

LEARN THE FACTS ABOUT REPLICATION IN MAINFRAME STORAGE

JOE AMATO
PRINCIPAL SOLUTIONS ARCHITECT
HITACHI DATA SYSTEMS

OCTOBER 26, 2011



Learn the Facts about Replication in Mainframe Storage

Business continuity is essential for today's enterprise computing environments, and protecting your data and information is key. However, the many myths associated with data replication can be confusing. How do you sort truth from fiction? Join Hitachi solution architect Joe Amato to learn about in-system replication as well as replication to remote locations, both synchronously and asynchronously. You'll come away equipped with valuable insight into the business continuity solutions available for mainframe storage.

Attend this WebTech to understand:

- The differences between 2-, 3- and 4-data-center solutions and the value of each
- How to avoid running out of unit control blocks and also address the problem of very large data sets required by application support with extended address volumes
- How to use your remote synchronous copy of data while automatically keeping the data at a secondary remote replication site in sync with the failover site

■ Mainframe Series

- See on-demand webcast: **How to Apply the Latest Advances in Mainframe Storage**, www.hds.com/webtech
- **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

- Introduction
- Hitachi Data Systems mainframe replication portfolio
- Hitachi Universal Replicator highlights
- Advanced replication configurations
- Replication management
- Summary

BUSINESS CONTINUITY AND DATA PROTECTION: ALIGN BUSINESS AND IT AVAILABILITY OBJECTIVES

Business and IT objectives

Cost



Operational resilience

100% availability and no data loss

Operational recovery

1 to few service failures

Local site(s) up

Disaster recovery

Site-wide (all) service failures

Site(s) down

Regulatory compliance

Retention
Timely search and retrieval

Risk



Recovery objectives

Tier-1

RTO

RPO

Zero

Zero

Mins-to-
hrs

Secs-to-
mins

Hrs-to-
days

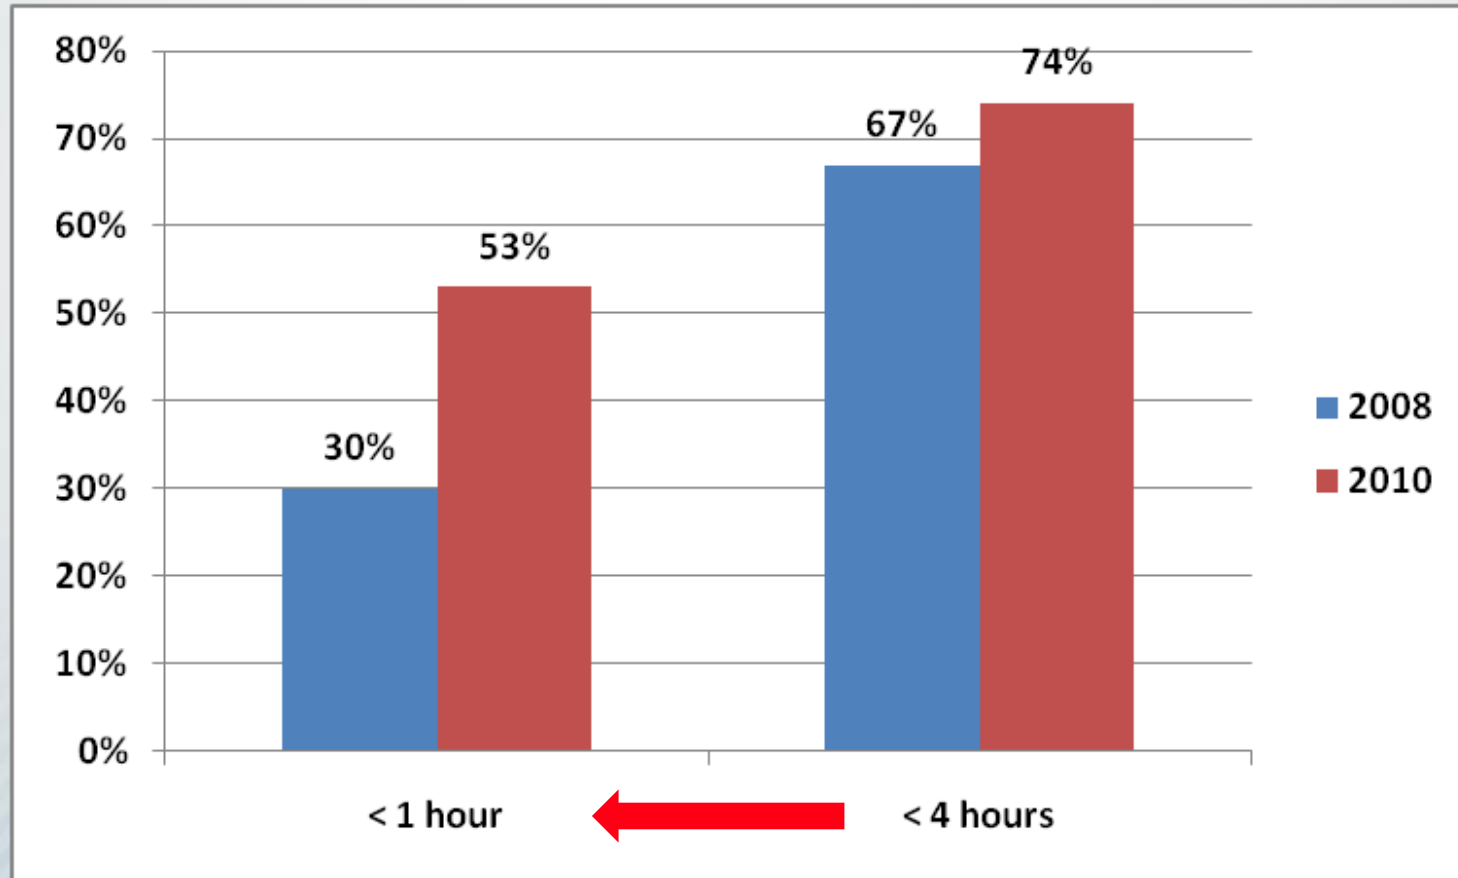
Hrs-to-
days

Tier-n



SHRINKING RTO FOR TIER-1 APPLICATIONS

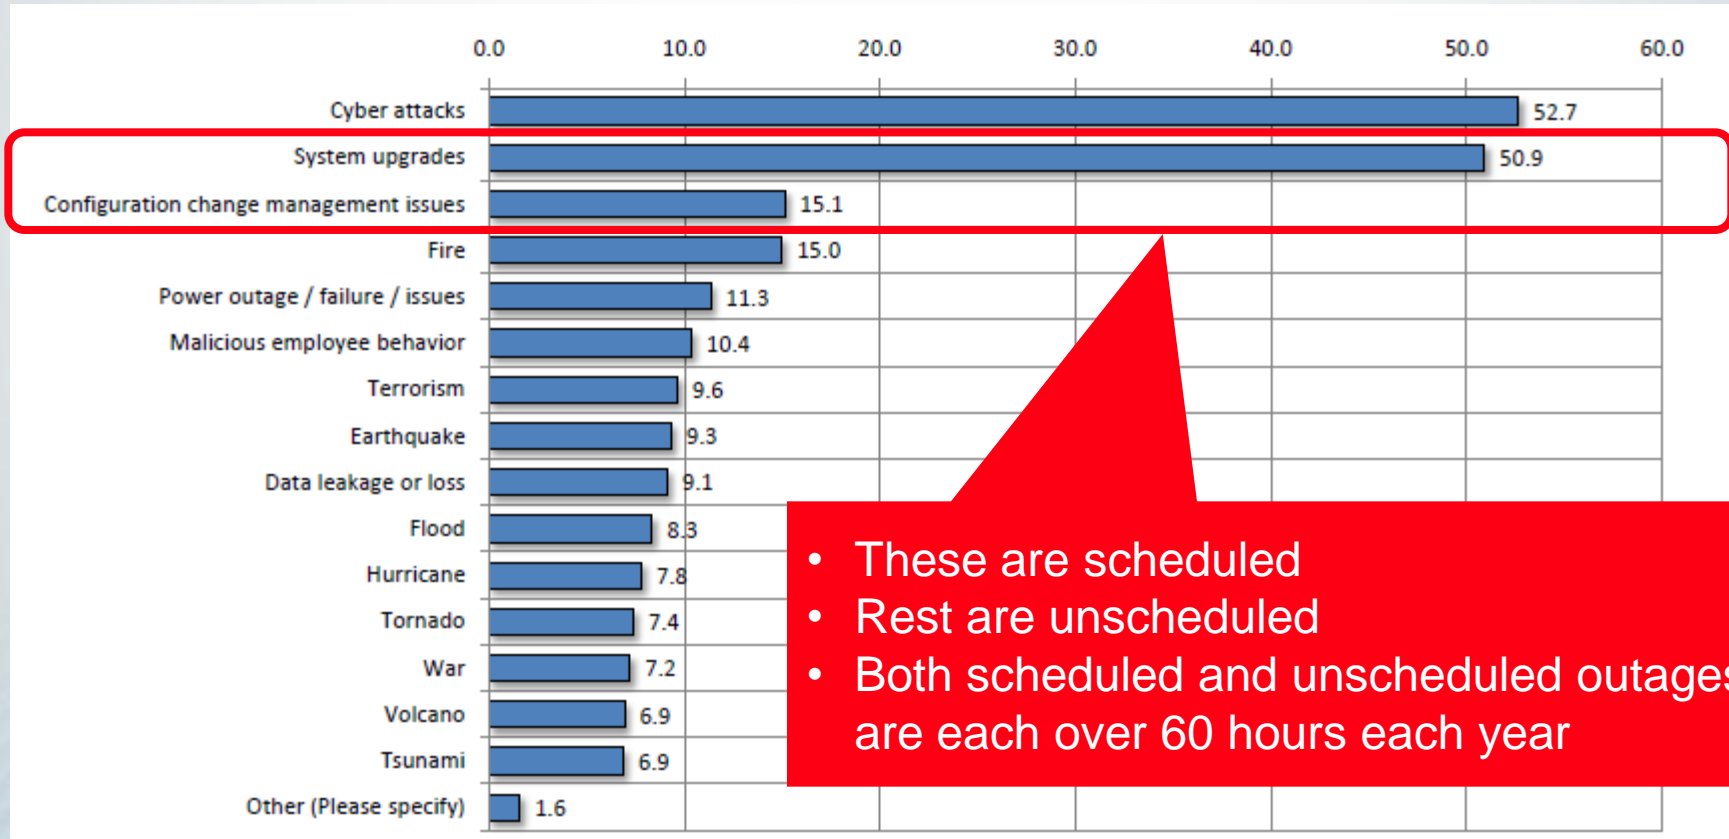
Percent of enterprise companies (> 1000 employees) with minimum RTO for most-critical applications



