



# Health Rule Conditions

**Related Pages:**

- [Health Rules](#)
- [Configure Health Rules](#)
- [Troubleshoot Health Rule Violations](#)
- [Policies](#)
- [Actions](#)



✱

This page provides an overview of health rule conditions in AppDynamics.

You define the acceptable range for a metric by establishing health rule conditions. A health rule condition sets the metric levels that constitute a **Warning** status or a **Critical** status.

A condition consists of a Boolean statement that compares the current value of a metric against one or more static or dynamic thresholds based on a selected baseline. If the condition is true, the health rule violates. The rules for evaluating a condition using multiple thresholds depend on configuration.

Static thresholds are straightforward. For example, is a business transaction's average response time greater than 200 ms? The condition is evaluated to 'true' if the average response time is greater than 200 ms and the health rule violates.

Dynamic thresholds are based on a percentage in relation to, or a standard deviation from, a baseline built on a rolled-up baseline trend pattern. A daily trend baseline rolls up values for a particular hour of the day during the last thirty days, whereas a weekly trend baseline rolls up values for a particular hour of the day, for a particular day of the week, for the last 90 days. For more information about baselines, see [Dynamic Baselines](#).

You can define a threshold for a health rule based on a single metric value or on a mathematical expression built from multiple metric values.

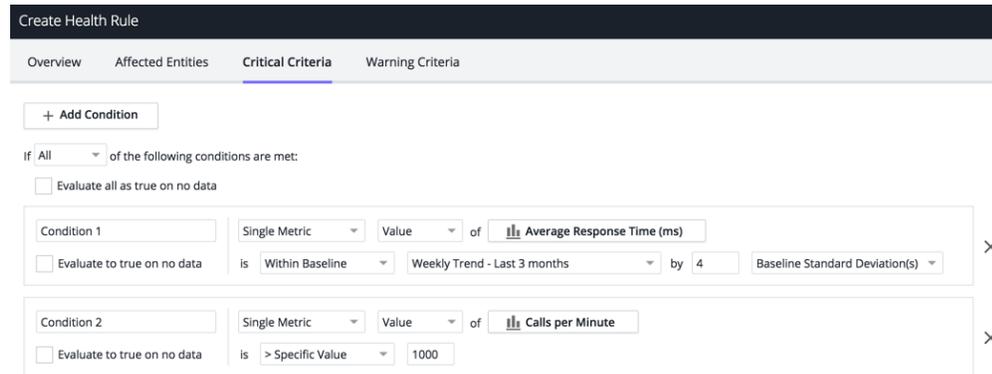
The following are typical health rule conditions:

- If the value of the Average Response Time is greater than the default baseline by 3 X the Baseline Standard Deviation . . .
- If the count of the Errors Per Minute is greater than 1000 . . .
- If the number of MB of Free Memory is less than 2 X the Default Baseline . . .
- If the value of **Errors per Minute/Calls per Minute** over the last 15 days > 0.2 . . .  
This example combines two metrics in a single condition. You can use the [expression builder](#) embedded in the health rules wizard to create conditions based on a complex expression comprising multiple interdependent metrics.
- If the (average response time > baseline OR errors per minute > baseline) **AND** (calls per minute > the defined threshold) . . .  
This example uses multiple conditions to evaluate the health rules. You can use the 'CUSTOM' option to define a boolean expression to evaluate the conditions.

# Critical and Warning Conditions

Conditions are classified as either critical or warning conditions.

Critical conditions are evaluated before warning conditions. If you have defined a critical condition and a warning condition in the same health rule, the warning condition is evaluated only if the critical condition is not true.



The configuration procedures for critical and warning conditions are identical, but you configure these two types of conditions in separate panels. You can copy a critical condition configuration to a warning configuration and vice-versa and then adjust the metrics in the copy to differentiate them. For example, in the **Critical Condition** panel you can create a critical condition based on the rule:

- If the Average Response Time is greater than 1000

Then from the Warning Condition panel, copy that condition and edit it to be:

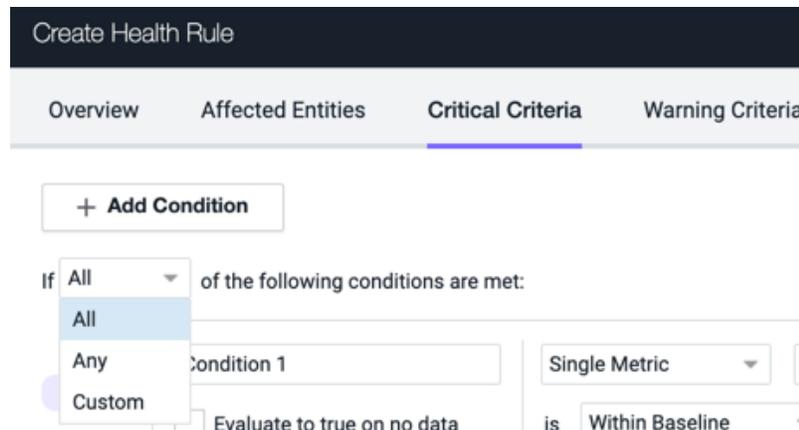
- If the Average Response Time is greater than 500

As performance changes, a health rule violation can be upgraded from warning to critical if performance deteriorates to the higher threshold or downgraded from critical to warning if performance improves to the warning threshold.

