



# Session B76

IMS DM Tools Recovery Solution - Database Recovery Facility:  
Features, Configuration, Usage & New Features

Allen Hall

**IMS**  
**Technical Conference**

**Sept. 27-30, 2004**

**Orlando, FL**



# Agenda

- IMS Database Recovery Options
- IMS Database Recovery Services - overview
- Installing DRF V2R1
- Using DRF V2R1
- What's new in DRF V2R1

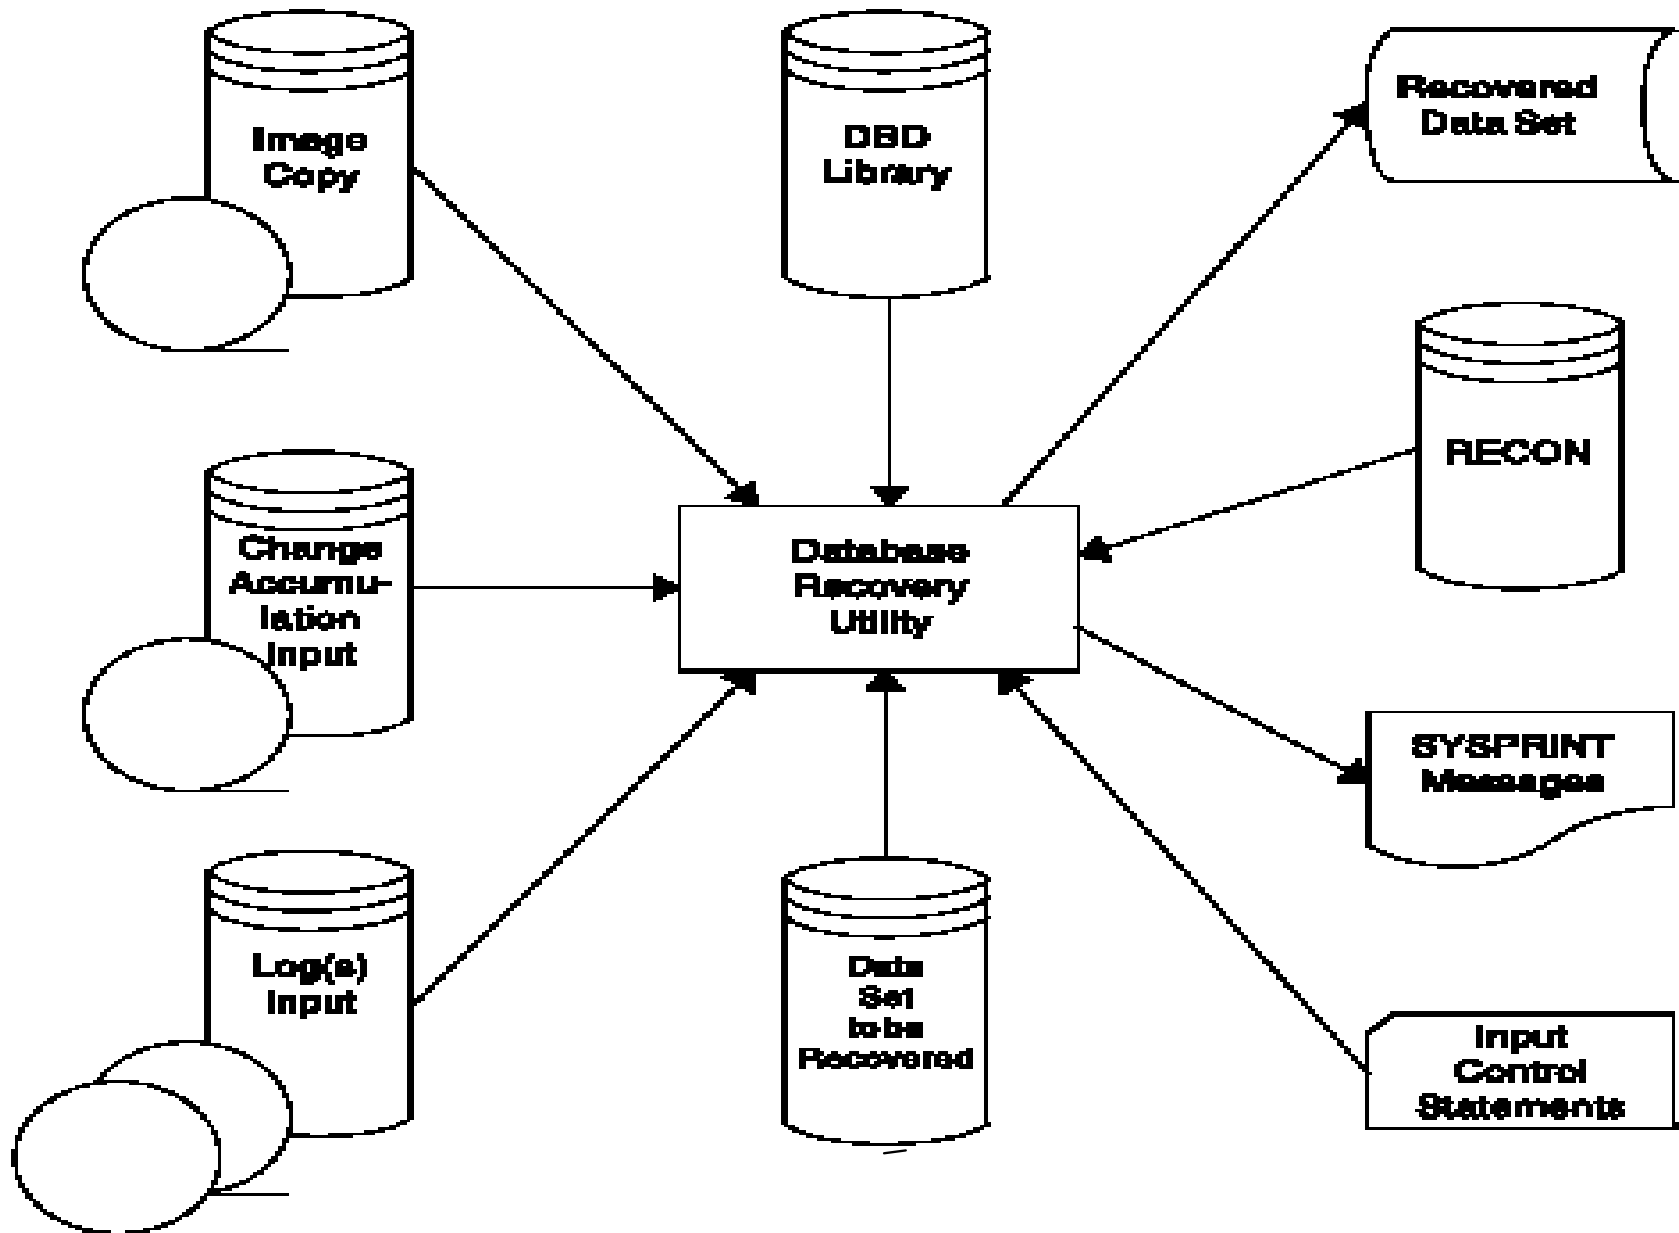
- IBM has several options for recovering IMS databases.
- DRF V2R1 is the most comprehensive of those options and is part of the IMS Database Recovery Services along with Online Recovery Services.
- We will take a look at how to install and use DRF V2R1.
- Then we will take a look at what features were added in DRF V2R1.

# IMS Database Recovery Options

- IMS standard database recovery utility (DFSURDB0)
- HP Image Copy extends DFSURDB0 functions
- IMS Database Recovery Services:
  - ▶ IMS Online Recovery Service for z/OS (ORS) - V1
  - ▶ IMS Database Recovery Facility (DRF) for z/OS - V2R1

- DFSURDB0 is provided with IMS. It does not allow for dynamic allocation of input data sets.
- IMS HPIC provides the same functions of DFSURDB0. This is accomplished by HPIC having a module that has the alias name DFSURDB0 and what HPIC calls 'IMS Compatible JCL.' The only difference between standard IMS DFSURDB0 JCL and HPIC's IMS Compatible JCL is the STEPLIB DD. HPIC extends the functions of DFSURDB0 by accepting compressed IC input and doing dynamic allocation, based on what is in the RECONS of the DBDS, IC, Logs & CA data sets. HPIC also checks the RECONS and can, by option, use the most recent batch IC taken. To use the extended properties of the HPIC recovery function, you must execute FABJMAIN instead of the alias for DFSURDB0.
- Neither DFSURDB0 nor HPIC can recover more than one database per execution.
- DRF V2R1 is version 2 of the product originally named Online Recovery Services for Z/OS V1. ORS was invoked only on-line with the /RECOVER command. DRF V2R1 includes ORS and extends it to allow for batch invocation; therefore the original product name no longer fits. The DRF V2R1 manual contains the information that was in the ORS V1 manual and the new and changed functionality.
- DRF can recover many databases in parallel in one execution and can accept many different types of IC input. It can also perform dynamic allocation of input data sets.