Source: Enterprise Strategy Group, 2008 and 2010

MAJOR CAUSES OF DOWNTIME

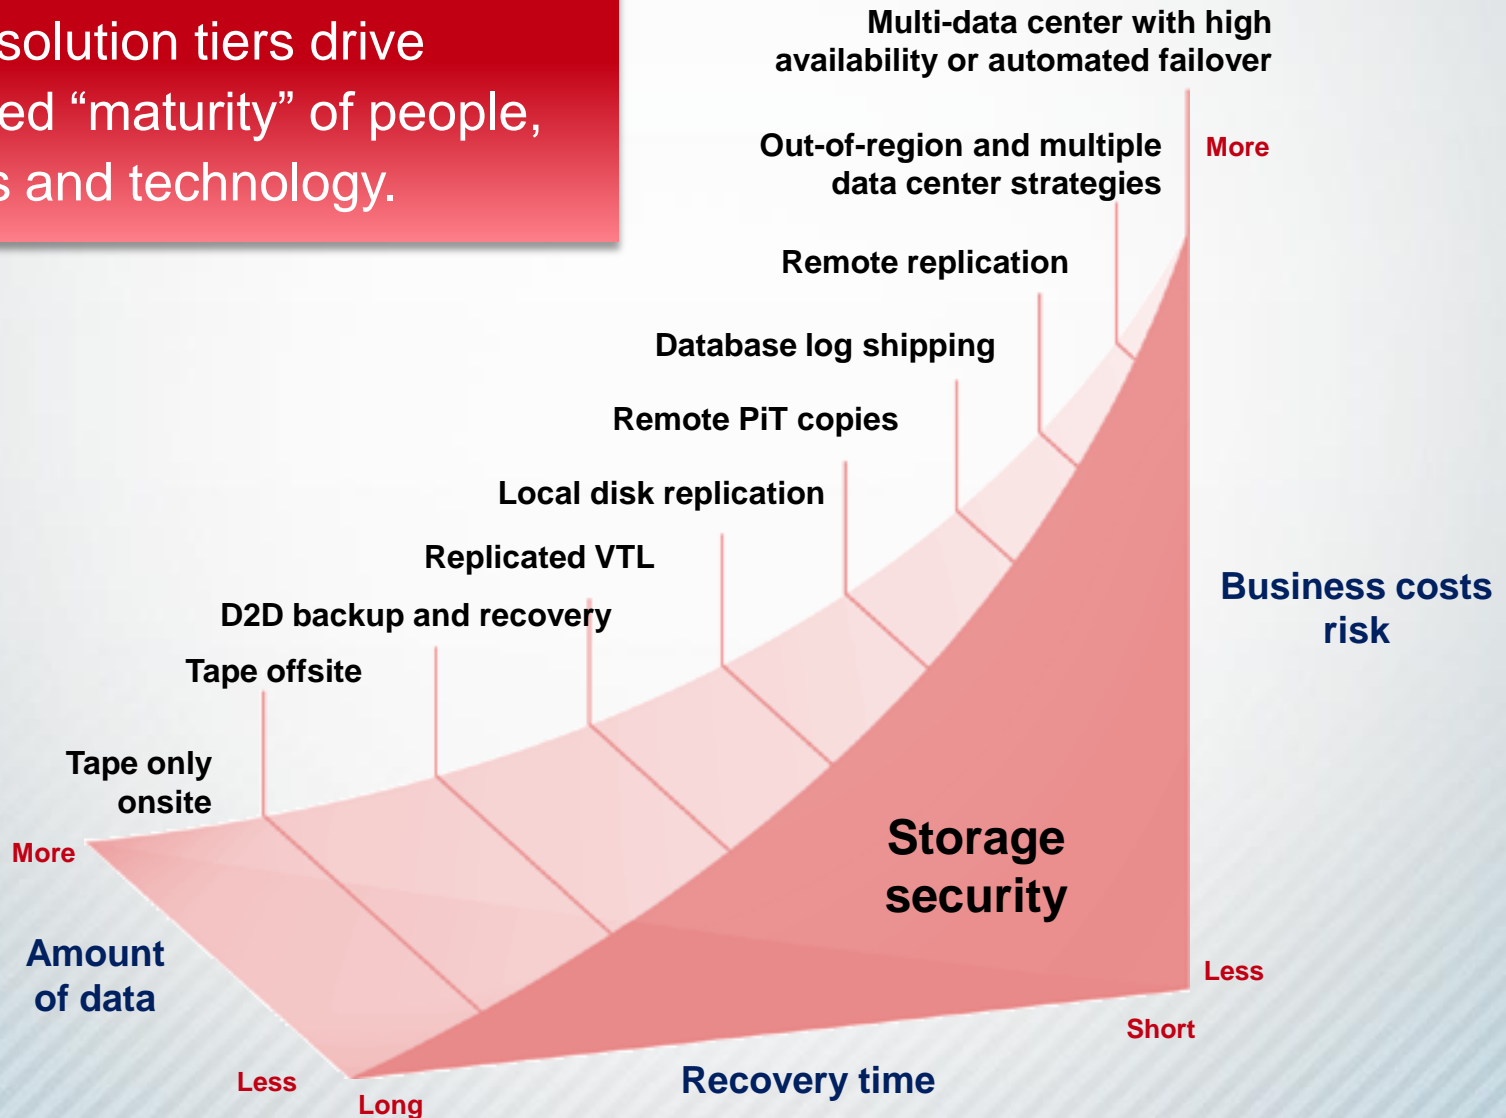
Average hours of downtime by cause over last 12 months



Source: Symantec 2010 Disaster Recovery Study – Global Results, November 2010
(n=1700, employees >5000, countries >18 worldwide)

SOLUTION TIERS TIED TO BUSINESS RISK

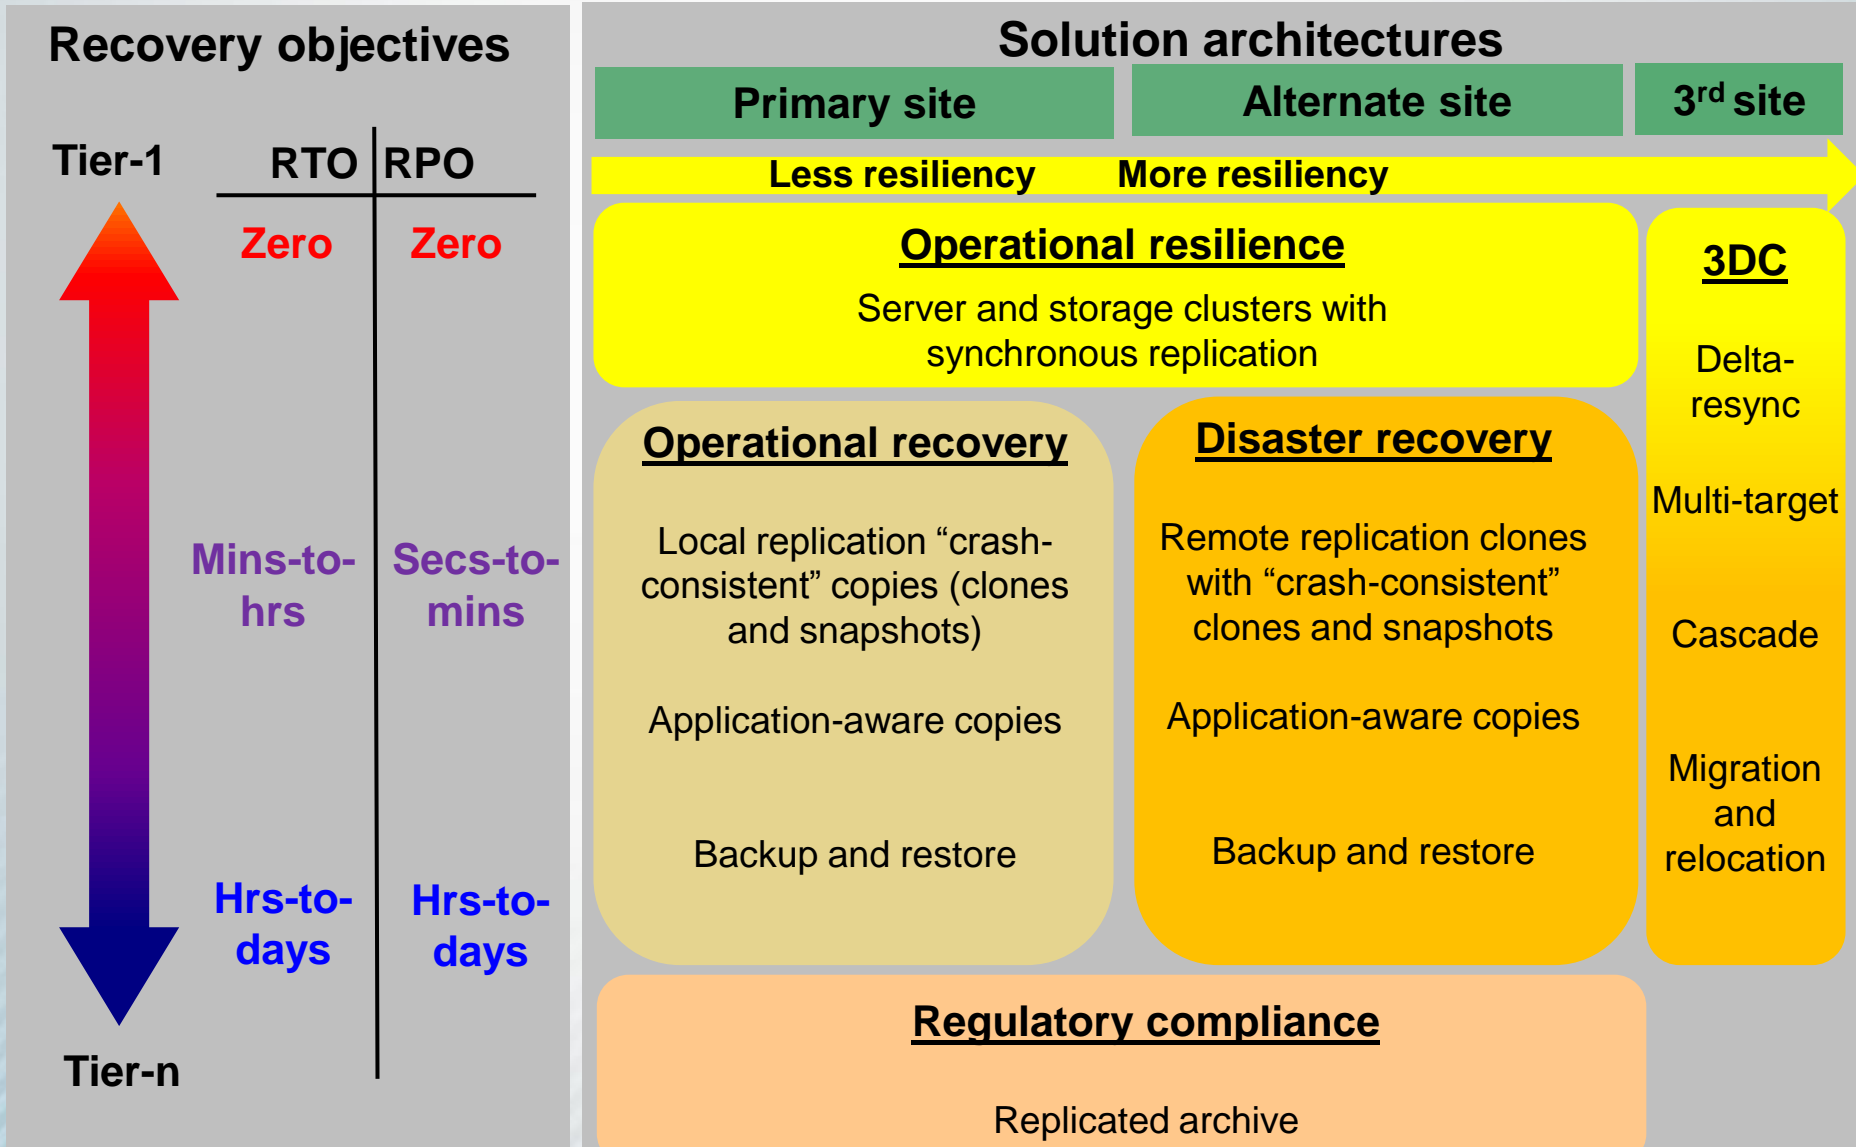
Higher solution tiers drive increased “maturity” of people, process and technology.



