

# Fast Path

**ON DEMAND BUSINESS**

© 2008 IBM Corporation



## Fast Path Enhancements

- Command Enhancements
- Shared VSO enhancements
- Capacity enhancements
- Usability enhancement
- EMH enhancement



# Command Enhancements

## Starting All Areas with UPDATE DB command

- Option to start all areas when starting a DEDB
  - ◆ UPDATE DB NAME(name) START(ACCESS) AREA(\*)
    - AREA(\*) starts all areas of the database
      - There is NOT a capability to specify an area with AREA(...) parameter
    - Example TSO SPOC command and response:

```
UPDATE DB NAME(DEBJN) START(ACCESS) AREA( * )
```

| DBName  | AreaName | MbrName | CC |
|---------|----------|---------|----|
| DEDEBJN | *        | IMS1    | 0  |

- Benefit
  - ◆ Separate UPDATE AREA commands are not required for each area

4

This enhancements allows you to start all of the areas in a database when you also start the database. The only valid specification for the AREA parameter is AREA(\*).

If an area cannot be started, a separate line is provided for it in the response. Here is an example where area ARJN20 was not started:

```
UPDATE DB NAME(DEBJN) START(ACCESS) AREA( * )
```

| DBName  | AreaName | MbrName | CC |
|---------|----------|---------|----|
| DEDEBJN | *        | IMS1    | 8  |
| DEDEBJN | ARJN20   | IMS1    | A5 |

The 8 completion code for the database indicates that the command completed with errors for one or more areas of the database. The A5 completion code for area ARJN20 indicates that the DBRC 'Prohibit Further Authorization' flag was on for this area.

## Keeping Randomizer Resident when Stopping DB

- Option to keep the randomizer resident when stopping access to DB
  - ◆ UPDATE DB NAME(name) STOP(ACCESS) OPTION(NORAND)
    - OPTION(NORAND) does not unload the randomizer

```
UPDATE DB NAME(DEBJN) STOP(ACCESS) OPTION(NORAND)
```




- Benefit
  - ◆ Avoids ECSA fragmentation from unloading and reloading randomizers

5

In previous releases if you made a DEDB inaccessible, its randomizer was unloaded from memory if the randomizer was not being used by any other DEDB. DEDBs are made inaccessible with either a /DBR DB command or an UPDATE DB STOP(ACCESS) command.

IMS V10 adds an OPTION(NORAND) parameter to the UPDATE DB STOP(ACCESS) command. This causes IMS to keep the randomizer module in memory. This helps to avoid fragmentation of ECSA which can occur when modules are unloaded and reloaded.

The /DBR DB command does not have a parameter that is equivalent to OPTION(NORAND) on the UPD DB STOP(ACCESS) command.

## Keeping Randomizer Resident when Stopping DB

- New status STOACC
  - ◆ Response to QUERY DB Command
    - Indicates that the database is not accessible
      - It has been /DBRed or an UPD DB STOP(ACCESS) command has been issued
- New implication for Status RNL
  - ◆ Response to QUERY DB Command
    - Indicates that the randomizer for the database is not loaded
    - This status previously also indicated that the database was not accessible
- Benefit
  - ◆ STOACC and RNL separately indicate that the database is not accessible or the randomizer is not loaded

6

In conjunction with the OPTION(NORAND) parameter on the UPDATE DB STOP(ACCESS) command, a new status is available with the QUERY DB command. In previous releases one could tell that a database was inaccessible because the RNL status was associated with the database. RNL indicated that the randomizer module was not loaded. The absence of the RNL status no longer indicates that the database is accessible. A new status is available for this. The STOACC status in the response to a QUERY DB command indicates that the database is inaccessible. That is, the database has been /DBRed or an equivalent UPDATE DB STOP(ACCESS) command has been processed for the database.

The RNL status is now used only to indicate that the randomizer module is not loaded. The STOACC status is used to indicate that the database is inaccessible. The STOACC status is not returned by the /DISPLAY DB command. It is only returned by QUERY DB commands.

After you issued the UPDATE DB NAME(DEBJN) STOP(ACCESS) command, you would expect that the response to a QRY DB NAME(DEBJN) SHOW(STATUS) command would include:

| DBName | AreaName | MbrName | CC | TYPE | LclStat              |
|--------|----------|---------|----|------|----------------------|
| DEDBJN |          | IMS1    | 0  | DEDB | NOTOPEN, STOACC, RNL |

After you issued the UPDATE DB NAME(DEBJN) STOP(ACCESS) OPTION(NORAND) command, you would expect that the response to a QRY DB NAME(DEBJN) SHOW(STATUS) command would include:

| DBName | AreaName | MbrName | CC | TYPE | LclStat         |
|--------|----------|---------|----|------|-----------------|
| DEDBJN |          | IMS1    | 0  | DEDB | NOTOPEN, STOACC |

RNL is not included in the second response.



