



TPF Users Group Fall 2007

Unplanned Module Down Concepts

Improve unplanned module down processing in a loosely coupled environment on z/TPF

Mark Lehrer / Chris Filachek

Hot Topic

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.

Disclaimer / Comments

- These are design concepts on how unplanned module down may be improved on z/TPF and are subject to change.
- These concepts apply only to loosely coupled customers. Base customers are not affected by these proposed changes.
- Please provide comments and feedback at the TPF Users Group or through your customer service representative.

Main Points of Improvement

1. Improved unplanned module down processing flow with a new Synchronization Point and new “Going-Offline” status.
2. Monitoring will be performed for module down / movelocks synchronization points to notify of processes stalled at a Sync Point.
3. Restart will check for modules going offline. A new command will be provided to leave modules online or take them offline.
4. IPL devices that were going offline at the time of a software IPL will not be used as the IPL device if a usable duplicate exists.
5. New flags and processing will replace the early file of Keypoint 6.

Module Down Processing Flow

Control Program Processing (CP):

- A. An unplanned module down condition is realized in the CP.
- B. NEW: Mark the module as “Going-Offline”.
- C. NEW: The module queue is left intact and is not purged. The module queue can be considered stalled at this point. It is probably not moving, and IOBs are being added by new work.
- D. NEW: The module is marked in lock suspense. Any duplicate module will be left alone until later in module down processing.
- E. Create a module down ECB for this module.

Module Down Processing Flow (cont)

Module Down ECB gets control:

- A. Call the module down user exit (UYEN)
- B. NEW: Record “Going-Offline” status and obtain unplanned module down ownership.
 - Records are found with hold, updated and filed to indicate the module is going offline.
 - The updating processor indicates itself as the owner.
 - The owning processor is responsible for notifying all processors of the condition and starting unplanned module down in the complex for this module.

Module Down Processing Flow (cont)

- If another processor already owns the module down for this module, then EXIT.
- C. NEW: Sync Point for all processors to agree module is going offline.
- The module down owning processor informs all processors of the module's "Going-Offline" status.
 - NOTE: As previously stated, the module queue is not purged at this point and is probably growing.
 - Start Sync Point Monitor
 - The module down owning processor waits for responses.
 - All processors must "agree" and respond to the unplanned module down owner before module down processing will continue.

Module Down Processing Flow (cont)

- D. NEW: Purge the module queue
- E. Notify all other processors to enter module down processing
- F. Perform movelocks processing
- G. NEW: Start Sync Point Monitor
- H. Movelocks Synchronization Point
- I. Dismount the module
- J. Update Keypoint 6 and Keypoint M

In summary, the unplanned module down process will:

- Keep the module queue intact for a longer period
- Add another Synchronization Point to the process
- Utilize DASD IO to record unplanned module down conditions.

Synchronization Point Monitor

TPF will check synchronization (sync) point status and issue warning messages.

- A Sync Point Monitor will be active if any module is going through the Module Down or movelocks Synchronization Point.
- The Monitor will check once every 10 seconds.
- The Sync Point Monitor will have an instance active in each MDBF subsystem in which modules are going offline.
- The Sync Point Monitor will scan through all modules going offline, looking for sync points that are incomplete. A list of unresponsive processors for all sync points will be displayed.

Synchronization Point Monitor (cont)

- Operations will now have the opportunity to react and remove unresponsive processors with the ZPSMS FORCE DEACTIVATE command.
- The timeout periods for module down and movelocks synchronization points will be lengthened. In order for ample time for the warning messages to be processed and action to be taken, the new timeouts will be 30 minutes.

Module Down Restart updates

Phase One:

- Performed for all processors in the complex:
 - If the IPL device for any TPF subsystem was going offline at IPL, then Restart is stopped.
- Performed for processors that are NOT the first to IPL in the complex:
 - If modules were going offline at IPL, then the Restart is stopped.

A processor stopped in Restart must be deactivated, with the ZPSMS FORCE DEACTIVATE command, to allow any possibly stalled Sync Points to complete.

Module Down Restart updates (cont)

Phase Two: Only processor in the complex

If this is the first processor to IPL (only processor in the complex) and modules were going offline at IPL:

- Display the list of modules that are “Going-Offline”
- Use a new command to decide if modules “Going-Offline” in this TPF subsystem should stay online or should be taken offline.
- If Operations decided to take modules offline, the Restart schedule will not continue and process the next subsystem until all module downs in this subsystem complete.

Module Down Restart updates (cont)

Module Down Restart will appear in these locations in the Restart Schedule:

CTIN calls Restart

- ...
- Keypoint Copy (CVRN)
- NEW: Module Down Restart Phase One
- Coupling Facility Phase One (CFLR)
- Commit/Rollback (CLM0)
- ...
- Locking Restart (CNPY)
- NEW: Module Down Restart Phase Two
- Validate Keypoint Pointer Record (CVZ6)
- CTK6 Synchronization Check (CVX3)
- ...

Software IPL updates

TPF should not attempt to use an IPL module that is going offline.

- At software IPL time (e.g. – ZRIPL or catastrophic), TPF will determine the status of the IPL device for each subsystem.
- If an IPL device is marked as “Going-Offline”, then TPF will attempt to use the duplicate module if it is usable. This information will be noted for IPL’s reference.
- Subsystems other than the BSS will read their IPL device from the file copy of Keypoint M (CTKM). If TPF noted a change in IPL device, then the new module will be used in place of what is in CTKM.



ZPSMS DEACTIVATE processing updates

Another processor will take over any module downs that were owned by a deactivated processor

- The module down owner is responsible for issuing the new IPC and synch point as well as issuing the broadcast IPC for all processors to enter unplanned module down.
- If a processor that owns a module down fails, then the process could stall.
- Upon the entry of a ZPSMS FORCE DEACTIVATE, either the ZPSMS issuing processor or a ZPSMS IPC receiving processor will takeover the module downs owned by the deactivated processor.

Questions????



Trademarks

- IBM is a trademark 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.