## Evaluation Criteria

When you define multiple conditions for a health rule, they are evaluated based on the criteria you define. You can use the following options to define the evaluation criteria:

- All: the health rule violates if all the conditions defined in the criteria evaluate to 'true'.
- Any: the health rule violates if one of the conditions defined in the criteria evaluates to 'true'.
- Custom: the health rule violates if the boolean expression with multiple conditions evaluates to 'true'.



For information on how to configure evaluation criteria, see [Configure Health Rule Evaluation Criteria](#).

The following table uses examples to illustrate how a health rule is evaluated based on the criteria and when is it considered to violate.

Health Rule Configuration	Evaluation	Example
Single condition	the condition evaluates to 'true'	A health rule that compares 'average response time' with a defined baseline.
Multiple conditions with 'ANY' evaluation criteria	one of the health rule conditions evaluates to 'true'	A health rule that monitors the health of business transaction may measure any of the following performance metrics: <ul style="list-style-type: none"> <li>• average response time or</li> <li>• errors per min</li> </ul>
Multiple conditions with 'ALL' evaluation criteria	all of the health rule conditions evaluate to 'true'	A health rule that monitors the health of business transaction measures all of the following metrics: <ul style="list-style-type: none"> <li>• response time</li> <li>• average response time greater than a baseline value, correlated with the application load</li> </ul> <p>For example, 50 concurrent users on the system. A policy is defined such that a remedial action is initiated only if the load (calls per minute) is high although the response time threshold is reached.</p> <p>The first part of the condition evaluates the response time and the second part ensures that the health rule is violated only when there is sufficient load.</p>
Multiple conditions with 'CUSTOM' evaluation criteria	the boolean expression with multiple conditions evaluates to 'true'	A health rule that monitors the health of a Business Transaction, measures the performance based on the following conditions: <ul style="list-style-type: none"> <li>• (average response time greater than baseline OR errors per min greater than baseline)</li> </ul> <p>AND</p> <ul style="list-style-type: none"> <li>• (calls per min greater than threshold)</li> </ul>

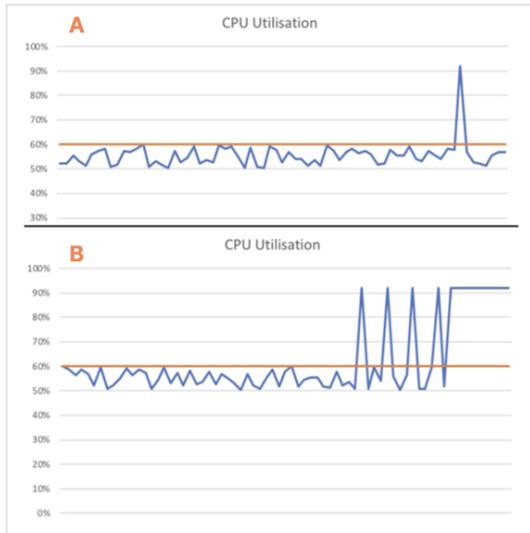
## Persistence Thresholds

Temporary spikes in metric performance data is a major cause of false alerts. Persistence thresholds allow you to define a 'sensitivity level' for a health rule and thereby reduce the number of false alerts. You can define the 'number of times metric performance data should exceed the defined threshold during the evaluation time frame' to constitute a violation and subsequently trigger an alert.



You can define a persistence threshold for a condition only if you have defined an [evaluation time frame](#) of 30 minutes or less.

For example, when monitoring the CPU utilization, you would not want to be reported of a single violation (section A in [figure](#)) of the threshold. However, if the violation of threshold continues to occur multiple times (section B in [figure](#)) during the evaluation time period, you would want to be alerted.



(Single spike)  
No Alert

(Multiple spikes)  
Alert

## Alert Sensitivity Tuning



- AST is available to SaaS customers only.
- AST feature is enabled only if you create a health rule to monitor a business transaction, a service endpoint, or a remote service.

It is important that you configure conditions appropriately to ensure that you do not miss any alerts or receive false alerts instead. With 'Alert Sensitivity Tuning' (AST), you can view historical data for metrics and baselines when you configure conditions. This data helps visualize the impact of the configuration you define and assists in fine-tuning the configuration.

