

z/TPFDF Update

Database / TPFDF Subcommittee

Claire Durant

2024 TPF Users Group Conference
May 5-8, New Orleans, LA

IBM Z



Agenda

- Efficiency

- Run CRUISE faster without impacting transactional work
- Reduce overhead of CRUISE keypointing
- Decrease CPU resources used by CRUISE

- Usability

- Improved z/TPFDF diagnostics
- Confidently run CRUISE as low priority
- Control CRUISE backward chain validation

- Scalability

- ☐ Improved statistics reporting

Efficiency

Run CRUISE faster without
impacting transactional work


Problem Statement

CRUISE runs on the same I-streams as transactional work.

As databases grow in size, CRUISE processing increasingly competes with transactional work for CPU processing time.



You can run CRUISE as a low-priority utility, but if CPU utilization is approaching 100%, CRUISE will not complete in a timely fashion.



You can wait to run CRUISE until there are low transaction volumes, but this delays database verification.

As-Is User Story

During the daily peak when little CPU capacity is available, the DBA asks to run the CRUISE utility so they can verify database integrity.


Carol, a z/TPF coverage programmer, notices that z/TPF is near its high-capacity threshold.

She wishes that she could start CRUISE right away and that CRUISE would **complete in a timely manner without impacting transactional work.**



Carol
Coverage Programmer

Value Statement




By running CRUISE on as many as 4 fenced I-streams, you can run CRUISE **without affecting transactional work and CRUISE processing time might be reduced.**

- Fenced I-streams are active I-streams that are not in use and have an I-stream number greater than the I-stream cap.
- Similar to the behavior of z/TPF recoup

APARs [PH54526](#) & [PJ47115](#) (Dec. 2023)

To-Be User Story



Carol
Coverage Programmer

During the daily peak when little CPU capacity is available, the DBA asks to run CRUISE so they can verify database integrity.

Carol, a z/TPF coverage programmer, knows CRUISE can run on fenced I-streams.

By running on fenced I-streams instead of in-use I-streams, the CRUISE verify operation **completes in a timely fashion without impacting transactional work.**

Technical Details

- To run CRUISE on fenced I-streams:
 1. Set the **ISTREAM** parameter in the CRUISE parameter table to **BALANCED**
 2. Run CRUISE as a low-priority utility
- CRUISE ECBs will run on as many as 4 fenced I-streams
 - CRUISE will run on the **4 highest fenced I-streams** defined
 - In-use I-streams might be used if the fenced I-streams are fully utilized

Efficiency

Reduce overhead of CRUISE
keypointing

Problem Statement

- The CRUISE keypoint is used to track the progress of CRUISE and restart a CRUISE operation.
- For some databases, frequent filing of the CRUISE keypoint record causes noticeable overhead.
 - For flat databases, as much as 2% of all CRUISE I/Os are to file the CRUISE keypoint (record ID X'FD18').
- Most z/TPF operators rarely restart CRUISE operations, which means this overhead isn't needed in most scenarios.

Value Statement

- Frequency of filing the CRUISE keypoint is reduced by 100x
 - For flat databases, only a fraction of a percent of all I/Os are for the CRUISE keypoint record.
- When CRUISE processes a large flat database, processing can **run as much as 40% faster** and have a negligible impact on the time it takes to restart CRUISE.
 - Results will vary based on your database structure and other factors.

Also in APARs [PH54526](#) & [PJ47115](#) (Dec. 2023)

Efficiency


Decrease CPU resources used
by CRUISE

As-Is User Story

When running CRUISE on a large flat database, Calvin, a z/TPF capacity planner, notices that each CRUISE ECB chain chases only 1 fixed record ordinal and its overflow records, and then exits.

Because each CRUISE ECB does a minimal amount of work, Calvin finds that a noticeable amount of the CPU resources used by CRUISE are from ECB create and exit processing.

