

Load Computation Metrics



Load Score

The load score is a time-averaged measure of the load on a given container. Load scores are published on the ‘collector’ containers and gives you an indication of the processing load of that specific collector. This is reported as a number between 1 and 100. The scores are mapped to low, medium or high severity zones as explained in the following table. A collector that is consistently operating in the High zone will mean that the collector has reached peak capacity for the given CPU/Memory resource profile.

Table 1. Load Score of a Collector

Severity	Load score
High	> 75
Medium	> 50
Low	< 50

Load Score

The load score for a collector is calculated using several metrics. The metrics and their values for each collector is displayed in the modal that pops up on clicking the icon in the Load column in the Cisco Crosswork UI.

Load score of each metric = $MetricScore \times MetricWeight$

Overall Load score of a collector = $(\text{Sum of load score of all the metrics})/10$

MetricWeight - A value between 1 to 100 representing the weight that the metric asserts towards overall load computation. While the number of metrics can change over releases, the sum of all the *MetricWeights* will always be equal to 100.

MetricScore - A value between 0 and 10 that represents the current rating of the metric in a scale between its pre-defined minimum and maximum values. Since each *MetricScore* is rated on a scale of 1 to 10, the overall load score is scaled down by 10.

The *MetricWeight* and *MetricScore* are configured in the [Load score Configuration map](#). These values are determined by empirical testing and cannot be changed.

Load Score Configuration Map

The load score configuration map assigns fixed values for different metrics as shown in the following tables.

Table 2. MetricScore of each metric

Metric	Description	MetricScore
CPU usage	<p>One minute CPU utilization of a collector as a percentage of maximum CPU resources assigned to it.</p> <p>CDG reports CPU percentage per CPU for each profile as:</p> <p>Standard CDG: If the CPU utilization is 800% (8vCPUs Standard Profile), CDG reports this as $800/8 = 100\%$</p> <p>Extended CDG: If the CPU utilization is 1600% (16vCPUs Extended Profiles), CDG reports this as $1600/16 = 100\%$</p>	<p>0 to 10%: 1</p> <p>10 to 20 %: 2</p> <p>20 to 30 %: 3</p> <p>30 to 40 %: 4</p> <p>40 to 50%: 5</p> <p>50 to 60%: 6</p> <p>60 to 70%: 7</p> <p>70 to 80 %: 8</p> <p>80 to 90%: 9</p> <p>90 to 100%: 10</p>
Memory usage	Memory consumed by collector as a percentage of maximum memory resource assigned to it.	<p>0 to 10 % : 1</p> <p>10 to 20 %: 2</p> <p>20 to 30 %: 3</p> <p>30 to 40 %: 4</p> <p>40 to 50%: 5</p> <p>50 to 60%: 6</p> <p>60 to 70%: 7</p> <p>70 to 80 %: 8</p> <p>80 to 90%: 9</p> <p>90 to 100%: 10</p>
Number of skipped cadence	<p>Number of jobs with a skipped cadence.</p> <p>Note: This metric applies only to poll-based (SNMP/CLI /NETCONF) collectors.</p>	<p>< 50: 3</p> <p>> 50: 7</p> <p>> 80: 10</p>
Number of control plane status updates	Number of control plane status updates sent by the collector.	<p>< 100: 1</p> <p>> 100: 8</p>
Number of control plane jobs received	Number of control plane jobs received from Magellan	<p>< 100: 2</p> <p>> 100: 8</p>

Table 3. MetricWeight for Poll-based collectors (CLI,SNMP and NETCONF)

Metric	MetricWeight
Skipped cadence	10

Memory Usage	40
CPU Usage	40
Control Plane status updates	5
Control plane jobs received	5

Table 4. MetricWeight for Event Based collectors(MDT , GNMI and SYSLOG):

Metric	MetricWeight
Memory Usage	50
CPU Usage	40
Control Plane status updates	5
Control plane Jobs Received	5