#SHAREorg

SHARE
Technology · Connections · Results

What's New in DFSMShsm

Glenn Wilcock
wilcock@us.ibm.com

March 14, 2012
Session 10950

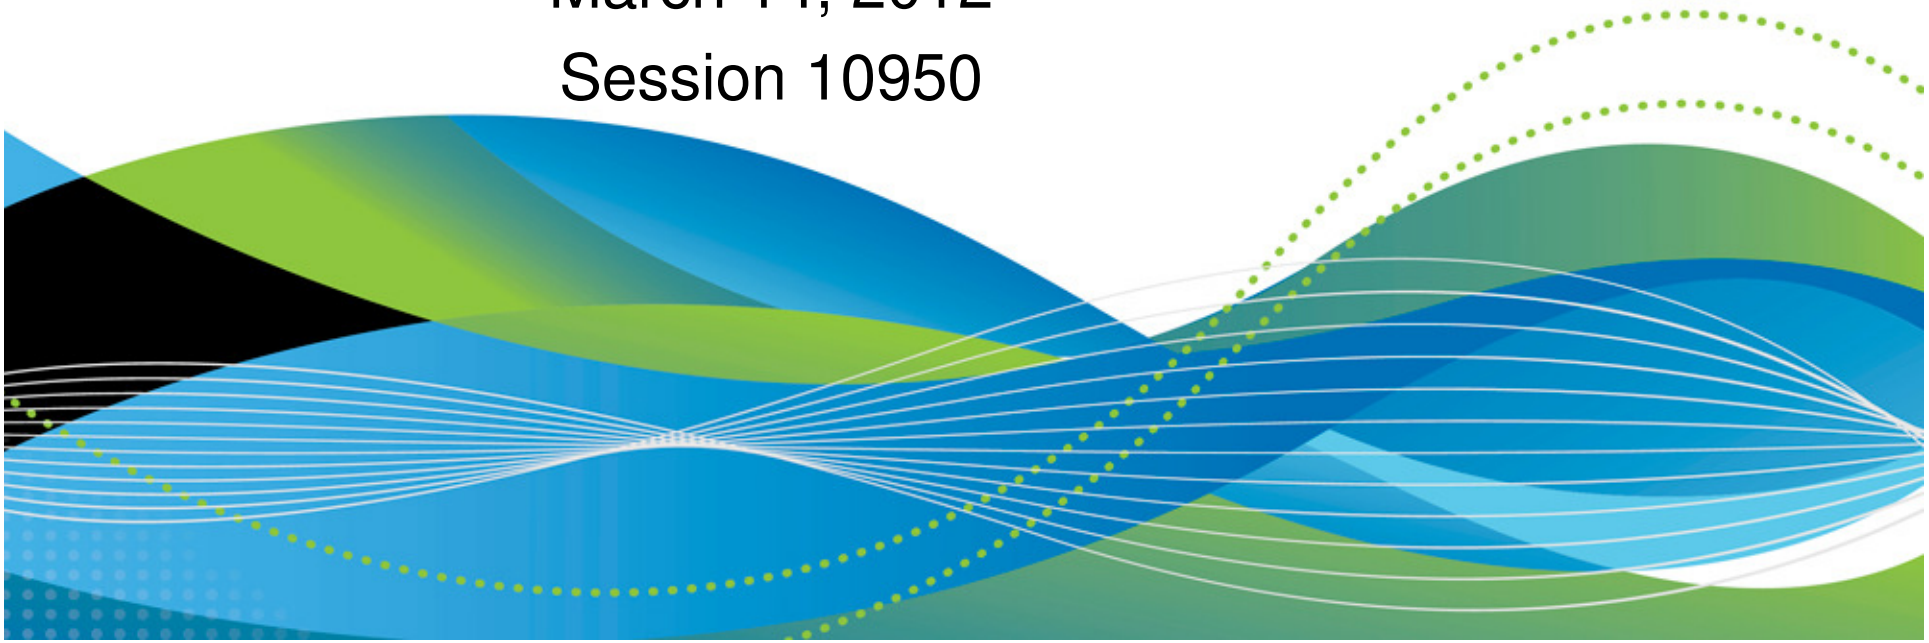

Legal Disclaimer

NOTICES AND DISCLAIMERS

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Trademarks

The following are trademarks of the *International Business Machines Corporation:*

IBM, DFSMS/MVS, DFSMShsm, DFSMSrmm, DFSMSdss, DFSMSopt, DFSMS Optimizer, z/OS, eServer, zSeries, MVS, FlashCopy®

The information contained in this presentation is distributed on an 'AS IS' basis without any warranty
either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. The use of this information is a customer responsibility and depends on the customer's ability to evaluate and integrate it into the customer's operational environment.

Agenda

- **General**
  - NonDisruptive CDS Backup
  - ONLYIF
  - Usability
  - CA Reclaim
  - Deduplication
  - EAV Support
  - Cross Memory
- **Migration**
  - On Demand Migration
  - Performance Improvements
- **Backup**
  - BD Exit
  - Dump Block Size
  - Dump Stacking Limit
  - Multitask Recovery from Dump Tapes
- **Fast Replication**
  - Copy Pool Recovery from Dump Tape
  - Fast Reverse Restore
  - Space Efficient
  - Preserve Mirror
  - UCB Refresh
- **Miscellaneous**

General: V1R13
Nondisruptive CDS Backup

- Today, CDS Backup can be very disruptive to other HSM activity
  - All other HSM activity must be quiesced before CDS Backup can start
    - Some customers HOLD all HSM activity prior to the start of CDS Backup to ensure that it can begin at its scheduled time
  - Functions that start while CDS Backup is waiting to start have to wait until the completion of CDS Backup
  - Higher impact in an RLS environment than nonRLS
- Journal is backed up using Standard I/O, even when Concurrent Copy is specified
  - Since there is a chance for Concurrent Copy to fail the physical copy after logical completion, Standard I/O is always used for the journal to ensure it is not nulled without being copied in its entirety
  - Outage for CDS backup is at least as long as the time it takes to backup the journal

General: V1R13
Nondisruptive CDS Backup

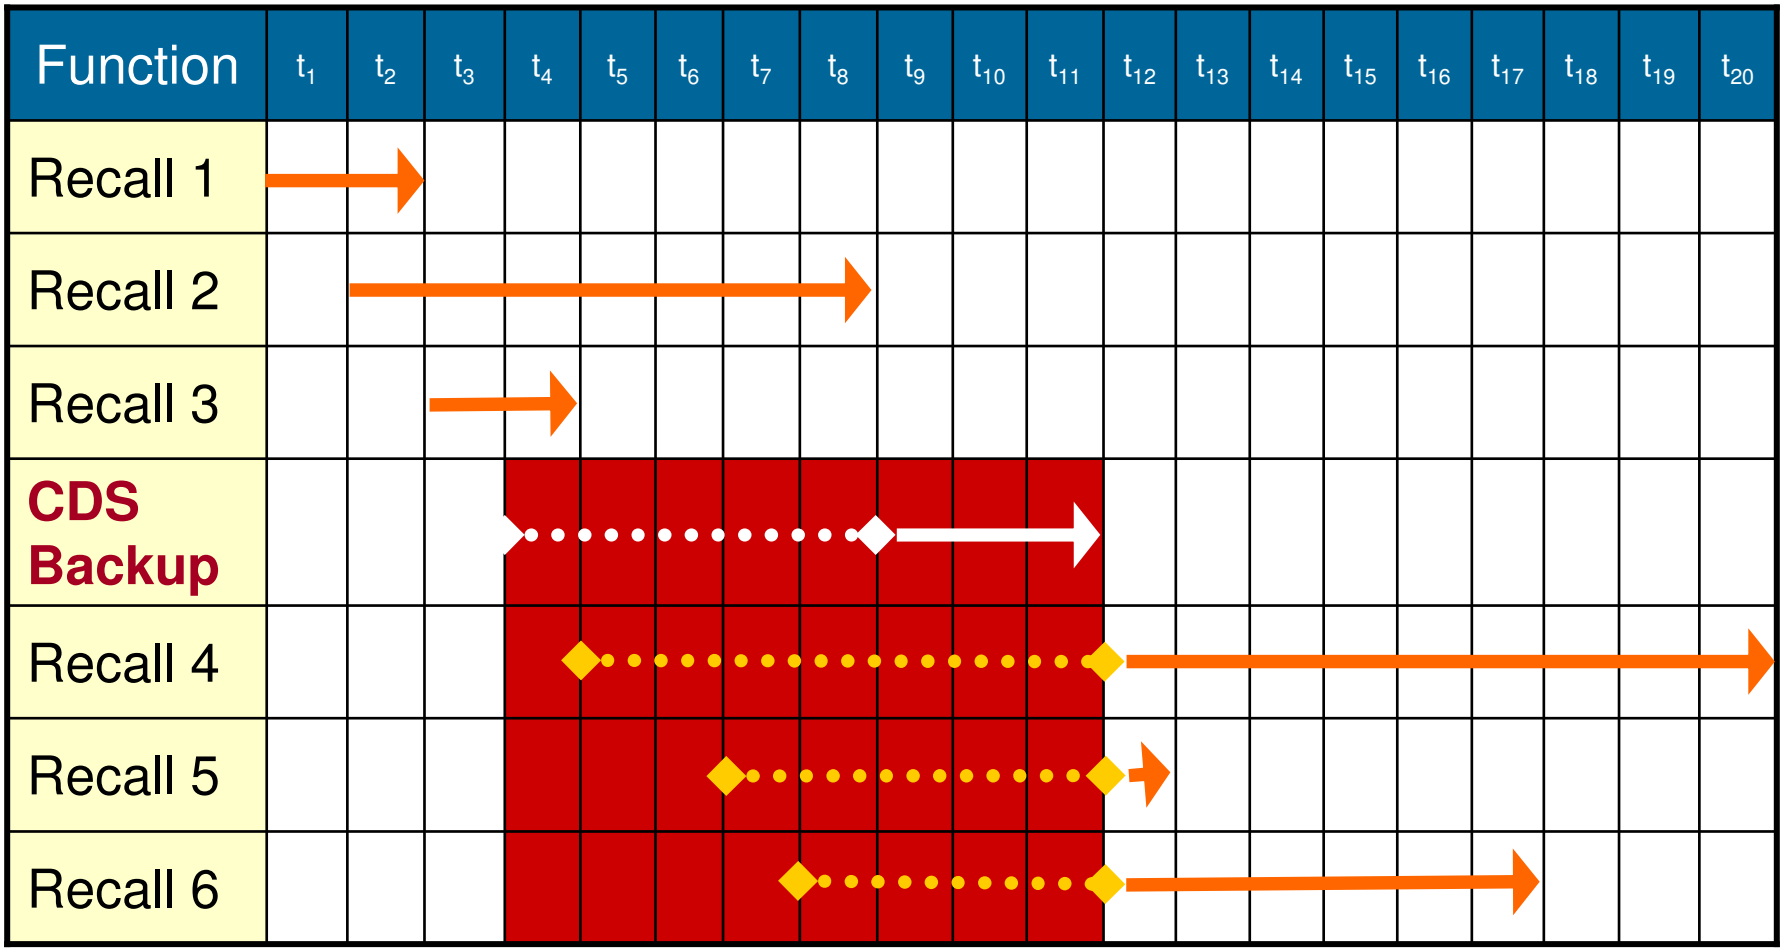

New HSM Activity Quiesced

Pre-V1R13

General: V1R13
Nondisruptive CDS Backup

- nonRLS Environment
  - Only HSM activity on the *same LPAR* impacts/is impacted by CDS Backup
    - Serialization scheme uses enqueue scope of SYSTEM and a Reserve to cover other systems
- RLS Environment
  - HSM activity on *any host* in the HSMplex impacts/is impacted by CDS Backup
    - Serialization scheme uses enqueue scope of SYSTEMS

General: V1R13
Nondisruptive CDS Backup

- V1R13 Enhancements
  - CDS Backup serialization scheme has been enhanced such that all active HSM activity *does not* have to complete before CDS Backup can begin
    - CDS and Journal I/O is quiesced before and during copy of control data sets and journal to ensure a data consistent backup
    - When concurrent copy is used, this is a brief disruption
  - The backup of the Journal will begin before the CDSes are quiesced
    - HSM activity can continue while the 'static' portion of the journal is backed up
    - Activity is quiesced during brief time required to backup the remainder of the journal
    - Requirements:
      - All CDS clusters are SMS-managed
      - Concurrent Copy specified
      - SETSYS CDSVERSIONBACKUP(DATAMOVER(DSS))
      - SETSYS JOURNAL(RECOVERY)