- Tighter RTOs are driving the need to *avoid downtime* and data loss via operational resilience solutions
- Local and remote replication are moving to replace backup and restore as *first-choice* for *most* operational and disaster recovery needs
- Shift continues from tape as backup target to disk as backup target
- Higher solution tiers drive higher organizational maturity levels in people and processes as well as in technology

- Introduction
- Hitachi Data Systems mainframe replication portfolio
- Hitachi Universal Replicator highlights
- Advanced replication configurations
- Replication management
- Summary

BUSINESS CONTINUITY AND DATA PROTECTION: ALIGNING RECOVERY OBJECTIVES AND ARCHITECTURES



In-system replication solutions	Remote replication solutions	Replication management solutions
<p>Hitachi ShadowImage® For full-volume clones of business data with consistency (ATTIME split with Hitachi Universal Replicator)</p>	<p>Hitachi TrueCopy® Synchronous, consistent clones at remote location up to 300km (~180 miles)</p>	<p>Hitachi Business Continuity Manager Replication management in z/OS® environments</p>
<p>Compatible with IBM® FlashCopy® Point-in-time volumes and datasets via IBM command set in z/OS environments</p>	<p>Universal Replicator Any distance, unique use of pull technology and journal volumes to accommodate link outages or interruptions</p>	<p>Hitachi Replication Manager with Mainframe Business Continuity Manager plus enterprise-wide GUI replication management and monitoring for mainframe and open systems</p>
	<p>Compatible with IBM XRC Asynchronous remote replication via compatibility with XRC in z/OS environment</p>	<p>IBM Basic HyperSwap® and GDPS HyperSwap certification and/or integration High-availability storage for z/OS® environments</p>

Replication assessment, migration and implementation services

Best-practice designs that ensure results and lower risks in implementation

Replication scorecard (health check) services

Periodic checkups that ensure continued solution value

- Introduction
- Hitachi Data Systems mainframe replication portfolio
- **Hitachi Universal Replicator highlights**
- Advanced replication configurations
- Replication management
- Summary

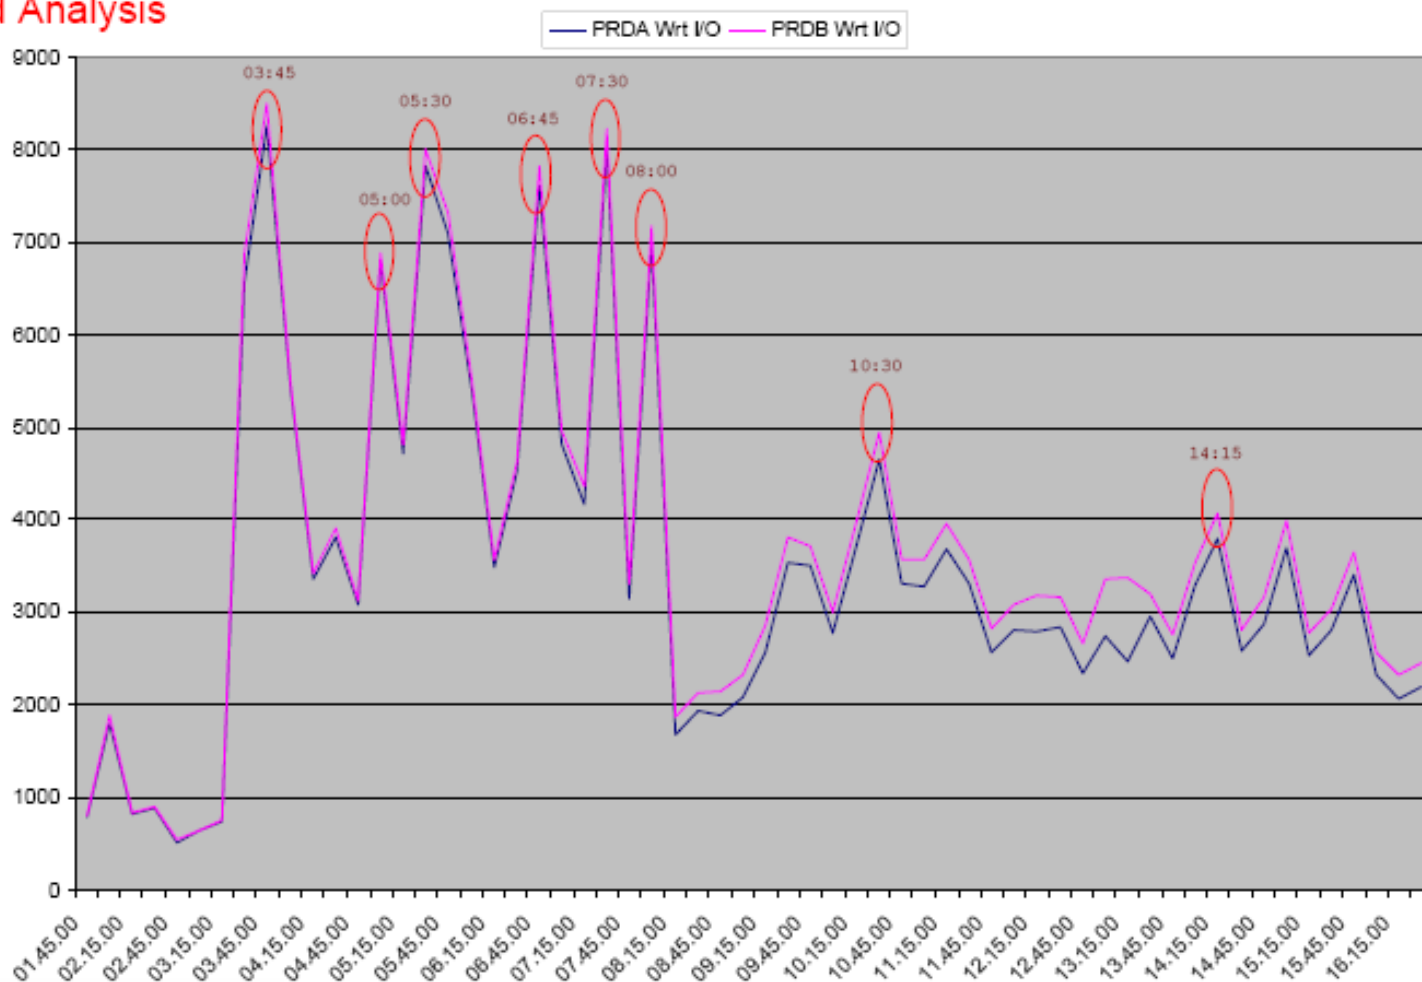
- **Hitachi Universal Replicator and data resilience**
 - Protects production performance during replication anomalies
 - Leverages Hitachi virtualization for lower-cost business continuity and disaster recovery solutions
 - Can improve customer's recovery point and recovery time objectives
 - Enables reductions in bandwidth requirements

- **Key functionality**
 - Asynchronous replication of customer's data
 - Unique cache- and disk-based journaling
 - Unique pull technology to move data to target storage system
 - Multi-volume and multi-storage-system consistency groups
 - Supports extended address volumes and Hitachi Dynamic Provisioning volumes
 - 3-data-center support: cascade, multi-target with delta-resync

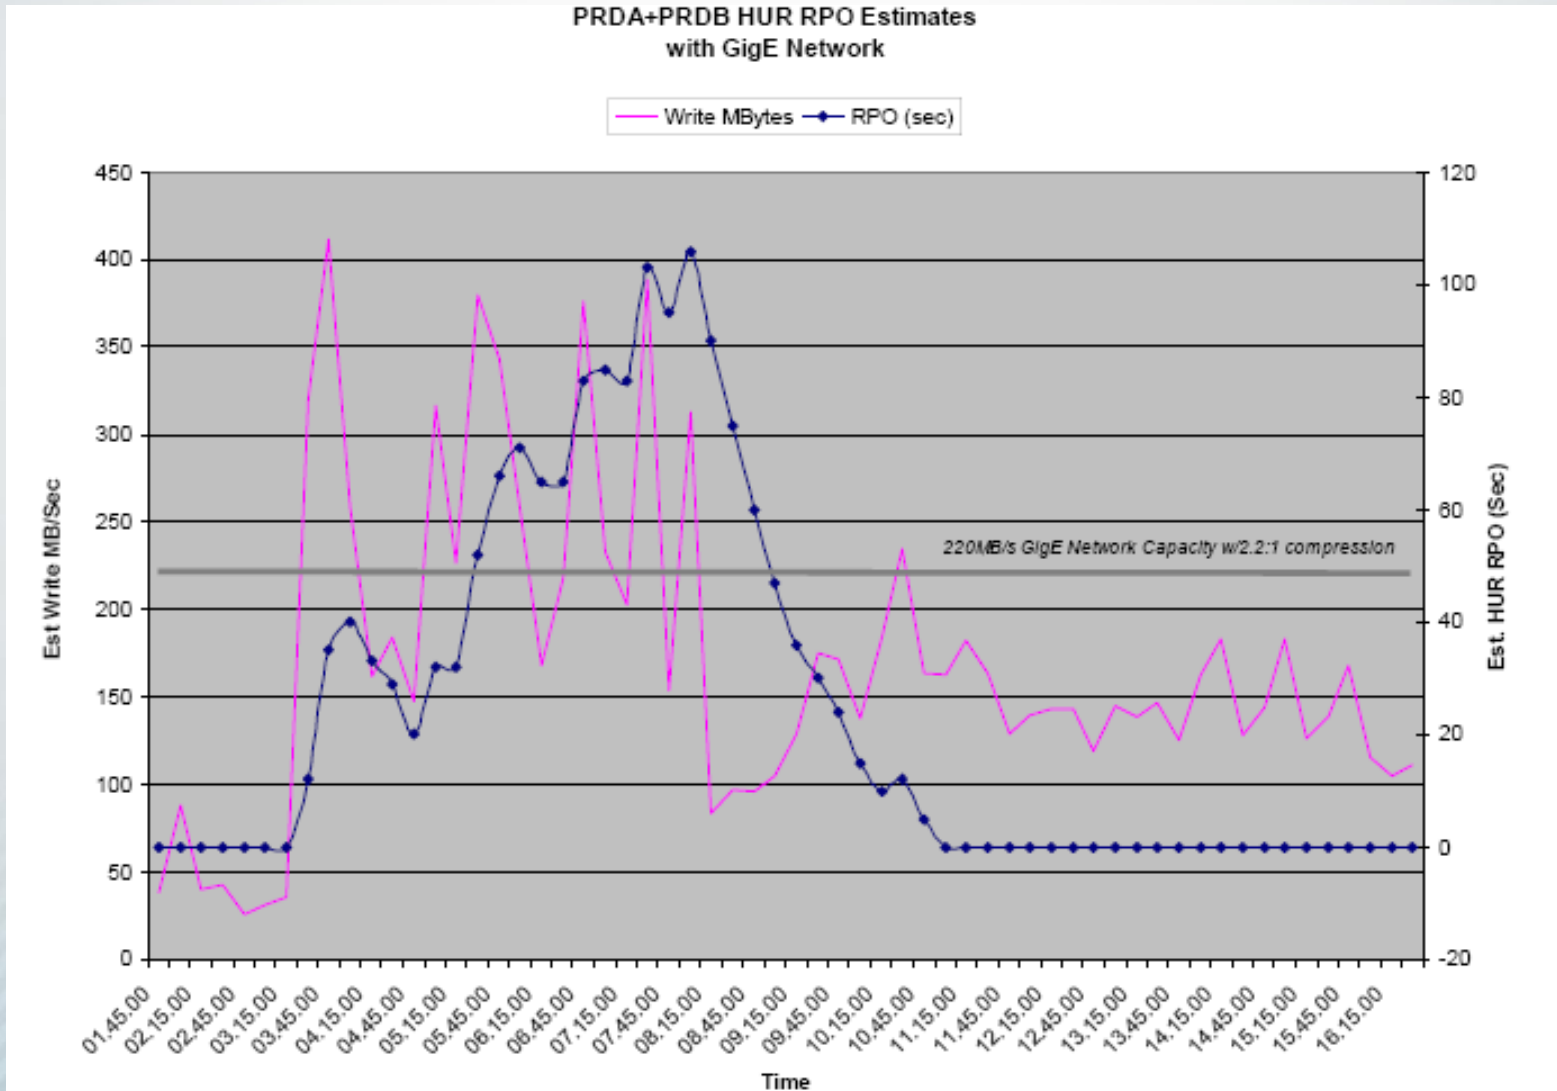
WRITE I/O ANALYSIS (PEAK WRITE LOAD IDENTIFIED)

