SPLK-4001 Dumps

Splunk O11y Cloud Certified Metrics User

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

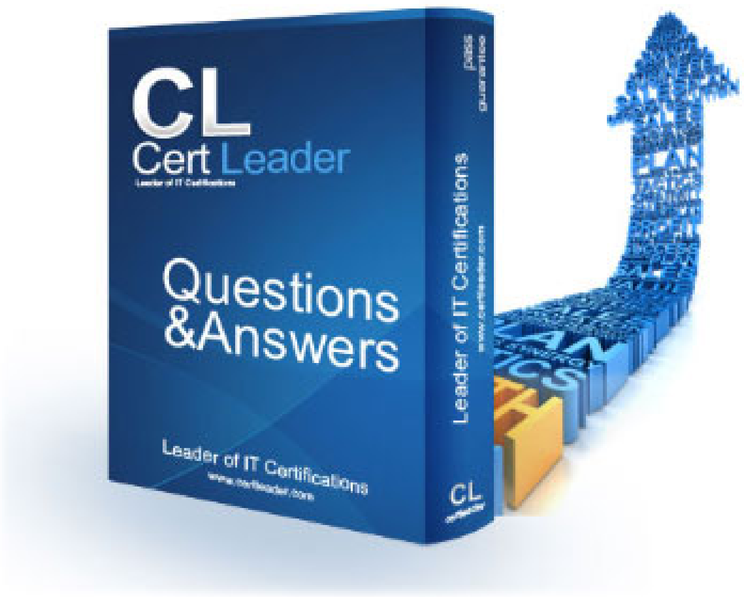

NEW QUESTION 1
Which analytic function can be used to discover peak page visits for a site over the last day?

A. Maximum: Transformation (24h)
B. Maximum: Aggregation (Id)
C. Lag: (24h)
D. Count: (Id)

Answer: A

Explanation:
According to the Splunk Observability Cloud documentation1, the maximum function is an analytic function that returns the highest value of a metric or a dimension over a specified time interval. The maximum function can be used as a transformation or an aggregation. A transformation applies the function to each metric time series (MTS) individually, while an aggregation applies the function to all MTS and returns a single value. For example, to discover the peak page visits for a site over the last day, you can use the following SignalFlow code:
maximum(24h, counters(“page.visits”))
This will return the highest value of the page.visits counter metric for each MTS over the last 24 hours. You can then use a chart to visualize the results and identify the peak page visits for each MTS.

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

A. Adding debug into the metrics receiver pipeline:

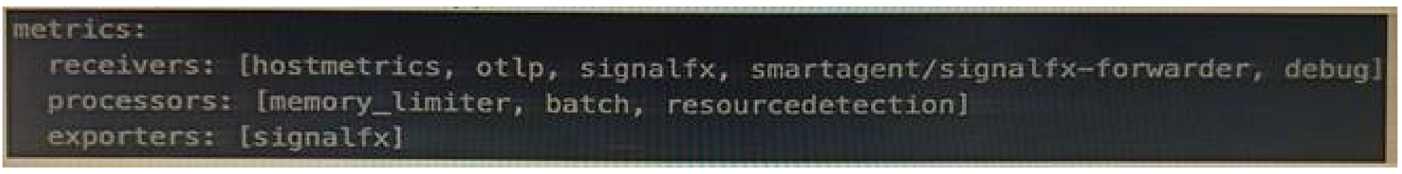

```
metrics:
  receivers: [hostmetrics, otlp, signalfx, smartagent/signalfx-forwarder, debug]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx]
```

B. Adding logging into the metrics receiver pipeline:

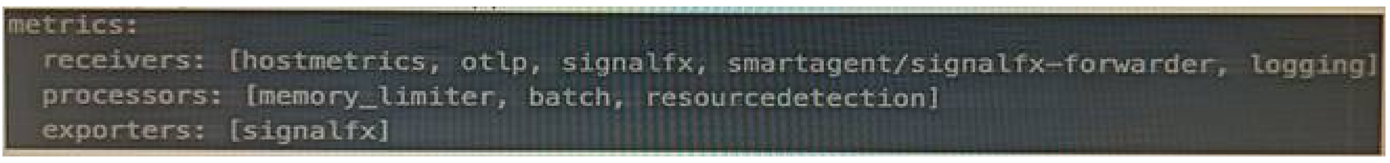

```
metrics:
  receivers: [hostmetrics, otlp, signalfx, smartagent/signalfx-forwarder, logging]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx]
```