# DFSURDB0 data flow diagram

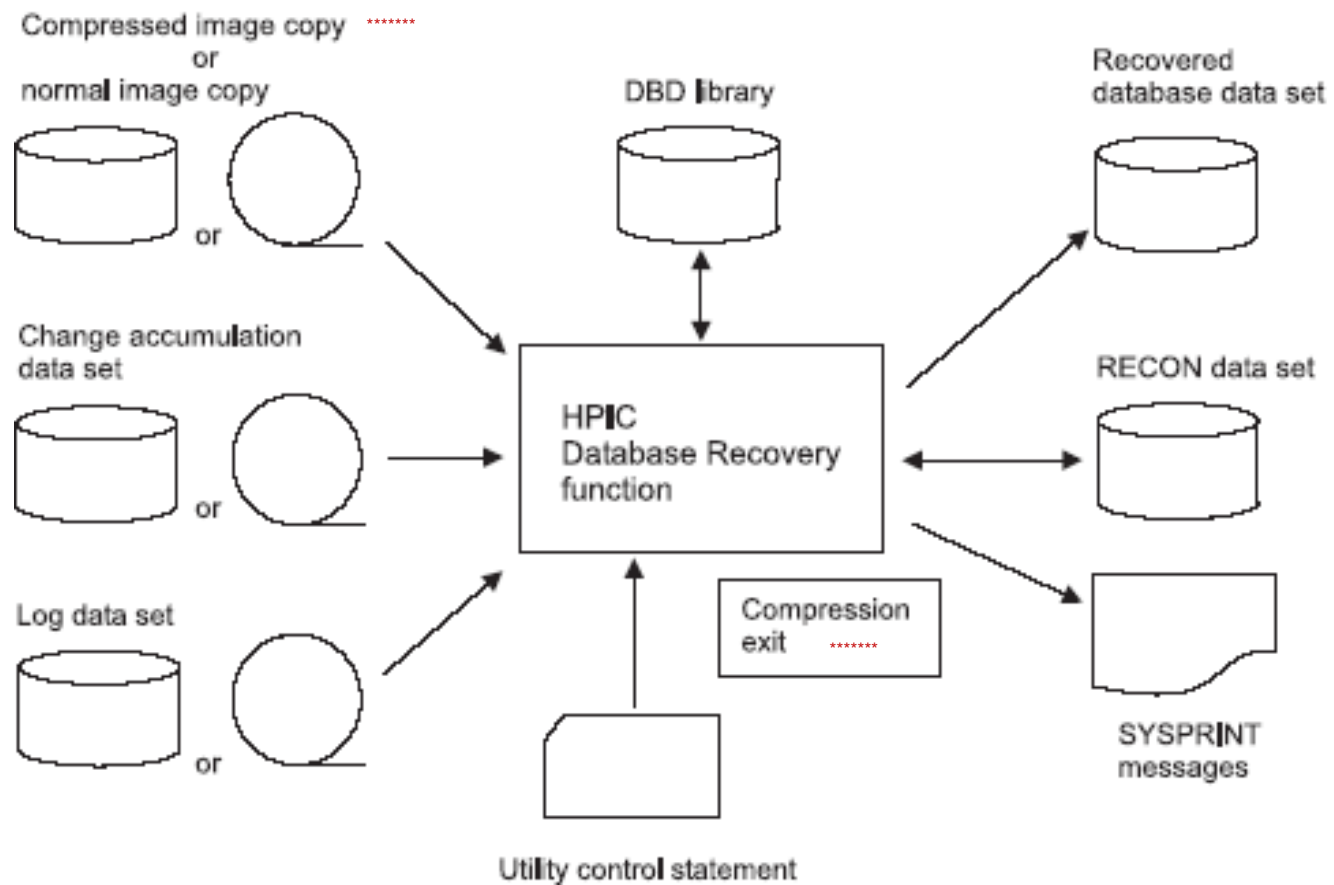


DFSURDB0 is executed as a special recovery region using the following EXEC statement:

```
//STEP EXEC PGM=DFSRRRC00,PARM='UDR,DFSURDB0,dbdname'
```

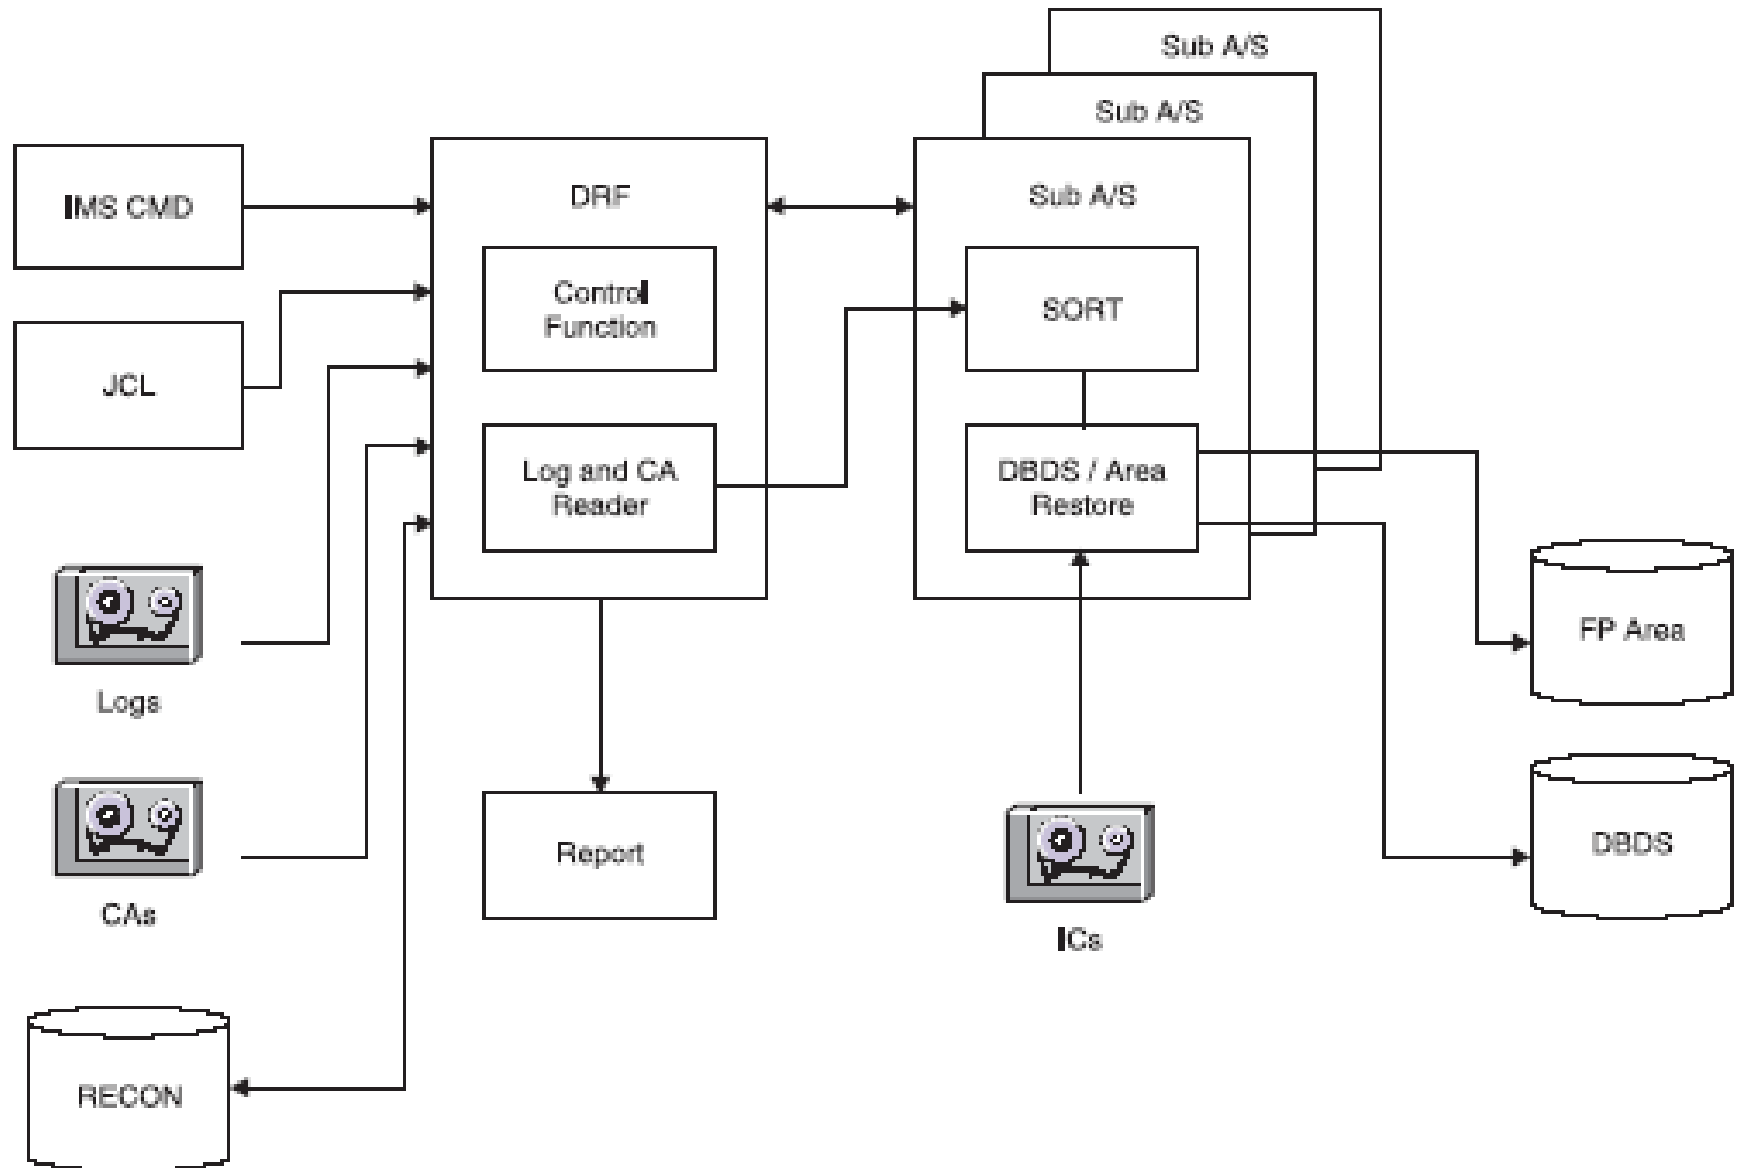


# HPIC recovery function data flow



- If the IC was created by HPIC using a compression routine (COMP=Y), then the same compression routine is automatically called by the HPIC Database Recovery function to decompress the compressed data.
- HPIC updates the RECONs automatically with a NOTIFY.RECOV command, as DFSURDB0 does.

# DRF data flow & process diagram



- Logs and Change Accumulation data sets are read only once by the DRF Master address space (AS). Pertinent records are sent to the subordinate address spaces.
- Subordinate address spaces read the appropriate image copies and apply the date-time-sorted Logs and CA records directly to the IC records resulting in updated database records. CAs are created by accumulating the database updates since the last image copy and they are known to DBRC.
- Multiple database data sets are recovered in parallel. You have control over the number of subordinate address spaces available to perform recovery plus the number of mountable input devices that are used to read the recovery assets.

