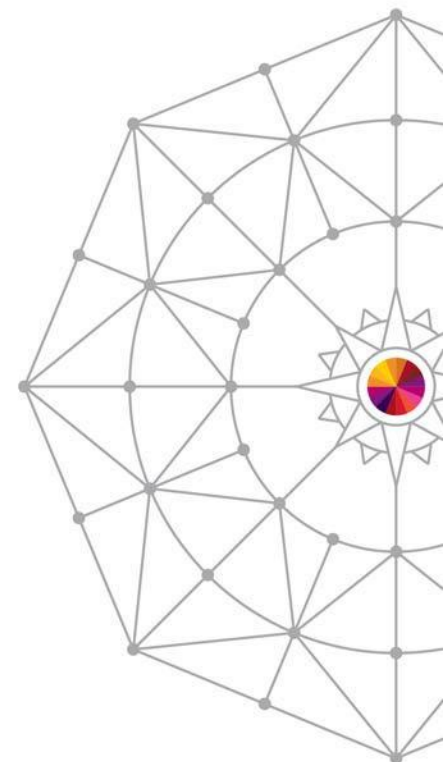# The Next Generation Storage Technology Today

William Smith

Enterprise Storage Product Manager

Hitachi Data Systems

August 6, 2014

Session # 15796

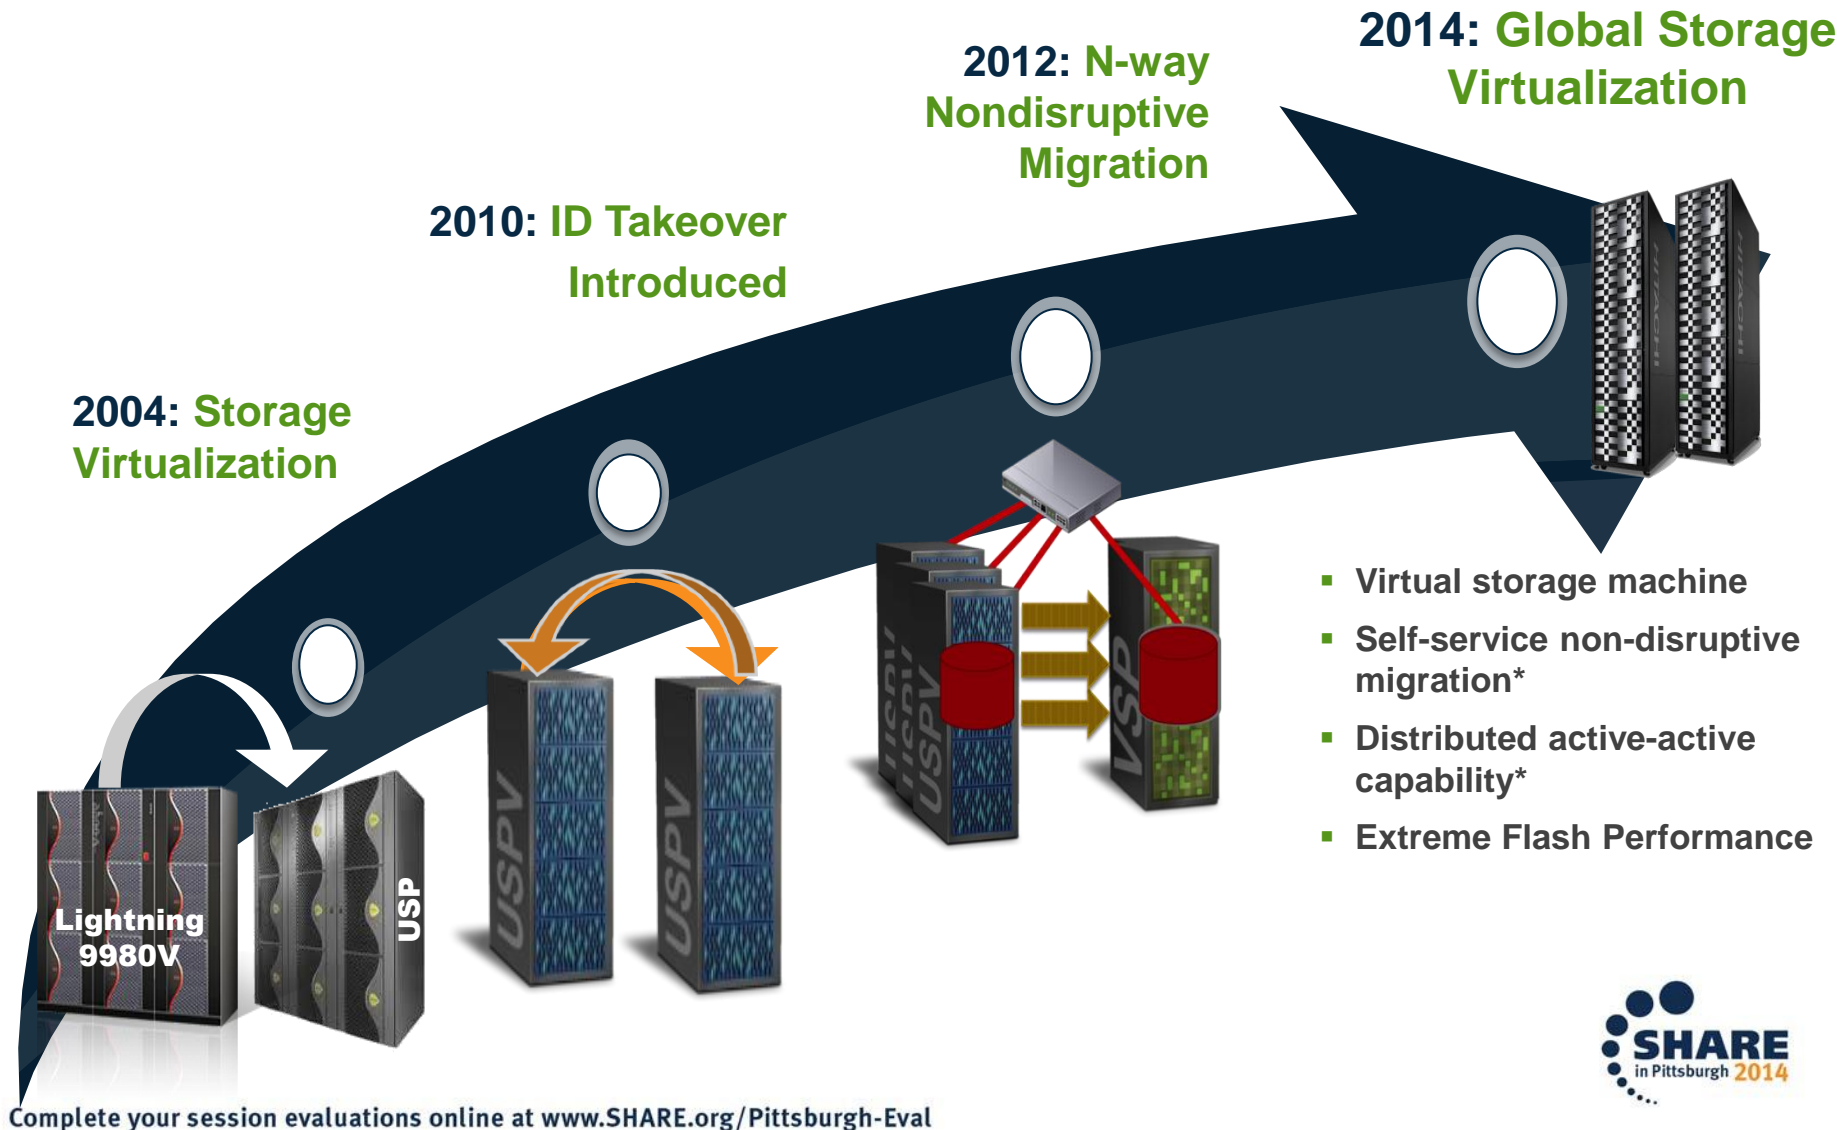SHARE is an independent volunteer-run information technology association
that provides **education**, professional **networking** and industry **influence**.

