

Splunk.SPLK-4001.by.An.32q

Number: SPLK-4001
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: SPLK-4001

Exam Name: Splunk O11y Cloud Certified Metrics User Exam



Exam A

QUESTION 1

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

Section:

Explanation:

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.

QUESTION 2

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

Section:

Explanation:

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 correctly¹

To learn more about how to send metrics to Splunk Observability Cloud, you can refer to this documentation².

1: <https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Metric-types> 2: <https://docs.splunk.com/Observability/gdi/metrics/metrics.html>

QUESTION 3

Which analytic function can be used to discover peak page visits for a site over the last day?

- A. Maximum: Transformation (24h)
- B. Maximum: Aggregation (1d)
- C. Lag: (24h)
- D. Count: (1d)

Correct Answer: A

Section:

Explanation:



According to the Splunk Observability Cloud documentation¹, the maximum function is an analytic function that returns the highest value of a metric or a dimension over a specified time interval. The maximum function can be used as a transformation or an aggregation. A transformation applies the function to each metric time series (MTS) individually, while an aggregation applies the function to all MTS and returns a single value. For example, to discover the peak page visits for a site over the last day, you can use the following SignalFlow code:

```
maximum(24h, counters('page.visits'))
```

This will return the highest value of the page.visits counter metric for each MTS over the last 24 hours. You can then use a chart to visualize the results and identify the peak page visits for each MTS.

QUESTION 4

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

Section:

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.

QUESTION 5

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'.

Correct Answer: C

Section:

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>

QUESTION 6

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

Correct Answer: A

Section:

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.

QUESTION 7

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

Section:

Explanation:

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 characters¹

To learn more about how to filter metrics in Splunk Observability Cloud, you can refer to this documentation².

1: <https://docs.splunk.com/Observability/gdi/metrics/search.html#Filter-metrics> 2: <https://docs.splunk.com/Observability/gdi/metrics/search.html>

QUESTION 8

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

Section:

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document¹, 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 9

Which of the following are correct ports for the specified components in the OpenTelemetry Collector?

- A. gRPC (4000), SignalFx (9943), Fluentd (6060)
- B. gRPC (6831), SignalFx (4317), Fluentd (9080)
- C. gRPC (4459), SignalFx (9166), Fluentd (8956)
- D. gRPC (4317), SignalFx (9080), Fluentd (8006)

Correct Answer: D

Section:

Explanation:

The correct answer is D. gRPC (4317), SignalFx (9080), Fluentd (8006).

According to the web search results, these are the default ports for the corresponding components in the OpenTelemetry Collector. You can verify this by looking at the table of exposed ports and endpoints in the first result¹.

You can also see the agent and gateway configuration files in the same result for more details.

1: <https://docs.splunk.com/observability/gdi/opentelemetry/exposed-endpoints.html>

QUESTION 10

When writing a detector with a large number of MTS, such as memory.free in a deployment with 30,000 hosts, it is possible to exceed the cap of MTS that can be contained in a single plot. Which of the choices below would most likely reduce the number of MTS below the plot cap?

- A. Select the Sharded option when creating the plot.
- B. Add a filter to narrow the scope of the measurement.
- C. Add a restricted scope adjustment to the plot.
- D. When creating the plot, add a discriminator.

Correct Answer: B

Section:

Explanation:

The correct answer is B. Add a filter to narrow the scope of the measurement.

A filter is a way to reduce the number of metric time series (MTS) that are displayed on a chart or used in a detector. A filter specifies one or more dimensions and values that the MTS must have in order to be included. For example, if you want to monitor the memory.free metric only for hosts that belong to a certain cluster, you can add a filter like cluster:my-cluster to the plot or detector. This will exclude any MTS that do not have the cluster dimension or have a different value for it¹

Adding a filter can help you avoid exceeding the plot cap, which is the maximum number of MTS that can be contained in a single plot. The plot cap is 100,000 by default, but it can be changed by contacting Splunk Support². To learn more about how to use filters in Splunk Observability Cloud, you can refer to this documentation³.