Calvin wishes that he could **reduce the amount of CPU resources that are used by CRUISE.**




Calvin
Capacity Planner

Value Statement

- For a flat (single-level) database, each CRUISE ECB chain chases multiple fixed record ordinals, so CRUISE can process the same database with fewer ECBs.
 - Use the [DBDEF RANGE](#) parameter to specify the number of ordinals chain chased by each ECB.
 - The **RANGE** parameter is already used by recoup for the same purpose.
- With a simple DBDEF change, the CPU utilization of CRUISE can be reduced by **up to 20%** for large, flat databases.
 - Results will vary based on the value of the **RANGE** parameter, the database structure, and other factors.

APAR [PJ47269](#) (Apr. 2024)

To-Be User Story



Calvin
Capacity Planner

Calvin, a z/TPF capacity planner, already has the **RANGE** parameter set on the **DBDEF** macros for large, flat databases.

Calvin is happy that CRUISE can chain chase these databases with fewer ECBs, reducing the amount of total CPU resource consumed by CRUISE.

Usability

Improved z/TPFDF diagnostics

Problem Statement



System errors detected by the CP while processing a z/TPFDF API do **not** include the trace name of the calling program in the console message.

- System errors detected by z/TPFDF include the trace name of the calling program in the console message.

The dump must be postprocessed and reviewed to determine the name of the program that called a z/TPFDF API.

As-Is User Story

Derrick, a z/TPF operator, observes a 000004 (OPR-4) system error on the system.

He needs to route the error to the application owner for investigation.

However, the console message does not indicate which application called the z/TPFDF API.

Derrick needs to **spend time reviewing the dump** so he can route the error to the application owner.



Derrick
Operator

As-Is Background

- Passing an invalid LREC address on a dfadd() or DBADD


```
CPSE0162E 15.18.42 IS-0001 SS-BSS SSU-BSS SE-000795 OPR-I000004
010000A TRC-UAA0 OWNER-drvrDFTDfrmCVZZ-ECBAdr:148BD000.
UTDF OBJ-uaa0 00001EA0 LOADSET-BASE
PSW 07150000 80000000 00000000 07E1CEA0 PIC 0011 ILC 0006 I-D50032B7E002
R0-2 00000000 14800420 00000000 0000002B 00000000 00000000
R3-5 00000000 19E02F00 00000000 148E2EA2 00000000 148E8320
R6-8 00000000 11223344 00000000 87E1B4D8 00000000 07E1B420
R9-11 00000000 14800000 00000000 07E1D1E4 00000000 00001000
R12-14 00000000 00002000 00000000 19C7F618 00000000 11223344
R15 00000000 00000000
```

Trace names (TRC-*) are from internal z/TPFDF routines

- Application close to timeout limit before calling z/TPFDF

```
CPSE0162E 15.42.54 IS-0001 SS-BSS SSU-BSS SE-000817 OPR-I000010
010000A TRC-UAM0 OWNER-drvrDFTDfrmCVZZ-ECBAdr:14818000.
CP OBJ-ccnucl 0002C2D8
```


Value Statement



By making the trace name of the calling program easily accessible, **valuable time and resources can be saved** during the initial investigation of system errors detected by the CP during z/TPFDF processing.

[APAR PJ47153](#) (Nov. 2023)

To-Be User Story



Derrick
Operator

Derrick, a z/TPF operator, observes an OPR-4 system error on the z/TPF system.

He is anxious to route the error to the application owner for investigation.

The trace name in the console message indicates which program called the z/TPFDF API.

Derrick quickly routes the error to the owner **without reviewing the dump!**

Technical Details

- Passing an invalid LREC address on a dfadd() or DBADD

```
CPSE0162E 15.19.04 IS-0001 SS-BSS  SSU-BSS  SE-002777 OPR-I000004
010000A TRC-QXAC          OWNER-drvrDFTDfrmCVZZ-ECBAdr:148C3000.
UTDF      OBJ-uaa0          00001EA0  LOADSET-BASE
PSW       07150000 80000000  00000000 07E1CEA0  PIC 0011 ILC 0006 I-D50032B7E002
R0-2     00000000 14800420  00000000 0000002B  00000000 00000000
R3-5     00000000 19DED300  00000000 14886EA2  00000000 1488A320
R6-8     00000000 11223344  00000000 87E1B4D8  00000000 07E1B420
R9-11    00000000 14800000  00000000 07E1D1E4  00000000 00001000
R12-14   00000000 00002000  00000000 19C1F618  00000000 11223344
R15      00000000 00000000
```