C. Adding logging into the metrics exporter pipeline:

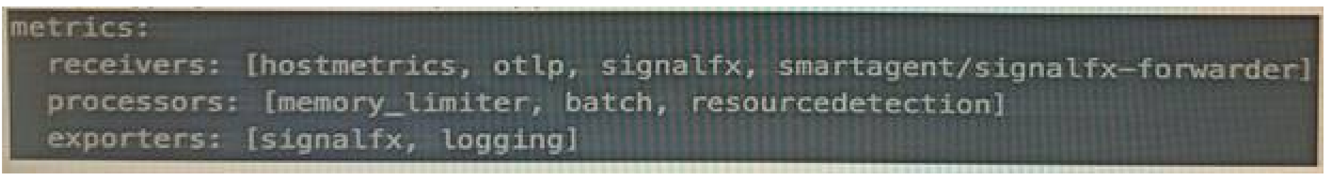

```
metrics:
  receivers: [hostmetrics, otlp, signalfx, smartagent/signalfx-forwarder]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx, logging]
```

D. Adding debug into the metrics exporter pipeline:

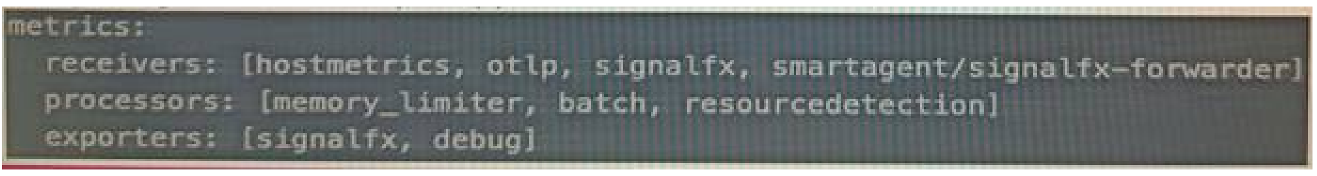

```
metrics:
  receivers: [hostmetrics, otlp, signalfx, smartagent/signalfx-forwarder]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx, debug]
```

Answer: B

Explanation:
The correct answer is B. Adding logging into the metrics receiver pipeline. The logging exporter is a component that allows the OpenTelemetry Collector to send traces, metrics, and logs directly to the console. It can be used to diagnose and troubleshoot issues with telemetry received and processed by the Collector, or to obtain samples for other purposes1
To activate the logging exporter, you need to add it to the pipeline that you want to diagnose. In this case, since you are experiencing issues with a new receiver for metrics, you need to add the logging exporter to the metrics receiver pipeline. This will create a new plot that shows the metrics received by the Collector and any errors or warnings that might occur1

The image that you have sent with your question shows how to add the logging exporter to the metrics receiver pipeline. You can see that the exporters section of the metrics pipeline includes logging as one of the options. This means that the metrics received by any of the receivers listed in the receivers section will be sent to the logging exporter as well as to any other exporters listed2
To learn more about how to use the logging exporter in Splunk Observability Cloud, you can refer to this documentation1.
1: https://docs.splunk.com/Observability/gdi/opentelemetry/components/logging-
exporter.html 2: https://docs.splunk.com/Observability/gdi/opentelemetry/exposed- endpoints.html

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
For which types of charts can individual plot visualization be set?

A. Line, Bar, Column
B. Bar, Area, Column
C. Line, Area, Column
D. Histogram, Line, Column

Answer: C

Explanation:
The correct answer is C. Line, Area, Column.
For line, area, and column charts, you can set the individual plot visualization to change the appearance of each plot in the chart. For example, you can change the color, shape, size, or style of the lines, areas, or columns. You can also change the rollup function, data resolution, or y-axis scale for each plot1
To set the individual plot visualization for line, area, and column charts, you need to select the chart from the Metric Finder, then click on Plot Chart Options and choose Individual Plot Visualization from the list of options. You can then customize each plot according to your preferences2
To learn more about how to use individual plot visualization in Splunk Observability Cloud, you can refer to this documentation2.
1: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Individual-plot-visualization
2: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Set-individual-plot- visualization

NEW QUESTION 6
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 7
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 8
Which of the following can be configured when subscribing to a built-in detector?

A. Alerts on team landing page.
B. Alerts on a dashboard.
C. Outbound notifications.
D. Links to a chart.

Answer: C

Explanation:
According to the web search results1, subscribing to a built-in detector is a way to receive alerts and notifications from Splunk Observability Cloud when certain criteria are met. A built-in detector is a detector that is automatically created and configured by Splunk Observability Cloud based on the data from your integrations, such as AWS, Kubernetes, or OpenTelemetry1. To subscribe to a built-in detector, you need to do the following steps:
? Find the built-in detector that you want to subscribe to. You can use the metric finder or the dashboard groups to locate the built-in detectors that are relevant to your data sources1.
? Hover over the built-in detector and click the Subscribe button. This will open a dialog box where you can configure your subscription settings1.
? Choose an outbound notification channel from the drop-down menu. This is where you can specify how you want to receive the alert notifications from the built-in detector. You can choose from various channels, such as email, Slack, PagerDuty, webhook, and so on2. You can also create a new notification channel by clicking the + icon2.
? Enter the notification details for the selected channel. This may include your email address, Slack channel name, PagerDuty service key, webhook URL, and so on2. You can also customize the notification message with variables and markdown formatting2.
? Click Save. This will subscribe you to the built-in detector and send you alert notifications through the chosen channel when the detector triggers or clears an alert.
Therefore, option C is correct.

