

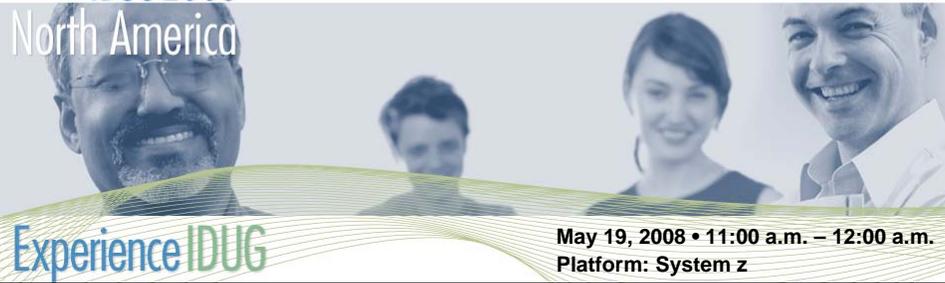


Session: K02

## IMS Application Dependent Disaster Recovery Solutions

Glenn Galler  
IBM SW IT Specialist (IMS)  
Advanced Technical Support  
Ann Arbor, Michigan  
gallerg@us.ibm.com

IDUG 2008  
North America



May 19, 2008 • 11:00 a.m. – 12:00 a.m.  
Platform: System z

## Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	GDPS*	S/390*
CICS*	HyperSwap	Sysplex Timer*
DB2*	IBM*	Tivoli*
e-business logo*	IBM eServer*	TotalStorage*
Enterprise Storage Server*	IBM logo*	z/OS*
ESCON*	NetView*	z/VM*
FICON	OS/390*	zSeries*
FlashCopy*	Parallel Sysplex*	

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of the Intel Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

UNIX is a registered trademark of The Open Group in the United States and other countries.

\* All other products may be trademarks or registered trademarks of their respective companies.

### Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

## Acknowledgments

- Peter Armstrong (BMC)
  - Book: “DBRC in Practice”
    - <http://www.dbazine.com/ofinterest/oi-articles/armstrong6>
- Rich Lewis (ATS)
  - Review Presentation and White Paper: Image Copies
- Helene Lyon (IBM Europe)
  - Review Presentation and Presentation: IMS Backup and Recovery
- Mitch Dooley (DM Tools Developer)
  - Review: IMS Recovery Expert
- Pat Schroeck (DBRC Developer), Rick Long (IBM Australia)
  - Review of presentation

## Agenda

- Concepts
- Two Disaster Recovery Strategies
  - Determines amount of data lost and time to recovery at Remote Site
  - Strategy Determines Data Sent to Remote Site
- Production Site
  - Log Management
  - Image Copies and RECON data set
- Remote Site
  - RECON cleanup work
  - Recovery strategies
  - Restart options

For Application Dependent Disaster Recovery, there are two common methods and then many variations of the two methods. In this presentation, we will focus on these two methodologies. The first provides a specific point in time recovery to an image copy of the data. There are no logs sent to the remote site so there is no forward or backward recovery of the database. With this solution, there is a great amount of data lost in the event of a disaster.

The second method includes sending image copies and the DBRC RECON data set to a remote site, but additionally, the user also sends the log data sets which allows the remote site to do forward and backward recoveries of the data. This means there is less data lost in the event of a disaster. The procedure to bring the remote site up to date is more complicated than the first solution.

For both methodologies, there is usually some cleanup work needed for the DBRC RECON data set.

## Concepts



- **Disaster recovery**
  - Process of recovering an environment after a major disaster
  - Bring to the point at which business can be conducted
- **Recovery Time Objective (RTO)**
  - How much time is available/allowed to recover the applications
  - All critical operations are up and running again
- **Recovery Point Objective (RPO)**
  - How much data can be lost in the event of disaster?
  - What is the last point-in-time when all data was current?

Disaster Recovery is the ability to restart a production environment in a remote site following a disaster that makes the primary site unusable. Customers must determine the amount of time and the amount of data they can afford to lose when they are setting their recovery goals.

The Recovery Time Objective (RTO) is the amount of time it takes to bring up the production system in the remote site. The Recovery Point Objective (RPO) is the amount of data that can be lost in such a disaster. Typically, most institutions with critical systems on mainframe computers have a RTO of under two hours. By comparison, most UNIX environments with less critical systems have an RTO of under eight hours. The RTO and RPO objectives need to be determined on an Application-by-Application basis.

The cost of the disaster recovery solution needs to be in proportion to the business value of IT. In other words, the disaster recovery solution should not cost more than the financial loss suffered from the disaster itself.

The amount of data mirrored to a remote site and the currency of the data will greatly affect both the RTO and RPO of the remote site. If the production system can simply be emergency restarted from a recent checkpoint, there will be less data lost (RPO) and the restart will take less time (RTO). The amount of time increases dramatically for IMS when the system needs to recovery data from older image copies. In disaster recovery terms, it is always better to do a Restart than a Recovery.

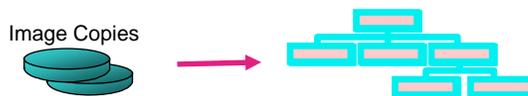
Once the RTO and RPO objectives are known, the customer is able to determine which disaster recovery methodology works best for their needs.

## Two Application Dependent DR Strategies



- **#1: Recover the Databases to an Earlier Image Copy**

- Image Copies and RECON are shipped to remote site
- Production activity is quiesced for image copy
  - Procedure may include IMS, DB2 and CICS
- Data since last image copy is lost
- RTO is low since recovery time is short
- RPO is high since the log updates are lost



IDUG 2008 North America

6

The first disaster recovery methodology is called “Recover the Databases to an Earlier Image Copy.” The image copies and DBRC RECON data set are the only items shipped from the Production site to the Remote site. There are no logs, Change Accumulation or, Recovery Log Data Sets shipped. This means that any data that was updated since the image copy was taken will be lost in the event of a disaster.

It is necessary to take a Clean Image Copy (which will be discussed later) as there are no log data sets to bring a Fuzzy Image Copy up-to-date. If the Production site includes other Database Managers like DB2 or CICS/VSAM along with the IMS, then a procedure to quiesce all of the systems is advisable.

For example, the following steps would quiesce the productions systems:

- DBD IMS Databases
- Suspend DB2 Logging
- Close CICS VSAM Data Sets
- Hold DFHSM Functions
- Take Clean Image Copies
- Release DFHSM Functions
- Open CICS VSAM Data Sets
- Resume DB2 Logging
- Start IMS Databases

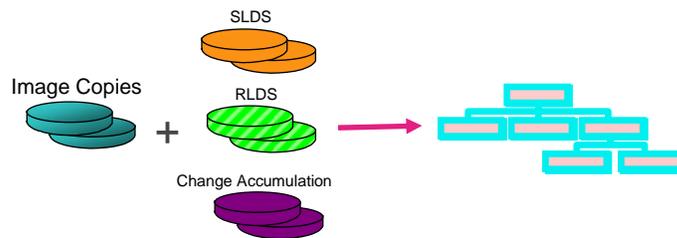
The RTO for this solution is very fast since there are no recoveries to be performed. The image copies represent the state of the data at a particular time. On the other hand, the RPO, or the amount of data lost, is very high.

## Two Application Dependent DR Strategies



- **#2: Recover the Databases to the Last Good Log Data Set**

- Image Copies, RECON and Log data sets are shipped to remote site
- Forward and Backward recovery is performed with the logs
- RTO is higher because recoveries are required
- RPO is lower because log updates are applied to image copy



IDUG 2008 North America

7

The second methodology is called “Recover the Databases to the Last Good Log Data Set.” With this solution, the logs, the Change Accumulation data sets and the Recovery Log Data Sets (RLDS) if they are used are shipped to the remote site along with the image copies and DBRC RECON data sets. In the event of a disaster, the image copies are restored and the logs are applied to them to produce a more recent database.

The RTO is higher with this methodology as there is some time to do forward recovery of the databases. The RPO is less with this solution as committed updates are applied to the image copies to bring the database up-to-date.

