[Q32-Q54 Pass Splunk SPLK-4001 Exam in First Attempt Guaranteed [Jun-2024

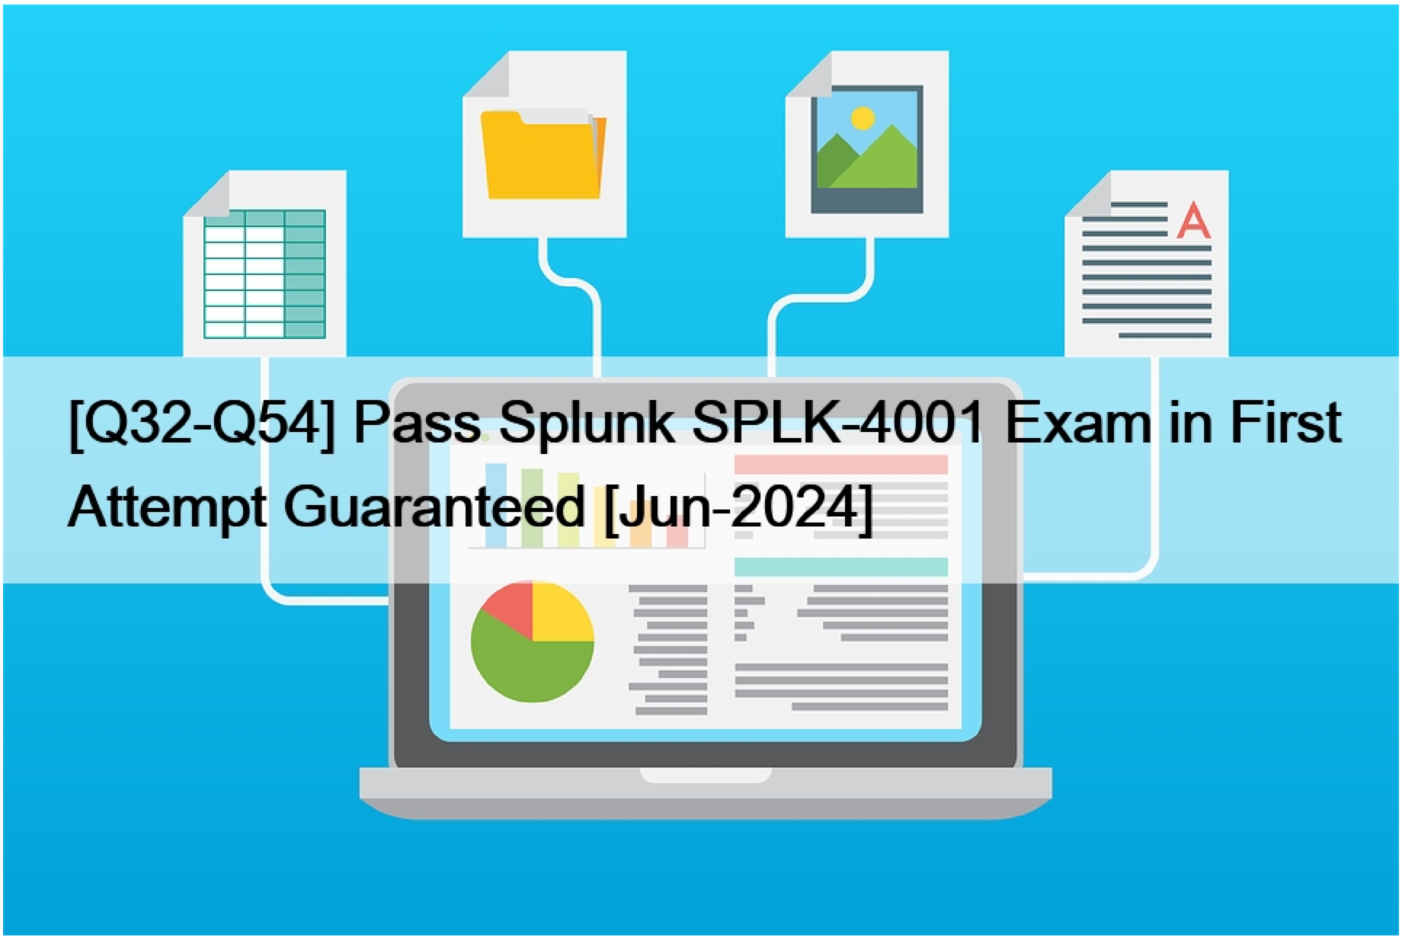

Pass Splunk SPLK-4001 Exam in First Attempt Guaranteed [Jun-2024]
Exam Sure Pass Splunk Certification with SPLK-4001 exam questions

Splunk SPLK-4001 exam is designed for individuals who are interested in obtaining the Splunk O11y Cloud Certified Metrics User certification. SPLK-4001 exam is intended for those who have a strong understanding of Splunk and its capabilities, as well as experience working with metrics data. Splunk O11y Cloud Certified Metrics User certification is ideal for professionals who want to demonstrate their expertise in using Splunk to monitor and analyze metrics data in cloud environments.

