Rule DAS300: Perhaps shared DASD conflicts caused performance problems

Finding: CPExpert believes that accessing conflicts caused by sharing DASD between systems or MVS images may have caused performance problems.

Impact: This finding is used to assess whether sharing DASD between systems or MVS images caused performance problems.

Logic flow: The following rules cause this rule to be invoked:

DAS100: Volume with the worst overall performance
DAS110: Seeking was major cause of I/O response delay
DAS120: Missed RPS reconnect was major cause of I/O delay
DAS130: Large PEND time was major cause of I/O delay
DAS150: Missed cache read hits was major cause of I/O delay

Discussion: DASD volumes can be shared between systems or between MVS images. Sharing of the DASD volumes might be implemented to allow backup of data, to facilitate recovery or restart, to permit transfer of data from one system to another, etc.

In some situations, sharing DASD volumes has no impact on performance (for example, few I/O operations might be directed to the shared volumes from potentially conflicting systems).

In other situations, sharing DASD volumes can have a significant impact on the performance of the shared volumes, and consequently, on the performance of the applications accessing the shared volumes.

CPExpert can perform an analysis of conflicts between DASD shared between systems or MVS images. The analysis performed by CPExpert is not intended to identify an isolated performance problem. Rather, CPExpert attempts to identify those problems which **continually** cause shared DASD performance problems.

Assume that CPExpert is analyzing DASD performance problems with System A, and some DASD devices are shared with System B. There are three major ways in which the shared DASD can cause performance problems for System A:

- **Case #1**. The arm of the device has been moved by System B when System A attempts to access a cylinder. The DASD I/O operation

from System A must move the arm to the desired cylinder, as a SEEK operation.

CPExpert computes the time required to perform SEEKs on System A. If the computed SEEK time is a major cause of performance problems, CPExpert analyzes the data from System B to determine whether System B generates a large number of I/O operations to the device.

- If System B does **not** generate a relatively large number of I/O operations to the device, CPExpert concludes that there is **not** a conflict. There is little doubt about the validity of this conclusion. If System B does not direct much I/O activity to the device, System B clearly cannot cause seek problems for System A.

- If System B **does** generate a relatively large number of I/O operations to the device, CPExpert concludes that there **is** a conflict caused by sharing the device.

  To be absolutely correct, CPExpert should process the configuration definitions for System B, process System B's channel and device information, and compute seek information for System B. CPExpert could then determine whether System B also experienced a high seek rate for the device.

  If both System A and System B experienced a high seek rate, CPExpert could be absolutely sure that there was a shared DASD conflict. This approach would unnecessarily use system resources and would be cumbersome to implement.

  Consequently, CPExpert makes the assumption that I/O operations to the device are random between System A and System B. CPExpert thus can conclude that if System A experiences a high seek rate and System B significantly uses the device (exhibited by a high I/O rate), then System B must also experience a high seek rate. To assume otherwise would require that I/O from System B be coordinated with the I/O from System A, such that System B does not experience seeking similar to System A.

  Since System A experiences a high seek rate and System B experiences a high I/O rate, CPExpert concludes that there is a conflict caused by the shared DASD.

- **Case #2**. The device (or the path to the device) is busy to System B when System A attempts to reference the device. The DASD I/O

operation from System A must wait until the device (or path) is free. This wait time is directly reflected as PEND time in the RMF measurement data.

CPExpert determines whether PEND time is a major cause of performance problems for the device (or devices) selected for detailed analysis. If PEND time is a major cause of performance problems, CPExpert analyzes the data from System B to determine whether DASD I/O operations from System B could cause the PEND delays for System A.

CPExpert analyzes the total CONN time and DISC time for the device from other systems sharing the device. CPExpert concludes that the PEND delay is caused by other systems if the total CONN time and DISC time is more than 25% of the PEND delay for the device experienced on System A.

This analysis does not detect PEND delays caused by path delays (e.g., the controller being busy). In order to analyze all path delays, CPExpert would have to process the IOCP macros for all systems and perform a complete analysis of all other I/O configurations. This seems to be an excessive amount of analysis for only marginal gain.

Not detecting PEND delays caused by path conflicts should not be a serious deficiency in the analysis. For most system, the PEND time is not caused by path delays (most systems have sufficient paths for shared DASD that path delays from another system are not serious). The PEND delays would normally be caused by the other system using the device. Consequently, the analysis which is performed should detect most shared DASD problems exhibited by PEND delays.

- **Case #3**. The device is cached, the cached controller is shared between System A and System B, and I/O operations from System B preempt data in the cache which is required by System A.

  This problem is evidenced by relatively large I/O times for System A. Since the data required by System A is not in the cache, the data must be acquired from the device. This potentially results in a seek operation to position the arm, and will require normal latency and data transfer from the device.

  In any of the above situations with cached shared devices, CPExpert analyzes the amount of I/O operations from System A and System B to the cached controller.

- If System B does **not** generate a relatively large number of I/O operations to the device, CPExpert concludes that there is **not** a conflict.