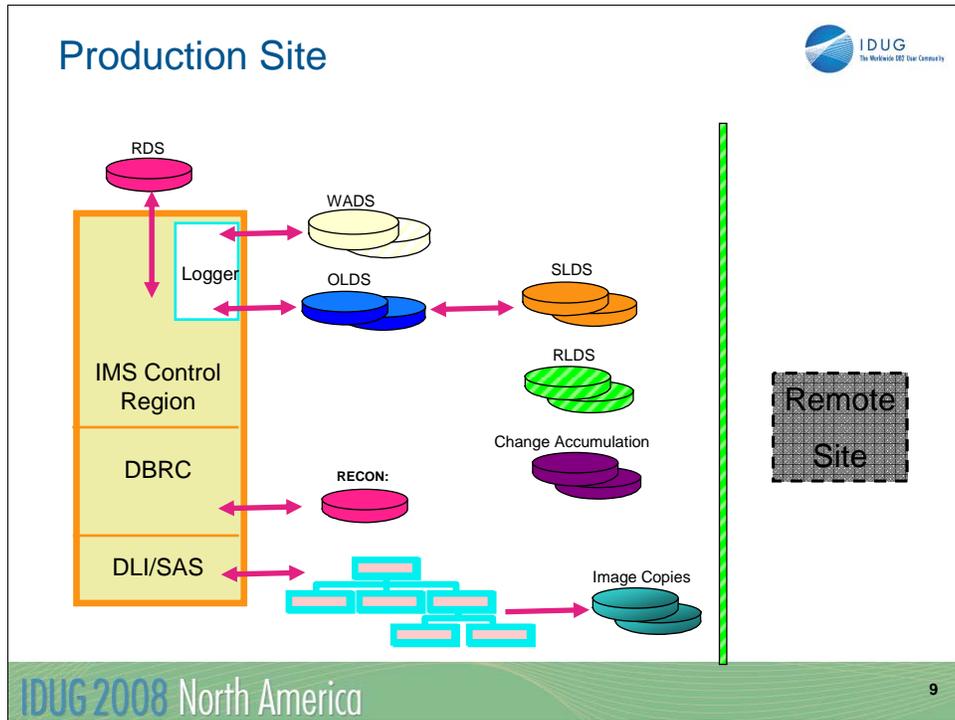
## IMS Production & Remote Site Responsibilities

- **Production Site**

- Database Recovery Control (DBRC)
  - The Supervisor of Recovery
  - Keep track of logs, CAs, and ICs needed for recovery
  - DBRC is optional, but strongly recommended
- Log Management Utilities
- Backup process
  - Create backup copy of the RECON data set
  - Create Image Copies of the data

The Production site is responsible for producing the recovery data for the remote site. While DBRC is optional, it is strongly recommended for use in recovery situations. All of the techniques described in this presentation can be done if the user does not register their databases with DBRC, however, the user will be required to manually build their JCL which is both time-consuming and error-prone.

The DBRC RECON data set keeps track of image copies, log data sets, Change Accumulation data sets and Recovery Log Data Sets and the GENJCL function can automatically create the recovery JCL that drives the recovery utilities. Mistakes are much less likely when DBRC is used in this manner. The Production site is also responsible for making backup copies of the RECON data set for the remote site.



IMS uses log data sets to record database activity. The logs contain before and after images of the data. For this reason, IMS is able to recover data in failure and disaster recovery situations.

The online IMS system writes Online Log Data Sets (OLDS) while IMS Batch jobs create System Log Data Sets (SLDS). When the OLDS becomes full, IMS switches from one OLDS to another OLDS. An archive of the OLDS data set is required to move the OLDS data into an SLDS data set. This frees up the OLDS data set so that it can be used again. DBRC is automatically notified when a new OLDS data set is being used.

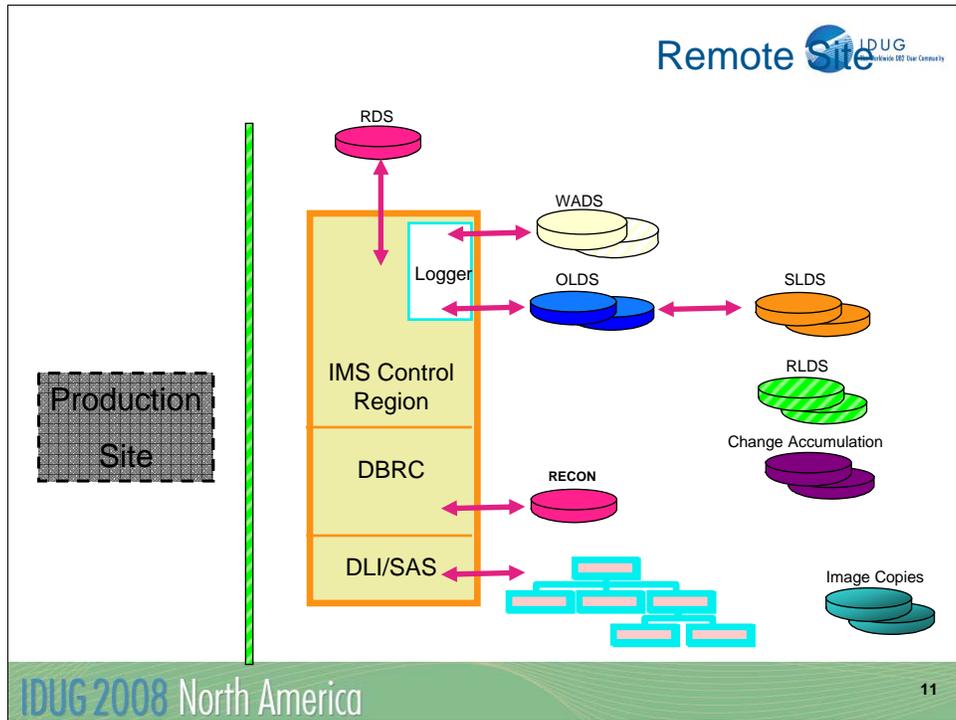
IMS writes system checkpoint information to the Restart Data Set (RDS). During each checkpoint, IMS creates or updates a checkpoint ID table; IMS uses this table during IMS restart to determine from which checkpoint to restart the system. If, for any reason, the RDS is not available at restart, IMS can obtain the required checkpoint information from the log. While the use of the log to recover this checkpoint information may lengthen the restart process, it does mean that the RDS is not required to be sent to the remote site.

## IMS Production & Remote Site Responsibilities

- **Remote Site**

- Clean up the Backup copy of RECON Data Set
- Recovery process (Forwards)
  - Restores followed by apply of the after-image updates
- IMS Batch Backout (DFSBB00)
  - Backout inflight UOWs after Full Database Recovery
- Restart Process

The Remote site is responsible for recovering the database data sets. Before this can be done, the DBRC RECON data set usually needs to be cleaned up. The recovery of the database data sets may be just to an earlier image copy or it could mean applying log updates to an image copy to bring the database data sets up-to-date. With full database recovery, it may also be necessary to Batch Backout the inflight UOWs.



Prior to the restart of the remote IMS systems, the only data sets at the remote site are those that will be used for recovering the database data sets. This means there will either be just image copies and the backup DBRC RECON data set or there will be a combination of log data sets that may also include the RLDS and Change Accumulation data sets. The WADS and OLDS from the Production site will not be at the remote, nor will the RDS be at the remote site.

# Production Site

In this next section, we will focus in more detail on the responsibilities of the Production site.

## DBRC

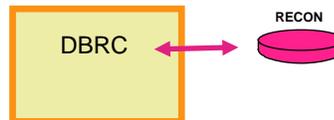


Skip to Page 17

- Skipping...
- Log Data Sets
- Log Utilities

- **DBRC Main Focuses**

- Control the availability of resources needed in recovery process
- Control database access
- GENJCL for various IMS Utilities
  - Data Base Recovery, Change Accumulation, Image Copy
  - Log Archive, Log Recovery
- RECON Data Sets
  - Keep track of status for the IMS System and IMS Database activity
  - DBRC's primary resource



IDUG 2008 North America

13

DBRC plays a key role in managing the log data needed to restart and recover IMS online subsystems. It controls the data sharing environment and simplifies the database recovery process. It is the primary IMS feature which makes recovery possible when IMS databases are registered in DBRC. The DBRC RECON data set is the primary resource used by DBRC. DBRC handles the following restart and recovery activities:

- Keeping status of the IMS System and database activity for Backup and Recovery
- Controlling access and activity to the IMS logs and IMS databases

The DBRC RECON data set is a Virtual Storage Access Method (VSAM) Key Sequenced Data Sets (KSDSs). Generally, users will have three RECON data sets to maintain dual-mode operation and also a spare RECON data set. If the databases are registered in the RECON data set, DBRC will also keep track of when the following utilities are executed:

- Database Image Copy, Online Database Image Copy, Database Image Copy 2
- Change Accumulation
- Database Recovery
- Log Recovery, Log Archive
- DEDB area data set create
- HD and HISAM Reorganization Unload and Reload
- HALDB Index/ILDS Rebuild

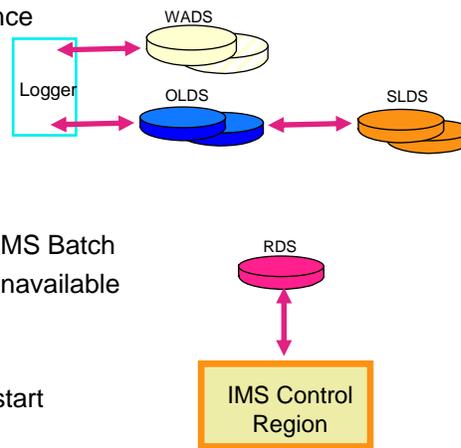
The DBRC GENJCL function can be used to create the JCL to Image Copy, create Change Accumulations or Recover database data sets. The information that is used is pulled from the RECON data set.

