

# E46

## IMS Online Recovery Service

Rick Long



Anaheim, California

October 23 - 27, 2000

# IMS Online Recovery Service

- Capability
  - ▶ Shared databases and areas need to be recovered quickly and easily
- Availability
  - ▶ Recovery can not leave databases and areas offline for extended periods of time
    - Log data sets should not be read sequentially
    - Multiple DBDS and areas should be recovered simultaneously

# IMS Online Recovery Service (continued)

- **Timestamp Recovery**
  - ▶ The recovery timestamp should not be limited to allocation boundaries
    - Does not allow sufficient database availability in modern computing environments

# Design Goals

- No impact to IMS activity that is not recovery related
- Simplify the recovery process for shared databases and areas
- Reduce the amount of time required for recovery of multiple databases and areas
- Allow recovery to any prior point in time

# Highlights

- Recover multiple DL/I DBDS and Fast Path areas in a single pass of the IMS log
- Timestamp recovery to allocation boundaries or any prior point in time
- A new facility executing in conjunction with the IMS control region

# Highlights (continued)

- Executes in parallel with online IMS activity
- Recovery is initiated via IMS commands
- Supports all IMS recoverable database types
  - ▶ Databases marked non-recoverable in RECON can have image copies restored if available
  - ▶ GSAM, HSAM, SHSAM, and MSDB database types are not supported

# ***Highlights (continued)***

- Shared DBDS and areas can be recovered directly from logs
  - ▶ Or a combination of change accumulation data and log record
  
- Incomplete change accumulation data sets can be used for recovery
  
- Databases and areas can be started automatically
  - ▶ On ALL applicable IMS systems
  - ▶ Only on the IMS systems performing recovery
  - ▶ Only on a full recovery

# Characteristics

- Users must restore nonstandard image copies prior to recovery
- Log data required for recovery must reside on SLDS or RLDS
- IMS Online Recovery Service is not restartable
- All resources required for recovery must be registered in RECON
- A DLI/SAS region is required to recover full function databases



# Timestamp Recovery

- Timestamp Recovery can be enforced to allocation boundaries
- Timestamp Recovery can be to any prior point in time (PITR)
  - ▶ The user determines the timestamp to be used
  - ▶ The timestamp is not restricted by allocation ranges
- All updates that are committed as of the specified timestamp are applied

# Timestamp Recovery (continued)

- All DBDS in a database or partition must be recovered to the same point in time before they can be used for online or batch processing
- Related DBDS and areas are optionally identified at the start of recovery
- Indirect List Entries and primary indices for partitioned databases must be rebuilt by the DFSPREC0 utility if PITR timestamp recovery is performed