The SPLK-4001 certification is highly valued in the IT industry, as it demonstrates a candidate's proficiency in using Splunk's Observability Cloud. It is a globally recognized certification that can help professionals advance their careers in cloud monitoring and analysis. By passing the SPLK-4001 exam, candidates can prove their expertise in using Splunk's Observability Cloud to monitor their organization's infrastructure and ensure its smooth operation. Splunk O11y Cloud Certified Metrics User certification is ideal for IT professionals who want to enhance their skills and knowledge in cloud monitoring and analysis.

QUESTION 32

What are the best practices for creating detectors? (select all that apply)
* View data at highest resolution.
* Have a consistent value.
* View detector in a chart.
* Have a consistent type of measurement.
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

QUESTION 33

Which of the following are supported rollup functions in Splunk Observability Cloud?
* average, latest, lag, min, max, sum, rate
* std_dev, mean, median, mode, min, max
* sigma, epsilon, pi, omega, beta, tau
* 1min, 5min, 10min, 15min, 30min
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 34

A customer is sending data from a machine that is over-utilized. Because of a lack of system resources, datapoints from this machine are often delayed by up to 10 minutes. Which setting can be modified in a detector to prevent alerts from firing before the datapoints arrive?

* Max Delay
* Duration
* Latency
* Extrapolation Policy

Explanation

The correct answer is A. Max Delay.

Max Delay is a parameter that specifies the maximum amount of time that the analytics engine can wait for data to arrive for a specific detector. For example, if Max Delay is set to 10 minutes, the detector will wait for only a maximum of 10 minutes even if some data points have not arrived. By default, Max Delay is set to Auto, allowing the analytics engine to determine the appropriate amount of time to wait for data points1 In this case, since the customer knows that the data from the over-utilized machine can be delayed by up to 10 minutes, they can modify the Max Delay setting for the detector to 10 minutes. This will prevent the detector from firing alerts before the data points arrive, and avoid false positives or missing data1 To learn more about how to use Max Delay in Splunk Observability Cloud, you can refer to this documentation1.

1: https://docs.splunk.com/observability/alerts-detectors-notifications/detector-options.html#Max-Delay

QUESTION 35

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?

* The detector has an incorrect alert rule.
* The detector has an incorrect signal,
* The detector is disabled.
* The detector has a muting rule.

