z/TPF V1.1

TPF Users Group - 2011
Title: z/TPFDF Status Update

Subtitle:

Name: Daniel Jacobs
Venue: Database Subcommittee

Agenda

- **Recent z/TPFDF Enhancements**
  - z/TPFDF PUT 07 Enhancements
  - z/TPFDF PUT 08 Enhancements
- **Future z/TPFDF Enhancements**
- **z/TPFDF TPFUG Requirements Update**
  - New z/TPFDF Requirements from Fall 2010 Users Group
  - Top 10 z/TPFDF Requirements from Fall 2010 Voting
  - z/TPFDF TPFUG Requirements Summary

Recent z/TPFDF Enhancements

Recent z/TPFDF Enhancements

- **Recent z/TPFDF enhancements have focused on:**
  - Migration assistance
    - For example, eliminating customer modifications
  - Performance improvements
  - Smaller items with faster turnaround to market
- **In past 2 years, 10 customer requirements have been made available**
  - Addresses requirements from 6 different customers
- **APARs to address 5 more customer requirements are in progress**

z/TPFDF PUT 07 Enhancements

- **PK90125: Expand z/TPFDF Data Collection Counters (PUT 6)**
- **PK95973: ZUDFM command to delete info from MLS DB**
- **PK97406: Formatted SW00SR in C/C++ application dumps**
- **PK60025: New Hashing Algorithm #TPFDB11**
- **PK97405: Central database routine path length reduction**

PK90125: Expand z/TPFDF Data Collection Counters

- **Previously, data collection rates for an API limited to 999,999.9 calls per second**
- **Multiple customers now exceeding that rate!**
- **Display changed to support up to 99,999,999 APIs per second**
  - Decimal place eliminated
- **Shipped on z/TPFDF PUT 06**

PK90125: Expand z/TPFDF Data Collection Counters

```
User:    ZUDFC DISPLAY

System:  UDFC0000I 16.41.23 TPFDF DATA COLLECTION
         02OCT09  09.30.00-09.34.59  - AVERAGE PER SECOND  - TOTAL

         DBADD:    64890 DBADR:        2 DBCKP:      523 DBCLS:    32495
         DBCPY:       41 DBCRE:      773 DBDEL:     6439 DBDIX:        0
         DBDSP:       -- DBIDX:       11 DBKEY:    20140 DBMOD:    30525
         DBMRG:        0 DBOPN:    32495 DBRED:   542822 DBREP:     4637
         DBRET:    25105 DBRST:       29 DBSPA:     5854 DBSRT:       39
         DBTLD:       -- DBTLG:       66 DBTRD:       -- DBUKY:      495

         FILNC:     1581 FILEC:        0 CFILE:     1050 PFILE:     4136
         TWRTC:      114 GETFC:     1841 RELFC:      986 DETAC:    73529
         ATTAC:    74686 GETCC:     8366 RELCC:    42468 CFIND:     9903
         PFIND:    33926


         CURRENT DATA COLLECTION COUNTS :

         DBADD:18558658 DBADR:      466 DBCKP:   149696 DBCLS: 9293619
         DBCPY:    11605 DBCRE:   220954 DBDEL: 1841490 DBDIX:      57
         DBDSP:       -- DBIDX:     3207 DBKEY: 5760123 DBMOD: 8730058
         DBMRG:        3 DBOPN: 9293619 DBRED:******** DBREP: 1326048
         DBRET: 7179922 DBRST:     8330 DBSPA: 1674165 DBSRT:   11207
         DBTLD:       -- DBTLG:    18924 DBTRD:       -- DBUKY:  141498

         FILNC:   452253 FILEC:       89 CFILE:   300185 PFILE: 1182816
         TWRTC:    32675 GETFC:   526491 RELFC:   281875 DETAC:21029152
         ATTAC:21360076 GETCC: 2392652 RELCC:12145885 CFIND: 2832157
         PFIND: 9702875


         *** END OF DISPLAY ***
```

PK95973: ZUDFM command to delete info from MLS DB

- **Previously, no way to remove DSECT information from a Macro Label Set (MLS) database**
- **New DEL parameter provided to delete info from database**
  - By DSECT name
  - By application type
  - All DSECTs

