
Rule DAS110: SEEKING WAS THE MAJOR CAUSE OF RESPONSE DELAY

Finding: Seeking was the major cause of the I/O response delay with the device.

Impact: This finding can have a MEDIUM IMPACT or HIGH IMPACT, depending upon the importance of the applications referencing the device and on the amount of seeking being done.

Logic flow: The following rule causes this rule to be invoked:
 DAS100: Volume with the worst overall performance

Discussion: As described in Chapter 1 of Section 5 of this User Manual, DISC time is composed of seeking, latency, and missed RPS reconnect. The discussion associated with RULE DAS100 describes how CPEXpert creates a model of the I/O configuration based upon the IOCP macros and RMF data.

CPEXpert applies queuing formulae to the model to estimate the amount of delay attributed to missed RPS reconnect (these delays are a function of the probability of a device finding all paths busy when the device tries to reconnect to the channel path).

The estimated missed RPS reconnect time is subtracted from the DISC time reported by RMF. Additionally, the average latency for the device type is subtracted from the DISC time. The resulting time is assumed to be the seek time. (Note the below discussion about why this assumption might not be correct.)

CPEXpert performs the above analysis for each measurement interval reflected in the data. RULE DAS110 is produced if seeking was the major problem for a majority of the measurement intervals.

There are potential problems with this approach, although the approach is generally used throughout the computer industry as a way of estimating missed RPS delays and of estimating seeking.

- The queuing formulae assume exponential interarrival times, exponential service distributions, and an infinite population (the M/M/c formula - Erlang's C formula - is used for the calculations). These assumptions may not be correct if, for example, the I/O activity is a function of a single application.

In his class "MVS I/O Configuration Management", Dr. Jeffery Buzen provides a Dump/Restore application as an excellent example of an application which does not follow standard queuing assumptions.

- The device may be cached, and it may be impossible to apportion the DISC time residual after subtracting missed RPS reconnect time. This time may represent a few missed cache read operations with long seek distances, or may represent a relatively large number missed cache read operations with little seeking but the standard latency for the device.

Thus, the seeking analysis can only show potential problems, rather be considered a definitive indication. However, it is usually a fairly accurate indication of the problem. If high average seeking is reported, you can be fairly certain that high seeking did occur. This is particularly true if the problem is reported throughout the measurement intervals. The "uncertainty" tends to be related to relatively low seeking or seeking reported for cached devices.

RULE DAS110 reports the overall average number of milliseconds out of each second in which the device was positioning the arm. Additionally, RULE DAS110 summarizes key information about the period of worst performance, if seeking was the major cause of delay during this period.

Suggestion: The seeks can be minimized by (1) rearranging files within the pack, (2) moving files from the pack to another actuator, (3) changing the application file accessing characteristics, or (4) possibly restricting the applications allowed to access the pack.