

Time is Money: How IMS can streamline business tasks to save you money

Greg Vance, IMS Senior Technical Staff Member

Jack Pinto, IMS Senior Software Engineer



Disclaimer

© Copyright IBM Corporation 2011. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, Information Management, IMS, CICS, DB2, WebSphere and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

Agenda

- Background
- Parallel RECON Access Overview
- Performance Overview



Why Parallel Recovery Control (RECON) Access Use cases

- I must restart several IMS images at the same time and open access to my databases in parallel
- I need to run a large number batch DB utilities at the same time
- I need to reduce my transaction response times when log switches occur
- **All can be addressed by parallel RECON access**

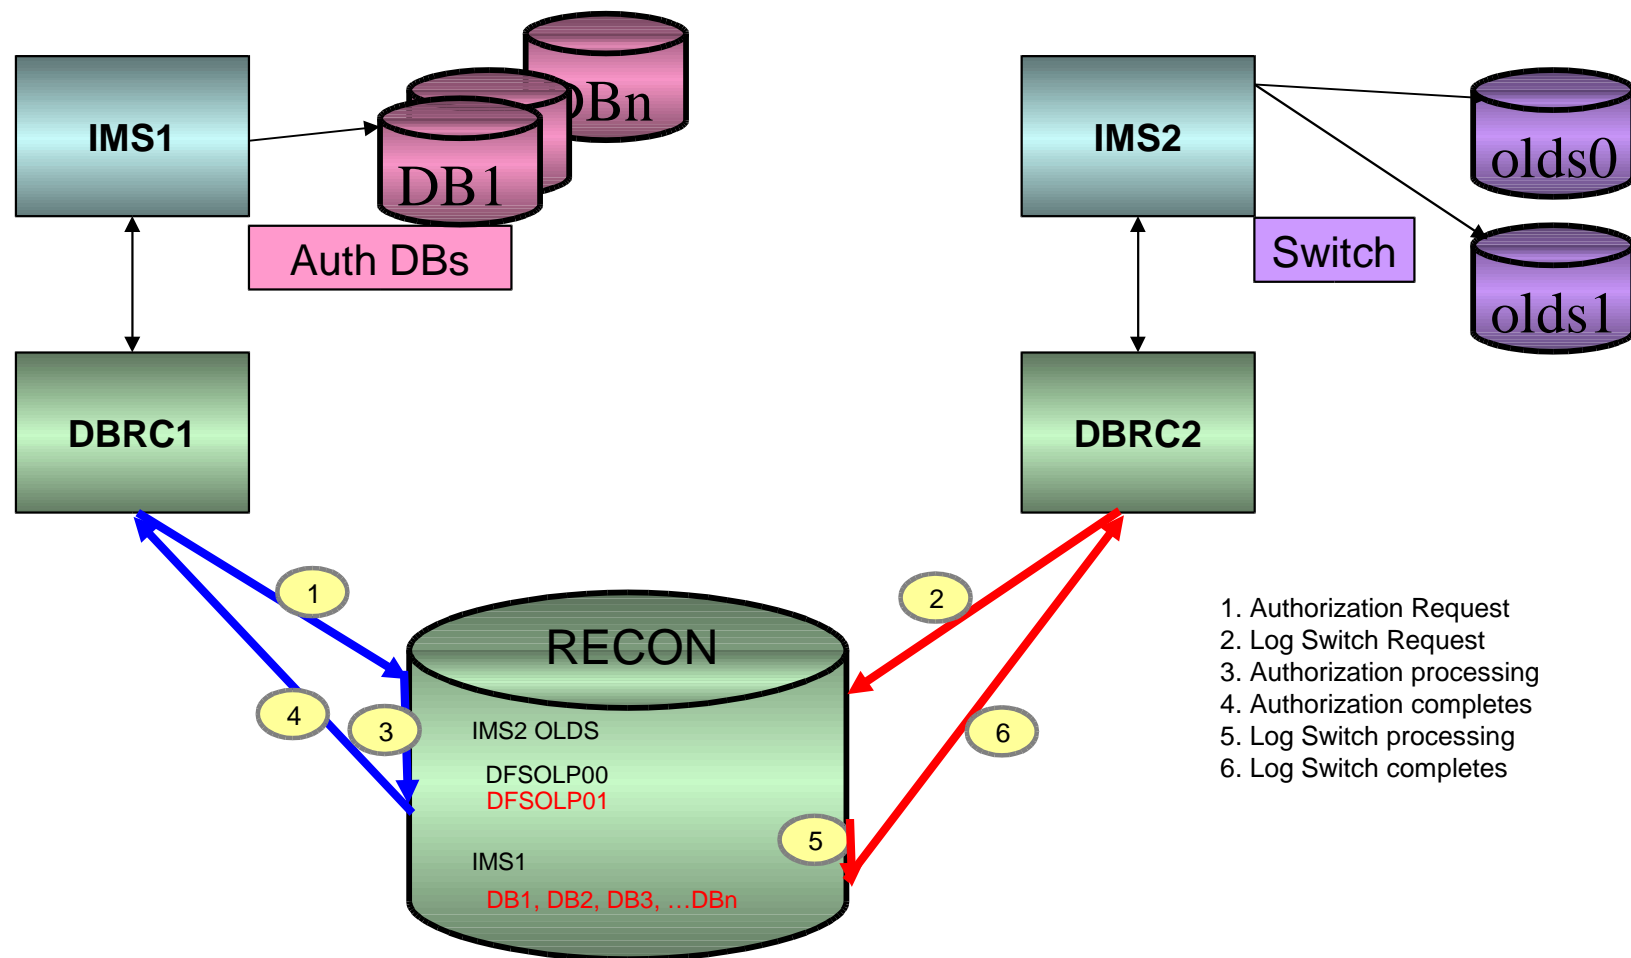


Parallel RECON Access

- Allows multiple DBRC instances to access the RECONs concurrently
 - DBRC instance: IMS Online subsystem, batch job, or utility
 - Parallel RECON Access is optional
 - Multiple requests from a given IMS are processed one at a time
- 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

