
Rule WLM366: Non-paging DASD IOSQ time was a major cause of DASD delay

Finding: CPExpert has determined that queuing in the I/O Supervisor (IOSQ) for non-paging DASD 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 Version 2 Release 4 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: IOSQ time is the time from the issuance of a STARTIO macro until the Start SubChannel (SSCH) instruction is issued. After the STARTIO macro is issued, the software determines whether the device is busy. If the device is not busy with this system, the SSCH instruction is issued. However, if the device is busy with this system, the I/O request is queued. Thus, IOSQ time always means that the device is unable to handle additional requests from this system.

Some small IOSQ time is often unavoidable. However, large IOSQ time imply a situation that should be examined. Large IOSQ times result from (1) too many I/O operations directed to the device or (2) lengthy device response times (perhaps caused by high seeking, high RPS delays, or high PEND time).

The following example illustrates the output from Rule WLM366:

```
RULE WLM366: NON-PAGING DASD IOSQ TIME WAS A MAJOR CAUSE OF DELAYS

  BATCHMED: A major part of the delay to the service class was due queuing
  in the I/O Supervisor (IOSQ) for non-paging DASD devices. IOSQ time is the
  time from the issuance of a STARTIO macro until the Start SubChannel (SSCH)
  instruction is issued. Please refer to the WLM Component User Manual
  for advice on how to minimize device IOSQ time.
```

Suggestion: Large IOSQ times usually involve the following situations:

-
- 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 for the single volume.
 - Multiple users may be using the same data set on the volume. Depending upon the data set characteristics, duplicate copies of the data set placed on different volumes may solve the IOSQ problems.
 - Multiple application systems may be using the volume experiencing high IOSQ times. In this case, perhaps application redesign or scheduling can solve the problem.
 - A particular application (or system function) may be executing I/O to the device faster than the device can respond.
 - The overall device response time (PEND, DISC, and CONN) times may be large, such that the device is unable to provide quick response to the I/O requests. This situation will be revealed by large values in the PEND, DISC, or CONN measures.

Depending on the amount of IOSQ time involved, on budget considerations, and on the business importance of the work being delayed, you might consider acquiring Parallel Access Volumes (PAV). The PAV design tends to eliminate IOSQ time.