
Rule WLM250: Service Class waited for access to CPU

Finding: CPExpert has determined that waiting for access to a CPU was a major cause of the service class not achieving its performance goal.

Impact: The impact of this finding depends upon the percent of time transactions in the service class were denied access to a CPU. A high percent denied CPU access means HIGH IMPACT while a low percent denied CPU access means LOW IMPACT. See the output associated with the rule that caused this rule to be invoked (Rule WLM101 to Rule WLM103, depending upon the type of service class and performance goal).

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

- Rule WLM101: Service Class did not achieve average response goal
- Rule WLM102: Service Class did not achieve percentile response goal
- Rule WLM103: Service Class did not achieve execution velocity goal

Discussion: As the System Resources Manager takes its samples of the state of address spaces, it examines whether a TCB or SRB associated with the address space is waiting for dispatching to a CPU, or whether a TCB is waiting for a local lock.

If an address space is waiting for dispatching, it is being denied access to a CPU because processors are active with higher priority address spaces or with address spaces at the same dispatching priority as the address space waiting for dispatching. Samples reflecting the time address spaces are denied access to a CPU are recorded by RMF in the SMF Type 72 delay samples, as CPU Delay (R723CCDE)¹.

CPExpert computes the percent of CPU Delay for the service class, as a function of the overall execution of transactions executing in the service class. CPExpert produces Rule WLM250 if the percent of CPU Delay for the service class is greater than the significance value specified in the **WLMSIG** guidance variable in USOURCE(WLMGUIDE).

With Rule WLM250, CPExpert provides the total number of ending transactions in the RMF measurement interval, the total CPU service units

¹The address space could also be waiting for dispatch because the Workload Manager has marked the TCB or SRB "non-dispatchable" because of CPU Capping. Please see Section 4 (Chapter 1.6) for a discussion of resource groups and how the Workload Manager implements the resource group specifications. The CPU Delay samples recorded in R723CCDE do **not** include any samples of waiting because of CPU Capping. CPU Capping Delay is recorded in a separate SMF Type 72 variable (R723CCA).

consumed by the service class, and the percent of active time when transactions in the service class were denied access to a CPUs.

Additionally, CPEXpert provides summary information about the CPU time used by service classes with higher importance, the same importance, and lower importance with respect to the service class failing to achieve its performance goal.

The CPU time used by other levels of importance can be used in association with the CPU USED by the service class missing its performance goal, to assess whether the problem is caused by service classes with a higher importance or service classes at the same level of importance.

The following example illustrates the output from Rule WLM250:

```
RULE WLM250: SERVICE CLASS WAITED FOR ACCESS TO CPU

Service Class TSO (Period 1) was delayed waiting for CPU
dispatching. During the following RMF measurement intervals, a TCB
or SRB was waiting to be dispatched, or a TCB was waiting for a local
lock. The "% Denied CPU" value represents the percent of TSO's
active time when TSO was waiting for access to a CPU. CPEXpert
will produce a report at the end of this analysis that shows the CPU
time used by all service class periods.
```

MEASUREMENT INTERVAL	CPU USED (TSO -- 1)	% DENIED CPU	CPU TIME USED BY OTHER ---LEVELS OF IMPORTANCE---		
			HIGHER	SAME	LOWER
13:02-13:07,21JUN1994	0:02:10	31.5	0:00:43	0:05:31	0:00:00
13:07-13:12,21JUN1994	0:02:09	29.6	0:00:51	0:05:30	0:02:14
13:17-13:22,21JUN1994	0:02:14	50.9	0:00:49	0:05:42	0:02:09
13:22-13:27,21JUN1994	0:02:09	35.9	0:00:45	0:05:25	0:02:10

Please note that CPEXpert does not produce Rule WLM250 for "served" service classes (e.g., a service class describing CICS transactions). The SRM does not collect resource information for "served" service classes. Rather, the SRM collects resource information at the "server" service class level (e.g., at the CICS region). CPEXpert will analyze the "server" service class to identify constraints and Rule WLM255 may result from this analysis.

