Free Questions for SPLK-4001 by go4braindumps

Shared by Richard on 04-10-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which of the following statements are true about local data links? (select all that apply)

Options:

A- Anyone with write permission for a dashboard can add local data links that appear on that dashboard.

B- Local data links can only have a Splunk Observability Cloud internal destination.

C- Only Splunk Observability Cloud administrators can create local links.

D- Local data links are available on only one dashboard.

Answer:

A, D

Explanation:

The correct answers are A and D.

According to the Get started with Splunk Observability Cloud document1, one of the topics that is covered in the Getting Data into Splunk Observability Cloud course is global and local data links. Data links are shortcuts that provide convenient access to related resources, such as Splunk Observability Cloud dashboards, Splunk Cloud Platform and Splunk Enterprise, custom URLs, and Kibana logs.

The document explains that there are two types of data links: global and local. Global data links are available on all dashboards and charts, while local data links are available on only one dashboard. The document also provides the following information about local data links:

Anyone with write permission for a dashboard can add local data links that appear on that dashboard.

Local data links can have either a Splunk Observability Cloud internal destination or an external destination, such as a custom URL or a Kibana log.

Only Splunk Observability Cloud administrators can delete local data links.

Therefore, based on this document, we can conclude that A and D are true statements about local data links. B and C are false statements because:

B is false because local data links can have an external destination as well as an internal one.

C is false because anyone with write permission for a dashboard can create local data links, not just administrators.

Question 2

Question Type: MultipleChoice

Which of the following are required in the configuration of a data point? (select all that apply)

Options:

A- Metric Name

B- Metric Type

C- Timestamp

D- Value

Answer:

A, C, D

Explanation:

The required components in the configuration of a data point are:

Metric Name: A metric name is a string that identifies the type of measurement that the data point represents, such as cpu.utilization, memory.usage, or response.time. A metric name is mandatory for every data point, and it must be unique within a Splunk Observability Cloud organization1

Timestamp: A timestamp is a numerical value that indicates the time at which the data point was collected or generated. A timestamp is mandatory for every data point, and it must be in epoch time format, which is the number of seconds since January 1, 1970 UTC1

Value: A value is a numerical value that indicates the magnitude or quantity of the measurement that the data point represents. A value is mandatory for every data point, and it must be compatible with the metric type of the data point1

Therefore, the correct answer is A, C, and D.

To learn more about how to configure data points in Splunk Observability Cloud, you can refer to this documentation1.

1: https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Data-points

Question 3

Question Type: MultipleChoice

The Sum Aggregation option for analytic functions does which of the following?

Options:

A- Calculates the number of MTS present in the plot.

B- Calculates 1/2 of the values present in the input time series.

C- Calculates the sum of values present in the input time series across the entire environment or per group.

D- Calculates the sum of values per time series across a period of time.

Answer:

C

Explanation:

According to the Splunk Test Blueprint - O11y Cloud Metrics User document1, one of the metrics concepts that is covered in the exam is analytic functions. Analytic functions are mathematical operations that can be applied to metrics to transform, aggregate, or analyze them.

The Splunk O11y Cloud Certified Metrics User Track document2states 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 Analytic Functions, which explains that analytic functions can be used to perform calculations on metrics, such as sum, average, min, max, count, etc. The document also provides examples of how to use analytic functions in charts and dashboards.

One of the analytic functions that can be used is Sum Aggregation, which calculates the sum of values present in the input time series across the entire environment or per group. The document gives an example of how to use Sum Aggregation to calculate the total CPU usage across all hosts in a group by using the following syntax:

sum(cpu.utilization) by hostgroup

Question 4

Question Type: MultipleChoice

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?

Options:

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.

Answer:

C

Explanation:

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 alert1

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 documentations12.

1: https://docs.splunk.com/Observability/alerts-detectors-notifications/detectors.html#Create-a-standalone-detector 2: https://docs.splunk.com/Observability/alerts-detectors-notifications/detector-options.html#Severity-levels

Question 5

Question Type: **MultipleChoice**

Which of the following are accurate reasons to clone a detector? (select all that apply)

Options:

A- To modify the rules without affecting the existing detector.

B- To reduce the amount of billed TAPM for the detector.

C- To add an additional recipient to the detector's alerts.

D- To explore how a detector was created without risk of changing it.

Answer:

A, D

Explanation:

The correct answers are A and D.

According to the Splunk Test Blueprint - O11y Cloud Metrics User document1, one of the alerting concepts that is covered in the exam is detectors and alerts. Detectors are the objects that define the conditions for generating alerts, and alerts are the notifications that are sent when those conditions are met.

The Splunk O11y Cloud Certified Metrics User Track document2 states that one of the recommended courses for preparing for the exam is Alerting with Detectors, which covers how to create, modify, and manage detectors and alerts.

In the Alerting with Detectors course, there is a section on Cloning Detectors, which explains that cloning a detector creates a copy of the detector with all its settings, rules, and alert recipients. The document also provides some reasons why you might want to clone a

detector, such as:

To modify the rules without affecting the existing detector. This can be useful if you want to test different thresholds or conditions before applying them to the original detector.

To explore how a detector was created without risk of changing it. This can be helpful if you want to learn from an existing detector or use it as a template for creating a new one.

