Rule WLM365: Non-paging DASD disconnect time was a major cause of DASD delay

Finding: CPExpert has determined that non-paging DASD disconnect (DISC) time was a major cause of delay in DASD response for the I/O operations of the service class.

Impact: This finding may have a MEDIUM IMPACT or HIGH IMPACT on the performance of the service class. This finding applies only with OS/390 Release 3 and subsequent versions.

Logic flow: The following rule causes this rule to be invoked:
Rule WLM361: Non-paging DASD I/O activity caused significant delays

Discussion: DISC time is the time (1) from when the controller initiates a SEEK Channel Command Word (and the seek requires an arm movement) on the device until the SEEK command is complete, (2) plus the time of the rotational delay while the SET SECTOR Channel Command Word is executing, and (3) plus the rotational position sensing (RPS) delay time required because of missed RPS reconnect. Please see Rule WLM350 for a discussion of DISC time.

CPExpert examines the non-paging DASD I/O disconnect time contained in SMF Type 72 records (field R723CIDT). CPExpert produces Rule WLM365 if the average DISC time accounted for a significant percent of the I/O time for transactions in the service class missing its performance goal.

If the service class missing its performance goal is a transaction service class (for example, composed of CICS 4.1 transactions), CPExpert will identify the server service class (for example, the CICS region). CPExpert will then analyze the DASD I/O times for the server.

The following example illustrates the output from Rule WLM365:

```
RULE WLM365: NON-PAGING DASD DISCONNECT TIME WAS A MAJOR CAUSE OF DELAYS

   CICSDEFA: A major part of the delay to the SYSSTC server was due to non-
   paging DASD device disconnect (DISC) time. Disconnect time is caused
   by normal rotation delay (latency), seeking, or missed rotational
   position reconnect (RPS).  Please refer to the WLM Component User Manual
   for advice on how to minimize device disconnect time.
```


Suggestion: Large DISC time is caused by device seek delay, device rotational delay, or device missed RPS reconnect delay.

- **Seek delay**. The SEEK command is responsible for positioning the arm to the proper cylinder. If no positioning is required (that is, the arm is already at the proper cylinder), the device is not disconnected. Seek operations occur because of accessing patterns with data sets and because of accessing patterns between data sets. Large seek times usually involve the following situations:

  - Multiple data sets being active on the volume. The data sets can be redistributed among different volumes, to eliminate the seeking on the single volume.

  - Multiple users using the same data set on the volume. While only one data set is involved, the user or application accessing patterns may require frequent arm movement. A partitioned data set in which several TSO users reference different members is a common situation.

    Depending upon the data set characteristics, duplicate copies of the data set placed on different volumes may solve the seeking problems.

- **Rotational delay**. The SET SECTOR command is responsible for locating the proper sector on the track as the disk rotates. (Actually, the SET SECTOR command locates a sector three sectors preceding the desired sector. This sector is called the *angular sector*.) The device is disconnected during the SET SECTOR command operation.

  The rotational delay may be from zero, to the total time required to rotate the disk to the required sector. It is possible that the required sector will be immediately under the head. In this case there is zero rotational delay. On the other hand, the sector could have just passed under the head before the SET SECTOR command was received by the drive. In this case, a full rotation must be accomplished before the required sector is located.

  On average, one-half of the rotation time will be required to locate the sector. This time is referred to as the **average latency** of the device. For example, IBM-3390 devices rotate every 14.1 milliseconds and the average latency is 7.1 milliseconds[1].

[1]These values apply only to Model 1, Model 2, and Model 3. IBM-3390 Model 9 devices rotate every 45.1 milliseconds and the average latency is 22.8 milliseconds.

It is important to realize that the latency is an average based upon many SET SECTOR commands. Any particular SET SECTOR command may have a latency ranging from zero to the maximum rotational delay.

If there are few I/O commands for a particular device in a given measurement interval, it is uncertain what the average latency will be. However, if there are many I/O commands for a particular device in a given measurement interval, the average latency will normally be one-half of the rotational delay.

The average latency may be (and should be) quite small with cached devices. This is because many I/O requests should be satisfied from the cache and have no latency.

If the average DISC per I/O operation is approximately equal to the average latency for the device type, the device would benefit from caching. If the average DISC per I/O operation is much larger than the average latency for the device type, it is unclear whether the device would benefit from caching.

- **Missed RPS reconnect**. The device attempts to reconnect to the path when the angular sector is reached (the angular sector is described above). If the reconnect attempt is successful before the desired sector is reached, then the device connects and the read or write operation can proceed.

  A path must be available for the device to reconnect successfully. If the reconnect is not successful before the desired sector is reached, then the device does not connect, and a complete revolution of the track must occur before the angular sector is again reached. This is called a *missed rotational position sensing reconnect (or missed RPS reconnect)* delay.

  There is no action which can alleviate the initial rotational positioning delay (aside from changing device characteristics, such as implementing caching or buffering at the device level). Over a large number of I/O operations, this initial delay will be one-half the rotation times.

  However, the missed RPS delay is a function of the probability that the path will be busy when the device attempts to reconnect; the busier the path(s), the more missed RPS delay. (Note that the path busy time is a function of the connect time of other actuators. The path cannot be busy from the device itself when the device attempts to reconnect.).