## Shared VSO Enhancements

## Shared VSO Private Buffer Pool Lookaside

- Previous releases for single area structures:
  - ◆ Lookaside must be specified in two places:
    - DBRC INIT.DBDS or CHANGE.DBDS command
      - LKASID or NOLKASID parameter
    - DEDB statement in DFSVSMxx member
    - If specifications conflict, DBRC value is used
- IMS V10
  - ◆ Lookaside specification in the DEDB statement is not required
- Multiple area structures have always ignored DBRC specification
  - ◆ Specification in DEDBMAS statement continues to apply
- Benefit
  - ◆ No longer required to specify parameter in DEDB statement that is ignored

8

In previous releases shared VSO users who had single area structures specified whether or not lookaside buffering would be used for an area in two places. It was specified in DBRC with an INIT.DBDS or CHANGE.DBDS command and also in the DEDB statement in the DFSVSMxx member as part of the private pool definition. If the specifications conflicted, the DBRC value was used. Only the DBRC value mattered, nevertheless, some value had to be specified in the DEDB statement. The requirement to specify a lookaside value in the DEDB statement has been eliminated in V10. You do not have to specify any value there.

The specification of lookaside for multiple area structures has not changed in V10. Multiple area structures were introduced in IMS V9. IMS continues to ignore any lookaside specification in DBRC for areas using multiple area structures. Only the specification in the DEDBMAS statement of the DFSVSMxx member is used.



## Shared VSO in XRF Tracking and FDBR Systems

- Previous releases:
  - ◆ XRF tracking systems and FDBR did not compress shared VSO private pools
    - Pools could expand but would not contract
- IMS V10
  - ◆ XRF tracking systems and FDBR compress shared VSO private pools
    - Compression done when expanded space is not used for two system checkpoint intervals
- Benefit
  - ◆ Reduced consumption of ECSA space

9

In previous releases of IMS when shared VSO was used with XRF or FDBR, the XRF tracking system or FDBR system did not compress the shared VSO private pools. These pools can expand as more buffers are required for the areas in the pools. When the buffers are no longer needed, XRF active systems and non-XRF systems compress the pool by releasing buffers that were acquired for expansion. IMS V10 has enhanced XRF tracking systems and FDBR by compressing their pools just as non-XRF and active XRF systems compress pools. This will potentially reduce the use of ECSA by XRF tracking systems and FDBR for shared VSO users.



## /START AREA GLOBAL Enhancement

- GLOBAL parameter enhancements for /START AREA command
  - ◆ For PREOPEN areas
    - Preopen will occur on all systems
      - Previous releases only preopened the system on which the command was entered
  - ◆ For PRELOAD areas
    - Sharing systems preload their buffer pool
      - Preload limited to the primary allocation
- Benefit
  - ◆ Simplifies operations for systems using data sharing



# Capacity Enhancements

## Maximum Number of FP Buffers Increased

- Previous releases:
  - ◆ Maximum number of FP buffers was 65,535
- IMS V10
  - ◆ Theoretical limit is 4,294,967,295 FP buffers
    - Execution parameter DBBF= accepts values up to 4,294,967,295
    - Some log records have been modified to accommodate the new maximum
      - IMS maintains number of buffers in full word (4 bytes)
- Benefit
  - ◆ Fast Path can exploit large capacities of new processors

12

The maximum number of fast path buffers that can be used has been increased from 65,535 to a theoretical value of 4,294,967,295. This is a theoretical value since a system is unlikely to have the storage for so many buffers. Any value up to the theoretical limit may be specified on the DBBF= execution parameter for the online system. In previous releases you could specify the number of Fast Path buffers on the FPCTRL macro. This macro is not used in IMS V10. The DBBF= execution parameter must be used to specify these buffers in V10.

The log records that have been modified include the following that are mapped by the DBFLSRT macro: 5950, 5951, 5952, 5953, 5953, 5956, and 5957.

Other modified records are mapped by the DBFLGRSD macro. They include: 5955 and 5958.

## Maximum Number of FP Output Threads Increased

- Previous releases:
  - ◆ Maximum number of FP output threads was 255.
    - OTHR cannot exceed MAXPST value.
- IMS V10
  - ◆ Maximum number of FP output threads is 32,767.
- Benefit
  - ◆ Fast Path can exploit large capacities of new processors

13

In previous releases the maximum number of Fast Path output threads was the smaller of the MAXPST value or 255. In V10 the maximum is the 32,767.

This allows IMS V10 users to have many more output threads. This could be desirable as processor capacities increase.

There is another enhancement for Fast Path users in IMS Version 10. In previous releases most Fast Path control blocks and buffers were created in one module, DBFCONT0. This module resides in ECSA and it could be very large. IMS V10 places these control blocks and buffers in five modules. This allows installations to more efficiently use ECSA storage since smaller contiguous areas of ECSA may be used.