NEW QUESTION 9
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 10
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 workflow1
? Send an SMS message. This option allows you to send a text message to one or more phone numbers when an alert is triggered or cleared. You can customize the message content and format using variables and templates2
? Send to email addresses. This option allows you to send an email notification to one or more recipients when an alert is triggered or cleared. You can customize the email subject, body, and attachments using variables and templates. You can also include information from search results, the search job, and alert triggering in the email3
Therefore, the correct answer is A, C, and D.
1: https://docs.splunk.com/Documentation/Splunk/latest/Alert/Webhooks 2:

https://docs.splunk.com/Documentation/Splunk/latest/Alert/SMSnotification 3: https://docs.splunk.com/Documentation/Splunk/latest/Alert/Emailnotification

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 Cloud1. Properties provide additional context and information about the metrics, such as the environment, role, or owner of the dimension. For example, you can add the property use: QA to the host dimension of your metrics to indicate that the host that is sending the data is used for QA. To add properties to MTS, you can use either the API or the UI. The API allows you to programmatically create, update, delete, and list properties for dimensions using HTTP requests2. The UI allows you to interactively create, edit, and delete properties for dimensions using the Metric Metadata page under Settings3. Therefore, option A and D are correct.

NEW QUESTION 11
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 goals1
A subscription is a way for a user to receive notifications when a detector triggers an alert. A user can subscribe to a detector by entering their email address in the Subscription tab of
the detector page. A user can also unsubscribe from a detector at any time2
When a user subscribes to an alert, they will receive an email notification that contains information about the alert, such as the detector name, the alert status, the alert severity, the alert time, and the alert message. The email notification also includes links to view the detector, acknowledge the alert, or unsubscribe from the detector2
To learn more about how to use detectors and subscriptions in Splunk Observability Cloud, you can refer to these documentations12.
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 16
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 17
An SRE creates an event feed chart in a dashboard that shows a list of events that meet criteria they specify. Which of the following should they include? (select all

that apply)

A. Custom events that have been sent in from an external source.
B. Events created when a detector clears an alert.
C. Random alerts from active detectors.
D. Events created when a detector triggers an alert.

Answer: ABD

Explanation:
According to the web search results1, an event feed chart is a type of chart that shows a list of events that meet criteria you specify. An event feed chart can display one or more event types depending on how you specify the criteria. The event types that you can include in an event feed chart are:
? Custom events that have been sent in from an external source: These are events that you have created or received from a third-party service or tool, such as AWS CloudWatch, GitHub, Jenkins, or PagerDuty. You can send custom events to Splunk Observability Cloud using the API or the Event Ingest Service.
? Events created when a detector triggers or clears an alert: These are events that are automatically generated by Splunk Observability Cloud when a detector evaluates a metric or dimension and finds that it meets the alert condition or returns to normal. You can create detectors to monitor and alert on various metrics and dimensions using the UI or the API.
Therefore, option A, B, and D are correct.

NEW QUESTION 19
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 20
Which of the following aggregate analytic functions will allow a user to see the highest or lowest n values of a metric?

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

Answer: D

Explanation:
The correct answer is D. Top / Bottom.
Top and bottom are aggregate analytic functions that allow a user to see the highest or lowest n values of a metric. They can be used to select a subset of the time series in the plot by count or by percent. For example, top (5) will show the five time series with the highest values in each time period, while bottom (10%) will show the 10% of time series with the lowest values in each time period1
To learn more about how to use top and bottom functions in Splunk Observability Cloud, you can refer to this documentation1.

NEW QUESTION 24
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 document1, one of the metrics concepts that is covered in the exam is analytic functions. Analytic functions are mathematical operations that can be applied to metrics to transform, aggregate, or analyze them.
The Splunk O11y Cloud Certified Metrics User Track document2 states that one of the recommended courses for preparing for the exam is Introduction to Splunk Infrastructure Monitoring, which covers the basics of metrics monitoring and visualization.
In the Introduction to Splunk Infrastructure Monitoring course, there is a section on Analytic Functions, which explains that analytic functions can be used to perform calculations on metrics, such as sum, average, min, max, count, etc. The document also provides examples of how to use analytic functions in charts and dashboards.
One of the analytic functions that can be used is Sum Aggregation, which calculates the
sum of values present in the input time series across the entire environment or per group. The document gives an example of how to use Sum Aggregation to calculate the total CPU usage across all hosts in a group by using the following syntax:
sum(cpu.utilization) by hostgroup

NEW QUESTION 26
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 27
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. Tlmeshift
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 29
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 32
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 33
......