## IMS Log Management ...



### • IMS Log Data Sets

- Write Ahead Data Set (WADS)
  - Log Write Ahead for performance
- Online Log Data Set (OLDS)
  - Pool of 3 to 100 OLDS pairs
  - Before image / After image
  - Commit Status
- System Log Data Set (SLDS)
  - Archived OLDS or created by IMS Batch
  - Used in IMS Restart if OLDS unavailable
  - Database Recovery
- Restart Data Set (RDS)
  - Checkpoint information for Restart
  - *Can be recreated from the log*



IDUG 2008 North America

14

### Online log data sets (OLDS)

The OLDS are the data sets which contain all of the log records required for restart and recovery. These data sets must be pre-allocated (but need not be pre-formatted) on DASD and will hold the log records until they are archived. The OLDS are made up of multiple data sets which are used in a wrap-around manner. There must be at least three data sets allocated for the OLDS to allow IMS to start. There is an upper limit of 100 OLDS data sets supported. All OLDS should be dynamically allocated, by using the DFSMDA macro, and not hard-coded in the IMS control region JCL.

### Write ahead data sets (WADS)

The WADS is a small direct access data set which contains a copy of committed log records which are in OLDS buffers, but have not yet been written to the OLDS. When IMS processing requires the writing of a partially filled OLDS buffer, a portion of the buffer is written to the WADS data set. If IMS fails, ESTAE processing closes the OLDS using the current log buffers. If ESTAE is not successful, or if it is not invoked (for example, an MVS failure), then the WADS will normally be used by IMS Emergency Restart to close the OLDS. However, in the event of a disaster, the WADS and OLDS will be unavailable and will not be used for the restart of the remote site. The WADS space is continually reused after the appropriate log data has been written to the OLDS. This data set is required for all IMS systems, and must be pre-allocated and formatted at IMS start-up when first used.

### System log data sets (SLDS)

The SLDS is created by the IMS log archive utility, often after an OLDS switch. It can contain the data from one or more OLDS data sets. The SLDS can also be used as input to all IMS log utilities and to IMS restart. Information about SLDS is maintained by DBRC in the RECON data set. Calls to DBRC are made by the Archive Utility identifying the OLDS being archived and the SLDS being created. The OLDS that have been archived are then available for reuse by IMS. The SLDS data sets are also used by Batch Programs.

### Restart Data Set (RDS)

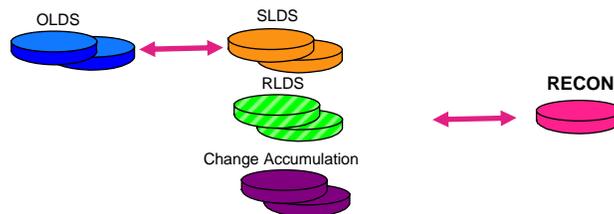
The RDS is a data set that contains information regarding Checkpoint IDs. This information is used during IMS Restart. It is faster to restart IMS from the information in the RDS, however, if it is unavailable, it can be recreated from the log data sets. For this reason, it is not necessary to ship the RDS to the remote site. In the event of a disaster, the remote site will find the Checkpoint ID information from the available SLDS data sets.

## IMS Log Management ...



- **IMS Log Data Sets**

- **Recovery Log Data Set (RLDS)**
  - An option of the Archive process
  - DB change information only for Database Recovery
- **Change Accumulation Data Set (CA)**
  - Contains condensed and sorted DB change records



### **Recovery Log Data Set (RLDS)**

When the IMS log archive utility is run, the user can request to have the creation of an output data set that contains all of the log records needed for database recovery. This is the RLDS and it is also recorded in the DBRC RECON data set. All database recoveries and change accumulation jobs will always use the RLDS if one exists. This speeds up these processes considerably because the RLDS contains only the database recovery log records. The IMS TM application scheduling and checkpoint log records are excluded from the RLDS's.

### **Change Accumulation Data Set (CA)**

The CA is a sequential data set that is used to streamline the recovery information needed for a database recovery. The information in the CA comes from the log data sets. It only contains database change records that have been sorted and merged for the same database physical record. In a data sharing environment, the CA data set is used to combine log data sets from several IMS subsystems. The production site records the CA data sets in the DBRC RECON data set so they can be shipped to the remote site. If they are not sent to the remote site, they need to be removed from the Backup RECON data set prior to creating the recovery jobs.

## IMS Log Management ...



- **Log Archive Utility (DFSUARC0)**
  - Archive one or several OLDSs to a single SLDS
  - Need “archive needed” status in DBRC
- **Log Recovery Utility (DFSULTR0)**
  - Correct errors in OLDSs, and batch or online SLDSs
  - Close (CLS) or Duplicate (DUP) option
- **Change Accumulation Utility (DFSUCUM0)**
  - Merge/Sort and Condense the DB updates for efficient recovery
  - Required in a data sharing environment
    - Unless using DRF

The **Log Archive utility (DFSUARC0)** produces an SLDS from a full OLDS or a batch IMS SLDS. The utility runs as an MVS batch job, and multiple log archive utility jobs can execute concurrently. When dual output is requested, the SLDS consists of primary and secondary data sets.

The **Log Recovery utility (DFSULTR0)** produces a usable log data set from a log data set that contains read errors or that was not properly terminated. After the log is corrected, it can be used by Change Accumulation, Batch Back out, and the Recovery utility. The Log Recovery utility can recover both OLDSs and batch or online SLDSs.

The **Change Accumulation utility (DFSUCUM0)** produces the Change Accumulation data set. It sorts and merges logs from one or more IMS subsystems to produce a data set that contains the log records needed for recovering a database. In a data sharing environment, it is required for merging logs from multiple IMS subsystems.

## DBRC RECON Backup



- **BACKUP.RECON**
  - Creates a copy of the RECON data set
  - Enqueue ensures no activity on RECON
- **Or, IMS Recovery Expert**
  - RECONCOPY command creates copy of RECON

**RECON data set must be same or later than Image Copy and Log data sets**

The **BACKUP.RECON** command creates a copy of the primary DBRC RECON data set. Before the copy is made, DBRC does an enqueue on the RECON data set to guarantee that no updates occur during the Backup.

The IMS Recovery Expert tool has a feature called **RECONCOPY** that creates a copy of the live DBRC RECON data set. It uses the **DBRC BACKUP.RECON RECON1** command.

As a rule, the backup of the DBRC RECON data set must be shipped at the same time as the Image Copy or SLDS log data set. There may be reasons why the RECON data set can not be shipped as often as the Image Copies or SLDS log data sets. However, to be useable, it can not be older than the Image Copy or SLDS log data set. For example, if the image copies and SLDS are from 1 am on Monday morning, then the RECON data set must be from after 1 am on Monday morning. It can not be from 12:30am on Monday morning.

## IMS Image Copy Solutions



- **Clean IC vs. Fuzzy IC vs. Clean Incremental IC**

- **Clean IC**

- Recovery may be done without logs
- Database is unavailable (/DBR or /DBD or /STA ACCESS RO)

- **Fuzzy IC**

- Recovery requires logs
- Database is available for update during the IC process



The Image Copy is the starting point for all database data set recoveries in IMS. There are several ways to take an image copy. Each option has benefits and some limitations. Besides choosing which method to use to take the image copy, the user must also choose whether the database will be quiesced during the image copy. When the database is taken offline for updating with a /DBR, /DBD or /STA ACCESS RO, then update activity is quiesced and the image copy is considered a *Clean Image Copy*. When a clean image copy is restored, the database data sets are consistent and do not need any additional log updates applied to them.