Eight intervals were selected for Detailed Analysis

PRDA & PRDB Total Write I/O Interval Selection
November 1st 2007

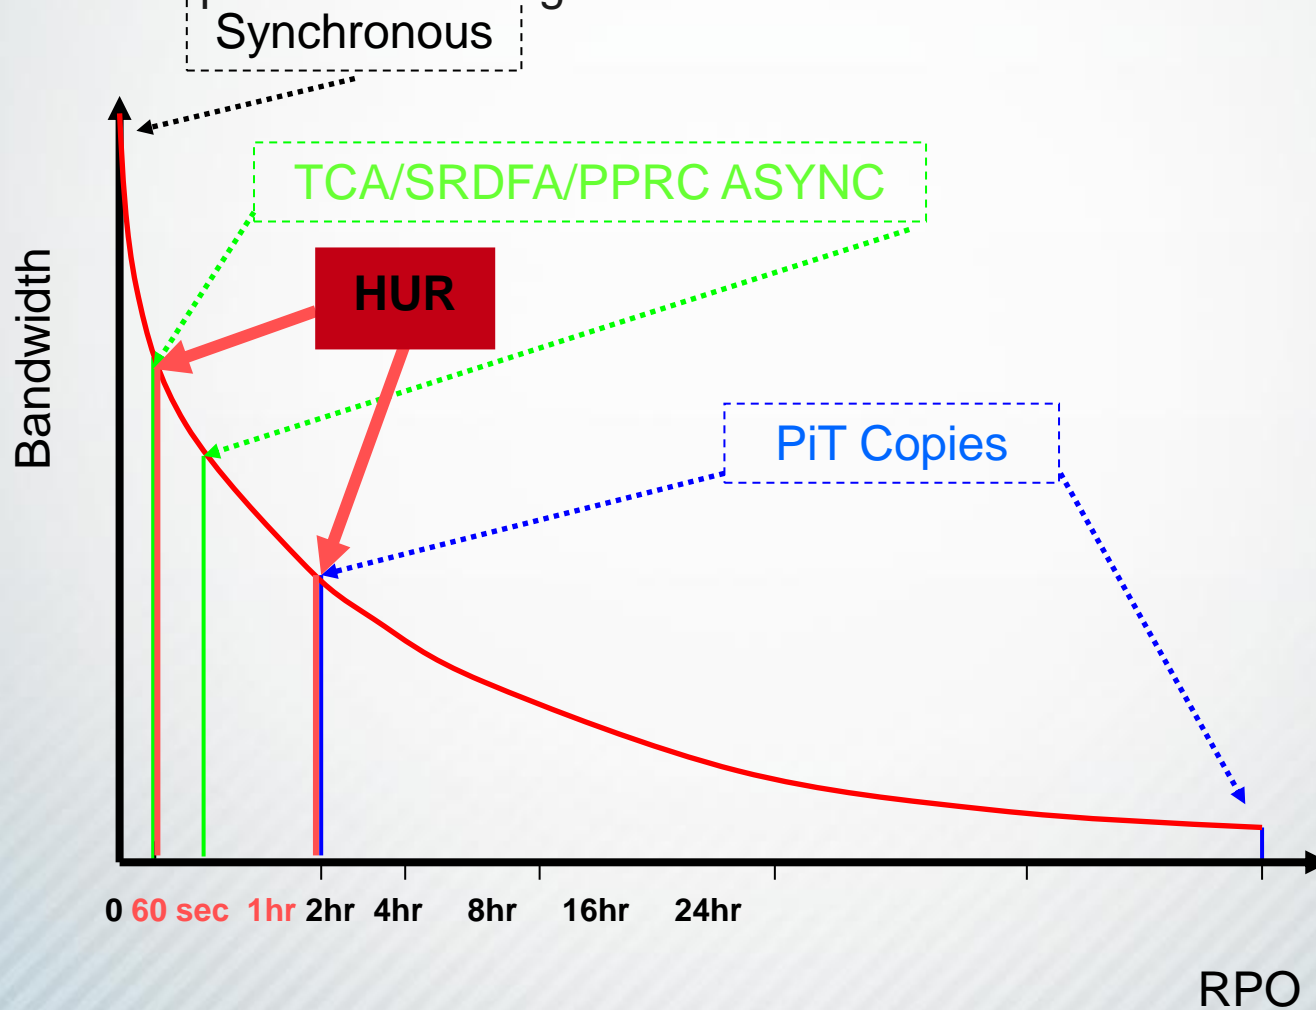


RECOVERY POINT OBJECTIVE ESTIMATE WITH LESS THAN PEAK BANDWIDTH



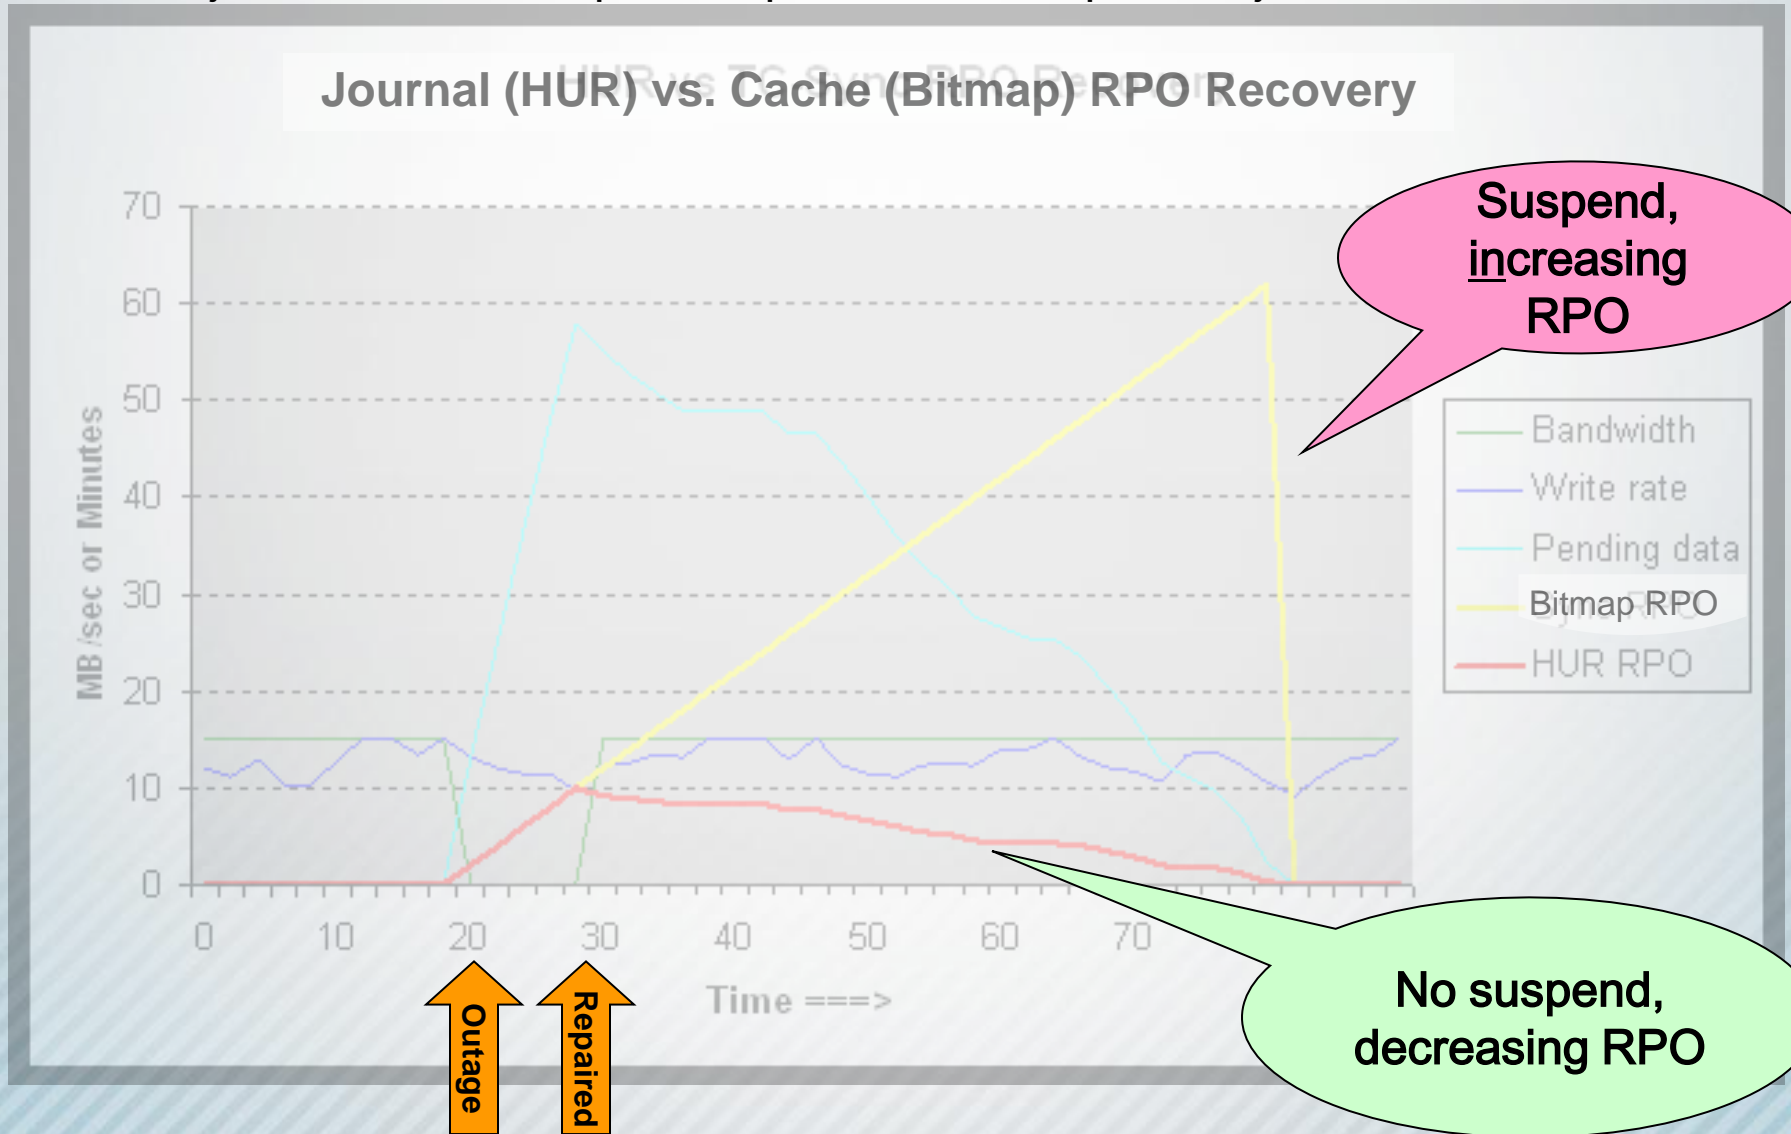
HITACHI UNIVERSAL REPLICATOR FILLS THE DATA-LOSS GAP

- Traditional replication solutions have a return point objective (RPO) gap that only Universal Replicator can bridge

