
Rule DAS233: PEND DELAY TIME WAS CAUSED BY CONTROLLER BUSY

Finding: A significant amount of the PEND time delay was caused by controller activity.

Impact: This finding may have a MEDIUM IMPACT or HIGH IMPACT on the performance of the device.

Logic flow: The following rules cause this rule to be invoked:
 DAS200: Volume with the worst overall performance
 DAS230: Major cause of I/O delay was PEND time

Discussion: PEND time is the time from the issuance of the SSCH instruction until the device is selected by the control unit. This time is caused by queuing for the path (wait for channel, wait for director port, wait for control unit, wait for device, or wait for other reasons).

Large PEND times for devices attached to cached controllers may imply a high percent of read miss operations, or non-volatile storage (NVS) writes.

To improve the probability of a read hit, the controller can *prestage* data into its cache. Prestaging means that data is read into the controller's cache ahead of its actually being required for use by an application. The amount of data that is prestaged depends on (1) whether the data is being accessed in a direct (random) mode or in a sequential mode and (2) the controller model and the enhancements made to the controller.

C For *direct mode*, after the record is located, the 3390-3 and 3990-6 (initial version) stages in the balance of the track being read.

The 3990 Model 6 (with record cache) stages only the records requested into cache, eliminating the balance of the track staging that is normal with track caching as was implemented on initial versions of 3990-6 and on the 3990-3. This improvement reduces the PEND time caused by the controller busy during track staging.

C As examples of prestaging for *sequential mode*, the 3990-3 reads up to two tracks into the cache¹ before they are required, while the ESS 2105 sequential staging reads up to two cylinders ahead.

¹With the Sequential Staging Performance Enhancement, the 3990-3 can prestage up to a full cylinder (15 tracks) into the cache.

During prestaging operations for sequential reads, the control unit regularly checks to see whether other I/O requests are waiting to be processed. If any are waiting, the control unit interrupts the prestage operation, processes the queued requests, and continues with the prestage.

In DASD Fast Write Mode, the data is stored simultaneously in cache storage and in nonvolatile storage (NVS). At some subsequent time, the data in NVS can be *destaged* to DASD.

In Cache Fast Write Mode, data is placed into cache immediately, and there is no interaction with the device nor with NVS. However, if cache memory is required (or if Cache Fast Write Mode is turned off), the data in cache is destaged to DASD.

Significant PEND time can result from destaging to DASD.