SPLK-4001$^{\text{Q\&As}}$

Splunk O11y Cloud Certified Metrics User

QUESTION 1

Which of the following is optional, but highly recommended to include in a datapoint?

A. Metric name

B. Timestamp

C. Value

D. Metric type

Correct Answer: D

The correct answer is D. Metric type. A metric type is an optional, but highly recommended field that specifies the kind of measurement that a datapoint represents. For example, a metric type can be gauge, counter, cumulative counter, or histogram. A metric type helps Splunk Observability Cloud to interpret and display the data correctly To learn more about how to send metrics to Splunk Observability Cloud, you can refer to this documentation.
https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Metric-types
https://docs.splunk.com/Observability/gdi/metrics/metrics.html

QUESTION 2

A customer operates a caching web proxy. They want to calculate the cache hit rate for their service. What is the best way to achieve this?

A. Percentages and ratios

B. Timeshift and Bottom N

C. Timeshift and Top N

D. Chart Options and metadata

Correct Answer: A

According to the Splunk O11y Cloud Certified Metrics User Track document, percentages and ratios are useful for calculating the proportion of one metric to another, such as cache hits to cache misses, or successful requests to failed

requests. You can use the percentage() or ratio() functions in SignalFlow to compute these values and display them in charts. For example, to calculate the cache hit rate for a service, you can use the following SignalFlow code:

percentage(counters("cache.hits"), counters("cache.misses")) This will return the percentage of cache hits out of the total number of cache attempts. You can also use the ratio() function to get the same result, but as a decimal value instead

of a percentage.

ratio(counters("cache.hits"), counters("cache.misses"))

QUESTION 3

A customer has a large population of servers. They want to identify the servers where utilization has increased the most since last week. Which analytics function is needed to achieve this?

A. Rate

B. Sum transformation

C. Tlmeshift

D. Standard deviation

Correct Answer: C

The correct answer is C. Timeshift.

According to the Splunk Observability Cloud documentation1, timeshift is an analytic function that allows you to compare the current value of a metric with its value at a previous time interval, such as an hour ago or a week ago. You can use

the timeshift function to measure the change in a metric over time and identify trends, anomalies, or patterns. For example, to identify the servers where utilization has increased the most since last week, you can use the following SignalFlow

code:

timeshift(1w, counters("server.utilization"))

This will return the value of the server.utilization counter metric for each server one week ago. You can then subtract this value from the current value of the same metric to get the difference in utilization. You can also use a chart to visualize

the results and sort them by the highest difference in utilization.

QUESTION 4

Which of the following can be configured when subscribing to a built-in detector?

A. Alerts on team landing page.

B. Alerts on a dashboard.

C. Outbound notifications.

D. Links to a chart.

Correct Answer: C

According to the web search results1, subscribing to a built-in detector is a way to receive alerts and notifications from Splunk Observability Cloud when certain criteria are met. A built-in detector is a detector that is automatically created and

configured by Splunk Observability Cloud based on the data from your integrations, such as AWS, Kubernetes, or OpenTelemetry1. To subscribe to a built-in detector, you need to do the following steps:

Find the built-in detector that you want to subscribe to. You can use the metric finder or the dashboard groups to locate the built-in detectors that are relevant to your data sources.

Hover over the built-in detector and click the Subscribe button. This will open a dialog box where you can configure your subscription settings1. Choose an outbound notification channel from the drop-down menu. This is where you can

specify how you want to receive the alert notifications from the built-in detector. You can choose from various channels, such as email, Slack, PagerDuty, webhook, and so on. You can also create a new notification channel by clicking the +

icon.

Enter the notification details for the selected channel. This may include your email address, Slack channel name, PagerDuty service key, webhook URL, and so on. You can also customize the notification message with variables and

markdown formatting.

Click Save. This will subscribe you to the built-in detector and send you alert notifications through the chosen channel when the detector triggers or clears an alert.

Therefore, option C is correct.

QUESTION 5

For a high-resolution metric, what is the highest possible native resolution of the metric?

A. 2 seconds

B. 15 seconds

C. 1 second

D. 5 seconds

Correct Answer: C

The correct answer is C. 1 second.

According to the Splunk Test Blueprint - O11y Cloud Metrics User document1, one of the metrics concepts that is covered in the exam is data resolution and rollups. Data resolution refers to the granularity of the metric data points, and rollups are the process of aggregating data points over time to reduce the amount of data stored. The Splunk O11y Cloud Certified Metrics User Track document2 states that one of the recommended courses for preparing for the exam is Introduction to Splunk Infrastructure Monitoring, which covers the basics of metrics monitoring and visualization. In the Introduction to Splunk Infrastructure Monitoring course, there is a section on Data Resolution and Rollups, which explains that Splunk Observability Cloud collects high- resolution metrics at 1-second intervals by default, and then applies rollups to reduce the data volume over time. The document also provides a table that shows the different rollup intervals and retention periods for different resolutions. Therefore, based on these documents, we can conclude that for a high-resolution metric, the highest possible native resolution of the metric is 1 second.