
Rule DAS130: MAJOR CAUSE OF I/O DELAY WAS PEND TIME

Finding: CPExpert has determined that excessive PEND time was a major cause of delay in DASD response for the device.

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:
DAS100: Volume with the worst overall performance

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).

CPExpert computes the average per-second PEND delay time, for each of the causes listed above. Rule DAS130 is produced if the average PEND time accounted for a significant percent of the device response time.

Suggestion: Large PEND times usually involve the following situations:

- **Shared devices.** If the device is shared with another system, PEND time may indicate contention with the other system. Large PEND times in shared-device environments usually involve situations very similar to those described under IOSQ time:

- Multiple data sets may be active on the volume. This situation is the most common and easiest to solve. The data sets can be redistributed among different volumes, to eliminate the queuing at the channel level (reflected as PEND time) for the single volume. If some of the data sets are not required to be shared, then the Data Base Administrator has complete flexibility to move these data sets (subject, of course, to the performance implications of the target devices). These data sets should be moved to a non-shared device.

If the data sets are required to be shared, then they must be relocated to shared devices.

- Multiple applications or users may be using the same data set on the volume. Depending upon the data set characteristics, duplicate

copies of the data set may be placed on different volumes.

This would solve the PEND problems caused by contending systems. If this option is feasible, the data sets could be placed on non-shared devices, likely resulting in even more performance improvement.

- Multiple application systems may be using the volume experiencing high PEND times. In this case, perhaps application redesign or scheduling can solve the problem.

Additionally, large PEND times for shared devices could be caused by RESERVE from the other system. The applications issuing the RESERVE should be examined to determine whether the RESERVE is required. If the RESERVE is required, the above situations should be reviewed to determine whether improvements can be achieved.

- **Non-shared devices.** Large PEND times for devices that are not shared may mean that there are insufficient paths available to the device. Too much I/O may be directed to many devices on the path, control unit, or head-of-string. The data sets can be redistributed among different volumes on different paths, control units, or heads-of-string. This will reduce the hardware-level queuing. Alternatively, the entire volume may be moved to a different (less busy) head-of-string or path.

If redistributing the data sets or moving the volume is not feasible, then the device should have more paths. Depending upon the existing configuration, this may involve re-configuring existing channel paths, or acquiring additional hardware.

- **Devices attached to cached controllers.** Large PEND times for devices attached to cached controllers may imply a high percent of read miss operations, or non-volatile storage (NVS) writes for IBM-3990-3 devices. Fairchild ¹ lists four ways in which staging in caching controllers can cause hidden device busy (with the device busy potentially reflected in high PEND time):
 - The normal (random) caching algorithm stages all records to the end of the track after a requested record is read.
 - The normal (random) caching algorithm stages all records from the beginning of the track to the requested record if a front-end miss occurs.

¹ Fairchild, Bill, "The Anatomy of an I/O Request", *Conference Proceedings*, CMG'90, the Computer Measurement Group, Chicago, IL.

-
- The sequential caching algorithm stages all records to the end of the track after the requested record is read, and stages in all of the next track. IBM-3990 (Model 3) controllers stages in all of the next three tracks.
 - Most writes to extended function IBM-3990 (Model 3) go into NVS with a subsequent destaging required.
 - **Dual Copy Initialize.** Large PEND times for IBM-3390 devices may be caused by dual copy initialize. In this case, the dual copy initialize should be turned off.