The *Fuzzy Image Copy* is an image copy taken when the databases are not quiesced, taken offline or changed for read only processing. A fuzzy image copy requires the logs to be processed for updates following the restore of the image copy.

The Image Copy utilities have a facility to make more than one Image Copy and each of these can be registered in the DBRC RECON data set. Generally, one of the secondary image copies are shipped to the remote site. However, it is not necessary to do this. The user can make their own copy of the Image Copy and ship this non-registered Image Copy to the remote site. At the remote site, it is possible to register this copy of the Image Copy with the DBRC RECON data set.

Generally, the indexes are not image copies and sent to the remote site. The indexes are rebuilt following recovery of the database data sets. The tool Index Builder is often used to rebuild the index. It is common to mark the secondary indexes as non-recoverable just for this situation.

## IMS Image Copy Solutions



Skip to Page 25

• Skipping...

• Image Copy Types

- **Clean Incremental IC**

- **Clean IC with IMS Recovery Expert, DRF and HPIC**
  - IMS RE Recovery Point Create (RPCR)
    - Quick /DBRs followed by /STAs
    - Creates Unallocation Point
  - HPIC creates Incremental Image Copy for DRF



IDUG 2008 North America

19

It is possible to create a Clean Image Copy with the IMS Recovery Expert tool without explicitly entering the /DBR, /DBD or /STA ACCESS RO commands. The IMS Recovery Expert Recovery Point Create (RPCR) feature can internally “bounce” one or more databases with /DBR or /DBD commands and then quickly restart them. This creates a clean unallocation point for these databases in the RECONS.

It is then possible to use the IMS Database Recovery Facility (DRF) feature called Incremental Image Copy (ICR) to create a Clean Image Copy up to this unallocation point in time just as if the database was down at the point the image copy was created. ICR uses the prior Image Copy as input and it will apply the necessary updates up to the unallocation point in time and this will create a Clean Image Copy.

- **Database Image Copy (DFSUDMP0)**
  - Clean and fuzzy image copies
  - Fuzzy IC also called “Concurrent Image Copy”
    - Fuzzy IC not valid for VSAM KSDSs
- **Online Database Image Copy (DFSUICP0)**
  - Fuzzy image copies
    - Updates allowed only in the same online system
      - Limited data sharing support
  - Runs in online system – special BMP
    - Uses online buffer pools – performance implications
  - Valid for all data set types: OSAM, VSAM ESDS and KSDS
  - Does not support Fast Path databases

The **Database Image Copy (DFSUDMP0)** utility may be used to copy any IMS database data set that supports standard recovery techniques. These are HDAM, HIDAM, HISAM, SHISAM, Index, PHDAM, PHIDAM, PSINDEX, and DEDB databases. The utility is run as a batch job and there is a Concurrent Image Copy (CIC) option which creates a fuzzy image copy. Updates are allowed while using the CIC option. The CIC option is not valid for KSDSs which are used for HISAM, SHISAM, Index, and PSINDEX database data sets.

The **Online Database Image Copy (DFSUICP0)** utility may be used to copy any full function database data set that may be copied with the Database Image Copy utility. However, it can not be used to copy Fast Path databases (DEDBs). This utility runs as part of an online system. It is a special type of BMP. It can create either clean or fuzzy image copies. There are no updates to KSDSs by other online systems or batch jobs allowed while this utility is running. Since the utility runs as part of the online system, it uses the online system’s database buffer pools and this tends to make it slower than the other utilities and causes it to interfere with other online work.

- **FP DEDB High-Speed Sequential Processing (HSSP)**

- Fast Path DEDBs only
- HSSP is an option for BMP application programs
  - High speed process to read the database
    - Anticipatory reads eliminate many read waits
- HSSP has an image copy option
  - Produces an image copy while application processes the area
  - Fuzzy image copy
    - Concurrent updates are allowed in any data sharing system

The **Fast Path DEDB High-Speed Sequential Processing (HSSP) Image Copy** is an option for non-message driven BMP application programs. It provides a high speed process to read DEDB areas which is done with anticipatory reads. The CIs are read before the application requests segments in them. HSSP is invoked by including H in the PROCOPT for a PCB. HSSP has an image copy option. HSSP image copies are extremely efficient. They are created as the application program is reading the data. A separate utility execution is not required.

HSSP uses its own private buffer pools which reduces the impact on the NBA/OBA buffers. It can be executed in both a mixed mode environment, concurrently with other programs, and in an IRLM global sharing environment. The image copy process can only be done if a database is registered with DBRC. In addition, image copy data sets must be initialized in DBRC. HSSP image copies can also be used for database recovery. However, the Database Recovery Utility must know that the image copy was created by HSSP.

- **Database Image Copy 2 Utility (DFSUDMT0)**

- Clean and fuzzy image copies
- Using DFSMS concurrent copy function (DUMP)
- IMS database data sets supported
  - HDAM, HIDAM, HISAM, DEDB, HALDB
  - OSAM, ESDS, KSDS
- Will produce 1-4 copies
  - 2 copies registered with DBRC at completion of DUMP
  - 2 copies available for shipment offsite
- Multiple ICs to single output data set
- With IMS 10, support for FlashCopy/SnapShot

The **Database Image Copy 2 (DFSUDMT0)** utility can copy any IMS database data set that can be copied with the Database Image Copy utility. It is also run as a batch job. IC2 invokes the concurrent copy capability of DFSMSdss DUMP. This requires that the database data set resides on a DASD system that supports Concurrent Copy. IC2 can create either a Clean or Fuzzy Image Copy. A fuzzy image copy may be taken for all IMS database data sets that support standard recovery techniques (including KSDSs).

In IMS 10, IC2 added support for FlashCopy and SnapShot which is invoked DFSMSdss COPY.

The Concurrent Copy function of DFSMS is a hardware and software solution that allows you to back up a database at a point in time and with minimum down time for the database. The database is unavailable only long enough for DFSMS to initialize a concurrent copy session for the data. This is generally a small fraction of the total time it takes to complete the backup of the data. Once the concurrent copy session has been established, the copy is said to be Logically Complete.

## IMS HP Image Copy Tool – DM Tools ...



- **Advanced Copy Services**

- For both Full Function DB (including HALDB) and DEDB
- Based on the concept of Logical Copy and Physical Copy
  - SnapShot support for RAMAC Virtual Array (RVA) devices
  - FlashCopy support
    - IBM TS Enterprise Storage Server (ESS) and IBM TS DS8000 devices (similar support for OEM devices)
- Implemented using DFSMSdss API to invoke DUMP and COPY
- Available in both batch and concurrent image copy processing

Like the IC2 support in IMS 10, the **IMS High Performance Image Copy (HPIC)** tool also supports Concurrent Copy using SnapShot and FlashCopy using the DFSMSdss API to invoke DUMP or COPY.

## IMS HP Image Copy Tool – DM Tools ...



- **Incremental Image Copy Support**

- With IMS Recovery Expert, DRF and HPIC
  - Image Copies are created without /DBR
- /DBR or /DBD commands are internally issued
- /STA database commands are internally issued
- Creates an “Unallocation Point”
- HPIC creates Clean Image Copy
  - Uses previous image copy + log updates to the Unallocation Point
- Incremental Image Copy is shipped to remote site
- DRF is required to restore Incremental Image Copy



IDUG 2008 North America

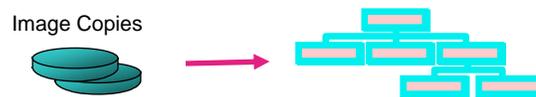
24

With the IMS Recovery Expert, IMS Database Recovery Facility (DRF), and IMS High Performance Image Copy (HPIC), a clean image copy can be created using the HPIC Incremental Image Copy feature without explicitly taking the IMS databases offline using /DBR or /DBD or /STA ACCESS RO commands.

IMS Recovery Expert creates an “Unallocation Point” for the databases in the RECON data set by internally issuing /DBR or /DBD commands followed quickly by /STA commands. When this point is created, the HPIC feature called “Incremental Image Copy” can run to create a Clean Image Copy using a previous image copy as input and the log updates from the time of the image copy to the time of the Unallocation Point. This Incremental Image Copy is a clean image copy that can be shipped to the remote site. At the remote site, it is necessary to use the DRF tool to restore an Incremental Image Copy.