# IMS Database Recovery Services

- Version 1 - IMS Online Recovery Service for z/OS
  - ▶ License Product ID: 5655-E50
  - ▶ General Availability: April 2001
- Version 2 - IMS Database Recovery Facility for z/OS
  - ▶ License Product ID: 5655-I44
  - ▶ General Availability: July 2003



# Database Recovery Services Objectives

- High Performance Recovery
  - ▶ Parallel I/O in DRF Master (Logs + CAs) and Subspaces (ICs)
  - ▶ Parallel Recovery Processing of Multiple DBDS or Areas
  - ▶ Single Log and CA Data Set Pass in DRF Master Address Space
- Ease of Use
  - ▶ DBRC Actuated
  - ▶ Minimal JCL - Dynamic Allocation
  - ▶ Results Recorded in RECON

- High Performance Recovery
  - ▶ The READNUM parameter in the FRXDRxx PROCLIB member limits the number of devices used to read in parallel the Logs, CAs, ICs.
  - ▶ The number of subspaces active to sort log records, read ICs and write DBDSs is limited by the NUM(nn) parameter.
  - ▶ The DRF Master AS passes Log and CA records to subspaces where the Log records are sorted and applied along with the CAs to ICs to create a recovered DBDS.
- Ease of Use
  - ▶ DBRC maintains records of recovery-related assets in the RECON data sets. All DBDSs to be recovered must be registered with DBRC. For resources migrated to different volumes, other than what is registered in the RECON data sets, DBRC determines their correct locations from a search of the CATALOG. DBRC restricts access to DBDSs and Areas that are undergoing recovery by granting exclusive authorization to DRF. DBRC CAGRPs, DBDSGRPSs, & RECOVGRPs can be referenced and used as input to DRF.
  - ▶ Dynamic allocation for Logs, CA, ICs, DBDSs, & AREAs is done based on the RECONs.
  - ▶ DRF issues a NOTIFY.RECOV command after all database recovery processing takes place. DRF also updates RECON information regarding the state of the database after the recovery, resets the RECOVER NEEDED flag, and sets the IC NEEDED flag. Also, if the recovery was a timestamp recovery, DRF records that time in the RECON. Users cannot run other applications until after DRF issues the NOTIFY.RECOV command.



## IMS Database Recovery Services - Common Attributes

- Log Data is read from RLDS & SLDS only, not OLDS
- DRF reads RLDS or SLDS Logs in Parallel with Change Accum. DS
- DRF accepts Change Accumulation Input - Complete or Incomplete
- DRF accepts all IMS Image Copy data sets:
  - ▶ Image Copy Utility (DFSUDMP0/DFSUICP0)
  - ▶ HPIC (a.k.a. ICE) Image Copy
  - ▶ IC2 (DFSMS format)
  - ▶ User Image Copy
    - User responsible for restoring non-standard Image Copy
    - notifying DBRC (NOTIFY.RECOV)
    - /RECOVER ADD dbds|area USEDBDS|USEAREA
    - /START RECOVER

- The USEDDBDS|USEAREA keyword on the /RECOVER ADD command tells Database Recovery Service whether or not it should restore the image copy before it applies log updates for a DBDS as part of the recovery process. This is used when user image copies are restored prior to recovery. Log and CA data are then applied to the recovered database instead of the IC.

## ... IMS Database Recovery Services - Common Attributes

- Recovery Time Options
  - ▶ Recover to End of Logs (Current time for ORS) - a.k.a. "full recovery"
  - ▶ Time Stamp Recovery (Recovery Point)
  - ▶ Point-In-Time Recovery (Arbitrary Point)
- Recovery of Multiple DBDS or Areas in one Run
- Runs in Parallel with IMS Transaction Processing
- Supports all Recoverable IMS Database Types  
(Not GSAM, HSAM, SHSAM, or MSDB types)

## ■ Recovery Time Options

- ▶ If neither Time Stamp Recovery nor Point In Time Recovery is specified, databases are recovered to the end of the log data.
- ▶ DRF supports time stamp recovery to DBRC/RECON database allocation boundaries or to any prior point-in-time recovery (PITR). None of the DBDSs or Areas being recovered can be in use for update by any IMS region or job. A /DBRECOVERY command must be issued to all online IMS systems that are using the DBDSs and Areas, and there must be no batch jobs updating the DBDSs and Areas. A listing of the RECON can help in determining valid time stamps for recovery.
- ▶ At the recovery time specified when PITR is selected, the DBDSs and Areas can have been at any allocation status in the RECON data set. It is only necessary that recovery resources are still available for the time stamp specified. There is no restriction on database allocation status if time stamp recovery with PITR is selected. In the recovery process only committed updates, up to and including the specified recovery time, are applied to the DBDSs and Areas in the recovery list. Change accumulation input is used under two conditions. First, the change accumulation data set must have a stop time prior to the PITR recovery time, and second, it must also be determined that the stop time does not fall within the same DB ALLOC/DEALLOC range that spans the recovery time. Only committed updates can be applied and this process assures that a change accumulation data set containing uncommitted data is not used.

## ... IMS Database Recovery Services - Common Attributes

- Support for VTS Data Staging or Caching
- Supports DFSMS DELCAT option
- Multi-Address Space Environment
- Optionally Automates DB Start Command (Local or Global)

- If you manage your image copy data sets through a virtual tape server (VTS), the system automatically moves the contents of the tape data set to a temporary cache data set (high speed DASD ~ Paging Dataset) when the data set is allocated. DRF can help you take advantage of this by pre-allocating the image copy data sets required for recovery as part of DRF initialization. DRF thereby facilitates efficient use of the VTS by overlapping the time required to stage image copy data from tape to cache with the operation of reading the log data sets.
- The DELCAT option on the /RECOVER ADD command should be used with extreme care. It provides the user with the ability to recover from a scenario in which an entire volume or volumes are lost but the catalog entries remain. When this option is specified, SMS deletes the prior catalog entries for the database data sets and areas being restored as part of recovery. Refer to the DFSMSdss Storage Administration Reference manual before using this option. Read the caution under the DELETCATALOGENTRY option of the RESTORE command.
- Whether DRF is started with a batch job or with a the /RECOVER command, a DRF Master address space is started. The DRF Master then starts one or more subordinate address spaces. The subordinate address space procedure is identified to the DRF Master on the DRFMBR parameter.
- The RECOVER ADD command provides the STAGLOBAL and STALOCAL parameters to automate the starting of the recovered database or area (/START DB or /START AREA) after the recovery is complete.

## Online Invocation - Recovery List Concepts

- Multiple active recovery lists are supported
- Each recovery list has a unique token supplied by the user
- Recovery in progress can be stopped with an option to save the recovery list
- Only one recovery list can be started per IMS region
- Entries are added to the recovery list by:
  - ▶ Individual DBDS or Area
  - ▶ Entire Database
  - ▶ Entire CA Group
  - ▶ Entire DBDS Group
  - ▶ Entire RECOV Group
- Entries are removable from the recovery list by any of the above entities
- Each recovery list entry is validated in the RECON data set
- Recovery lists are deleted at the end of recovery

- /REC RCVTOKEN token /REC ADD RCVTOKEN=token DB /REC REMOVE RCVTOKEN=token DB
- /REC STOP RCVTOKEN token SAVE specifies that the list is retained in main storage even if the /REC STOP command results in all of the DBDSs in the recovery list being stopped. Use the /REC STOP command to stop recovery processing for all data sets or areas. If, for instance, you discover that one or more database data sets were inadvertently omitted from the recovery list, you can stop execution of the list with the SAVE option, add the missing data sets, and start processing again from the start.
- /REC START RCVTOKEN token



## DRF (V2) - Evolution from ORS (V1)

- Key features added:
  - ▶ Batch Invocation
  - ▶ Full Formatted Recovery Report
  - ▶ Report Only Execution (VERIFY option)
  - ▶ Distinguishes Mountable and Non-mountable Devices
  - ▶ Parallel Outputs
- Eliminated:
  - ▶ The need for an Active IMS region in batch
  - ▶ Using IMS resources for DBDS updates
  - ▶ All entry options for STOP except ALLENTRIES

- Key Features Added:

- ▶ Submit batch job with ADD control statements in SYSIN as opposed to issuing /REC ADD online.
- ▶ Report Only Execution is handy for testing "what if" disaster recovery scenarios. /START VERIFY command.
- ▶ You can control the number of both mountable (tape) and non-mountable (DASD) devices separately via the values specified with the READNUM parameter in the FRXDRFxx PROCLIB member or in the START control statement. The number of input devices has a direct impact on the performance of the IMS Database Recovery Facility. The READNUM(tn,total) parameter specifies the number of image copy, change accumulation, and log data sets that can be read in parallel. The time required for recovery of databases and areas is directly proportional to this value. If the value of the READNUM parameter is too low, reading image copies and log data sets becomes a bottleneck during recovery. If the image copies and log data sets reside on tape, specifying too high a value for the READNUM parameter impacts the availability of tape drives for other operations during the recovery process. There may be more log/CA read tasks than (total) there are total mountable devices (tn) allowed because some of the logs/CAs may be on non-mountable devices (DASD).
- ▶ Each DRF subspace creates recovered DBDSs. You can set the maximum number of subordinate address spaces to be used by setting the NUM(nn) parameter under the SORTPARM keyword in the FRXDRFxx member. The actual number of address spaces may be less depending on the number of image copy input volumes.

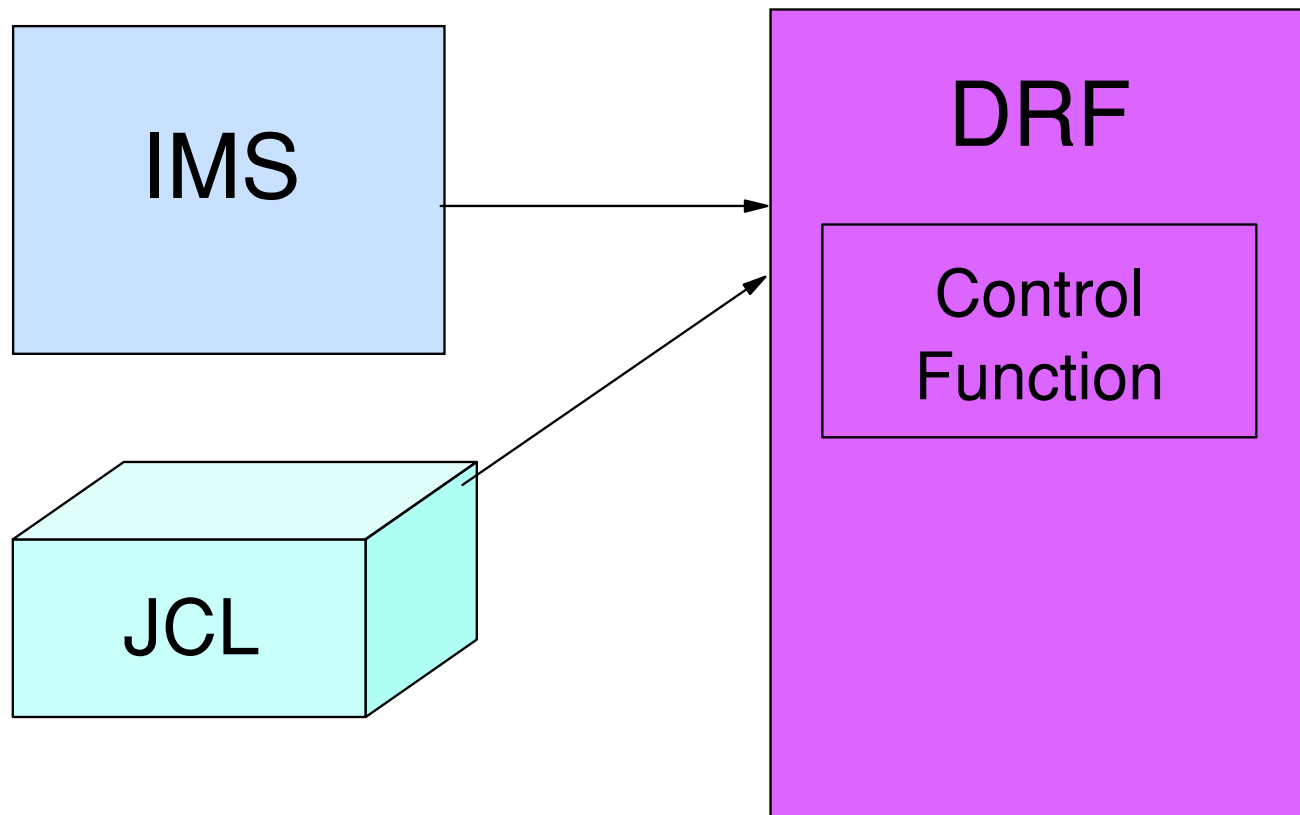
- Eliminated:

- ▶ Database record changes are all done in the recovery subspaces.
- ▶ DL/I calls are not required because the records are updated in place based on Log or CA records.
- ▶ /RECOVER STOP ALLENTIRES is the only command that is supported. Stopping individual database recoveries in the recovery list is no longer supported.

## DRF Control Flow

**Invocation  
Choices**

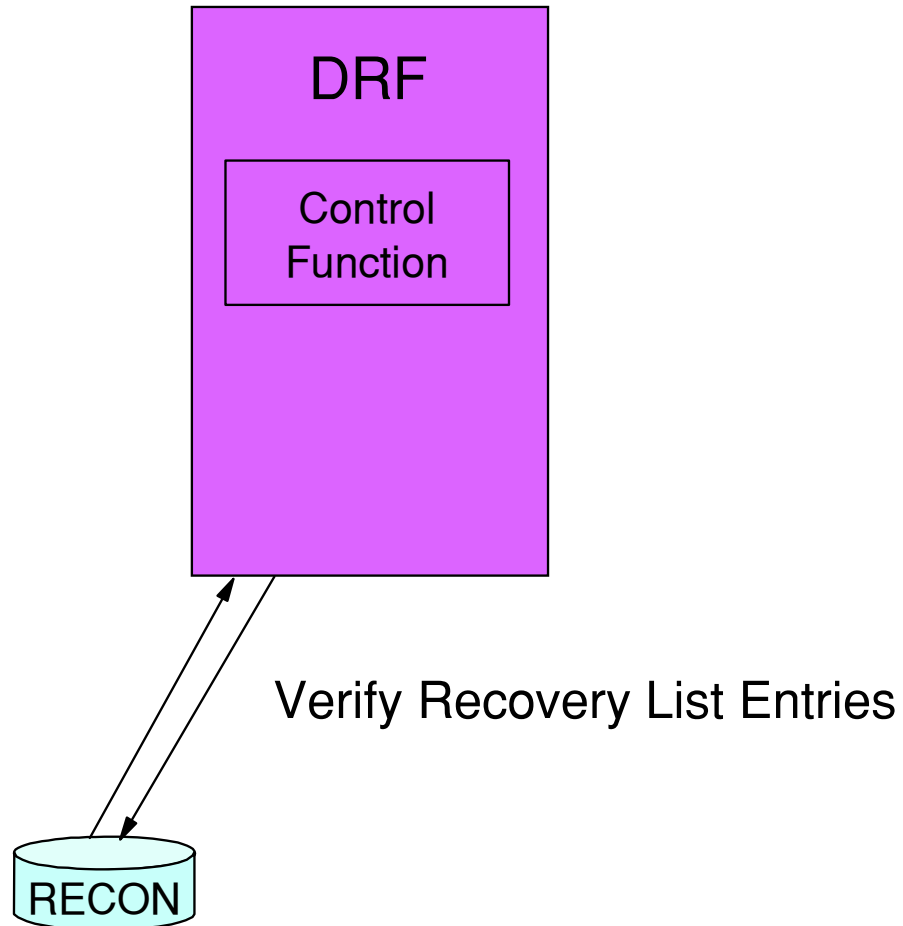
**DRF Master  
Address Space**



- Both the IMS invocation (/RECOVER START) and submission of a DRF batch job cause the creation of the DRF Master AS which contains the Control Function.
- The Control Function manages the creation and termination of the subspaces and communication with the subspaces.

## ... DRF Control Flow

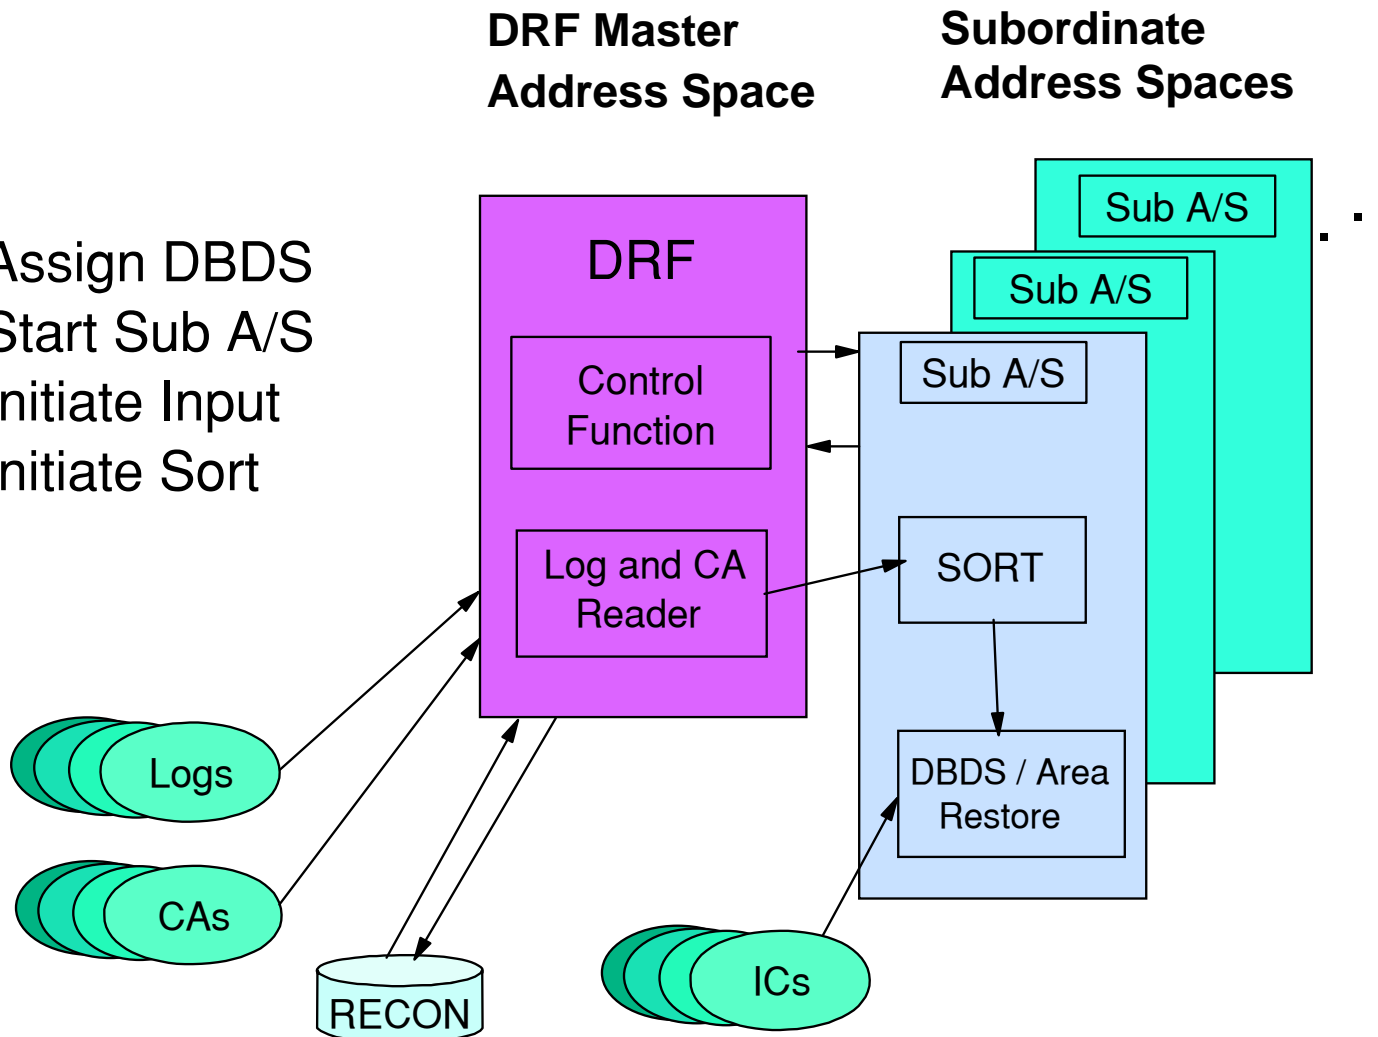
**DRF Master  
Address Space**



- The Control Function verifies the recovery assets & recovery list entries in the RECONs.
- Recovery List Entries include RECON groups including: recovery groups (RECOVGRP parameter on the /RECOVER ADD), Change Accumulation groups (CAGRP parameter) and Database Data set Groups (DBDSGRP parameter).
- DBRC is the facility that maintains records of recovery related resources in the RECON data sets.