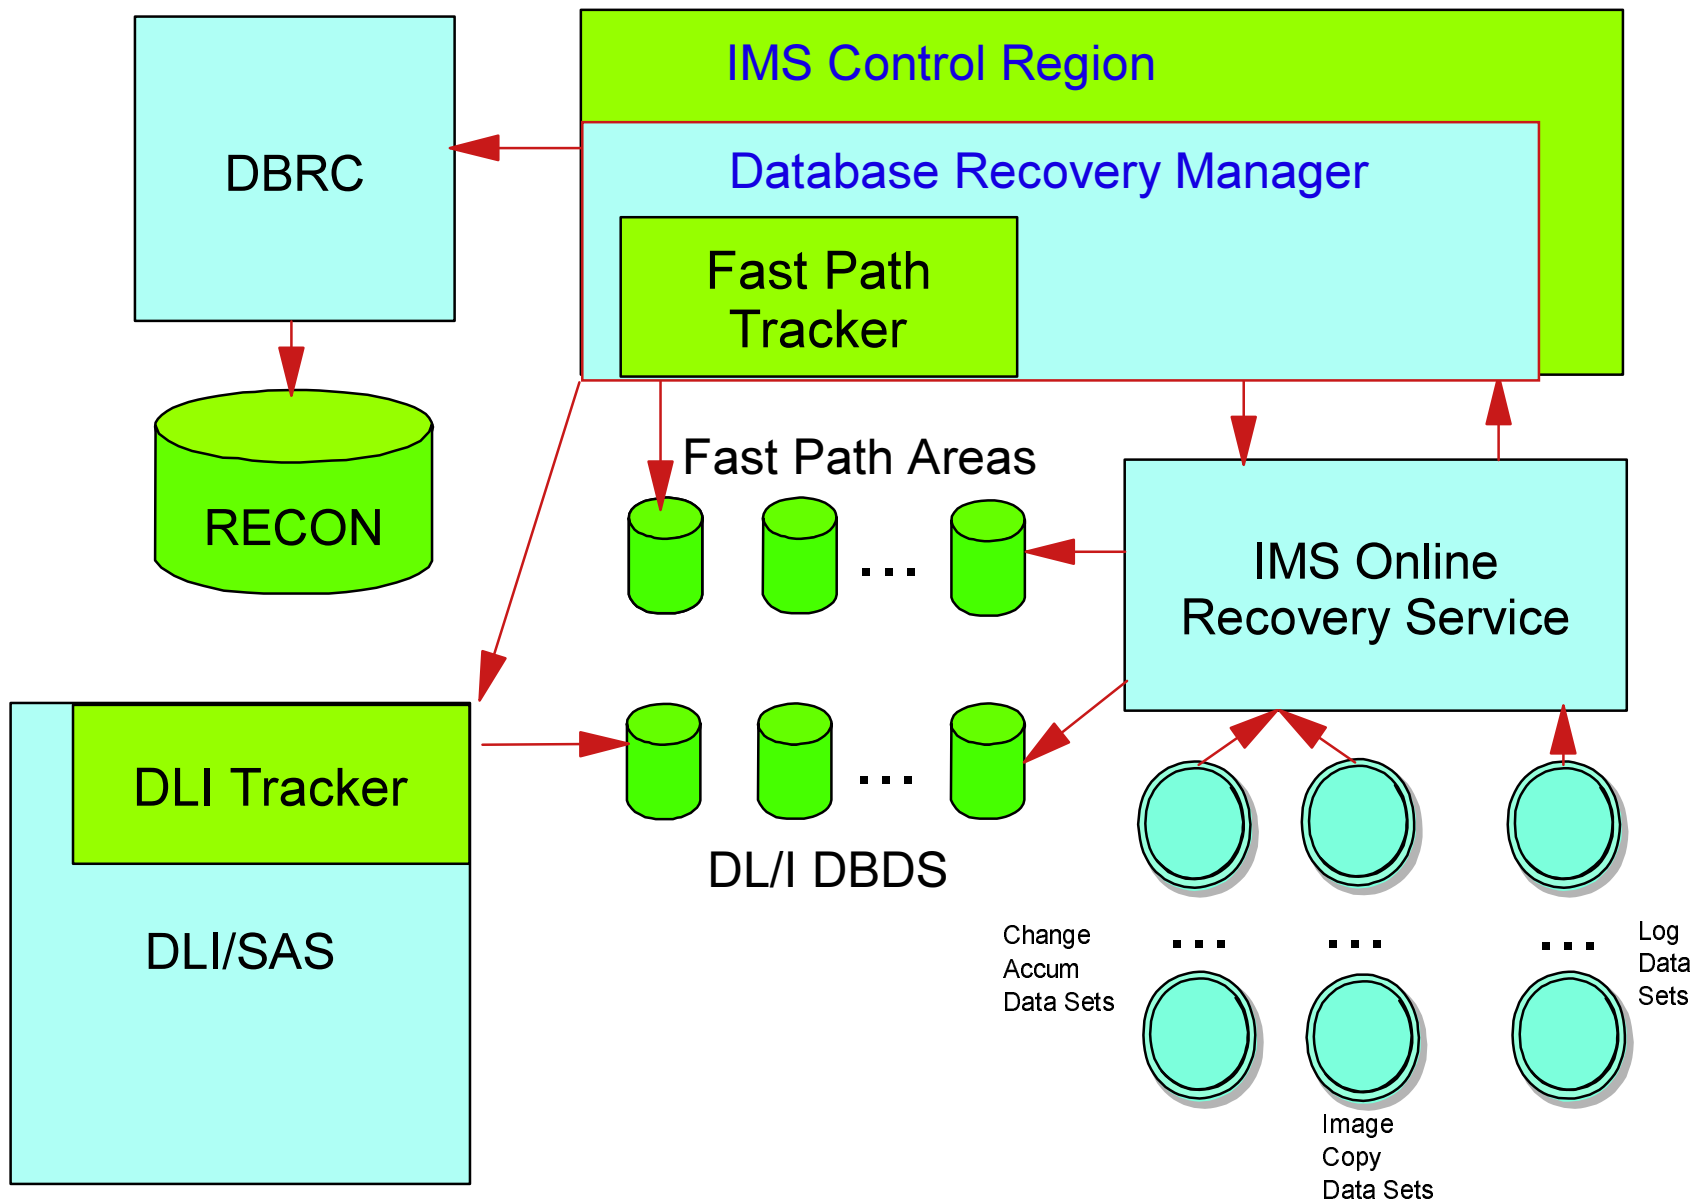
# Timestamp Recovery (continued)

