



We offer free update service for one year!
<https://www.certqueen.com>

Exam : SPLK-4001

**Title : Splunk O11y Cloud Certified
Metrics User Exam**

Version : DEMO

1.What are the best practices for creating detectors? (select all that apply)

- A. View data at highest resolution.
- B. Have a consistent value.
- C. View detector in a chart.
- D. Have a consistent type of measurement.

Answer: ABCD

Explanation:

The best practices for creating detectors are:

View data at highest resolution. This helps to avoid missing important signals or patterns in the data that could indicate anomalies or issues¹

Have a consistent value. This means that the metric or dimension used for detection should have a clear and stable meaning across different sources, contexts, and time periods. For example, avoid using metrics that are affected by changes in configuration, sampling, or aggregation²

View detector in a chart. This helps to visualize the data and the detector logic, as well as to identify any false positives or negatives. It also allows to adjust the detector parameters and thresholds based on the data distribution and behavior³

Have a consistent type of measurement. This means that the metric or dimension used for detection should have the same unit and scale across different sources, contexts, and time periods. For example, avoid mixing bytes and bits, or seconds and milliseconds.

1: <https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors> 2:

<https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors> 3:

<https://docs.splunk.com/Observability/gdi/metrics/detectors.html#View-detector-in-a-chart> :

<https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors>

2.An SRE came across an existing detector that is a good starting point for a detector they want to create. They clone the detector, update the metric, and add multiple new signals.

As a result of the cloned detector, which of the following is true?

- A. The new signals will be reflected in the original detector.
- B. The new signals will be reflected in the original chart.
- C. You can only monitor one of the new signals.
- D. The new signals will not be added to the original detector.

Answer: D

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document¹, cloning a detector creates a copy of the detector that you can modify without affecting the original detector. You can change the metric, filter, and signal settings of the cloned detector. However, the new signals that you add to the cloned detector will not be reflected in the original detector, nor in the original chart that the detector was based on. Therefore, option D is correct.

Option A is incorrect because the new signals will not be reflected in the original detector. Option B is incorrect because the new signals will not be reflected in the original chart. Option C is incorrect because you can monitor all of the new signals that you add to the cloned detector.

3.Which of the following are supported rollup functions in Splunk Observability Cloud?

- A. average, latest, lag, min, max, sum, rate

- B. std_dev, mean, median, mode, min, max
- C. sigma, epsilon, pi, omega, beta, tau
- D. 1min, 5min, 10min, 15min, 30min

Answer: A

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document¹, Observability Cloud has the following rollup functions: Sum: (default for counter metrics): Returns the sum of all data points in the MTS reporting interval. Average (default for gauge metrics): Returns the average value of all data points in the MTS reporting interval. Min: Returns the minimum data point value seen in the MTS reporting interval. Max: Returns the maximum data point value seen in the MTS reporting interval. Latest: Returns the most recent data point value seen in the MTS reporting interval. Lag: Returns the difference between the most recent and the previous data point values seen in the MTS reporting interval. Rate: Returns the rate of change of data points in the MTS reporting interval. Therefore, option A is correct.

4.A Software Engineer is troubleshooting an issue with memory utilization in their application. They released a new canary version to production and now want to determine if the average memory usage is lower for requests with the 'canary' version dimension. They've already opened the graph of memory utilization for their service.

How does the engineer see if the new release lowered average memory utilization?

- A. On the chart for plot A, select Add Analytics, then select MeanTransformation. In the window that appears, select 'version' from the Group By field.
- B. On the chart for plot A, scroll to the end and click Enter Function, then enter 'A/B-I'.
- C. On the chart for plot A, select Add Analytics, then select Mean:Aggregation. In the window that appears, select 'version' from the Group By field.
- D. On the chart for plot A, click the Compare Means button. In the window that appears, type 'version1'.

Answer: C

Explanation:

The correct answer is C. On the chart for plot A, select Add Analytics, then select Mean: Aggregation. In the window that appears, select 'version' from the Group By field.

This will create a new plot B that shows the average memory utilization for each version of the application. The engineer can then compare the values of plot B for the 'canary' and 'stable' versions to see if there is a significant difference.

To learn more about how to use analytics functions in Splunk Observability Cloud, you can refer to this documentation¹.

1: <https://docs.splunk.com/Observability/gdi/metrics/analytics.html>

5.One server in a customer's data center is regularly restarting due to power supply issues.

What type of dashboard could be used to view charts and create detectors for this server?

- A. Single-instance dashboard
- B. Machine dashboard
- C. Multiple-service dashboard
- D. Server dashboard

Answer: A

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document¹, a single-instance dashboard is a type of dashboard that displays charts and information for a single instance of a service or host. You can use a single-instance dashboard to monitor the performance and health of a specific server, such as the one that is restarting due to power supply issues. You can also create detectors for the metrics that are relevant to the server, such as CPU usage, memory usage, disk usage, and uptime. Therefore, option A is correct.