

Agenda

- **What's New in IMS Database Recovery Control (DBRC)**
 - Support for Image Copy 2 and Fuzzy User Image Copies
 - DBRC Timestamp Precision
 - RECON READONLY Access
 - Improved SAF Support for RECONs
 - DBRC API Enhancements
 - Parallel RECON Access
 - Latest DBRC SPEs

- **What's New in Database Recovery Facility (DRF)**

Image Copy Enhancements

- **Image Copy 2 enhancements**
 - Support for fast replication
 - Control statements for the SET PATCH interface
- **User image copy enhancements**
 - DBRC support for fuzzy user image copies

Image Copy 2 Fast Replication

- Uses DFSMSdss COPY command with FASTREP(REQ) parameter
 - Invokes FlashCopy on ESS, DS8000
 - Invokes SnapShot on RVA
- Copy is done in one phase
 - Time is comparable to the logical copy time for concurrent copy
- Copies database data sets to other data sets on the same storage system
 - Output is in same format as database data set
- Supports both fuzzy and clean image copies
- DD statements for the output data sets are not required
- Support in Database Recovery utility for these image copies
 - Restores the data set with a DFSMSdss COPY command with FASTREP(PREF) parameter
- GENJCL.IC and GENJCL.RECOV commands support provided

Image Copy 2 SET PATCH Support

- DFSMSDss has SET PATCH capabilities
 - Changes default actions
 - Examples:
 - Change wait/retry values
 - Change actions when orphaned CAs are found
- Image Copy 2 has added control statements for SET PATCH
 - Previous IMS releases required module zap to use SET PATCH
 - GENJCL.IC includes support for generating the SET PATCH control statement
- Benefits
 - User friendly interface for changing DFSMSDss defaults

Fuzzy User Image Copy Support

- DBRC support for fuzzy user image copies
 - Fuzzy image copies taken by utility or tool without a DBRC interface
 - Pack dump, DFSMSdss DUMP or COPY not invoked by IC2, etc.
 - Support:
 - NOTIFY.UIC can specify a fuzzy user image copy
 - With BATCH and STOPTIME(time) parameters
 - GENJCL.RECOV can be used to generate recovery from logs after fuzzy user image copy has been restored
- Benefits
 - Integration of fuzzy user image copies into DBRC environment

DBRC Timestamp Precision

- **DBRC timestamps will be used to the microsecond**
 - Previously timestamps used to tenth of a second
 - Could lead to duplicate timestamps (log open, log close, allocation)
 - Increased precision not in effect until MINVERS('10.1') is specified
 - For compatibility with previous releases
 - Abbreviated timestamps still supported
 - Unspecified part of time will be padded with zeros
 - DBRC commands that use timestamps to locate specific records in RECON must use a timestamp with microsecond precision with MINVERS('10.1')

- **Benefits**
 - Avoids possible duplicate timestamps

DBRC Command and Message Output

- Output includes full precision timestamps

```
LIST.LOG OLDS(DFSOLP00) SSID(SYS3) TIMEFMT(L,O,P,4)
```

```
2009.114 18:31:24.167342 -07:00 LISTING OF RECON PAGE 0002
```

```
PRIOLD
```

```
SSID=SYS3 # DD ENTRIES=1  
EARLIEST CHECKPOINT = 2009.114 16:01:03.123456 -07:00
```

```
DDNAME=DFSOLP00 DSN=IMSTESTL.IMS01.OLDSP0  
START = 2009.114 16:00:59.123456 -07:00 FIRST DS LSN= 0000000000000001  
STOP = 2009.114 16:13:10.185501 -07:00 LAST DS LSN= 0000000000000042C  
LOCK SEQUENCE# = 000000000000  
STATUS=ARC COMPLT FEOV=NO AVAIL  
PRILOG TIME=2009.114 16:00:59.123456 -07:00 ARCHIVE JOB NAME=JT161314  
VERSION=10.1
```

```
DSP0180I NUMBER OF RECORDS LISTED IS 1  
DSP0203I COMMAND COMPLETED WITH CONDITION CODE 00  
DSP0220I COMMAND COMPLETION TIME 2009.114 18:31:27.167342 -07:00  
IMS VERSION 10 RELEASE 1 DATA BASE RECOVERY CONTROL PAGE 0003  
DSP0211I COMMAND PROCESSING COMPLETE  
DSP0211I HIGHEST CONDITION CODE = 00
```