1: <https://docs.splunk.com/Observability/gdi/metrics/search.html#Filter-metrics> 2: <https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Plot-cap> 3: <https://docs.splunk.com/Observability/gdi/metrics/search.html>

QUESTION 11

An SRE creates a new detector to receive an alert when server latency is higher than 260 milliseconds. Latency below 260 milliseconds is healthy for their service. The SRE creates a New Detector with a Custom Metrics Alert Rule for latency and sets a Static Threshold alert condition at 260ms.

How can the number of alerts be reduced?

- A. Adjust the threshold.
- B. Adjust the Trigger sensitivity. Duration set to 1 minute.
- C. Adjust the notification sensitivity. Duration set to 1 minute.
- D. Choose another signal.

Correct Answer: B

Section:

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document¹, trigger sensitivity is a setting that determines how long a signal must remain above or below a threshold before an alert is triggered. By default, trigger sensitivity is set to Immediate, which means that an alert is triggered as soon as the signal crosses the threshold. This can result in a lot of alerts, especially if the signal fluctuates frequently around the threshold value. To reduce the number of alerts, you can adjust the trigger sensitivity to a longer duration, such as 1 minute, 5 minutes, or 15 minutes. This means that an alert is only triggered if the signal stays above or below the threshold for the specified duration. This can help filter out noise and focus on more persistent issues.

QUESTION 12

Where does the Splunk distribution of the OpenTelemetry Collector store the configuration files on Linux machines by default?

- A. /opt/splunk/
- B. /etc/otel/collector/
- C. /etc/opentelemetry/
- D. /etc/system/default/

Correct Answer: B

Section:

Explanation:

The correct answer is B. /etc/otel/collector/

According to the web search results, the Splunk distribution of the OpenTelemetry Collector stores the configuration files on Linux machines in the /etc/otel/collector/ directory by default. You can verify this by looking at the first result¹, which explains how to install the Collector for Linux manually. It also provides the locations of the default configuration file, the agent configuration file, and the gateway configuration file.

To learn more about how to install and configure the Splunk distribution of the OpenTelemetry Collector, you can refer to this documentation².

1: <https://docs.splunk.com/Observability/gdi/opentelemetry/install-linux-manual.html> 2: <https://docs.splunk.com/Observability/gdi/opentelemetry.html>

QUESTION 13

A customer is experiencing issues getting metrics from a new receiver they have configured in the OpenTelemetry Collector. How would the customer go about troubleshooting further with the logging exporter?

- A. Adding debug into the metrics receiver pipeline:

```
metrics:
  receivers: [hostmetrics, otlp, signalfx, smartagent/signalfx-forwarder]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx, debug]
```

- B. Adding logging into the metrics receiver pipeline:

- C. Adding logging into the metrics exporter pipeline:

```
metrics:
  receivers: [hostmetrics, otlp, signalfx, smartagent/signalfx-forwarder]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx, debug]
```

- D. Adding debug into the metrics exporter pipeline:

Correct Answer: B

Section:

Explanation:

The correct answer is B. Adding logging into the metrics receiver pipeline.

The logging exporter is a component that allows the OpenTelemetry Collector to send traces, metrics, and logs directly to the console. It can be used to diagnose and troubleshoot issues with telemetry received and processed by the Collector, or to obtain samples for other purposes¹

To activate the logging exporter, you need to add it to the pipeline that you want to diagnose. In this case, since you are experiencing issues with a new receiver for metrics, you need to add the logging exporter to the metrics receiver pipeline. This will create a new plot that shows the metrics received by the Collector and any errors or warnings that might occur¹

The image that you have sent with your question shows how to add the logging exporter to the metrics receiver pipeline. You can see that the exporters section of the metrics pipeline includes logging as one of the options. This means that the metrics received by any of the receivers listed in the receivers section will be sent to the logging exporter as well as to any other exporters listed2 To learn more about how to use the logging exporter in Splunk Observability Cloud, you can refer to this documentation1.