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 36

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

* Public dashboards
* Dashboard groups
* Chart exporter
* Shared charts

Explanation

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 37

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?
* Rate/Sec
* Median
* Mean (by host)
* Mean (Transformation)
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(&#8220;cpu.utilization&#8221;))
```

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 38

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 &#8216;canary&#8217; version dimension. They&#8217;ve already opened the graph of memory utilization for their service.

How does the engineer see if the new release lowered average memory utilization?
* On the chart for plot A, select Add Analytics, then select MeanrTransformation. In the window that appears, select &#8216;version&#8217; from the Group By field.
* On the chart for plot A, scroll to the end and click Enter Function, then enter &#8216;A/B-l&#8217;.
* On the chart for plot A, select Add Analytics, then select Mean:Aggregation. In the window that appears, select &#8216;version&#8217; from the Group By field.
* On the chart for plot A, click the Compare Means button. In the window that appears, type &#8216;version1.
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 &#8216;version&#8217; 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 &#8216;canary&#8217; and &#8216;stable&#8217; 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 39

Which of the following chart visualization types are unaffected by changing the time picker on a dashboard?

(select all that apply)
* Single Value
* Heatmap
* Line
* List
Explanation

The chart visualization types that are unaffected by changing the time picker on a dashboard are:

Single Value: A single value chart shows the current value of a metric or an expression. It does not depend on the time range of the dashboard, but only on the data resolution and rollup function of the chart1 List: A list chart shows the values of a metric or an expression for each dimension value in a table format. It does not depend on the time range of the dashboard, but only on the data resolution and rollup function of the chart2 Therefore, the correct answer is A and D.

To learn more about how to use different chart visualization types in Splunk Observability Cloud, you can refer to this documentation3.

1: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Single-value 2:

https://docs.splunk.com/Observability/gdi/metrics/charts.html#List 3:

https://docs.splunk.com/Observability/gdi/metrics/charts.html

QUESTION 40

Which of the following are required in the configuration of a data point? (select all that apply)
* Metric Name
* Metric Type
* Timestamp
* Value
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 41

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?
* The new signals will be reflected in the original detector.
* The new signals will be reflected in the original chart.
* You can only monitor one of the new signals.
* The new signals will not be added to the original detector.
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.

QUESTION 42

Which of the following is optional, but highly recommended to include in a datapoint?
* Metric name
* Timestamp
* Value

* Metric type
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 43

Given that the metric demo. trans. count is being sent at a 10 second native resolution, which of the following is an accurate description of the data markers displayed in the chart below?

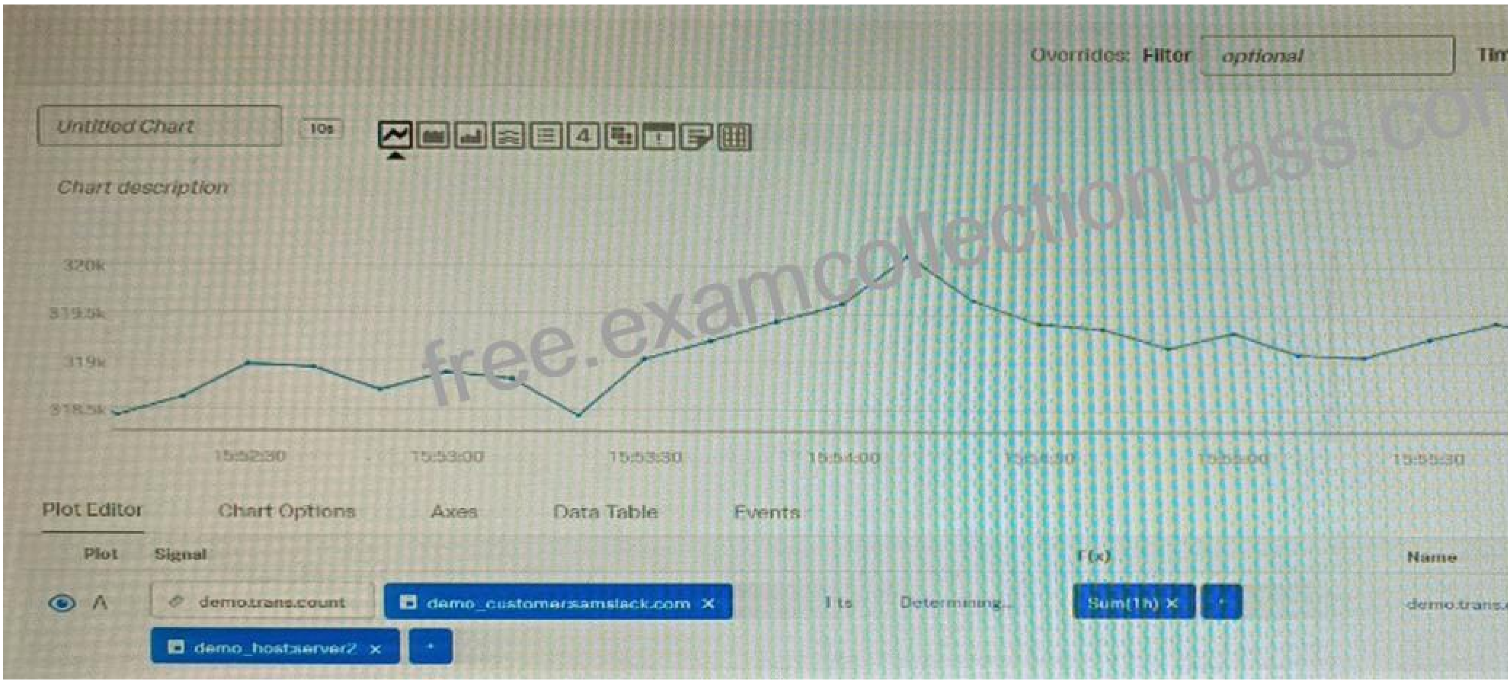

* Each data marker represents the average hourly rate of API calls.
* Each data marker represents the 10 second delta between counter values.
* Each data marker represents the average of the sum of datapoints over the last minute, averaged over the hour.
* Each data marker represents the sum of API calls in the hour leading up to the data marker.
Explanation

The correct answer is D. Each data marker represents the sum of API calls in the hour leading up to the data marker.

The metric demo.trans.count is a cumulative counter metric, which means that it represents the total number of API calls since the start of the measurement. A cumulative counter metric can be used to measure the rate of change or the sum of events over a time period1 The chart below shows the metric demo.trans.count with a one-hour rollup and a line chart type. A rollup is a way to aggregate data points over a specified time interval, such as one hour, to reduce the number of data points displayed on a chart. A line chart type connects the data points with a line to show the trend of the metric over time2 Each data marker on the chart represents the sum of API calls in the hour leading up to the data marker. This is because the rollup function for cumulative counter metrics is sum by default, which means that it adds up all the data points in each time interval. For example, the data marker at 10:00 AM shows the sum of API calls from 9:00 AM to 10:00 AM3 To learn more about how to use metrics and charts in Splunk Observability Cloud, you can refer to these documentations123.

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

https://docs.splunk.com/Observability/gdi/metrics/charts.html#Data-resolution-and-rollups-in-charts 3:

https://docs.splunk.com/Observability/gdi/metrics/charts.html#Rollup-functions-for-metric-types

QUESTION 44

Which of the following statements about adding properties to MTS are true? (select all that apply)
* Properties can be set via the API.