General: V1R13
Nondisruptive CDS Backup

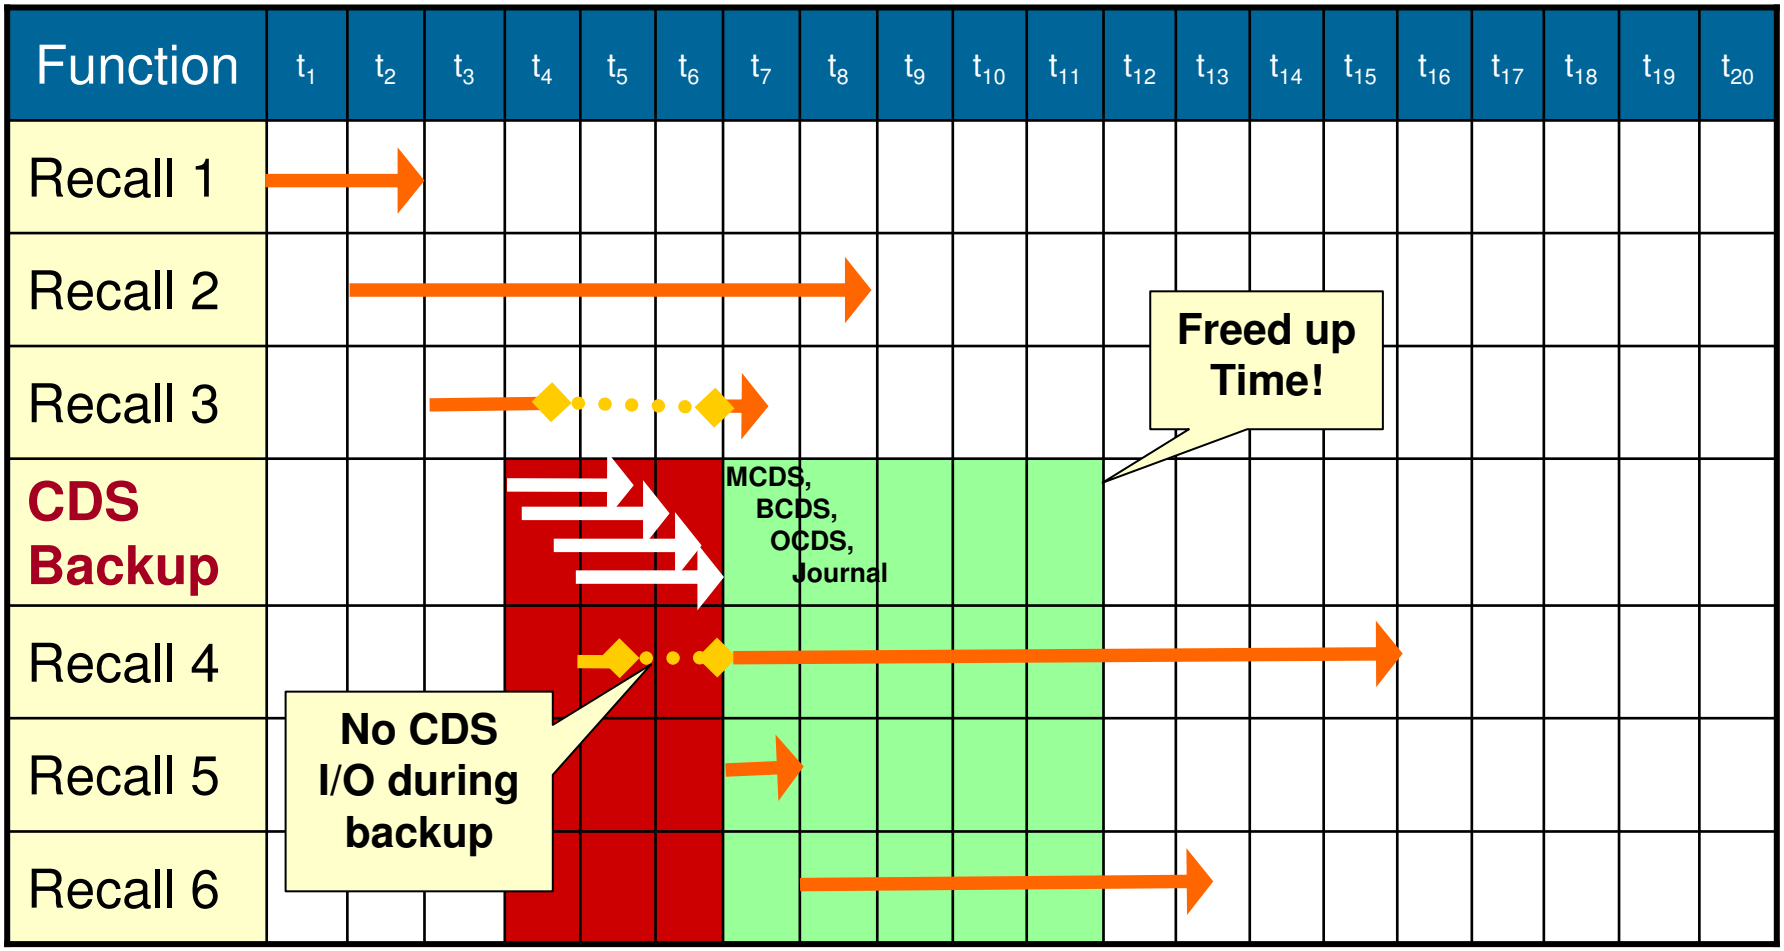

HSM Activity Quiesced

V1R13: Not CC

General: V1R13
Nondisruptive CDS Backup

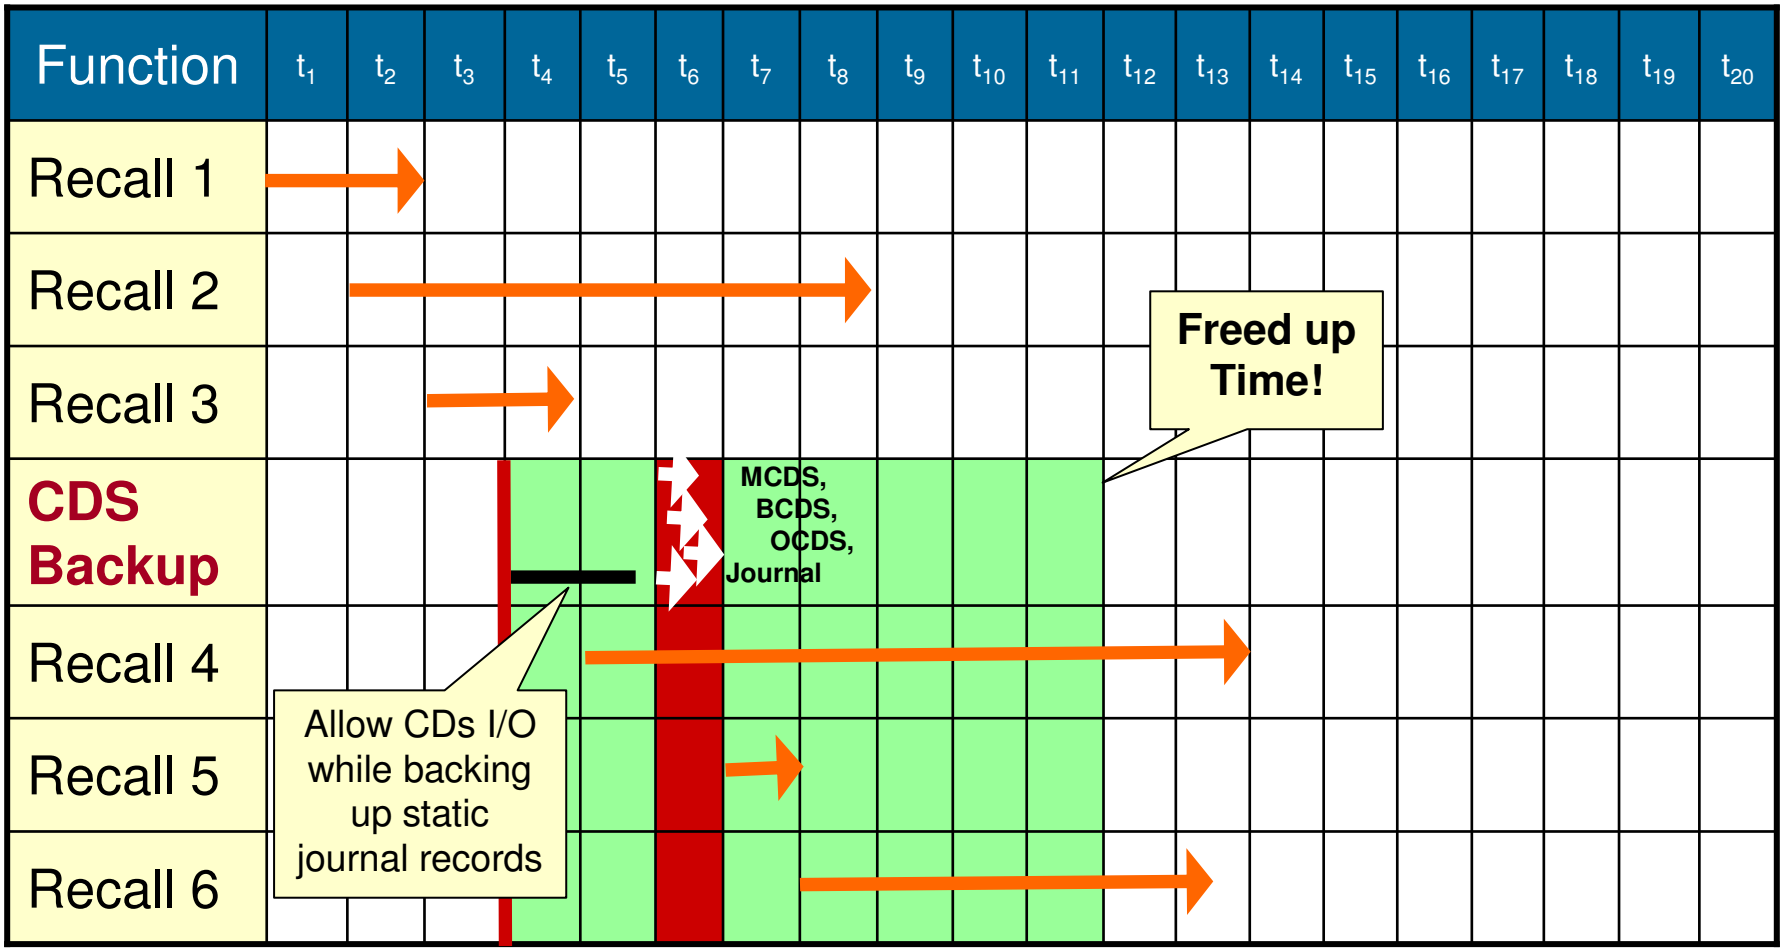

HSM Activity Quiesced

V1R13: CC

General: V1R13
Nondisruptive CDS Backup

Journal Backup Detail

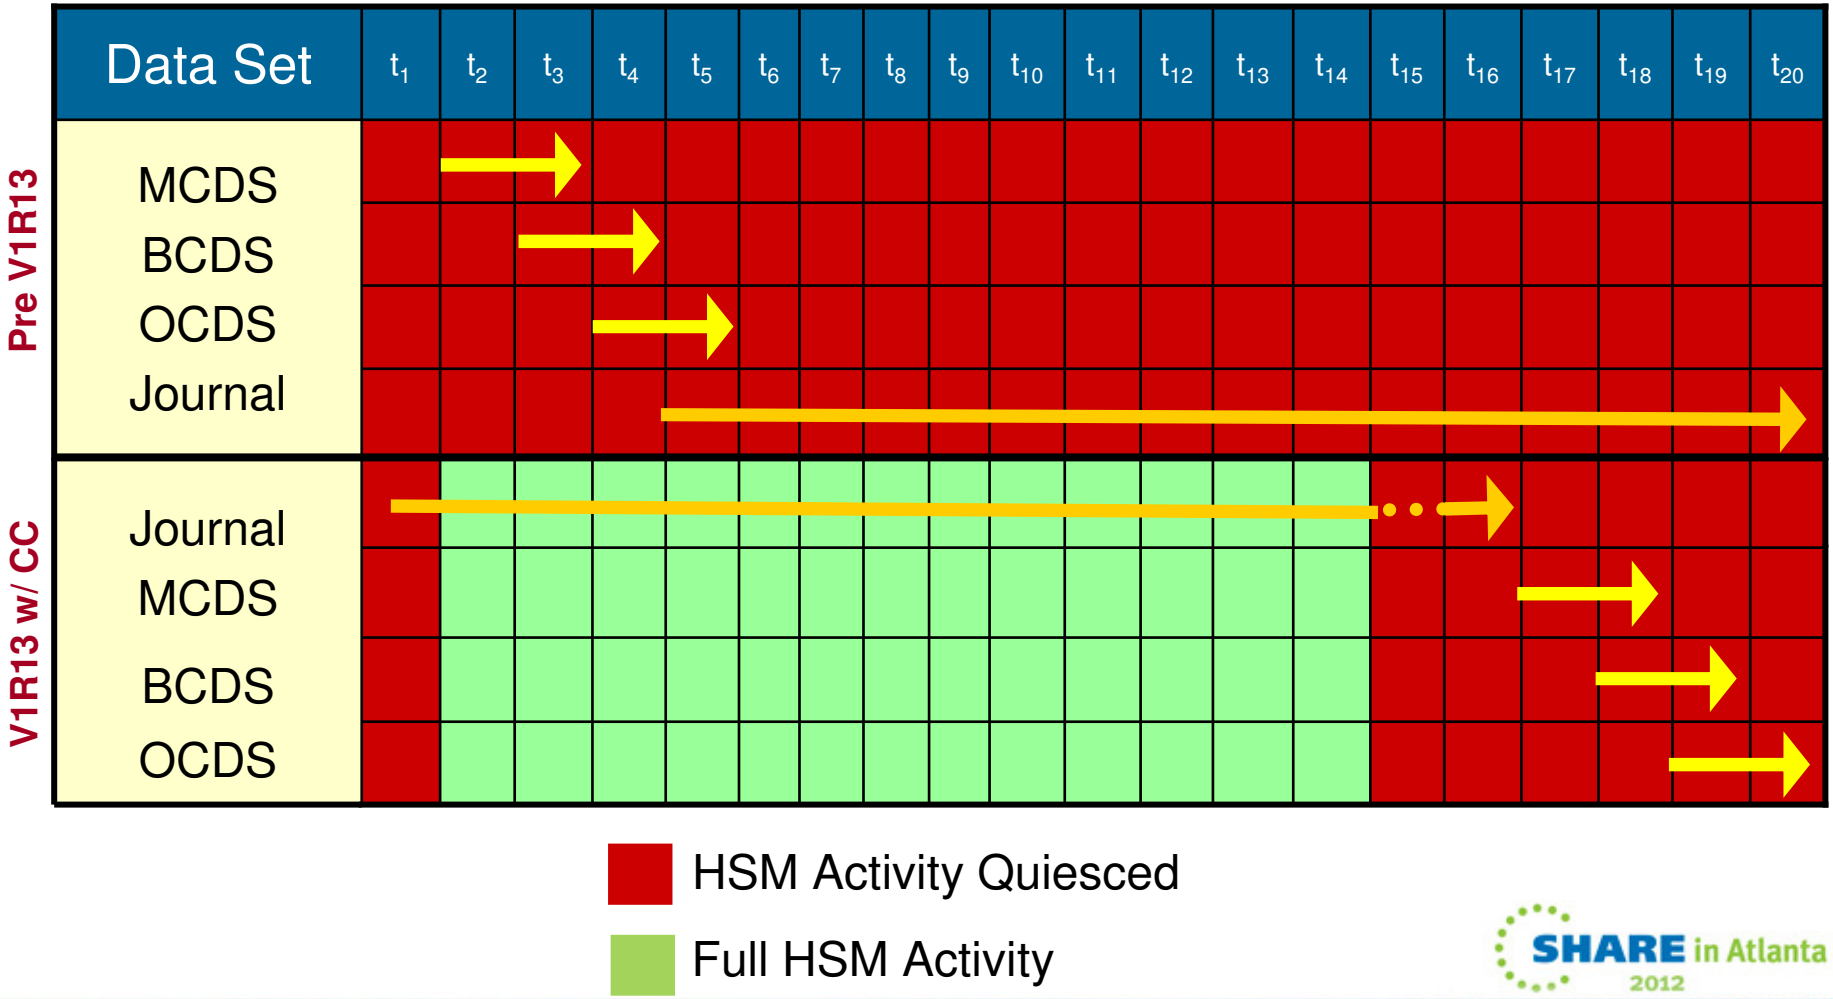

General: V1R13
Nondisruptive CDS Backup

- CDS Recovery process is unchanged
- Migrating to new function
  - The journal backup enhancement requires the coexistence APAR to be active for a full backup cycle before becoming fully enabled
  - If you use the function on a V1R13 system that never had the coexistence applied, then not until the *second* and subsequent CDS backups will the function be fully enabled
  - Remove DFSMShsm HOLD commands scheduled before the start time of CDS Backup
- Coexistence
  - When using RLS, the full benefit isn't seen until *all* hosts are at V1R13

General: V1R13 ONLYIF

- IBM recommends a single, shared ARCCMDnn member for all DFSMShsm hosts
  - Eliminates chance of discrepancies between ARCCMDnn members across systems
- Pre-V1R13 ONLYIF support is *very* basic
  - Used when commands are unique to one or more DFSMShsm hosts
  - Example: Only hosts A & B run Automatic Dump:

```
ONLYIF HSMHOST(A)
    SETSYS ADSTART(1800 2400)
ONLYIF HSMHOST(B)
    SETSYS ADSTART(1800 2400)
ONLYIF HSMHOST(A)
    SETSYS MAXDUMPTASKS(7)
ONLYIF HSMHOST(B)
    SETSYS MAXDUMPTASKS(7)
```

General: V1R13 ONLYIF

- V1R13
  - Support for BEGIN and END operators
  - Enables multiple host IDs on the HSMHOST keyword

```
ONLYIF HSMHOST(A,B)
  BEGIN