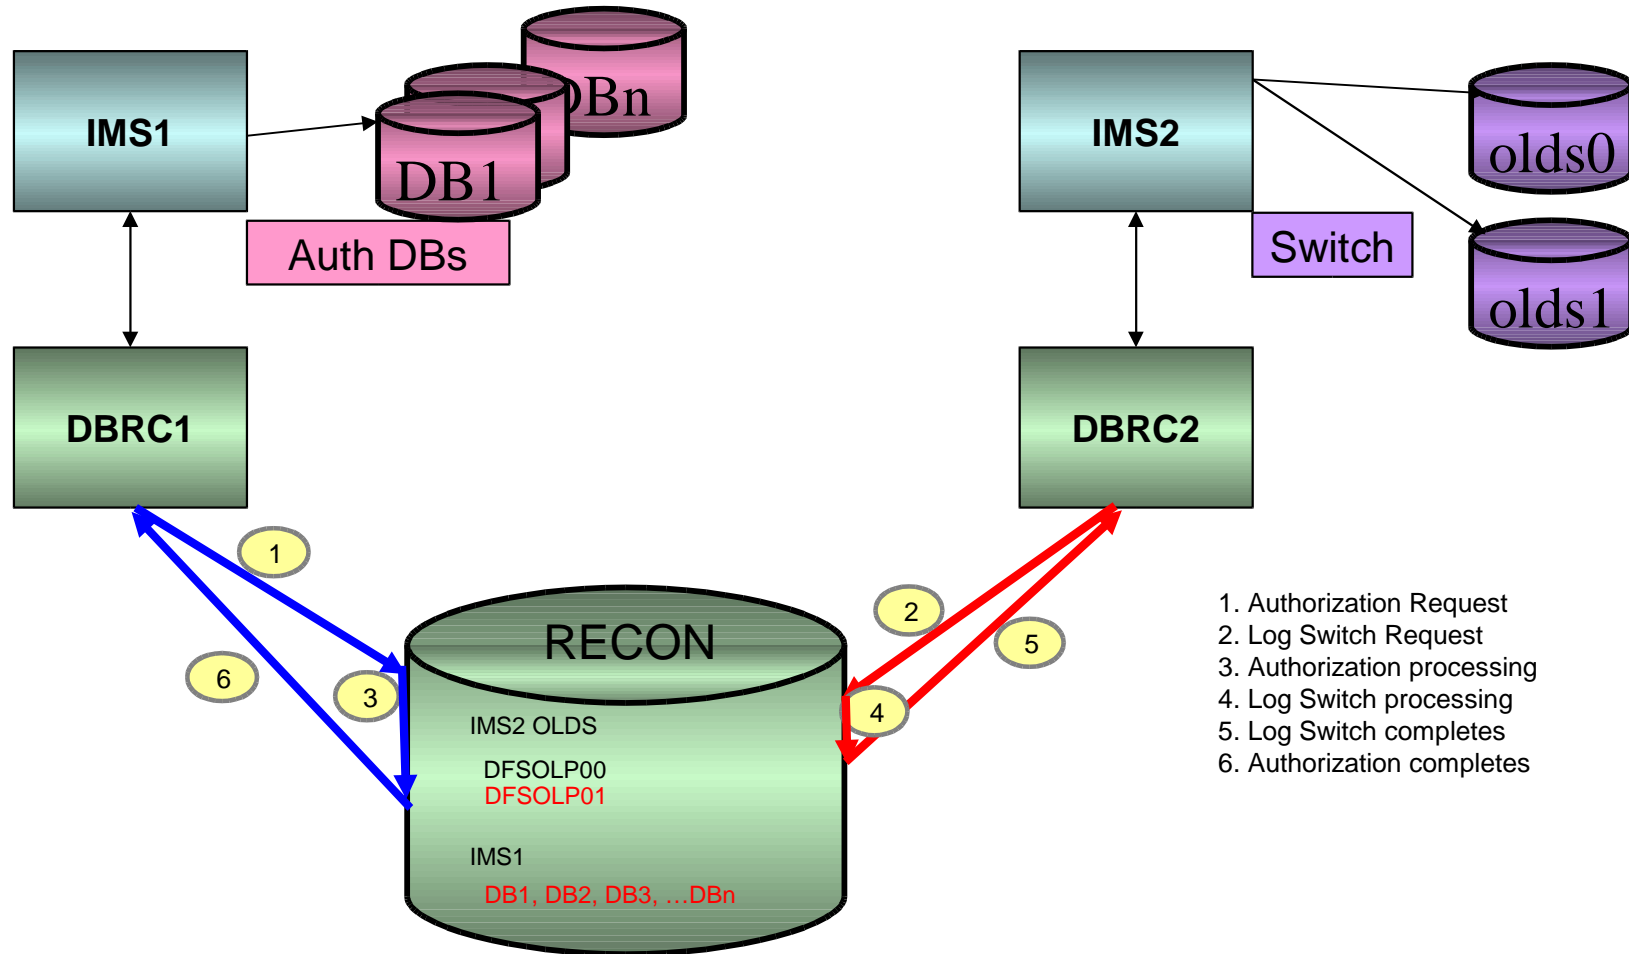


Serial Request Processing



1. Authorization Request
2. Log Switch Request
3. Authorization processing
4. Authorization completes
5. Log Switch processing
6. Log Switch completes