PK97406: Formatted SW00SR in C/C++ application dumps

- **Previously, formatted SW00SR only included in dump when R3 contained valid SW00SR address**
  - Dumps occurring in C/C++ applications don't have R3 set up
- **Dump formatter now also checks if the last used SW00SR is set up in the DBIFB block**
- **Co-requisite z/TPF APAR: PJ36878**

PK60025: New Hashing Algorithm #TPFDB11

- **Existing #TPFDB09 is an 8-byte hashing algorithm**
  - May not provide even distribution if:
    - Algorithm argument often contains small subset of possible hexadecimal data
    - Large number of hash slots exist
  - For example:
    - Using pairs of 4-byte location codes (like “MIAMSANF”)
    - Only 100 different cities
    - 10,000 hash slots
    - Results in only a small subset of the 10,000 slots being used

PK60025: New Hashing Algorithm #TPFDB11

- **New #TPFDB11 also is an 8-byte hashing algorithm**
  - 8-byte hashing algorithm to address above issue
  - May not provide even distribution if:
    - Algorithm string has disproportionately high number of same value for either low or high-order 4 bytes
    - Example: 8-byte file addresses in 4x4 mode
  - Users should run tests to find best algorithm for their needs

PK97405: Central DB Routine Path Length Reduction

- **On TPF 4.1, several central DB routines exceeded 4K**
  - Required overflow segments
- **On z/TPFDF, many overflow segments recombined into a single routine**
  - Reduces path length on various API calls
  - Exploits baseless technology
- **May improve performance**

z/TPFDF PUT 08 Enhancements

- **PK80800: Alter ZUDFM DISP LAST & ZUDFM PACK output**
- **PM15783: Reduce heap storage usage for SW00SR**
- **PM05719: Switch z/TPFDF SDO DAS to APACHE Tuscany**
- **PM30434: C/C++ API path length reduction**
- **PM26537: Allow ZUDFM INIT to run faster**
- **PM26891: Change pack inhibit function for recoup**
- **PM34934: Create header file for SR05SR**
- **PM38645: B+Tree FARF6 Support**
- **PM37761: B+Tree/LLR file/find counters in Data Collection**

PK80800: Alter ZUDFM DISP LAST & ZUDFM PACK output

- **Requirements opened to:**
  - Display LREC number on ZUDFM DISPLAY LAST
    - Allows user to see how many LRECs are in subfile
  - Don’t display all LRECs in subfile after ZUDFM PACK
- **With this APAR:**
  - Display LREC number when ZUDFM DISPLAY LAST entered
  - Display “LAST” when last LREC included in ZUDFM DISPLAY
  - Display relative numbers of added/replaced LRECs in subfile with default keys on subsequent ZUDFM re-display
  - Repeated ZUDFM modification entries are permitted
  - Display “PACK COMPLETE” after ZUDFM PACK instead of re-displaying entire subfile

PM15783: Reduce heap storage usage for SW00SR

- **Previously, DBIFB always obtained the storage needed for the max number of SW00SR entries**
  - Based on #TPFDBMO value in macro ACPDBE
  - Every ECB would have to allocate ***and*** clear all of that storage
  - Wastes a lot of storage and is inefficient

PM15783: Reduce heap storage usage for SW00SR

- **More efficient and dynamic approach taken**
  - Initial number of SW00SR slots allocated
    - Based on new equate #TPFDBMI in macro ACPDBE
    - User settable
    - Default value is 8
  - When additional slots are needed, more slots allocated based on #TPFDBMI
  - MALOC used instead of CALOC
    - Storage only cleared as necessary
  - May improve performance

PM05719: Switch z/TPFDF SDO DAS to Apache Tuscany

- **Previously, z/TPFDF Data Access Service (DAS) used Service Data Objects (SDO) implementation based on 1.0 specifications**
- **z/TPFDF DAS client code now uses Apache Tuscany SDO implementation**
  - Implements SDO 2.1 specification
    - Improved capabilities over 1.0
    - Makes development simpler and maintenance easier
  - Utility APIs further extend the functionality
- **No migration impact for existing applications**

PM30434: C/C++ API path length reduction

- **Previously, z/TPFDF C/C++ interface code (CTDF) was defined as a standard library**
- **Has been redefined as a special linkages library**
  - Streamlines processing done on every z/TPFDF C/C++ API call
  - Path length of these API calls is reduced
  - May improve performance
- **Co-requisite z/TPF APAR: PJ38520**

