
Rule WLM252: CPU access might be denied because of Resource Group minimum

Finding: CPExpert believes that CPU access might have been denied for the service class missing its performance goal because some other service class was assigned to a Resource Group with a **minimum CPU service** specification.

Impact: This finding should be viewed as generally having a HIGH IMPACT on the performance of your computer system.

Logic flow: The following rule causes this rule to be invoked:
Rule WLM250: Service class waited for access to CPU

Discussion: When CPExpert determines that a service class waited for access to a CPU, CPExpert continues to analyze the data trying to identify why the service class was denied access.

From a simplistic view, the service class was denied access because address spaces with a higher dispatching priority used the CPU. With SP5 (Goal Mode), user do not assign dispatching priority to workloads. The Workload Manager assigns dispatching priority based on user performance goals and goal importance for different service classes, and based on how well the service classes meet their performance goals.

Additionally, users can assign service classes to resource groups, and this assignment can cause the Workload Manager to grant or deny access to a CPU by address spaces.

A resource group is simply a "named" description of the total minimum and maximum **unweighted** CPU service units per second that may be used by one or more service classes assigned to the resource group. A resource group is defined using the *Create a Resource Group* panel in the Workload Manager ISPF application. A resource group applies across an entire sysplex. Service classes¹ are assigned to a resource group using the *Create a Service Class* panel in the Workload Manager ISPF application.

¹A resource group may not be associated with a service class representing subsystem transactions (e.g., a service class defined for transactions executing under CICS or under IMS). This is because CPU resources are not monitored by the SRM for the transactions; the CPU resources are monitored at the **address space level** (e.g., the CICS region or IMS message processing region). Further, CPU dispatching occurs at the address space level, rather than at the transaction level. Since CPU usage is not collected at the transaction level and CPU dispatching is at the address space level, the Workload Manager cannot control the amount of CPU resources allocated to service classes that represent transactions.

The Workload Manager will attempt to provide the specified minimum CPU service to the resource group. The Workload Manager attempts to provide the specified minimum CPU service to the resource group by adjusting the dispatching priority of service classes assigned to the resource group. The Workload Manager will restrict service classes assigned to the resource group from using more than the specified maximum CPU service. The Workload Manager uses "CPU capping" to restrict the total amount of CPU service used by service classes assigned to the resource group.

There are potentially serious effects of specifying a minimum CPU service for a resource group. The effect is caused by the order in which the Workload Manager selects service classes for policy adjustment.

- The Workload Manager first determines whether any resource group is below the **minimum** CPU service specification. If the minimum CPU specification is not being provided, the Workload Manager takes the following actions in an attempt to provide the minimum CPU service:
 - The Workload Manager determines whether any service class assigned to the resource group is not meeting its performance goal. If any service class is not meeting its performance goal, the Workload Manager increases the dispatching priority (if appropriate) of the service class.
 - If no service classes assigned to the resource group were missing their performance goal, the Workload Manager increases the dispatching priority (if appropriate) of all service classes assigned to the resource group. The dispatching priority of all service classes assigned to the resource group (including those service classes with a discretionary goal²) may be increased.
- After the Workload Manager performs the above tasks, the Workload Manager may examine service classes based on the Goal Importance of the service classes.

The result of the above process can be that service classes with a low importance (or even service classes with a discretionary goal) can be assigned CPU dispatching priority above that that is assigned to the service classes with the highest Goal Importance! The resulting CPU dispatching priorities and CPU demands can result in service classes with high Goal Importance missing their performance goals.

²Please note that the *MVS/ESA SP5 Planning: Workload Management* document is incorrect. This document states in the *Defining Resource Groups* section that "If there is a resource group defined for a service class with a discretionary goal, workload management achieves the minimum as long as the goals of work running in any other service class are not impacted. If other performance goals are impacted, then workload management does not maintain the minimum." Based on personal communication with the Workload Manager developer who wrote the specific code that attempts to provide the minimum specified CPU service, these statements are incorrect in the IBM document and the description provided above is what actually transpires.

This might not be the effect you wish, but the Workload Manager simply follows the specific direction provided for the resource group, namely, that a minimum CPU service was specified for the resource group and this minimum should be provided.

When a service class missed its performance goal and the service class was denied access to a CPU, CPExpert determines whether any service classes were assigned to resource groups with a **minimum CPU** specification. If so, CPExpert computes the CPU service used by service classes that were assigned to each resource group with a minimum CPU specification. The computations are done separately for service classes at a lower goal importance or at the same goal importance.

The purpose of the computations is to estimate whether resource group minimum CPU specifications might have caused the service class to be denied access. The result is simply an estimate of the potential impact; the SMF Type 72 records do not contain dispatching priority for service classes (the dispatching priority is dynamically adjusted by the Workload Manager).

CPExpert produces Rule WLM252 if any service classes were assigned to a resource group with a minimum CPU specification and these service classes actually used CPU service.

Suggestion: CPExpert suggests that you verify the minimum CPU service specification for resource groups defined in the service policy. Unless there are unique requirements for the minimum CPU service specifications, CPExpert suggests that the minimum be changed to zero.

Reference: MVS Planning: Workload Management

MVS/ESA(SP 5):	Chapter 7: Defining Resource Groups
OS/390 (V1R1):	Chapter 7: Defining Resource Groups
OS/390 (V1R2):	Chapter 7: Defining Resource Groups
OS/390 (V1R3):	Chapter 7: Defining Resource Groups
OS/390 (V2R4):	Chapter 7: Defining Resource Groups
OS/390 (V2R5):	Chapter 7: Defining Resource Groups
OS/390 (V2R6):	Chapter 7: Defining Resource Groups
OS/390 (V2R7):	Chapter 7: Defining Resource Groups
OS/390 (V2R8):	Chapter 7: Defining Resource Groups
OS/390 (V2R9):	Chapter 7: Defining Resource Groups
OS/390 (V2R10):	Chapter 7: Defining Resource Groups
z/OS (V1R1):	Chapter 7: Defining Resource Groups
z/OS (V1R2):	Chapter 7: Defining Resource Groups
z/OS (V1R3):	Chapter 7: Defining Resource Groups
z/OS (V1R4):	Chapter 7: Defining Resource Groups