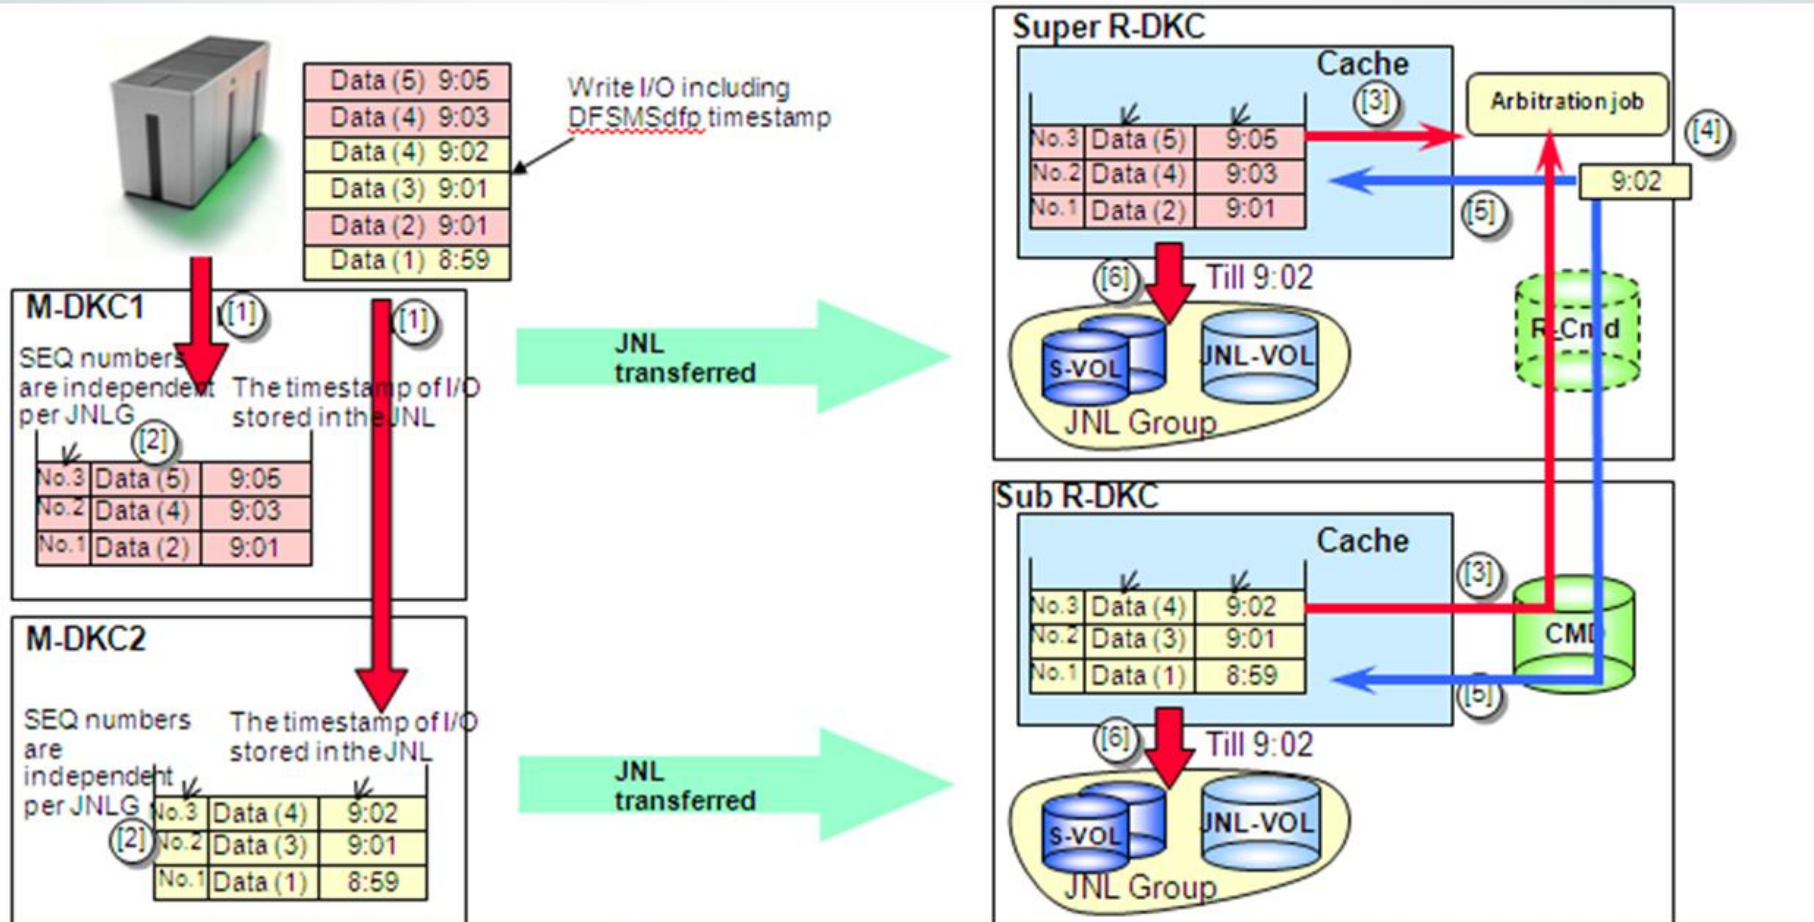


HITACHI UNIVERSAL REPLICATOR RETURN TO NORMAL OPERATION

- Disk journals avoid suspend impacts to return point objective



CROSS-CONTROLLER CONSISTENCY

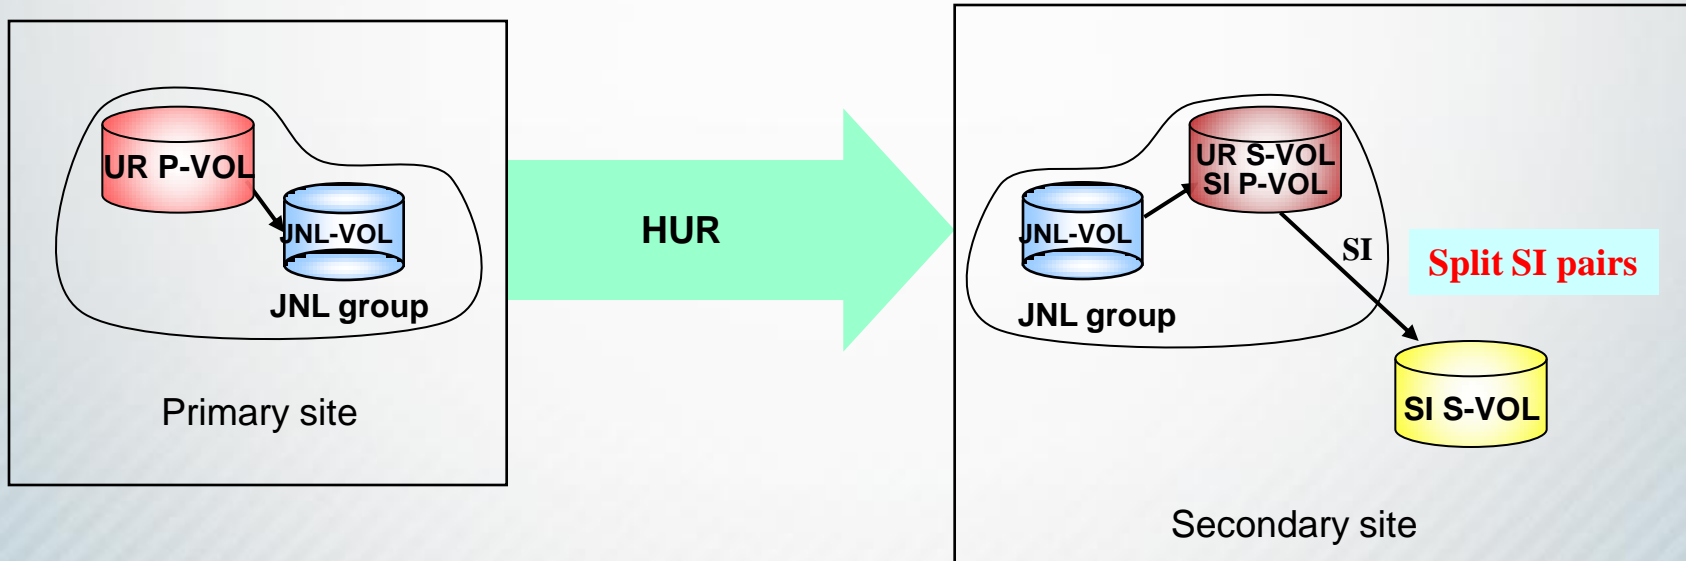


- [1] The mainframe host adds the timestamp to write data.
- [2] Each DKC also stores the timestamp to write data in the JNL. A sequence number is individually given to each JNLG.

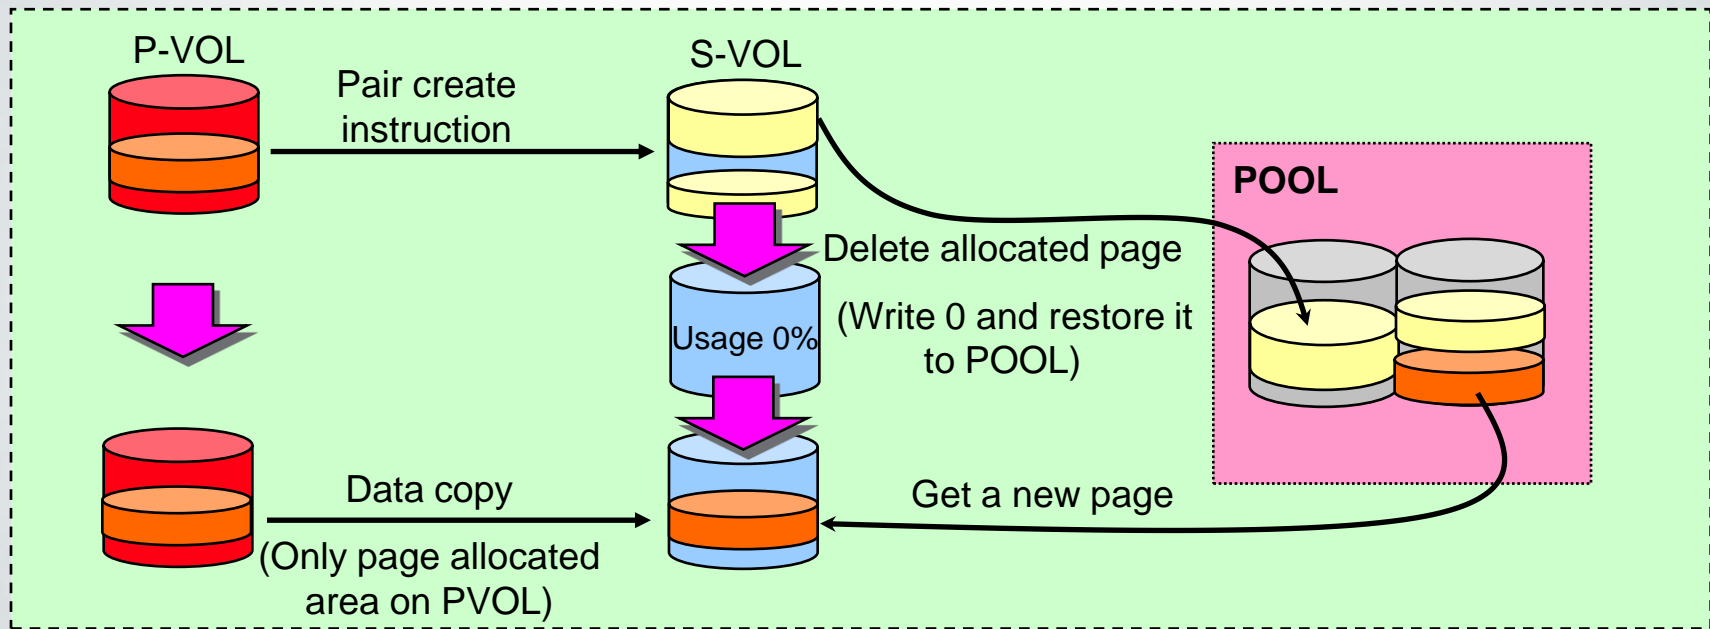
- [3] Super DKC obtains the latest timestamp (9:05 and 9:02 in the example above) of the JNL that has arrived from the R-JNLG in the EXCTG (Super DKC reads the timestamp).
- [4] Super DKC selects the oldest timestamp of all the obtained ones. (9:02 in the example)
- [5] Super DKC instructs each R-JNLG to restore data before 9:02.
- [6] Each of the R-JNLGs writes the JNL before the instructed timestamp 9:02 in the order of sequence number to the S-VOL.

