SPLK-4001 Dumps

Splunk O11y Cloud Certified Metrics User

https://www.certleader.com/SPLK-4001-dumps.html

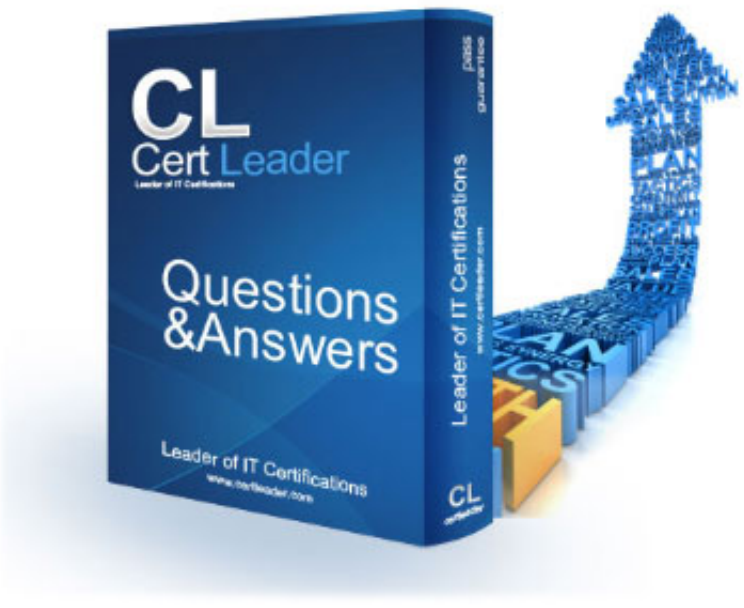

NEW QUESTION 1
Which of the following are correct ports for the specified components in the OpenTelemetry Collector?

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)

Answer: D

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

NEW QUESTION 2
A customer has a very dynamic infrastructure. During every deployment, all existing instances are destroyed, and new ones are created Given this deployment model, how should a detector be created that will not send false notifications of instances being down?

A. Create the detecto
B. Select Alert settings, then select Auto-Clear Alerts and enter an appropriate time period.
C. Create the detecto
D. Select Alert settings, then select Ephemeral Infrastructure and enter the expected lifetime of an instance.
E. Check the Dynamic checkbox when creating the detector.
F. Check the Ephemeral checkbox when creating the detector.

Answer: B

Explanation:
According to the web search results, ephemeral infrastructure is a term that describes instances that are auto-scaled up or down, or are brought up with new code versions and discarded or recycled when the next code version is deployed1. Splunk Observability Cloud has a feature that allows you to create detectors for ephemeral infrastructure without sending false notifications of instances being down2. To use this feature, you need to do the following steps:
? Create the detector as usual, by selecting the metric or dimension that you want to monitor and alert on, and choosing the alert condition and severity level.
? Select Alert settings, then select Ephemeral Infrastructure. This will enable a special mode for the detector that will automatically clear alerts for instances that are expected to be terminated.
? Enter the expected lifetime of an instance in minutes. This is the maximum amount of time that an instance is expected to live before being replaced by a new one. For example, if your instances are replaced every hour, you can enter 60 minutes as the expected lifetime.
? Save the detector and activate it.
With this feature, the detector will only trigger alerts when an instance stops reporting a metric unexpectedly, based on its expected lifetime. If an instance stops reporting a metric within its expected lifetime, the detector will assume that it was terminated on purpose and will not trigger an alert. Therefore, option B is correct.

NEW QUESTION 3
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?


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

Answer: D

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

NEW QUESTION 4
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

Answer: C

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.

NEW QUESTION 5
Which of the following statements is true of detectors created from a chart on a custom dashboard?

A. Changes made to the chart affect the detector.
B. Changes made to the detector affect the chart.
C. The alerts will show up in the team landing page.
D. The detector is automatically linked to the chart.

Answer: D

Explanation:
The correct answer is D. The detector is automatically linked to the chart. When you create a detector from a chart on a custom dashboard, the detector is automatically linked to the chart. This means that you can see the detector status and alerts on the chart, and you can access the detector settings from the chart menu. You can also unlink the detector from the chart if you want to1
Changes made to the chart do not affect the detector, and changes made to the detector do not affect the chart. The detector and the chart are independent entities that have their own settings and parameters. However, if you change the metric or dimension of the chart, you might lose the link to the detector1
The alerts generated by the detector will show up in the Alerts page, where you can view, manage, and acknowledge them. You can also see them on the team landing page if you assign the detector to a team2
To learn more about how to create and link detectors from charts on custom dashboards, you can refer to this documentation1.
1: https://docs.splunk.com/observability/alerts-detectors-notifications/link-detectors-to-charts.html 2: https://docs.splunk.com/observability/alerts-detectors-notifications/view- manage-alerts.html

NEW QUESTION 6
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.

Answer: D

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.

NEW QUESTION 7
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

Answer: D

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.

NEW QUESTION 8
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

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.

NEW QUESTION 9
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

Answer: A

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”))

NEW QUESTION 10
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)

Answer: D

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

NEW QUESTION 10

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.

Answer: AD

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.

NEW QUESTION 13
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.

Answer: B

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

NEW QUESTION 14
When installing OpenTelemetry Collector, which error message is indicative that there is a misconfigured realm or access token?

A. 403 (NOT ALLOWED)
B. 404 (NOT FOUND)
C. 401 (UNAUTHORIZED)
D. 503 (SERVICE UNREACHABLE)

Answer: C

Explanation:
The correct answer is C. 401 (UNAUTHORIZED).
According to the web search results, a 401 (UNAUTHORIZED) error message is indicative that there is a misconfigured realm or access token when installing OpenTelemetry Collector1. A 401 (UNAUTHORIZED) error message means that the request was not authorized by the server due to invalid credentials. A realm is a parameter that specifies the scope of protection for a resource, such as a Splunk Observability Cloud endpoint. An access token is a credential that grants access to a resource, such as a Splunk Observability Cloud API. If the realm or the access token is misconfigured, the request to install OpenTelemetry Collector will be rejected by the server with a 401 (UNAUTHORIZED) error message.
Option A is incorrect because a 403 (NOT ALLOWED) error message is not indicative that there is a misconfigured realm or access token when installing OpenTelemetry Collector. A 403 (NOT ALLOWED) error message means that the request was authorized by the server but not allowed due to insufficient permissions. Option B is incorrect because a 404 (NOT FOUND) error message is not indicative that there is a misconfigured realm or access token when installing OpenTelemetry Collector. A 404 (NOT FOUND) error message means that the request was not found by the server due to an invalid URL or resource. Option D is incorrect because a 503 (SERVICE UNREACHABLE) error message is not indicative that there is a misconfigured realm or access token when installing OpenTelemetry Collector. A 503 (SERVICE UNREACHABLE) error message means that the server was unable to handle the request due to temporary overload or maintenance.

NEW QUESTION 18
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

Answer: A

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.

NEW QUESTION 22
Which component of the OpenTelemetry Collector allows for the modification of metadata?

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

NEW QUESTION 25
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?

A. The new signals will be reflected in the original detector.
B. The new signals will be reflected in the original chart.
C. You can only monitor one of the new signals.
D. The new signals will not be added to the original detector.

Answer: D

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.

NEW QUESTION 28
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SPLK-4001 Exam with Our Prep Materials Via below:

https://www.certleader.com/SPLK-4001-dumps.html