- If System B **does** generate a relatively large number of I/O operations to the device, CPExpert concludes that there **is** a conflict caused by sharing the device. The rationale for these conclusions is the same as was discussed under Case #1 (seek conflicts).

The standard analysis performed by CPExpert may not detect a cache problem under two possible scenarios.

- If only one cached controller is attached to System A, CPExpert may not detect a problem with the device. This is because the logic employed by CPExpert selects devices with the most performance improvement within each **type of device** and then selects the **overall** "worst" devices for detailed analysis.

  CPExpert considers cached devices to be a unique type. If all devices on the cached controller received bad service caused by shared cached problems, CPExpert may not detect a performance problem with the cached devices. This is because all devices in the "device type" could have roughly equal poor service and no device would be **significantly** worse than the other devices in the device type. Consequently, CPExpert might not select any of the cached devices for detailed analysis

  If there are multiple cached controllers, there would be a larger number of "candidate" devices, and the standard analysis performed by CPExpert is more likely to identify any problem caused by shared cache controllers.

- The analysis performed by CPExpert may not detect a cache problem if the cache is being replaced with data from another volume. It is possible that System B could cause data from another volume to be loaded into the controller's cache. This volume could be a volume not be flagged as the "worst" performing volume when analyzing performance from the perspective of System A. From System A's perspective, accesses to the worst volume would simply not find required data in cache.

  Without analyzing the IOCP information for all systems, CPExpert cannot determine which volumes are attached to which controllers. Consequently, CPExpert cannot at present relate (1) poor performance for one volume on System A and (2) I/O operations to a different volume by System B.

If you suspect problems because of shared cached devices, you can direct CPExpert to analyze the specific devices (using the SELECT option in DASGUIDE). CPExpert will then analyze only the devices selected. The analysis of shared DASD is accomplished you have specified **%LET SHARED = YES;** in USOURCE(DASGUIDE) to direct CPExpert to analyze shared DASD.

CPExpert performs the following processing if you have indicated that an analysis of potential conflicts between shared DASD should be performed:

- CPExpert determines whether the IOCP has indicated that the "worst" device selected for detailed analysis is attached to a control unit shared with another system[1]. If so, CPExpert performs an analysis of potential conflicts caused by shared DASD.

- CPExpert identifies other systems which reference the "worst" device. This identification is accomplished by analyzing the SMF Type 74 data in the performance data base relating to all other systems. The SMF Type 74 data contain the VOLSER for each device referenced. CPExpert simply selects SMF Type 74 information for the systems which reference the VOLSER of the "worst" device. This information is retained for more detailed analysis about potential conflicts.

  There is a potential problem with this method of identifying devices shared between systems. Multiple systems in the performance data base could use the same VOLSER to identify different devices. This could happen if the devices were not shared between systems.

  For example, suppose that CPExpert had identified PAGE01 as the "worst" device. Several system in the performance data base could reference VOLSER PAGE01, but the devices with VOLSER PAGE01 could be unique to each system. CPExpert would assume that all references by other systems to PAGE01 applied to the "worst" device being analyzed. The references could apply to a totally different device, and the other systems might not even share DASD with the system being analyzed.

  If this should be a problem (that is, if the DASD Component reports shared DASD conflicts with systems which do not share the device being analyzed), simply ignore the analysis produced by CPExpert[2].

[1]CPExpert tests whether the SHARED keyword of the CNTLUNIT macro in the IOCP was specified as either SHARED=Y or SHARED=YB.

[2]We do not feel that this problem will be common. It is described only to alert you to a potential incorrect analysis. If any user encounters this problem and it becomes annoying, code can be implemented to allow users to identify specific systems which share DASD with the system being analyzed. At present, this option seems to add unnecessary complexity to the user options.

- Once CPExpert has identified all systems which reference the "worst" device, CPExpert analyzes the DASD I/O characteristics of these systems with respect to the "worst" device. As described earlier, the analysis makes a basic assumption that the I/O activity from the different systems is random among the systems (that is, the I/O activity of System B is independent from the I/O activity of System A).

  If the I/O activity is random, CPExpert can conclude that delays on System A are caused by shared DASD conflicts only if (1) System B experienced the same delays with the "worst" device (in the case of PEND delays) or (2) System B generated a high I/O activity to the "worst" device.

  CPExpert can conclude that delays on System A are not caused by shared DASD conflicts if the data from System B does not exhibit these characteristics. The earlier discussion explains the rationale for these conclusions.

CPExpert will list statistics relating to potential conflicts, by system, by volume, and by RMF measurement interval. Additionally, CPExpert will identify the major cause of performance problems with the shared volume (i.e., seeking, PEND time, missed RPS delays, or missed cache hits).

Suggestion: You should use the information displayed by Rule DAS300 to assess the significance of the performance problems caused by shared DASD.

- You may conclude that the performance delays are acceptable, in which case you may take no action.

- You may conclude that the performance delays are not acceptable, in which case you may wish to take actions such as:

  - Schedule conflicting applications so that they do not access the volume at the same time.

  - Duplicate data or files and placing the duplicate information onto different volumes. This may be feasible for read-only data.

  - Move conflicting files to another shared volume.