Database Change Accumulation Utility

- DB0 and DB1 control statements
 - Purge time is expanded
 - Does not depend on MINVERS value
 - Number of DD names allowed reduced from 4 to 3

Version 9 example:

```

-----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6-----+-----7-----
//SYSIN      DD *
DB0DI32DB012009.175 12:00:00.1 -08:00DI320101 DI320102 DI320103 DI320104
DB0AB77DB012009.175 12:00:00.1 -08:00AB770101 AB770102
/*

```

Version 10 example:

```

-----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6-----+-----7-----
//SYSIN      DD *
DB0DI32DB012009.175 12:00:00.123456 -08:00DI320101 DI320102 DI320103
DB0DI32DB012009.175 12:00:00.123456 -08:00DI320104
DB0AB77DB012009.175 12:00:00.1 -08:00      AB770101 AB770102
/*

```

- ◆ GENJCL.CA has been updated to create V10 control statements

Database Recovery Utility

- S control statement
 - Timestamp recovery time is expanded
 - Does not depend on MINVERS value
 - Increased precision is not required

Version 9 example:

```

-----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6-----+-----7-----
//SYSIN      DD *
S  DI32DB01  DI320101          2009.175 12:00:00.1-08:00 C
/*
  
```

Version 10 examples:

```

-----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6-----+-----7-----
//SYSIN      DD *
S  DI32DB01  DI320101          2009.175 12:00:00.123456 -08:00 C
/*
  
```

```

//SYSIN      DD *
S  DI32DB01  DI320101          2009.175 12:00:00.1 -08:00      C
/*
  
```

- ◆ GENJCL.RECOV has been updated to create V10 control statements

READONLY Support for RECONs

- V10 READONLY support
 - Specification:
 - PARM(READONLY) on DSPURX00 EXEC statement
 - READONLY=YES on DBRC API FUNC=STARTDBRC macro
 - Use for users with only READ authority
 - Causes RECONs to be opened for input

- Benefit
 - Users need only READ authority to list RECON contents

Improved SAF Support for RECONs

- **IMS V10 SAF authority for RECONs**
 - READ is sufficient for readers
 - UPDATE is sufficient for accesses except DELETE and DEFINE
 - ALTER required for DELETE and DEFINE
 - CONTROL is never required

- **Previous IMS releases**
 - Required CONTROL authority for all RECON access

- **Benefits**
 - Users need only UPDATE authority for DBRC commands

DBRC API Enhancements

- RECON update capability via DBRC command support
 - INIT, CHANGE, and NOTIFY
- QUERY enhancements
 - Queries for DBDS, Partition, Log
 - Wildcard support
- Alternate RECON and IMS DD names
 - May be used to access multiple sets of RECONs and ACBLIBs easily
- Application may register as subsystem and authorize databases
 - Allows application to do utility functions with authorization integrity
- SAF(RACF) invocation for API security
 - Extension of DBRC command authorization
- Benefits
 - Complete API interface for users and IMS tools

Parallel RECON Access

- Allows multiple DBRC instances to access the RECONs concurrently
 - DBRC instance: IMS Online subsystem, batch job, or utility

- Eliminates serialization of accesses between DBRC instances
 - Data set RESERVE (or global enqueue) eliminated

- Reduces RECON contention
 - Could provide better responsiveness from IMS online and batch
 - Removes growth constraint

- Parallel RECON Access is optional
 - Specified by DBRC command for a set of RECONs

Parallel RECON Access

- **Uses Transactional VSAM**
 - System facility that provides locking, logging, caching, and commit for concurrent updates to VSAM data sets (RECONs)
 - Exploits Parallel Sysplex
- **Prerequisites**
 - Hardware
 - Parallel Sysplex environment
 - Requires Coupling Facility
 - Software
 - z/OS DFSMS Transactional VSAM (DFSMSStvs)
 - Requires RRS for DFSMSStvs (IMS use of RRS is not required)
 - DFSMSStvs is an optional feature
 - Software license required
 - Special bids will be considered