Therefore, based on these documents, we can conclude that A and D are accurate reasons to clone a detector. B and C are not valid reasons because:

Cloning a detector does not reduce the amount of billed TAPM for the detector. TAPM stands for Tracked Active Problem Metric, which is a metric that has been alerted on by a detector. Cloning a detector does not change the number of TAPM that are generated by the original detector or the clone.

Cloning a detector does not add an additional recipient to the detector's alerts. Cloning a detector copies the alert recipients from the original detector, but it does not add any new ones. To add an additional recipient to a detector's alerts, you need to edit the alert settings of the detector.

Question 6

Question Type: **MultipleChoice**

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

Options:

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.

Answer:

C

Explanation:

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 goals1

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 time2

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 detector2

To learn more about how to use detectors and subscriptions in Splunk Observability Cloud, you can refer to these documentations12.

1: https://docs.splunk.com/Observability/alerts-detectors-notifications/detectors.html 2: https://docs.splunk.com/Observability/alerts-detectors-notifications/subscribe-to-detectors.html

Question 7

Question Type: MultipleChoice

Which component of the OpenTelemetry Collector allows for the modification of metadata?

Options:

A- Processors

B- Pipelines

C- Exporters

D- Receivers

Answer:

A

Explanation:

The component of the OpenTelemetry Collector that allows for the modification of metadata is A. Processors.

Processors are components that can modify the telemetry data before sending it to exporters or other components. Processors can perform various transformations on metrics, traces, and logs, such as filtering, adding, deleting, or updating attributes, labels, or resources. Processors can also enrich the telemetry data with additional metadata from various sources, such as Kubernetes, environment variables, or system information1

For example, one of the processors that can modify metadata is the attributes processor. This processor can update, insert, delete, or replace existing attributes on metrics or traces. Attributes are key-value pairs that provide additional information about the telemetry data, such as the service name, the host name, or the span kind2

Another example is the resource processor. This processor can modify resource attributes on metrics or traces. Resource attributes are key-value pairs that describe the entity that produced the telemetry data, such as the cloud provider, the region, or the instance type3

To learn more about how to use processors in the OpenTelemetry Collector, you can refer to this documentation1.

1: https://opentelemetry.io/docs/collector/configuration/#processors 2: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/processor/attributesprocessor 3: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/processor/resourceprocessor

Question 8

Question Type: MultipleChoice

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

Options:

A- 2 seconds

B- 15 seconds

C- 1 second

D- 5 seconds

Answer:

C

Explanation:

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.

Question 9

Question Type: **MultipleChoice**

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?

Options:

A- Outlier Detection

B- Static Threshold

C- Calendar Window

D- Historical Anomaly

Answer:

D

Explanation:

historical anomaly is a detector condition that allows you to trigger an alert when a signal deviates from its historical pattern1.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 pattern1.You can use historical anomaly to detect unusual changes in a signal that are not explained by seasonality, trends, or cycles1.

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 accordingly1. 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 unexpectedly1.

Question 10

Question Type: MultipleChoice

A DevOps engineer wants to determine if the latency their application experiences is growing fester after a new software release a week ago. They have already created two plot lines, A and B, that represent the current latency and the latency a week ago, respectively. How can the engineer use these two plot lines to determine the rate of change in latency?

Options:

A- Create a temporary plot by dragging items A and B into the Analytics Explorer window.

B- Create a plot C using the formula (A-B) and add a scale:percent function to express the rate of change as a percentage.

C- Create a plot C using the formula (A/B-l) and add a scale: 100 function to express the rate of change as a percentage.

D- Create a temporary plot by clicking the Change% button in the upper-right corner of the plot showing lines A and B.

Answer:

C

Explanation:

The correct answer is C. Create a plot C using the formula (A/B-l) and add a scale: 100 function to express the rate of change as a percentage.

To calculate the rate of change in latency, you need to compare the current latency (plot A) with the latency a week ago (plot B). One way to do this is to use the formula (A/B-l), which gives you the ratio of the current latency to the previous latency minus one. This ratio represents how much the current latency has increased or decreased relative to the previous latency. For example, if the current latency is 200 ms and the previous latency is 100 ms, then the ratio is (200/100-l) = 1, which means the current latency is 100% higher than the previous latency1

To express the rate of change as a percentage, you need to multiply the ratio by 100. You can do this by adding a scale: 100 function to the formula. This function scales the values of the plot by a factor of 100. For example, if the ratio is 1, then the scaled value is 100%2

To create a plot C using the formula (A/B-l) and add a scale: 100 function, you need to follow these steps:

Select plot A and plot B from the Metric Finder.

Click on Add Analytics and choose Formula from the list of functions.

In the Formula window, enter (A/B-l) as the formula and click Apply.

Click on Add Analytics again and choose Scale from the list of functions.

In the Scale window, enter 100 as the factor and click Apply.

You should see a new plot C that shows the rate of change in latency as a percentage.

To learn more about how to use formulas and scale functions in Splunk Observability Cloud, you can refer to these documentations34.

1: https://www.mathsisfun.com/numbers/percentage-change.html 2: https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Scale 3: https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Formula 4: https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Scale

Question 11

Question Type: MultipleChoice

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

Options:

A- Line, Bar, Column

B- Bar, Area, Column

C- Line, Area, Column

D- Histogram, Line, Column

Answer:

C

Explanation:

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 plot1

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 preferences2

To learn more about how to use individual plot visualization in Splunk Observability Cloud, you can refer to this documentation2.

1: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Individual-plot-visualization 2: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Set-individual-plot-visualization