## ... DRF Control Flow

1. Assign DBDS
2. Start Sub A/S
3. Initiate Input
4. Initiate Sort

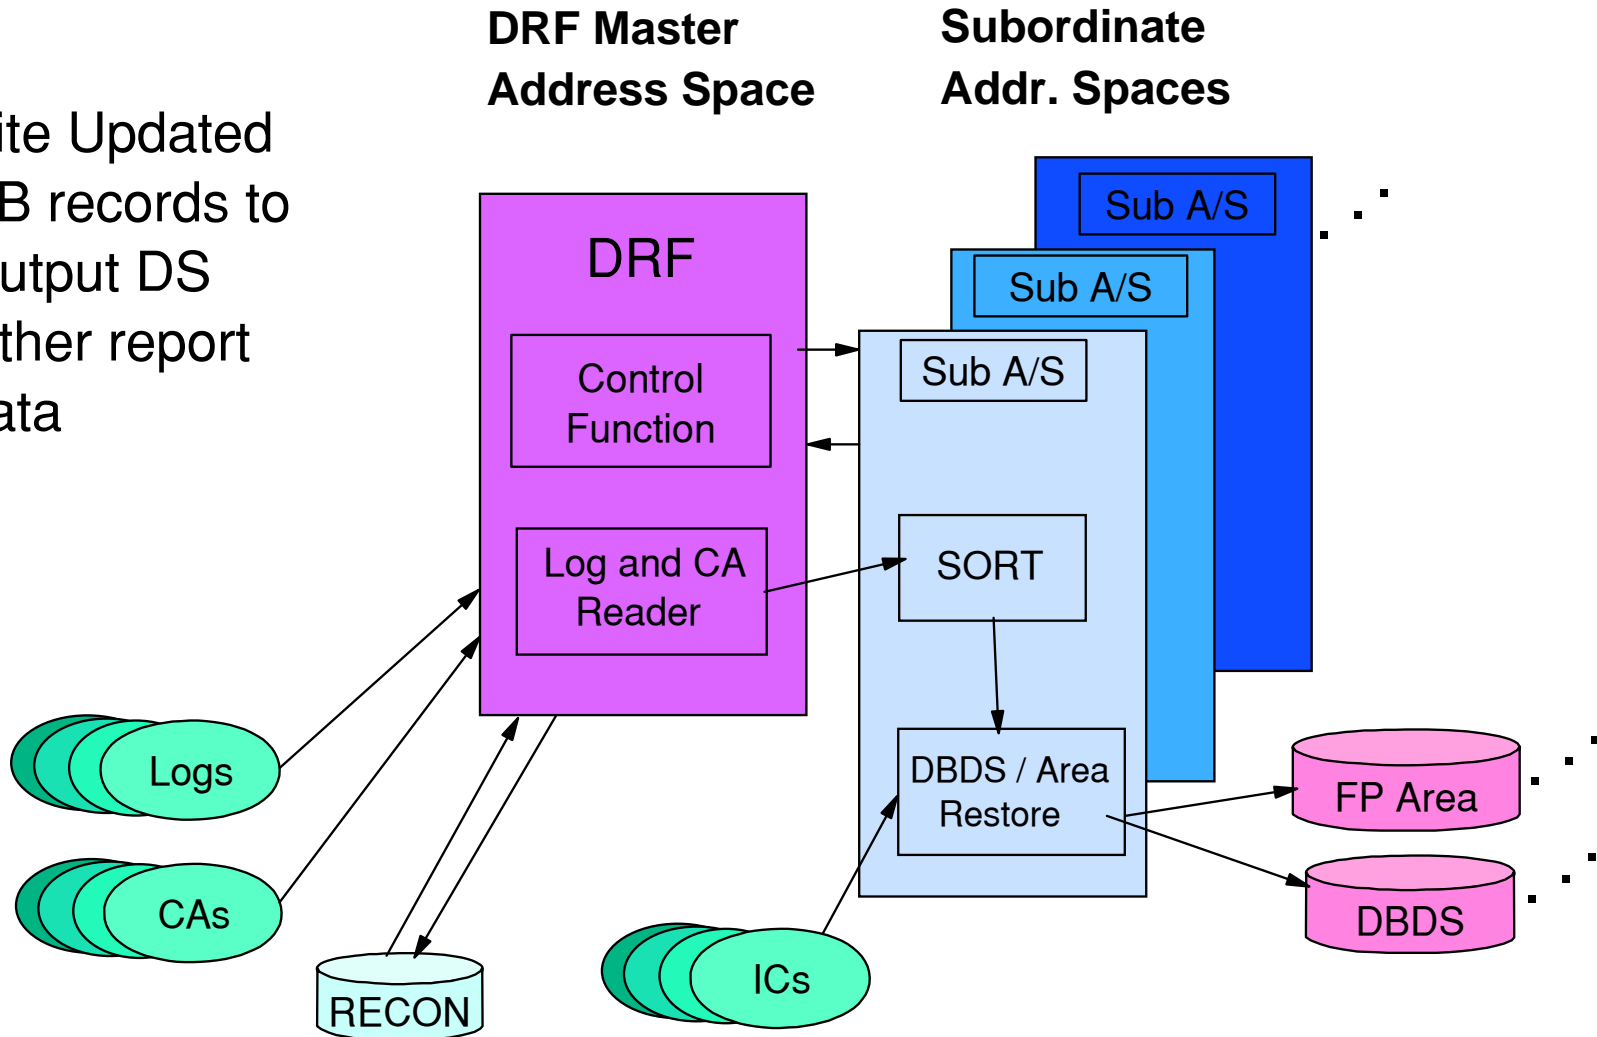


- The DRF Master determines from the RECONs the DBDSs that are to be recovered based on the Recovery List Entries.
- The Master assigns DBDSs to the subspaces about to be created by the MVS ASCRE macro based on the NUM(nn) of subspaces set by the user.
- Next the subspaces are created.
- The Master initiates input by starting the Log or CA reader tasks (max input devices = READNUM)
- The Master instructs the subspace to initiate sort to prepare for the coming Log records.
- The Master begins to send to each subspace only the Log or CA records it needs to recover the DBDSs it is assigned.
- After the Log records are sorted, the subspaces begin to read the Image Copy data set of the DBDS or Area it is restoring, and applies the Log or CA records to the applicable database blocks.



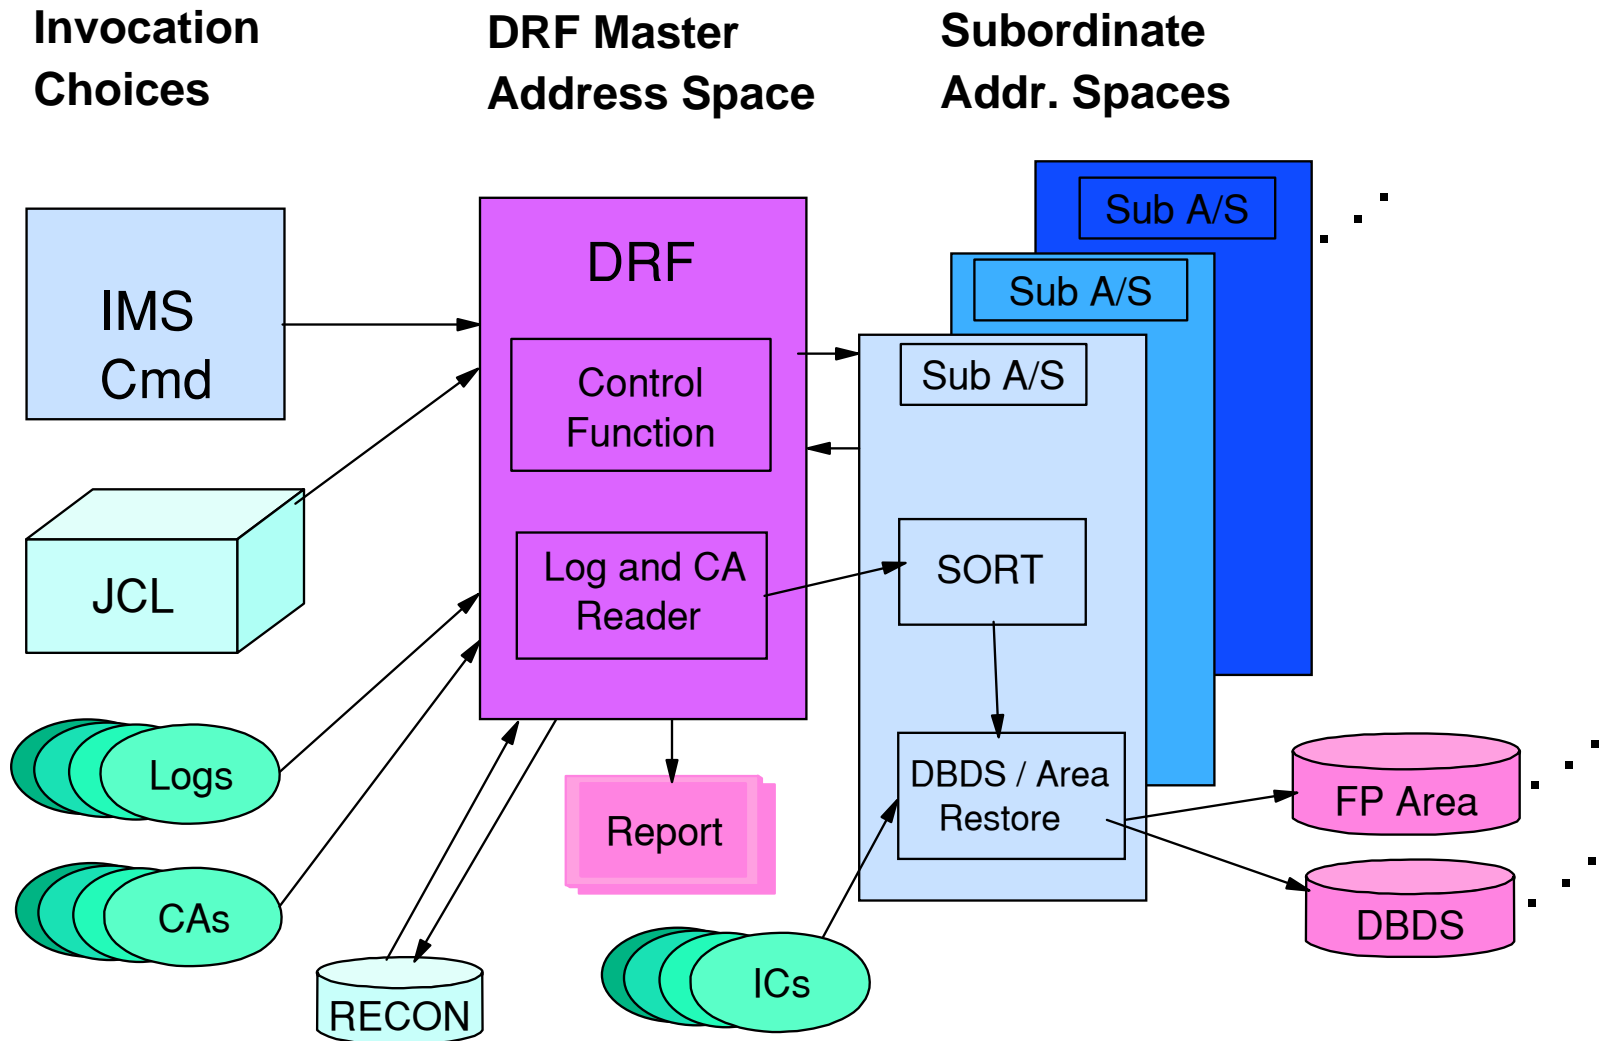
# ... DRF Control Flow

- . Write Updated DB records to Output DS
- . Gather report data



- Next the updated records are written to the recovered database.
- Reports are collected from each subspace regarding the recovery of each DBDS.
- This process continues until all DBDSs assigned to the subspace have been recovered and then the subspace is terminated by the MVS ASDES macro.
- Finally, the collective report is written to the REPORT file by the Master AS.