    SETSYS ADSTART(1800 2400)
    SETSYS MAXDUMPTASKS(7)
  END
```

- Coexistence
  - Pre-V1R13 systems will correctly parse new syntax introduced on a V1R13 system

General: V1R13 Usability

- RELEASE RECALL(DASD)
  - If there is an issue with tape drives and HOLD RECALL is issued, there used to be no way to release just ML1 recalls without also releasing tape recalls
  - New RELEASE RECALL(DASD) converts a HOLD RECALL state to a HOLD RECALL(TAPE) state
- QUERY COMMONQUEUE(RECALL) output shows the host that initiated each recall request
  - If a request in the CRQ needs to be canceled, the cancel must be done on the host that initiated the Recall
  - New information on ARC1543I eliminates need to issue QUERY on each HSM host

General: V1R12
CA Reclaim

- Necessity to reorganize the DFSMShsm CDSs impacts…
  - Availability – Must stop all DFSMShsm hosts
  - Performance – increased CI/CA splits afterwards
  - Integrity - Doing it incorrectly is a common cause for CDS breakage
- **VSAM** provided a **CA (Control Area) Reclaim** function
  - Reclaims empty CAs that remain after all records are deleted
  - Common for DFSMShsm
    - Many records written with the date in the key
- What this means for DFSMShsm…
  - Significantly *reduces* the need for CDS reorgs
  - Significantly *reduces* the need for SDSP reorgs
  - *Doesn't eliminate* need to reorg
- Session 10964 (Wednesday 1:30pm)

General: V1R12
CA Reclaim

CDS w/o CA Reclaim

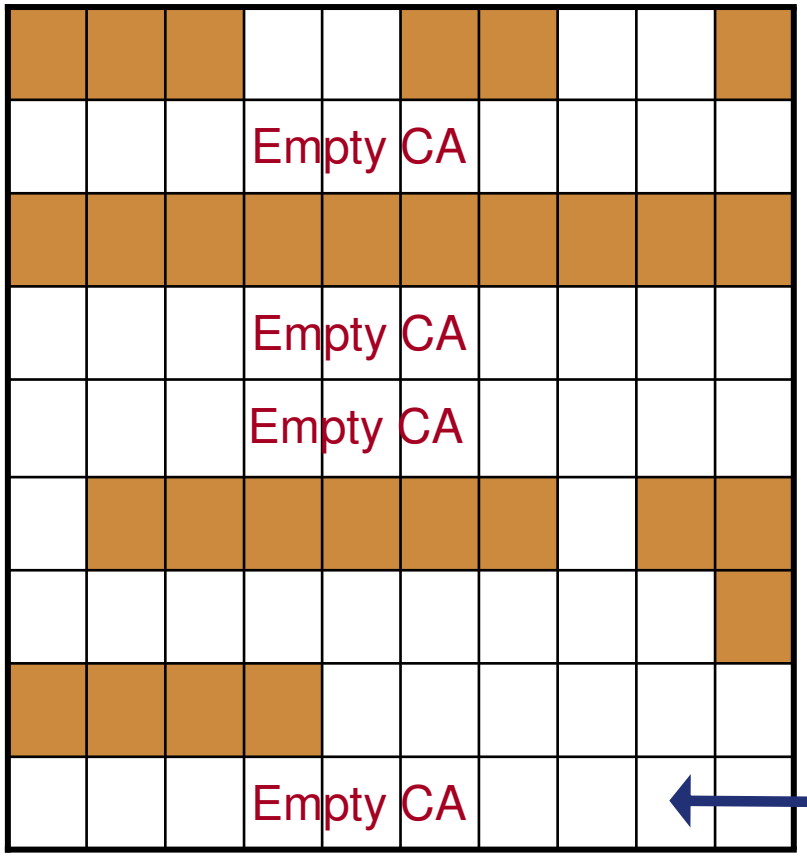

CDS w/ CA Reclaim Enabled

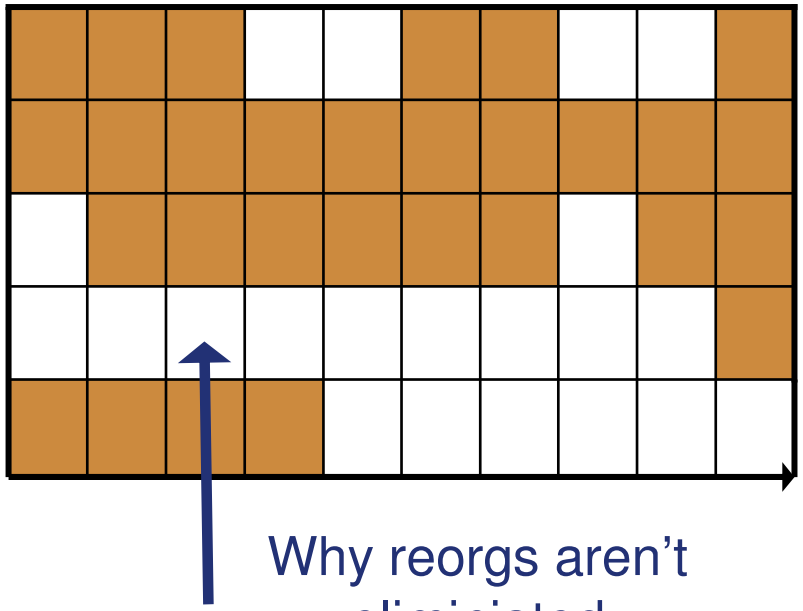

Why reorgs aren't eliminiated

Reorg required to reclaim empty CAs!