Parallel Request Processing



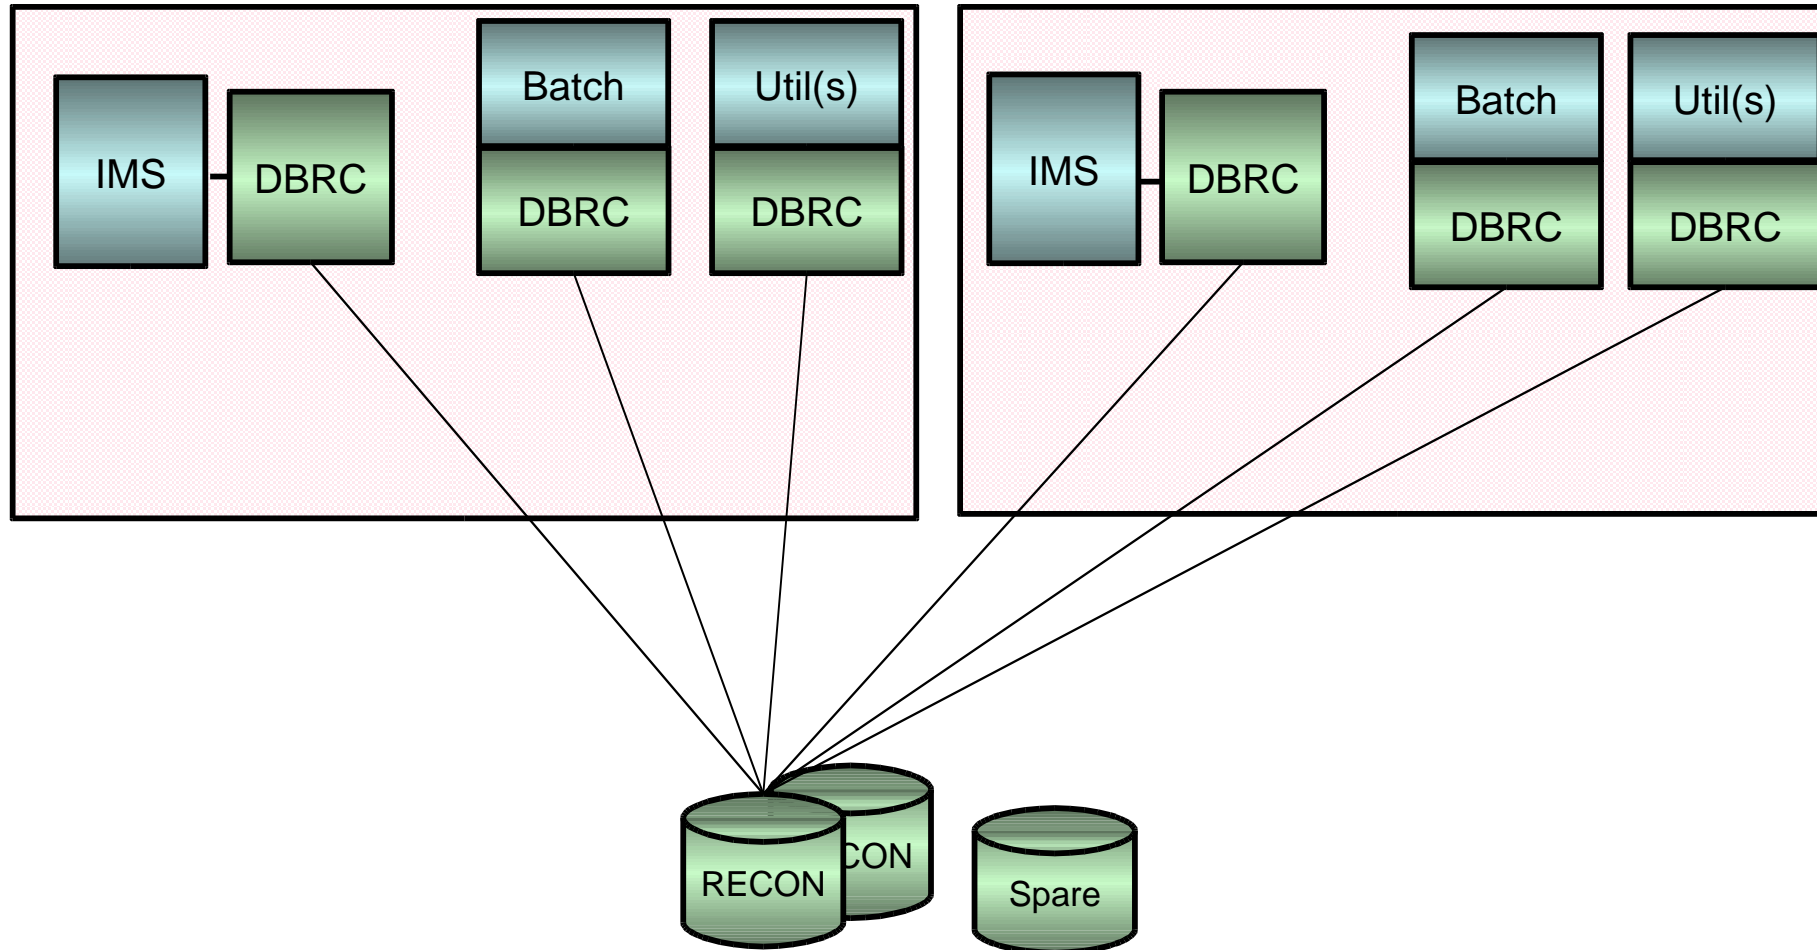
1. Authorization Request
2. Log Switch Request
3. Authorization processing
4. Log Switch processing
5. Log Switch completes
6. Authorization completes

