



| z/TPF V1.1

TPF Users Group – 2011

SCP Updates

| Lisa Banks

| SCP Subcommittee

AIM Enterprise Platform Software
IBM z/Transaction Processing Facility Enterprise Edition 1.1.0

Any reference to future plans are for planning purposes only. IBM reserves the right to change those plans at its discretion. Any reliance on such a disclosure is solely at your own risk. IBM makes no commitment to provide additional information in the future.

© 2011 IBM Corporation

Agenda

- Key Enhancements
- Online Database Reorganization
- General Enhancements

PUT 8 Key Enhancements

- High Performance FICON Support
 - PJ38360
- Dump Buffer Statistics
 - PJ38649
- Input List Monitor
 - PJ37952
- Online Database Reorganization
 - PJ39198

High Performance FICON (HPF) for DASD

• Functions

- Command mode structures (CCWs) are replaced by Transport mode structures (TCWs) in HPF.
- These new structures block commands, data, control information and status together to reduce channel and controller overhead.
- z/TPF will employ a translator program to convert CCW format I/O operations into TCW format when the hardware supports HPF

• Benefit

- High Performance FICON offers a 2x to 3x performance improvement over traditional FICON
- Channel and CU performance improvements.
- Full exploitation of Solid State DASD when using High Performance FICON
- Improved error recovery (concurrent sense buffer returned on error)
- Device status for missing interrupts now available via interrogate command

• Dependencies

- DS8000 with HPF Phase 0 micro code
- z10 or higher processor
- FICON Express-2 or higher

Dump Buffer Area (DBA) Statistics

- **Functions**

- Statistical information about the usage of the DBA

- **Benefit**

- Ability to appropriately size the DBA
- Improved availability by reducing time needed to write a dump when there is adequate space available in the DBA

Input List Monitor

- **Functions**

- Detects when system is in input list shutdown for a user-specified time
- Sends a warning message to Prime CRAS

- **Benefit**

- Will allow customers to react proactively to determine the cause of the shutdown, which should improve system uptime

Online Database Reorganization (ODBR)

TPFUG Requirement # D06001F

- **PJ39198**
- Online database reorganization (ODBR) supports the relocation of online data records resident on a z/TPF complex while maintaining system availability.
- This was the number one database requirement when last on the ballot

Online Database Reorganization (ODBR)

Key Features

- Ability to move data while the system is online and running traffic on all processors in the complex
- Shadow copy for updates made to data
- Validation of moved data

Online Database Reorganization (ODBR)

How does it help?

- Reduction in cost - minimize the need for downtime during the performance of a database reorganization .
- Reduction in cost - do not need to purchase DASD over what you are adding to the system.
- Reduction in system resources - eliminates the need to use tapes for record capture, and exception recording and logging during the reorganization process.
- Reduces risk associated with DASD expansions by separating hardware installation from database reorganization.

Online Database Reorganization (ODBR)

How do you use it?

• STEP 1

- Build an alternate FACE table (FCTB)
 - Select records would wish to move
 - Modify RAMFIL statements in SIP input deck
 - Target areas must be defined as IFORMAT or SPARE
 - Target areas must have a record size that is the same as the records being moved to that area.

Online Database Reorganization (ODBR)

STEP 1

Base FCTB

```
RAMFIL TYPE=4SA,RECNO=00256,RECID=#HOTREC,DUPE=YES,  
          UFTI4=(61,117),UFTI5=(51,117),EQU=216  
RAMFIL TYPE=4SA,RECNO=005000,RECID=#IBMLOST,DUPE=YES  
          UFTI4=(12,42),UFTI5=(11,42),EQU=315  
RAMFIL TYPE=4SA,RECNO=056777,RECID=SPARE,DUPE=YES
```

Alternate FCTB

```
RAMFIL TYPE=4SA,RECNO=00256,RECID=SPARE,DUPE=YES  
RAMFIL TYPE=4SA,RECNO=005000,RECID=#IBMLOST,DUPE=YES  
          UFTI4=(12,42),UFTI5=(11,42),EQU=315  
*RAMFIL TYPE=4SA,RECNO=056777,RECID=SPARE,DUPE=YES  
RAMFIL TYPE=4SA,RECNO=056521,RECID=SPARE,DUPE=YES  
RAMFIL TYPE=4SA,RECNO=00256,RECID=#HOTREC,DUPE=YES,  
          UFTI4=(61,117),UFTI5=(51,117),EQU=216
```

Online Database Reorganization (ODBR)

How do you use it?

- **STEP 2**

- Compare base and alternate FCTB
 - Offline tool (**fctbval**)

- **STEP 3**

- Load the alternate FCTB (**ZODBR LOAD**)
 - Identifies the data that is being moved
 - Validates the consistency between the FCTBs
 - Validates the track formats of the new areas

Online Database Reorganization (ODBR)

How do you use it?

- **STEP 4**

- Ensure all processors are on the same image

- **STEP 5**

- Start the ODBR utility (**ZODBR START**)
 - Moves all data whose location has changed in the FCTBs
 - Ensures data integrity
 - Shadow copy
 - Serialization of data updates
 - Validation of the moved data

Online Database Reorganization (ODBR)

How do you use it?