PM26537: Allow ZUDFM INIT to run faster

- **ZUDFM INIT processing can take a long time for large files**
  - Bottleneck is obtaining holds on records before initialization
- **NOHOLD option provided with ZUDFM INIT**
  - Uses GETCC/FILEC processing instead of FIWHC/FILUC
  - Applicable for new databases or when there is no application activity
  - Default value is HOLD

PM26891: Change pack inhibit function for recoup

- **Packing is inhibited during recoup if:**
  - PACKINHI=YES is specified in DBDEF, or
  - PACKINHI=COND is specified in DBDEF and the file contains forward references
- **Previously, packing would be inhibited for the duration of TPFDF recoup**
  - Was only a part of Phase 1 recoup on TPF 4.1
  - Now spans entire duration of Phase 1 recoup on z/TPF

PM26891: Change pack inhibit function for recoup

- **More granular mechanism introduced to reduce time of pack inhibiting**
  - 2 new DBDEF parameters:
    - DBTAG: identifies z/TPFDF files that are part of same database structure
      - Defined in macro idbtag.mac
    - MULTI: Used with MPNXTD and MPRECD to process all top level index files with same DBTAG value sequentially with packing inhibited
  - With DBTAG specified, packing will only be inhibited while the database structure the file is part of is chain-chased

PM26891: Change pack inhibit function for recoup

- **DBTAG=0 results in no change from previous z/TPFDF behavior**
  - This is the default
- **Co-requisite z/TPF APAR: PJ38381**

PM26891: Change pack inhibit function for recoup

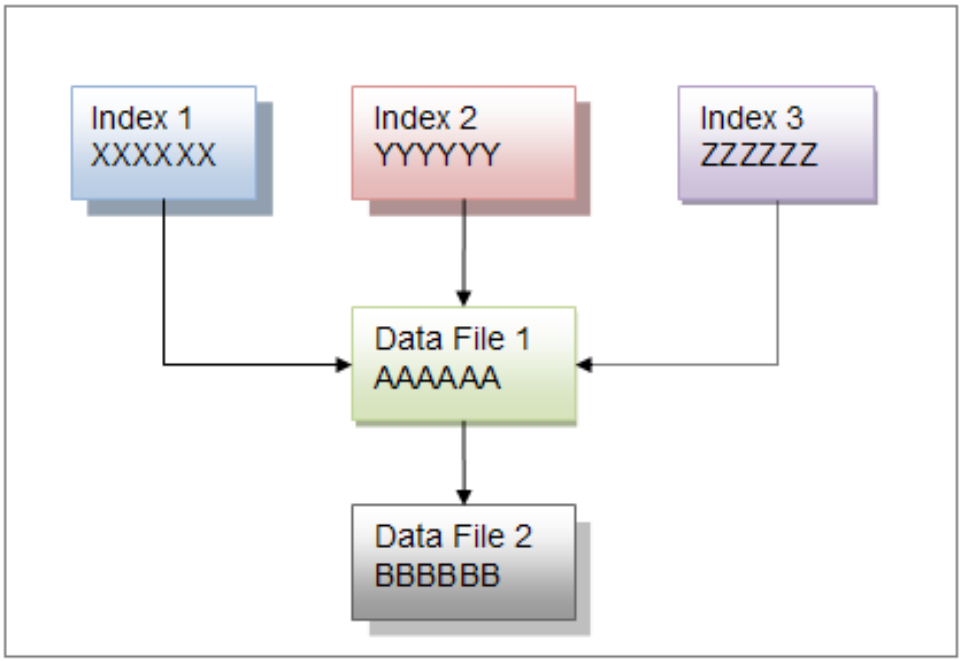

- User wants each index and data file 1 to be pack inhibited
- Assign the same database tags to these files
- Ensure indexes are processed sequentially