# ... DRF Control Flow



- The complete big picture.

## DRF Installation: Procedures & Parameters

- IMS Startup Parameters
  - ▶ IMS JCL
  - ▶ DFSORSxx Member
  - ▶ BPE Configuration File
- Database Recovery Facility Start Up Parameters
  - ▶ FRXDRFxx Member
  - ▶ DRF Master Region Procedure
  - ▶ DRF Subordinate Region Procedure

## ■ IMS Startup Parameters

- ▶ To identify the DRF environmental parameter member that is required by the IMS subsystem, include the parameter ORSMBR=xx in your DBCTL or DB/TM EXEC statement or in the parm default block (DFSPBxxx) member of IMS PROCLIB.
- ▶ The ORSMBR parameter points to the DFSORSxx member in which you specify the member name of the JCL cataloged procedure that is used to start the DRF Master AS on the DRFNAME(jclname) statement (any other statements in DFSORSxx are ignored). If the ORSMBR parameter is not specified, the cataloged procedure name defaults to FRXJCLDF. The DFSORSxx member can be reloaded without terminating the IMS. To reload the DFSORSxx member identified by the ORSMBR= startup parameter, issue a /RECOVER TERMINATE command, followed by another /RECOVER ADD command.
- ▶ The DRF address spaces are managed through the services of IMS's Base Primitive Environment. You can change the default settings for BPE such as storage management, internal tracing, dispatching, and other system service functions. DRF supplies a configuration data set member for the DRF Facility BPE system service functions that you can modify, named FRXBPECF. You can change the BPE configuration member with the EXEC PARM='BPECFG=bpembr' statement for the Master AS and/or the Recovery subspace.

## ■ DRF Startup Parameters

- ▶ The FRXDRFxx PROCLIB member provides DRF the parameters to establish its environment. The FRXDRFxx member is identified by the EXEC PARM='DRFMBR=xx' in the JCL procedure that is used to start DRF and must reside in the PDS identified by the procedure's PROCLIB DD statement. The parameters control the level of parallelism in execution, the log data set selection option of primary or secondary, and the generation of data set names for the new recovered output data sets.
- ▶ The DRF Master AS / Region procedure FRXJCLDF invokes program FRXSDR00 and passes several execution parameters (PARM=) to it. The procedure contains DDs that point to the RECONS, IMSMDALIB, DBDLIB, REPORT file, and others. One of the execution parameters is DRFPROC=nnn.
- ▶ DRFPROC=nnn points to the cataloged procedure used for execution when the DRF Master AS starts a subordinate AS (subspace). The subspace, which is involved in sorting among other functions, has DDs to point to the SORTLIB and IMSMDALIB for dynamic allocation of Image Copy and output recovered databases. If this parameter is omitted, the default name is FRXJCLSB. The FRXDRFxx PROCLIB member may also contain DRFPROC=nnn, but the execution parameter is the override.

## DRF Installation Steps

- **Create Members in PROCLIB**
  - ▶ Master and Subordinate Cataloged Procedures
  - ▶ See Samples in SAMPLIB on Product Tape
- **If Use is ONLINE then**
  - ▶ Add ORSMBR=xx to DFSPBxxx or IMS PROCs
  - ▶ Add the DRF Load Library to the IMS STEPLIB Concatenation
- **If Use is Batch then**
  - ▶ Create JCL members with SYSIN control statements including ADD DB() and START
- **Ensure that FRXSDR00 is in the MVS Program Properties Table**
- **Must be in an APF Authorized Library. It runs Key-7 Sup. State**

- Updating the MVS Program Properties Table: The IMS Database Recovery Facility runs as an authorized address space. Its entry point, FRXSDR00, must be listed in the MVS program properties table (PPT). The steps for adding the FRXSDR00 entry point to the MVS PPT are:
  1. Edit the SCHEDxx member of the SYS1.PARMLIB data set.
  2. The following program properties table entry can be used as an example for adding an entry to the SCHEDxx member of SYS1.PARMLIB:
    - ▶ PPT PGMNAME(FRXSDR00) /\* PROGRAM NAME = FRDRVS00 \*/
    - ▶ CANCEL /\* PROGRAM CAN BE CANCELLED \*/
    - ▶ KEY(7) /\* PROTECT KEY ASSIGNED IS 7 \*/
    - ▶ NOSWAP /\* PROGRAM IS NON-SWAPPABLE \*/
    - ▶ NOPRIV /\* PROGRAM IS NOT PRIVILEGED \*/
    - ▶ DSI /\* REQUIRES DATA SET INTEGRITY \*/
    - ▶ PASS /\* CANNOT BYPASS PASSWORD PROTECTION \*/
    - ▶ SYST /\* PROGRAM IS A SYSTEM TASK \*/
    - ▶ AFF(NONE) /\* NO CPU AFFINITY \*/
    - ▶ NOPREF /\* NO PREFERRED STORAGE FRAMES \*/
  3. Choose one of the following actions to make the SCHEDxx changes effective:
    - A) Re-IPL the MVS system where SCHEDxx is identified in the IAESYS IPL member. See z/OS MVS Initialization and Tuning Reference for more information.
    - B) Issue the MVS SET SCH= command specifying the suffix for the SCHEDxx member which was changed.



## FRXDRFxx Member

- READNUM (tn,total) - number of log / CA reader tasks
  - ▶ tn=tasks allowed to read mountable devices
  - ▶ total=tn+tasks to read non-mountable devices
- SORTPARM (parms)
  - ▶ NUM(nn)
    - Controls the number of subordinate address spaces
  - ▶ HIPRMAX(vvvvv)
  - ▶ MAINSIZE(mmmm)
  - ▶ ASPREF(cccc)
  - ▶ AVGRLEN(nnnnn)
  - ▶ FILSZ(nnnnnnnn)
- LBI (y | n)
- SPSIZE (yyyyy)

- HIPRMAX is the maximum size of hiperspace to be utilized in the sort process.
- MAINSIZE is the main storage size to be used by SORT running in the Recovery subordinate AS.
- ASPREF(cccc) is used to construct the STC name of the Recovery subspace whose JCL is in the PROCLIB member identified by the DRFPROC=nnn execution parameter. If ASPREF is not specified, the default used is IDRF. A four digit number is concatenated at the end of the supplied parameter to form a unique address space name. For example, the first subordinate address space has the name "IDRF0001," , using the default The second subordinate address space has the name "IDRF0002."
- AVGRLEN is a 1-to-5 digit integer numerical value (4 through 32766; in bytes) that specifies the average record length of records to be sorted. The default value is half of the maximum record length. The range includes the 4-byte record descriptor word (RDW).
- FILSIZ is a 1-to-9 digit integer numerical value specifying the estimated number of records to be sorted. This parameter is optional, and if it is not included, DRF makes an estimate based on the RLDS record sequence number range provided by DBRC. The minimum estimate applied is 20000.
- LBI: This parameter indicates whether large block size is in use at this installation.
- SPSIZE is the size in megabytes of each data space used by DRF for managing various storage. When data space usage reaches the value specified, another data space is obtained. The range of valid values is: 1 to 2047. The default is 1024 megabytes. Some installations limit the size of data spaces, so a greater number of smaller dataspace may be required.

# DRF Master Address Space JCL

```
//*****  
//* PROCEDURE FOR DATABASE RECOVERY FACILITY (DRF) MASTER REGION *  
//*****  
//FRXJCLDF PROC RGN=0M,SOUT=A,/*from DFSORSxx mbr DRFMBR(jclname)parm  
// RESLIB='DRF.SFRXRESL',  
// BPECFG=FRXBPECF, /* BPE configuration (trace parameters, etc.)  
// DRFMBR=ZZ, /* FRXDRFxx contains parameters (parallelism, etc.)  
// GSGNAME=,  
// PLEXNAME=,  
// DRFPROC=FRXJCLSB,/* subordinate AS procedure name */  
//STEP1 EXEC PGM=FRXSDR00,  
// PARM=('DRF,BPECFG=&BPECFG,DRFMBR=&DRFMBR,DRFPROC=&DRFPROC',  
// 'GSGNAME=&GSGNAME,PLEXNAME=&PLEXNAME'),  
// REGION=&RGN,TIME=1440  
//STEPLIB DD DSN=&RESLIB,DISP=SHR /* Must be APF authorized  
// DD DSN=IMS.SDFSRESL,DISP=SHR  
//IMSDALIB DD DSN=IMS.DYNALLOC,DISP=SHR  
//PROCLIB DD DSN=IMS.PROCLIB,DISP=SHR  
//DBDLIB DD DSN=IMS.DBDLIB,DISP=SHR  
//SYSIN DD DUMMY  
//SYSPRINT DD SYSOUT=&SOUT  
//SYSUDUMP DD SYSOUT=&SOUT  
//REPORT DD SYSOUT=&SOUT,DCB=(RECFM=FBA,LRECL=133,BLKSIZE=6118)  
//*  
//*
```



# DRF Subordinate Address Space JCL

```
//*-----  
//* DRF SUBORDINATE REGION PROCEDURE  
//*-----  
//FRXJCLSB PROC RGN=0M,SOUT=A, /*From DRFPROC=nnn on PARM= on FRXDRFJCL master.  
// RESLIB='DRF.SFRXRESL',  
// BPECFG=BPEDRFCG  
//*****  
//* BRING UP A SUBORDINATE RECOVERY REGION *  
//*****  
//STEP1 EXEC PGM=FRXSDR00,  
// PARM='DRF,BPECFG=&BPECFG',  
// REGION=&RGN,TIME=1440  
//STEPLIB DD DSN=&RESLIB,DISP=SHR /* Must be APF authorized.  
//IMSDALIB DD DSN=IMS.DYNALLOC,DISP=SHR  
//PROCLIB DD DSN=IMS.PROCLIB,DISP=SHR /*Contains FRXDRFxx meber with SORTPARM  
//SYSPRINT DD SYSOUT=&SOUT  
//SYSUDUMP DD SYSOUT=&SOUT  
//SYSOUT DD SYSOUT=&SOUT  
//SORTLIB DD DSN=DFSORT.PGMLIB,DISP=SHR
```



## Using DRF

- Sub-Topics:
  - ▶ Pre-Recovery Actions
  - ▶ On-Line Recovery Control
  - ▶ Batch Notes
  - ▶ The Batch Recovery Report
  - ▶ Miscellaneous Items
  - ▶ Post Recovery Activities

- What needs to be done before DRF is executed to recover a DBDSGRP, CAGRP, RECOVGRP, DB, DBDS or AREA.
- What commands are used to control ORS recovery.
- Notes on Batch invocation.
- Format and contents of the Recovery Report (REPORT file).
- What might need to be done after the recovery completes.