Parallel RECON Access


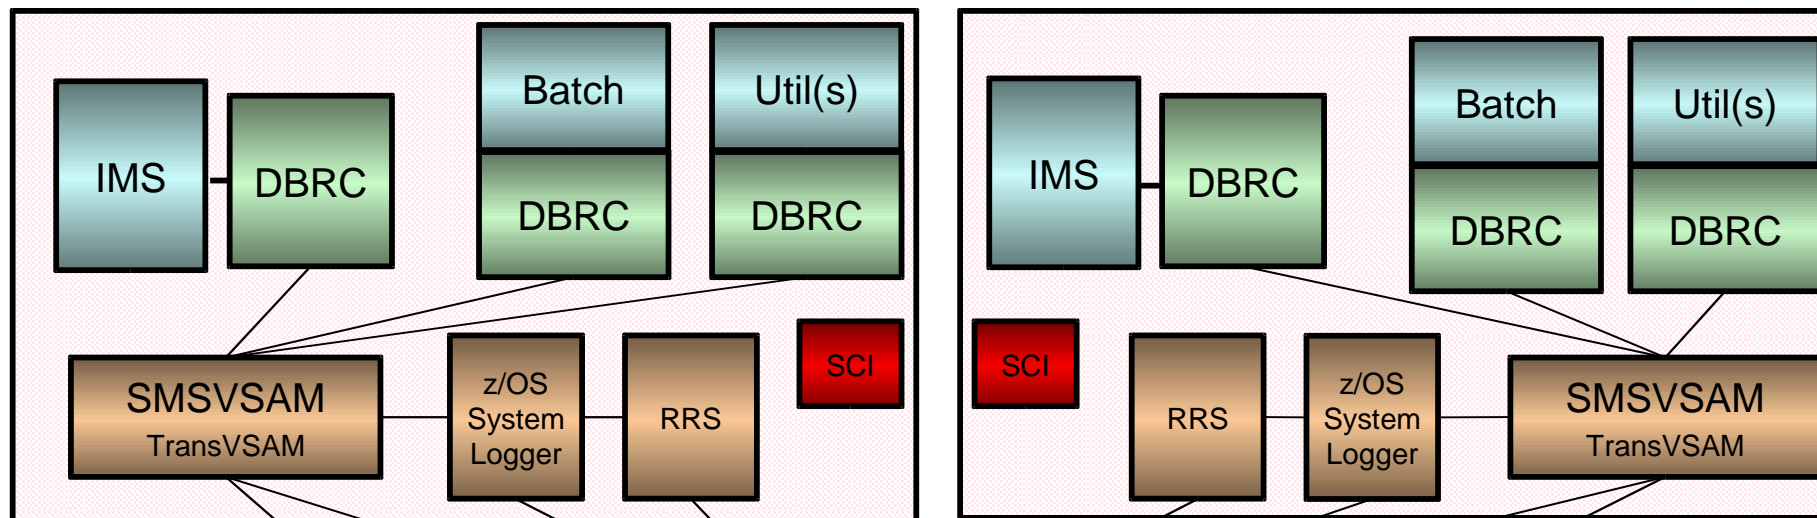
- Uses DFSMS Transactional VSAM (DFSMSStvs)
 - System facility that provides locking, logging, caching, and commit for concurrent updates to VSAM data sets (RECONs)
 - Exploits Parallel Sysplex
- Uses Structured Call Interface (SCI) for DBRC communication
- Prerequisites
 - Hardware
 - Parallel Sysplex environment
 - Requires Coupling Facility
 - Software
 - z/OS DFSMS Transactional VSAM (DFSMSStvs)
 - Requires RRS
 - DFSMSStvs is a priced feature




Serial RECON access environment



Parallel RECON access environment



Coupling Facility



- SMSVSAM Coupling Facility Structures
- Lock structure
 - Cache structure
 - Logging structures
 - RRS logging structures

Parallel RECON Access Definition and Set Up

- IMS and IMSplex definitions
- RECON definitions
- MVS/DFSSMS related Setup
 - SYS1.PARMLIB member definitions
 - SHCDS data set
 - Log stream definitions
 - Structure definitions
- Security specifications



Parallel RECON access processing

- External Changes
 - Commands
 - User Exits

- Some behavioral differences from serial access processing
 - Initialization
 - Locking
 - DBRC activity quiesce
 - Failure Handling



Turning on parallel access

- Parallel RECON access is turned on and off by a DBRC command
 - *IMS systems do not have to be shut down*
- CHANGE.RECON ACCESS(PARALLEL) turns on parallel access
 - There must be two active RECON data sets and a spare RECON
 - RECON associated with an IMSplex
 - No RSR tracking subsystem