General Deduplication

- TS7680 ProtecTier Deduplication
  - Virtual Tape Library
  - Inline data deduplication
  - Replication – Only deduplicated data is transmitted

- DFSMShsm data
  - DFSMShsm ‘wraps’ blocks of native data with meta data, making all blocks unique
  - TS7680 has logic specific for DFSMShsm blocks so that it can deduplicate them

General
Extended Address Volumes

V1R12

- Support for all data set types
- ★ ML1 / $ML2_{disk}$ and Backup volumes may now use all space

SETSYS USECYLINDERMANAGEDSPACE(Y|N)

! DO NOT specify (Y) until ALL hosts are V1R12 or higher

★ Good candidates for ML1 Overflow volumes

! Use caution when using as standard ML1

- All of the unused space cannot be released from migration data sets because allocation is done in Multi-Cylinder Units
- Leaves free space in migration copy

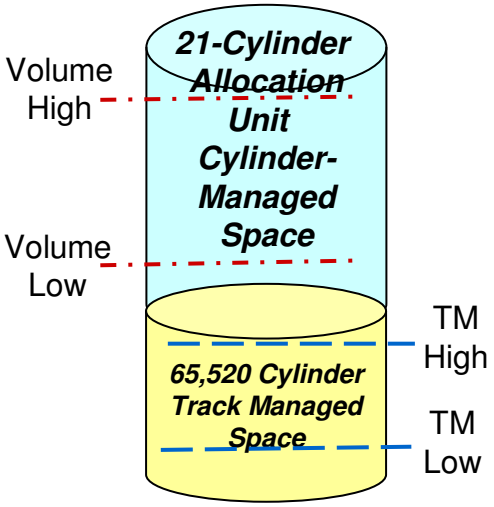

General: V1R12
DFSMSdss Cross Memory

More granular control of which functions use DFSMSdss Cross Memory

- All functions can be turned on or off using the existing

SETSYS DSSXMMODE(Y|N)…

or individual functions can be specified using

SETSYS DSSXMMODE(BACKUP(Y|N) CDSBACKUP(Y|N) DUMP(Y|N) MIGRATION(Y|N) RECOVERY(Y|N))

- Must be specified in ARCCMDxx member

Agenda

- **General**
  - NonDisruptive CDS Backup
  - ONLYIF
  - Usability
  - CA Reclaim
  - Deduplication
  - EAV Support
  - Cross Memory
- **Migration**
  - On Demand Migration
  - Performance Improvements
- **Backup**
  - BD Exit
  - Dump Block Size
  - Dump Stacking Limit
  - Multitask Recovery from Dump Tapes
- **Fast Replication**
  - Copy Pool Recovery from Dump Tape
  - Fast Reverse Restore
  - Space Efficient
  - Preserve Mirror
  - UCB Refresh
- **Miscellaneous**

General: V1R13
On Demand Migration

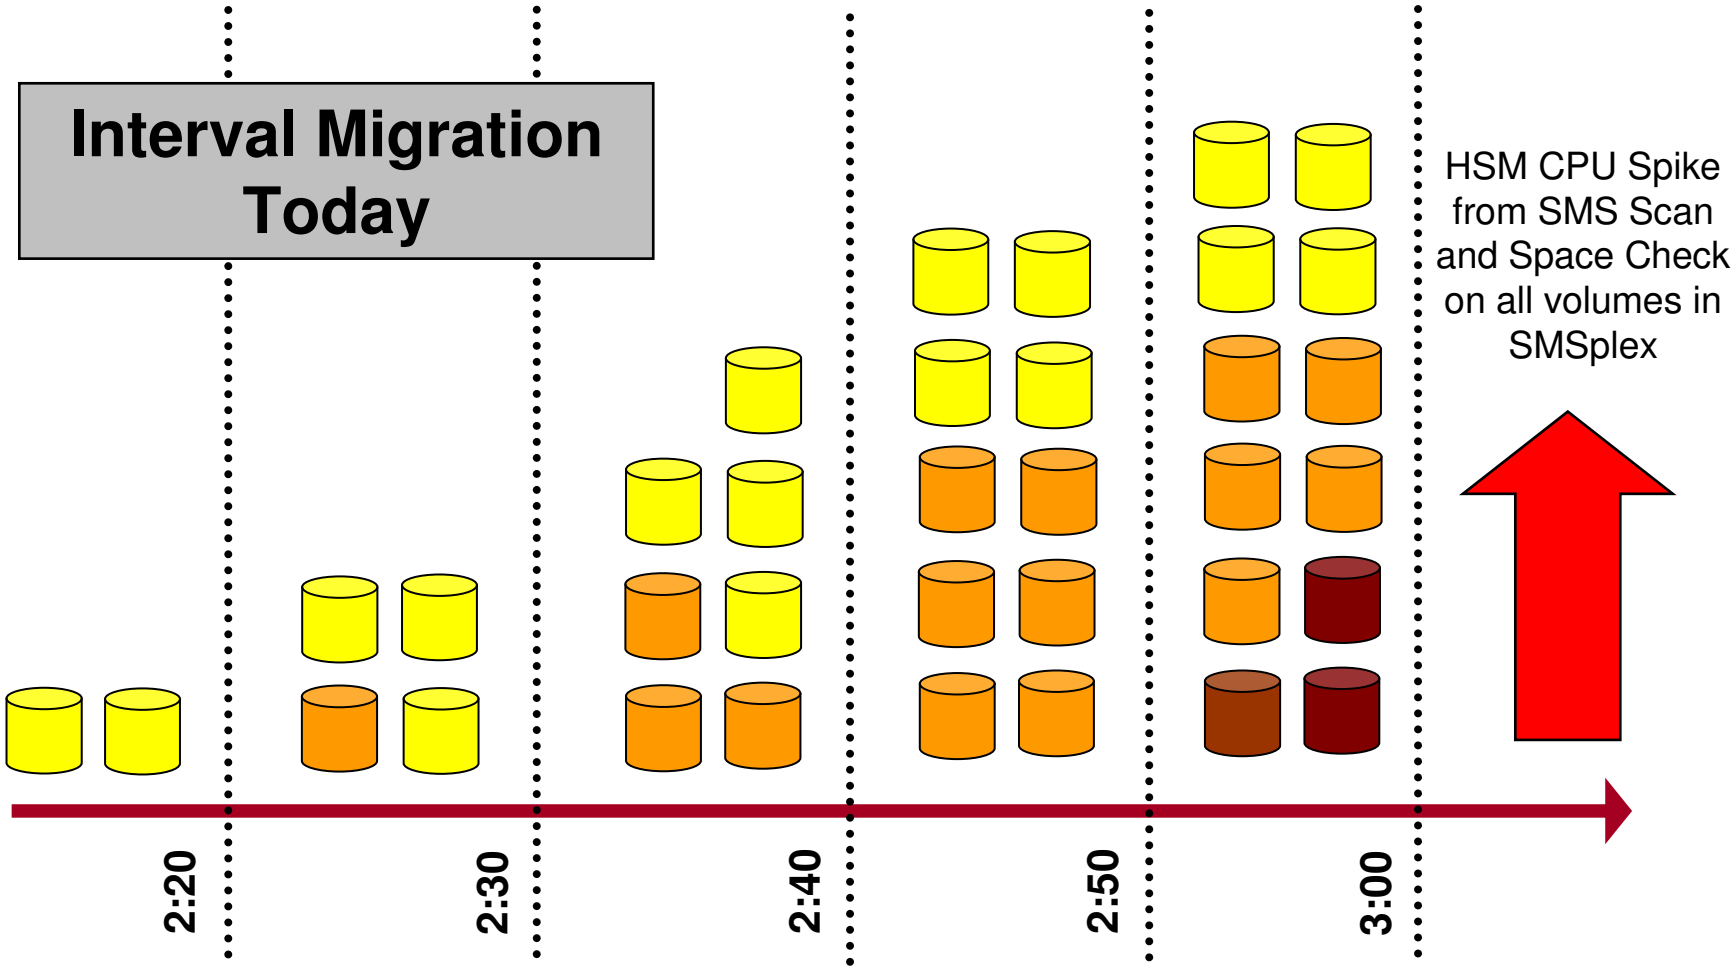

General: V1R13
On Demand Migration

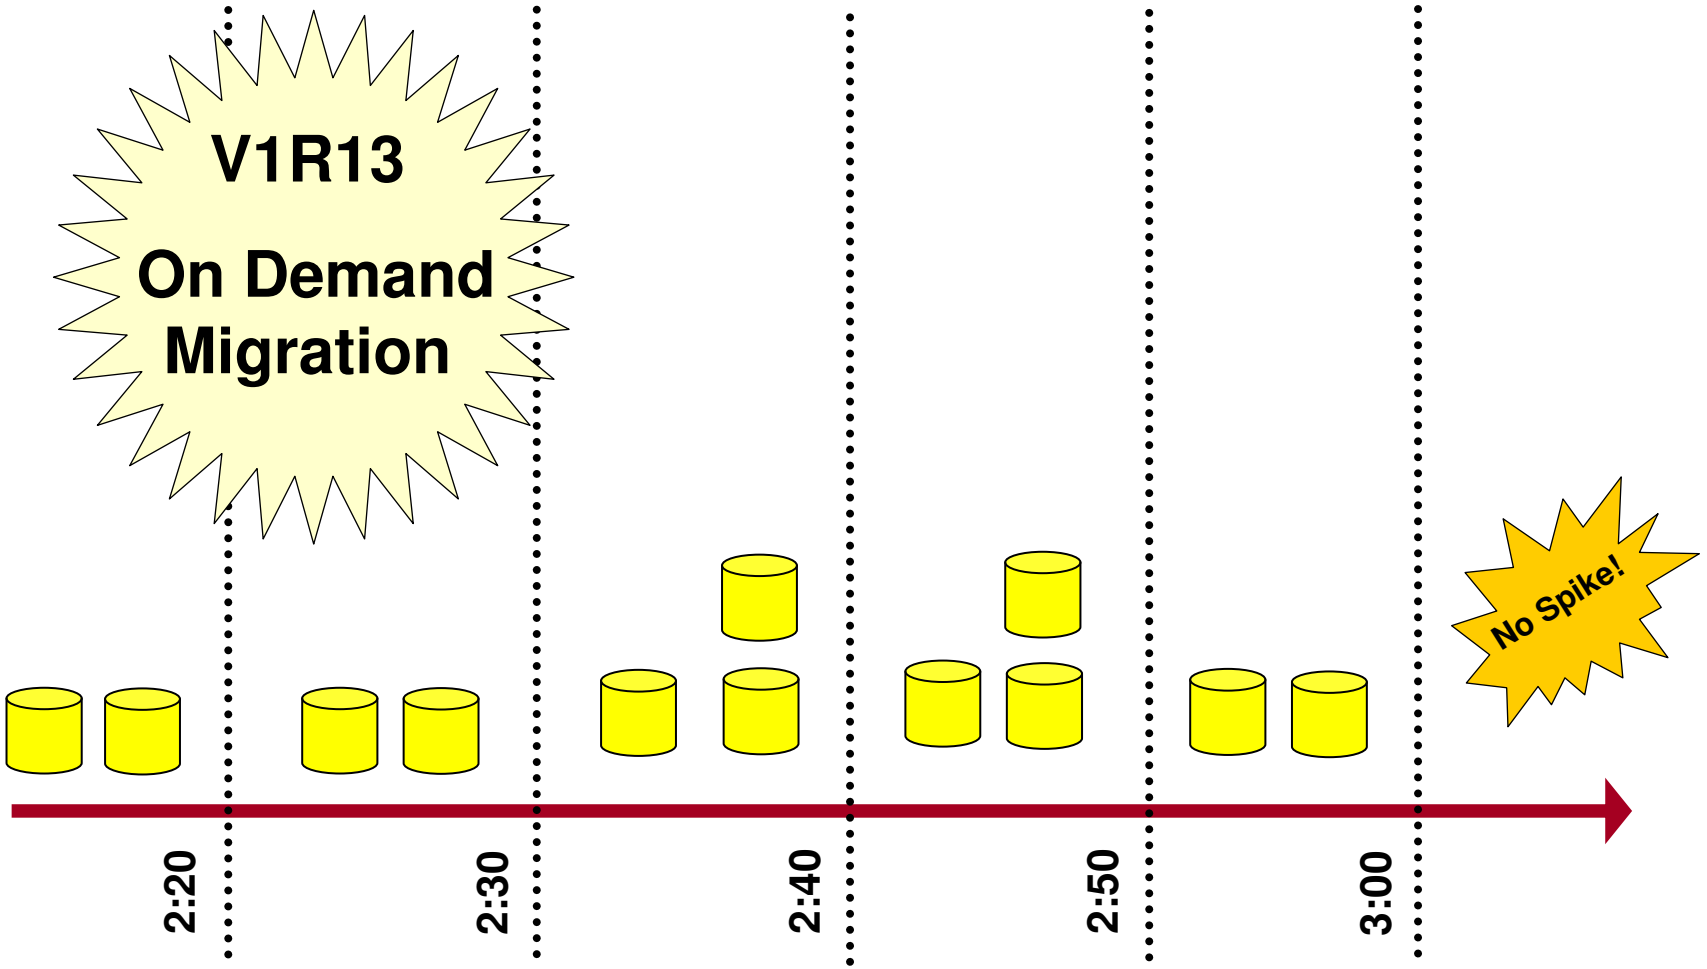

General: V1R13
On Demand Migration

- New SMS ENF 72
  - When an allocation (new data set or extent) causes a volume to go over its high threshold, SMS issues ENF 72
    - For EAVs, issued for both track-managed and cylinder-managed space
  - Additional ENFs issued at each 25% increment of (100% - High threshold)
    - If the high threshold is 80%, then an ENF is issued at 80%, 85%, 90%, …
    - If the high threshold is 92%, then an ENF is issued at 92%, 94%, 96%, …
    - ENF issued for every allocation after threshold exceeds 97% increment of high threshold
  - SMSplex wide ENF

