
Rule DAS220: MISSED RPS RECONNECT WAS MAJOR CAUSE OF I/O DELAY TO CRITICAL WORKLOAD

Finding: CPEXpert has determined that missed Rotational Position Sensing (RPS) reconnects 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. *This finding applies only to legacy systems (e.g., 3380 devices attached to 3990-2 controllers).*

Logic flow: The following rules cause this rule to be invoked:
DAS200: Volume with the worst overall performance

Discussion: CPEXpert computes the average channel path busy for paths to the device. SMF Type 78 information is used to acquire the channel path utilization, and the IOCP macro information is used to determine the channel paths to the device.

If the device is capable of DLS mode, two channel paths can be concurrently busy to the string. If the device is capable of DLSE mode, four channel paths can be concurrently busy to the string.

CPEXpert use an M/M/C queuing model (Erlang's C formula) to compute the probability that all paths to the device were busy when the device attempted to reconnect. The M/M/C queuing model is adjusted depending upon the number of concurrent paths. (Please refer to *Probability, Statistics, and Queuing Theory* by Arnold O. Allen for a description of the M/M/C queuing model.)

The average number of missed RPS reconnect attempts is given by the formula:

$$N = \frac{U}{1 - U}$$

where U is the probability that a reconnect attempt will find all paths busy

The average time spent attempting to reconnect is simply the rotational time of the device multiplied by the average number of missed RPS reconnect

attempts. This yields an estimate of the average I/O delay caused by missed RPS reconnect attempts.

Rule DAS220 is produced if the average time spent attempting to reconnect to the channel path accounts for a significant percent of the device response time **to the "loved one" workload**.

The I/O delay caused by missed RPS reconnect may be under-estimated if the device is shared between systems. This is because the controller may be busy to a different system when the device attempts to reconnect. There is no information in RMF records to show the time the controller was busy to another system. RMF data from all systems sharing the controller must be consolidated and analyzed in order to obtain this information. This analysis is performed by the DASD Component only if you have selected the "analysis of shared DASD conflicts" option (see Section 3, Chapter 2.5).

Suggestion: Missed RPS reconnect time is caused by too much activity on the channel paths to the device. This problem can be corrected by:

- Removing I/O data transfer from the path(s). When devices are seeking or searching for a sector, the channel path is not busy. Therefore, high channel activity is primarily due to transferring data (controller protocol accounts for some path busy, but this time generally is quite small). Removing I/O data transfer from the path(s) can be accomplished by:
 - Moving data sets from the paths. This often is the easiest solution, since data sets with high data transfers can be relocated to other volumes on different channel paths.
 - The volumes responsible for high data transfer time could be moved to other strings on different channel paths.
- Rescheduling workloads to minimize the contention. For example, some batch jobs may be performing heavy I/O to volumes on the string. These jobs may be rescheduled to a time when their I/O would not cause problems.
- Adding paths to the device. If the device is not capable of DLS or DLSE mode, this will require upgrading the device (e.g., from IBM-3380 Model A04 to a more recent model). If the device is capable of DLS mode (that is, it can dynamically reconnect to either of two paths), you may consider upgrading the device to one capable of DLSE mode (that is, it can dynamically reconnect to any of four paths). This action may require that you upgrade your controller from an IBM-3880 storage controller to an IBM-3990 storage controller. Since upgrades are relatively expensive,

you should first assess the feasibility of moving active data sets or volumes from the path.

- Adding channel paths, acquiring additional controllers, and moving some volumes to the new controllers. This would be a more expensive option, and may not be feasible (depending upon the processor and I/O configuration).