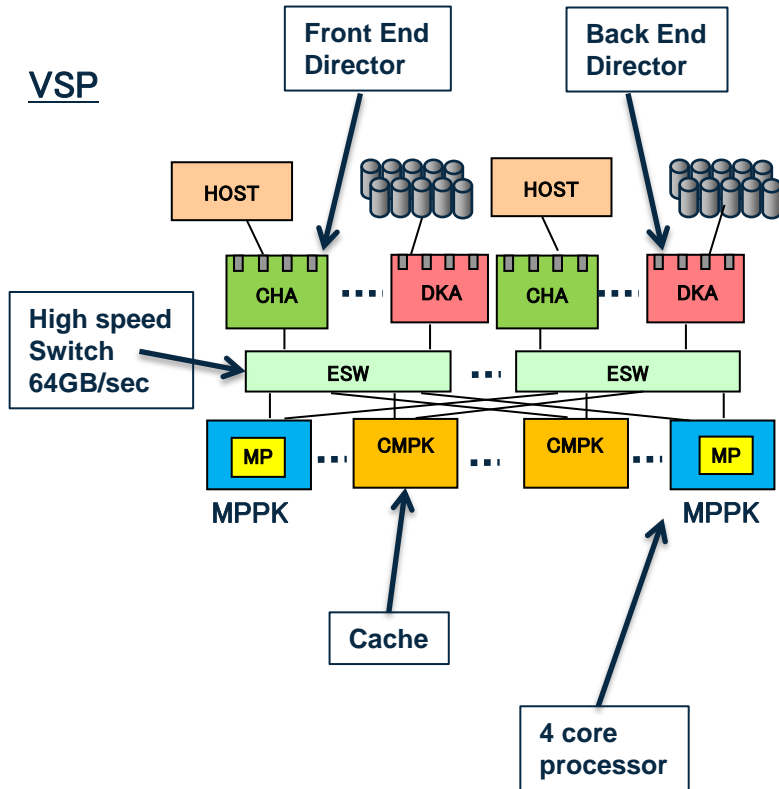SHARE in Pittsburgh 2014

Agenda

- **Hardware Architecture for VSP G1000**

- Features for Mainframe and OPEN

- Mainframe Specific

- OPEN Specific

- Summary

Continuous Innovation In Enterprise Storage

2014: Global Storage Virtualization

2012: N-way Nondisruptive Migration

2010: ID Takeover Introduced

2004: Storage Virtualization

Lightning 9980V

USP

USPV

USPV

USPV

VSP

- Virtual storage machine
- Self-service non-disruptive migration*
- Distributed active-active capability*
- Extreme Flash Performance

SHARE
in Pittsburgh 2014

Hardware Architecture Overview

VSP

Front End Director

Back End Director

HOST

CHA ···· DKA CHA ···· DKA

High speed Switch 64GB/sec

ESW ··· ESW

MP ··· CMPK ··· CMPK ··· MP

MPPK   MPPK

Cache

4 core processor

VSP G1000

Front End Director 2x paths 4x Faster Processors

Back End Director 2x paths 4x Faster Processors

HOST   HOST

MP
MPPK

CHA ···· DKA CHA ···· DKA

CMPK ··· CMPK

MP   MP

MPPK   MPPK

Above reflects that either a Front-end (CHA) or VSD pair (MP) may be installed in this location in the controller chassis

Cache and High speed Switch 420GB/sec

8 core processor

1 Side of Controller Shown for Simplicity

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

VSP G1000: Data Center Planning

Traditional System Layout

- **Increase floor-space efficiency**
- **Eliminate data center hotspots**

Today:
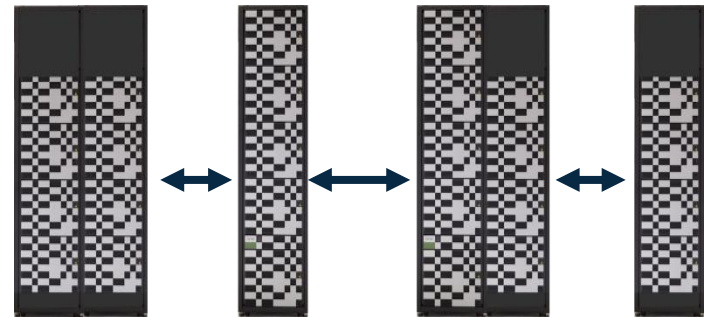**VSP G1000 Flexible Deployment**
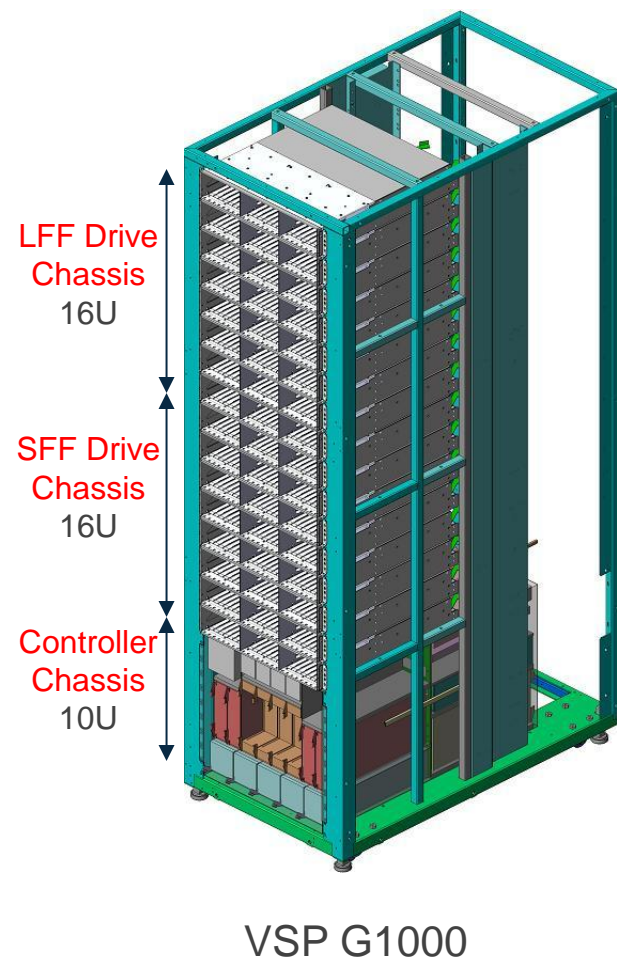
Separate Storage Controllers

5M, 30M, 100M Options

Coming Soon:
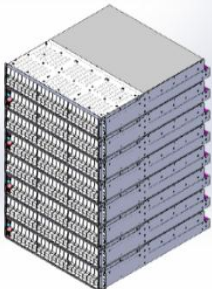**Ultimate Deployment Flexibility**
Separate Controller Racks and Disks Racks
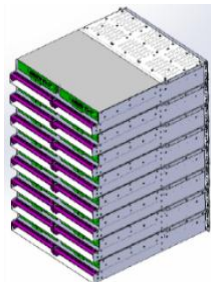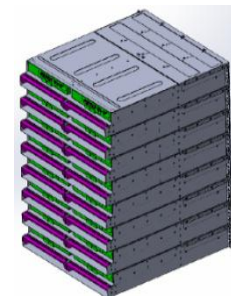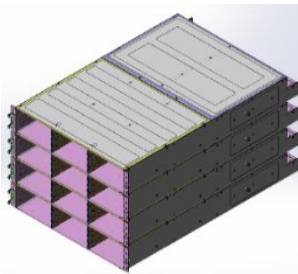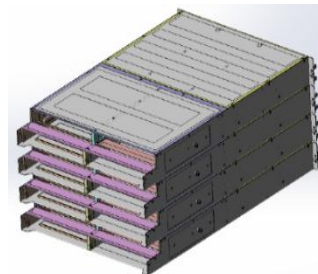
Drive Chassis Information

LFF Drive Chassis 16U

SFF Drive Chassis 16U

Controller Chassis 10U

VSP G1000

	Front view	Rear view	Details
SFF Drive Chassis DKC-F810-SBX			Maximum of <u>192 SFF drives</u> **Height: 16U**
LFF Drive Chassis DKC-F810I-UBX			Maximum of <u>96 LFF drives</u> **Height: 16U**
FMD Chassis DKC-F810I-FBX			Maximum of <u>48 FMDs</u> **Height: 8U**