- RECOVERY NEEDED is turned on in RECON for all DBDSs in a DB or partition that need to be recovered to the same point in time
- DBDS and areas are marked "image copy needed" in RECON if PITR timestamp recovery is performed
- ICs prior to recovery any point in time are not allowed for later (non-PITR) timestamp recoveries once the PITR timestamp recovery is hardened

# Components

- IMS Online Recovery Service is a new facility
  - ▶ A new address space which interacts with the IMS control region
- Changes in IMS to support IMS ORS consist of a new component and changes to DBRC and RSR Database Trackers:
  - ▶ Database Recovery Manager
    - Executes in the IMS control region
  - ▶ Fast Path Database Tracker
  - ▶ Full Function Database Tracker

# Components and Flow



# IMS Online Recovery Service

- Started when the first /RECOVER command is entered
  - ▶ Terminates on command
- Restores image copies, if required, with change accumulation data, if available
  - ▶ Image copies are restored in parallel
  - ▶ Change accumulation data sets are read in parallel

# IMS Online Recovery Service (continued)

- Reads log data and processes records related to this recovery
  - ▶ RLDS are read if available
- Merges log data into a single recovery stream
  - ▶ Sends the stream to the Database Recovery Manager

# Database Recovery Manager

- Driven by the IMS Command Processor to process recovery related commands
- Creates the IMS Online Recovery Service address space
  - ▶ When the first /RECOVER command is processed
- Establishes communication with the IMS Online Recovery Service address space



# Database Recovery Manager (continued)

- Receives log data from the IMS Online Recovery Service address space
- Passes updates to the database trackers
- Coordinates recovery termination and IMS Online Recovery Service participation in IMS shutdown processing

# DBRC

- Maintains recovery related information for
  - ▶ DL/I database data sets
  - ▶ Fast path areas
- All resources required for recovery must be registered with DBRC
- Validates recovery initiation and results
- Maintains a new group type: RECOVGRP

# PROCLIB Changes

- The IMS Online Recovery Service address space procedure must be specified in a PROCLIB available to the IMS control region
  - ▶ The IMS Online Recovery Service parmlib member is identified on the control region EXEC statement:  
ORSMBR=xx
  - ▶ The IMS Online Recovery Service RESLIB must be part of the IMS control region STEPLIB concatenation
    - The IMS Online Recovery Service RESLIB must be APF authorized

# PARMLIB Changes

- The DFSORSxx parmlib member contains specifications on
  - ▶ Number of available input devices
  - ▶ IMS Online Recovery Service proclib member name
  - ▶ DL/I data space size
  - ▶ FP data space size
  - ▶ Spill data space maximum size
- The BPECFG member contains specifications for IMS Online Recovery Service trace levels
- The IMS Online Recovery Service entry point must have an entry in the MVS Program Properties Table

