



# Splunk

## Exam Questions SPLK-4001

Splunk O11y Cloud Certified Metrics User

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

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

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

**Answer:** A

#### Explanation:

According to the Splunk Observability Cloud documentation<sup>1</sup>, the maximum function is an analytic function that returns the highest value of a metric or a dimension over a specified time interval. The maximum function can be used as a transformation or an aggregation. A transformation applies the function to each metric time series (MTS) individually, while an aggregation applies the function to all MTS and returns a single value. For example, to discover the peak page visits for a site over the last day, you can use the following SignalFlow code:

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

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

### NEW QUESTION 2

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 result<sup>1</sup>. 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 3

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

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

**Answer:** B

#### Explanation:

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

According to the web search results, the Splunk distribution of the OpenTelemetry Collector stores the configuration files on Linux machines in the /etc/otel/collector/ directory by default. You can verify this by looking at the first result<sup>1</sup>, which explains how to install the Collector for Linux manually. It also provides the locations of the default configuration file, the agent configuration file, and the gateway configuration file.

To learn more about how to install and configure the Splunk distribution of the OpenTelemetry Collector, you can refer to this documentation<sup>2</sup>.

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

### NEW QUESTION 4

Clicking a metric name from the results in metric finder displays the metric in Chart Builder. What action needs to be taken in order to save the chart created in the UI?

- A. Create a new dashboard and save the chart.
- B. Save the chart to multiple dashboards.
- C. Make sure that data is coming in for the metric then save the chart.
- D. Save the chart to a dashboard.

**Answer:** D

#### Explanation:

According to the web search results, clicking a metric name from the results in metric finder displays the metric in Chart Builder<sup>1</sup>. Chart Builder is a tool that allows you to create and customize charts using metrics, dimensions, and analytics functions<sup>2</sup>. To save the chart created in the UI, you need to do the following steps:

? Click the Save button on the top right corner of the Chart Builder. This will open a

dialog box where you can enter the chart name and description, and choose the dashboard where you want to save the chart.

? Enter a name and a description for your chart. The name should be descriptive and unique, and the description should explain the purpose and meaning of the chart.

? Choose an existing dashboard from the drop-down menu, or create a new dashboard by clicking the + icon. A dashboard is a collection of charts that display metrics and events for your services or hosts<sup>3</sup>. You can organize and share dashboards with other users in your organization using dashboard groups<sup>3</sup>.

? Click Save. This will save your chart to the selected dashboard and redirect you to the dashboard view. You can also access your saved chart from the Dashboards menu on the left navigation bar.

### NEW QUESTION 5

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 preserved<sup>1</sup>

? 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 preserved<sup>2</sup>

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 documentations<sup>12</sup>.

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 6**

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 timestamp<sup>1</sup>

**NEW QUESTION 7**

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 latency<sup>1</sup>

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%<sup>2</sup> 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 documentations<sup>34</sup>.

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 8**

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 document<sup>1</sup>, 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 document<sup>2</sup> 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 9**

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

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

**Answer:** ACD

**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 workflow<sup>1</sup>

? 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 templates<sup>2</sup>

? 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 email<sup>3</sup>

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>

**NEW QUESTION 10**

Which of the following statements about adding properties to MTS are true? (select all that apply)

- A. Properties can be set via the API.
- B. Properties are sent in with datapoints.
- C. Properties are applied to dimension key:value pairs and propagated to all MTS with that dimension
- D. Properties can be set in the UI under Metric Metadata.

**Answer:** AD

**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 Cloud<sup>1</sup>. 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 requests<sup>2</sup>. The UI allows you to interactively create, edit, and delete properties for dimensions using the Metric Metadata page under Settings<sup>3</sup>. Therefore, option A and D are correct.

**NEW QUESTION 10**

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 documentation<sup>1</sup>, 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 14

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

- A. To determine the root cause of the Issue triggering the detector.
- B. To perform transformations on the data used by the detector.
- C. To receive an email notification when a detector is triggered.
- D. To be able to modify the alert parameters.

**Answer: C**

#### Explanation:

One reason a user of Splunk Observability Cloud would want to subscribe to an alert is C. To receive an email notification when a detector is triggered.

A detector is a component of Splunk Observability Cloud that monitors metrics or events and triggers alerts when certain conditions are met. A user can create and configure detectors to suit their monitoring needs and goals<sup>1</sup>

A subscription is a way for a user to receive notifications when a detector triggers an alert. A user can subscribe to a detector by entering their email address in the Subscription tab of

the detector page. A user can also unsubscribe from a detector at any time<sup>2</sup>

When a user subscribes to an alert, they will receive an email notification that contains information about the alert, such as the detector name, the alert status, the alert severity, the alert time, and the alert message. The email notification also includes links to view the detector, acknowledge the alert, or unsubscribe from the detector<sup>2</sup>

To learn more about how to use detectors and subscriptions in Splunk Observability Cloud, you can refer to these documentations<sup>1,2</sup>.

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

#### NEW QUESTION 19

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 document<sup>1</sup>, 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 document<sup>2</sup> 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 20

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 it<sup>1</sup>

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 Support<sup>2</sup>

To learn more about how to use filters in Splunk Observability Cloud, you can refer to this documentation<sup>3</sup>.

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 25

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 characters<sup>1</sup>

To learn more about how to filter metrics in Splunk Observability Cloud, you can refer to this documentation<sup>2</sup>.

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

#### NEW QUESTION 28

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 points<sup>1</sup>

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 data<sup>1</sup>

To learn more about how to use Max Delay in Splunk Observability Cloud, you can refer to this documentation<sup>1</sup>.

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

#### NEW QUESTION 29

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 document<sup>1</sup>, 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 34

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

- A. Calculates the number of MTS present in the plot.
- B. Calculates 1/2 of the values present in the input time series.
- C. Calculates the sum of values present in the input time series across the entire environment or per group.
- D. Calculates the sum of values per time series across a period of time.

**Answer:** C

#### Explanation:

According to the Splunk Test Blueprint - O11y Cloud Metrics User document<sup>1</sup>, 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 document<sup>2</sup> 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
```

#### NEW QUESTION 37

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 MeanTransformatio
- 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 documentation<sup>1</sup>.

1: <https://docs.splunk.com/Observability/gdi/metrics/analytics.html>

**NEW QUESTION 38**

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)<sup>1</sup>

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 name<sup>1</sup>

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 changed<sup>2</sup>

To learn more about how to use metadata in Splunk Observability Cloud, you can refer to this documentation<sup>2</sup>.

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 41**

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 documentation<sup>1</sup>, 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 45**

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 documentation<sup>1</sup>, 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 48**

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 document<sup>1</sup>, 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 51**

.....

## Relate Links

**100% Pass Your SPLK-4001 Exam with ExamBible Prep Materials**

<https://www.exambible.com/SPLK-4001-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>