VSP G1000: Media Options

Hitachi Flash Module Drive	SSD SFF (2.5-inch)	HDD SFF (2.5-inch)	HDD LFF (3.5-inch)
1.6TB	400GB	300GB 15K	4TB 7.2K
3.2TB	800GB	600GB 10K	
	800GB (3.5")	900GB 10K	
		1.2TB 10K	

Hitachi Accelerated Flash Controller ASIC

HIGHLY PARALLELIZED ARCHITECTURE

- 8 Lanes of PCIe 2.0
- PCIe Root Complex
- 4 Core 1.0GHz ARM9 CPU
- Integrated DDR-3 Interface
- Integrated Flash Controller logic
- Support for 32 paths to the Flash Array

60+ Related Patents Granted or Pending

PCIe 2.0

X8 PCIe PHYs

PCIe Root Complex

| **ARM9 CPU** | **ARM9 CPU** | DDR3 | To/From DDR |
| **ARM9 CPU** | **ARM9 CPU** | Misc. Interfaces | UART, UI2C, Ethernet |

X16 Flash Controllers

X32 Flash PHYs

Flash Memory interface
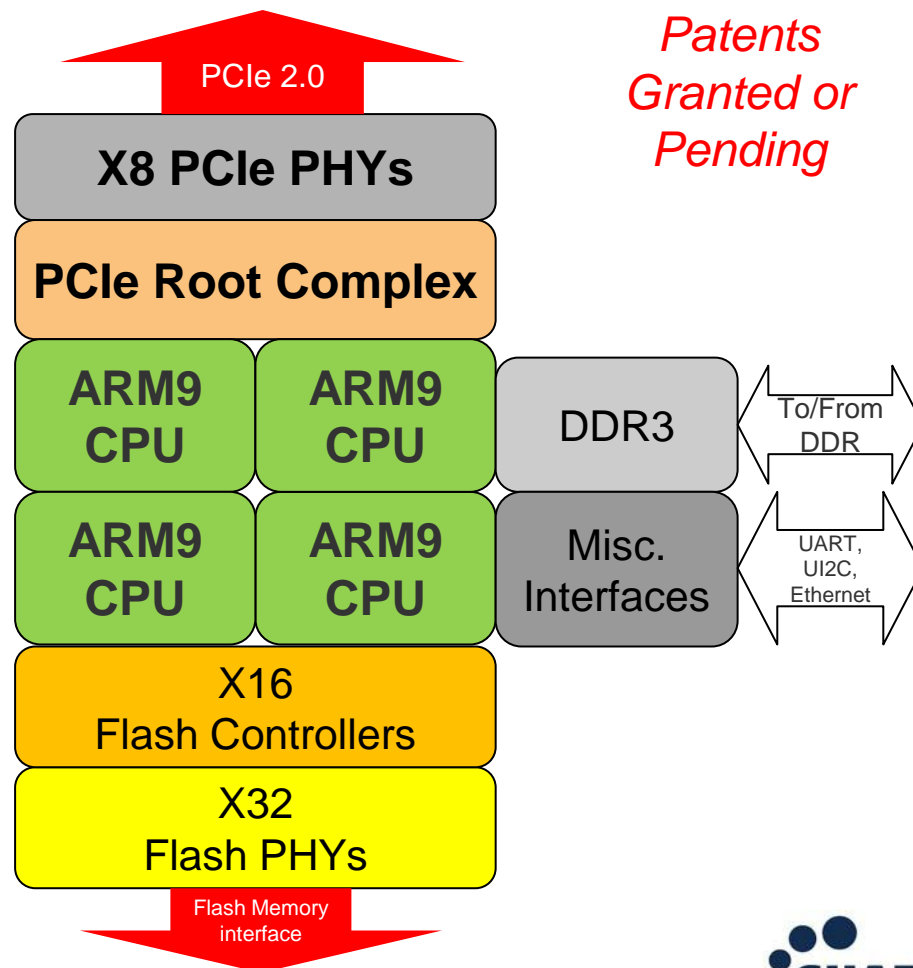
VSP G1000 to VSP Compared

	VSP G1000 1-Controller Chassis	VSP G1000 2-Controller Chassis	VSP Two-Controller Chassis
Maximum flash devices	96 Hitachi FMD 192 SSD	192 Hitachi FMD 384 SSD	192 Hitachi FMD 256 SSD
Maximum internal disks	1152 2.5-inch HDD 576 3.5-inch HDD	2304 2.5-inch HDD 1152 3.5-inch HDD	2048 2.5-inch HDD 1280 3.5-inch HDD
Maximum cache path bandwidth (GB/sec)	420	840	128
Maximum virtual storage director pairs	4 (64)	8 (128)	4 (32)
Maximum host ports	96 Fibre Channel 80 IBM® FICON® 80 FCoE*	192 Fibre Channel 176 FICON 176 FCoE	192 Fibre Channel 176 FICON 88 FCoE
Maximum cache	1TB	2TB	1TB
Fully configured power consumption (2.5" drives)	18.5 KVA	37 KVA	41.4 KVA
Maximum local copy pairs	32K	32K	16K
Maximum remote copy pairs	64K	64K	32K

Agenda

- Hardware Architecture for VSP G1000

- **Features for Mainframe and OPEN**

- Mainframe Specific

- OPEN Specific

- Summary

Hitachi Storage Virtualization Operating System
Virtualizes Externally Attached Storage

Automated Data Mobility with Dynamic Tiering

AUTOMATED

- Automated page-based data movement for performance and cost efficiency
- 38MB page for Mainframe (42MB for Open)
- Frees users from hands-on tier management and data layout

Large Financial Services Customer

Capacity Distribution

33%
64%
3%

Workload Distribution

19%
51%
30%

■ SAS ■ Flash ■ NL-SAS

Automated Data Relocation

Most Accessed

Less Accessed

Least or Never Accessed
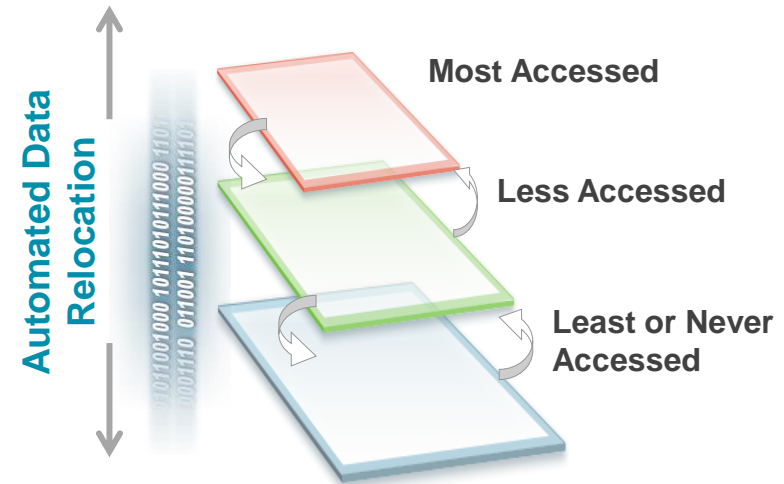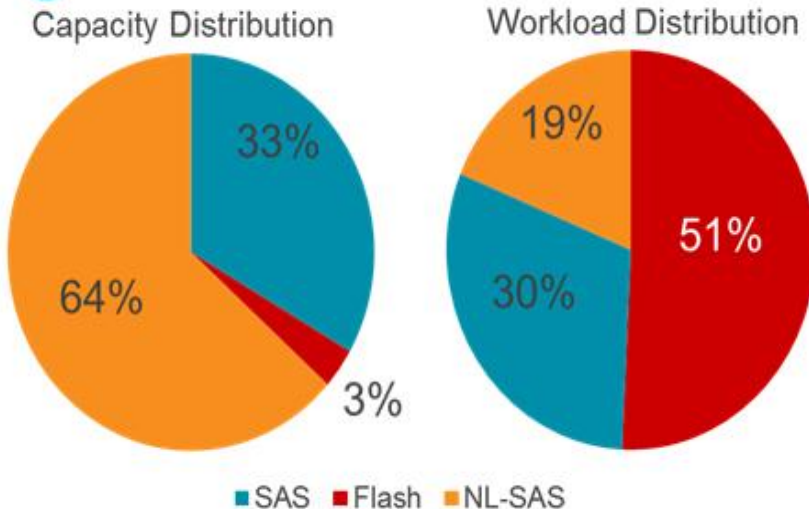
Dynamic Tiering Goals