General: V1R13
On Demand Migration

- DFSMShsm SETSYS command to enable function

  SETSYS ONDEMANDMIGRATION (Y|N)

  - Y- Enables ODM
    - ODM performed for volumes in storage group with AM=Y
    - SETSYS INTERVALMIGRATION is ignored for SMS volumes
      - *Except, storage groups with AM=I*
    - Interval Migration will still run for nonSMS volumes
  - If you use AM=I, consider changing to ODM and AM=Y
    - **AM=I**: Volume eligible when space exceeds midpoint between low and high threshold (used for TMM)
    - **AM=Y**: Volume eligible when space exceeds high threshold

General: V1R13
On Demand Migration

- When ODM is enabled, DFSMShsm host will listen for ENF 72
  - When ENF 72 is received, volume is added to work queue
  - Multiple hosts should be enabled. Only one host will process the volume.
    - Don't setup a single point of failure (single host enabled)
- Standard volume-level space management is performed
  - Volume processed below low threshold or until no more eligible data sets
  - If no data sets are processed, then there are no eligible data sets
    - To prevent the volume from being repetitively processed, the volume will not be reselected for 24 hours

      PATCH .MGCB.+138 X '00015180' /* 24 hours in seconds – default value*/
- In order to catch exception cases where a large number of volumes are concurrently going over threshold, DFSMShsm issues a highlighted message to the console

  ARC1901E NUMBER OF VOLUMES ELIGIBLE FOR ON DEMAND MIGRATION HAS REACHED nnnnn

  - Updated as number of volumes increases/decreases. Removed from console after the number of volumes drops below the specified number

  SETSYS ODMNOTIFICATIONLIMIT(nnnnn)
  (100 is default)

General: V1R13
Space Management

- **PreV1R13**
  - During every interval migration window and before every Primary Space Management window, DFSMShsm scanned every volume in SMSplex to look for changes
- **V1R13**
  - DFSMShsm now listens for ENF 15 – SMS configuration change
  - Scan is only done if there was a configuration change
    - If you use this patch to turn off the scan, you can remove it:

      PATCH .MCVT.+C8 BITS(1…….)

    - If running ODM, you can remove any patches for changing the frequency of running Interval Migration
      - There are several. They are documented in the *DFSMShsm Implementation and Customization Guide*

Migration: V1R12
Performance Improvement

To improve performance, two phases of volume migration process are overlapped

- **Phase 1**
  - Obtain a list of all data sets on the volume
  - Process the data set list
    - Expire
    - Partial Release
    - Reconnect
    - Generate migration/extent reduction Queue
- **Phase 2**
  - Process Queue

- **New task started for Phase 2**
- **Phase 1 task continues to the next volume**

Migration: V1R12 & 13

Performance Improvement *(cont)*

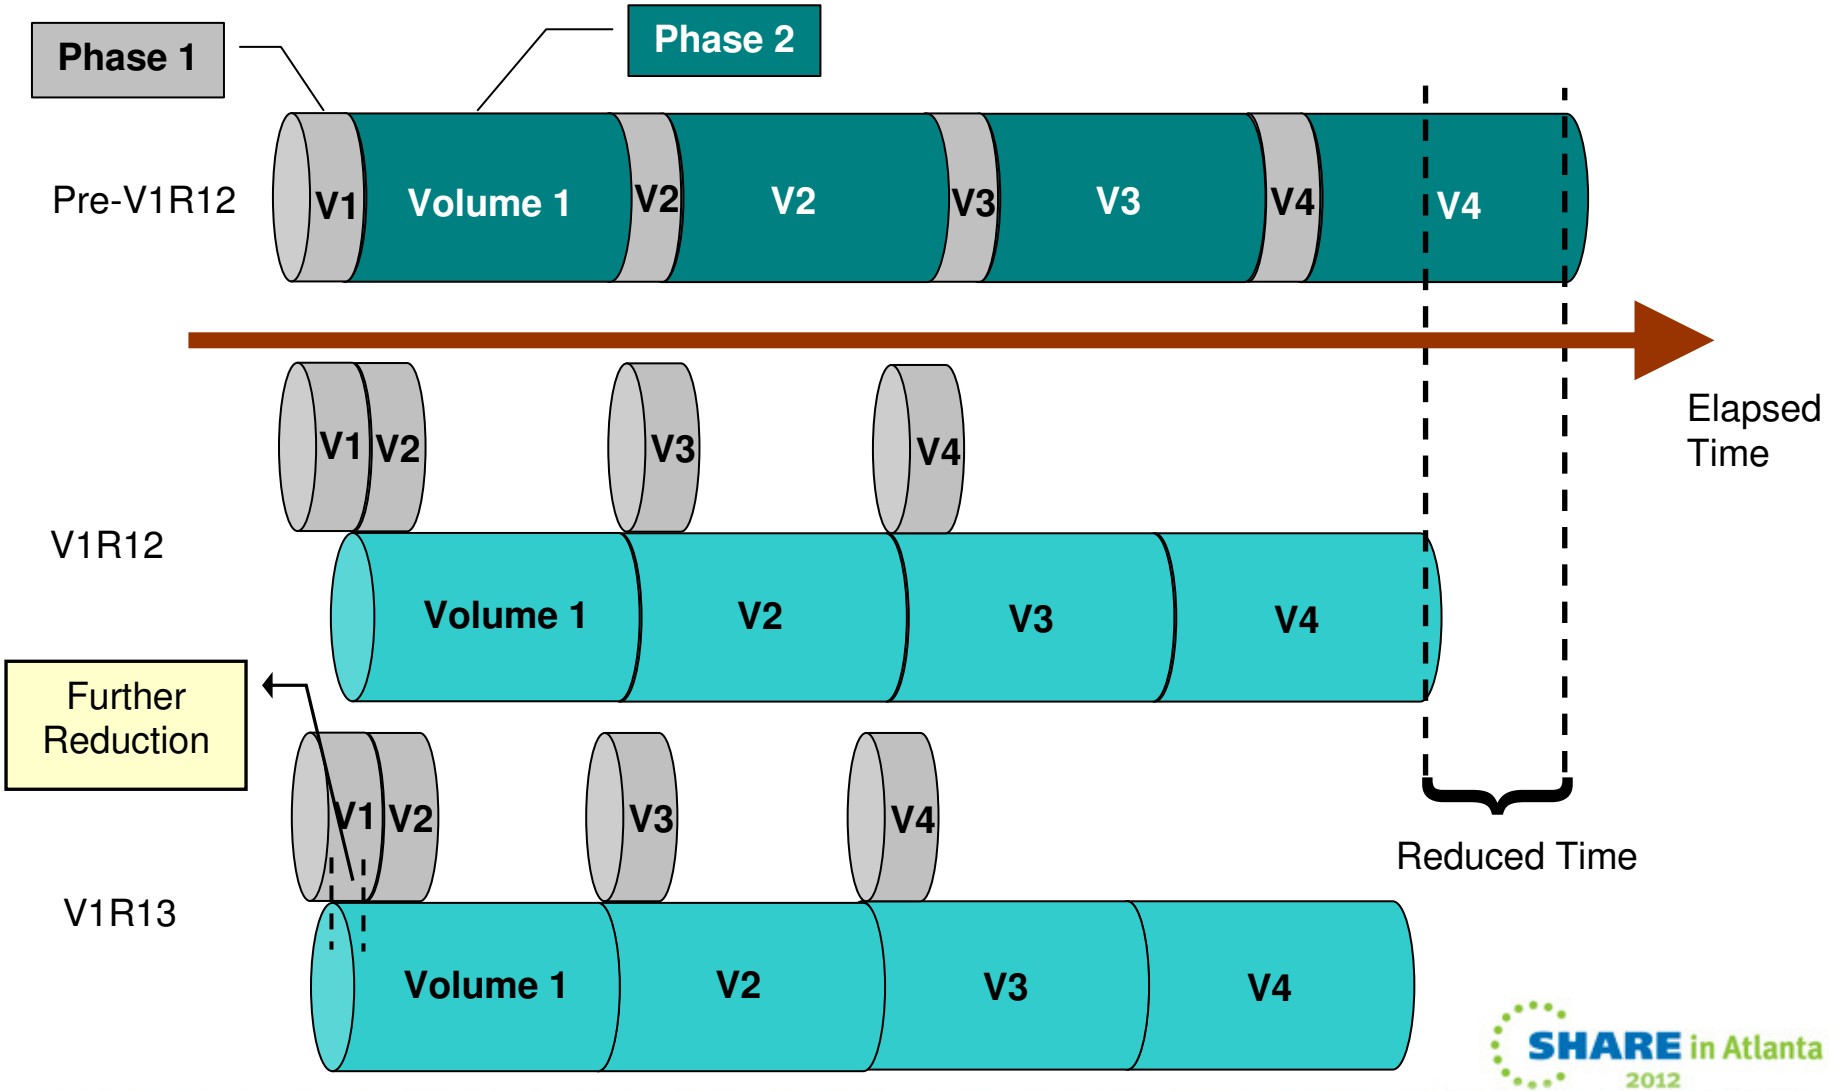

Migration: V1R12

Performance Improvement *(cont)*

Supported for Volume Level Space Management

- Primary Space Management
- On Demand Migration
- Interval Migration
- MIGRATE PRIMARY command (nonSMS volumes)

Agenda

- General
  - NonDisruptive CDS Backup
  - ONLYIF
  - Usability
  - CA Reclaim
  - Deduplication
  - EAV Support
  - Cross Memory
- Migration
  - On Demand Migration
  - Performance Improvements
- Backup
  - BD Exit
  - Dump Block Size
  - Dump Stacking Limit
  - Multitask Recovery from Dump Tapes
- Fast Replication
  - Copy Pool Recovery from Dump Tape
  - Fast Reverse Restore
  - Space Efficient
  - Preserve Mirror
  - UCB Refresh
- Miscellaneous

Backup: V1R11

OA38152 – RETAINDAYS in BD Exit

Today: BACKDS RETAINDAYS enables a specific retention period to be specified on a data set backup command

- Minimum number of days to retain the backup copy
- RETAINDAYS(99999) indicates 'Never Expire'
- Not supported for Automatic or Volume Backup

OA38152: Enables a RETAINDAYS value to be specified in the BD (Backup Dataset) Exit

- Greatly extends the ability to specify a retention period for a data set
  - Dataset backup, Volume backup, Automatic backup

Backup: V1R12

Full Volume Dump – Optimal Block Size

DFSMSdss enhanced DUMP and RESTORE to use the optimal block size for the output tape device

- Up to 256KB for newer devices

DFSMShsm BACKVOL/FRBACKUP DUMP, AUTOMATIC DUMP and RECOVER/FRRECOV FROMDUMP will take advantage of this enhancement

- If you use these functions, examine your setting for SETSYS DUMPIO(*n,m*)
- *n* indicates the number of tracks that DFSMSdss reads at a time for DUMP.
  - 1 is the default for 1 track
  - 4 specifies that a cylinder should be read – requires much more storage

Backup: V1R12

Multitask Recovery From Dump Tape

Full Volume Recovery from Dump Tape enhanced to support multiple concurrent tasks

SETSYS MAXDUMPRECOVERTASKS(*1*-*64*)

★ Multiple concurrently queued volume recovery requests on the same tape will be processed with a single tape mount

Stacking

• Use a stacking value that will enable you to maximize the parallelism of the recovery