## Recover the Databases to an Earlier Image Copy

- **Create a Clean Image Copy**
  - DBD or DBR the database and Image copy the database
- **Or, Create Incremental IC with RE, HPIC and DRF tool**
  - Incremental IC is a clean image copy
- **Ship Image Copy and RECON to remote site**

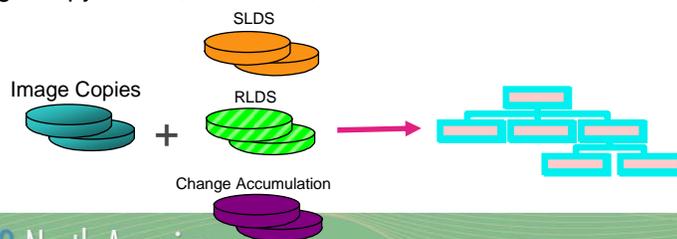


There are two disaster recovery solutions discussed in this presentation. The first is to recover the database data sets to an earlier image copy and not provide any forward recoveries using the logs. With this solution, only the image copies and the DBRC RECON data set are shipped to the remote site.

For this solution, the image copy should be a Clean Image Copy using the appropriate image copy tool or utility described in the previous slides. Alternatively, the image copy could be a HPIC Incremental Image Copy if the DRF tool is used at the remote site to restore the image copy. The advantage of using the HPIC Incremental Image Copy is that the IMS databases do not need to be taken offline (/DBR) or made unavailable for database updates (/DBD or /STA ACCESS RO) and it is still considered a Clean Image Copy.

## Recover Databases to Last Good Log Data Set

- **Create a Clean Image Copy**
  - DBD or DBR the database and Image copy the database
- **Or, Create an Incremental IC with RE, HPIC and DRF**
  - Incremental IC is a Clean Image Copy
- **Or, Create a Fuzzy Image Copy**
- **Ship to remote site**
  - Image Copy, SLDS, CA/RLDS, RECON data set to remote site



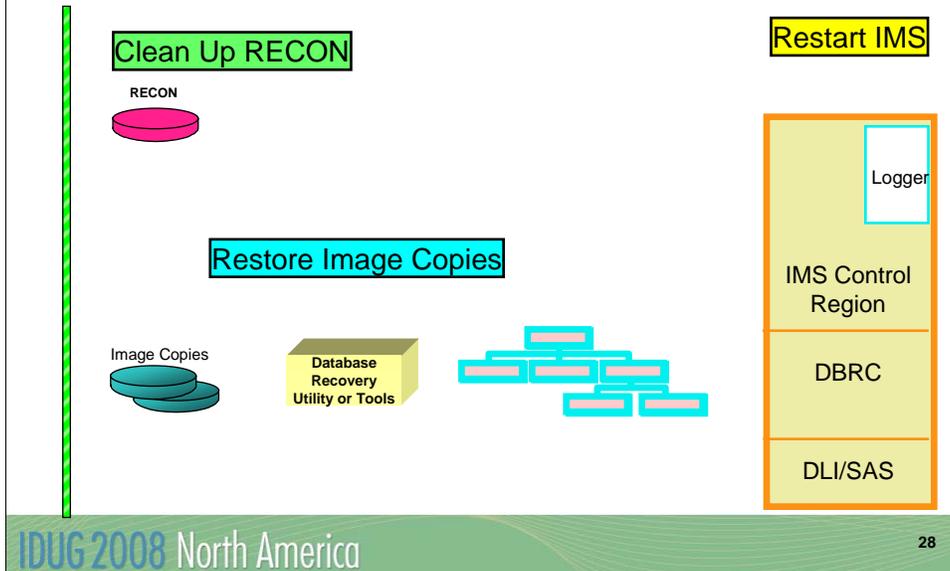
For the second disaster recovery strategy, the remote site database data sets are recovered using the Image Copy, the SLDS up to the last good log data set and the DBRC RECON data set. The image copy can be a Clean or Fuzzy image copy or it can be an HPIC Incremental Image Copy since the log data sets are shipped to the remote site and will be used to forward recover the database data sets. If RLDS and CA data sets exist and are shipped to the remote site, they can be used to recover the database data sets, too.

# Remote Site

In this next section, we will focus in more detail on the responsibilities of the Remote site.

## Recover Databases to Earlier Image Copy

## Remote Site



The first disaster recovery strategy is to recover the database data sets from an earlier image copy. The first step is to clean up the Backup RECON data set for recovering the database data sets. The second step is to restore the image copies. Since there are no log data sets at the remote site, there are no forward or backward recovery steps with this strategy. The final step is to restart the IMS subsystems at the remote site.

## Recover Databases to Earlier Image Copy



### • **Step 1: Restore and Clean the Backup RECON Data Set**

- Active Subsystems
  - If RECON shows active SUBSYSTEMs
    - Use LIST.SUBSYS to show active subsystems
    - Issue DBRC Commands:
      - CHANGE.SUBSYS SSID(ssidname) ABNORMAL
      - CHANGE.SUBSYS SSID(ssidname) STARTRCV
      - CHANGE.SUBSYS SSID(ssidname) ENDRECOV
      - DELETE.SUBSYS SSID(ssidname)
- Secondary Image Copies
  - If Image Copy at Remote site is a Secondary Image Copy
    - Flag the Primary Image Copy in the RECON as Invalid
- Fastpath DEDBs
  - Flag DEDBs AREAs as Recovery Needed to enable recoveries

The first step in this three-step process is to clean up the Backup RECON data set so that it can be used at the remote site. The RECON data set may show active subsystems and these need to be deleted since there are no active subsystems at the remote site. If the plan is to do a Emergency Restart, then the online IMS system SUBSYS record should not be deleted. If it is deleted, then a Cold Start is required. A LIST.SUBSYS command will show if there are active subsystems and there are several DBRC commands that will allow these to be deleted. They are:

- CHANGE.SUBSYS SSID(ssidname) ABNORMAL
- CHANGE.SUBSYS SSID(ssidname) STARTRCV
- CHANGE.SUBSYS SSID(ssidname) ENDRECOV
- DELETE.SUBSYS SSID(ssidname)

The second step is to recognize if the Image Copies at the remote site are secondary image copies and not the primary image copies. If this is the case, then you have to flag the primary image copy as Invalid. The command to do this is:

- CHANGE.IC DBD(dbname) DDN(ddname) INVALID

The third step is to flag the AREAs of the DEDBs (if any) as Recovery Needed so that they can be recovered by the recovery tools. The command for this is:

- CHANGE.DBDS DBD(dedbname) AREA(areaname) RECOV

## Recover Databases to Earlier Image Copy



- **Step 1: OR, IMS Recovery Expert Tool**

- RECON Clean Up (RCU) Feature
  - Input is “Clean Up Time”
    - Closes open PRILOG, PRIOLD and SECSLD records
    - Deletes PRIOLD, SECOLD and SUBSYS records
    - Updates or deletes ALLOC and LOGALL records
    - Deletes IC and CA records past the Clean Up Time
    - And many other tasks
  - Health Check (HCHECK) Feature
    - Identifies issues that could affect database recovery
  - Verify (VERIFY) Feature
    - Ensures data sets exist prior to performing database recovery

The **IMS Recovery Expert** tool can be used as an alternative to manually cleaning up the DBRC RECON data set. The RECON Clean Up (RCU) function takes a “clean up time” as input to the RCU function. All indications of IMS activity equal to or later than the clean up time are removed from the RECON data set. The RECON is not modified to appear that all activity ended normally at the clean up time. Instead, it appears as if time simply stopped at the clean up time.

The following list of items are just a subset of what the IMS Recovery Expert tool can do during RCU:

- Close open PRILOG, PRISLD and SECSLD, records
- Delete PRIOLD, SECOLD and SUBSYS records
- Update or delete ALLOC and LOGALL records
- Delete IC and CA records that are past the RECON clean up time
- Mark CA runs invalid if log data spans past the RECON clean up time
- Mark DBs as Recov Needed if they are affected by RCU clean up
- Increment Recovery Needed Counts if marked for recovery

Once the RECON data sets are cleaned up, the IMS Recovery Expert Health Checker feature working together with the VERIFY feature can ensure that all of the data sets needed for database recovery are available at the remote site.