- Reduce costs with self-managed and self-optimized storage tiers
- Most efficient use of flash ensures that investments are fully utilized

VSP G1000 Data-at-Rest Encryption

- Optional back-end director
- Third generation of controller based encryption
- Encryption is performed in hardware with no degradation in throughput
- Unique encryption key per piece of media. Associated key is deleted when media is removed
- Full KMIP support

Key Management Partners

- Key Management Interoperability Protocol
- Centralized Key Management
- OASIS Interoperability Protocol

Hitachi Virtual Storage Platform G1000
Performance Comparison: Backend 8kB [IOPS]

Reads
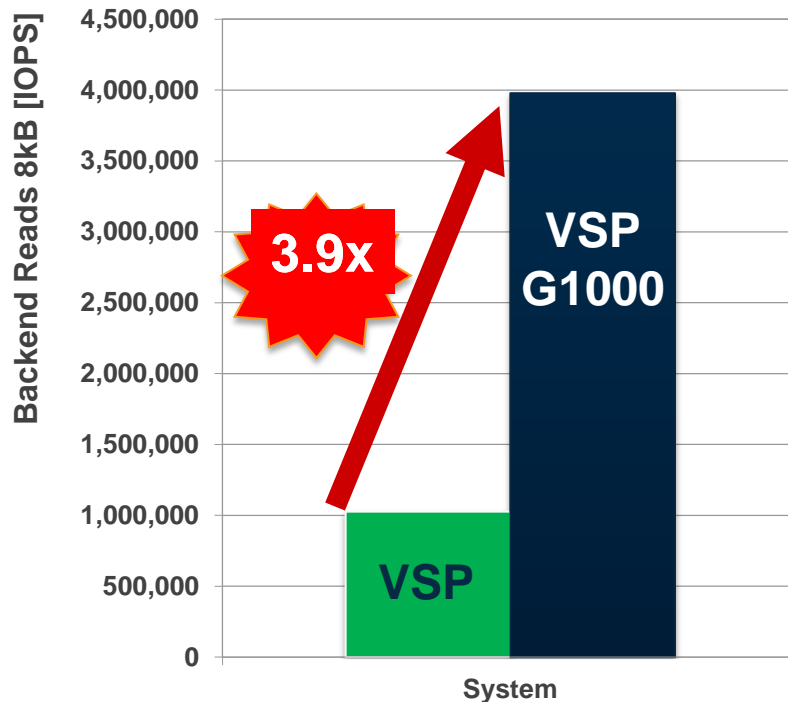
Writes

Preliminary internal testing

Hitachi Virtual Storage Platform G1000 Performance Comparison: Sequential 256kB [MB/sec]

Reads

Writes

3.7x

3.8x
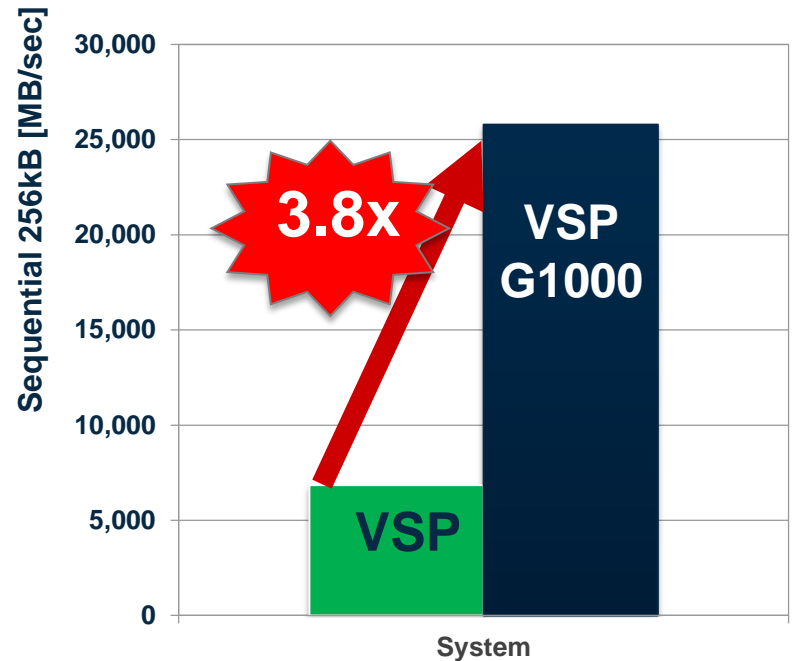
Preliminary internal testing

Agenda

- Hardware Architecture for VSP G1000

- Features for Mainframe and OPEN

- **Mainframe Specific**

- OPEN Specific

- Summary

Hitachi Mainframe Storage Strategy

ESSENTIAL COMPATIBILITY – LEVERAGING IBM COMPATIBILITY WITH HITACHI VALUE-ADD

Middleware and Application Level	IMS™, CICS®, DB2®, MQSeries®, etc.	**HITACHI STRATEGY** (Tivoli®-compatible, DB2 performance extensions)
BC/DR Solution Level	IBM portfolio - Global Mirror, PPRC - HyperSwap® Manager, GDPS, Basic HyperSwap, TPC-R　　HITACHI Portfolio - HUR, BCM, FDM - FC UVM, HDP, HDT, HTSM	**IBM Compatibility** **COMPLIMENTED BY HITACHI VALUE-ADDED SOLUTIONS**
Storage-Compatible Functions	Concurrent Copy, FlashCopy®, FlashCopy SE, Metro Mirror (PPRC), Global Copy (PPRC-XD)[1]	
Operating Systems Compatibility	z/OS® + DFSMSdfp, Linux for System z®, z/VM®, zVSE®, zTPF	
Hardware Interface Compatibility	FICON®, zHPF, FCP, MA, PAV, HyperPAV	

VSP G1000 Mainframe Features

FEATURE	DESCRIPTION
FLASHCOPY PERFORMANCE IMPROVEMENTS	FlashCopy and related TrueCopy and HUR performance improvements.
TPC-R AND BASIC HYPERSWAP	The ability to configure 2 and 3 data center environments using Tivoli Storage Productivity Center for Replication (TPC-R) and automate HyperSwap between sites and control units within the same center.
TPC-R SUPPORT FOR METRO MIRROR FAILOVER /FAILBACK	Enable Hitachi storage devices to participate in the automation of HyperSwaps between local sites (synchronous distances) and within the same data center.

SHARE
in Pittsburgh 2014

VSP G1000 Mainframe Features

FEATURE	DESCRIPTION
HTSM Support	Microcode release adds support for HTSM
Hierarchical Memory (VMA)	**Enables larger number of bitmaps as well as support larger volumes and LUNs by using a new hierarchical approach. Bitmap information will spread across local processor memory, (traditional share memory), and disk drives**
1TB (EAV) EXTENDED ADDRESS VOLUME	1TB EAV is a volume with more than 65,520 cylinders. *EAV increases the amount of addressable DASD storage per volume beyond 65,520 cylinders by changing how tracks on ECKD volumes are addressed*

VSP G1000 Mainframe Features

FEATURE	DESCRIPTION
Z/HPF BSAM QSAM	System z High-Performance FICON (zHPF) supports additional workloads using QSAM and BSAM access methods.
Z/HPF FORMAT WRITE	zHPF format writes − this function speeds DB2 loads, reorganizations, index rebuilds, and database restores. Additionally, DB2 load throughput with DB2 9 and 10 increases as much as 52 percent using 4K pages.
Z/HPF AND DB2 I/O	All DB2 I/O can convert to zHPF. When z/OS preformats DB2 data sets, zHPF enables a 15-to-2 reduction in the number of I/Os—an especially signification reduction when used with synchronous replication technologies such as peer-to-peer remote copy