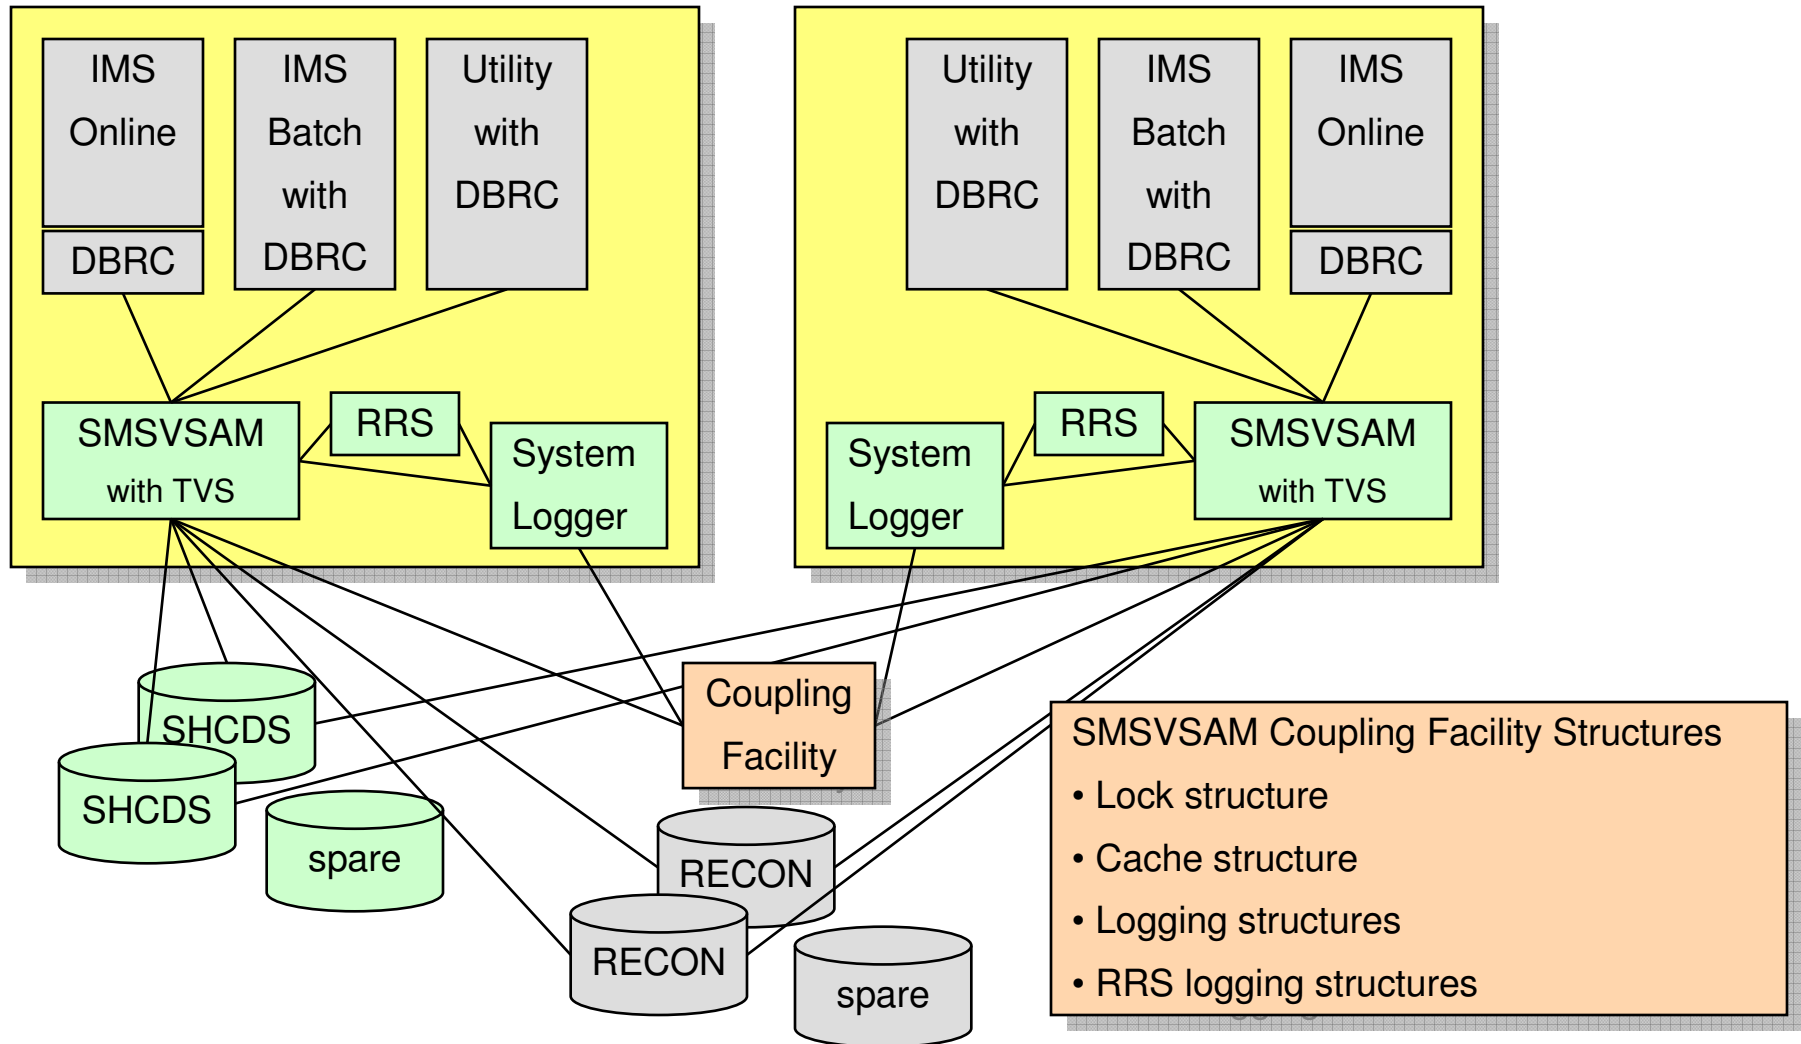
Parallel RECON Access

- Requests are processed in parallel
 - Reduces RECON contention
 - Service times for individual requests may be increased
 - Due to locking, logging, CF accesses, etc.
 - Service times for some requests may be decreased
 - Due to elimination of I/Os and caching
- Requires new operational procedures
 - New failure scenarios
 - New recovery procedures
- Parallel RECON Access is optional
 - A set of RECONs is defined to use it

Transactional VSAM (DFSMStvs) Overview

- **DFSMStvs is an enhancement to VSAM RLS (record level sharing)**
 - RLS uses locking and data caching in CFs
 - Typically used by CICS to allow concurrent updates to VSAM files by multiple CICSs
 - DFSMStvs adds logging, commit, and backout processing to RLS
 - Typically used to allow concurrent batch and CICS updates to VSAM files
 - Used by IMS V10 DBRC to provide Parallel RECON Access
- **SMSVSAM address space**
 - System address space which is started at z/OS IPL
 - Provides RLS and TVS services
 - One address space per LPAR