# Pre-Recovery Actions

- /DBR command to take DB offline
- Delete and Redefine the DBDS where necessary
- Define the Recovery Scope
  - ▶ Build the Recovery List
    - Online - Use /RECOVER ADD & /RECOVER REMOVE
    - Batch - Use SYSIN Recovery Control Statements (ADD & REMOVE)
- STAGLOBAL or STALOCAL on /RECOVER ADD or START

- Issue the /DBRECOVERY command to take the DB offline to prepare for recovery. DBRC deallocates the database from an IMS perspective so that no further updates can be made. This command prevents further authorizations or allocations except for the reorganization and recovery utilities. Specify the GLOBAL keyword on the /DBRECOVERY command to have the command apply to all subsystems sharing the database. The IRLM must be active if you use the GLOBAL keyword. The default is LOCAL, which specifies that the command applies only to the subsystem on which you enter the command.
- To change the size of or volume on which the dataset is allocated, delete and redefine (IDCAMS) the DBDSs.
- If the STALOCAL or STAGLOBAL parameters are used on the /RECOVER START or ADD commands, all full-function databases and fast path areas for which recovery completed successfully will be started. If the fast path areas are brought offline with the /DBR DB command, STALOCAL or STAGLOBAL does not start the fast path database. If the STALOCAL or STAGLOBAL parameters are not used, then the databases are ready to be started in the normal manner.

## Online Recovery Control: Building the Online Recovery List

- /RECOVER ADD command
- /RECOVER REMOVE Command
- Multiple List Identification
  - ▶ RCVTOKEN parameter on /RECOVER commands
- /DISPLAY RECOVERY - Monitor progress
  - ▶ /DISPLAY RECOVERY RCVTOKEN token
  - ▶ /DISPLAY RECOVERY ALL
- Recovery List Disposition
  - ▶ Forming
  - ▶ Started
  - ▶ Unknown (not found)
  - ▶ Normal Completion
  - ▶ Abnormal Completion
- /RECOVER STOP with the SAVE option

- The RCVTOKEN on the various /RECOVER commands allows the activation of multiple recovery lists. Each recovery list is given a token. The default is DFSxxxxx starting at 00001. At IMS restart time, the default is reset to 00001. If no token is specified, a default of DFSxxxxx is used, where xxxxx is the next available integer. If a token is specified and the recover list is already defined, then database records are added to the list. If two ADD commands are submitted without specifying a token for each, then the entries go into two separate recovery lists. Many lists can be built, but only one at a time can be started. The lists are kept while the current IMS system is active. The recovery list for the token is kept until one of the following occurs:
  - a) Recovery is run successfully
  - b) IMS is terminated
- Recovery lists are not retained after recovery completes normally.
- If the recovery abnormally completes the recover list is saved and put in FORMING status if /RECOVER START ERRORABORT is specified.
- Use the /RECOVER STOP command to stop recovery processing for all data sets or areas. If, for instance, you discover that one or more database data sets were inadvertently omitted from the recovery list, you can stop execution of the list with the SAVE option, add the missing data sets, and start processing again from the start.

## ... Online Recovery Control: /RECOVER START command

- /RECOVER START Command - Online
  - ▶ Identify the correct Recovery List with RCVTOKEN
  - ▶ VERIFY Option
  - ▶ ERRORCONT or ERRORABORT
  - ▶ READNUM(n) Override
  - ▶ RCVTIME - time stamp
    - TSR versus PITR
    - CHECK versus NOCHECK

- /RECOVER START DRF with the VERIFY (report only) option. This is done to obtain a report showing all of the input log, change accumulation, and image copy data sets required for the recovery process.
- ERRORABORT specifies that the recovery is to be terminated at the first error. The recovery list is kept and put back into FORMING state. ERRORCONT specifies that the recovery process is to continue for those DBDSs without errors. Any DBDS in error remains in the RECOVERY NEEDED state in the RECON data set. The recovery list is purged at the end of processing. The recovery list will not start until all the DBDSs in the list can be authorized for the DRF recovery if the ERRORABORT parameter is in effect. This is not true if the ERRORCONT parameter is specified on the /RECOVER START command. In this case, the recovery continues.
- READNUM specifies the number of concurrent read operations that DRF is allowed to start. The read operations include the restores of the image copy, log, and change accumulation data sets. If the number of devices available is less than the value of n, DRF uses only the available devices. READNUM is defined in the FRXDRFxx PROCLIB member. If you specify READNUM on the /RECOVER START command, the command value takes precedence.
- The RCVTIME parameter must be specified whenever a time stamp recovery is performed. If TSR or PITR is specified, this is a required parameter. If the PITR parameter is not specified at the recovery time, none of the DBDSs and AREAs being recovered were able to be in use for update at the recovery time by any IMS subsystem.
- TSR: Use of this parameter causes time stamp recovery to be performed. A /DBR command must have been issued to all online IMS systems which had the databases or areas in use and there must not be any batch jobs updating them at the time specified by the RCVTIME parameter. A listing of the RECON can help in determining valid time stamps. There must be no ALLOC record for any DBDS or AREA being recovered that spans the recovery time.
- PITR specifies point-in-time recovery is to be performed. At the recovery time specified, database data sets and areas were able to be in any allocation state. That is, there is no restriction on database allocation status if PITR recovery 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. If you specify the PITR parameter, the recovery time can be specified to any time prior to the current time.
- CHECK specifies DRF will automatically check to ensure whether a DBDS is in a recovery group, a full-function database, or a HALDB partition. Recovery will not start unless all DBDS, areas, or both are either being recovered by this recovery or are already recovered to an equivalent time. You will be notified of DBDSs and areas that might be inconsistent with recovery of database data sets and areas on the recovery list. NOCHECK puts this responsibility on the user.

## ... Online Recovery Control: /RECOVER STOP or TERMINATE

- /RECOVER STOP ALLENTRIES Command
- /RECOVER TERMINATE Command

- `/RECOVER STOP ALLENTRIES` specifies that recovery is to be stopped for all databases, data sets, and areas (all entries) in the recovery list. Processing stops immediately. The list is deleted unless you also specify the `SAVE` parameter. If you specify the `SAVE` parameter the recovery list is placed back in `FORMING` status.
- `/RECOVER TERMINATE` will delete all lists in `FORMING` status and terminate the DRF address space. If a recovery is in progress the recovery will ignore the `/RECOVER TERMINATE` command and continue to process. When it completes, another `/RECOVER TERMINATE` command is required. It will not take effect automatically. If you would like to force DRF down while a recovery is running, you must issue the `/RECOVER STOP ALLENT` command first. This will stop the recovery and then you might enter the `/RECOVER TERMINATE` command to terminate the DRF address space.



## Online Recovery Control: VERIFY Option

- (works in Batch, too on START control statement)
- /RECOVER START VERIFY performs:
  - ▶ Recovery List Syntax Validation
  - ▶ Obtaining of Recovery Input File Information from DBRC
  - ▶ Listing of All Input Data Sets in Formatted Report
- VERIFY Does not perform:
  - ▶ Recovery Processing
  - ▶ Allocation of Input Data Sets (Logs, CA, IC)
  - ▶ Opening of Input Data Sets
  - ▶ Purge of Recovery List



# Batch Notes

- SYSIN Recovery Control Statements
  - ▶ ADD
  - ▶ REMOVE
  - ▶ START
- IMS Control Region is not Required as in Online
- DBRC / RECON Access is Required (via DD)
- Runs Authorized, Key 7, from APF Auth. Library
- User ID requires RACF Access to All Protected Resources
- Single Recovery List - No Token Entered
- ERRORABORT -> ERROR(STOP) Online-> Batch
- ERRORCONT -> ERROR(CONT) Online-> Batch

- When submitting a batch job to invoke DRF, the userid associated with the JOB statement must have any passwords required by the (RACF or security management software) for authorized database data set access. In the Online environment, when invoking DRF with the z/OS START command or with the IMS /RECOVER command, the userid associated with the primary address space standard task control (STC) name, as identified via association with the RACF® class STARTED or the started task table (SYS1.LPALIB(ICHRIN03)), must have appropriate database data set access authorized through the security management software. In both batch and online invocations of DRF, subordinate address spaces will be initiated to perform sort-related functions. The STC names for the subordinate address spaces must also be identified and allowed resource access in a manner similar to the primary address space. The STC name prefix for the subordinate address spaces is specified by the DRF ASPREF environment control parameter or by the default value, IDRF, if the parameter is not specified. It is recommended that the same userid that is associated with the IMS control region be associated with the started tasks and/or batch job for the IMS Database Recovery Facility.