Trace names (TRC-*) are from programs that called z/TPFDF

- Application close to timeout limit before calling z/TPFDF

```
CPSE0162E 15.19.30 IS-0001 SS-BSS  SSU-BSS  SE-002778 OPR-I000010
010000A TRC-QYSF          OWNER-drvrDFTDfrmCVZZ-ECBAdr:14842000.
CP       OBJ-ccnucl        0002C2D8
```

Usability

Confidently run CRUISE as
low priority

As-Is User Story

Derrick, a z/TPF operator, wants to always run CRUISE as low priority to guarantee that CRUISE processing does not compete with transactional workload.

Derrick is concerned that he might accidentally forget to use the `-LP/` prefix when starting CRUISE.

If he forgets the `-LP/` prefix, CRUISE **will not run as low priority, will not run on fenced I-streams, and would compete with transactional work.**




Derrick
Operator

Value Statement



An operator can set a CRUISE option so that CRUISE runs as a low-priority utility. With this option, operators can **confidently set up CRUISE to run as low priority without using command prefixes.**

To-Be User Story



Derrick
Operator

Derrick, a z/TPF operator, wants to always run CRUISE as low priority.

Derrick uses the **ZFCRU SETUP** command to set the **LP** parameter in the CRUISE parameter table to **YES**. This causes CRUISE to run as low priority without having to use or remember the **-LP/** prefix.

Derrick is happy because he can run CRUISE as low priority **without having to remember the prefix or which CRUISE command to use it on.**

Technical Details

- Use the [ZFCRU SETUP](#) command to set low priority for existing parameter tables

ZFCRU SETUP LP-YES

- Use the [ZFCRU DEFAULT](#) to set low priority for future parameter tables (default settings)

ZFCRU DEFAULT LP-YES

APARs [PH54526](#) & [PJ47115](#) (Dec. 2023)

Usability

Control backward chain validation

Problem Statement

- APAR PH44215 (2022) introduced validation of backward chains on CRUISE operations.
- Certain databases can trigger backward chain error messages while running CRUISE.
 - Legacy databases created before backward chaining existed
 - Databases that were restored but not packed
- You can limit the number of messages generated by CRUISE, but that might hide other errors.

Value Statement

- By disabling backward chain validation, CRUISE does not issue any backward chain validation messages and you can **focus on other CRUISE messages without having to pack the database.**
- Use the [ZFCRU SETUP](#) command to disable backward chain validation.

ZFCRU SETUP BCHVAL-NO

Also in APARs [PH54526](#) & [PJ47115](#) (Dec. 2023)

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.



Scalability

Improved statistics reporting

Problem Statement

- As the number of subfiles, overflow records, LRECs, and LLRs in a z/TPFDF database grow above 4 billion, **it becomes difficult to track the database size and characteristics.**
 - z/TPFDF recoup and CRUISE statistics count the number of records, LRECs, LLRs, and more for each z/TPFDF file.
 - Almost all statistics use a 4-byte counter and are limited to 4 billion.
 - After a counter reaches 4 billion, z/TPFDF recoup and CRUISE stop counting that statistic and display “GT 4B”.
 - An exception is the z/TPFDF recoup statistic for the number of LRECs in a file, which uses an 8-byte counter and displays counts up to 999 billion.

Value Statement

- z/TPFDF recoup and CRUISE will be enhanced to provide counts above 4 billion, so that you can **accurately monitor your database size and characteristics**.
 - You can define over 1 million cylinders per DASD volume and many FARF6 pool sections already have 10's of billions of allocated pools.
 - As your databases grow into this space, you can continue to “see”, track, and manage your database.
- Plan to update the applicable z/TPFDF recoup and CRUISE statistics to use 8-byte counters and display values as large as 999 billion.

Be a sponsor user

Sponsor users assist in design and implementation, and your feedback drives our development cycle.

Target personas

- Database administrator
- Capacity planner

Begins

Now!

Interested? Contact

Chris Filachek (filachek@us.ibm.com)



Thank you

© Copyright IBM Corporation 2024. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM 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 at [Copyright and trademark information](#).