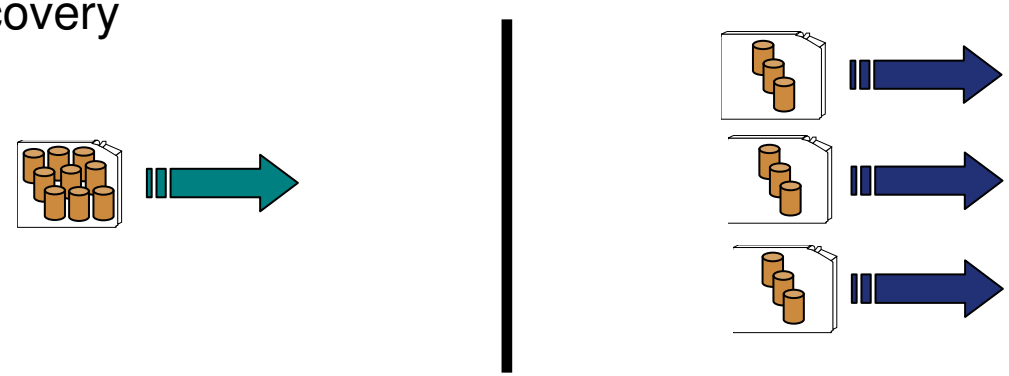

Backup: V1R12
Dump Stacking

Dump Stacking limit increased from 99 to 255

DEFINE DUMPCLASS(*class* STACK(***nnn***))

BACKVOL SGROUP(*sgroup*) DUMP(DCLASS(*class*) STACK(***nnn***))

Coexistence

- Increased values can only be specified on V1R12 and higher
- Lower level releases will honor a Dump Class stack value > 99

Agenda

- General
  - NonDisruptive CDS Backup
  - ONLYIF
  - Usability
  - CA Reclaim
  - Deduplication
  - EAV Support
  - Cross Memory
- Migration
  - On Demand Migration
  - Performance Improvements
- Backup
  - BD Exit
  - Dump Block Size
  - Dump Stacking Limit
  - Multitask Recovery from Dump Tapes
- Fast Replication
  - Copy Pool Recovery from Dump Tape
  - Fast Reverse Restore
  - Space Efficient
  - Preserve Mirror
  - UCB Refresh
- Miscellaneous

Fast Replication Overview

HSM function that manages Point-in-Time copies

• Combined with DB2 BACKUP SYSTEM, provides non-disruptive backup and recovery to any point in time for DB2 databases and subsystems (SAP)

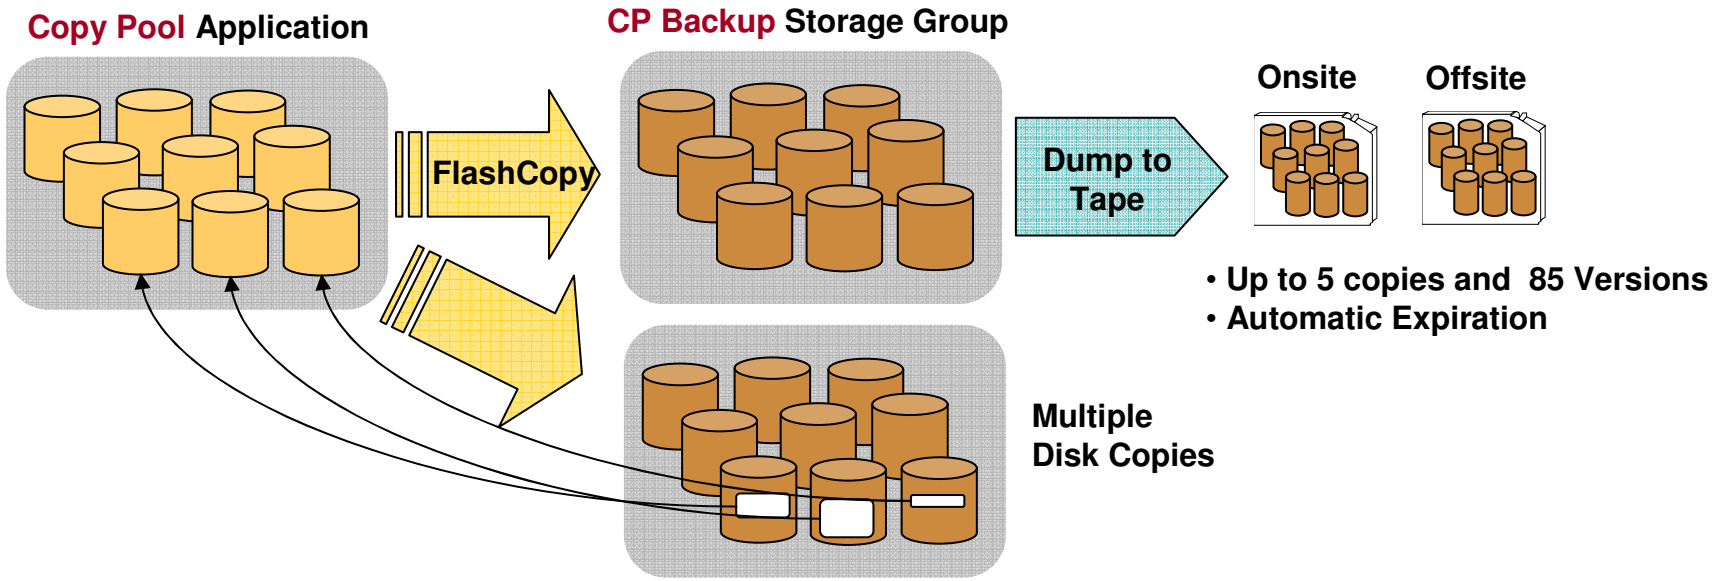

★**Recovery at all levels from either disk or tape!**

- Entire copy pool, individual volumes and …
- Individual data sets

Fast Replication: V1R12
Copy Pool Recovery from Tape

Copy Pool recovery from tape

- Prior to V1R12, HSM could only recover an entire copy pool from disk
- DB2 or native DSS had to be used to recover from tape

With this enhancement…

- DFSMShsm supports recovery of entire copy pool from tape
  - Made possible with the multitask recovery from dump tape support
  - Fast Replication tape copies are full-volume dump copies that represent the time that FlashCopy disk copy was created
- Enables solution to be used for Disaster Recovery
- Once again, stacking level of dump tapes may limit recovery parallelism

Fast Replication: V1R12
Copy Pool Recovery from Tape

Retry Logic

- If there is an error recovering one or more volumes…

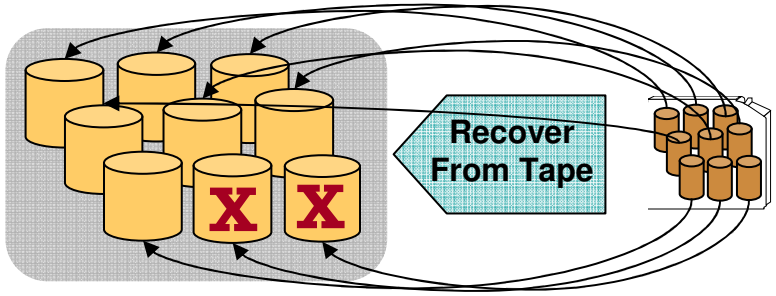

- Correct the problem
- Reissue the FRRECOV command

★ DFSMShsm will only process those volumes that previously failed

★LIST COPYPOOL output indicates if a particular version recovery is resumable

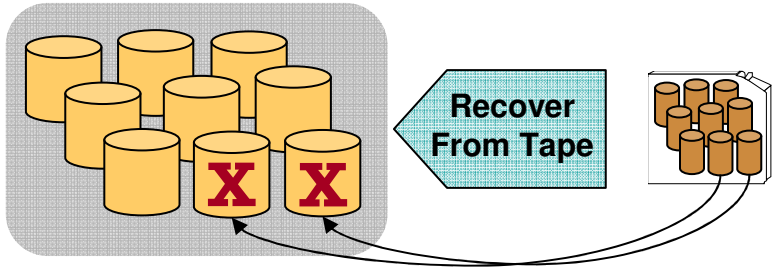

RESUME(NO) option to retry all volumes

(YES) is the default

Fast Replication: V1R12
Copy Pool Recovery from Tape

QUERY ACTIVE indicates the progress of the recovery…

ARC1822I FRRECOV OF COPY POOL *cpname* FOR USER *userid*, REQUEST *request-number* ON HOST *host_id* IS IN PROGRESS: NOT PROCESSED = *xx*, TOTAL = *yy*

- TOTAL indicates the total number of volumes being recovered
- NOT PROCESSED indicates the number of volumes that have not yet been recovered

Helpful messages

- ARC1802 will list overall return code when all volumes complete
- ARC1803 will summarize failed volumes with DFSMSdss message ids

CANCEL REQUEST(*reqnum*)

- All volumes being recovered will have the same request number
- A single CANCEL command will cancel all volumes

Fast Replication: V1R12
Fast Reverse Restore

Prior to V1R12

- Disk recovery could not be performed until the physical background copy was complete (*hours* after the logical backup was created)
- Disk recovery from a NOCOPY version was not supported

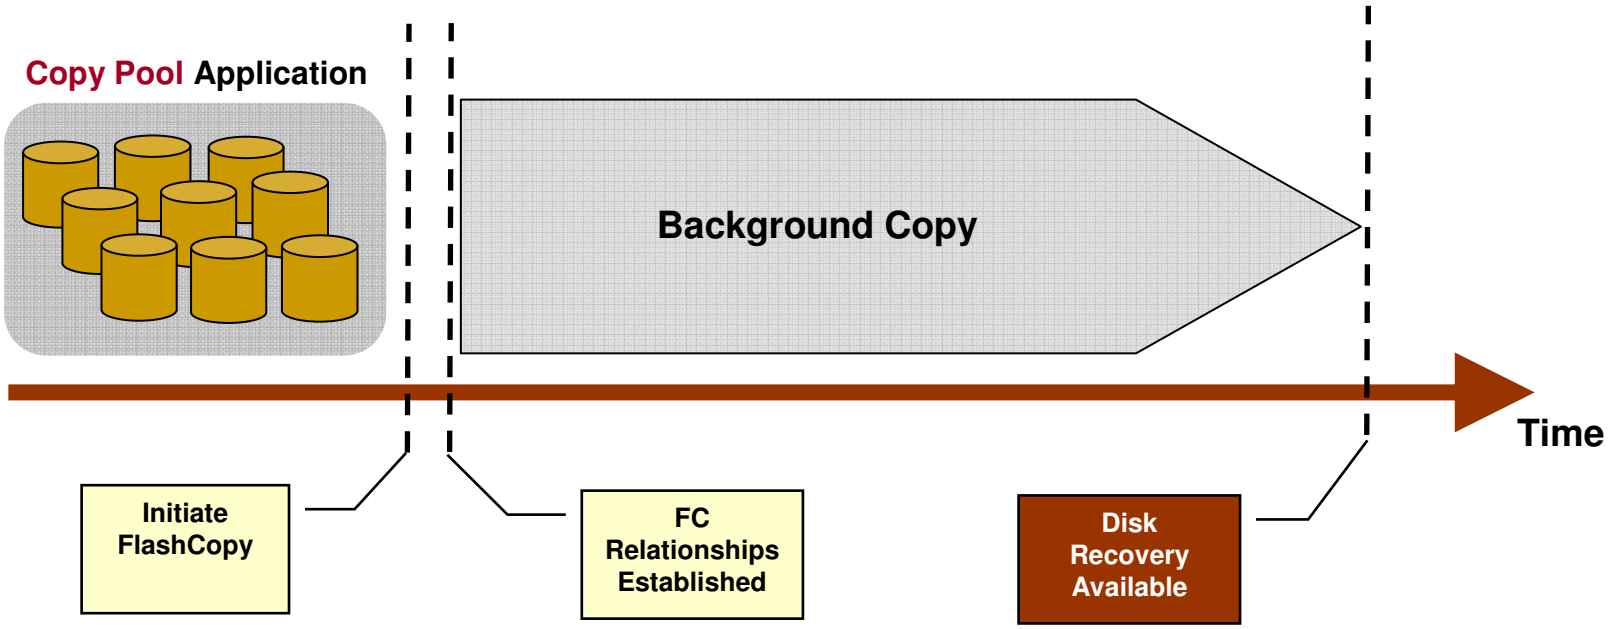

Fast Replication: V1R12
Fast Reverse Restore

Fast Reverse Restore

- ★ Enables DFSMShsm to FlashBack for recovery even though the background copy has not completed
- ★ Enables recovery from NOCOPY versions