DBRC RECONs - DFSMStvs Environment

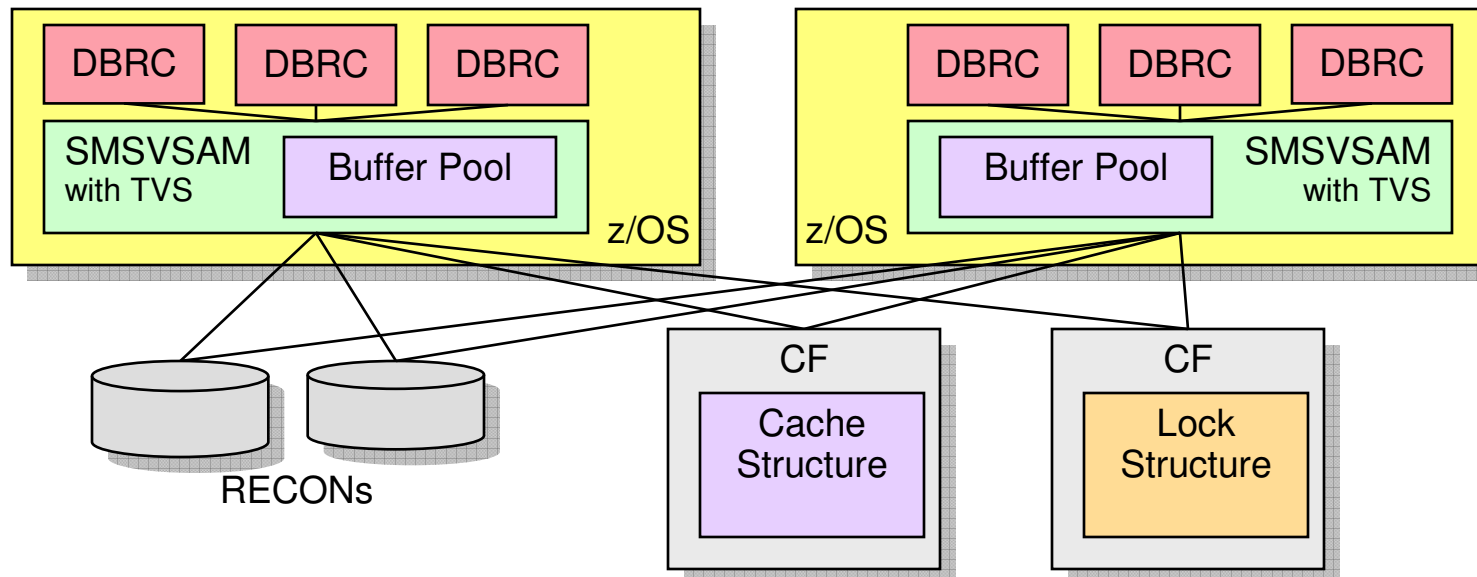


Transactional VSAM (DFSMStvs) Overview

- Recovery of failed users
 - Each DFSMStvs instance has an undo log
 - Used for backout after failures
- Recovery for failed SMSVSAM address space
 - Restarted automatically if it fails
 - Backs out in-flight work and releases retained locks
- Recovery for failed z/OS system
 - Peer recovery
 - Back outs done by another SMSVSAM address space
 - Locks released

Transactional VSAM (DFSMStvs) Overview

- TVS uses a cache structure in CF and a buffer pool in SMSVSAM address space
 - When a buffer in one SMSVSAM is updated, buffers with the same record in other SMSVSAM address spaces are invalidated
- VSAM record is locked when accessed by a user of TVS
 - SMSVSAM has its own lock manager
 - RECON record locked by a DBRC instance

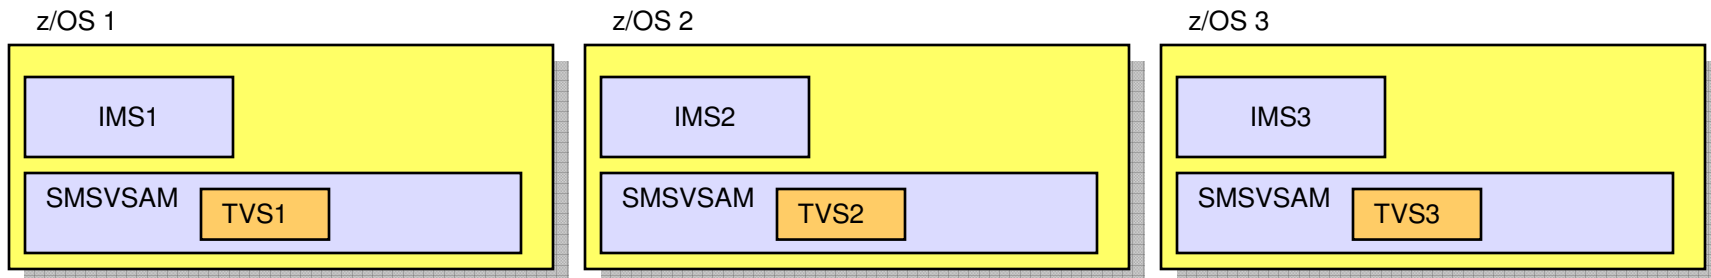


Recovery

- Peer recovery for z/OS system failures
 - In case of system failures, peer instance of DFSMStvs may be started on another z/OS
 - This z/OS must have a primary instance of DFSMStvs on it
 - ARM is highly recommended
 - Peer uses log to backout in-flight work and release retained locks
 - Peer does not accept new work