VSP G1000 Mainframe Features

| FEATURE	DESCRIPTION	
Z/HPF DB2 LIST PREFETCH	With FICON Express8S, zHPF DB2 list prefetch reduces channel connect time by up to 2.5 times. DB2 10 uses list prefetch for disorganized index scans.	
Z/HPF BI-DIRECTIONAL CHANNEL PROGRAM	DB2 exploitation of zHPF will improve performance by allowing Media Manager to exploit bidirection zHPF I/O support	
Z/HPF LIST PREFETCH OPTIMIZER	zHPF List Prefetch Optimizer is used by DB2 technology to exploit the ability of the System z I/O architecture to read discontinuous disk segments in single I/O operations	

VSP G1000 Mainframe Features

FEATURE	DESCRIPTION	
GDPS – PPRCSUM	When compared to reporting suspensions on a per-devices basis, the Summary Event Notification for PPRC Suspends (PPRCSUM) dramatically reduces the message traffic and extraneous processing associated with PPRC suspension events and freeze processing	
GDPS HYPERSWAP Storage Control Health Message	This new attention message will be generated from the hardware, 1 per corresponding logical storage system, to alert the operating system of a condition that in the past would have surfaced as a general equipment check. This message will give more details and is intended to reduce the number of false HyperSwap events that have occurred with the less descriptive equipment check	

VSP G1000 Mainframe Features

FEATURE	DESCRIPTION	
GDPS/HYPERSWAP **QUALIFICATION TESTING WITH IBM**	Hitachi works with IBM to perform • Qualify GDPS/HyperSwap with PPRC with Global Mirror (XRC) with Hitachi compatible FlashCopy and FlashCopy Space Efficient, HDP, and HDT • Qualify GDPS/HyperSwap with 3 Data Center with Delta Resync and BCM along FlashCopy and FlashCopy Space Efficient, HDP, and HDT	

Hitachi Tiered Storage Manager for Mainframe – Z/OS HDT management

HTSM for Mainframe

Dynamic Tiering Volume

Data Center
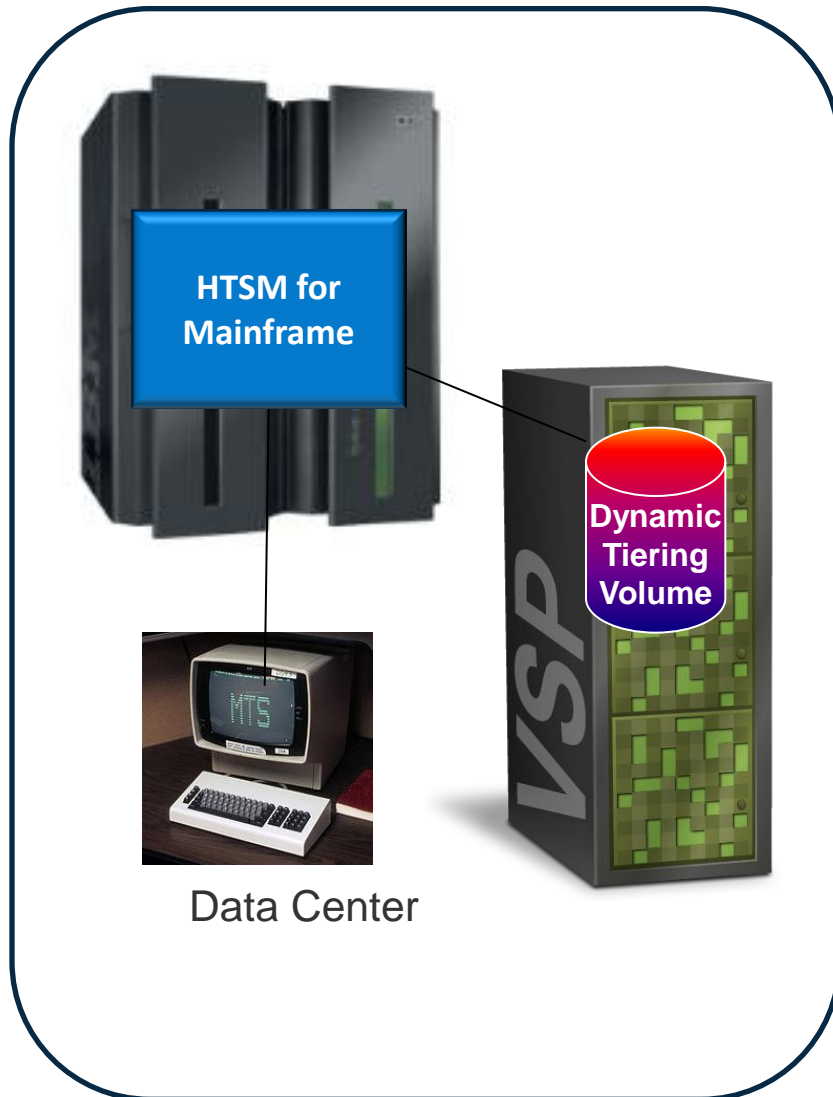
Host-based software that provides:

- Centralized and unified mainframe management of Hitachi Dynamic Tiering
  - Automation
  - Integration with DFSMS and storage groups
- Online storage service level controls
  - Increase application performance
  - Improves problem avoidance
- Single, consistent interface
  - Command based, script driven
  - ISPF interface

- Auto-discovery eliminates errors
  - Accelerates deployment
- Enables reporting and automatic notifications

ISPF ease of use with Point and Shoot

```
Command ===> _                                              Scroll ===> PAGE
                                                         2014/01/29 16:47:35
 Install   Defaults   Storage   Policy   TPG   Exit

 +=================== Installation Management 'Install'===================
 -=================== Set Defaults 'Defaults'============================
  Configuration file prefix . . : VAREND.HTSM80
  ISPF log max   . . . . . . . . : 0
  Pool usage threshold  . . . . : 80  %
  Capacity unit . . . . . . . . : Page
 -=================== Storage System List 'Storage'=====================
  <Scan>
  AC  SN      Status
  _     53004
 -=================== Policy CSV List 'Policy'==========================
  <Create>
  AC   PolicyID   Status
  _      DB2PROD
  _      DB2TEST
 -=================== Tiering Policy Group List 'TPG'====================
  COMMAND        TPGID          Status
  _____      DB2PROD
  _____      DB2TEST

All Rights Reserved. Copyright (c) 2013, 2014, Hitachi, Ltd.
Copyright (c) 2013-2014 Hitachi Data Systems Corporation. All rights reserved.

                                         Version 8.0.0-00
```

- **TPG_QUERY_STATISTICS** Sample Output

HTSM Mainframe Reporting


```
          Tiering Policy ID: PROD
                    Date: 10 Jul 2013
                    Time: 00:07:23

***************** Query TPG Tier Metrics *******************

=============================================================
TPG Total or         Tier1      Tier2      Tier3      Total
SN:PoolID or         Used       Used       Used       Used
*StorGrp* or         Pages      Pages      Pages      Pages
Volser    or       / UsedGB   / UsedGB   / UsedGB   / UsedGB
Volser Prefix      / Used%    / Used%    / Used%    / Used%
=============================================================
TPG Total              340          0          0        340
                     12.9GB       0GB        0GB      12.9GB
                      100%         0%         0%        100%
-------------------------------------------------------------
SN53004:81             340          0          0        340
                     12.9GB       0GB        0GB      12.9GB
                      100%         0%         0%        100%

  Used% of Pool       25.4%        0%         0%       8.67%

    Pool Pages        1340       1240       1340       3920
    Pool GB          50.9GB     47.1GB     50.9GB      149GB
  Tier% of Pool       34.2%      31.6%      34.2%       100%
-------------------------------------------------------------
*ALPHA*                340          0          0        340
                     12.9GB       0GB        0GB      12.9GB
                      100%         0%         0%        100%
-------------------------------------------------------------
GSE*                   340          0          0        340
                     12.9GB       0GB        0GB      12.9GB
                      100%         0%         0%        100%
-------------------------------------------------------------
10 Jul 2013 00:07:23 *** Action TPG_QUERY_TIERS Successful
```