Suggestion: When a service class fails to achieve its goal because it is denied access to a CPU, you have several alternatives:

- **Increase the importance of the service class.** The Workload Manager attempts to achieve the performance goal for each service class. When the Workload Manager detects that a service class is not achieving its performance goal, the Workload Manager will assess whether changing

the existing distribution of system resources will help a service class achieve its performance goal².

The Workload Manager examines (and attempts to help) service classes in descending order of importance. Importance levels may be specified as values of 1 to 5, with Importance 1 being the most important and Importance 5 being the least important. Importance 0 is an implied importance level for system tasks, and Importance 6 is an implied importance level for service classes with a Discretionary performance goal.

If you increase the importance of a service class, the Workload Manager will give a higher priority to the service class when resources are allocated. Of particular relevance to the problem of a service class being denied access to a CPU is that the Workload Manager may assign a higher dispatching priority to address spaces in the service class if the service class is missing its goal. With a higher dispatching priority, the service class will be less likely to be denied access to a CPU.

- **Decrease the importance of another service class.** The Workload Manager will attempt to provide resources to help service classes missing their performance goal. As described above, the Workload Manager examines (and attempts to help) service classes in descending order of importance.

You should examine the importance specified for (1) service classes with a higher importance and (2) service classes at the same importance as the service class missing its performance goal. Determine whether these importance levels match the management objectives of your installation.

- **Alter the performance goal specified for the service class.** You should assess whether the performance goal is appropriate for the applications assigned to the service class. Perhaps the performance achieved is adequate, or perhaps the specified performance goal can be altered so that the service class meets its objective at the existing level of service. That is, the delivered service may be adequate for management objectives and you may need to change the performance goal specified to the Workload Manager.
- **Alter the performance goal specified for another service class.** You should assess whether the performance goal is appropriate for the applications assigned to other service classes. The Workload Manager attempts to achieve the performance goal for each service class. When the Workload Manager detects that a service class is not achieving its

²Please refer to Section 4 for a more comprehensive discussion of the Workload Manager's algorithms.

performance goal, the Workload Manager will assess whether changing the existing distribution of system resources will help a service class achieve its performance goal.

As described above, the Workload Manager first examines service classes based on importance. However, if several service classes are of the *same* importance, the Workload Manager will attempt to help the service class having the *worst* performance (as measured by the performance index).

You should assess whether appropriate performance goals have been specified for other service classes at a higher importance or at the same importance.

- **Evaluate the effects of HiperDispatch.** With z10, HiperDispatch attempts to minimize the number of logical processors in use by the LPAR. To accomplish this, work can be dispatched to fewer logical processors. Logical processors in excess of the number required to actually achieve the share of processor capacity guaranteed to the LPAR, are placed in a “parked” status. Parked logical processors are not considered available for dispatching by the z/OS Dispatcher.

Please refer to Rule WLM254 for additional information about the effect of HiperDispatch on the performance of individual work units.

- **Reschedule workloads.** Your organization may be able to reschedule conflicting workloads to another system to eliminate the conflicts for processor access.
- **Add another processor.** You may be able to add another processor (potentially not so difficult in an LPAR environment). Adding another processor will provide another "CPU server" from a queuing model view; having another "CPU server" significantly reduces the probability that an address space will be denied access to a CPU³.
- **Acquire faster processors.** If the service class missing its performance goal is sufficiently important and it is being denied access to a CPU, you may be able to solve the problem by acquiring faster processors.
- **Ignore the finding.** There may be situations in which you wish to simply ignore CPEXpert's finding. You might not care that a low priority batch service class is denied access to the CPU. If this is the case, perhaps you should not have a performance goal associated with the workload.

³Please refer to *Probability, Statistics, and Queuing Theory* by Arnold O. Allen for a description of the M/M/C queuing model that can be used to assess the effect of changing the number of processors.

However, you may wish to have a performance goal (and have CPEXpert perform analysis) simply to assess other delays. For example, you may wish to assess the auxiliary paging delays experienced by the workload.

Another (and potentially more common) reason a service class period is denied access to a CPU is caused by the inherent processing characteristics of the workload, along with the MVS dispatching algorithms. Please refer to Rule WLM251 for a discussion of this situation. Rule WLM251 will be produced if CPEXpert believes that the service class period is denied access to a CPU because of this situation.

