Rule WLM480: Target Multiprogramming Level was a major cause of delay

Finding: CPExpert has determined that the target multiprogramming level (MPL) was a major cause of the service class not achieving its performance goal.

Impact: This finding can have a LOW IMPACT, MEDIUM IMPACT, or HIGH IMPACT on performance of your computer system. The impact of this finding depends upon the percent of time transactions in the service class were waiting for target MPL before address space swap-in began. A high percent waiting for MPL means HIGH IMPACT while a low percent waiting for MPL means LOW IMPACT.

Please note that the percentages reported by CPExpert are computed as a function of **the active time of the transactions**, rather than percentages of RMF measurement interval time. The percentages show the impact of MPL delay **on the transactions**, rather than the impact of MPL delay from an overall system view. This data presentation approach is significant when the service class being delayed is a **server** service class; the MPL delays[1] represent delays to the response times of the served transaction!

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
Rule WLM104:	Subsystem Service Class did not achieve average response goal
Rule WLM105:	Subsystem Service Class did not achieve percentile response goal
Rule WLM150:	Server Service Class delays
Rule WLM151:	Server Service Class delays

Discussion: Domains are maintained by the SRM in Goal Mode, even though the specification and control of domains has been removed from the user interface. The Workload Manager creates domain control table entries for each service class period so long as the service class period is associated with address spaces (that is, the Workload Manager does not create domain control table entries for service classes representing CICS or IMS transactions).

[1]In practice, MPL delay should rarely occur for server service classes. The address spaces associated with server service classes usually are non-swappable, although some organizations do make test CICS regions swappable. If the address spaces associated with a server service class are non-swappable, the address spaces will not normally incur MPL delays.

The Workload Manager defines an MPL-in target and MPL-out target for each service class period which is assigned to a domain[2]. The MPL-in target represents the number of address spaces that must be in the swapped-in state for the service class to meet its performance goal. The MPL-out target is the maximum number of address spaces allowed to be in the swapped-in state.

The Workload Manager adjusts the MPL-in target and MPL-out target, if necessary to achieve performance goals. The adjustments are made during the policy adjustment interval.

The Workload Manager adjusts the system-wide target MPL and subsequently may adjust the target MPL for individual service class periods. The Workload Manager adjusts the target MPL for the following reasons:

- **CPU is over-utilized**. The Workload Manager will decrease the system target MPL if the Workload Manager detects that the CPU is over-utilized.

- **CPU is under-utilized**. The Workload Manager will increase the system target MPL if the Workload Manager detects that the CPU is under-utilized.

- **Too much auxiliary storage paging**. The Workload Manager will decrease the system target MPL if the Workload Manager detects that too much auxiliary storage paging occurred.

- **Too much unmanaged paging**. The Workload Manager will decrease the system target MPL if the Workload Manager detects that too much unmanaged paging occurred.

- **Storage shortage below 16 megs**. The Workload Manager will decrease the system target MPL if the Workload Manager detects that there is a shortage of storage below 16 megs.

The above system utilization actions normally will cause the target MPLs for a service class period to be adjusted.

In addition to the system utilization actions, the Workload Manager may take action because a service class period did not meet its performance goal. These actions also are taken at the policy adjustment interval. If a

[2]Actually, the targets are associated with the domain, but it is easier to think of them as being associated with the service class period since domains are no longer part of the user interface

service class period is not meeting its performance goal, the Workload Manager may increase the target MPLs for the service class period. If appropriate, the Workload Manager may concurrently decrease the target MPLs for a service class period at a lower importance.

The IBM *MVS/ESA Programming: Workload Management Services* document (see the "MPL Policy Example" in the “Examples of interpreting SMF Type 99 data” section) provides an excellent example of the Workload Manager's decision process in adjusting the MPLs to meet performance goals.

The Workload Manager may also decrease the MPLs of a service class period if required by working set management. This action would normally be taken only after the working set manager had decided that there were no opportunities to decrease system paging by managing a particular address space.

Additionally, the Workload Manager may decrease or increase the target MPLs for service class periods during "time driven housekeeping". The Workload Manager will implement time driven housekeeping to make periodic adjustments of various resource policies. During time driven housekeeping, the Workload manager may increase or decrease the target MPLs for service class periods based on its analysis of the impact the MPL "slots" have had on the performance of the service class period.