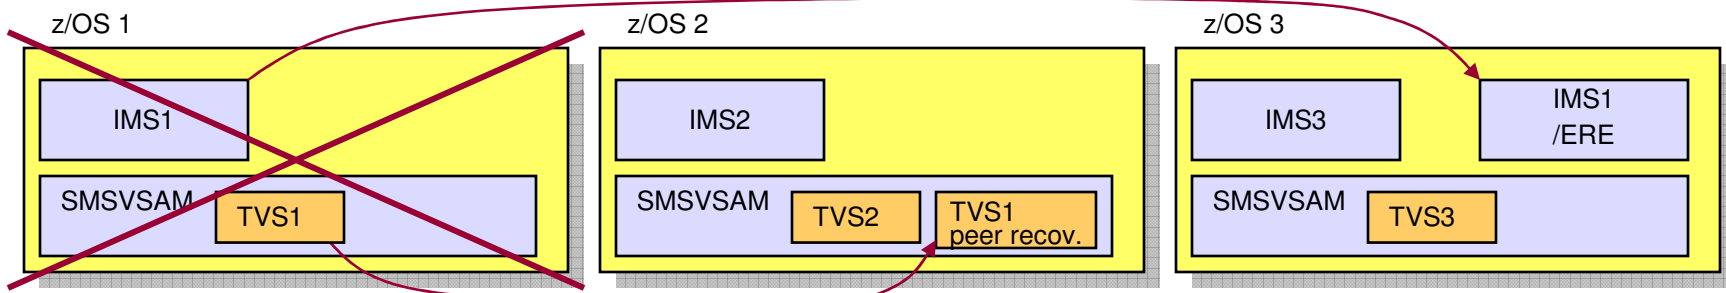
Peer Recovery Example

- Before failure:



- After failure:

- TVS1 peer recovery may be done on z/OS2 or z/OS3
 - It does not require that IMS1 be restarted
- IMS1 may be restarted on a different z/OS



IMS Definitions for Parallel RECON Access

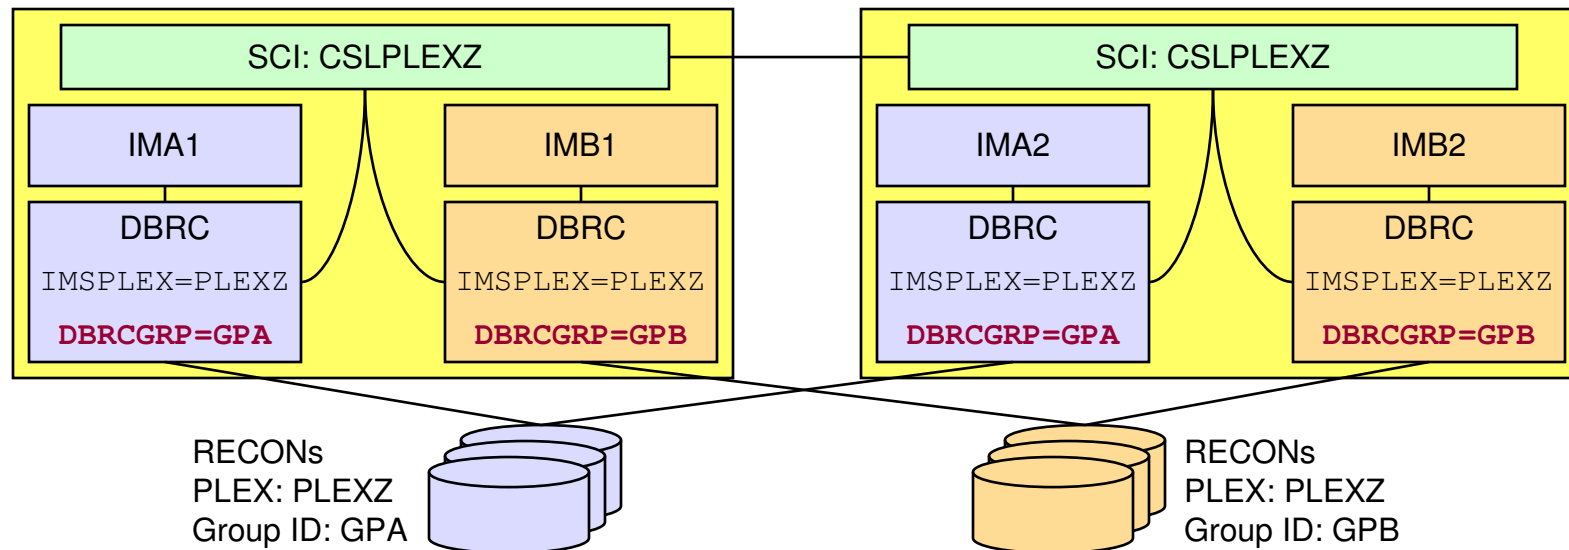
- All IMSs must be V10
 - MINVERS in RECON must be '10.1'

