
Rule DAS134: PENDING DELAY TIME WAS CAUSED BY DEVICE BUSY

Finding: A significant amount of the PENDING time delay was caused by device busy.

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
DAS130: Major cause of I/O delay was PENDING time

Discussion: PENDING 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). Some of the causes of PENDING for device busy are listed below:

C PENDING time for device busy can be caused by other systems in the sysplex that issue a RESERVE for the device. While the RESERVE is held, I/O operations to the device will be held in a PENDING for device busy state.

Additional information will be available if CPExpert is performing shared DASD analysis (shared DASD analysis will be done if %LET SHARED=Y; was specified in USOURCE(DASGUIDE) and if the SMF data indicates that the device was shared). Rule DAS300 will be produced to show the effects of activity from other systems on the device being analyzed. Included in the output from Rule DAS300 will be the amount of RESERVE time from each system, and the total from all systems sharing the device.

C Multiple Allegiance allows multiple active concurrent I/O operations on a particular device when the I/O requests originate from different systems. With Multiple Allegiance, there is complete access with read I/O operations. For write I/O operations, there is concurrent access unless there is a conflicting extent¹. If there is a conflicting extent, the controller holds the I/O operation in a PENDING state for the device.

C After an I/O operation, the device will read the remainder of the track into its device-level buffer. This is done to prevent delay for rotational positioning. If a new I/O operation is attempted while data is being read

¹A conflicting extent is one in which the write operation attempts to update an extent.

into the device cache buffer, the I/O operation will be in a PEND state for device busy.

Suggestion: If Rule DAS134 is produced frequently, you should consider the following alternatives.

C If shared device analysis is specified as **%LET SHARED = Y**; in USOURCE(DASGUIDE), CPEXpert will analyze potential problems caused by sharing DASD. Rule DAS300 will be produced for all systems that share the device with high PEND time, if CPEXpert concludes that other systems could cause performance problems. The RESERVE time will be included in the output from Rule DAS300.

If this RESERVE time is high for the device, you should consider whether high activity data sets can be moved among different volumes on different paths.

C Alternatively, determine whether the data sets can be moved to a controller that supports Multiple Allegiance². Multiple Allegiance allows multiple active concurrent I/O operations on a particular device when the I/O requests originate from different systems. With Multiple Allegiance, there is complete access with read I/O operations. For write I/O operations, there is concurrent access unless there is a conflicting extent.

C Alternatively, consider whether workload scheduling can eliminate the conflicts between the data access requirements between systems.

²Multiple Allegiance is available with IBM's Enterprise Storage Server (ESS) subsystems.