Agenda

- Hardware Architecture for VSP G1000

- Features for Mainframe and OPEN
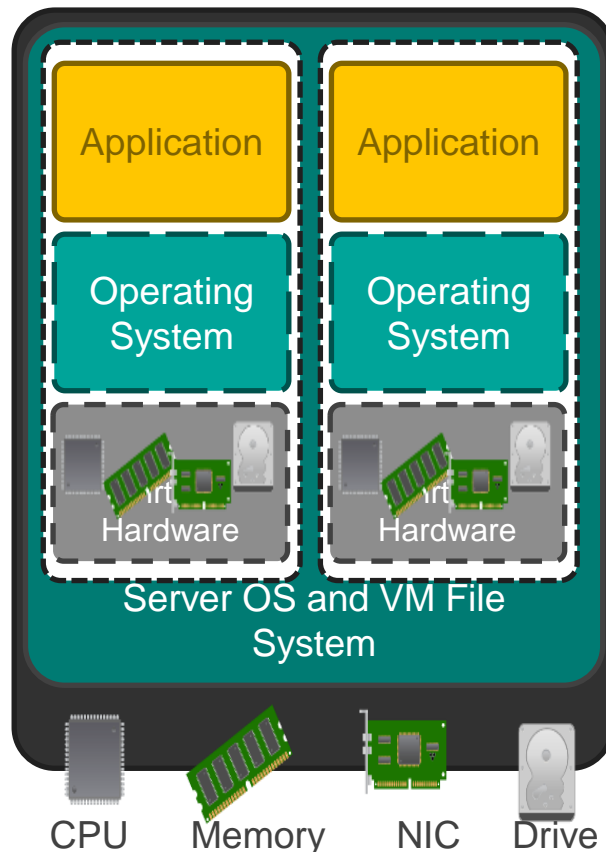
- Mainframe Specific

- **OPEN Specific**

- Summary

Hitachi Storage Virtualization Operating System:
Introducing Global Storage Virtualization

Virtual Server Machines FOREVER CHANGED the way we see DATACENTERS

Hitachi VIRTUAL STORAGE MACHINES will do the SAME

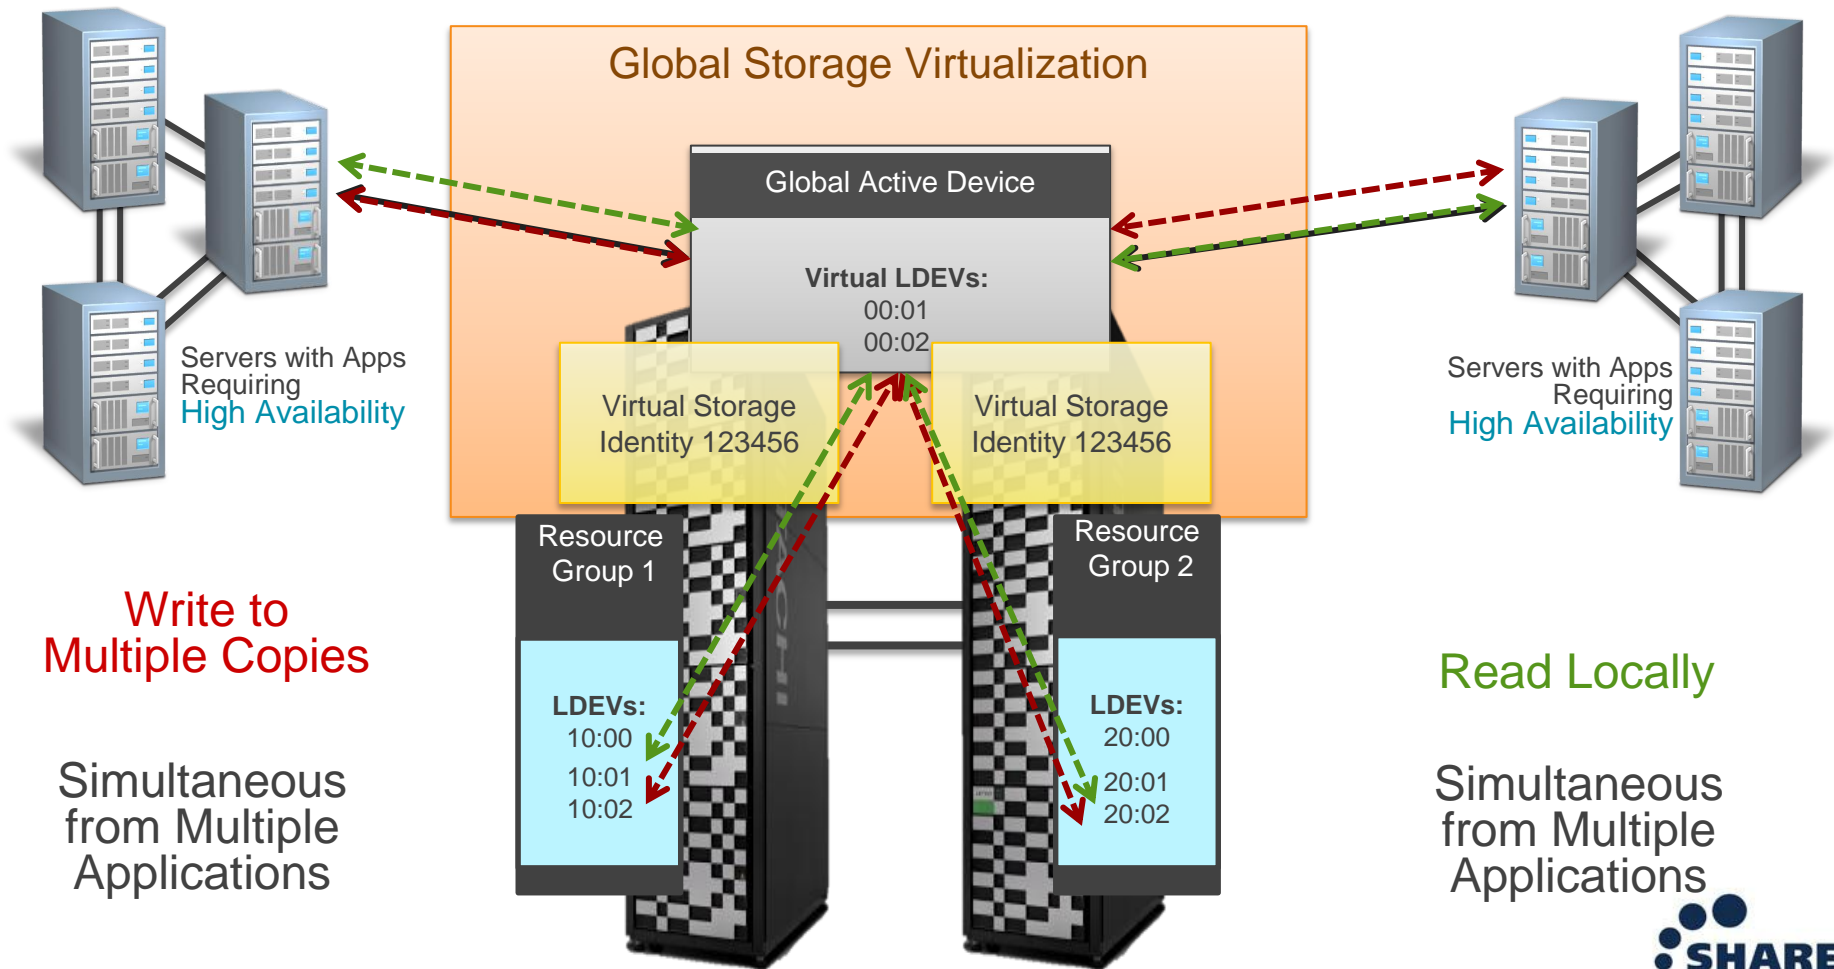| Application	Application
Operating System	Operating System
Virtual Hardware	Virtual Hardware

