



# z/TPFDF Status Update

Database Subcommittee

**Chris Filachek**

z/TPF and z/TPFDF Architecture & Development

IBM **z/TPF**

April 12, 2016

©Copyright IBM Corporation 2016.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Disclaimer

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.

15 Minutes

**Delivered Enhancements**

1 Minute

**Future Enhancements**

1 Minute

**z/TPFDF TPFUG Requirements**

# **z/TPFDF Delivered** **Enhancements**

# Continuously Monitor z/TPFDF Activity

- A system administrator can gain insight about z/TPFDF activity in an IBM Tivoli Monitoring environment **so they can properly monitor and tune their system.**
- View data as near real-time data in default or customized views
  - Data is reported by subsystem and file ID
  - Uses same counters as shown by ZUDFC command
  - New counters added (HOLD / DETAC combinations)
- Use the data to set user-defined situations to generate events
- Store data in the Tivoli Data Warehouse
- APAR PI30702 on PUT 12

**Navigator** View: Physical

- Enterprise
  - Windows Systems
    - STRAITG
      - WIN-NSJT7L2CGJQ
        - z/TPF
          - GERRY::PP
            - GSTRAIT::PP
              - CDC Session Info
              - Channel Utilization
              - Communications
              - Common Deploy Files
              - DASD
              - LPAR Utilization
              - Predefined User Data
              - Pools
              - System
              - Tape
              - TPPDF**
              - VFA
            - TPFP1-EC12::PP
            - TPFP3-EC12::PP

**TPPDF Rates Based on TOPn DBRED**

SS Name	File ID	DBRED per Sec	DBADD per Sec	DBADR per Sec	DBCKP per Sec	DBCLS per Sec	DBCPY per Sec	DBCRE per Sec	DBDEL per Sec	DBDIX per Sec	DBDSP per Sec	DBIDX per Sec	DBKEY per Sec	DB pe
BSS	B212	14.7	0.0	0.0	0.0	0.2	0.0	0.0	3.8	0.1	0.0	0.0	18.1	0.0
BSS	B214	9.2	0.0	0.0	0.0	1.1	0.0	0.0	1.7	0.1	0.0	0.0	5.1	0.0
BSS	B211	2.0	0.0	0.0	0.0	0.5	0.0	0.0	0.1	0.0	0.0	0.0	0.5	0.1
BSS	B243	0.2	0.0	0.0	0.0	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.0
BSS	B241	0.1	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0
BSS	Total	26.1	0.0	0.0	0.0	2.1	0.0	0.0	5.5	0.1	0.0	0.0	24.0	0.1

**TPPDF Counts Based on TOPn DBRED**

SS Name	File ID	DBRED	DBADD	DBADR	DBCKP	DBCLS	DBCPY	DBCRE	DBDEL	DBDIX	DBDSP	DBIDX	DBKEY	DBMC
BSS	B212	430	0	0	0	6	0	0	110	2	0	0	530	0
BSS	B214	270	0	0	0	32	0	0	50	2	0	0	150	0
BSS	B211	58	0	0	0	16	0	0	2	0	0	0	16	2
BSS	B243	5	0	0	0	5	0	0	0	0	0	0	5	0
BSS	B241	4	0	0	0	4	0	0	0	0	0	0	4	0
BSS	Total	767	0	0	0	63	0	0	162	4	0	0	705	2

## PI30702 – Continuously Monitor z/TPPDF Activity

# Expanded NODUMP Support



- An application programmer can suppress all z/TPFDF system errors for the currently open subfile **with a single API call**.
- New API allows all z/TPFDF dumps for the currently open subfile to be suppressed
  - `DBOPT` and `dfopt()`
  - All z/TPFDF system errors that return to the application are suppressible
- APAR PI40719 on PUT 12

# Save Subfiles as Part of Initialization



- An operator can backup a z/TPFDF file and initialize it **with a single command**.
- Single command ensures the file being initialized is also saved
  - `ZUDEFM INIT C1C2 FILE-myfile.fil`
- Prime blocks are saved to a file before being initialized
- z/TPF ZCFIL RESTORE command can be used to restore a file
- APAR PI45863 on PUT 12



