
Rule DAS130: PEND TIME WAS MAJOR CAUSE OF I/O DELAY

Finding: CPExpert has determined that PEND time 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.

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

Discussion: PEND time is the time from the issuance of the StartSubChannel (SSCH) instruction until the device is selected by the control unit and physical positioning commands (such as seek and set sector, or define extent) are transferred to the device.

With modern fixed block architecture (FBA) devices, the PEND time ends when the physical positioning commands are presented to the *logical volume control block* within the control unit. The PEND time is caused by queuing for the path (wait for channel, wait for director port, wait for control unit, or wait for device, or wait for “other” reasons)¹.

PEND is measured by the channel subsystem. After IOS issues the Start Subchannel command, the channel subsystem may not be able to initiate the I/O operation if any path or device busy condition is encountered:

C The channel selected for the I/O operation could be busy with another I/O operation from another system image in the same CEC. This time is not reflected in the SMF data.

C The director port could be busy with another I/O operation². This time is reflected in SMF data as SMF74DPB.

C The control unit could be busy with another I/O operation from another system. This time is reflected in SMF data as SMF74CUB.

¹PEND time is significantly reduced with FICON channels. FICON channels can have multiple I/O operations concurrently active, which reduces the potential PEND time caused by channel busy. There is no port busy time with FICON switches, and control unit time is significantly reduced. This statement regarding PEND time is not necessarily correct if a large number (more than 5) I/O operations are concurrently executing on a FICON channel. Dr. H. Pat Artis and Mr. Robert Ross have presented the results of research indicating that performance can degrade significantly when more than 5 I/O operations (Open Exchanges) are concurrently active on a FICON channel (see “Understanding FICON Channel Path Metrics” at www.perfassoc.com).

²Director port busy can occur only on an ESCON channel. The use of buffer credits on a FICON native channel eliminates director port busy.

C The device could busy with I/O from another system. This time is reflected in SMF data as SMF74DVB.

There can be “other” PEND time not reflected in the above descriptions. For many systems, “other” PEND time is zero or very small. For some systems, the “other” PEND time is dramatically large (often, 75% or more of the average response time).

One possible cause of the “other” PEND time is PEND for channel busy. If all channels between the MVS image and the device are busy, the channel subsystem must wait until a channel becomes available. This wait for channel is reflected in PEND time. Depending on the number of MVS images using the channels to the device, channel activity could be high. This activity could (and often would) be caused by activity to other logical volumes, rather than the device exhibiting poor performance.

As mentioned earlier, PEND for channel busy is not reflected in the SMF data³. However, CPEXpert calculates an estimated PEND for channel busy based on I/O configuration information.

When CPEXpert creates the model of the I/O configuration, it retains information about each path to a device. Included in this path information is the physical path busy at the CEC level, for each path. Consequently, CPEXpert has an overall view of all physical paths to the device, and can calculate overall channel activity for all channels to the device.

After computing an estimated PEND for channel busy, CPEXpert computes “other” PEND by subtracting estimated PEND for channel, PEND for director port, PEND for control unit, and PEND for device from the total Device Pending time contained in SMF Type 74 records for a particular device.

At present, there is only conjecture⁴ about additional cause of this “other” PEND time. Perhaps either IBM will better describe this “other” PEND time in future, or perhaps research will reveal likely causes of the “other” PEND time.

PEND time can be significant with shared systems. If one system does an I/O request to a device while the storage subsystem is already processing an I/O to that device that came from another system, then the storage

³MXG contains a variable AVGPNCHA (titled “AVG (MS)*PEND DUE TO*CHANNEL BUSY”). However, the MXG AVGPNCHA variable is simply created from the AVGPNDIR (“AVG (MS)*PEND DUE TO*DIRECTOR PORT”) variable. MICS does not contain a “PEND CAUSED BY CHANNEL BUSY” variable.

⁴According to MXG (ADOC74 comments), Dr. H. Pat Artis believes that the “other” PEND is often the internal response time of the subsystem, i.e., the time it takes the subsystem to accept, validate, and acknowledge the first Channel Control Word (CCS) of the channel program.

subsystem will send back a *device busy* indication, resulting in PEND time. This delays the new request and adds to processor and channel overhead.

CPEXpert computes the average per-second PEND delay time, for each of the causes listed above. Rule DAS130 is produced if the average PEND time accounted for a significant percent of the device response time.

The following example illustrates the output from Rule DAS130:

```
RULE DAS130: MAJOR CAUSE OF I/O DELAY WAS PEND TIME.

A major cause of the I/O delay with VOLSER PPVOL1 was PEND time. The
average per-second PEND delay for I/O is shown below:
```

MEASUREMENT INTERVAL	PEND		PEND		PEND		TOTAL
	CHAN	DIR PORT	CONTROL	DEVICE	OTHER	PEND	
8:30- 8:45,22OCT2001	0.013	0.000	0.000	0.003	0.142	0.158	
8:45- 9:00,22OCT2001	0.018	0.000	0.000	0.004	0.205	0.226	

Suggestion: No suggestions are associated with this finding. CPEXpert will analyze the high PEND delay time. The following rules will be produced to indicate which is the major cause:

C DAS131: PEND time was caused by channel busy

C DAS132: PEND time was caused by director port busy

C DAS133: PEND time was caused by controller busy delays

C DAS134: PEND time was caused by device busy delays

C DAS135: PEND time was caused by other delays

Please note that the DAS132-DAS135 rules are "LEVEL-2" rules, which means that they will not be produced unless you have specified %LET VERBOSE=VERBOSE in USOURCE(DASGUIDE). If the rules are not produced, you can simply examine data associated with Rule DAS130 to select the major cause. Then you can examine the documentation associated with the major cause.

Rule DAS131 is a "LEVEL-1" rule because it provides additional information on physical channel busy times, by CHPID.