## Recover Databases to Earlier Image Copy...



### • **Step 2: Recover the Databases from Image Copies**

- Use Cleaned Up RECON to GENJCL.RECOV Recovery JCL
- Recover database data sets:
  - Standard IMS Recovery Utility (DFSURDB0)
  - Or, IMS Database Recovery Facility (DRF)
- HPIC Incremental Image Copies
  - Must be recovered with IMS Database Recovery Facility (DRF)
- HPIC COPY Image Copies
  - Exact copy of database, no recovery is needed
  - Ensure Database Data Set name matches



After the RECON data sets have been cleaned up at the Remote site, the database data sets can be recovered. The JCL to create do this can be created using the RECON data set using the DBRC GENJCL.RECOV command. The RECON data set knows all of the relevant image copies.

To recover the database data sets, the standard IMS Recovery Utility (DFSURDB0) can be used or the IMS Database Recovery Facility (DRF). If the IMS High Performance Image Copy (HPIC) tool was used to create Incremental Image Copies, then the DRF tool must be used to restore them at the remote site. The HPIC tool offers another option which is the DFSMSdss COPY function. This creates an exact copy of the database and no additional processing is required to recover it. The user just needs to ensure the database data set name of this copy is correct at the remote site.

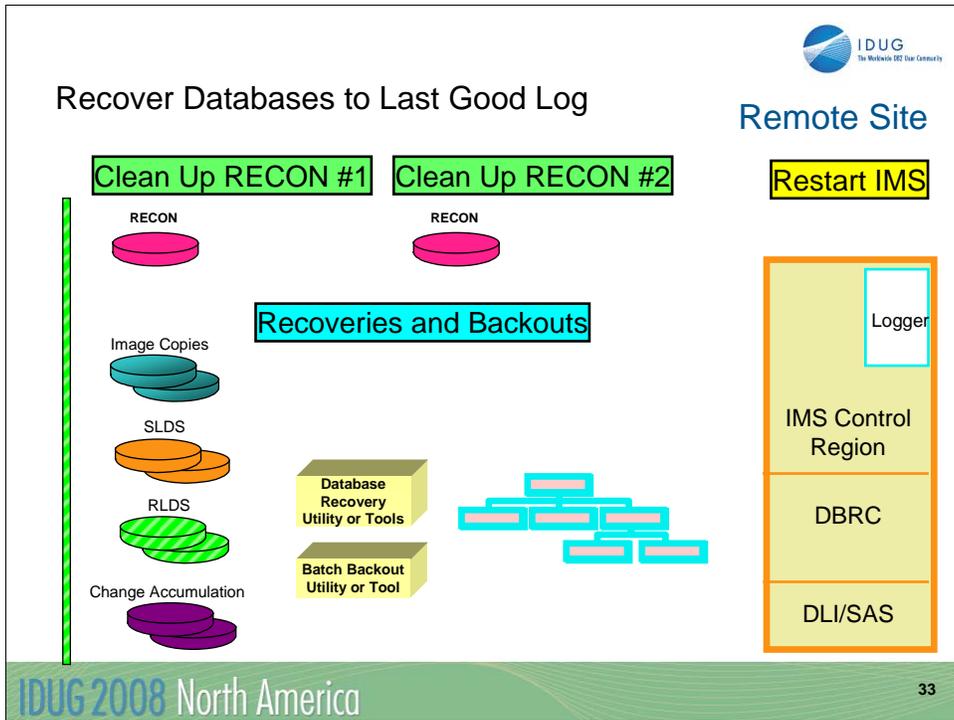
## Recover Databases to Earlier Image Copy...



- **Step 3: Cold Start Online and Batch Systems**

- Databases are consistent with earlier Image Copy

Finally, the IMS and Batch subsystems can be restarted at the remote site. The restart will be a Cold Start. The data lost in this Disaster Recovery strategy is the data created since the time the image copy was shipped to the remote site.



The second disaster recovery strategy is to recover the database data sets to the last good log data set. The first step is to clean up the Backup RECON data set for recovering the databases. The second step is to restore the image copies and do forward recovery with the log data sets. The third step is to do some clean up on the Backup RECON data set for restarting IMS. The final step is to restart the IMS subsystems at the remote site.

## Recover Databases to Last Good Log



- **Step 1A: Restore and Clean Backup RECON data set**
  - To use DBRC RECON for Recovery
    - No Databases can be Allocated on an OPEN log
  - Close and Archive OLDS in RECON data
    - When OLDS are not at the Remote Site
    - All PRILOG, PRISLD and PRIOLD need non-zero Stop Times
  - Two Methods:
    1. Add dummy SLDS log entry and Start Archive
    2. Close and flag INUSE OLDS as Archive Needed and Start Archive

In order to use DBRC and the RECON data set for building the recovery JCL, the Backup RECON data set needs to have some initial cleaning. There can be no databases allocated on an open log data set. The OLDS data sets will not be available at the remote site. When the OLDS is still Open, the PRILOG, PRISLD and PRIOLD records in the RECON data set all have valid Start times, but their Stop times are all zeros. When the OLDS are closed and archived, the Stop Times of these three records are set to non-zero values. The Log Recovery Utility (DFSULTR0) will not put the non-zero Stop times into these three records because it is expecting an Archive to be run which will handle the Stop times.

There are two methods for closing and archiving the OLDS log data sets in the RECON data set which ultimately creates the non-zero Stop times in the PRILOG, PRISLD and PRIOLD records. The first method is to add a dummy SLDS record which is associated with a VB null data set. This can be done with the following DBRC commands:

- NOTIFY.PRILOG SLDS SSID(..) STARTIME(starttime) DSN(....) -VOLSER(...)
- NOTIFY.PRILOG SLDS SSID(..) STARTIME(starttime) -RUNTIME(starttime+1)

When the Archive utility is run, it will find there are no updates in this dummy record and will put non-zero Stop times in the PRILOG, PRISLD and PRIOLD records.

The second method is to identify the active OLDS data set in the RECON data set. This can be done with the following DBRC command:

- %SELECT OLDS INUSE

This OLDS data set (ex. DFSOLP01) is flagged as ACTIVE and has a Stop time of zeros. This OLDS should be closed and flagged as ARCHIVE NEEDED. This can be done with the DBRC command:

- NOTIFY.PRILOG OLDS(oldsdd) SSID(newssid) -STARTIME(start) RUNTIME(start+1)

Prior to running the Archive utility, an empty OLDS data set should be allocated. When the Archive utility is run, it will find the OLDS is empty and it will put non-zero Stop times in the PRILOG, PRISLD and PRIOLD records.

When the PRILOG, PRISLD and PRIOLD records all have non-zero Stop times, DBRC can be used to create the recovery JCL.

## Recover Databases to Last Good Log



- **Step 1B: Continue to Clean the Backup RECON data set**

- Active Subsystems
  - If RECON shows active SUBSYSTEMs
    - Use LIST.SUBSYS to show active subsystems
    - Issue DBRC Commands:
      - CHANGE.SUBSYS ABNORMAL
      - CHANGE.SUBSYS STARTRCV
      - CHANGE.SUBSYS ENDRECOV
      - DELETE.SUBSYS
- Secondary Image Copies
  - If Image Copy at Remote site is a Secondary Image Copy
    - Flag the Primary Image Copy in the RECON as Invalid
- Fastpath DEDBs
  - Flag DEDBs AREAs as Recovery Needed to enable recoveries

The next step in the DBRC RECON cleanup is to delete any active subsystems. The DBRC command LIST.SUBSYS will show if there are active subsystems and there are several DBRC commands to delete them:

- CHANGE.SUBSYS SSID(ssidname) ABNORMAL
- CHANGE.SUBSYS SSID(ssidname) STARTRCV
- CHANGE.SUBSYS SSID(ssidname) ENDRECOV
- DELETE.SUBSYS SSID(ssidname)

The next step is to recognize if the Image Copies at the remote site are secondary image copies and not the primary image copies. If this is the case, then you have to flag the primary image copy as Invalid. The command to do this is:

- CHANGE.IC DBD(dbname) DDN(ddname) INVALID

The next step is to flag the DEDB AREAs (if any) as Recovery Needed so that they can be recovered by the recovery tools. The command for this is:

- CHANGE.DBDS DBD(dedbname) AREA(areaname) RECOV

## Recover Databases to Last Good Log



- **Step 1C: Continue to Clean the Backup RECON data set**
  - If copies of Change Accumulation data sets are at remote site
    - RECON will show CA data sets from Production site
    - Use CHANGE.CA to point to copy of CA data set at remote site
  - If CA data sets are not at remote site, but are in the RECON
    - Flag the RECON to show CA is INVALID
      - DBRC will use the logs instead of the CA data sets

