

z/TPF EE V1.1

z/TPFDF V1.1

TPF Toolkit for WebSphere® Studio V3

TPF Operations Server V1.2



IBM Software Group

TPF Users Group Spring 2007

TPFDF and z/TPFDF Update

Name: Kevin Jones

Venue: Database Subcommittee

AIM Enterprise Platform Software

IBM z/Transaction Processing Facility Enterprise Edition 1.1.0

© IBM Corporation 2007

Any references 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.

TPFDF PUT 22

- Scheduled to be available in June 2007
- Approximately 48 APARs including 4 enhancements:
 - PK19745 – MLS support for the High Level Assembler Release 5 (HLASM R5)
 - PK26312 – Automatic de-indexing
 - PK34634 – Support for z/OS 1.8 compiler
 - PK33468 – Reuse pools when replacing a Large Logical Record (LLR)

Automatic De-indexing Enhancement

- Automatic de-indexing is an enhancement to basic indexing
- Prior to this support, when a detail subfile became empty, it was the application's responsibility to de-index on each path and release the pool records
 - De-indexing requires the algorithm string that was used when the subfile was indexed. TPFDF and z/TPFDF did not maintain these strings, making any type of automatic de-indexing impossible.
- Automatic de-indexing introduces an optional *Algorithm Information File* to maintain algorithm strings

Automatic De-indexing Enhancement

- New DBDEF parameters are available to associate a detail file with an *Algorithm Information File*, and to enable automatic de-indexing
- It is possible for an *Algorithm Information File* to be maintained without enabling automatic de-indexing, allowing other applications to access the algorithm strings
- IBM is not providing a means of populating the *Algorithm Information File* for existing databases
 - recommended only for new databases and applications

Automatic De-indexing Enhancement

- Also included is a series of new DBDEF parameters that allow DSECT options to be more easily overridden
 - For example, FULLBCH overrides the SW02OP1 #BIT0 setting in the application DSECT, or the OP1 parameter on the DBDEF
- Support is now available electronically
 - APAR PK26312 planned for TPFDF PUT 22
 - APAR PK25978 on z/TPFDF PUT 3
- Satisfies TPFUG Requirement DF00073

Reuse Pools When Replacing LLRs

- Previously, when a Large Logical Record (LLR) was replaced (DBREP/dfrep), all associated pool records would be released, and new pools obtained
 - a potentially large and unnecessary use of long-term pools
- TPFDF and z/TPFDF have been enhanced to reuse pool records when replacing LLRs
 - If the LLR is smaller, excess pools are still released
 - If the LLR is larger, additional pools are obtained
- Does not apply to DETAC mode
 - reusing pools would prevent updates from being rolled back using “close abort”

Reuse Pools When Replacing LLRs

- A sequence number in each block is used to ensure that any ECB reading the LLR receives consistent data
 - If a sequence number mismatch is detected while reading an LLR, then the record is being replaced
 - TPFDF will reattempt the “read” up to 5 times before returning to the application with an error (SW00RT3, #BIT6)
 - TPFDF APIs that flag LLR errors will detect this new condition upon application reassembly or recompilation

Reuse Pools When Replacing LLRs

- Support is now available electronically
 - TPFDF 1.1.3 customers must install:
 - APAR PK33468 planned for TPFDF PUT 22, and
 - APAR PJ31705 planned for TPF PUT 21
 - z/TPFDF customers must install:
 - APARs PK33480 and PK42403 planned for z/TPFDF PUT 4, and
 - APARs PJ31724 and PJ31945 planned for z/TPF PUT 4

Question and Answer



Trademarks

IBM is a trademark of International Business Machines Corporation in the United States, other countries, or both.

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.