The major new modules and their contents are:

DBFCONT1: contains the ECNTs, MSDBs, and MSDB blocks

DBFCONT3: contains buffer headers (DHMRs) and buffers

DBFCONT4: contains DEDB blocks

DBFCONT5: contains output threads

DBFCONT6: contains BALGs (balancing groups)

## FDBR FPBUFF= Parameter Default

- FPBUFF= parameter in DFSFDRxx member controls location of FP buffers
  - ◆ Previous releases
    - FPBUFF=LOCAL placed DBFCONT0, HSSP, and HSRE buffers in extended private
      - DBFCONT0 includes database buffer pool (DBBF) and SVSO buffer pools
      - Recommended
    - Default was to place these blocks and buffers in ECSA
  - ◆ IMS V10
    - FPBUFF=LOCAL is the default
      - Places DBFCONTx modules and HSSP and HSRE buffers in extended private
      - Recommended
    - FPBUFF=ECSA places these modules and buffers in ECSA
      - Not recommended

14

APARs PQ97745 for IMS V9 and PQ80264 and PQ97043 for IMS V8 added a new parameter in the DFSFDRxx PROCLIB member for FDBR. FPBUFF=LOCAL could be specified to place module DFSCONT0 and HSSP and HSRE buffers in extended private. There was no other valid value for FPBUFF=. If the parameter was not specified, the module and buffers were placed in ECSA.

IMS V10 changes the default. If the FPBUFF= parameter is not specified, DBFCONT1, DBFCONT3, DBFCONT4, DBFCONT5, DBFCONT6 and HSSP and HSRE buffers placed in extended private. A specification of FPBUFF=ECSA places these modules and buffers in ECSA. A specification of FPBUFF= LOCAL and allowing the parameter to default is recommended. There is no advantage to placing the modules and buffers in ECSA.



# Usability Enhancement

## DFSVSMxx Specification for QUITCI

- New DFSVSMxx parameter
  - ◆ SDEPQCI= S | D | A
    - S - Force QUITCI by Scan utility
    - D - Force QUITCI by Delete utility
    - B - Force QUITCI by both Scan and Delete utility
  - ◆ If SDEPQCI= is not specified, utility control statement determines the use of QUITCI
    - No change from previous releases
- Benefit
  - ◆ Eliminates need to specify a QUITCI control statement on each Scan or Delete utility execution
  - ◆ May be used to enforce the use of QUITCI

16

QUITCI is an optional parameter specified in the SYSIN data set for the DEDB SDEP Scan and Delete utilities. It is used in data sharing environments to cause other IMS systems to release their current SDEP CI and any preallocated SDEP CIs. Any partially filled SDEP CIs are written to the area data set. This simplifies the running of the Scan and Delete utilities by making it easy to ensure that they process all committed SDEPs without waiting on CIs to be filled.

Some users want to ensure that QUITCI is used by all executions of the Scan and/or Delete utilities. The new SDEPQCI parameter in the DFSVSMxx member may be used to enforce the use of QUITCI, even when it is not specified on the SYSIN input to the utilities.





# EMH Enhancement



## EMH enhancement in IMS V10

- Command support to reset Fast Path Response Mode
  - ◆ In previous releases, hung FP response mode nodes and users typically require an IMS restart (or complex sequence of commands)
  - ◆ IMS 10 allows hung nodes and users to be reset without bringing IMS down
    - Applies to input response mode or output response mode with either local or shared EMH
    - Commands can be issued at any time
      - /STO NODE and /STA NODE for static nodes
      - /STO USER and /STA USER for dynamic users
        - Reset performed during /STA processing
- Enhanced DISPLAY command to show Fast Path input response mode
  - ◆ Hung users are typically in 'input' response mode, not 'output' response mode
  - ◆ New status – RESP-INP-FP

18

This EMH enhancement addresses a rarely occurring situation; however, when it does occur, an IMS outage is likely needed to correct the problem.

Full function users have been able to reset response mode for hung terminals via the /RST, /STO, and /STA commands; this enhancement extends this support to Fast Path users. You cannot use the /RST command to reset response mode for Fast Path nodes/users as you can for full function users. This is because Fast Path requires a session to be in a quiesced state and /RST does not have this requirement.

RESP-INP-FP is a new status in IMS Version 10 that may be returned by the following /DIS commands

/DIS NODE, /DIS NODE RESPINP

/DIS USER, /DIS USER RESPINP

/DIS STATUS, /DIS STATUS NODE, /DIS STATUS USER

Input response mode is now indicated by two statuses, RESP-INP-FP for Fast Path, and RESP-INP for full function. If you have automation dependent on the RESP-INP status and use Fast Path EMH, your automation needs to be changed to handle both input response mode statuses.

## Fast Path Enhancements Summary

- Command Enhancements
  - ◆ Improved UPD DB and QRY DB commands
  - ◆ Improved use of ECSA
- Shared VSO enhancements
  - ◆ Simplified definition
  - ◆ Improved use of ECSA by FDBR and XRF tracking systems
- Capacity enhancements
  - ◆ Increased output thread and buffers
  - ◆ Improved use of ECSA
- Usability enhancement
  - ◆ Simple enforcement of QUITCI for SDEP Scan and Delete utilities
- EMH enhancement
  - ◆ Capability to reset Fast Path response mode