PM26891: Change pack inhibit function for recoup

```
DBDEF FILE=XXXXXX,PACKINHI=YES,DBTAG=(#UTAG257,MULTI),          X
      RCIDID=AA,MPNXTD=YY,MPPRCD=PRIME,                         X
      (ITK=X'80',ID2=(RCI),INDEX=(AAAAAA,0))
DBDEF FILE=YYYYYY,PACKINHI=YES,DBTAG=(#UTAG257,MULTI),          X
      RCIDID=AA,MPNXTD=ZZ,MPPRCD=PRIME,MPRECD=XX,               X
      (ITK=X'80',ID2=(RCI),INDEX=(AAAAAA,0))
DBDEF FILE=ZZZZZZ,PACKINHI=YES,DBTAG=#UTAG257,                  X
      RCIDID=AA,MPPRCD=PRIME,MPRECD=YY,                         X
      (ITK=X'80',ID2=(RCI),INDEX=(AAAAAA,0))
DBDEF FILE=AAAAAA,PACKINHI=YES,DBTAG=#UTAG257,                  X
      (IID=XXXXXX,IKY=80,PTH=0,IPA=0,ILA=0,IPK=0),              X
      (IID=YYYYYY,IKY=80,PTH=1,IPA=0,ILA=0,IPK=0),              X
      (IID=ZZZZZZ,IKY=80,PTH=2,IPA=0,ILA=0,IPK=0),              X
      (ITK=X'80',ID2=,RID=BBBBBB,ADR=4,RCP=8)
DBDEF FILE=BBBBBB
```

PM34934: Provide header file for SR05SR

- **Previously, z/TPFDF provided DSECT and DBDEF for SR05SR**
  - Accessed by applications as a sorting workfile
- **Corresponding C/C++ header file not available for C/C++ applications**
  - Customers had to create / maintain their own versions
- **C/C++ header file c_sr05sr created**

PM38645: B+Tree FARF6 Support

- **Previously, z/TPFDF had restriction that FARF6 addresses could not be used in B+Tree indexes or data files**
- **With this APAR:**
  - FARF6 addresses can be used in B+Tree data files
  - FARF6 addresses can be used in B+Tree index (node) files
- **Restrictions:**
  - The B+Tree data file and corresponding node file must use the same file address header format
  - Migration of existing B+Tree files is not supported

PM37761: B+Tree/LLR file/find counters in Data Collection

- **Previously, z/TPFDF data collection did not collect find/file information for:**
  - B+Tree index (node) blocks
  - Large logical record (LLR) index and data blocks (LLIBs/LLDBS)
- **Actual I/O's could be significantly higher than indicated in z/TPFDF data collection**
- **Four new counters created**
  - BTFIL: Number of files of a B+Tree index block
  - BTFND: Number of finds of a B+Tree index block
  - LRFIL: Number of files of an LLIB/LLDB
  - LRFND: Number of finds of an LLIB/LLDB

PM37761: B+Tree/LLR file/find counters in Data Collection

```
User:    ZUDFC DISPLAY

System:  23SEP11  09.30.00-09.34.46  - AVERAGE PER SECOND - TOTAL

         DBADD:    64890 DBADR:        2 DBCKP:      523 DBCLS:    32495
         DBCPY:       41 DBCRE:      773 DBDEL:     6439 DBDIX:        0
         DBDSP:       -- DBIDX:       11 DBKEY:    20140 DBMOD:    30525
         DBMRG:        0 DBOPN:    32495 DBRED:   542822 DBREP:     4637
         DBRET:    25105 DBRST:       29 DBSPA:     5854 DBSRT:       39
         DBTLD:       -- DBTLG:       66 DBTRD:       -- DBUKY:      495

         FILNC:     1581 FILEC:        0 CFILE:     1050 PFILE:     4136
         TWRTC:      114 GETFC:     1841 RELFC:      986 DETAC:    73529
         ATTAC:    74686 GETCC:     8366 RELCC:    42468 CFIND:     9903
         PFIND:    33926 BTFIL:      254 BTFND:      368 LRFIL:       12
         LRFND:       37


         CURRENT DATA COLLECTION COUNTS :

         DBADD:18558658 DBADR:      466 DBCKP:   149696 DBCLS: 9293619
         DBCPY:    11605 DBCRE:   220954 DBDEL:  1841490 DBDIX:       57
         DBDSP:       -- DBIDX:     3207 DBKEY:  5760123 DBMOD: 8730058
         DBMRG:        3 DBOPN:  9293619 DBRED:********* DBREP: 1326048
         DBRET:  7179922 DBRST:     8330 DBSPA:  1674165 DBSRT:    11207
         DBTLD:       -- DBTLG:    18924 DBTRD:       -- DBUKY:   141498
         FILNC:   452253 FILEC:       89 CFILE:   300185 PFILE: 1182816
         TWRTC:    32675 GETFC:   526491 RELFC:   281875 DETAC:21029152
         ATTAC:21360076 GETCC:  2392652 RELCC:12145885 CFIND: 2832157
         PFIND: 9702875 BTFIL:    72644 BTFND:   105248 LRFIL:     3432
         LRFND:    10582


         *** END OF DISPLAY ***
```