You can view a graphical representation of the metric data, threshold value, standard deviation, and the baseline. The graphical view is instantly updated when you update any configuration. You can also view granular details by modifying the graphical view. To view granular details, you can:

- increase the time period of the data capture to 1 day or 3 days.



If data is not available for the selected time period, only available data is presented.

- adjust the time range in the graph to view the metric data details.
- hover over the metric data to view metric and baseline values at any given time.

You can analyze the data presented and then make adjustments to your configuration accordingly. For more information on fine-tuning a condition, see [Create a BT Health Rule](#) and [Fine-tune Metric Evaluation](#).

For example, if you select **Average CPU Used (ms)** (1) as a metric to be monitored for a health rule condition, you can view the past metric data for a time period of 8 hours (2). The graphical representation of the metric data indicates the baseline (3) and the baseline standard deviation (4). Based on this data, you can fine-tune the evaluation of the condition so that you avoid false alerts and receive the alerts only when the health rule violates the conditions you define.

Health Rule will violate if the conditions above evaluate to true for:

The BT Average (average for all Nodes in the BT)

Any Node

Cancel Save

## Moving Average

If you define a persistence threshold to evaluate a condition, the metric data for every minute is compared to the baseline and plotted on the AST graph. However, if you do not define the persistence threshold, a 'moving average' for the selected metric is plotted as follows:

1. Depending on the 'Use data from last' value 'X', the metric data for 'X' minutes is considered.
2. On the 'X+1'th minute, the average of the past 'X' minutes is computed and denoted as the first point on the graph.
3. Similarly, the average for the following 'X' minutes is computed and points are denoted on the graph for the rest of the time range.

## Why Use Moving Average?

Unless persistence thresholds are used, health rules compare the moving average of a metric to a threshold or a baseline. Thus, representing the moving average in the graph is appropriate. For more information, see [Create a Health Rule and Fine-tune Metric Evaluation](#).

Health Rule will violate if the conditions above evaluate to true for:

The BT Average (average for all Nodes in the BT)

Any Node

Cancel Save

## How are Conditions Evaluated if No Data is Reported?

The Evaluate to true on no data option controls the evaluation of the condition in cases where any metric on which the condition is based, does not return any data. The condition evaluates to 'unknown' (default) when no data is returned. If the health rule is based on all the conditions evaluating to true, having no data returned may affect whether the health rule triggers an action.

When you define a health rule evaluation time frame, reference data is collected for each data point. If the configured metric fails to report data during the time frame, the health rule condition is evaluated as follows:

Evaluate to true on no data	Trigger only when violation occurs x times in the last ymin (s)	Condition Evaluation
Enabled	Enabled	The condition is evaluated for each data point in the evaluation time frame. The condition evaluates to 'true' when metric fails to report any data for a given data point.  For example, when you set the persistence threshold, X = 3 for an evaluation time frame, Y = 5. This means that 5 data points are required to evaluate the condition. Data is reported for 4 data points, no data is reported for 1 data point and the metric exceeds the threshold twice. The condition evaluates to 'true' for the minute when no data is reported.
Enabled	Disabled	The condition evaluates to 'true' if the configured metric fails to report data for any data point during the evaluation time frame.
Disabled	Disabled	The condition does not evaluate to 'true' if the configured metric fails to report data for any data point during the evaluation time frame.

## Custom Boolean Expression

A condition consists of single or multiple statements that evaluate different metrics. You can define a single condition or multiple conditions to evaluate the performance metrics of your application. When you define multiple conditions, you may want to define an evaluation criteria using a boolean expression.

Advantages of using a boolean expression are:

- eliminates the need to create multiple health rules to monitor various performance metrics. Using a boolean expression allows you to evaluate complex criteria for multiple conditions in one go.
- well-calibrated boolean expression ensures reduced false alerts.
- easy to create and maintain health rules with complex evaluation criteria using simple condition names. Conditions are named as **A**, **B**, **C** and so on.
- allows the use of **AND** and **OR** operators to define a highly complex boolean expression. You can use a maximum of 8 operators in your boolean expression.

## Evaluation Scope

The health rule evaluation scope defines how many nodes in the affected entities must violate the condition before the health rule is considered violated.

Evaluation scope applies only to business transaction performance type health rules and node health type health rules in which the affected entities are defined at the tier level.

For example, you may have a critical condition in which the condition is unacceptable for any node, or you may want to consider the condition a violation only if the condition is true for 50% or more of the nodes in a tier.

Options for this evaluation scope are:

- The tier average: Evaluation is performed on the tier average instead of the individual nodes.
- Any node: If any node exceeds the thresholds, the rule is violated.
- Percentage of the nodes: If x% of the nodes exceed the thresholds, the rule is violated.
- Number of nodes: If x nodes exceed the thresholds, the rule is violated.