- **Exclude the service class from analysis.** If none of the above alternatives apply and if Rule WLM250 continually produces for the service class, you may wish to exclude the service class from CPEXpert's analysis. There is little point in having findings produced that cannot be acted upon. Please see Section 2 of this CPEXpert User Manual for information on how to exclude service classes from analysis.

After CPEXpert has completed its analysis of performance constraints, a summary of CPU time used by each service class period is produced for any measurement interval in which a service class did not achieve its performance goal and the service class was significantly denied access to a processor.

The following example illustrates the report that is produced:

SUMMARY OF SERVICE CLASS CPU TIME CAPTURED IN TYPE 72 RECORDS						
MEASUREMENT INTERVAL	SERVICE CLASS	CLASS PERIOD	GOAL TYPE	GOAL IMPORT	CPU USED	% CPU
21JUN1994:13:07:01	SYSSTC	1	SYSTEM TASKS	0	0:00:36.29	6.6
21JUN1994:13:07:01	SYSTEM	1	SYSTEM TASKS	0	0:00:14.57	2.6
21JUN1994:13:07:01	CICSRGN	1	SERVER CLASS	2	0:02:00.11	21.8
21JUN1994:13:07:01	IMSCTL	1	SERVER CLASS	2	0:00:37.87	6.9
21JUN1994:13:07:01	IMSMP	1	SERVER CLASS	2	0:01:18.97	14.3
21JUN1994:13:07:01	TSO	1	AVG RESPONSE	2	0:01:08.24	12.4
21JUN1994:13:07:01	TSO	2	AVG RESPONSE	2	0:00:19.16	3.5
21JUN1994:13:07:01	TSO	3	AVG RESPONSE	2	0:00:42.02	7.6 DENIED CPU (67%)
21JUN1994:13:07:01	BATCHHI	1	EX. VELOCITY	3	0:02:02.73	22.3
21JUN1994:13:07:01	BATCHLOW	1	EX. VELOCITY	3	0:00:11.40	2.1
TOTAL SERVICE CLASS CPU TIME CAPTURED IN TYPE 72 RECORDS:					0:09:22.75	

The CPU USED column reflects the total TCB and SRB CPU time used by the service class during the measurement interval. The "% CPU USED" reflects the percent of "TOTAL SERVICE CLASS CPU TIME CAPTURED IN TYPE 72 RECORDS" that was used by the service class.

Not all CPU time is accounted for by MVS. As much as 25% CPU time has been documented in the literature as "unrecovered" CPU time - the CPU time that is not included in TCB or SRB CPU time recorded by SMF in Type 72 records. Consequently, the "TOTAL SERVICE CLASS CPU TIME CAPTURED" may be significantly less than the CPU time actually used by service classes.

CPEXpert annotates any service class that was denied access to the CPU as a primary or secondary cause of the service class failing to achieve its performance goal. Along with the annotation, CPEXpert shows the percent of service class active time when an address space was denied access to a processor.

This report will allow you to assess the CPU time used by different service classes, by level of importance. To facilitate this review, the service class information is ordered by Importance associated with each service class.

Please note that the distribution of CPU time may include CPU time associated with SERVER service classes. The goal importance of the SERVER service classes is ignored after address space start-up. The importance of the SERVER service classes is a function of the service classes being served. Consequently, the CPU times may be misleading, as the CPU times shown for SERVER service classes may be at a higher or lower importance than that defined for the SERVER service class.

CPEXpert identifies the **highest** goal importance of any served service class, and displays this highest goal importance for the server service class. **This goal importance may be different from the goal importance that was defined for the server service class using the Workload Manager ISPF panel.**

No information is available to identify the CPU time used by the server to support different served service classes. Consequently, if the served service classes have different goal importance, you may be unable to determine whether the distribution of CPU time properly reflects what was actually required to support different goal importance levels. On the other hand, if the served service classes have the same goal importance, then the report properly reflects the CPU time used at the specified goal importance level.