1: <https://docs.splunk.com/Observability/gdi/opentelemetry/components/logging-exporter.html> 2: <https://docs.splunk.com/Observability/gdi/opentelemetry/exposed-endpoints.html>

QUESTION 14

What information is needed to create a detector?

- A. Alert Status, Alert Criteria, Alert Settings, Alert Message, Alert Recipients
- B. Alert Signal, Alert Criteria, Alert Settings, Alert Message, Alert Recipients
- C. Alert Signal, Alert Condition, Alert Settings, Alert Message, Alert Recipients
- D. Alert Status, Alert Condition, Alert Settings, Alert Meaning, Alert Recipients

Correct Answer: C

Section:

Explanation:

According to the Splunk Observability Cloud documentation1, to create a detector, you need the following information:

Alert Signal: This is the metric or dimension that you want to monitor and alert on. You can select a signal from a chart or a dashboard, or enter a SignalFlow query to define the signal.

Alert Condition: This is the criteria that determines when an alert is triggered or cleared. You can choose from various built-in alert conditions, such as static threshold, dynamic threshold, outlier, missing data, and so on. You can also specify the severity level and the trigger sensitivity for each alert condition.

Alert Settings: This is the configuration that determines how the detector behaves and interacts with other detectors. You can set the detector name, description, resolution, run lag, max delay, and detector rules. You can also enable or disable the detector, and mute or unmute the alerts.

Alert Message: This is the text that appears in the alert notification and event feed. You can customize the alert message with variables, such as signal name, value, condition, severity, and so on. You can also use markdown formatting to enhance the message appearance.

Alert Recipients: This is the list of destinations where you want to send the alert notifications. You can choose from various channels, such as email, Slack, PagerDuty, webhook, and so on. You can also specify the notification frequency and suppression settings.

QUESTION 15

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. TImeshift
- D. Standard deviation

Correct Answer: C

Section:

Explanation:

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 16

A customer is experiencing an issue where their detector is not sending email notifications but is generating alerts within the Splunk Observability UI. Which of the below is the root cause?

- A. The detector has an incorrect alert rule.
- B. The detector has an incorrect signal,
- C. The detector is disabled.
- D. The detector has a muting rule.

Correct Answer: D

Section:

Explanation:

The most likely root cause of the issue is D. The detector has a muting rule.

A muting rule is a way to temporarily stop a detector from sending notifications for certain alerts, without disabling the detector or changing its alert conditions. A muting rule can be useful when you want to avoid alert noise during planned maintenance, testing, or other situations where you expect the metrics to deviate from normal¹

When a detector has a muting rule, it will still generate alerts within the Splunk Observability UI, but it will not send email notifications or any other types of notifications that you have configured for the detector. You can see if a detector has a muting rule by looking at the Muting Rules tab on the detector page. You can also create, edit, or delete muting rules from there¹

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

QUESTION 17

To smooth a very spiky cpu.utilization metric, what is the correct analytic function to better see if the cpu. utilization for servers is trending up over time?

- A. Rate/Sec
- B. Median
- C. Mean (by host)
- D. Mean (Transformation)

Correct Answer: D

Section:

Explanation:

The correct answer is D. Mean (Transformation).

According to the web search results, a mean transformation is an analytic function that returns the average value of a metric or a dimension over a specified time interval¹. A mean transformation can be used to smooth a very spiky metric, such as cpu.utilization, by reducing the impact of outliers and noise. A mean transformation can also help to see if the metric is trending up or down over time, by showing the general direction of the average value. For example, to smooth the cpu.utilization metric and see if it is trending up over time, you can use the following SignalFlow code:

```
mean(1h, counters("cpu.utilization"))
```

This will return the average value of the cpu.utilization counter metric for each metric time series (MTS) over the last hour. You can then use a chart to visualize the results and compare the mean values across different MTS.

