



Splunk

Exam Questions SPLK-4001

Splunk O11y Cloud Certified Metrics User

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

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 3

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)

- A. Build a global data link.
- B. Add a link to the Runbook URL.
- C. Add a link to the field.
- D. Add the link to the alert message body.

Answer: AC

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>

NEW QUESTION 4

What constitutes a single metrics time series (MTS)?

- A. A series of timestamps that all reflect the same metric.
- B. A set of data points that all have the same metric name and list of dimensions.
- C. A set of data points that use different dimensions but the same metric name.
- D. A set of metrics that are ordered in series based on timestamp.

Answer: B

Explanation:

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

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

MTS1: Gauge metric cpu.utilization, dimension "hostname": "host1" MTS2: Gauge metric cpu.utilization, dimension "hostname": "host2" MTS3: Gauge metric memory.usage, dimension "hostname": "host1"

A metric is a numerical measurement that varies over time, such as CPU utilization or memory usage. A dimension is a key-value pair that provides additional information about the metric, such as the hostname or the location. A data point is a combination of a metric, a dimension, a value, and a timestamp¹

NEW QUESTION 5

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 6

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-I) and add a scale: 100 function to express the rate of change as a percentage.
- D. Create a temporary plot by clicking the Change% button in the upper-right corner of the plot showing lines A and B.

Answer: C

Explanation:

The correct answer is C. Create a plot C using the formula (A/B-I) 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-I), 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>

NEW QUESTION 7

Which of the following chart visualization types are unaffected by changing the time picker on a dashboard? (select all that apply)

- A. Single Value
- B. Heatmap
- C. Line

D. List

Answer: AD

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>

NEW QUESTION 8

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-](https://docs.splunk.com/observability/alerts-detectors-notifications/link-detectors-to-charts.html)

[charts.html](https://docs.splunk.com/observability/alerts-detectors-notifications/view-manage-alerts.html) 2: <https://docs.splunk.com/observability/alerts-detectors-notifications/view-manage-alerts.html>

NEW QUESTION 9

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 10

Which of the following is optional, but highly recommended to include in a datapoint?

- A. Metric name
- B. Timestamp
- C. Value
- D. Metric type

Answer: D

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>

NEW QUESTION 10

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.

Answer: AD

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.

NEW QUESTION 15

What Pod conditions does the Analyzer panel in Kubernetes Navigator monitor? (select all that apply)

- A. Not Scheduled
- B. Unknown
- C. Failed
- D. Pending

Answer: ABCD

Explanation:

The Pod conditions that the Analyzer panel in Kubernetes Navigator monitors are:

? Not Scheduled: This condition indicates that the Pod has not been assigned to a Node yet. This could be due to insufficient resources, node affinity, or other scheduling constraints¹

? Unknown: This condition indicates that the Pod status could not be obtained or is not known by the system. This could be due to communication errors, node failures, or other unexpected situations¹

? Failed: This condition indicates that the Pod has terminated in a failure state. This could be due to errors in the application code, container configuration, or external factors¹

? Pending: This condition indicates that the Pod has been accepted by the system, but one or more of its containers has not been created or started yet. This could be due to image pulling, volume mounting, or network issues¹

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

To learn more about how to use the Analyzer panel in Kubernetes Navigator, you can refer to this documentation².

1: <https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#pod-phase> 2: <https://docs.splunk.com/observability/infrastructure/monitor/k8s-nav.html#Analyzer-panel>

NEW QUESTION 17

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 sensitivit
- C. Duration set to 1 minute.
- D. Adjust the notification sensitivit
- E. Duration set to 1 minute.
- F. Choose another signal.

Answer: B

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.

NEW QUESTION 20

How is it possible to create a dashboard group that no one else can edit?

- A. Ask the admin to lock the dashboard group.
- B. Restrict the write access on the dashboard group.
- C. Link the dashboard group to the team.
- D. Hide the edit menu on the dashboard group.

Answer: B

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 set permissions for each dashboard group, such as who can view, edit, or manage the dashboards in the group¹. To create a dashboard group that no one else can edit, you need to do the following steps:

? Create a dashboard group as usual, by selecting Dashboard Group from the

Create menu on the navigation bar, entering a name and description, and adding dashboards to the group¹.