Linkage between Hitachi Universal Replicator (HUR) and Hitachi ShadowImage®

- Enables consistent split of SI pairs without suspending HUR replication
- Facilitates disaster recovery testing or data repurposing
- Maintains constant and consistent disaster recovery readiness on HUR secondary c



ONE EXAMPLE OF THIN (DYNAMIC) PROVISIONING “AWARENESS” (NOW FOR MAINFRAME!)



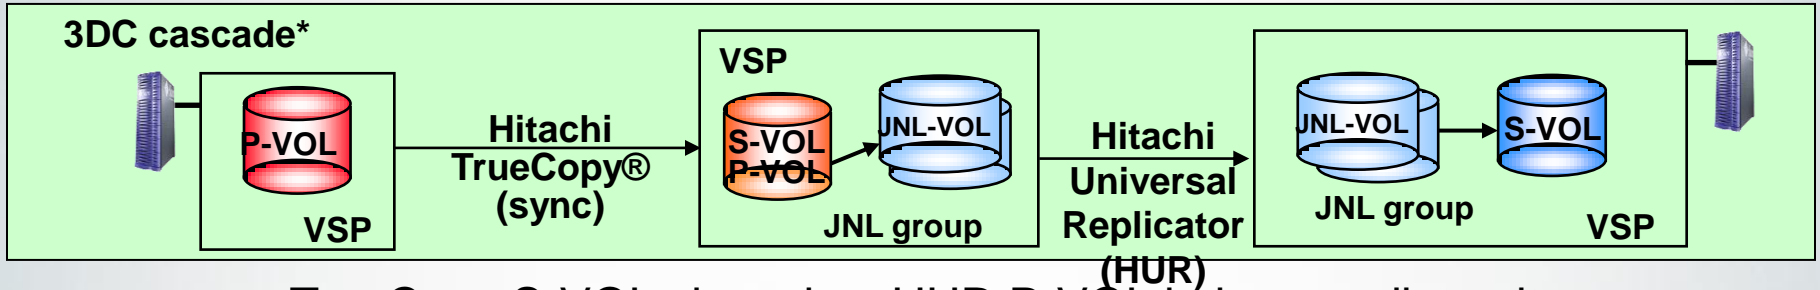
- Pages allocated to the volume used as S-VOL are deleted (0 data is written, the areas are allocated to POOL free area) when a pair is created.
- At initial copy, only the data in the page allocated area in P-VOL is transferred to S-VOL.
- Because a page deletion and initial copy run in parallel, in the case pages are allocated to a volume used as S-VOL and when the total pool capacity is not enough, initial copy progresses as the pages are deleted.

- Introduction
- Hitachi Data Systems mainframe replication portfolio
- Hitachi Universal Replicator highlights
- **Advanced replication configurations**
- Replication management
- Summary

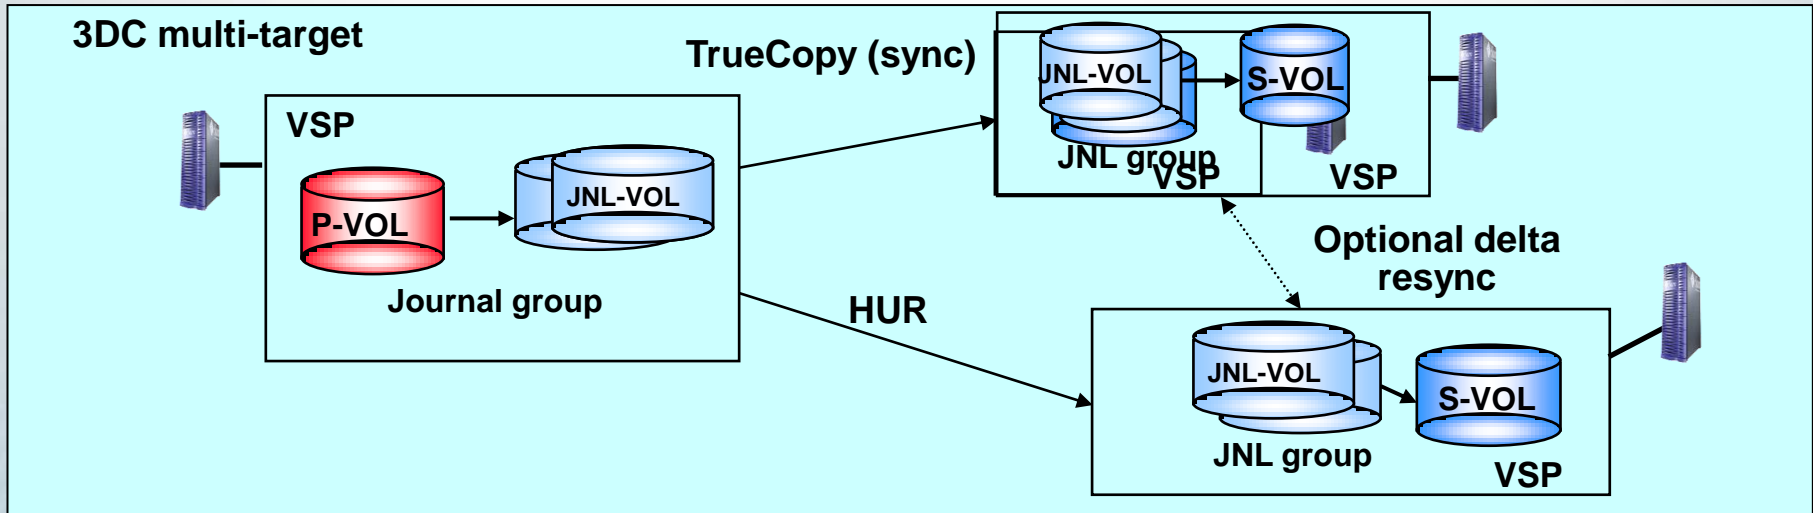
INTRODUCTION

- Primary use cases are for enhanced business continuity such as recovery from regional disasters
- Secondary use cases are for relocation or migration of primary, secondary or tertiary data centers
- Configurations discussed here allow primary and secondary volumes and consistency groups of volumes to be in more than one remote replication relationship at a time
- Replicating different volumes from multiple frames or sites to a single frame or site (i.e., fan-in) or vice versa (i.e., fan-out) is also supported
- Although these can “link” multiple frames or sites, from a replication architecture standpoint they are 2-data-center configurations, as volumes or groups of volumes are only in one remote replication relationship at a time
- Additional configurations that combine in-system replication (e.g., clones) with remote replication for either data protection or data distribution use cases are also available

3-DATA-CENTER REPLICATION SOLUTIONS: MAINFRAME



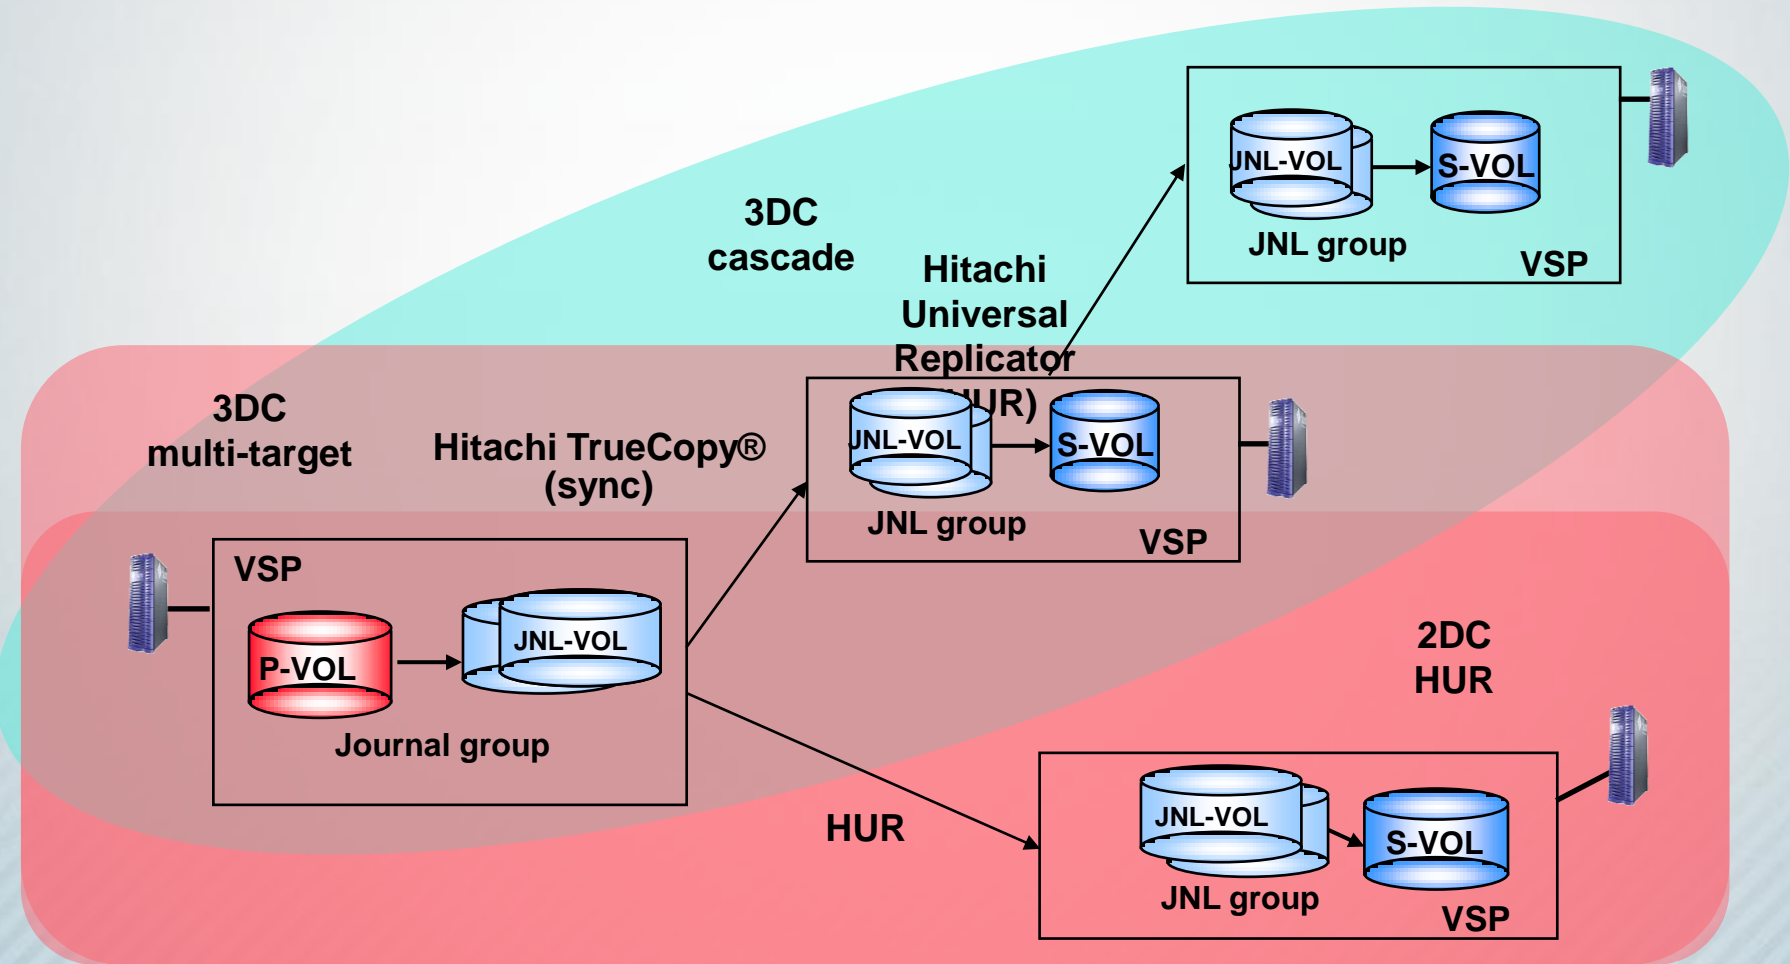
TrueCopy S-VOL shared as HUR P-VOL in intermediate site