The objectives of these adjustments to the system MPL or the target MPLs of individual service class periods are to (1) allow sufficient workload into the multiprogramming set such that the system is adequately used, (2) exclude workload if necessary to prevent the system from being over-utilized, and (3) make sure that the performance goals of individual service class periods are achieved.

The Workload Manager cannot always achieve these objectives for every service class. The Workload Manage might have to exclude lower importance service classes from the multiprogramming set in order to achieve the performance goals of higher importance service classes.

CPExpert produces Rule WLM480 when the target MPL was a major cause of delay to a service class not achieving its performance goal.

The following example illustrates the output from Rule WLM480:


```
RULE WLM480: TARGET MULTIPROGRAMMING LEVEL WAS A MAJOR CAUSE OF DELAY

   The target multiprogramming level maintained by the System Resources
   Manager for Service Class BATCH (Period 1) delayed swap-in of address
   spaces in Service Class BATCH.  This finding means that an address
   space became READY, but the SRM did not start swap-in of the address
   space because of target MPL constraints.  The below information shows
   the average number of address spaces in the system, by category.
   CPExpert will produce a report at the end of this analysis which shows
   the average MPL for all service class periods.

                          AVERAGE MPL  AVG    AVG    AVG     AVG    AVG
   MEASUREMENT INTERVAL   (BATCH--1)   STC    TSO   BATCH   APPC  OPEN/MVS
   15:00-15:16,01MAR1994      8.7     97.2   51.9   13.0    0.0    0.0
```

Rule WLM480 shows the average MPL for the service class missing its performance goal, and shows the average MPL for various categories of work.

Please note that CPExpert does not produce Rule WLM480 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 WLM250 may result from this analysis.

Suggestion: Rule WLM480 should never be produced for important service classes. The Workload Manager adjusts the target MPLs every 10 seconds, if necessary. The RMF measurement interval typically is at least 15 minutes. In order for Rule WLM480 to be produced, MPL delay must be a major cause of delay for the entire RMF measurement interval. This implies that higher importance work prevented swap-in of the service class period being delayed for MPL, for the entire RMF measurement interval. Such lengthy delay without Workload Manager action would be cause for considerable alarm; your system would be significantly overloaded and able to process only the higher importance work.

You may see Rule WLM480 produced often for less important service class periods. In the above example, the service class missing its performance goal consisted of batch work, and the batch work was delayed because of MPL. You may find that this delay is acceptable for such work.

When a service class fails to achieve its goal because of MPL delay, you have several alternatives:

- **Increase the importance of the service class**. The Workload Manager attempts to achieve the performance goal for each service class period.

When the Workload Manager detects that a service class period is not achieving its performance goal, the Workload Manager will assess whether changing the existing distribution of system resources will help a service class period achieve its performance goal[3].

The Workload Manager examines (and attempts to help) service class periods 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 class periods with a Discretionary performance goal.

If you increase the importance of a service class period, the Workload Manager will give a higher priority to the service class period when resources are allocated. Of particular relevance to the problem of a service class period being denied access to the multiprogramming set is that the Workload Manager may increase the target MPLs for the service class period if the service class period is missing its goal. With higher target MPLs, the service class period will be less likely to be delayed for MPL.

- **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 service classes with a higher importance and 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

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

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 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.

- **Reschedule workloads**. Your organization may be able to reschedule conflicting workloads to another system to eliminate the conflicts for processor access.

- **Improve your paging subsystem.** This option should be considered only if your system experiences significant paging delays (recall that the Workload Manager will decrease the system MPL if the paging is excessive). You can assess the paging levels by examining output from RMF or from other monitoring tools. Additionally, CPExpert will identify common problems with the paging subsystem.

- **Add another processor or acquire a faster processor**. This option should be considered only if your system is over-utilized (recall that the Workload Manager will decrease the system MPL if the system is over-utilized). You can assess the CPU utilization levels by examining output from RMF or from other monitoring tools.