The next step in cleaning the DBRC RECON data set at the remote site is to determine if the primary site was using Change Accumulation (CA) data sets and if these data sets were sent to the remote site. If there are CA data sets in the Backup RECON data set, they will need to be changed to point at the copy of the CA data set at the remote site. The DBRC command to do this is:

•CHANGE.CA CADSN(cacopydatasetname)

Alternatively, the CA data sets in the RECON data set can be flagged as INVALID, in which case, DBRC will not use the CA data sets for recovery and it will only use the available log data sets. The DBRC command to invalidate a CA data set is:

•CHANGE.CA CADSN(datasetname) INVALID

## Recover Databases to Last Good Log



- **Step 1: OR, IMS Recovery Expert Tool**

- RECON Clean Up (RCU) Feature
  - Input is “Clean Up Time”
    - Closes open PRILOG, PRIOLD and SECSLD records
    - Deletes PRIOLD, SECOLD and SUBSYS records
    - Updates or deletes ALLOC and LOGALL records
    - Deletes IC and CA records past the Clean Up Time
    - And many other tasks
  - Health Check Feature
    - Identifies issues that could affect database recovery
  - Verify Feature
    - Ensures data sets exist prior to performing database recovery

The **IMS Recovery Expert** tool can be used as an alternative to manually cleaning up the DBRC RECON data set. The RECON Clean Up (RCU) function takes a “clean up time” as input to the RCU function. All indications of IMS activity equal to or later than the clean up time are removed from the RECON data set. The RECON is not modified to appear that all activity ended normally at the clean up time. Instead, it appears as if time simply stopped at the clean up time.

The following list of items are just a subset of what the IMS Recovery Expert tool can do during RCU:

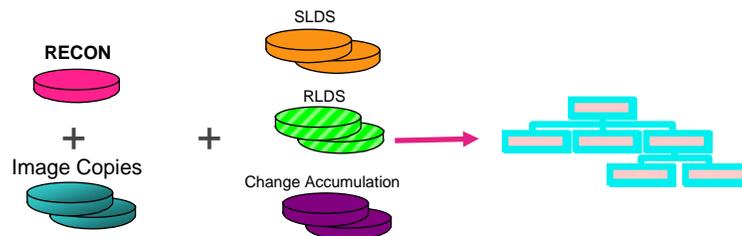
- Close open PRILOG, PRISLD and SECSLD, records
- Delete PRIOLD, SECOLD and SUBSYS records
- Update or delete ALLOC and LOGALL records
- Delete IC and CA records that are past the RECON clean up time
- Mark CA runs invalid if log data spans past the RECON clean up time
- Mark DBs as Recov Needed if they are affected by RCU clean up
- Increment Recovery Needed Counts if marked for recovery

Once the RECON data sets are cleaned up, the IMS Recovery Expert Health Checker feature working together with the VERIFY feature can ensure that all of the data sets needed for database recovery are available at the remote site.

## Recover Databases to Last Good Use

Skip to Page 43

- **Step 2: Recover the Databases from**
  - Use Cleaned Up RECON to GENJCL.RECOV
  - Recover database data sets from Image Copies and Logs:
    - Standard IMS Recovery Utility (DFSURDB0)
    - Or, IMS Database Recovery Facility (DRF)
      - HPIC Incremental Image Copies
        - Must be recovered with IMS Database Recovery Facility (DRF)



Once the Backup RECON data set has had some initial cleaning, it can be used to create the recovery JCL using the DBRC command GENJCL.RECOV. The recovery JCL will include the image copies, the log data sets and it may also include the RLDS and CA data sets if they exist at the remote site. The standard IMS Recovery Utility (DFSURDB0) or the IMS Database Recovery Facility (DRF) can be used to recover the database data sets. If the IMS High Performance Image Copy (HPIC) tool was used to create an Incremental Image Copy, then the DRF tool must be used to recover it.

## • Step 2: Option 1: Full Database Recovery

- Uses Image Copy + Log records
- Supported by DBRC



DBRC supports two types of database recovery for DBDSs and area data sets (ADSs):

- **Full Recovery** recovers a DBDS or ADS to its current state as recorded in the RECON data set.
- **Time-stamp Recovery** recovers a DBDS or ADS to the contents it had at the time you specify for the recovery.

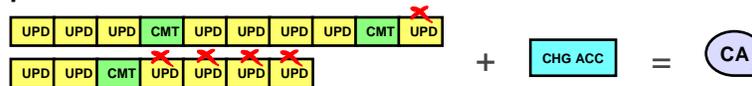
You perform a forward recovery of a database data set by restoring a backup copy of the database and then reapplying all of the logged database changes that occurred after the backup copy was made. While these two basic steps are common to all forward recoveries of databases, the specific steps of a forward recovery can vary depending on the type of database you are recovering and whether you are operating in a data-sharing environment or a non-data-sharing environment.

A backup copy of an IMS database is usually an image copy of the database data set. Image copies are created by one of the IMS image copy utilities or a separate image copy tool. The IMS image copy utilities automatically register the image copies they create with the Database Recovery Control facility (DBRC). Image copies created by some other means might need to be registered with DBRC before DBRC and the Database Recovery utility (DFSURDB0) can use them for recovery.

## • Step 2: Option 2: Full DB Recovery with Data Sharing

- Multiple concurrent logs
- Merge updates on logs using Change Accumulation (CA)
- Uses: Last image copy + change accumulation + 2 backouts
- Supported by DBRC

### Step 1



### Step 2



Continued...

The logs that contain the database changes can be processed before recovery by the Change Accumulation utility (DFSUCUM0), to consolidate all of the database change records and optimize them for recovery. In a data-sharing environment, you are required to use the Change Accumulation utility unless the full database recovery is being done by the IMS Database Recovery Facility (DRF). In a non-data-sharing environment, using the Change Accumulation utility is optional, although it can be beneficial to use for performance reasons.

DBRC manages the forward recovery process and can generate the required JCL for all of the backup and recovery related utilities, except in cases in which the database being recovered is not registered with DBRC. The JCL that is generated by the DBRC command GENJCL.RECOV for the Database Recovery utility automatically identifies and includes the correct image copies, logs, DBD names, ddnames and timestamps.

The process for recovering the database data sets, that is, the data sets that contain the database records, is the same across most database types, including full-function, HALDB partitioned full-function, and Fast Path DEDB database types. For HALDB databases, however, indirect list data sets and the PHIDAM primary indexes are not backed up or recovered; instead, after the database data sets they support are recovered, HALDB primary indexes and indirect list data sets (ILDSs) are typically rebuilt by using the HALDB Index/ILDS Rebuild utility (DFSPREC0). You recover a primary index of a HIDAM database in the same way as you recover a database data set: using image copies and logs.

Another difference about recovering HALDB databases is that the DBRC command GENJCL.RECOV can generate JCL to recover all of the partitions of the HALDB database at once or to recover only a single partition. The GENJCL.RECOV command also provides a similar option for recovering entire Fast Path DEDB databases or individual DEDB areas.

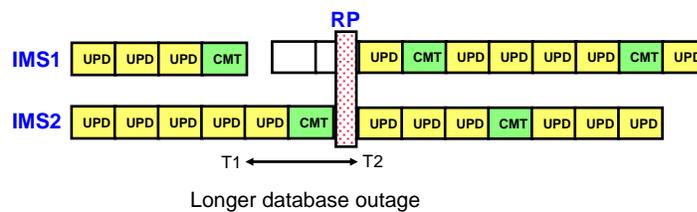


## Recover Database to Last Good Log



### • Step 2: Option 4: Timestamp Recovery with Data Sharing

- Change Accumulation needed to merge logs from multiple systems
- Recovery point must be created on each IMS
  - /DBR database on all systems before /START on any system
    - All ALLOC records must be closed at the same time
- Recover to point when database not allocated to any IMS
- Supported by DBRC



With Timestamp Recovery where the primary site was executing in a data sharing environment, Change Accumulation will be required to merge the updates from multiple log data sets. In this environment, it is necessary to create Recovery Points (RP) on all systems at the same time. That is, each system must have done the /DBR and all ALLOC records should be closed during that time before any /STA commands are issued. In this way, the databases can be recovered to a point when they were not allocated to any IMS system in the data sharing environment. If the plan is to recover the remote site to a RP, then the production site should try to create a consistent RP across all databases and IMS subsystems.

## Recover Database to Last Good Log



Skip to Page 45

• Skipping...