- SCI registration is required
 - DBRC SCI Registration Exit routine is recommended
 - Ensures consistent registration by all DBRC instances
 - Updates required to existing registration exits
 - DBRC group ID is assigned by the exit
 - DBRC group ID has been added to the RECON header

- DBRC I/O exit routine updates required for exit users
 - Meanings of fields have changed

DBRC Sharing Group ID

- DBRC Sharing Group ID is stored in RECONs
 - Used to distinguish multiple sets of RECONs in the same IMSplex
 - Set with `CHANGE.RECON IMSPLEX(plexname, groupid)`
 - Specified on `DBRCGRP= EXEC` parameter
 - Value for an IMS can be set with DBRC SCI user exit (DSPSCIX0)



Parallel RECON Access Implementation

- PRA is turned on with a RECON setting

```
CHANGE.RECON ACCESS (SERIAL | PARALLEL)
```

```
INIT.RECON ACCESS (SERIAL | PARALLEL)
```

- PARALLEL turns on PRA
 - Requires DFSMStvs environment
 - Requires "pair and spare" RECONs
- SERIAL turns off PRA
- IMS does not have to be shut down to change access

LIST.xxx Command Options

- **LIST.xxx CONCURR|STATIC [QUIESCE]**
 - CONCURR - updates may occur during list process
 - Output may not be consistent
 - STATIC - produces output that is consistent as of a point in time
 - Locks held until list processing completes
 - STATIC QUIESCE - RECONs quiesced during list processing
 - No concurrent activity to the RECONs is allowed
 - Batch and utility jobs wait (even if not doing RECON activity)
 - Locking is not used
 - Consider using this option when static is required and many records will be listed (e.g. LIST.RECON)
- **CHANGE.RECON LIST(STATIC|CONCURR)**
 - Sets default for LIST commands

LIST.RECON output

This listing was produced with the CONCURR option.

```
LIST.RECON CONCURR
09.152 10:43:55.891200          LISTING OF RECON (CONCURRENT)          PAGE 0002
```

RECON

```
RECOVERY CONTROL DATA SET, IMS V10R1
DMB#=357                      INIT TOKEN=05306F1839312F
NOFORCER LOG DSN CHECK=CHECK44  STARTNEW=NO
TAPE UNIT=3490      DASD UNIT=SYSDA      TRACEOFF  SSID=**NULL**
LIST DLOG=NO        CA/IC/LOG DATA SETS CATALOGED=NO
MINIMUM VERSION = 10.1      CROSS DBRC SERVICE LEVEL ID= 00001
LOG RETENTION PERIOD=00.001 00:00:00.0
COMMAND AUTH=NONE  HLQ=**NULL**
```

ACCESS(PARALLEL) has been specified for these RECONS

```
ACCESS=PARALLEL  LIST=STATIC
SIZALERT DSNUM=15  VOLNUM=16  PERCENT= 95
LOGALERT DSNUM=3   VOLNUM=16
```

List default for the RECONS is STATIC

TIME STAMP INFORMATION:

```
TIMEZIN = %SYS
```

```
OUTPUT FORMAT:  DEFAULT = LOCORG NONE  PUNC YY
                 CURRENT = LOCORG NONE  PUNC YY
```

```
IMSPLEX = PLEX1          GROUP ID = GPB
```

The DBRC Group ID is GPB

High Availability Large Database (HALDB) Online Reorganization (OLR) SPEs

- **HALDB Online Reorganization SPEs**
 - Retain Bytes/Segments Moved Across Term/Init OLR (PK54615)
 - Query OLREORG and/or DIS DB OLR to get OLR statistics (PK54616)
- **DBRC Changes**
 - Partition DB Record
 - Lists OLR bytes, segments, root segments moved
 - Online Reorganization Record
 - Total number of RAPS processed (PHDAM)
 - Total number of ROOTS to process (PHIDAM)
 - CHANGE.DB DBD(PartDB) OLRBYTES() OLRSEGS() OLRROOTS()
 - DBRC API returns new fields
- **Benefits**
 - Users have information for bytes/segments/roots when OLRs are terminated to allow better decision making when restarting the OLR

DBRC Change Accumulation SPE - PK53223

- Removes Change Accumulation utility failures due to RECON contents changing between the time the CA JCL is generated and executed
 - New Image Copy
 - OLDS archived which changes purgetime
 - New HALDB Online Reorganization
 - Offline Reorganization

- DBRC will allow the JCL generated for Change Accumulation to execute using the purgetime in the JCL as long as input logs verify

- Benefit
 - Minimize disruption to 24 X 7 operations due to Change Accumulation Utility failures.