## ... Batch Notes: Examples

- Batch Control Statements
  - ▶ ADD DB(DBOVLFPC)
  - ▶ ADD DB(HIDAM001)
  - ▶ START
- Verify Option
  - ▶ ADD DB(HDAMW440)
  - ▶ ADD DB(HIDAM001)
  - ▶ START VERIFY



## Batch Recovery Report

- Display of Batch Control Statements
- Display of Processing Options
- Lists of All Input Data Sets with Statistics
  - ▶ Log Data Sets
  - ▶ Change Accumulation Data Sets
  - ▶ Image Copy Data Sets
- List of all DBDS and Areas Recovered and Status
- List of Open UORs in the Case of PITR Recovery

- The Open Unit Of Work/Unit Of Recovery (UOW/UOR) Report for PITR processing:  
When the user invokes a point-in-time recovery (PITR) there is a possibility that there are open UORs at the recovery time. The changes associated with the open UOR are not yet applied to the database because the changes were not committed at the recovery time. PITR recoveries include a trailing “PITR Open UOW/UOR Report” following the other sections; non-PITR recoveries do not include this section. This report includes information about PSBs which had activity related to the databases being recovered, where some of the activity is not included in the input (log records) to the recovery because the point in time selected was in the middle of an active unit of recovery. The Recoveries specifying VERIFY contain the value “N/A” for statistics in the “Facility Summary Report,” zeros for statistics in the “Data Set I/O Report,” “No open UOWs...” in the “PITR Open UOW/UOR Report”; actual recoveries contain valid statistics in these reports. UORs are identified by the 'begin UOR' x'5607' log record and the 'End UOR' x'07' log record.



# ... Batch Recovery Report

===>

FRD0000I D A T A B A S E R E C O V E R Y F A C I L I T Y 05/21/2003 08:57  
1

Page

D A T A B A S E R E C O V E R Y F A C I L I T Y C O M M A N D S / C O N T R O L S T A T E M E N T S

FRD7201I ADD DB(DBOVLFPC)  
FRD7201I ADD DB(DEDBDD01)  
FRD7201I START /\* VERIFY \*/

D A T A B A S E R E C O V E R Y F A C I L I T Y R E C O V E R Y P A R A M E T E R S

PROCESS : RCVDBDS RECOVERY OPTIONS  
SOURCE : PRI RCVTIME : Not Specified  
READNUM : 06,20 TYPE : Full  
ERROR : STOP CHECK : Yes  
RCVTOKEN: DRFBASE1

D A T A B A S E R E C O V E R Y F A C I L I T Y S U M M A R Y R E P O R T

Database Name	DD/Area Name	DSID	IC	Records Read	CA	LOG	Records Written	Subord. Reg Name	Final Status
DBOVLFPC encountered	VLOSAM01	1	1011	0		4229	1011	DRFI0001	No errors
DEDBDD01 encountered	DD01AR0	1	84	0		205	84	DRFI0002	No errors

FRD0000I D A T A B A S E R E C O V E R Y F A C I L I T Y 05/21/2003 08:57  
Page 1

- Notice 'echoed' control statements and the commented-out VERIFY.
- SOURCE: PRImary IC data sets are used. READNUM: 06,20 The max number of log read tasks for tape is 6 and the max number of concurrent read threads is 20 (in case some inputs are on devices other than tape). ERROR: STOP If there is an error, DRF will stop all tasks and terminate. RCVTOKEN: DRFBASE1 is a user assigned token. RCVTIME not specified, so this is known as a 'full recovery', which is a recovery to the end of Logs time. DRF will use all Log and CA updates between the latest IC stop time and the end of Logs time.
- Database Name: If there were more than one data set or area in the database then more than one DD or area name listed per database name, thus monotonically ascending DSID numbers for each DD name.
- Records Read: No applicable CA records. Subord. Reg Name: Indicates two subspaces ran in parallel.

# ... Batch Recovery Report

FRD0000I DATABASE RECOVERY FACILITY 05/21/2003 08:57

Page 1

DATABASE RECOVERY FACILITY DATA SET I/O REPORT

Recover to point: Not Specified

Image Copy Data Set Name -----	Volume Serial	IC DS Read Count	IC Type	----- Time Stamp Range ----- 1st Record	----- Last Record
IMSVS.DEDBDD01.DD01AR0.IC.IC163902	222222	84	STD		
IMSVS.DBOVLFPC.VLOSAM01.IC.IC163932	222222	1011	STD		

Change Accum Data Set Name -----	Volume Serial	CA DS Read Count	----- Time Stamp Range ----- 1st Record	----- Last Record
-------------------------------------	------------------	---------------------	--------------------------------------------	----------------------

No data available for this type data set

Log Data Set Name -----	Volume Serial	Log DS Read Count	IMS SYSID	----- Time Stamp Range ----- 1st Record	----- Last Record
IMSVS.RLDSP.IMS1.D03140.T1641368.V00 Prilog: 2003.140 23:37:58.7	000000	41	IMS1	2003.140 23:41:36.8	2003.140 23:42:59.7
IMSVS.RLDSP.IMS1.D03140.T1638265.V00 Prilog: 2003.140 23:37:58.7	000000	4472	IMS1	2003.140 23:38:26.5	2003.140 23:41:36.8

- Although the Time Stamp Range appears for IC and CA data sets, it has no significance.
- Notice the Prilog record (DBRC record indicating log allocation time) times are earlier than the first and Last Record times for the Log Data Sets.

# Miscellaneous Items

- Sort Work Space
- Parallelism affected by:
  - ▶ Stacked ICDS
  - ▶ Tape Drive Availability
  - ▶ DFSMSdss SAMEDS usage
- On-line AOI Programs - review required:
  - ▶ Some messages are different
  - ▶ DATAGROUP Support Removed

- Sort work space consists of the SORTPARM environmental parameters and the more accurately these are chosen the more efficient the call to sort from the Recovery subspaces. NUM(nn) for the number of subspaces is the most important. HIPRMAX should be allowed to default to OPTIMAL in most cases. MAINSIZE should be as large as the installation can afford to have. AVGRLEN is usually an educated guess, but the more accurate the more efficiently SORT uses storage; SORT statistics produced during the recovery can be helpful. DRF does a good job of setting FILSIZ dynamically, so this is not as important as the others.
- Parallelism: Stacked Image Copy Data Sets (ICDS) are on the same Tape Volume therefore only one device is able to read this volume at a time rather than two volumes reading two or more separate stacked ICs in parallel. If there are not enough tape drives to satisfy the READNUM environmental parameter then parallelism is not maximized. DFSMSDSS SAMEDS is similar to Stacked ICDS in that, a single subspace is used to call DFSMSDSS to restore IC2 Image Copies created by SMS using the SAMEDS option.
- AOI programs that look for certain messages may need to be changed.
- DATAGROUP was removed because it provided Fast Path area names without the DB name for those areas. DRF needs to have the DB names for registering the results of recovery with DBRC. DBRC would have to look at each area record to find the DB name for each area in the recovery list and this would impact performance. ORS had the Fast Path area list available to process outside of DBRC & DRF does not have that luxury - that is why DATAGROUP support was removed. A recovery group (RECOVGRP) is a group type in the RECON data set. It is a kind of database group that can be used with DRF and DBRC commands anywhere a database group (DATAGROUP) can be used. It differs from a regular database group in these ways:
  - A) All members must be registered with DBRC.
  - B) A member can belong to only one RECOVGRP. If it belongs to a RECOVGRP, the group name is stored in the DB record (for FF) or the AAUTH record (for an Area).
  - C) A member's presence in a RECOVGRP can affect the DRF process. When DRF is about to perform a time stamp recovery, either to an allocation boundary or to any-point-in-time, for a recovery group member, it ensures that all members of the recovery group are recovered in the same recovery operation unless otherwise specified by the user via the NOCHECK parameter.

# Post Recovery Activities

- Rebuild HALDB Partitions' Indirect List Data Set (ILDS) and Primary Index using DFSPREC0.
  - ▶ After time stamp recovery (TSR, PITR, RCVTIME)
- Rebuild Index Data Sets
  - ▶ After a time stamp recovery or if media containing index also failed.
  - ▶ IMS Tools Index Builder
- Run Pointer Checker
  - ▶ To validate internal database pointers - ensures correct recovery
  - ▶ IMS Tools High Performance Pointer Checker
- Run Image Copies
  - ▶ To reset recovery point for less complicated future recoveries
  - ▶ IMS Tools High Performance Image Copy (a.k.a. ICE)

- After a time stamp recovery, the ILDS data sets and a PHIDAM primary index data set must be rebuilt. Use the HALDB Index/ILDS Rebuild Utility (DFSPREC0) is used to perform this function. If a full recovery is run on the ILDS and index data sets, there is no need to perform the rebuild.
- Index data sets can be recoverable or non-recoverable. Recoverable index data sets can be recovered as part of the IMS Database Recovery Facility recovery process by adding them to the recovery list. Alternatively, they can be rebuilt after recovery is complete through an index rebuild utility. Non-recoverable index data sets might need to be rebuilt. The index data sets must be rebuilt or recovered if a time stamp recovery has been performed. If time stamp recovery has not been performed and if the media that the index data sets are on did not experience failure, then the index data sets might not need to be rebuilt.
- Optionally, you can ensure that the recovery has accomplished its specific task by running a pointer checker to validate the internal IMS pointers. If you normally run a pointer checker in the image copy process (HPIC can do hash pointer checking), then this step might be omitted.
- Running an image copy for those databases which have just been recovered will reset the recovery point and avoid a more complicated recovery if one is needed in the near future. If PITR was used on a time stamp recovery the IMAGE COPY NEEDED flag has been set on. The databases will not be available to be authorized for update processing until the image copy. Resetting the IMAGE COPY NEEDED flag after a PITR recovery is very dangerous, because the PITR recovery processed only some records on a log file and subsequent recoveries will not handle those log files correctly.




# What's New in DRF

- **Secondary Sources (Disaster Recovery)**
  - ▶ Primary copies are ignored during recovery
    - RLDS, SLDS, and image copies
    - PQ79162 / UQ81789
- **Large Block Image Copy**
  - ▶ Image copies with >32K blocks are supported
  - ▶ LBI(Y or N) as a start up parameter
    - PQ79327 / UQ82548
- **VTS Support for image copy**
  - ▶ Virtual tape support for image copy restore
    - PQ79326 / UQ81798
- **Environmental Parameter Override in SYSIN**
  - ▶ All FRXDRFxx parms in batch is now also available in SYSIN DD statements
    - PQ77559 / UQ81788



## ... What's new in DRF

- PITR with Change Accum. input
  - ▶ Change accum. data sets stopped prior to a DB allocation range are used



No CADS is used for recovery to T1 because it is within an ALLOC boundary  
CADS1 is used for recovery to T2 because it is outside an allocation boundary  
CADS1 is used for recovery to T3 because CADS2 is within an allocation boundary  
CADS1 is used for recovery to T4 because it is outside an allocation boundary

## ■ V9 Run-Under Support

- ▶ All IMS releases in service are supported