Option A is incorrect because rate/sec is not an analytic function, but rather a rollup function that returns the rate of change of data points in the MTS reporting interval¹. Rate/sec can be used to convert cumulative counter metrics into counter metrics, but it does not smooth or trend a metric. Option B is incorrect because median is not an analytic function, but rather an aggregation function that returns the middle value of a metric or a dimension over the entire time range¹. Median can be used to find the typical value of a metric, but it does not smooth or trend a metric. Option C is incorrect because mean (by host) is not an analytic function, but rather an aggregation function that returns the average value of a metric or a dimension across all MTS with the same host dimension¹. Mean (by host) can be used to compare the performance of different hosts, but it does not smooth or trend a metric.

Mean (Transformation) is an analytic function that allows you to smooth a very spiky metric by applying a moving average over a specified time window. This can help you see the general trend of the metric over time, without being distracted by the short-term fluctuations¹

To use Mean (Transformation) on a cpu.utilization metric, you need to select the metric from the Metric Finder, then click on Add Analytics and choose Mean (Transformation) from the list of functions. You can then specify the time window for the moving average, such as 5 minutes, 15 minutes, or 1 hour. You can also group the metric by host or any other dimension to compare the smoothed values across different servers²

To learn more about how to use Mean (Transformation) and other analytic functions in Splunk Observability Cloud, you can refer to this documentation².

1: <https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Mean-Transformation>2: <https://docs.splunk.com/Observability/gdi/metrics/analytics.html>

QUESTION 18

What happens when the limit of allowed dimensions is exceeded for an MTS?

- A. The additional dimensions are dropped.



- B. The datapoint is averaged.
- C. The datapoint is updated.
- D. The datapoint is dropped.

Correct Answer: A

Section:

Explanation:

According to the web search results, dimensions are metadata in the form of key-value pairs that monitoring software sends in along with the metrics. The set of metric time series (MTS) dimensions sent during ingest is used, along with the metric name, to uniquely identify an MTS¹. Splunk Observability Cloud has a limit of 36 unique dimensions per MTS². If the limit of allowed dimensions is exceeded for an MTS, the additional dimensions are dropped and not stored or indexed by Observability Cloud². This means that the data point is still ingested, but without the extra dimensions. Therefore, option A is correct.

QUESTION 19

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

Section:

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>

QUESTION 20

A DevOps engineer wants to determine if the latency their application experiences is growing faster 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?

- 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 $(A/B-1)$ 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.

Correct Answer: C

Section:

Explanation:

The correct answer is C. Create a plot C using the formula $(A/B-1)$ 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 $(A/B-1)$, 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 $(200/100-1) = 1$, 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-I) 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-I) 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 21

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

Section:

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 22

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

Section:

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 23

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

- A. Processors
- B. Pipelines
- C. Exporters
- D. Receivers

Correct Answer: A

Section:

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 24

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

Section:

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 documentations^{1,2}.

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 25

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

- 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.

Correct Answer: A, D

Section:

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 26

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

Section:

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 27

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

- 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.

Correct Answer: C

Section:

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 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 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 28

The alert recipients tab specifies where notification messages should be sent when alerts are triggered or cleared. Which of the below options can be used? (select all that apply)

- A. Invoke a webhook URL.
- B. Export to CSV.
- C. Send an SMS message.
- D. Send to email addresses.

Correct Answer: A, C, D

Section:

Explanation:

The alert recipients tab specifies where notification messages should be sent when alerts are triggered or cleared. The options that can be used are:

Invoke a webhook URL. This option allows you to send a HTTP POST request to a custom URL that can perform various actions based on the alert information. For example, you can use a webhook to create a ticket in a service desk system, post a message to a chat channel, or trigger another workflow¹

Send an SMS message. This option allows you to send a text message to one or more phone numbers when an alert is triggered or cleared. You can customize the message content and format using variables and templates²

Send to email addresses. This option allows you to send an email notification to one or more recipients when an alert is triggered or cleared. You can customize the email subject, body, and attachments using variables and templates. You can also include information from search results, the search job, and alert triggering in the email³

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

1: <https://docs.splunk.com/Documentation/Splunk/latest/Alert/Webhooks> 2: <https://docs.splunk.com/Documentation/Splunk/latest/Alert/SMSnotification> 3: <https://docs.splunk.com/Documentation/Splunk/latest/Alert/Emailnotification>

QUESTION 29

With exceptions for transformations or timeshifts, at what resolution do detectors operate?

