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
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 2
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 3
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 4
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 5
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 6
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 7
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 8
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 9
A customer operates a caching web proxy. They want to calculate the cache hit rate for their service. What is the best way to achieve this?

A. Percentages and ratios
B. Timeshift and Bottom N
C. Timeshift and Top N
D. Chart Options and metadata

Answer: A

Explanation:
According to the Splunk O11y Cloud Certified Metrics User Track document1, percentages and ratios are useful for calculating the proportion of one metric to another, such as cache hits to cache misses, or successful requests to failed requests. You can use the percentage() or ratio() functions in SignalFlow to compute these values and display them in charts. For example, to calculate the cache hit rate for a service, you can use the following SignalFlow code:
percentage(counters(“cache.hits”), counters(“cache.misses”))
This will return the percentage of cache hits out of the total number of cache attempts. You can also use the ratio() function to get the same result, but as a decimal value instead of a percentage.
ratio(counters(“cache.hits”), counters(“cache.misses”))

NEW QUESTION 10
To smooth a very spiky cpu.utilization metric, what is the correct analytic function to better see if the cpu. utilization for servers is trending up over time?

A. Rate/Sec
B. Median
C. Mean (by host)
D. Mean (Transformation)

Answer: D

Explanation:
The correct answer is D. Mean (Transformation).
According to the web search results, a mean transformation is an analytic function that returns the average value of a metric or a dimension over a specified time interval1. A mean transformation can be used to smooth a very spiky metric, such as cpu.utilization, by reducing the impact of outliers and noise. A mean transformation can also help to see if the metric is trending up or down over time, by showing the general direction of the average value. For example, to smooth the cpu.utilization metric and see if it is trending up over time, you can use the following SignalFlow code:
mean(1h, counters(“cpu.utilization”))
This will return the average value of the cpu.utilization counter metric for each metric time series (MTS) over the last hour. You can then use a chart to visualize the results and compare the mean values across different MTS.
Option A is incorrect because rate/sec is not an analytic function, but rather a rollup function that returns the rate of change of data points in the MTS reporting interval1. Rate/sec can be used to convert cumulative counter metrics into counter metrics, but it does not smooth or trend a metric. Option B is incorrect because median is not an analytic function, but rather an aggregation function that returns the middle value of a metric or a dimension over the entire time range1. Median can be used to find the typical value of a metric, but it does not smooth or trend a metric. Option C is incorrect because mean (by host) is not an analytic function,

but rather an aggregation function that returns the average value of a metric or a dimension across all MTS with the same host dimension1. Mean (by host) can be used to compare the performance of different hosts, but it does not smooth or trend a metric.
Mean (Transformation) is an analytic function that allows you to smooth a very spiky metric by applying a moving average over a specified time window. This can help you see the general trend of the metric over time, without being distracted by the short-term fluctuations1
To use Mean (Transformation) on a cpu.utilization metric, you need to select the metric from the Metric Finder, then click on Add Analytics and choose Mean (Transformation) from the list of functions. You can then specify the time window for the moving average, such as 5 minutes, 15 minutes, or 1 hour. You can also group the metric by host or any other dimension to compare the smoothed values across different servers2
To learn more about how to use Mean (Transformation) and other analytic functions in Splunk Observability Cloud, you can refer to this documentation2.
1: https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Mean-Transformation 2: https://docs.splunk.com/Observability/gdi/metrics/analytics.html

NEW QUESTION 10
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 12
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 17
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 18
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 23
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-l'.
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 28
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 30
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 34
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 39
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 42
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 44
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 49
.....

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/