Vendor: Splunk

Exam Code:SPLK-4001

Exam Name: Splunk O11y Cloud Certified Metrics User

Version: Demo

QUESTION 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.

Correct Answer: ABCD

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 issues 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 aggregation 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 behavior 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. https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best- practices-for-detectors 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

---

QUESTION 2

A customer deals with a holiday rush of traffic during November each year, but does not want to be flooded with alerts when this happens. The increase in traffic is expected and consistent each year. Which detector condition should be used when creating a detector for this data?

A. Outlier Detection

B. Static Threshold

C. Calendar Window

D. Historical Anomaly

Correct Answer: D

historical anomaly is a detector condition that allows you to trigger an alert when a signal deviates from its historical pattern. Historical anomaly uses machine learning to learn the normal behavior of a signal based on its past data, and then compares the current value of the signal with the expected value based on the learned pattern. You can use historical anomaly to detect unusual changes in a signal that are not explained by seasonality, trends, or cycles. Historical anomaly is suitable for creating a detector for the customer\\'s data, because it can account for the expected and consistent increase in traffic during November each year. Historical anomaly can learn that the traffic pattern has a seasonal component that peaks in November, and then adjust the expected value of the traffic accordingly. This way, historical anomaly can avoid triggering alerts when the traffic increases in November, as this is not an anomaly, but rather a normal variation. However, historical anomaly can still trigger alerts when the traffic deviates from the historical pattern in other ways, such as if it drops significantly or spikes unexpectedly.

QUESTION 3

To refine a search for a metric a customer types host: test-*. What does this filter return?

A. Only metrics with a dimension of host and a value beginning with test-.

B. Error

C. Every metric except those with a dimension of host and a value equal to test.

D. Only metrics with a value of test- beginning with host.

Correct Answer: A

The correct answer is A. Only metrics with a dimension of host and a value beginning with test-. This filter returns the metrics that have a host dimension that matches the pattern test-. For example, test-01, test-abc, test-xyz, etc. The asterisk () is a wildcard character that can match any string of characters To learn more about how to filter metrics in Splunk Observability Cloud, you can refer to this documentation.
https://docs.splunk.com/Observability/gdi/metrics/search.html#Filter-metrics
https://docs.splunk.com/Observability/gdi/metrics/search.html

---

QUESTION 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 MeanrTransformation. 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-l\\'.

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.

Correct Answer: C

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

---

QUESTION 5

When creating a standalone detector, individual rules in it are labeled according to severity. Which of the choices below represents the possible severity levels that can be selected?

A. Info, Warning, Minor, Major, and Emergency.

B. Debug, Warning, Minor, Major, and Critical.

C. Info, Warning, Minor, Major, and Critical.

D. Info, Warning, Minor, Severe, and Critical.

Correct Answer: C

The correct answer is C. Info, Warning, Minor, Major, and Critical. When creating a standalone detector, you can define one or more rules that specify the alert conditions and the severity level for each rule. The severity level indicates how urgent or important the alert is, and it can also affect the notification settings and the escalation policy for the alert Splunk Observability Cloud provides five predefined severity levels that you can choose from when creating a rule: Info, Warning, Minor, Major, and Critical. Each severity level has a different color and icon to help you identify the alert status at a glance. You can also customize the severity levels by changing their names, colors, or icons2 To learn more about how to create standalone detectors and use severity levels in Splunk Observability Cloud, you can refer to these documentations. https://docs.splunk.com/Observability/alerts-detectors- notifications/detectors.html#Create-a-standalone-detector https://docs.splunk.com/Observability/alerts-detectors-notifications/detector- options.html#Severity-levels

---

QUESTION 6

What constitutes a single metrics time series (MTS)?

A. A series of timestamps that all reflect the same metric.

B. A set of data points that all have the same metric name and list of dimensions.

C. A set of data points that use different dimensions but the same metric name.

D. A set of metrics that are ordered in series based on timestamp.

Correct Answer: B

The correct answer is B. A set of data points that all have the same metric name and list of dimensions.

A metric time series (MTS) is a collection of data points that have the same metric and the same set of dimensions. For example, the following sets of data points are in three separate MTS:

MTS: Gauge metric cpu.utilization, dimension "hostname": "host" MTS: Gauge metric cpu.utilization, dimension "hostname": "host" MTS: Gauge metric memory.usage, dimension "hostname": "host"

A metric is a numerical measurement that varies over time, such as CPU utilization or memory usage. A dimension is a key-value pair that provides additional information about the metric, such as the hostname or the location. A data point is

a combination of a metric, a dimension, a value, and a timestamp

---

QUESTION 7

Clicking a metric name from the results in metric finder displays the metric in Chart Builder. What action needs to be taken in order to save the chart created in the UI?