# CRUISE'ing Faster than Ever



- An operator can restore a z/TPFDF database to new pools with a rebuilt data structure **in less time**.
- Redesigned tape restore process makes more efficient use of multiple ECBs
- Multiple ECBs are used to exchange embedded file address references in parallel
- APAR PI44954 on PUT 12

# A Faster z/TPFDF Recoup



- An operator can recoup a z/TPFDF file with a large number of fixed file ordinals and few or no embedded references **in less time**.
- A new DBDEF parameter (RANGE=) allows fixed file ordinals to be processed with a single ECB
- Ideal for very wide (lots of subfiles) and very shallow (few references or overflows) files
- APAR PI46835 on PUT 12
- Requires co-requisite APAR PJ43492

# z/TPFDF Cache Support

- A database administrator can define read-dominant z/TPFDF files to be retrieved from a local memory cache in a single operation **without any application changes and without any DASD or VFA I/O.**
  - Stores and retrieves highly accessed, infrequently updated subfiles from local processor memory
  - Avoids limitations of VFA Sync locks in an MPLF environment
  - Maintains data integrity in a loosely coupled environment
- Memory management handled automatically
  - Only set total cache size and maximum number of blocks to cache for each subfile



# Set up the z/TPFDF cache

1. Define the z/TPFDF control format-2 global

```
ZGLBL GLOBAL DEF IDFCNTRL PROT=YES IS-NO  
SYNC=YES LOC-31BIT SSU-N
```

2. Initialize the z/TPFDF control format-2 global

```
ZGLBL GLOBAL INIT IDFCNTRL SO-ZERO SI-hexsize
```



# Activate cache support

1. Determine amount of data to cache and set the total size of the cache

```
ZUDFM CACHE SYSTEM SIZE-size
```

2. Activate z/TPFDF cache support

```
ZUDFM CACHE SYSTEM ACTIVATE
```

```
ZUDFM CACHE SYSTEM ACTIVATE CONTINUE
```



# Enable caching for a file

1. Define the maximum number of blocks to cache per subfile:  
`ZUDFM CACHE FILE LIMIT-blockLimit ID-fileid`
2. Enable caching for a z/TPFDF file  
`ZUDFM CACHE FILE ENABLE ID-fileID`

*No application changes required!*

# z/TPFDF Cache Requirements

- In a loosely coupled environment, you must attach one or more coupling facilities to the z/TPF complex
- Any z/TPFDF R-type file may use cache support, except those with the following characteristics
  - B+Tree files
  - Block index files
  - Uses algorithm #TPFBD0D
  - See Knowledge Center for other minor restrictions...
- APAR PI44953 on PUT 13
  - Requires co-requisite APAR PJ43474

# z/TPFDF **Future** Enhancements





# z/TPFDF Encryption Support

A database administrator can encrypt data-at-rest in z/TPFDF files and protect sensitive customer information without requiring any application changes.

# **z/TPFDF TPFUG Requirements Update**

# Requirements with Changed Status

Requirement	Description	Was	Now
DF14202	CRUISE Restore Improvement	Accepted	Partly Available (PI44954)
DF00165	Integrate TPFDF Data Collection and reporting	Accepted	Partly Available (PI30702 / PJ42742)
DF05182F	Export/Import TPFDF LRECs to/from XML Format	Likely	Available (PI33010)

# Summary of Delivered Enhancements



- PI30702 – Continuously Monitor z/TPFDF Activity
- PI40719 – Expanded NODUMP support
- PI45863 – Save Subfiles as Part of Initialization
- PI44954 – CRUISE'ing Faster than Ever
- PI46835 – A Faster z/TPFDF Recoup
- PI44953 – z/TPFDF cache support

**Thank you!**

Questions or comments?

# Trademarks

- IBM, the IBM logo, ibm.com and Rational are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “[Copyright and trademark information](#)” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

## 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.