DBRC Change Accumulation SPE - PK73660

- Provide NOTIFY.CA PITCA support to add a CA that was generated outside of IMS
 - New type of CA - PITCA
 - Error flag is set on and DBRC will not use it for recovery or
 - Runtime and stoptime are not switched when a PITCA is recorded in RECON
 - Later CAs are not deleted
 - If a CA record with the same stoptime as specified in the command already exists, the command fails with DSP0025I RECON RECORD ALREADY EXISTS
 - If a CHANGE.CA INVALID|VALID command is attempted, it fails with new message DSP0096E THE STATUS OF THE SPECIFIED CA RECORD CANNOT BE CHANGED
 - The DBRC API also returns the new flag to indicate the CA is a PITCA

- Benefit
 - Allow PITCA CAs to be recorded in RECON with real times without erasing later CAs

Migration Considerations for IMS V10

Prerequisites for IMS V10

- z/OS V1.7 or later
- DBRC Parallel RECON Access requires z/OS DFSMSStvs
 - Special bids will be considered for those who do not already have DFSMSStvs
- Image Copy 2 fast replication support requires z/OS V1.8



IMS

What's New in IMS Database Recovery Facility (DRF)

Mitchell Dooley
IMS Tools Developer
mdooley@us.ibm.com

Information Management software



IMS Database Recovery Facility

- High Performance IMS Database Recovery Solution
 - Recover multiple DBDS in a single job step
 - Multiple address space architecture
 - Parallel processing of Input / Output
 - Single pass of log and CA data sets

- Integrated with other IMS Tools products
 - IMS High Performance Image Copy
 - IMS High Performance Pointer Checker / IMS HP Fast Path Utilities
 - IMS Index Builder / DFSPREC0
 - One step processing for all utilities instead of one step for each utility

IMS Database Recovery Facility

Recovery Options

- Recover to End of Logs
 - Full recovery to current time
 - Uses IC + logs / CA
 - DBDS must be unallocated at recover-to time
- Time Stamp Recovery
 - Recover to a specific time
 - Uses IC + logs / CA
 - DBDS must be unallocated at recover-to time
- Point-In-Time Recovery
 - Recover to any point in time
 - Uses IC + logs / CA
 - DBDS may be allocated or unallocated at recover-to time
 - Only committed updates are recovered
 - List of Open UOW reported

IMS Database Recovery Facility

- A 'one-button' recovery solution!
- During Recovery phase
 - IMS HP Image Copy
 - Produce image copies as part of recovery processing
 - Image copies are registered with DBRC
 - Block by block image copy
 - Pointer Checker using IMS HP PC or IMS HP Fast Path Utilities DEDB PC
 - Hash check pointer validation run as part of recovery processing
 - Block by block pointer check, not post processed
- In Post Recovery phase
 - IMS Index Builder
 - Build primary and secondary index data sets during the recovery process
 - IMS Index/ILDS Rebuild utility (DFSPREC0)
 - HALDB Primary Index and Indirect List Data Set (ILDS) rebuild

IMS Database Recovery Facility

- **Database Copy Generation**
 - Create copies of database data sets without accessing production copies
 - Input is prior IC + logs / CA
 - Copies can be created to any point in time
- **Incremental Image Copy**
 - Create a new image copy using a prior IC + logs / CA
 - Production DB is not accessed
 - IC is either batch or concurrent, depending on state of DB
 - IC can be created to any point in time
 - Requires IMS High Performance Image Copy
- **Automatic Delete/Define of Database Data Sets**
 - Output data sets are created as part of the recovery process
- **Allocate/Open Option on START VERIFY**
 - Logical and physical validation of data set availability prior to running the actual recovery job

Thank You for Joining Us today!

Go to www.ibm.com/software/systemz to:

- ▶ Replay this teleconference
- ▶ Replay previously broadcast teleconferences
- ▶ Register for upcoming events