• Last Recovery Type

- **Step 2: Option 5: IMS Database Reco**
  - Recovery Options
    - Full Database Recover to End of Logs
    - Timestamp Recovery (Recovery Point)
    - Point-In-Time Recovery (Arbitrary Point)
  - Recovery of Multiple DBDS or Areas in one step
    - Parallel Input/Output
    - Parallel Processing / Sorting
    - Single archived Log and CA Data Set Pass
    - Single pass DB write
    - DBRC Controlled

The IMS Database Recovery Facility (DRF) provides a significant boost to performance in recovering database data sets over the Database Recovery Utility. DRF can provide all three types of recovery: 1) Full database recovery to the end of the log, 2) Timestamp recovery to a specific Recovery Point (RP), and 3) Point-In-Time Recovery to an arbitrary point in time.

The tool performs well because it is able to process data sets in parallel. With DRF, multiple database data sets can be processed at the same time. The log data sets are processed in a single pass. This parallelism separates this tool from the Database Recovery Utility which processes one database data set at a time.

## Recover Database to Last Good Log



### • Step 2: Option 6: Point-In-Time Recovery

- IMS Database Recovery Facility (DRF) Required
- Recover IMS Databases to any timestamp
  - Timestamp may be in middle of UOW
  - Only committed updates applied
- Could coordinate recovery with DB2 Quiesce Point



With the IMS Database Recovery Facility (DRF), it is possible to do a **Point-In-Time Recovery** which is a recovery to any prior point in time. The user determines a timestamp without being restricted by allocation ranges or Recovery Points. All updates that are committed as of the specified timestamp are applied. The data that is not committed at that time will not be part of the recovery.

At the recovery time specified, database data sets and areas can be at any state as long as an image copy and log data are available. That is, there is no restriction on database allocation status if timestamp recovery with PITR is selected. All committed updates up to and including the specified recovery time are applied to the database data sets and areas in the recovery list.

The database data sets and areas are not made available for processing by other IMS subsystems or batch jobs until they have been successfully image copied and the image copies registered with DBRC.

With the Point-In-Time Recovery option, it is possible to coordinate the recovery with DB2 which also has a Point-In-Time Recovery function.

## Recover Database to Last Good Log



- **Step 3: Batch Backout**
  - **Batch Backout Utility (DFSBB000)**
    - Following Full Database Recovery
      - Backout Inflight UOWs
      - Cold Start or ERE COLDSYS or ERE COLDBASE
    - Following Timestamp (RP) or Point-In-Time Recovery
      - No Inflight UOWs to backout
    - Cold Start

The next step is to do run the Batch Backout (BBO) Utility (DFSBB000) if it is needed. BBO is required if there are inflight UOWs on the log and full database recovery is used to recover the database data sets. After BBO is done, IMS can be restarted with a Cold Start or an Emergency Restart with the COLDSYS or COLDBASE option.

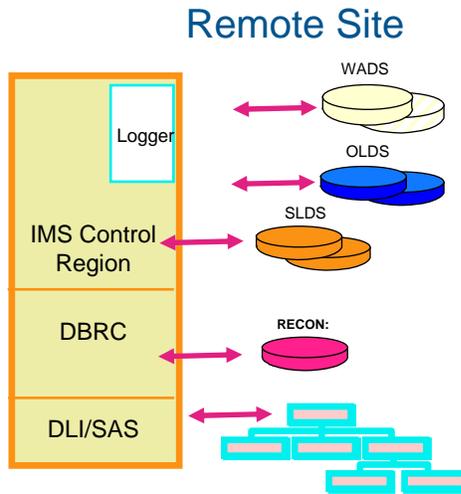
For Timestamp Recovery to a Recovery Point (RP) or Point-In-Time Recovery using the DRF Tool, BBO is not needed because there will not be any inflight UOWs at the end of the log to back out. A Cold Start can be run after these latter two types of recovery.

## Recover Database to Last Good Log



- **Step 4: Restart IMS at Remote Site**

- Option 1: Cold Start
- Option 2: ERE from SLDS
  - Clean RECONS #2
  - Restart is from SLDS



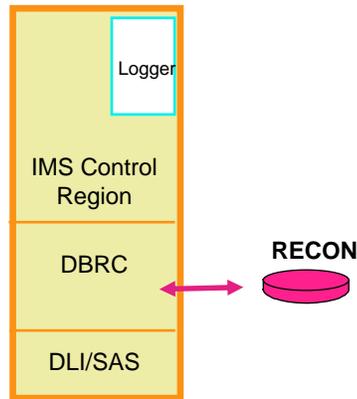
There are two Restart options for the remote site for the disaster recovery strategy of recovering the databases to the last good log data set. The first is a Cold Start. The second is to do some further clean up of the RECON data sets and then do an Emergency Restart from the SLDS.

## Recover Database to Last Good Log



- **Step 4: Option 1: Cold Start**

- Following Recoveries and Batch Backout Utility
- Cold Start



Once the database data sets have been restored using the image copies, forward recovery and batch backout, the IMS subsystems can be restarted using a Cold Start.

## Recover Databases to Last Good Log



- **Step 4: Option 2: ERE from SLDS and Existing RECONS**

- Ensure SUBSYS records have been cleaned up
  - Online IMS SUBSYS record should not be deleted
- Force IMS to restart from the SLDS
  - Last OLDS are not available due to Production Site Failure
  - Must delete the PRIOLD and SECOLD in RECON
    - NOTIFY.PRILOG OLDS(..) STARTTIME(..) RUNTIME(..) SSID(..)
    - CHANGE.PRILOG OLDS(..) SSID(..) ARCHIVED
    - DELETE.LOG OLDS(..) SSID(..)
    - DELETE.LOG OLDS(..) SSID(..) LASTCLOS
- IMS will dynamically allocate SLDS
  - Needs DFSMDA macro definition

The second restart option is to do an Emergency Restart from the SLDS data sets using the existing RECON data sets. In order to do this, some further cleaning is required on the RECON data set. In Step 1B above, the DBRC SUBSYS records were deleted from the Backup RECON data set for non-online IMS systems.

Now, the RECON needs to have some further cleaning to allow the ERE to be done from the SLDS. Since the OLDS are not available at the remote site, the PRIOLD and SECOLD records must be deleted from the Backup RECON data set. The following DBRC commands can be used to do this:

- NOTIFY.PRILOG OLDS(oldsname) STARTIME(starttime) RUNTIME(runtime) SSID(ssidname) LASTREC(lastrec)
- CHANGE.PRILOG OLDS(oldsname) SSID(ssidname) ARCHIVED
- DELETE.LOG OLDS(oldsname) SSID(ssidname)
- DELETE.LOG OLDS(oldsname) SSID(ssidname) LASTCLOS

The NOTIFY.PRILOG command closes the last open OLDS data set. The CHANGE.PRILOG flags the OLDS data set as ARCHIVED. The first DELETE.LOG deletes all of the PRIOLD and SECOLD records except the last one. The second DELETE.LOG deletes the last PRIOLD and SECOLD record.

Once the PRIOLD and SECOLD records are deleted in the DBRC RECON data set, IMS will dynamically allocate the last SLDS data set for restart. There must be a DFSMDA dynamic macro definition set up ahead of time for the SLDS data set to allow this to happen.

## Summary



- **Two Disaster Recovery Strategies**
  - Recover to Earlier Image Copy
  - Recover to Last Good Log
- **Choice of Strategy**
  - Determines amount of data lost and time to recovery at Remote Site
- **Strategy Determines Data Sent to Remote Site**
  - Image Copies, Change Accumulations, RLDS, Logs, RECONs
  - Type of recovery (image copy vs. log recovery/backout)
- **DBRC**
  - Crucial element of Disaster Recovery
  - RECON needs Cleaning at the Remote Site

In this presentation, we discussed two Disaster Recovery strategies. The first was to ship only the image copies and DBRC RECON data set to the remote site and in the event of a disaster at the primary site, the recovery of the database data sets will be back to the time of the image copy. The second strategy was to ship image copies, SLDS log data sets, the DBRC RECON data set and Change Accumulation and RLDS data sets if they exist to the remote site. In the event of a disaster, the remote site would do forward recovery and backout using the image copies, the logs and the RECON data set.

The choice of disaster recovery strategy is dependent on the amount of data that can be lost (RPO) and the amount of time to recover and restart the remote site (RTO). For the first strategy the RPO is high and the RTO is low. For the second strategy, the RPO is low and the RTO is slightly higher because the remote site must perform forward and backward recoveries.

DBRC is crucial for determining what data sets need to be involved in recovering the databases at the remote site. However, there is some cleaning required of the Backup RECON data set before it can be used at the remote site.