
Rule WLM150: Server service class delays (single transaction service class)

Finding: CPEXpert has identified delays for the server service class which provided service to a subsystem transaction service class.

Impact: This finding is provided for information purposes.

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

Rule WLM104: Subsystem Service Class did not achieve average response goal

Rule WLM105: Subsystem Service Class did not achieve percentile response goal

Discussion: When CPEXpert produces Rule WLM104 or Rule WLM105 to indicate that a subsystem service class did not achieve its performance goal, the logic of these rules tries to identify the cause of the delay. The cause of the delay initially is analyzed from the "served" service class view. Please refer to Rule WLM120 to Rule WLM132 for a discussion of the delays from the served service class.

After analyzing the **served** service class delays, CPEXpert identifies the **server** service class. The server service class normally will be one or more CICS regions or IMS regions.

CPEXpert analyzes the following possible delays to response time¹:

- **CPU Using delay**
- **Denied CPU delay**
- **CPU Capping delay**
- **Swap-in delay**
- **MPL delay**
- **Page-in delay**
- **I/O delay**

¹Please see Section 4 (Chapter 3.3) for a description of these delays.

- **Unknown delay**

CPEXpert produces Rule WLM150 to provide a summary of the delay for the server service class.

The output from Rule WLM150 does not contain the MPL delay or swap-in delay. In most environments, server service classes are non-swappable and the MPL delay and swap-in delay columns would always show zero. Consequently, CPEXpert does not clutter up the output with columns which almost always would be zero. However, CPEXpert **does** analyze these delays if any are non-zero.

More than one server service class might serve the subsystem transaction service class which missed its performance goal. In this case, CPEXpert produces multiple Rule WLM150 findings - one for each server service class. CPEXpert then analyzes the delays for each server service class.

The following example illustrates the output from Rule WLM150:

```
RULE WLM150:  SERVER SERVICE CLASS DELAYS

The IMS Service Class was served by the IMSCTL Service Class.
The IMSCTL Service Class experienced the following delays during the
measurement intervals when the IMS Service Class missed its
performance goal (the delays are shown relative to the active time
of IMSCTL):
```

MEASUREMENT INTERVAL	PCT CPU USING	PCT CPU DELAYED	PCT CPU CAPPING	PCT PAGING WAIT	PCT UNKNOWN WAIT
13:07-13:12,21JUN1994	13.3	86.7	0.0	0.0	0.0

```

RULE WLM150:  SERVER SERVICE CLASS DELAYS

The IMS Service Class also was served by the IMSMP Service Class.
The IMSMP Service Class experienced the following delays during the
measurement intervals when the IMS Service Class missed its
performance goal (the delays are shown relative to the active time
of IMSMP):
```

MEASUREMENT INTERVAL	PCT CPU USING	PCT CPU DELAYED	PCT CPU CAPPING	PCT PAGING WAIT	PCT UNKNOWN WAIT
13:07-13:12,21JUN1994	0.5	24.4	0.0	0.0	75.1

In the above example, IMS transactions were placed in the IMS Service Class. The IMS Service Class was served by an IMS control region (the IMSCTL Service Class) and IMS message processing region (the IMSMP Service Class). Rule WLM150 is produced for both servers, to show the delays to the servers.

The delay information is shown relative to the active time of the server service class, and the percentages will total 100%.

There is no information in SMF Type 72 records which shows how much of the response time of the **served** service class (e.g., the IMS Service Class) could be attributed to delays in the **individual servers** (e.g., IMSCTL or IMSMP).

If the individual servers serve more than one service class, there is information in the SMF Type 72 records to show how many **times** an address space in the server was observed to be providing service to the served service class. In this example, the IMSMP Service Class could have served (1) the IMS Service Class and (2) several other service classes². SMF data would show how many times the IMSMP Service Class provided service to IMS, IMS1, IMS2, etc.

The WLM counts each time the server issues the IWMRPT macro to indicate that a transaction has completed. This count lets the WLM know how many times the server (e.g., a CICS region) provided service to the served service class (e.g., CICS transactions).

Additionally, every 250 milliseconds, the WLM samples server service classes to see which served service classes they are serving. The sampling process ensures that the WLM keeps track of service provided to long-running transactions.

SMF field R723SCS# contains a summary of the count and samples³. This field can be used to apportion the service provided by the server to the various transaction service classes being served.

Suggestion: There are no suggestions directly associated with this finding. CPExpert will continue analysis of the server service class(es), and other rules should be produced to provide suggestions.

²For example, IMS1 Service Class, IMS2 Service Class, etc.

³SMF field R723SCS# was improperly described in early versions of the SMF manual. The field has been modified to conform with the above description after CPExpert advised IBM of the error.