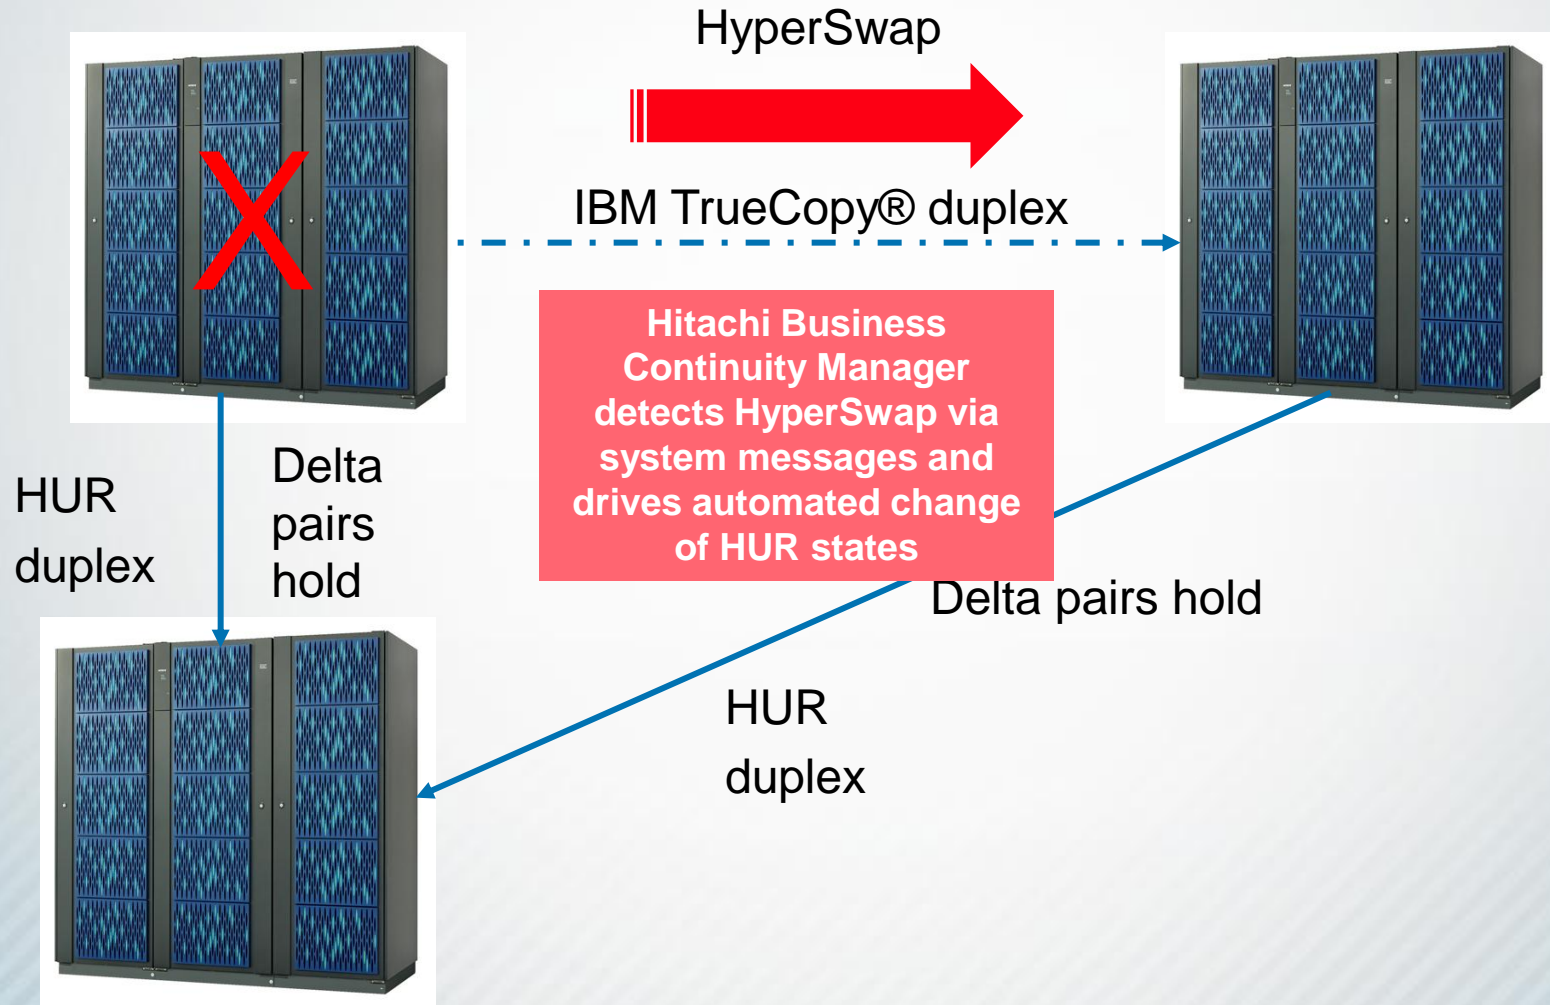
Primary volume is shared P-VOL for 2 remote systems

Both support up to 12x12x12 controller configurations

4-DATA-CENTER REPLICATION SOLUTIONS: MAINFRAME




HITACHI UNIVERSAL REPLICATOR (HUR) INTEGRATION WITH IBM® HYPERSWAP®: HIGH AVAILABILITY WITH DISASTER RECOVERY

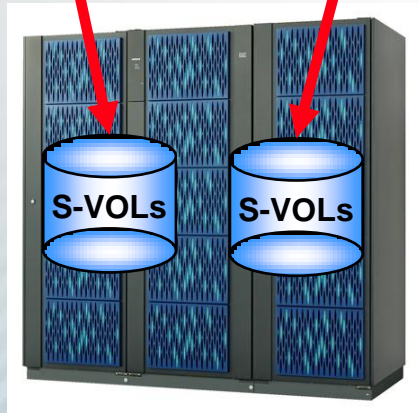


REMOTE REPLICATION FAN-IN AND FAN-OUT


FAN-IN




IBM® TrueCopy® or Hitachi Universal Replicator (HUR)



FAN-OUT

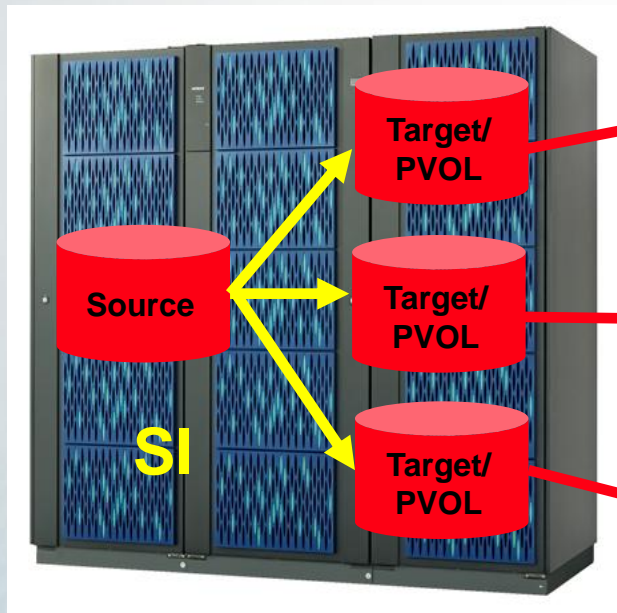


TrueCopy or HUR

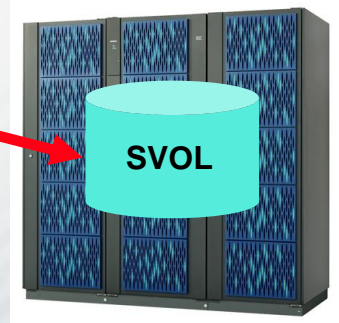
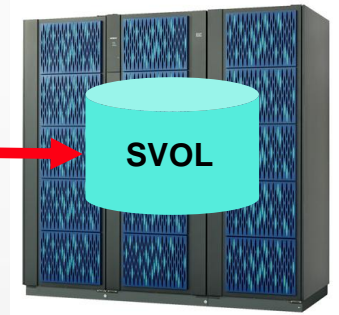
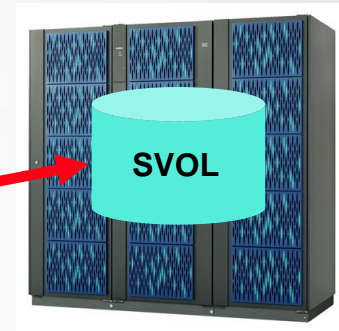


COMBINING IN-SYSTEM REPLICATION (CLONES) WITH REMOTE REPLICATION FOR FAN-OUT

Clones can be taken at
«application-aware» points in
time



**IBM® TrueCopy® or
Hitachi Universal
Replicator**



Restore of a point-in-time target to source
provides operational recovery benefit

- Introduction
- Hitachi Data Systems mainframe replication portfolio
- Hitachi Universal Replicator highlights
- Advanced replication configurations
- **Replication management**
- Summary

- **Hitachi-IBM® Compatibility**

- Hitachi storage has essential compatibility with IBM z/OS®
- IBM replication management compatibility
 - Hitachi storage controllers support PPRC, XRC® and FlashCopy® commands (assuming applicable software licensing) without additional Hitachi-provided software
 - Hitachi TrueCopy® synchronous is compatible with PPRC
 - For full volumes, IBM commands operate on individual replication pairs
 - But IBM commands do not support Hitachi Universal Replicator

- **Hitachi Business Continuity Manager (BCM) value-add**

- Reduce complexity, simplify management
- Provide enhanced control and visibility
 - (e.g., support gen'd and ungen'd devices)
- Reduce risk

- **Automatically discovers Hitachi storage systems and internal copy volumes**
- **Uses a single command to manage groups of copy volumes with common attributes**
- **Receives auto notification of key copy “state transitions”**
- **Supports advanced Hitachi TrueCopy®, Hitachi ShadowImage® and Hitachi Universal Replicator features**
- **Real-time view of Universal Replicator**
- **Single, consistent interface based on TSO/ISPF**
- **Standard REXX scripting to customize and automate**
- **Optional integration with Replication Manager for combined mainframe and open systems replication management**



- Hitachi Business Continuity Manager automatically scans and maps storage devices from MVS to their respective storage controllers

