
Rule WLM106: Response time distribution for service class with average response performance goal

Finding: This rule provides information about the distribution of response times during those intervals when the identified service class missed its performance goal.

Impact: This finding has NO IMPACT on performance of your computer system. The finding is provided to allow you to assess the overall performance of service classes having an average response time performance goal.

Logic flow: The following rule causes this rule to be invoked:
Rule WLM101: Service Class did not achieve average response goal

Discussion: For service classes with response goals, RMF includes in SMF Type 72 records a count of transactions that completed in varying percentages of the response goal. These transaction counts are recorded by RMF as the "Response Time Distribution Count Table" contained in SMF Type 72(Subtype 3) records. Section 4 describes the percentages recorded by RMF;

When CPExpert produces Rule WLM101, CPExpert automatically produces Rule WLM106 to provide a *summary* distribution of the response information. The purpose of Rule WLM106 is to allow you to assess whether the average response finding is meaningful, or whether there are some transactions that skew the averages.

The following example illustrates the output from Rule WLM106.

In the example, notice that 0.7% of the transactions had a response of over 400% of the 0.200 second goal. The data do not show the actual response time, but over 400% of the goal corresponds to at least 0.800 second response ($0.200 \text{ second goal} * 400\% = 0.800$). In this example, 0.7% of 137 transactions represents only 1 transaction. Thus, 1 transaction had an extremely long response, while most of the transactions experienced a response of less than 50% of the goal (or less than 0.100 seconds).

RULE WLM106: RESPONSE TIME DISTRIBUTION FOR SERVICE CLASS

Service Class TSOUSERS (Period 2) did not achieve its average response goal during the measurement intervals shown below. The response goal was 0.200 second average response. Average response can be misleading, since extremes can skew the average. The below information shows the distribution of response times:

		--PERCENT COMPLETIONS RELATIVE TO GOAL--							
		TOTAL	<50%	50-90%	90-100%	100-110%	110-200%	200-400%	>400%
MEASUREMENT INTERVAL	TRANS	GOAL	GOAL	GOAL	GOAL	GOAL	GOAL	GOAL	
12:00-12:15,08NOV1994	137	98.5	0.0	0.7	0.0	0.0	0.0	0.7	

Suggestion: If you find that some transactions skewed the findings, you may wish to consider the following alternatives:

- Since you specified an **average response goal** for the service class, perhaps you can change the goal to a **percentile response goal**. With a percentile goal, the Workload Manager would not be as concerned about the few transactions that used significantly more resources and consequently skewed the average response. Rather, the Workload Manager would base its workload management decisions on the percent of transactions that met the response goal.
- If you can identify the transactions, perhaps you can use Workload Categorization to place the transactions into a different service class. You may wish to specify a different importance and different performance goal for this new service class.
- You can simply ignore the findings that CPExpert made associated with this service class for the interval. You may decide that the transaction response is an anomaly and not take any further action. In the example shown above, only one transaction had a response significantly over the goal. It may be unnecessary to take action based on a small number of transactions exceeding the performance goal.