Future z/TPFDF Enhancements*

- **PM48151: CRUISE enhancements**
  - CRUISE should support I-stream affinity options
  - Reduce CRUISE System Usage
  - When using SSU-ALL option, CRUISE should not ABORT
  - Allow override of standard pool type evaluation
- **PM48823: Allow named LOCTRs and z/TPFDF calls**
- **PM49973: Reduce size of z/TPFDF listings**
- **PM50234: Increase size of ZUDFM log**

* All plans subject to change

z/TPFDF TPFUG Requirements Update

- **Requirements with Changed Status**
- Requirements Summary

Requirements with Changed Status

- **#1: Export TPFDF LRECs to XML Format (DF05182F)**
  - Was: Not Likely
  - Now: Likely

Requirements with Changed Status

- **#2: TPFDF C API improvements (DF00152)**
  - Was: Likely
  - Now: Accepted
  - Requirement contains multiple independent requests
    - Supporting T-types in C
    - In-lining functions such as dfifb
    - Improving setup of key lists
    - Key lists should save search argument values internally
    - C APIs do not support all functionality as HLASM APIs
  - Would like to receive feedback on which sub items are of greatest interest

Requirements with Changed Status

- **#5: RELFC API for TPFDF (DF08189S)**
  - Was: Likely
  - Now: Accepted
- **#7: TPFDF Display Command Directed to the TPF File Systems (DF08187S)**
  - Was: Likely
  - Now: Accepted

Requirements with Changed Status

- **#11: Need a tool for DBDEF check-out and verification (DF10193)**
  - Was: New
  - Now: Not Likely

z/TPFDF TPFUG Requirements Update

- Requirements with Changed Status
- **Requirements Summary**

z/TPFDF TPFUG Requirements Summary

Rank	Req Num	Description	Was	Now
1	DF05182F	Export LRECs to XML	**Not Likely**	**Likely**
2	DF00152	C API improvements	**Likely**	**Accepted**
3	DF00079	In-core records	Not Likely	Not Likely
4	DF08186S	Processor unique struct	Likely	Likely
5	DF08189S	RELFC API for TPFDF	**Likely**	**Accepted**
5	DF08191F	Error checking	Likely	Likely
7	DF08187S	DBDSP to file system	**Likely**	**Accepted**
7	DF08188S	New CRUISE targets	Likely	Likely
9	DF00165	Integrate data collection	Not Likely	Not Likely
9	DF00096	DBDEF/RIAT in 1 place	Not Likely	Not Likely

z/TPFDF TPFUG Requirements Summary

Rank	Req Num	Description	Was	Now
11	DF10193	DBDEF verification	**New**	**Not Likely**
12	DF00170	DBDEF load changes	Not Likely	Not Likely
13	DF08190F	User defined record IDs	Likely	Likely
13	DF00098	Support GDS	Not Likely	Not Likely
15	DF00153	C++ APIs	Likely	Likely
15	DF00021	More efficient sort	Not Likely	Not Likely
17	DF00156	Auto refresh of DBDEF	Not Likely	Not Likely

z/TPFDF Requirements Receiving Less Than 1% of Votes

Rank	Req Num	Description	Was	Now
18	DF09192F	DBUKY file retrieval	Not Likely	Not Likely
18	D94004	Performance measure.	Not Likely	Not Likely
18	DF00176	Transmit metadata	Not Likely	Not Likely
18	DF00010	Block size expansion	Not Likely	Not Likely
18	DF00179	PACK user exits	Not Likely	Not Likely
23	DF00012	Safer initialization	Not Likely	Not Likely
23	DF00011	Recoup user exit updts.	Not Likely	Not Likely
23	DF05180F	String key feature	Not Likely	Not Likely
23	DF05181F	ZUDFM undo	Not Likely	Not Likely
23	DF00140	Add new key -> reindex	Not Likely	Not Likely
23	DF000178	User exits for all APIs	Not Likely	Not Likely

Trademarks

- IBM is a trademarks of International Business Machines Corporation 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.