none"> • 1-cylinder units 	<ul style="list-style-type: none"> • Not rounded up

SIMPLIFIED VOLUME DEFINITION, EXPANSION AND OPERATION

WHEN DYNAMIC PROVISIONING IS USED

- Virtual volume can be expanded using DVE without stopping host I/O
- Volume redefinition is unnecessary

Virtual volume




When expanding capacity, volume definition is unnecessary. (DP-VOL or pool capacity can be expanded without stopping host I/O.)

SIMPLIFIED DESIGN FOR PERFORMANCE LEVELING

WHEN DYNAMIC PROVISIONING IS NOT USED

- One RAID group has a heavy workload that may become a performance bottleneck
- Resolution: manually create a configuration that balances workloads on RAID groups

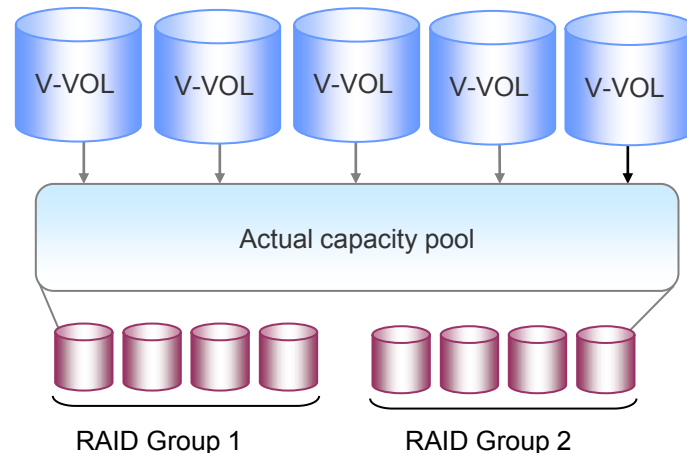


RAID Group 1 has heavy workloads that may become a performance bottleneck.

SIMPLIFIED DESIGN FOR PERFORMANCE LEVELING

WHEN DYNAMIC PROVISIONING IS USED

- Actual capacity of the pool is created from multiple RAID groups
- Volume is allocated from actual capacity of pool and is spread across multiple RAID groups
 - Distributes workloads across RAID groups
 - Reduces the possibility that a specific RAID group becomes a performance bottleneck




MAINFRAME STORAGE MANAGEMENT CONSIDERATIONS

- DFSMS enables fully automated and highly efficient storage management
 - Automatic class selection routines provide automation, which drives very high-capacity utilization
 - Naming conventions (data class = DC)
 - Backup and migration policies (management class = MC)
 - Pooling (storage class = SC and storage group = SG)
- Mainframe host storage group view and HDP pool view of free space are not aligned
 - Updates and deletes to catalog, VTOC, VTOC-IX are not reflected in the DP-Pool.
- Value of Hitachi Dynamic Provisioning for Mainframe

– Space Utilization Host SG View

- Foundation for compatibility with 80%, DVE and FCSE
- Prerequisite for Hitachi Dynamic Tiering for Mainframe



– Space Utilization DP-Pool View

- Foundation for compatibility with 80%, DVE and FCSE
- Prerequisite for Hitachi Dynamic Tiering for Mainframe



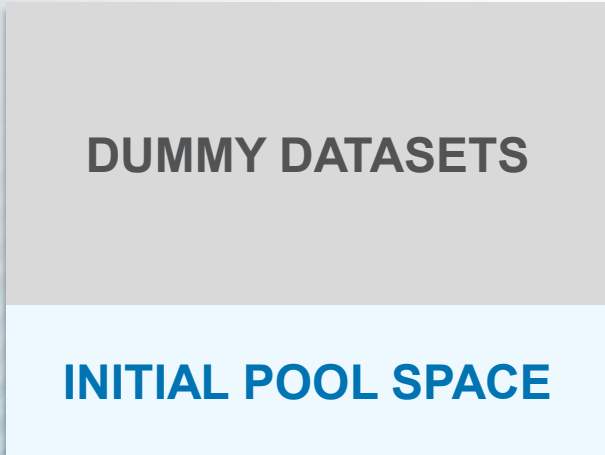
HITACHI DYNAMIC PROVISIONING FOR MAINFRAME STORAGE OVERPROVISIONING CONSIDERATIONS

- Overprovisioning: NOT recommended for mainframe environments
 - If used, customer must have a way to guarantee never to run out of space
 - Option 1
 - Use dynamic volume expansion (DVE): IBM limitations on DVE require an initial copy with replication products when volume size changes
 - Option 2
 - Allocate dummy datasets to cover the overprovisioned space that is physically not available in the pool

Total UCB capacity = 30TB
Initial pool space = 10TB
Dummy DSN space = 20TB


Pool space increase: Plus 10TB
Dummy space decrease: Minus 10TB
(Dummy must occupy same cylinders)

30TB PROVISIONED CAPACITY



VIRTUAL

PHYSICAL



DSN = data set name

SPACE RECLAMATION



- **Mainframe storage**
 - Uses complex count key data (CKD) structure and cylinder-head-record (CCHHR) addressing
 - CKD-formatted storage contains control information even when there is no data
 - Hitachi Virtual Storage Platform monitors existence of user records on track and maintains status in track metadata

- **Recovers space and restores it to the dynamic provisioning pool when it is no longer used, making it available for use by other volumes**
 - If no user records exist on track, track is eligible for reclamation
 - All tracks in DP page must not have any user records