- A. 10 seconds
- B. The resolution of the chart
- C. The resolution of the dashboard
- D. Native resolution

Correct Answer: D

Section:

Explanation:

According to the Splunk Observability Cloud documentation¹, detectors operate at the native resolution of the metric or dimension that they monitor, with some exceptions for transformations or timeshifts. The native resolution is the frequency at which the data points are reported by the source. For example, if a metric is reported every 10 seconds, the detector will evaluate the metric every 10 seconds. The native resolution ensures that the detector uses the most granular and accurate data available for alerting.

QUESTION 30

Which of the following are true about organization metrics? (select all that apply)

- A. Organization metrics give insights into system usage, system limits, data ingested and token quotas.
- B. Organization metrics count towards custom MTS limits.
- C. Organization metrics are included for free.
- D. A user can plot and alert on them like metrics they send to Splunk Observability Cloud.

Correct Answer: A, C, D

Section:

Explanation:

The correct answer is A, C, and D. Organization metrics give insights into system usage, system limits, data ingested and token quotas. Organization metrics are included for free. A user can plot and alert on them like metrics they send to Splunk Observability Cloud.

Organization metrics are a set of metrics that Splunk Observability Cloud provides to help you measure your organization's usage of the platform. They include metrics such as:

Ingest metrics: Measure the data you're sending to Infrastructure Monitoring, such as the number of data points you've sent.

App usage metrics: Measure your use of application features, such as the number of dashboards in your organization.

Integration metrics: Measure your use of cloud services integrated with your organization, such as the number of calls to the AWS CloudWatch API.

Resource metrics: Measure your use of resources that you can specify limits for, such as the number of custom metric time series (MTS) you've created¹

Organization metrics are not charged and do not count against any system limits. You can view them in built-in charts on the Organization Overview page or in custom charts using the Metric Finder. You can also create alerts based on organization metrics to monitor your usage and performance¹

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

1: <https://docs.splunk.com/observability/admin/org-metrics.html>

QUESTION 31

Which of the following aggregate analytic functions will allow a user to see the highest or lowest n values of a metric?

- A. Maximum / Minimum
- B. Best/Worst
- C. Exclude / Include
- D. Top / Bottom

Correct Answer: D

Section:

Explanation:

The correct answer is D. Top / Bottom.

Top and bottom are aggregate analytic functions that allow a user to see the highest or lowest n values of a metric. They can be used to select a subset of the time series in the plot by count or by percent. For example, top (5) will show the five time series with the highest values in each time period, while bottom (10%) will show the 10% of time series with the lowest values in each time period¹

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

QUESTION 32

Which of the following are ways to reduce flapping of a detector? (select all that apply)

- A. Configure a duration or percent of duration for the alert.
- B. Establish a reset threshold for the detector.
- C. Enable the anti-flap setting in the detector options menu.
- D. Apply a smoothing transformation (like a rolling mean) to the input data for the detector.

Correct Answer: A, D

Section:

Explanation:

According to the Splunk Lantern article [Resolving flapping detectors in Splunk Infrastructure Monitoring](#), flapping is a phenomenon where alerts fire and clear repeatedly in a short period of time, due to the signal fluctuating around the threshold value. To reduce flapping, the article suggests the following ways:

Configure a duration or percent of duration for the alert: This means that you require the signal to stay above or below the threshold for a certain amount of time or percentage of time before triggering an alert. This can help filter out noise and focus on more persistent issues.

Apply a smoothing transformation (like a rolling mean) to the input data for the detector: This means that you replace the original signal with the average of its last several values, where you can specify the window length. This can reduce the impact of a single extreme observation and make the signal less fluctuating.