# DFSORSxx PARMLIB Example

- **READNUM(10)**
- **RDMNM(DFSRDM00)**
- **DLIDSIZE(DSIZE(1000) REDO(1000))**
- **FPDSIZE(1000)**
- **SPSIZE(1000)**
  
- **READNUM** default is 3, range is 1-99
- **RDMNM** default is RDM
- **DLIDSIZE(DSIZE)** default is 15, range is 15 to 1600
- **DLIDSIZE(REDO)** default is 256, range is 128 to 4096
- **SPSIZE** default is 1000, range is 15 to 1600

# Sample PPT Entry

```
▪ PPT PGMNAME (FRDRVS00) /* IMS ONLINE RECOVERY SERVICE */
▪      CANCEL      /* PROGRAM CAN BE CANCELED */
▪      KEY (7)     /* PROTECT KEY ASSIGNED IS 7 */
▪      NOSWAP     /* PROGRAM IS NOT-SWAPPABLE */
▪      NOPRIV     /* PROGRAM NOT PRIVILEGED */
▪      DSI        /* DOES REQUIRE DATA SET INTEGRITY */
▪      SYST       /* PROGRAM IS A SYSTEM TASK */
▪      NOPASS     /* CAN BYPASS PASSWORD PROTECTION */
▪      AFF (NONE) /* NO CPU AFFINITY */
▪      NOPREF     /* NO PREFERRED STORAGE FRAMES */
```

# Sample RDM JCL

```
▪ //FRDRDM00 PROC
▪ //*****
▪ /*      IMS ONLINE RECOVERY SERVICE JCL      *
▪ //*****
▪ //STEP1      EXEC PGM=FRDRVS00 , PARM='ORS , BPECFG=ORSCONFG '
▪ //          REGION=2048K , TIME=1440
▪ //STEPLIB DD   DSN=IMSORS . SFRDRESL , DISP=SHR
▪ //PROCLIB DD   DSN=IMSORS . PROCLIB , DISP=SHR
▪ //SYSPRINT DD   SYSOUT=*
▪ //*
```

# Recovery via Command

- Recovery is initiated via command:
  - ▶ /RECOVER
- Commands can be issued from
  - IMS Master Terminal
  - MVS System Console
  - Logical Terminal
  - AOI Program
  - OTMA
  - APPC



# Recover Command

- /RECOVER ADD identifies DBDS and areas and builds a "recovery list"
  - ▶ A recovery list is the set of the DBDSs and Areas being recovered by one recovery instance
  - ▶ Multiple recovery lists can exist at the same time on the same IMS
- /RECOVER REMOVE eliminates database data sets and areas from the recovery list before recovery is started

# Recover Command (continued)

- /RECOVER START initiates recovery
  - ▶ Only one recovery list can be active (being recovered) at any one time
- /RECOVER STOP aborts recovery for one or more database data sets and/or areas
- /RECOVER TERMINATE shuts down the recovery environment

# Changed Commands

- /DISPLAY
  - ▶ /DISPLAY RECOVERY provides information on recovery activity
  - ▶ /DISPLAY DATABASE indicates to the user that the database is being recovered
  - ▶ /DISPLAY AREA notifies the user that the area is being recovered
  
- /STA, /DBD, /DBR, /LOCK
  - ▶ Rejected for DBDS and areas if they are being recovered by Online Recovery Service

# Changed Commands (continued)

- NOTIFY.RECOV
  - ▶ Allows the user to add information about a point in time recovery for a specific DBDS or area to RECON
  
- INIT.DBDSGRP
  - ▶ Specifies the RECOVGRP members for a recovery group
  
- CHANGE.DBDSGRP
  - ▶ Add or delete members of a recovery group

# Coexistence

- Log data used as input to recovery must be from IMS 5.1 or later
- Change accumulation data sets must be created by the IMS version 7 Change Accumulation utility