Run Recoup on fenced I-streams

Michael Shershin

IBM Z

IBM

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.

Problem Statement

Typically, Recoup has a time window when it can run. If Recoup exceeds the time window, other work is impacted. As databases grow larger, it becomes difficult to keep Recoup run time within the allowed time window.

Pain Points

A balancing act is performed when Recoup is run.

- Transactional work should have higher priority. If there is not enough capacity, Recoup slows down and takes longer to complete.
- Other business required utilities are usually not run while Recoup is running. Historically, there is concern that running multiple utilities together will cause system problems.
- If Recoup does not complete in its allotted time, the ability to run other utilities is impacted. This might cause business impacts.
- As the size of databases grow, will Recoup be able to complete in its allotted time?

Technical Details – Terminology review

Active I-stream – an I-stream that is defined to the LPAR and is available for use by z/TPF.

In-use I-stream – an active I-stream that application work can be dispatched to.

I-stream cap – the highest I-stream number of any I-stream that application work can be dispatched to.

Fenced I-stream – an active I-stream that is not in use and that has an I-stream number greater than the I-stream cap

As-Is: Example configuration

z/TPF LPAR

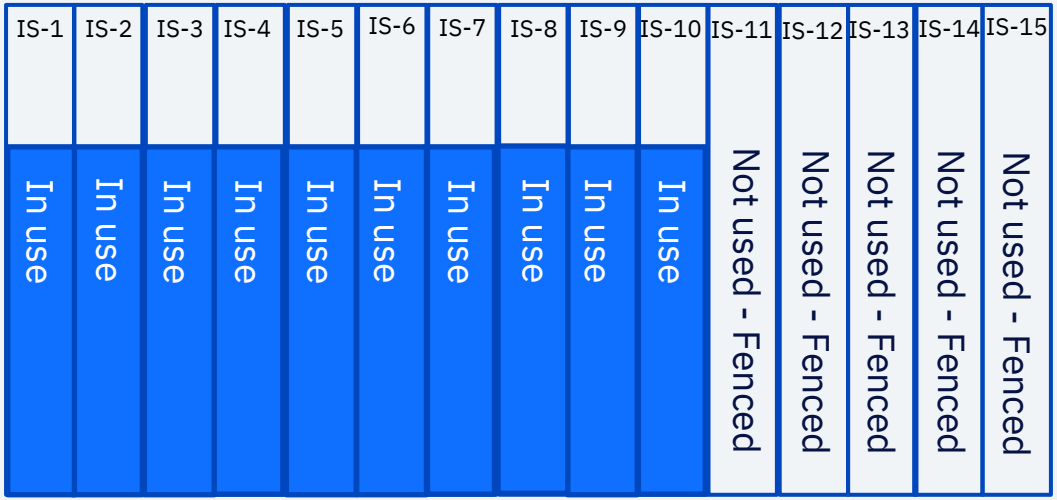

I-stream cap = 10
In-use I-streams (ISTUSEIS) = 10
Active I-streams (ISTACTIS) = 15
Fenced I-streams = IS-11 through IS-15

- **Transactions** run on in use I-streams: IS-1 through IS-10
- **Recoup ECBs** run on in use I-streams

As-Is: Recoup

Recoup work distribution

- Recoup shares processing power of in-use I-streams with transactional work.
- Recoup primary ECB runs on IS-1.
- Recoup child ECBs are load balanced across all in-use I-streams.

To-Be: Recoup

Provide option to allow Recoup to run on fenced I-streams

- Dynamic CPU must be in use.
  - An I-stream CAP must be set and fenced I-stream must exist
- Recoup must be run as a low priority utility.
- Up to 4 fenced I-streams will be used.

To-Be: Recoup

Recoup work distribution when fenced I-streams are used

- Recoup primary ECB will continue to run on IS-1.
- Recoup child ECBs will be automatically routed to fenced I-streams.
- If the 4 fenced I-streams are running at 100% utilization, additional Recoup child ECBs will be load balanced across in-use I-streams.

To-Be: Example configuration with Recoup using fenced I-streams

z/TPF LPAR

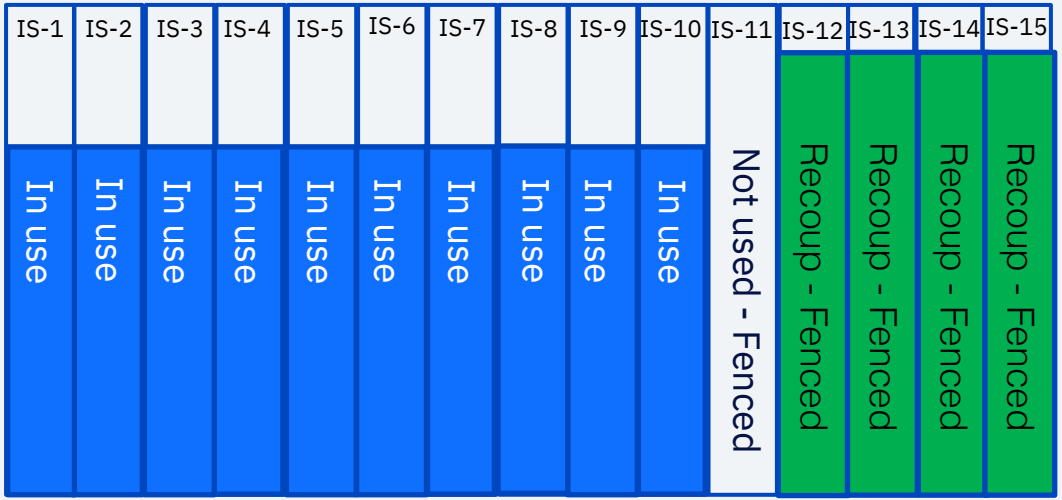

I-stream cap = 10
In-use I-streams (ISTUSEIS) = 10
Active I-streams (ISTACTIS) = 15
Fenced I-streams = IS-11 through IS-15
Recoup use = IS-12 through IS-15