Server OS and VM File System

CPU — Memory — NIC — Drive

Virtual Storage Software

Global Virtualization

Virtual Storage Director — Cache — Front-End Ports — Media

Hitachi Global Storage Virtualization
Clustered Active-Active Systems

Global Storage Virtualization

Global Active Device

Virtual LDEVs:
00:01
00:02

Servers with Apps Requiring
High Availability

Virtual Storage Identity 123456

Virtual Storage Identity 123456

Servers with Apps Requiring
High Availability

Resource Group 1

Resource Group 2

Write to Multiple Copies

Read Locally

LDEVs:
10:00
10:01
10:02

LDEVs:
20:00
20:01
20:02

Simultaneous from Multiple Applications

Simultaneous from Multiple Applications

Global-Active Device Supported Configurations
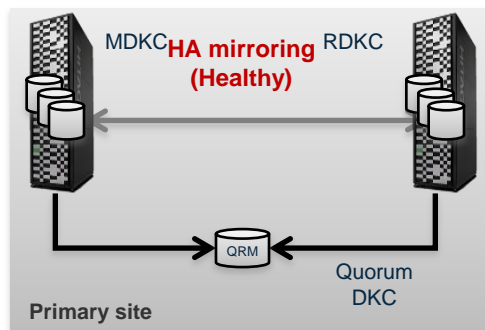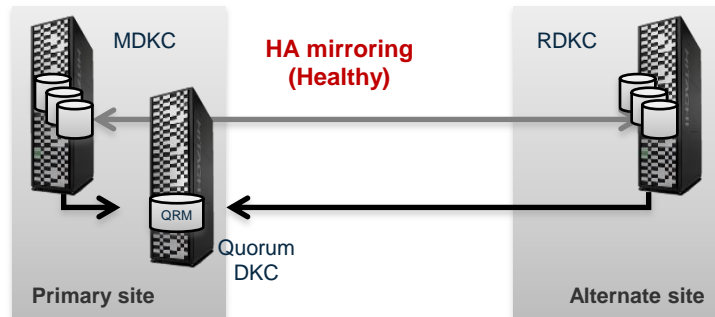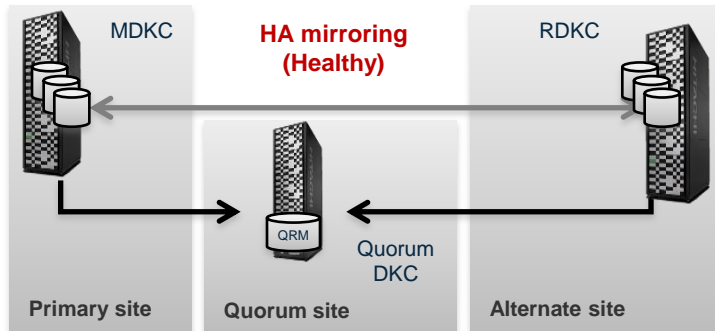
3 Sites

- Each DKC is located on a separate site
- Provides maximum level of business continuity for any type of DKC failures, or site failures (Primary site, Alternate site, Quorum site)
- Quorum disk can reside on a HUR target DKC connected via UVM and FCIP

2 Sites

- Reserve and quorum DKC is located on the primary site
- Provides moderate level of business continuity for any type of DKC failures, or alternate site failure
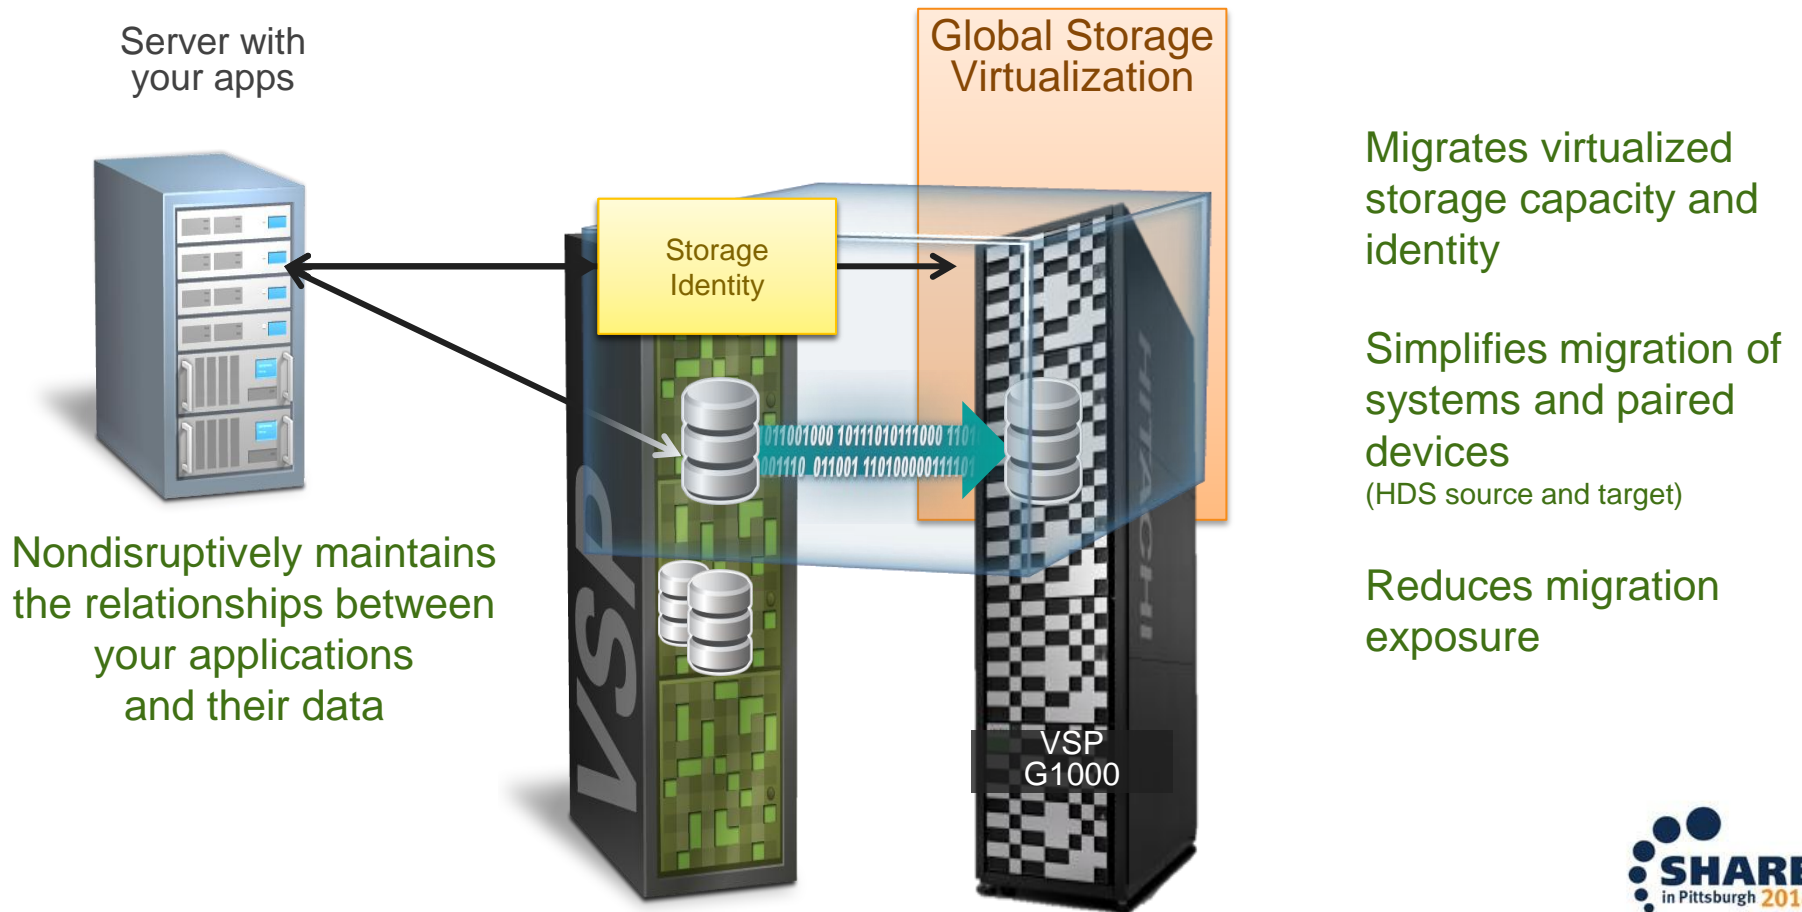
Single data center