- CHANGE.RECON ACCESS(SERIAL) turns off parallel access



User interfaces - Commands

- **CHANGE.RECON new/change keywords**
- **LIST.XXX new keywords**
- **BACKUP.RECON new behavior with PRA**
- **LIST.RECON Output changes**



User interfaces – RECON I/O Exit

- The RECON I/O exit routine (DSPCEXT0) interface is extended to support parallel RECON access
 - A **Request Token** replaces the RESERVE Sequence Number
 - This token is a STCK for parallel access
 - A **DBRC Instance Token** is passed on every exit invocation
 - May be used to qualify the Request Token for parallel access in case of duplicate STCK values
 - 16-byte area
 - The Instance Token is zeroes for serial RECON access
 - The **Changed Record Count** is always zero for parallel access
 - A **new flag** indicates parallel RECON access
 - There are new exit invocations with parallel access
 - **Commit exit call** - invoked when the previous updates for the current series have been committed to the RECON data set
 - **Backout exit call** - invoked when the previous updates for the current series have been backed out



RECON I/O Exit (cont.)

- Multiple instances of the I/O exit can be running in parallel
 - May cause different behavior if exits write to common data sets.
- With PRA, Request Token/Changed Record count cannot be used to determine if updates were missed.
- Updates during backout not passed to exit
 - Done by TVS
 - Should save updates until commit before (discard if backout)



User interfaces – DBRC SCI Registration Exit...