- **Transactions** run on in use I-streams: IS-1 through IS-10
- **Recoup primary ECB** runs on IS-1
- **Recoup child ECBs** run on fenced I-streams

To-Be: Recoup

I-stream utilization

- Drive the 4 fenced I-streams as hard as possible.
- Utilization on in-use I-streams will likely be different than utilization on the 4 fenced I-streams.

Technical Details

Why must Recoup run as low priority?

- If the I-stream cap changes and a fenced I-stream that is currently doing Recoup work is now an in-use I-stream, transactional work must be higher priority.
  - Purpose of changing the I-stream cap is to get more capacity for transactional work.
  - By having Recoup as low priority, Recoup will allow transactional work to use the added capacity immediately.
- Even as low priority, all Recoup work has the same priority when running on fenced I-streams.

To-Be: Example configuration with Recoup using fenced I-streams after I-stream cap change

z/TPF LPAR

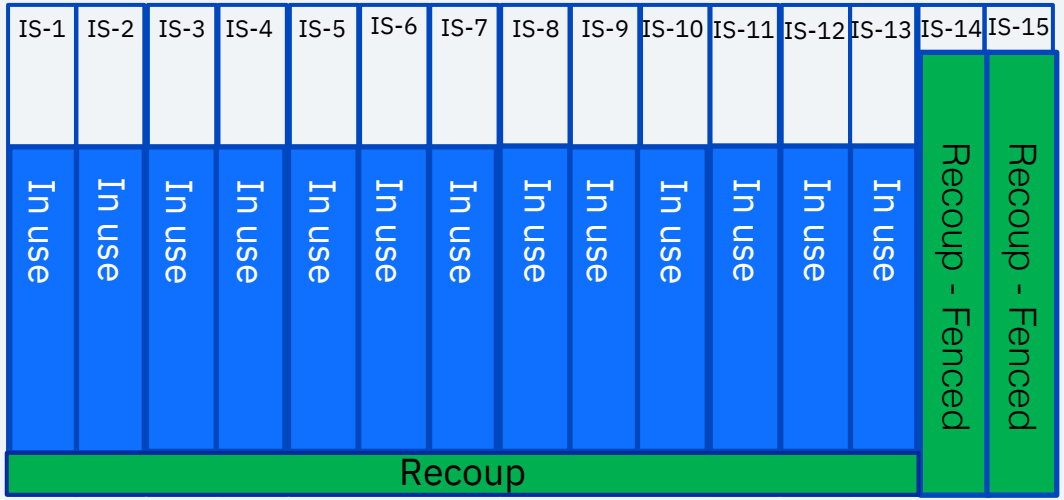

I-stream cap = 13 (was 10)
In-use I-streams (ISTUSEIS) = 13
Active I-streams (ISTACTIS) = 15
Fenced I-streams = IS-14 through IS-15
Recoup use = IS-14 through IS-15

- **Transactions** run on in use I-streams: IS-1 through IS-13
- **Recoup primary ECB** runs on IS-1
- **Recoup child ECBs** run on fenced I-streams and overflow to in-use I-streams

To-Be: Recoup

Overflow from fenced I-streams to in-use I-streams

- If the Recoup child ECBs overflow from fenced I-streams to in-use I-streams, transactional ECBs will be higher priority and they will not be impacted by the Recoup child ECBs.
- The intent of overflowing to in-use I-streams is to exploit unused capacity on in-use I-streams.
  - If in-use I-streams are not running at 100%, spare cycles are available.
  - These cycles are highly pershible.
  - Best to use spare cycles as long as there is no impact to transactions.
  - No additional software licensing costs to run Recoup on in-use I-streams if SCRT is used.

Concern

- DASD I/Os are naturally limited when Recoup is running on in use I-streams. Recoup shares in use I-streams with transactional workload. Available CPU time to allow Recoup to run is limited and this limits Recoup DASD I/O requests.
- DASD I/Os can increase substantially when Recoup is running on fenced I-streams. Recoup will get all the CPU time on fenced I-streams.
- Will the DASD subsystem be able to handle the added load when Recoup is running on fenced I-streams?

To-Be: New IOPS control

DASD I/Os per second (IOPS) control

- User would set the maximum allowed IOPS via Z command
  - IOPS are physical I/O (not satisfied in VFA)
  - DASD IOPS are a summation of physical I/O across all DASD modules on a subsystem
  - Maximum allowed IOPS value should be less than the DASD subsystem capacity
- New API will be provided
  - Return indicator that says whether I/Os are higher than the maximum allowed IOPS
  - Return a time period to wait if I/Os are too high

To-Be: New IOPS control

DASD I/Os per second (IOPS) control users

- Recoup will be updated to use the new API
- New API can be used by any utility
  - Utilities that do significant amounts of physical I/Os should use the API
- Transactional workload should NOT use the API
  - Transactions can drive IOPS above the maximum allowed IOPS that is used by the API

Value Statement

Additional processing power can be used for Recoup so that its intersection with transactional work and other business required utilities is reduced.

Allows Recoup to complete in much less time than today.

Thank you