Number: SPLK-4001
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: SPLK-4001

Exam Name: Splunk O11y Cloud Certified Metrics User Exam

Exam A

QUESTION 1
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
Section:
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.

QUESTION 2
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.

Correct Answer: A
Section:
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

QUESTION 3
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

Correct Answer: A
Section:

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''))

QUESTION 4
Which of the following are correct ports for the specified components in the OpenTelemetry Collector?

A. gRPC (4000), SignalFx (9943), Fluentd (6060)
B. gRPC (6831), SignalFx (4317), Fluentd (9080)
C. gRPC (4459), SignalFx (9166), Fluentd (8956)
D. gRPC (4317), SignalFx (9080), Fluentd (8006)

Correct Answer: D
Section:
Explanation:
The correct answer is D. gRPC (4317), SignalFx (9080), Fluentd (8006).
According to the web search results, these are the default ports for the corresponding components in the OpenTelemetry Collector. You can verify this by looking at the table of exposed ports and endpoints in the first result1. You can also see the agent and gateway configuration files in the same result for more details.
1: https://docs.splunk.com/observability/gdi/opentelemetry/exposed-endpoints.html

QUESTION 5
When writing a detector with a large number of MTS, such as memory. free in a deployment with 30,000 hosts, it is possible to exceed the cap of MTS that can be contained in a single plot. Which of the choices below would most likely reduce the number of MTS below the plot cap?

A. Select the Sharded option when creating the plot.
B. Add a filter to narrow the scope of the measurement.
C. Add a restricted scope adjustment to the plot.
D. When creating the plot, add a discriminator.

Correct Answer: B
Section:
Explanation:
The correct answer is B. Add a filter to narrow the scope of the measurement.
A filter is a way to reduce the number of metric time series (MTS) that are displayed on a chart or used in a detector. A filter specifies one or more dimensions and values that the MTS must have in order to be included. For example, if you want to monitor the memory.free metric only for hosts that belong to a certain cluster, you can add a filter like cluster:my-cluster to the plot or detector. This will exclude any MTS that do not have the cluster dimension or have a different value for it1
Adding a filter can help you avoid exceeding the plot cap, which is the maximum number of MTS that can be contained in a single plot. The plot cap is 100,000 by default, but it can be changed by contacting Splunk Support2
To learn more about how to use filters in Splunk Observability Cloud, you can refer to this documentation3.
1: https://docs.splunk.com/Observability/gdi/metrics/search.html#Filter-metrics 2: https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Plot-cap 3: https://docs.splunk.com/Observability/gdi/metrics/search.html

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 sensitivity. Duration set to 1 minute.
C. Adjust the notification sensitivity. Duration set to 1 minute.
D. Choose another signal.

Correct Answer: B
Section:
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.

QUESTION 7
Where does the Splunk distribution of the OpenTelemetry Collector store the configuration files on Linux machines by default?

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

Correct Answer: B
Section:
Explanation:
The correct answer is B. /etc/otel/collector/
According to the web search results, the Splunk distribution of the OpenTelemetry Collector stores the configuration files on Linux machines in the /etc/otel/collector/ directory by default. You can verify this by looking at the first result1, which explains how to install the Collector for Linux manually. It also provides the locations of the default configuration file, the agent configuration file, and the gateway configuration file.
To learn more about how to install and configure the Splunk distribution of the OpenTelemetry Collector, you can refer to this documentation2.
1: https://docs.splunk.com/Observability/gdi/opentelemetry/install-linux-manual.html 2: https://docs.splunk.com/Observability/gdi/opentelemetry.html

QUESTION 8
Which of the following rollups will display the time delta between a datapoint being sent and a datapoint being received?

A. Jitter
B. Delay
C. Lag
D. Latency

Correct Answer: C
Section:
Explanation:

According to the Splunk Observability Cloud documentation1, lag is a rollup function that returns the difference between the most recent and the previous data point values seen in the metric time series reporting interval. This can be used to measure the time delta between a data point being sent and a data point being received, as long as the data points have timestamps that reflect their send and receive times. For example, if a data point is sent at 10:00:00 and received at 10:00:05, the lag value for that data point is 5 seconds.

QUESTION 9
Which of the following is optional, but highly recommended to include in a datapoint?

A. Metric name
B. Timestamp
C. Value
D. Metric type

Correct Answer: D
Section:
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

QUESTION 10
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)

Correct Answer: A
Section:
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.