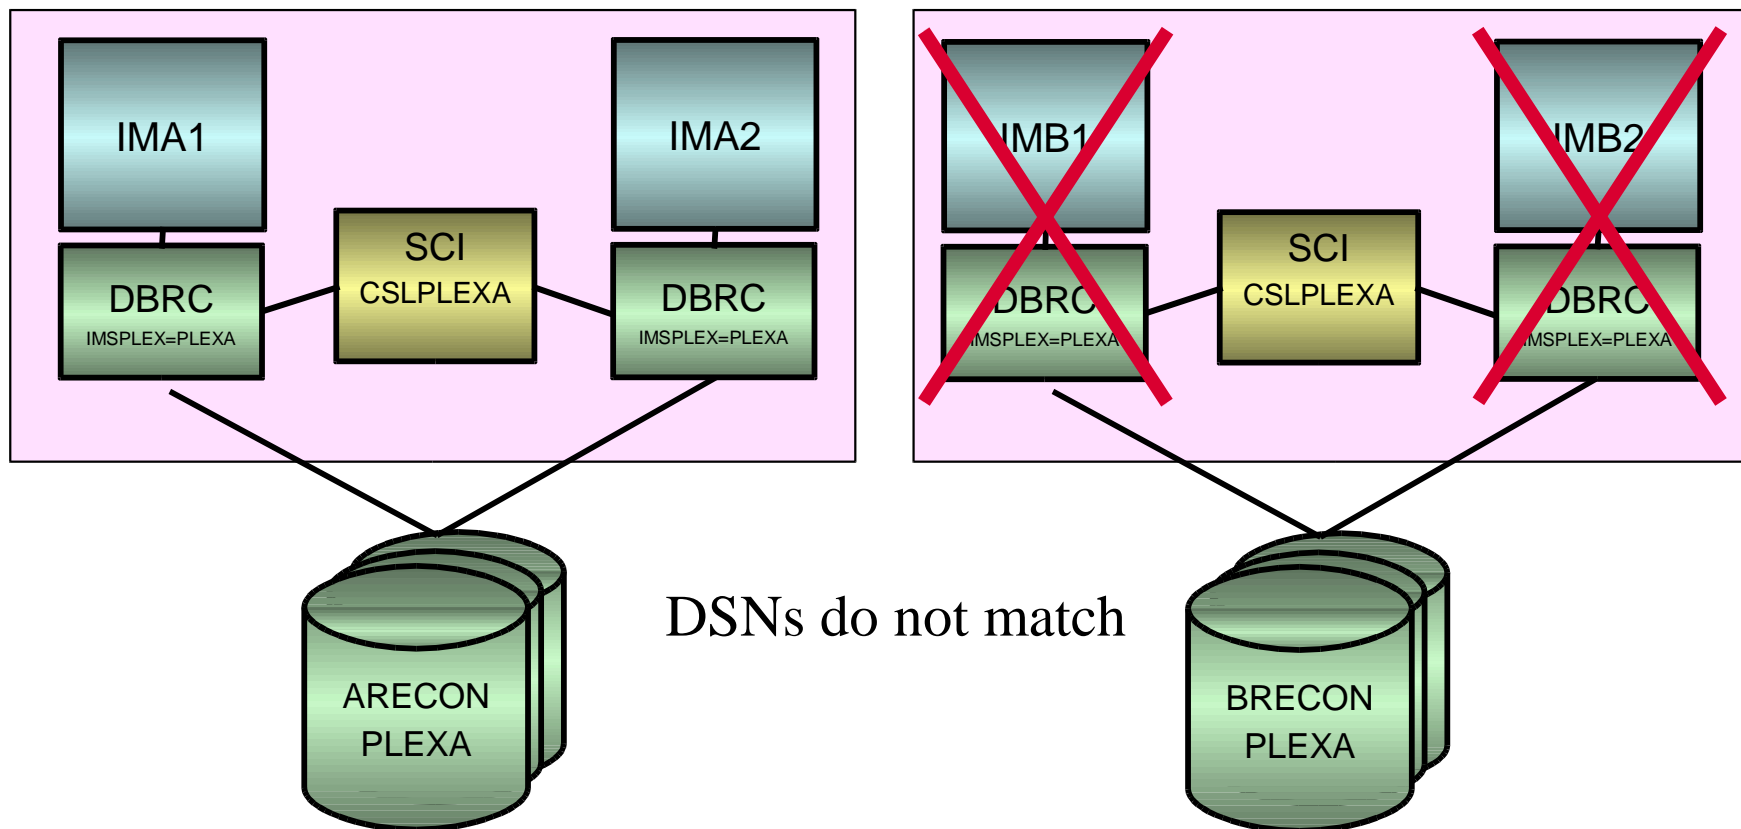
- The DBRC SCI Registration exit routine (DSPSCIX0) interface is extended to support the DBRC group ID
 - A RECON **Data Set Name**
 - An **IMSplex name return area**
 - 5-byte area in which DSPSCIX0 returns an IMSplex name
 - The **IMSPLEX EXEC parameter value**
 - if specified as an EXEC parameter, the value of the IMSPLEX parameter
 - The **parameter list version (new)**
 - A **DBRC Group ID return area (new)**
 - 3-byte area in which DSPSCIX0 returns a DBRC Group ID
 - The **DBRCGRP EXEC parameter value (new)**
 - if specified as an EXEC parameter, the value of the DBRCGRP parameter
- Sample Exit updated to support DBRC Sharing Group ID