- **Host initiated reclaim for mainframe storage**
 - Reclaims pages that only have tracks with no user records
 - Tracks metadata used to determine if user records exist on track
 - Done during CCW execution

- Mainframe storage with count key data (CKD) format and control data may not contain zeros even after datasets have been deleted
 - Must have techniques to erase records and update metadata to indicate that track contains no user records and is eligible for reclamation

- Techniques to make mainframe storage eligible to be reclaimed:
 - ICKDSF TRKFMT

 - IBM® Resource Access Control Facility with erase-on-scratch option enabled (overhead will elongate batch elapsed times)

 - Other CCW sequences that follow our host-initiated reclaim rules will also reclaim capacity with far less overhead

INTEROPERABILITY WITH OTHER PROGRAM PRODUCTS



DYNAMIC PROVISIONING FOR MAINFRAME EAV AND DVE INTEROPERABILITY

SUPPORTED IN V2+1 ¹

- Customized Volume Size ²
- Hitachi Compatible PAV for Mainframe
- IBM® Concurrent Copy
- IBM XRC
- Hitachi TrueCopy® ³
- Hitachi Universal Replicator ²
- IBM FlashCopy® v2
- Hitachi ShadowImage®
- Hitachi Volume Migration
- Hitachi Universal Volume Manager

SUPPORT PLANNED

- Dynamic cache residency
- Logical device security
- Hitachi Volume Retention Utility for IBM zOS®

NO SUPPORT PLANNED

- IBM FICON® data migration
- Hitachi Cross-OS File Exchange

Notes:

1. Any combination of 3390-A fixed and 3390-A DP-VOLs
- 2: To use 3390-A emulation with volumes smaller than 262,668 (max) cylinders, select desired size among 3390-3/9/L/M equivalent size. This does not require the CVS PP.
- 3: Three data centers will be supported with a later release of microcode (V3+1). Older storage models (Hitachi Universal Storage Platform® V, USP, etc.) will not be supported.

REPLICATION PRODUCTS: 3390 INTEROPERABILITY

Supported Copy Combinations	Hitachi TrueCopy®, Hitachi Universal Replicator P-VOL → S-VOL			Hitachi ShadowImage®, IBM® FlashCopy®, VM, XRC S-VOL → T-VOL			
	TrueCopy		HUR	SI	FC	VM	XRC
	VSP → VSP	USP, USP V → VSP					
3390-3/9/L/M ↓ 3390-A (Fixed)	Yes	Planned (V03)	No	No	Yes	No	Yes
3390-3/9/L/M ↓ 3390-A (DP-Vol)	Yes	Planned (V03)	No	No	Yes	Planned	Yes
3390-A (Fixed) ↓ 3390-3/9/L/M	No	No	No	No	Yes	No	Yes
3390-A (DP-Vol) ↓ 3390-3/9/L/M	No	No	No	No	Yes	Planned	Yes

- Three types of mainframe volumes
 - Traditional 3390: 1,3,9,27,54 volumes – USP, USP V and VSP
 - Fixed size EAV volumes: 3390-A – VSP with V02+1 microcode
 - DVE-capable EAV volumes: 3390-A DP-VOLs – VSP with V02+1 microcode

- HDPM supports interoperability between 3390-A volumes
 - Complete interoperability between 3390-A volumes with Hitachi ShadowImage®, Hitachi TrueCopy®, Hitachi Universal Replicator, etc.
 - Limited 3-data-center scalability – currently only 1x1x1 (MxNxO planned in V3+1)
 - Hitachi TrueCopy and Hitachi Compatible Mirroring for IBM® FlashCopy® supports interoperability between some combinations of all 3 emulation types.

- Migration options from traditional 3390-3/9/L/M to 3390-A DP volumes
 - Host-based migration using FDR/PAS or DFDSS, etc. (supported!)
 - Use TrueCopy or XRC to migrate from USP or USP V to VSP, or VSP to VSP

SUMMARY



- IBM compatibility advantage
 - Hardware interface compatibility minimizes risk: no host software required
 - Interoperability testing with IBM ensures highest quality
 - Intellectual property agreements enable continuous support

- Hitachi added-value advantage
 - Hitachi IP leadership in open systems storage becoming available to mainframe space
 - More customer choices to meet increasing business requirements
 - Automation moving into storage subsystems to improve productivity and quality of service

- Performance advantage
 - Technology innovation and hardware architecture improvements as strong foundation
 - Lower response time and higher throughput enhances storage utility value
 - Larger subsystems lower capital expenditure

QUESTIONS?

- **MAINFRAME SERIES**

- **Find the Facts and Lies about Replication in Mainframe Storage**, Oct. 26, 9 a.m. PT, 12 p.m. ET
- **Why Hitachi Virtual Storage Platform Does So Well in a Mainframe Environment**, Nov. 2, 9 a.m. PT, 12 p.m. ET

- **November and December**

- **Increase Your IT Agility and Cost-efficiency with HDS Cloud solutions**, Nov. 9, 9 a.m. PT, 12 p.m. ET
- **Best Practices for Upgrading to Hitachi Device Manager v7**, Nov. 16, 9 a.m. PT, 12 p.m. ET
- **Hitachi Clinical Repository**, Dec. 7, 9 a.m. PT, 12 p.m. ET

THANK YOU