QUESTION 11
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]
  processors: [memory_limiter, batch, resourcedetection]
  exporters: [signalfx, debug]
```

B. Adding logging into the metrics receiver pipeline:

C. Adding logging into the metrics exporter pipeline:

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

D. Adding debug into the metrics exporter pipeline:

Correct Answer: B
Section:
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

QUESTION 12
What information is needed to create a detector?

A. Alert Status, Alert Criteria, Alert Settings, Alert Message, Alert Recipients
B. Alert Signal, Alert Criteria, Alert Settings, Alert Message, Alert Recipients
C. Alert Signal, Alert Condition, Alert Settings, Alert Message, Alert Recipients
D. Alert Status, Alert Condition, Alert Settings, Alert Meaning, Alert Recipients

Correct Answer: C
Section:
Explanation:

According to the Splunk Observability Cloud documentation1, to create a detector, you need the following information:
Alert Signal: This is the metric or dimension that you want to monitor and alert on. You can select a signal from a chart or a dashboard, or enter a SignalFlow query to define the signal.
Alert Condition: This is the criteria that determines when an alert is triggered or cleared. You can choose from various built-in alert conditions, such as static threshold, dynamic threshold, outlier, missing data, and so on. You can also specify the severity level and the trigger sensitivity for each alert condition.
Alert Settings: This is the configuration that determines how the detector behaves and interacts with other detectors. You can set the detector name, description, resolution, run lag, max delay, and detector rules. You can also enable or disable the detector, and mute or unmute the alerts.
Alert Message: This is the text that appears in the alert notification and event feed. You can customize the alert message with variables, such as signal name, value, condition, severity, and so on. You can also use markdown formatting to enhance the message appearance.
Alert Recipients: This is the list of destinations where you want to send the alert notifications. You can choose from various channels, such as email, Slack, PagerDuty, webhook, and so on. You can also specify the notification frequency and suppression settings.

QUESTION 13
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
Section:
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.

QUESTION 14
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.

Correct Answer: A, C, D
Section:
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

QUESTION 15
With exceptions for transformations or timeshifts, at what resolution do detectors operate?

A. 10 seconds
B. The resolution of the chart
C. The resolution of the dashboard
D. Native resolution

Correct Answer: D
Section:
Explanation:
According to the Splunk Observability Cloud documentation1, detectors operate at the native resolution of the metric or dimension that they monitor, with some exceptions for transformations or timeshifts. The native resolution is the frequency at which the data points are reported by the source. For example, if a metric is reported every 10 seconds, the detector will evaluate the metric every 10 seconds. The native resolution ensures that the detector uses the most granular and accurate data available for alerting.

QUESTION 16
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.

Correct Answer: A, C, D
Section:
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

QUESTION 17
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

Correct Answer: D
Section:
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.

QUESTION 18
Which of the following are ways to reduce flapping of a detector? (select all that apply)

A. Configure a duration or percent of duration for the alert.
B. Establish a reset threshold for the detector.
C. Enable the anti-flap setting in the detector options menu.
D. Apply a smoothing transformation (like a rolling mean) to the input data for the detector.

Correct Answer: A, D
Section:
Explanation:
According to the Splunk Lantern articleResolving flapping detectors in Splunk Infrastructure Monitoring, flapping is a phenomenon where alerts fire and clear repeatedly in a short period of time, due to the signal fluctuating around the threshold value. To reduce flapping, the article suggests the following ways:
Configure a duration or percent of duration for the alert: This means that you require the signal to stay above or below the threshold for a certain amount of time or percentage of time before triggering an alert. This can help filter out noise and focus on more persistent issues.
Apply a smoothing transformation (like a rolling mean) to the input data for the detector: This means that you replace the original signal with the average of its last several values, where you can specify the window length. This can reduce the impact of a single extreme observation and make the signal less fluctuating.

QUESTION 19
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.

Correct Answer: D
Section:
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.

QUESTION 20
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
Section:
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-Transformation2: https://docs.splunk.com/Observability/gdi/metrics/analytics.html

QUESTION 21
What happens when the limit of allowed dimensions is exceeded for an MTS?

A. The additional dimensions are dropped.
B. The datapoint is averaged.
C. The datapoint is updated.
D. The datapoint is dropped.

Correct Answer: A
Section:
Explanation:
According to the web search results, dimensions are metadata in the form of key-value pairs that monitoring software sends in along with the metrics.The set of metric time series (MTS) dimensions sent during ingest is used, along with the metric name, to uniquely identify an MTS1.Splunk Observability Cloud has a limit of 36 unique dimensions per MTS2.If the limit of allowed dimensions is exceeded for an MTS, the additional dimensions are dropped and not stored or indexed by Observability Cloud2. This means that the data point is still ingested, but without the extra dimensions. Therefore, option A is correct.

QUESTION 22
Changes to which type of metadata result in a new metric time series?

A. Dimensions
B. Properties
C. Sources
D. Tags

Correct Answer: A
Section:
Explanation:
The correct answer is A. Dimensions.
Dimensions are metadata in the form of key-value pairs that are sent along with the metrics at the time of ingest. They provide additional information about the metric, such as the name of the host that sent the metric, or the location of the server. Along with the metric name, they uniquely identify a metric time series (MTS)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 host2, 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 changed2
To learn more about how to use metadata in Splunk Observability Cloud, you can refer to this documentation2.
1: https://docs.splunk.com/Observability/metrics-and-metadata/metrics.html#Dimensions 2: https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html

QUESTION 23
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
Section:
Explanation:
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 cluster1
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 node2
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 workload2
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 pod2
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 container2
To learn more about how to use Kubernetes Navigator in Splunk Observability Cloud, you can refer to this documentation3.
1: https://docs.splunk.com/observability/infrastructure/monitor/k8s-nav.html#Kubernetes-Navigator 2: https://docs.splunk.com/observability/infrastructure/monitor/k8s-nav.html#Detail-pages 3: https://docs.splunk.com/observability/infrastructure/monitor/k8s-nav.html

QUESTION 24
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 detector. Select Alert settings, then select Auto-Clear Alerts and enter an appropriate time period.
B. Create the detector. Select Alert settings, then select Ephemeral Infrastructure and enter the expected lifetime of an instance.
C. Check the Dynamic checkbox when creating the detector.
D. Check the Ephemeral checkbox when creating the detector.

Correct Answer: B
Section:
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.

QUESTION 25
A customer wants to share a collection of charts with their entire SRE organization. What feature of Splunk Observability Cloud makes this possible?

A. Dashboard groups
B. Shared charts
C. Public dashboards
D. Chart exporter

Correct Answer: A
Section:
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 create dashboard groups based on different criteria, such as service, team, role, or topic. You can also set permissions for each dashboard group, such as who can view, edit, or manage the dashboards in the group. Dashboard groups make it possible to share a collection of charts with your entire SRE organization, or any other group of users that you want to collaborate with.

QUESTION 26
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.

Correct Answer: D
Section:
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

QUESTION 27
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
Section:
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