DBRC Initialization – Configuration matching

- During initialization, if another DBRC is active, the new one will be given its configuration information
 - Access mode
 - COPY1|COPY2 Data set names
- If the access mode is PRA, then the new DBRC must match the COPY1|COPY2 configuration
- WTOR issued if configurations do not match for online DBRC region
- User can CANCEL or determine the cause and RETRY
- Matching also done on end-quiesce processing

Configuration Matching – Multiple RECONS in IMSplex



DBRC Initialization fails with DSP1175W & DSP1176A messages

Quiesce Overview

- DBRC quiesces RECON activity to allow certain processing, for example:
 - BACKUP.RECON
 - copying to the spare data set
 - I/O error processing
 - CHANGE.RECON REPLACE
 - LIST command processing
 - turning on/off parallel RECON access
- Allows requestor to have sole ownership of RECON data sets

Locking Overview

- Locking is used to prevent access to a resource
- For serial access we lock at the data set level by RESERVING the RECON data sets
 - the 'lock' is held for the entire DBRC request
- For PRA access, TVS locks at a record level
 - the locks are held for the duration of the VSAM access or until a commit point
- Record level locking introduces new 'failures'
 - Deadlocks
 - Timeouts (waiting to get a lock took too long)
 - Retained Locks (TVS unable to backout updates)
 - Lost Locks (let's hope never see this!)
- Access to RECONs w/o DBRC must understand locking implications



Locking of VSAM Records with PRA

- Locks are either shared or exclusive
 - Shared for get requests
 - Exclusive for 'get for update' requests

- Locks held during access or until a commit point
 - Locks for updates held until commit
 - Locks for read depends on options
 - LIST command with CONCURR
 - Lock released after record is read
 - LIST command with STATIC
 - Locks held until commit
 - Other reads hold locks until commit
 - Except reads used to browse multiple records looking for a specific record



Failure Handling - Retry Overview

- With TVS, new VSAM request errors can occur that DBRC needs to tolerate
 - Deadlock/Timeout/TVS failure for example
- Will not ABEND or return to the DBRC request caller with a bad RC since these errors can be recovered from by DBRC
 - In general the error requires DBRC to backout the current UR
 - DBRC then needs to retry the DBRC request
- DSP1184W message if same error hit 5 times
- I/O errors require the DBRC request be retried with PRA if in-flight updates exist.
- Certain Quiesce failures will also be tolerated by retrying the DBRC request



Automated Operation Issues

- **RECON Reconfiguration**
 - New set of messages (quiesce related) issued when using parallel access
 - Key messages: DSP1130I and DSP1132I
 - No automatic loss messages (DSP1141I, DSP1145I)
 - DSP0388I message still issued
- **Repeated Retry**
 - Key messages: DSP1184W
 - Messages issued of same error encountered several times
 - Could mean TVS is unavailable or application holding lock
- **Configuration Matching**
 - Key messages: DSP1175E and DSP1176A
 - Check that job has all 3 RECON DDs if DD allocation done
 - Check that all volumes for RECONs are available

Serial vs. Parallel RECON Access Comparison

Serial	PRA
<ul style="list-style-type: none"> • Hardware reserves (or global serialization) serialize access to the entire RECONS 	<ul style="list-style-type: none"> • Individual RECON records are locked for serialization • Deadlocks and lock timeouts may occur
<ul style="list-style-type: none"> • RECON status maintained in RECON header/header extension • RECON header/header ext. read during "logical open" to determine RECON status 	<ul style="list-style-type: none"> • RECON configuration information maintained in RECON header extension • RECON header read during "logical open" to get any updates due to CHANGE.RECON
<ul style="list-style-type: none"> • "Multiple update" processing used for recoveries from failures 	<ul style="list-style-type: none"> • TVS log used for recoveries from failures
<ul style="list-style-type: none"> • Uses VSAM LSR pools 	<ul style="list-style-type: none"> • Uses VSAM RLS pools and cache structures
<ul style="list-style-type: none"> • SCI registration and ARLN are optional 	<ul style="list-style-type: none"> • SCI registration and ARLN are required
<ul style="list-style-type: none"> • I/O Error results in reconfiguration and retry of last I/O 	<ul style="list-style-type: none"> • I/O Error results in reconfiguration and retry of entire request