- **Add additional processor storage.** This option should be considered only if your system experiences significant paging delays (recall that the Workload Manager will decrease the system MPL if the paging is excessive). You can assess the paging levels by examining output from RMF or from other monitoring tools.

- **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 delayed for target MPL. If this is the case, perhaps you should not have a performance goal associated with the workload.

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.

- **Exclude the service class from analysis**. If none of the above alternatives apply and if Rule WLM480 continually is produced for the service class, you may wish to exclude the service class from CPExpert's analysis. There is little point in having findings produced which cannot be acted upon. Please see Section 3 (Chapter 1.1.8) for information on how to exclude service classes from analysis.

After CPExpert has completed its analysis of performance constraints, a summary of MPL levels 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 delayed for MPL reasons. The following example illustrates the report which is produced:

The AVG MPL column reflects the average number of address spaces concurrently executing during the RMF measurement interval.

CPExpert annotates any service class which was delayed for target MPL 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 delayed for MPL.

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 average MPLs may include 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.

SUMMARY OF MPL LEVELS WHEN A SERVICE CLASS WAS DELAYED FOR MPL REASONS

MEASUREMENT INTERVAL	SERVICE CLASS	CLASS PERIOD	GOAL TYPE	GOAL IMPORT	AVG MPL	
01MAR1994:15:00:16	SYSSTC	1	SYSTEM TASKS	0	60.0	
01MAR1994:15:00:16	SYSTEM	1	SYSTEM TASKS	0	15.0	
01MAR1994:15:00:16	MVSSUBSY	1	EX. VELOCITY	1	18.0	
01MAR1994:15:00:16	ST_USER	1	AVG RESPONSE	1	2.1	
01MAR1994:15:00:16	APPCFEED	1	EX. VELOCITY	2	1.0	
01MAR1994:15:00:16	BATCH	1	EX. VELOCITY	2	8.7	MPL DELAY(8%)
01MAR1994:15:00:16	BERDFEED	1	EX. VELOCITY	2	0.0	
01MAR1994:15:00:16	MONITORS	1	EX. VELOCITY	2	2.0	
01MAR1994:15:00:16	ST_TOOLS	1	EX. VELOCITY	2	4.1	
01MAR1994:15:00:16	TPNSBATC	1	% RESPONSE	2	1.1	
01MAR1994:15:00:16	TPNSEVEN	1	AVG RESPONSE	2	12.8	
01MAR1994:15:00:16	TPNSEVEN	2	AVG RESPONSE	2	1.5	
01MAR1994:15:00:16	TPNSFEED	1	EX. VELOCITY	2	1.0	
01MAR1994:15:00:16	TPNSODD	1	AVG RESPONSE	2	15.1	
01MAR1994:15:00:16	TPNSODD	2	AVG RESPONSE	2	0.3	
01MAR1994:15:00:16	TPNSODD	3	AVG RESPONSE	2	0.4	
01MAR1994:15:00:16	TPNSODD	4	AVG RESPONSE	2	0.2	
01MAR1994:15:00:16	APPC	1	EX. VELOCITY	3	1.0	
01MAR1994:15:00:16	ASCH	1	EX. VELOCITY	3	2.0	
01MAR1994:15:00:16	TPNSEVEN	3	AVG RESPONSE	3	6.7	
01MAR1994:15:00:16	TPNSODD	5	AVG RESPONSE	3	1.9	
01MAR1994:15:00:16	TPNSODD	6	AVG RESPONSE	3	0.4	
01MAR1994:15:00:16	TPNSODD	7	AVG RESPONSE	3	0.1	
01MAR1994:15:00:16	VEL3	1	EX. VELOCITY	4	2.0	
01MAR1994:15:00:16	TPNSEVEN	4	DISCRETIONARY	-	2.3	
01MAR1994:15:00:16	TPNSODD	8	DISCRETIONARY	-	2.5	

CPExpert identifies the **highest** goal importance of any served service class which had active transactions during the RMF measurement interval, and displays this highest goal importance for the server service class. **This goal importance may be different from the goal importance which was defined for the server service class using the Workload Manager ISPF panel.**

In practice, the average MPL should be relatively constant for server service classes. The address spaces associated with server service classes usually are non-swappable and typically are running for considerable periods.