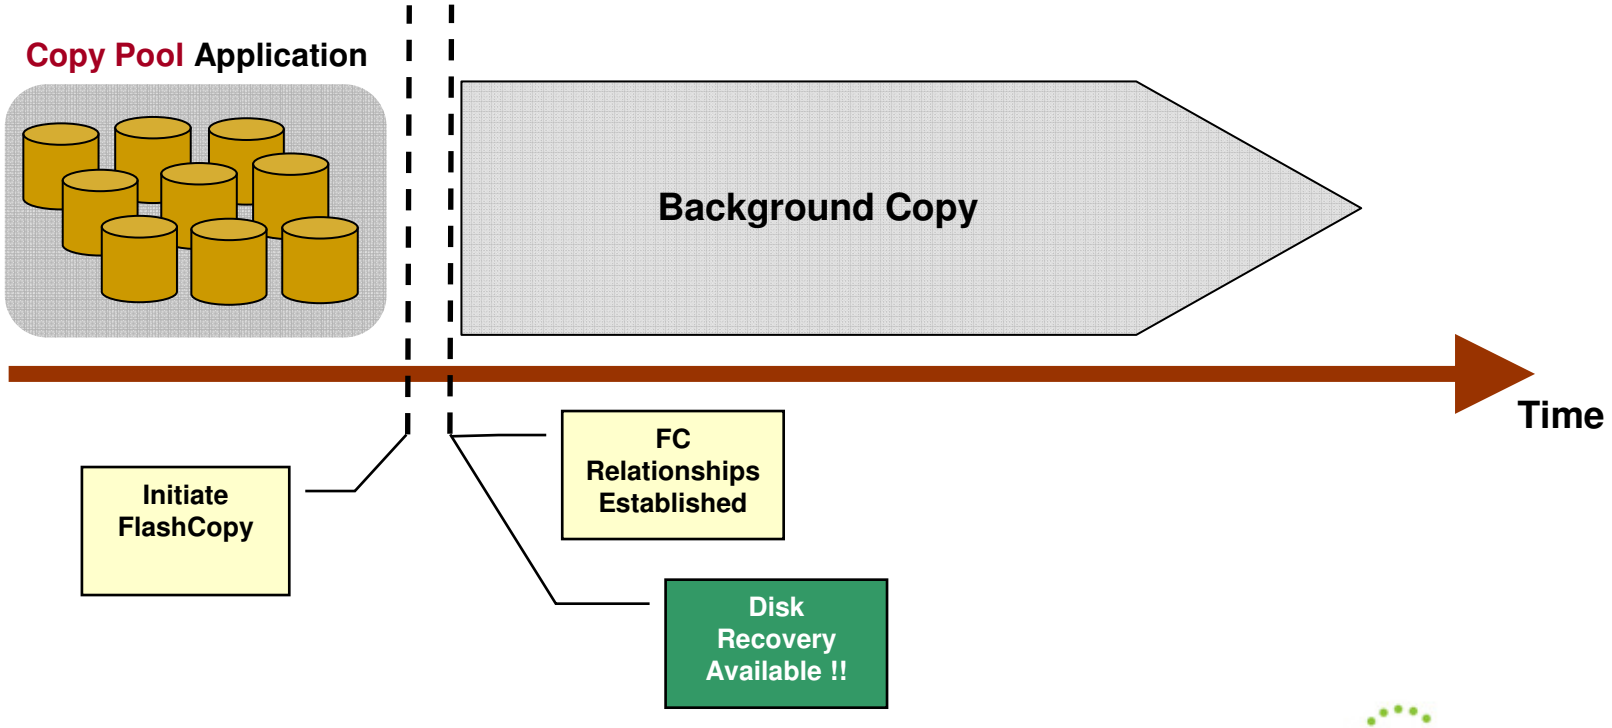

Fast Replication: V1R12

Fast Reverse Restore

Restrictions

! Since the background copy does not complete, the backup is invalidated as part of the recovery

! Source volumes can't be in any other active relationships

! Only supported for recovery of entire copy pool

Resume after failure

• If one or more volumes fail, similar to recovery from tape, just reissue the command and only the set of previously failed volumes is retried

Fast Replication: V1R12

Fast Reverse Restore

- New copy pool setting indicates if FRR is enabled for the copy pool
  - YES indicates that it is acceptable to recover a version before the background copy is complete and for the backup to become invalidated
- QUERY COPYPOOL indicates percent complete
  - Determine if you should just wait for background copy to complete

```
ARC1820I THE FOLLOWING VOLUMES IN COPY POOL CP1, VERSION 003,
         HAVE AN ACTIVE FLASHCOPY BACKGROUND COPY
ARC1820I (CONT.) SGNAME   FR-PRIMARY FR-BACKUP PCT-COMP
ARC1820I (CONT.) SGRP1    SRC01B     TGT01B     70
ARC1820I (CONT.) SGRP1    SRC02B     TGT02B     80
```

- FORCE keyword if the disk copy is currently being dumped to tape
  - Incomplete dump volumes will be discarded

Fast Replication: V1R12

Fast Reverse Restore

- Hardware Dependencies
  - DS8000 series storage servers with the following **minimum** LMC:
    - 5.3.1.450
    - 5.4.21.540
    - 5.4.30.253
    - 6.5.0.220
- Software Dependencies
  - PK99337 ICKDSF – Support DFSMShsm invocation of INIT

Fast Replication: OA30816
Space Efficient Volumes

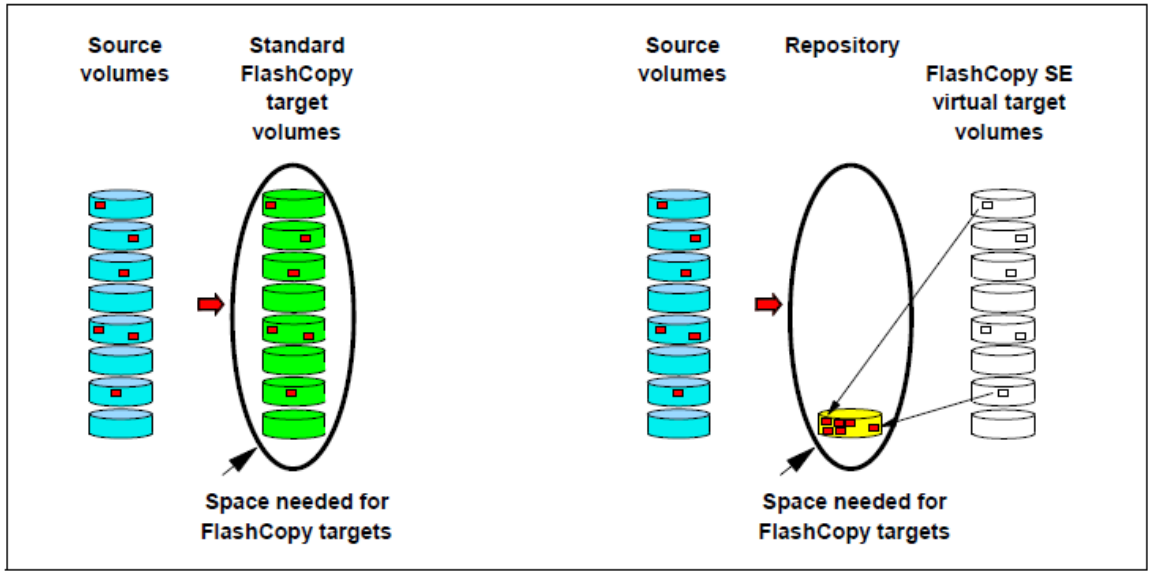

Space Efficient Volumes

- Space Efficient volumes don't take physical space until the space is actually needed
- Only a fraction of the space is required for target volumes
  - ✓ License is required
- Valid when **NOCOPY** is the FlashCopy technique

Fast Replication: OA30816
Space Efficient FlashCopy

HSM selects available Space Efficient target volumes when NOCOPY (VERSIONS=0) is selected

(V1R9 and higher)

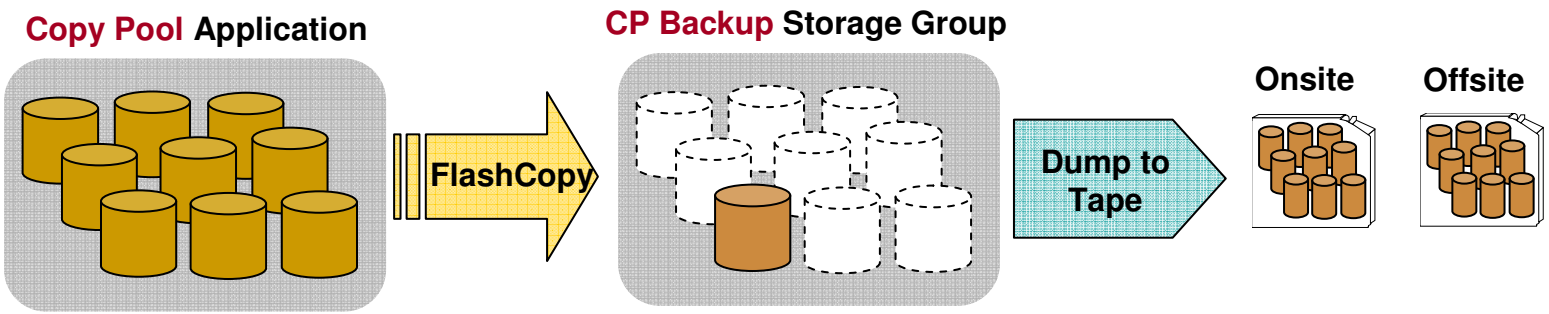

Fully Provisioned

- 1 TByte copy pool needs 1 TByte of copy space

Space Efficient Target Volumes

- ★1 TByte copy pool that changes <10% in-between copies only needs 100 GBytes of target space

Fast Replication: V1R12
Space Efficient FlashCopy

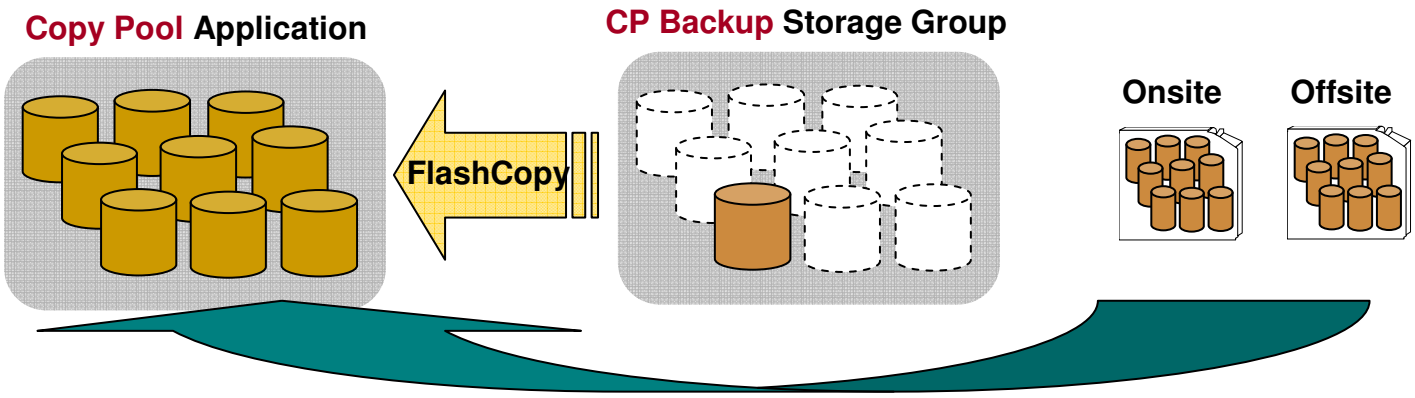

With Fast Reverse Restore…

- In addition to being able to recover from tape, a disk recovery can be performed!