The screenshot displays the Hitachi Business Continuity Manager (BCM) interface. The window title is "GS01-1 - EXTRA! Enterprise 2000". The main display area shows a list of storage devices discovered from MVS. The list is organized into columns: SSID, CCA, Device Number, Volume, and Serial Number. The devices are listed in a table format, with SSID ranging from 7710-00 to 7710-1F. The CCA column shows values like CC00, CC01, etc., and the Device Number column shows values like GSC400, GSC401, etc. The Serial Number column shows values like PICC05, PICC06, etc. The interface also shows a command prompt at the top with "Command ===>" and "Scroll ===> PAGE". The bottom status bar indicates "Connected to host 172.17.51.13 (TCP00022)" and "7:58 AM".

```
Command ===> _____ Command Control Address (CCA) Selection Lis Row 1 to 32 of 84
                                          Scroll ===> PAGE
                                          2010/04/01 04:58:04

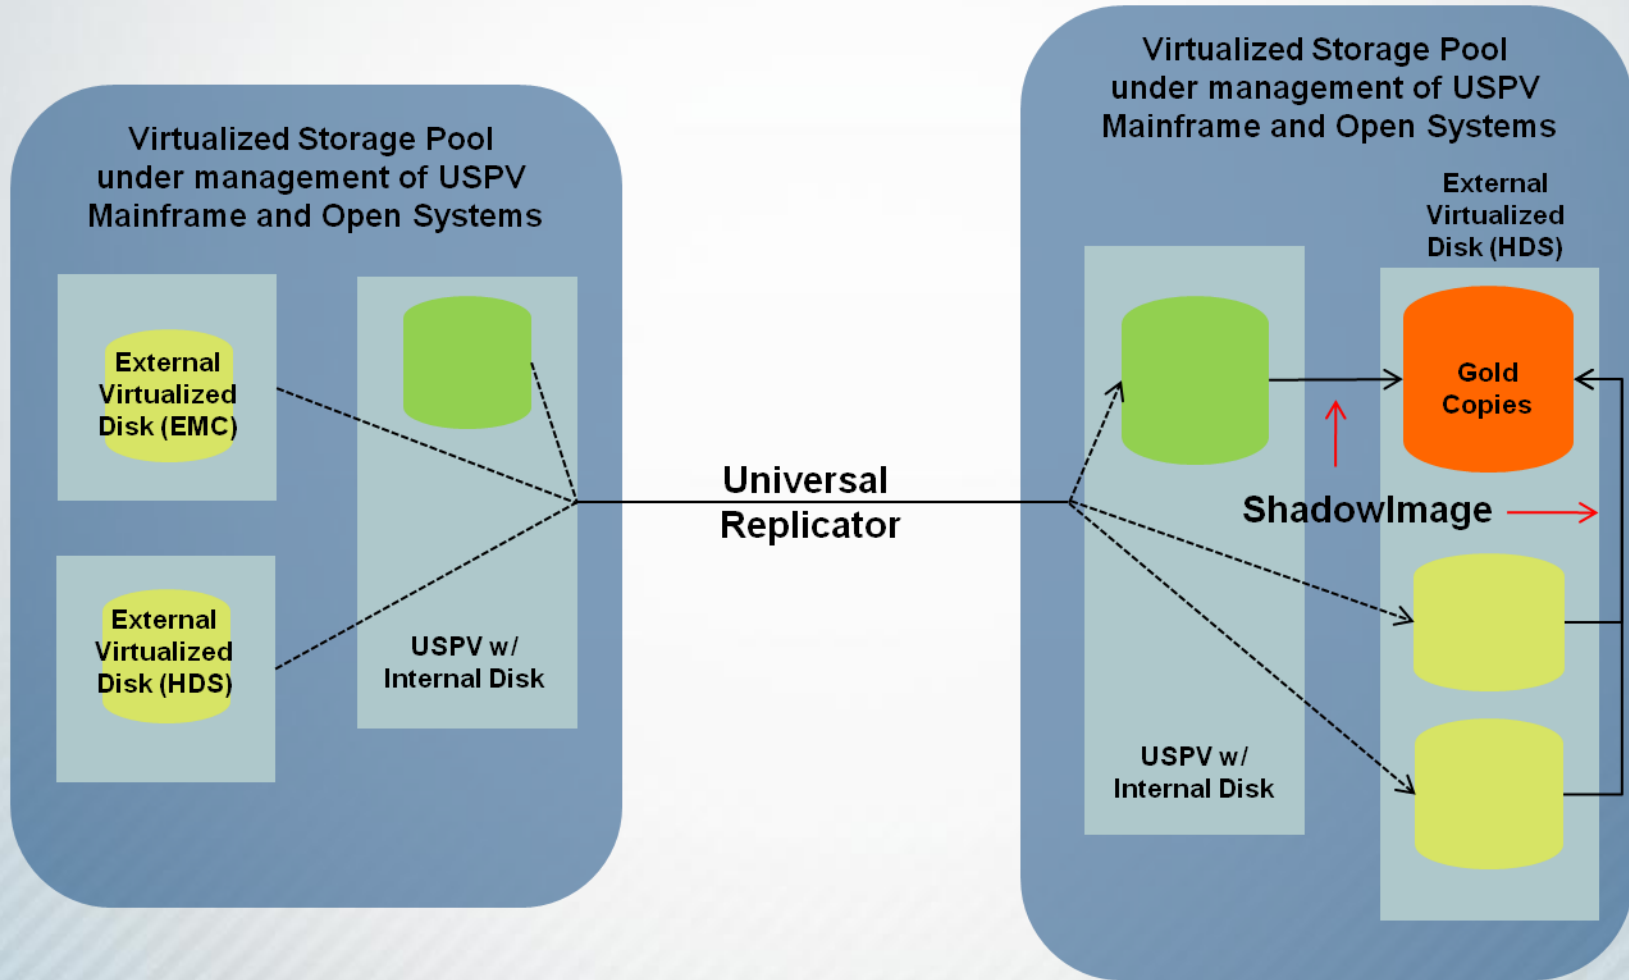
Device Address Domain : DAD
Description :
DKC S/N : 22471 Model : USP uCode : 50090490 IFTYPE : 1414
Control Unit : 10

SSID CCA Device Number Volume Serial Number
7710 00- CC00 GSC400
7710 01- CC01 GSC401
7710 02- CC02 GSC402
7710 03- CC03 GSC403
7710 04- CC04 GSC404
7710 05- CC05 PICC05
7710 06- CC06 PICC06
7710 07- CC07 GSC407
7710 08- CC08 GSC408
7710 09- CC09 GSC409
7710 0A- CC0A
7710 0B- CC0B
7710 0C- CC0C PICC0C
7710 0D- CC0D
7710 0E- CC0E
7710 0F- CC0F
7710 10- CC10
7710 11- CC11
7710 12- CC12
7710 13- CC13
7710 14- CC14
7710 15- CC15
7710 16- CC16
7710 17- CC17
7710 18- CC18
7710 19- CC19
7710 1A- CC1A
7710 1B- CC1B
7710 1C- CC1C
7710 1D- CC1D
7710 1E- CC1E
7710 1F- CC1F

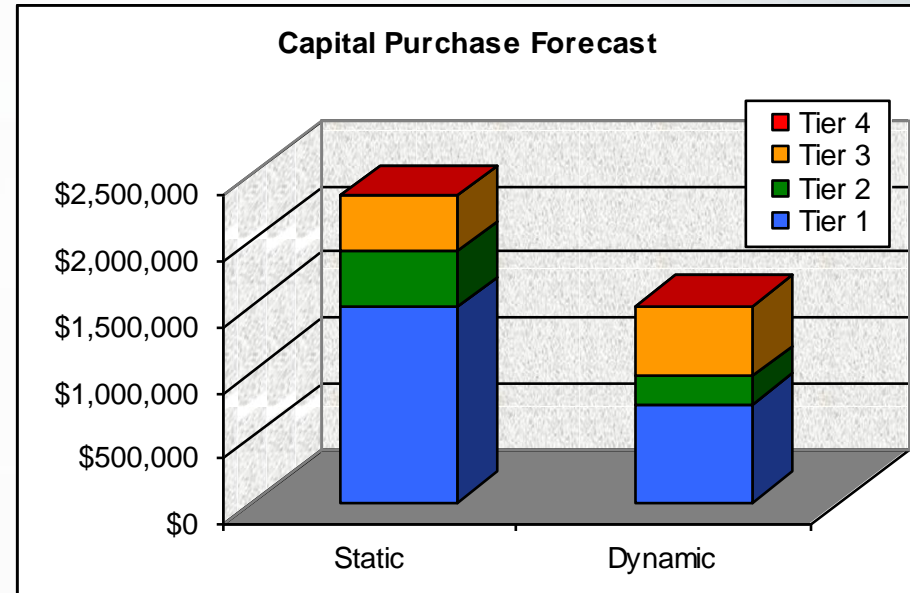
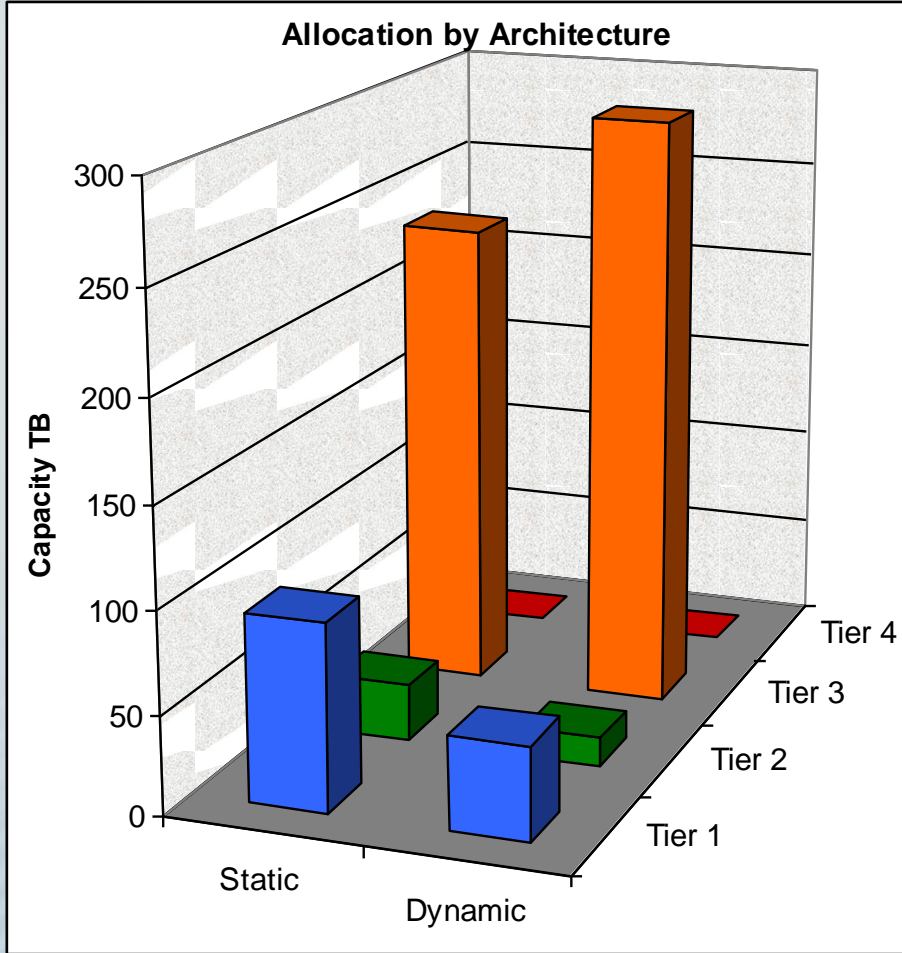
F1=Help F3=Exit F7=Backward F8=Forward F12=Cancel F17=DispConf
```


- Introduction
- Hitachi Data Systems mainframe replication portfolio
- Hitachi Universal Replicator highlights
- Advanced replication configurations
- Replication management
- **Summary**

HDS DISASTER RECOVERY REPLICATION SOLUTION (LEVERAGING EXTERNAL VIRTUALIZATION)



STATIC “IN-THE-BOX” VERSUS DYNAMIC “EXTERNAL” QUANTIFIED CAPEX AVOIDANCE (BASED ON HDS)



**Total additional savings
over 4 years**

\$ 841,353

This is the effect of putting local
and/or replication gold copies for
Tier-1 and Tier-2 in Tier-3

■ Final thoughts

- Hitachi Data Systems (HDS) is rooted in, and committed to, strong IBM compatibility
- HDS has a complete replication solution portfolio for mainframe storage
- Hitachi Universal Replicator is a mature product (2004) with over 250 worldwide mainframe customers, and is on its third generation of hardware platform support
- HDS storage provides one platform for virtualization, dynamic (thin) provisioning and tiered storage management

■ Benefits:

- Deliver SLAs that are better aligned with business objectives
- Reduced risk via proven, enterprise-class data protection
- Cost savings through
 - Best-of-breed asynchronous replication
 - Seamless tiered storage management
 - Protection of existing investments with external virtualization
 - Dynamic provisioning (now for mainframe)

QUESTIONS?

- **Mainframe Series**
 - **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!