SPLK-4001 Q&As

Splunk O11y Cloud Certified Metrics User

QUESTION 1

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

Correct Answer: C

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.

QUESTION 2

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

Correct Answer: A

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.

QUESTION 3

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.

Correct Answer: AD

The correct answers are A and D.

According to the Splunk Test Blueprint - O11y Cloud Metrics User document, 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 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.

QUESTION 4

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.

Correct Answer: C

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 sources.

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 on. You can also create a new notification channel by clicking the +

icon.

Enter the notification details for the selected channel. This may include your email address, Slack channel name, PagerDuty service key, webhook URL, and so on. You can also customize the notification message with variables and

markdown formatting.

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.

QUESTION 5

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 MeanrTransformation. In the window that appears, select \\'version\\' from the Group By field.

B. On the chart for plot A, scroll to the end and click Enter Function, then enter \\'A/B-l\\'.

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.

D. On the chart for plot A, click the Compare Means button. In the window that appears, type \\'version1.

Correct Answer: C

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

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.

Correct Answer: B

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:

MTS: Gauge metric cpu.utilization, dimension "hostname": "host" MTS: Gauge metric cpu.utilization, dimension "hostname": "host" MTS: Gauge metric memory.usage, dimension "hostname": "host"

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

QUESTION 7

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)

Correct Answer: D

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 interval. 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 fluctuations 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 documentation. https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Mean-Transformation https://docs.splunk.com/Observability/gdi/metrics/analytics.html

QUESTION 8

The built-in Kubernetes Navigator includes which of the following?

A. Map, Nodes, Workloads, Node Detail, Workload Detail, Group Detail, Container Detail

B. Map, Nodes, Processors, Node Detail, Workload Detail, Pod Detail, Container Detail

C. Map, Clusters, Workloads, Node Detail, Workload Detail, Pod Detail, Container Detail

D. Map, Nodes, Workloads, Node Detail, Workload Detail, Pod Detail, Container Detail

Correct Answer: D

The correct answer is D. Map, Nodes, Workloads, Node Detail, Workload Detail, Pod Detail, Container Detail. The built-in Kubernetes Navigator is a feature of Splunk Observability Cloud that provides a comprehensive and intuitive way to monitor the performance and health of Kubernetes environments. It includes the following views: Map: A graphical representation of the Kubernetes cluster topology, showing the relationships and dependencies among nodes, pods, containers, and services. You can use the map to quickly identify and troubleshoot issues in your cluster Nodes: A tabular view of all the nodes in your cluster, showing key metrics such as CPU utilization, memory usage, disk usage,

and network traffic. You can use the nodes view to compare and analyze the performance of different nodes1 Workloads: A tabular view of all the workloads in your cluster, showing key metrics such as CPU utilization, memory usage, network traffic, and error rate. You can use the workloads view to compare and analyze the performance of different workloads, such as deployments, stateful sets, daemon sets, or jobs1 Node Detail: A detailed view of a specific node in your cluster, showing key metrics and charts for CPU utilization, memory usage, disk usage, network traffic, and pod count. You can also see the list of pods running on the node and their status. You can use the node detail view to drill down into the performance of a single node Workload Detail: A detailed view of a specific workload in your cluster, showing key metrics and charts for CPU utilization, memory usage, network traffic, error rate, and pod count. You can also see the list of pods belonging to the workload and their status. You can use the workload detail view to drill down into the performance of a single workload Pod Detail: A detailed view of a specific pod in your cluster, showing key metrics and charts for CPU utilization, memory usage, network traffic, error rate, and container count. You can also see the list of containers within the pod and their status. You can use the pod detail view to drill down into the performance of a single pod Container Detail: A detailed view of a specific container in your cluster, showing key metrics and charts for CPU utilization, memory usage, network traffic, error rate, and log events. You can use the container detail view to drill down into the performance of a single container To learn more about how to use Kubernetes Navigator in Splunk Observability Cloud, you can refer to this documentation.
https://docs.splunk.com/observability/infrastructure/monitor/k8s-nav.html#Kubernetes- Navigator:
https://docs.splunk.com/observability/infrastructure/monitor/k8s- nav.html#Detail-pages
https://docs.splunk.com/observability/infrastructure/monitor/k8s- nav.html

QUESTION 9

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

A. Dimensions

B. Properties

C. Sources

D. Tags

Correct Answer: A

The correct answer is A. Dimensions. Dimensions are metadata in the form of key-value pairs that are sent along with the metrics at the time of ingest. They provide additional information about the metric, such as the name of the host that sent the metric, or the location of the server. Along with the metric name, they uniquely identify a metric time series (MTS)1 Changes to dimensions result in a new MTS, because they create a different combination of metric name and dimensions. For example, if you change the hostname dimension from host1 to host, you will create a new MTS for the same metric name1 Properties, sources, and tags are other types of metadata that can be applied to existing MTSes after ingest. They do not contribute to uniquely identify an MTS, and they do not create a new MTS when changed To learn more about how to use metadata in Splunk Observability Cloud, you can refer to this documentation.
https://docs.splunk.com/Observability/metrics-and-metadata/metrics.html#Dimensions
https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html

QUESTION 10

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.

Correct Answer: C

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 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 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 To learn more about how to use detectors and subscriptions in Splunk Observability Cloud, you can refer to these documentations. https://docs.splunk.com/Observability/alerts-detectors-notifications/detectors.html https://docs.splunk.com/Observability/alerts-detectors-notifications/subscribe-to-detectors.html