? Select Alert settings from the Dashboard actions menu () on the top right corner of the dashboard group. This will open a dialog box where you can configure the permissions for the dashboard group¹.

? Under Write access, select Only me. This will restrict the write access to the

dashboard group to yourself only. No one else will be able to edit or delete the dashboards in the group¹.

? Click Save. This will create a dashboard group that no one else can edit.

NEW QUESTION 24

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 29

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.

Answer: A

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>

NEW 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?

A. Max Delay

B. Duration

C. Latency

D. Extrapolation Policy

Answer: A

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>

NEW QUESTION 38

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

Answer: A

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.

NEW QUESTION 43

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

- A. Anyone with write permission for a dashboard can add local data links that appear on that dashboard.
- B. Local data links can only have a Splunk Observability Cloud internal destination.
- C. Only Splunk Observability Cloud administrators can create local links.
- D. Local data links are available on only one dashboard.

Answer: AD

Explanation:

The correct answers are A and D.

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

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

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

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

? Only Splunk Observability Cloud administrators can delete local data links. Therefore, based on this document, we can conclude that A and D are true statements about local data links. B and C are false statements because:

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

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

NEW QUESTION 46

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 MeanrTransformatio
- B. In the window that appears, select 'version' from the Group By field.
- C. On the chart for plot A, scroll to the end and click Enter Function, then enter 'A/B-!'.
- D. On the chart for plot A, select Add Analytics, then select Mean:Aggregatio
- E. In the window that appears, select 'version' from the Group By field.
- F. On the chart for plot A, click the Compare Means butto
- G. In the window that appears, type 'version1'.

Answer: C

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>

NEW QUESTION 51

Changes to which type of metadata result in a new metric time series?

- A. Dimensions
- B. Properties
- C. Sources
- D. Tags

Answer: A

Explanation:

The correct answer is A. Dimensions.

Dimensions are metadata in the form of key-value pairs that are sent along with the metrics at the time of ingest. They provide additional information about the metric, such as the name of the host that sent the metric, or the location of the server. Along with the metric name, they uniquely identify a metric time series (MTS)¹

Changes to dimensions result in a new MTS, because they create a different combination of metric name and dimensions. For example, if you change the hostname dimension from host1 to host2, you will create a new MTS for the same metric name¹

Properties, sources, and tags are other types of metadata that can be applied to existing MTSES after ingest. They do not contribute to uniquely identify an MTS, and they do not create a new MTS when changed²

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

1: <https://docs.splunk.com/Observability/metrics-and-metadata/metrics.html#Dimensions> 2: <https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html>

NEW QUESTION 52

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 57

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

Answer: C

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.

NEW QUESTION 61

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

Answer: D

Explanation:

historical anomaly is a detector condition that allows you to trigger an alert when a signal deviates from its historical pattern¹. Historical anomaly uses machine learning to learn the normal behavior of a signal based on its past data, and then compares the current value of the signal with the expected value based on the learned pattern¹. You can use historical anomaly to detect unusual changes in a signal that are not explained by seasonality, trends, or cycles¹.

Historical anomaly is suitable for creating a detector for the customer's data, because it can account for the expected and consistent increase in traffic during November each

year. Historical anomaly can learn that the traffic pattern has a seasonal component that peaks in November, and then adjust the expected value of the traffic accordingly¹. This way, historical anomaly can avoid triggering alerts when the traffic increases in November, as this is not an anomaly, but rather a normal variation. However, historical anomaly can

still trigger alerts when the traffic deviates from the historical pattern in other ways, such as if it drops significantly or spikes unexpectedly¹.

NEW QUESTION 63

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.

Answer: ACD

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>

NEW QUESTION 65

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

Answer: C

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.

NEW QUESTION 69

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

SPLK-4001 Practice Exam Features:

- * SPLK-4001 Questions and Answers Updated Frequently
- * SPLK-4001 Practice Questions Verified by Expert Senior Certified Staff
- * SPLK-4001 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SPLK-4001 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SPLK-4001 Practice Test Here](#)