Fast Replication: New Function APARs
FlashCopy to PPRC Primary

- Without new support, DFSMShsm does not support FlashCopy to mirrored volumes
- This means that you have to break pairs before each FlashCopy and re-establish pairs after FlashCopy

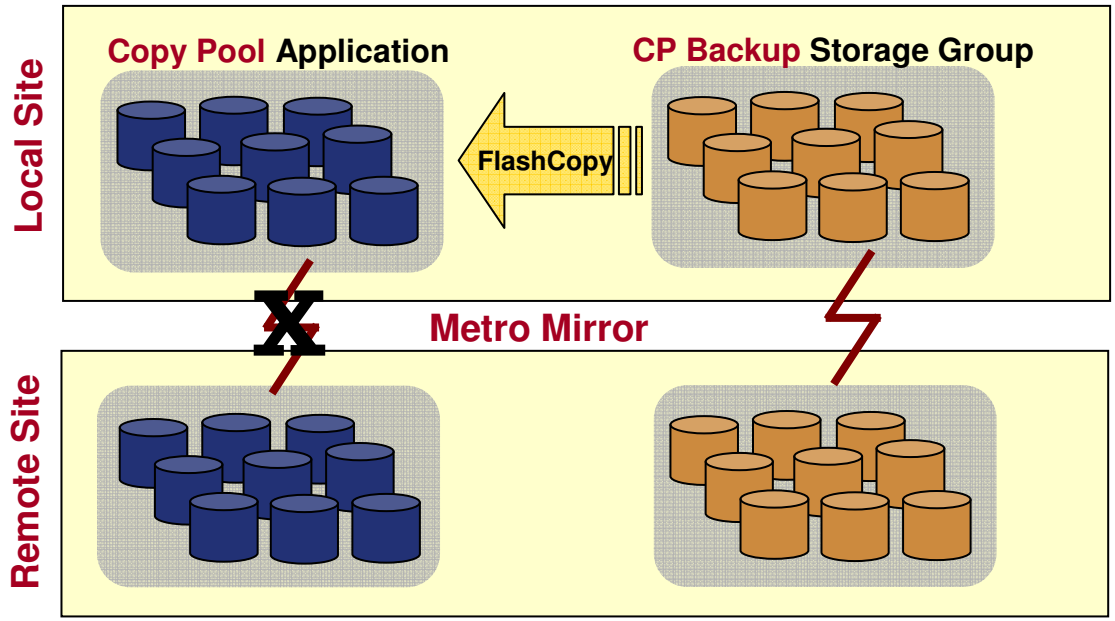

Fast Replication: New Function APARs
FlashCopy to PPRC Primary *(Cont)*

OA23849 (V1R9) – Enable FlashCopy to a PPRC Primary Volume

- Enables easier implementation of backup volumes at remote site *(Metro Mirror Only)*
- Pair is put into a *'Duplex Pending'* state until volumes re-synched
- Enabled for Recovery, but you generally don't want production volumes to go in a 'Duplex Pending' state

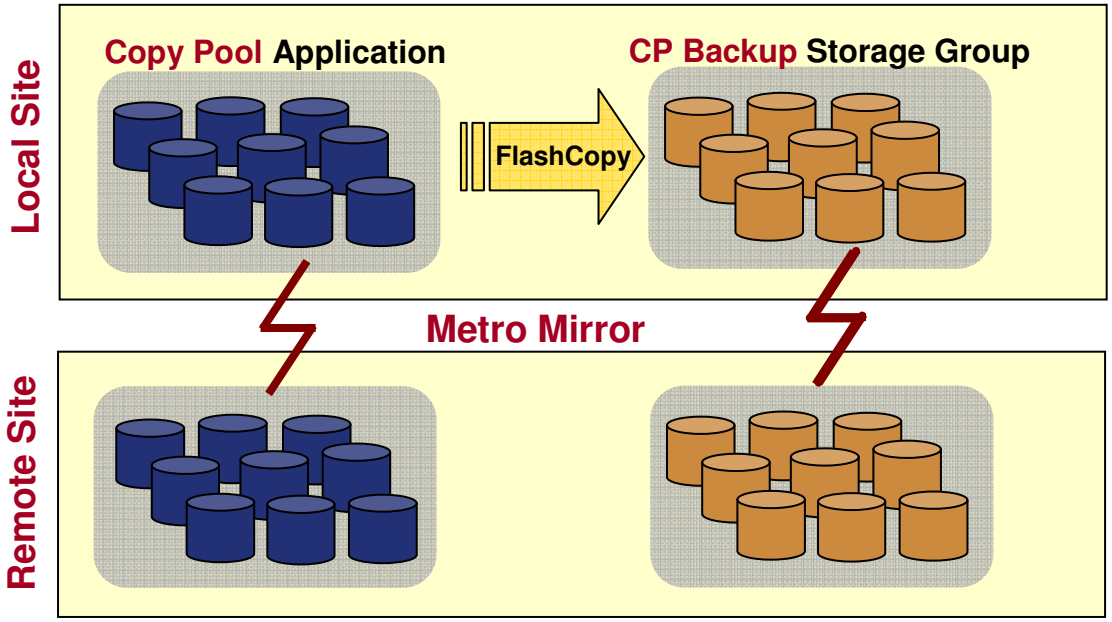

Fast Replication: New Function APARs
FlashCopy to PPRC Primary *(Cont)*

OA24814 (V1R9) – Preserve Mirror Support SPE

- "Remote Pair FlashCopy" DS8000 Support
  - Technical Information section of announcement letter ENUS109-119, located at http://www-01.ibm.com/common/ssi
- FlashCopy to PPRC Primary volume *while maintaining Full Duplex*

! *Fast Reverse Restore and Space Efficient cannot be used in this environment*

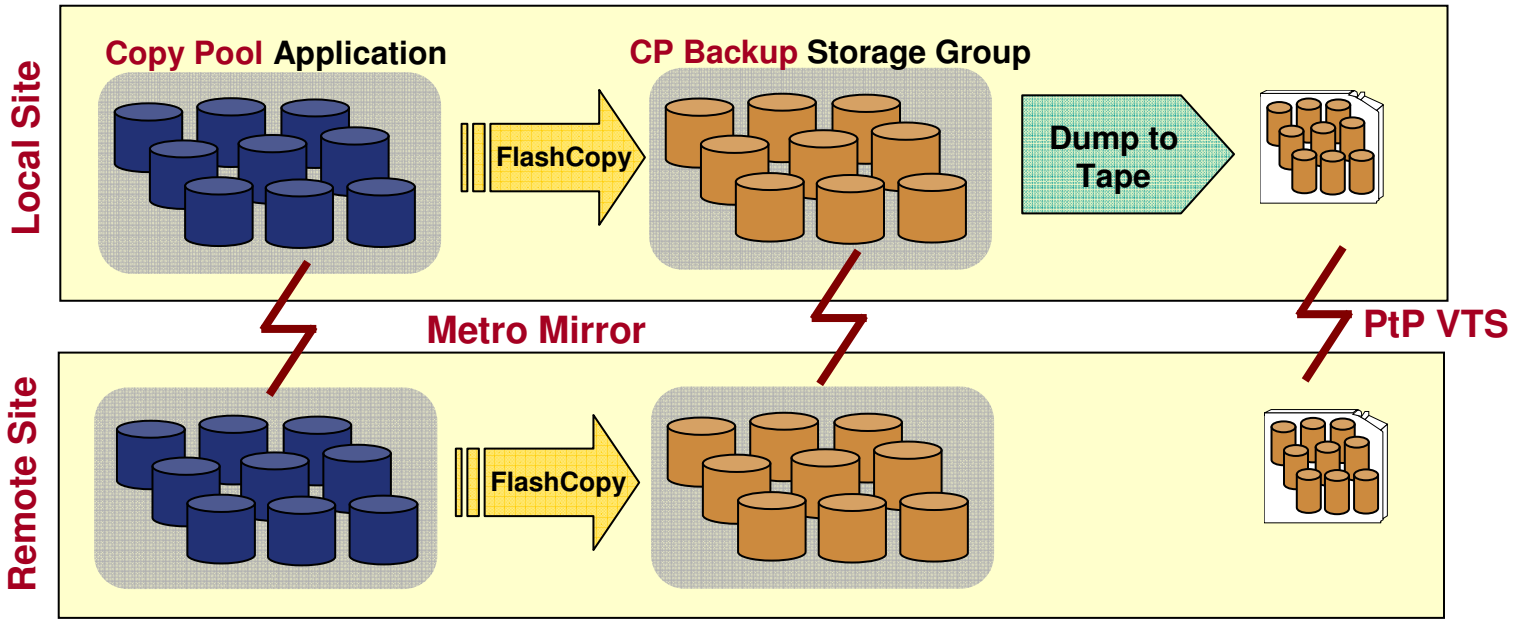

Fast Replication: OA32494
Volume Preferencing

DFSMShsm will select target volumes in same cluster

- FlashCopy background copy is more efficient when the source and target are within the same DS8000 cluster
- When selecting target volumes, DFSMShsm will prefer target volumes in the same cluster as the source

- Pre-V1R12, enable with PATCH .FRGCB.+A BITS(....1...)
- V1R12+, enabled by default

Fast Replication: V1R13

Usability

SETSYS replaces patch for volume pairing messages

SETSYS FASTREPLICATION(VOLUMEPAIRMESSAGES(YES|NO))
SETSYS FR(VPM(Y|N))

- Controls issuing ARC1809I message for volume pairing failures
- Replaces PATCH .FRGCB.+9 BITS(.1……)

Reduced messaging

- ARC1809I RC2 used to be issued for every source volume that attempted to be paired to a target volume that was already paired
- Now the message is only issued once per target volume for each storage group in the copy pool

Fast Replication: V1R13 Usability

FASTREPLICATION(DataSetRecovery) default changed to NONE

- When fast replication is used for data set recovery, a subsequent FRBACKUP cannot be performed until the background copy of the recovery completes
  - FRBACKUP will fail, but the reason will just say that there is an existing relationship
  - When standard I/O is used for the recovery, HSM will fail the FRBACKUP with a specific message indicating that another function hasn't completed

AUDIT COPYPOOLCONTROLS detects and fixes orphaned records

- 'Orphaned' target volume records (FRTV) will prevent a target volume from being selected
- AUDIT will now identify these records and optionally delete them

Fast Replication: V1R13
UCB Refresh

Pre-V1R13

- If the FlashCopy moves the VTOC (source VTOC is in a different location than the target VTOC), then the target volume has to be varied offline/online before it can be accessed on systems other than the one that did the FlashCopy

V1R13

- DSS builds an ENF64 that can be used to automatically refresh the UCB on all systems in sysplex
- To enable the function
  - PARMLIB member DEVSUPxx: ENABLE(REFUCB)

    or
  - MODIFY DEVMAN …
- To disable the function
  - DISABLE(REFUCB)

Agenda

- **General**
  - NonDisruptive CDS Backup
  - ONLYIF
  - Usability
  - CA Reclaim
  - Deduplication
  - EAV Support
  - Cross Memory
- **Migration**
  - On Demand Migration
  - Performance Improvements
- **Backup**
  - BD Exit
  - Dump Block Size
  - Dump Stacking Limit
  - Multitask Recovery from Dump Tapes
- **Fast Replication**
  - Copy Pool Recovery from Dump Tape
  - Fast Reverse Restore
  - Space Efficient
  - Preserve Mirror
  - UCB Refresh
- **Miscellaneous**

Miscellaneous

(V1R11) IEFBR14 DELETE will not recall a migrated data set before it is deleted

- Change made to Allocation
  - SETALLOC SYSTEM,IEFBR14_DELMIGDS=NORECALL

  or
  - ALLOCxx keyword: IEFBR14_DELMIGDS(NORECALL)

(V1R12) DELETE GDG FORCE will not recall a migrated GDS before it is deleted

- Change made to Catalog

Miscellaneous

(V1R12) GDG Serializaton

- For existing Migration, GDG base is serialized when migrating a generation
- Serialization scheme is changed to not require the GDG base to be serialized
    - Apply OA36664

(V1R12) 100% Threshold

- SMS allows value of 100 for High Threshold
- Primary Space Management and Interval Migration are unchanged
    - DFSMShsm rounds up values, so 99.1% and higher are considered 100%

(V1R12) WTOR support

- Based on z/OS support, default responses have been established for DFSMShsm WTOR messages

Miscellaneous

(V1R13) Migration and Backup Dataset Names

- Today, migration/backup copies are named:

```
prefix.HMIG.Tssmmhh.user1.user2.Xyddd
prefix.BACK.Tssmmhh.user1.user2.Xyddd
```

- When duplicate names are generated, '**T**' is incremented to the next letter until a unique name is generated
- In a busy environment with many like-named datasets being processed, this can result in dozens of extra CDS I/Os per data set
  - Examine your backup/migration data set names to determine how many non-'T' names you have
- **OA38729**: Changes time field to use hundreds of seconds to make name more unique

```
prefix.HMIG.Tthssmh.user1.user2.Xyddd
prefix.HMIG.Tthssmh.user1.user2.Xyddd
```

- '**mh**' is an alphanumeric conversion of '`mmhh`'

Summary

- Significant number of new and enhanced functions
  - Emphasis on addressing customer requirements, increasing throughput and reducing MIPS
- Some critical new functions provided as “Development APARS”
- Expanded functionality of DFSMShsm Fast Replication Support
  - These items demonstrate IBM’s commitment to continual improvement of the DFSMShsm product

Session 10950