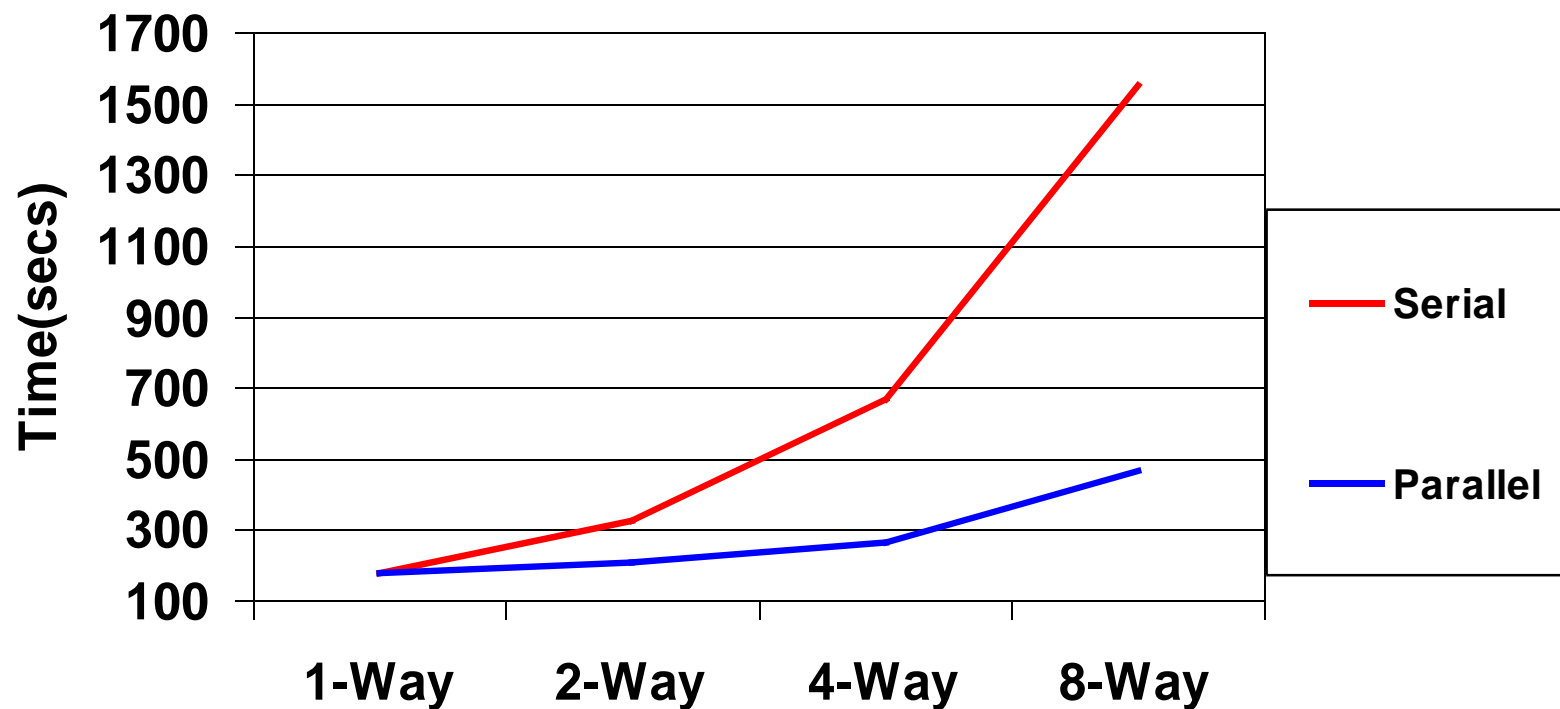
Performance Considerations

- Transactional VSAM access has extra overhead vs. traditional LSR access
 - Cross-address space access to SMSVSAM server
 - Overhead of CF cache access
 - Overhead of record locking
 - Overhead of logging
- However... access to the RECON is no longer serialized across the entire sysplex! Also, no logging in RECON
- Performance considerations
 - SMSVSAM local buffer pool sizes
 - Cache structure sizing for VSAM RLS buffer caching
 - RLS lock structure sizes
 - Log stream structure sizes
 - DEADLOCK_DETECTION settings
 - Etc.



IMS FP Area PREOPEN Comparison

4000 Fast Path Areas



Time savings and ECSA savings translate into \$\$ saved

- PRA: DEDB areas Pre-open faster
- PRA: Batch Image Copy jobs complete faster
- PRA: Database Recoveries complete quicker
- WADS Channel program: 10% reduction in response time
- Logging: Extended Format OLDS (striping): 35% log rate gain
- Logging: 64-bit log buffers: ECSA reduction
- FastPath 64 bit buffers: ECSA reduction
- FastPath Secondary Index: 60% ITR improvement
- APPC/OTMA Synch SMQ with XCF: 40% to 50% ITR gain



Thank you for joining us today!

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

Replay this teleconference

Replay previously broadcast teleconferences

Register for upcoming events