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

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

Answer: D

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

1: https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Metric-types 2: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Data-resolution-and-rollups- in-charts 3: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Rollup-functions- for-metric-types

NEW QUESTION 4
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 5
For a high-resolution metric, what is the highest possible native resolution of the metric?

A. 2 seconds
B. 15 seconds
C. 1 second
D. 5 seconds

Answer: C

Explanation:
The correct answer is C. 1 second.
According to the Splunk Test Blueprint - O11y Cloud Metrics User document1, one of the metrics concepts that is covered in the exam is data resolution and rollups. Data resolution refers to the granularity of the metric data points, and rollups are the process of aggregating data points over time to reduce the amount of data stored.
The Splunk O11y Cloud Certified Metrics User Track document2 states that one of the recommended courses for preparing for the exam is Introduction to Splunk Infrastructure Monitoring, which covers the basics of metrics monitoring and visualization.
In the Introduction to Splunk Infrastructure Monitoring course, there is a section on Data Resolution and Rollups, which explains that Splunk Observability Cloud collects high- resolution metrics at 1-second intervals by default, and then applies rollups to reduce the data volume over time. The document also provides a table that shows the different rollup intervals and retention periods for different resolutions.
Therefore, based on these documents, we can conclude that for a high-resolution metric, the highest possible native resolution of the metric is 1 second.

NEW QUESTION 6
What is the limit on the number of properties that an MTS can have?

A. 64
B. 36
C. No limit
D. 50

Answer: A

Explanation:
The correct answer is A. 64.
According to the web search results, the limit on the number of properties that an MTS can have is 64. A property is a key-value pair that you can assign to a dimension of an existing MTS to add more context to the metrics. For example, you can add the property use: QA to the host dimension of your metrics to indicate that the host is used for QA1
Properties are different from dimensions, which are key-value pairs that are sent along with the metrics at the time of ingest. Dimensions, along with the metric name, uniquely identify an MTS. The limit on the number of dimensions per MTS is 362
To learn more about how to use properties and dimensions in Splunk Observability Cloud, you can refer to this documentation2.
1: https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html#Custom-properties 2: https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html

NEW QUESTION 7
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 8
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 9
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 10
What are the best practices for creating detectors? (select all that apply)

A. View data at highest resolution.
B. Have a consistent value.
C. View detector in a chart.
D. Have a consistent type of measurement.

Answer: ABCD

Explanation:
The best practices for creating detectors are:
? View data at highest resolution. This helps to avoid missing important signals or patterns in the data that could indicate anomalies or issues1
? Have a consistent value. This means that the metric or dimension used for detection should have a clear and stable meaning across different sources, contexts, and time periods. For example, avoid using metrics that are affected by changes in configuration, sampling, or aggregation2
? View detector in a chart. This helps to visualize the data and the detector logic, as
well as to identify any false positives or negatives. It also allows to adjust the detector parameters and thresholds based on the data distribution and behavior3
? Have a consistent type of measurement. This means that the metric or dimension
used for detection should have the same unit and scale across different sources, contexts, and time periods. For example, avoid mixing bytes and bits, or seconds and milliseconds.
1: https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-
detectors 2: https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-
practices-for-detectors 3: https://docs.splunk.com/Observability/gdi/metrics/detectors.html#View-detector-in-a-chart :
https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for- detectors

NEW QUESTION 10
Which of the following are required in the configuration of a data point? (select all that apply)

A. Metric Name
B. Metric Type
C. Timestamp
D. Value

Answer: ACD

Explanation:
The required components in the configuration of a data point are:
? Metric Name: A metric name is a string that identifies the type of measurement that the data point represents, such as cpu.utilization, memory.usage, or response.time. A metric name is mandatory for every data point, and it must be unique within a Splunk Observability Cloud organization1
? Timestamp: A timestamp is a numerical value that indicates the time at which the data point was collected or generated. A timestamp is mandatory for every data point, and it must be in epoch time format, which is the number of seconds since January 1, 1970 UTC1
? Value: A value is a numerical value that indicates the magnitude or quantity of the
measurement that the data point represents. A value is mandatory for every data point, and it must be compatible with the metric type of the data point1
Therefore, the correct answer is A, C, and D.
To learn more about how to configure data points in Splunk Observability Cloud, you can refer to this documentation1.
1: https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Data-points

NEW QUESTION 13
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 16
An SRE came across an existing detector that is a good starting point for a detector they want to create. They clone the detector, update the metric, and add multiple new signals. As a result of the cloned detector, which of the following is true?

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

Answer: D

Explanation:
According to the Splunk O11y Cloud Certified Metrics User Track document1, cloning a detector creates a copy of the detector that you can modify without affecting the original detector. You can change the metric, filter, and signal settings of the cloned detector.
However, the new signals that you add to the cloned detector will not be reflected in the original detector, nor in the original chart that the detector was based on.
Therefore, option D is correct.
Option A is incorrect because the new signals will not be reflected in the original detector. Option B is incorrect because the new signals will not be reflected in the original chart. Option C is incorrect because you can monitor all of the new signals that you add to the cloned detector.

NEW QUESTION 17
.....

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/