- **STEP 6**

- Optionally display schedule (**ZODBR DISPLAY**)

- **STEP 7**

- Once complete, deactivate other processors

- **STEP 8**

- Accept the alternate FCTB (**ZODBR ACCEPT**)
 - Automatic IPL

Online Database Reorganization (ODBR)

Additional Features

- Ability to pause and restart
- Ability to cancel prior to accepting alternate FCTB
- Perform data validation automatically or manually (**ZODBR VALIDATE**)
- User profile settings (**ZODBR PROFILE**)

Online Database Reorganization (ODBR)

Considerations

- Data moving and validation is done by one processor
- An IPL is performed when alternate FCTB is accepted
- Shadow copy applies to updates made using standard write methods
- The following are mutually exclusive with ODBR
 - Database reorganization (DBR)
 - Online DASD formatter
 - File copy/capture/restore
- For the best results avoid running during peak hours or DASD-intensive operations

Online Database Reorganization (ODBR)

Where do we go from here?

- Remove IPL requirement
- Allow manual validation to be paused, resumed and canceled
- Automatic pause of ODBR in the event of errors

PUT 8 General Enhancements

- New ZASER option
 - **PJ38179**
 - The ZASER HEAPBUF command provides a way to limit the number of heap buffers in a dump

PUT 8 General Enhancements

- Dump enhancement
 - **PJ38500**
 - The failing ECB trace name needs to be added to the dump

PUT 8 General Enhancements

- Improve response time and overhead for ZDHST
 - **PJ39329**
 - An entry in the database no longer uses text keywords. Instead, an entry is binary data.
 - Reduce the number of mallinfo() calls.
 - Give up control more frequently to reduce the impact on other work that is being done
 - Processing has been moved off the main I-stream to the highest I-stream.
 - Some filtering is now done at file selection time for
 - ZDHST DUMP/NODUMP/TYPE/PROC

PUT 8 General Enhancements

- Performance enhancement for next generation processors
 - **PJ39153**
 - Convert several STCK instructions to STCKF instructions if store-clock-fast facility is installed

PUT 8 General Enhancements

- Performance Analyzer (PA) enhancement
 - **PJ38334**
 - To improve ease of transferring the PA trace file from the z/TPF file system to the TPF Toolkit, the PA was changed to create the trace file in the subsystem of the application being traced.

PUT 8 General Enhancements

- New C/C++ function for ECB heap information
 - **PJ38648**
 - A new function that is named `tpf_eheapc_info_fordblks()` has been created to return the largest `malloc()` request that can be done for this ECB
 - This function returns the same information returned by the `fordblks` value from `mallinfo()`
 - When `fordblks` is the only data needed from a `mallinfo()`, using `tpf_eheapc_info_fordblks()` in frequently used code can improve performance.

PUT 8 General Enhancements

- New user exit in record hold table (RHT) processing
 - **PJ39312**
 - Provide a choice in how to proceed when the RHT wait queue threshold is exceeded
 - User can choose to exit the ECB or allow the request to be queued

PUT 8 General Enhancements

- New user exit in multiple release detection (MRD)
 - **PJ38981**
 - Provide the ability for users to prevent MRD processing from updating the record that is being released

PUT 8 General Enhancements

- ECB Resource Monitor enhancements
 - **PJ38501**
 - An ECB resource monitor user exit macro, UECBRC, was added to internal macro ECBRC that you can use to change the logic for counting resource use to meet your installation needs
- Add new macros to “CREATE” type
 - SWISC CREATE
 - CRESC
 - FORKC

PUT 8 General Enhancements

- Support special Server Time Protocol (STP) configuration
 - **PJ39209**
 - The special STP configuration provides an alternate configuration that you can use to keep the time-of-day (TOD) clock synchronous between all processors in a loosely coupled complex.
 - In this configuration, each physical processor in the loosely coupled complex uses STP, but the z/TPF system is not aware that the processors are using STP.
 - This allows each z/TPF system logical partition (LPAR) to be in local timing mode.

PUT 8 General Enhancements

- Enhance ZRFPC to better detect and repair corruption in the short-term common control record (STCCR) and the short-term processor control record (STPCR) record
 - **PJ37533**
 - ZRFPC processing repairs the control records when a checksum error is detected, so that the system can be cycled up immediately.
 - The STCCR records will be repaired if the 'duplicate directories found' or 'timestamps lost' errors are detected.

PUT 8 General Enhancements

- Update to I/O lock caching
 - **PJ38788**
 - If record locks are kept in DASD control units, there is logic to cache an unlock request (UNFRC) rather than driving the unlock I/O immediately. On the next I/O request to DASD, the CCW chain will include the cached unlock. This has the benefit of reducing I/Os to DASD
 - Update made to DASD read queue logic for cached unlocks to primes and dupes

PUT 8 General Enhancements

- Vendor enhancements
 - **PJ38312**
 - Provide a 16MEG area below the bar for vendor use.
 - Add a vendor user exit during restart
 - Add a copy member (CVSR) to the CCUEXT CSECT for vendor use.

QUESTIONS?

Trademarks

- IBM and TPF Toolkit are trademarks of International Business Machines Corporation in the United States, other countries, or both.
- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Other company, product, or service names may be trademarks or service marks of others.
- **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.
- 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.