A. Create a new dashboard and save the chart.

B. Save the chart to multiple dashboards.

C. Make sure that data is coming in for the metric then save the chart.

D. Save the chart to a dashboard.

Correct Answer: D

According to the web search results, clicking a metric name from the results in metric finder displays the metric in Chart Builder1. Chart Builder is a tool that allows you to create and customize charts using metrics, dimensions, and analytics

functions2. To save the chart created in the UI, you need to do the following steps:

Click the Save button on the top right corner of the Chart Builder. This will open a dialog box where you can enter the chart name and description, and choose the dashboard where you want to save the chart.

Enter a name and a description for your chart. The name should be descriptive and unique, and the description should explain the purpose and meaning of the chart.

Choose an existing dashboard from the drop-down menu, or create a new dashboard by clicking the + icon. A dashboard is a collection of charts that display metrics and events for your services or hosts. You can organize and share

dashboards with other users in your organization using dashboard groups. Click Save. This will save your chart to the selected dashboard and redirect you to the dashboard view. You can also access your saved chart from the Dashboards

menu on the left navigation bar.

QUESTION 8

For which types of charts can individual plot visualization be set?

A. Line, Bar, Column

B. Bar, Area, Column

C. Line, Area, Column

D. Histogram, Line, Column

Correct Answer: C

The correct answer is C. Line, Area, Column. For line, area, and column charts, you can set the individual plot visualization to change the appearance of each plot in the chart. For example, you can change the color, shape, size, or style of the lines, areas, or columns. You can also change the rollup function, data resolution, or y-axis scale for each plot To set the individual plot visualization for line, area, and column charts, you need to select the chart from the Metric Finder, then click on Plot Chart Options and choose Individual Plot Visualization from the list of options. You can then

customize each plot according to your preferences To learn more about how to use individual plot visualization in Splunk Observability Cloud, you can refer to this documentation.
https://docs.splunk.com/Observability/gdi/metrics/charts.html#Individual-plot-visualization
https://docs.splunk.com/Observability/gdi/metrics/charts.html#Set-individual-plot- visualization

---

QUESTION 9

A customer wants to share a collection of charts with their entire SRE organization. What feature of Splunk Observability Cloud makes this possible?

A. Dashboard groups

B. Shared charts

C. Public dashboards

D. Chart exporter

Correct Answer: A

According to the web search results, dashboard groups are a feature of Splunk Observability Cloud that allows you to organize and share dashboards with other users in your organization1. You can create dashboard groups based on different criteria, such as service, team, role, or topic. You can also set permissions for each dashboard group, such as who can view, edit, or manage the dashboards in the group. Dashboard groups make it possible to share a collection of charts with your entire SRE organization, or any other group of users that you want to collaborate with.

QUESTION 10

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

Correct Answer: A

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.

QUESTION 11

What is one reason a user of Splunk Observability Cloud would want to subscribe to an alert?

A. To determine the root cause of the Issue triggering the detector.

B. To perform transformations on the data used by the detector.

C. To receive an email notification when a detector is triggered.

D. To be able to modify the alert parameters.

Correct Answer: C

One reason a user of Splunk Observability Cloud would want to subscribe to an alert is C. To receive an email notification when a detector is triggered. A detector is a component of Splunk Observability Cloud that monitors metrics or events and triggers alerts when certain conditions are met. A user can create and configure detectors to suit their monitoring needs and goals A subscription is a way for a user to receive notifications when a detector triggers an alert. A user can subscribe to a detector by entering their email address in the Subscription tab of the detector page. A user can also unsubscribe from a detector at any time When a user subscribes to an alert, they will receive an email notification that contains information about the alert, such as the detector name, the alert status, the alert severity, the alert time, and the alert message. The email notification also includes links to view the detector, acknowledge the alert, or unsubscribe from the detector To learn more about how to use detectors and subscriptions in Splunk Observability Cloud, you can refer to these documentations. https://docs.splunk.com/Observability/alerts-detectors-notifications/detectors.html https://docs.splunk.com/Observability/alerts-detectors-notifications/subscribe-to-detectors.html

---

QUESTION 12

Which of the following rollups will display the time delta between a datapoint being sent and a datapoint being received?

A. Jitter

B. Delay

C. Lag

D. Latency

Correct Answer: C

According to the Splunk Observability Cloud documentation1, lag is a rollup function that returns the difference between the most recent and the previous data point values seen in the metric time series reporting interval. This can be used to measure the time delta between a data point being sent and a data point being received, as long as the data points have timestamps that reflect their send and receive times. For example, if a data point is sent at 10:00:00 and received at 10:00:05, the lag value for that data point is 5 seconds.