* Properties are sent in with datapoints.
* Properties are applied to dimension key:value pairs and propagated to all MTS with that dimension
* Properties can be set in the UI under Metric Metadata.
Explanation

According to the web search results, properties are key-value pairs that you can assign to dimensions of existing metric time series (MTS) in Splunk Observability Cloud1. Properties provide additional context and information about the metrics, such as the environment, role, or owner of the dimension. For example, you can add the property use: QA to the host dimension of your metrics to indicate that the host that is sending the data is used for QA.

To add properties to MTS, you can use either the API or the UI. The API allows you to programmatically create, update, delete, and list properties for dimensions using HTTP requests2. The UI allows you to interactively create, edit, and delete properties for dimensions using the Metric Metadata page under Settings3.

Therefore, option A and D are correct.

QUESTION 45

The Sum Aggregation option for analytic functions does which of the following?
* Calculates the number of MTS present in the plot.
* Calculates 1/2 of the values present in the input time series.
* Calculates the sum of values present in the input time series across the entire environment or per group.
* Calculates the sum of values per time series across a period of time.
Explanation

According to the Splunk Test Blueprint &#8211; 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 46

Which of the following are accurate reasons to clone a detector? (select all that apply)
* To modify the rules without affecting the existing detector.
* To reduce the amount of billed TAPM for the detector.
* To add an additional recipient to the detector&#8217;s alerts.
* To explore how a detector was created without risk of changing it.
Explanation

The correct answers are A and D.

According to the Splunk Test Blueprint &#8211; 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&#8217;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&#8217;s alerts, you need to edit the alert settings of the detector.

QUESTION 47

Which of the following aggregate analytic functions will allow a user to see the highest or lowest n values of a metric?
* Maximum / Minimum
* Best/Worst
* Exclude / Include
* Top / Bottom
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 48

Which component of the OpenTelemetry Collector allows for the modification of metadata?
* Processors
* Pipelines
* Exporters
* Receivers
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 49

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?
* Percentages and ratios
* Timeshift and Bottom N

* Timeshift and Top N
* Chart Options and metadata
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(&#8220;cache.hits&#8221;), counters(&#8220;cache.misses&#8221;))

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(&#8220;cache.hits&#8221;), counters(&#8220;cache.misses&#8221;))

QUESTION 50

Which of the following are ways to reduce flapping of a detector? (select all that apply)
* Configure a duration or percent of duration for the alert.
* Establish a reset threshold for the detector.
* Enable the anti-flap setting in the detector options menu.
* Apply a smoothing transformation (like a rolling mean) to the input data for the detector.
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.

QUESTION 51

A user wants to add a link to an existing dashboard from an alert. When they click the dimension value in the alert message, they are taken to the dashboard keeping the context. How can this be accomplished? (select all that apply)
* Build a global data link.
* Add a link to the Runbook URL.
* Add a link to the field.
* Add the link to the alert message body.
Explanation

The possible ways to add a link to an existing dashboard from an alert are:

Build a global data link. A global data link is a feature that allows you to create a link from any dimension value in any chart or table to a dashboard of your choice. You can specify the source and target dashboards, the dimension name and value, and the query parameters to pass along. When you click on the dimension value in the alert message, you will be taken to the dashboard with the context preserved1 Add a link to the field. A field link is a feature that allows you to create a link from any field value in any search result or alert message to a dashboard of your choice. You can specify the field name and value, the dashboard name and ID, and the query parameters to pass along. When you click on the field value in the alert message, you will be taken to the dashboard with the context preserved2 Therefore, the correct answer is A and C.

To learn more about how to use global data links and field links in Splunk Observability Cloud, you can refer to these documentations12.

1: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Global-data-links 2:

https://docs.splunk.com/Observability/gdi/metrics/search.html#Field-links

QUESTION 52

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)
* Invoke a webhook URL.
* Export to CSV.
* Send an SMS message.
* Send to email addresses.
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 53

What happens when the limit of allowed dimensions is exceeded for an MTS?
* The additional dimensions are dropped.
* The datapoint is averaged.
* The datapoint is updated.
* The datapoint is dropped.
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.

Real Splunk SPLK-4001 Exam Questions Study Guide:
https://www.examcollectionpass.com/Splunk/SPLK-4001-practice-exam-dumps.html]