- All the DKCs are located on the same site
- Provides business continuity for DKC failures, but cannot maintain the business for a site failure

Nondisruptive Migration
Through Global Storage Virtualization

AVAILABLE

Move Data and Refresh Systems as Needed

Server with
your apps

Global Storage
Virtualization

Storage
Identity

Migrates virtualized
storage capacity and
identity

Simplifies migration of
systems and paired
devices
(HDS source and target)

Reduces migration
exposure

Nondisruptively maintains
the relationships between
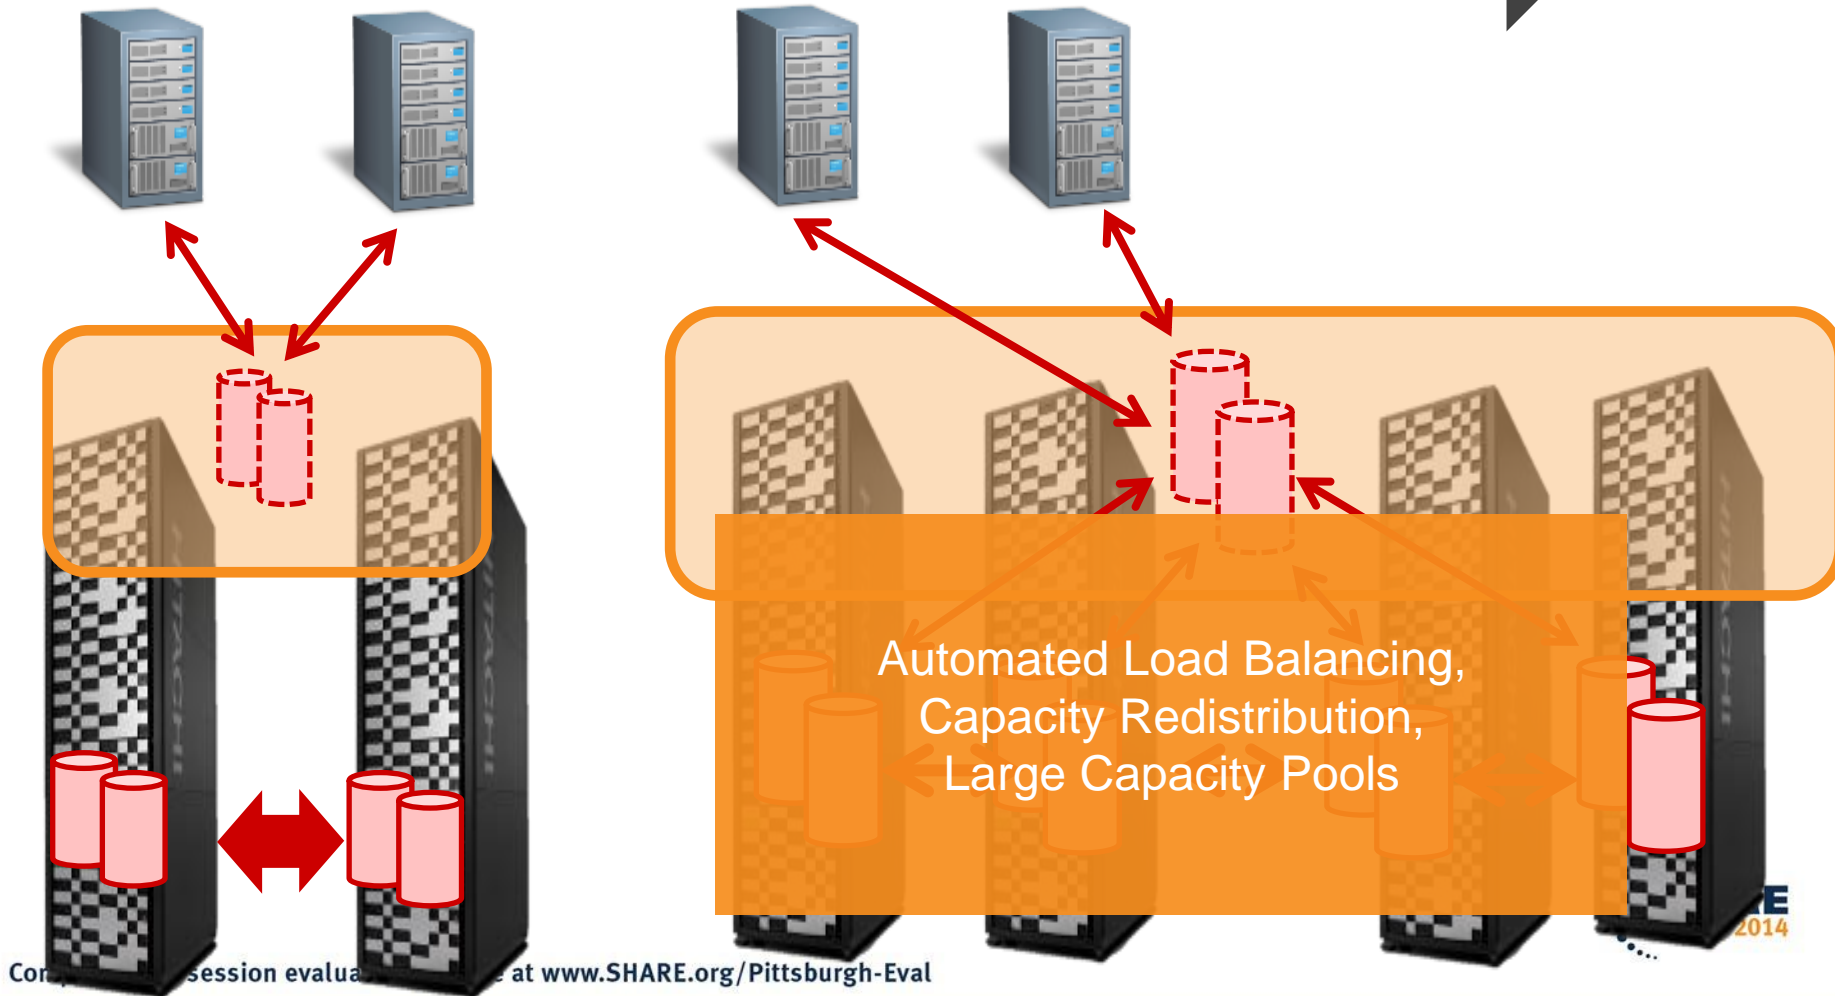your applications
and their data

VSP
G1000

Agenda

- Hardware Architecture for VSP G1000

- Features for Mainframe and OPEN

- Mainframe Specific

- OPEN Specific

- **Summary**

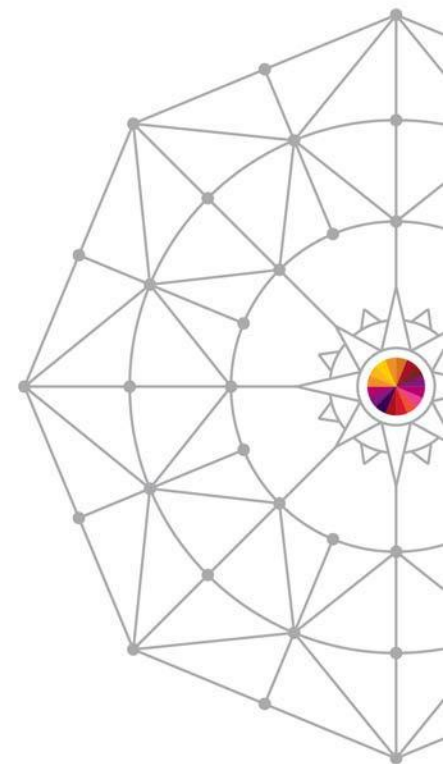# Where Might We Go Together?

2014

Potential Future Capabilities

Automated Load Balancing,
Capacity Redistribution,
Large Capacity Pools

Thank You