SPLK-4001 Dumps

Splunk O11y Cloud Certified Metrics User

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

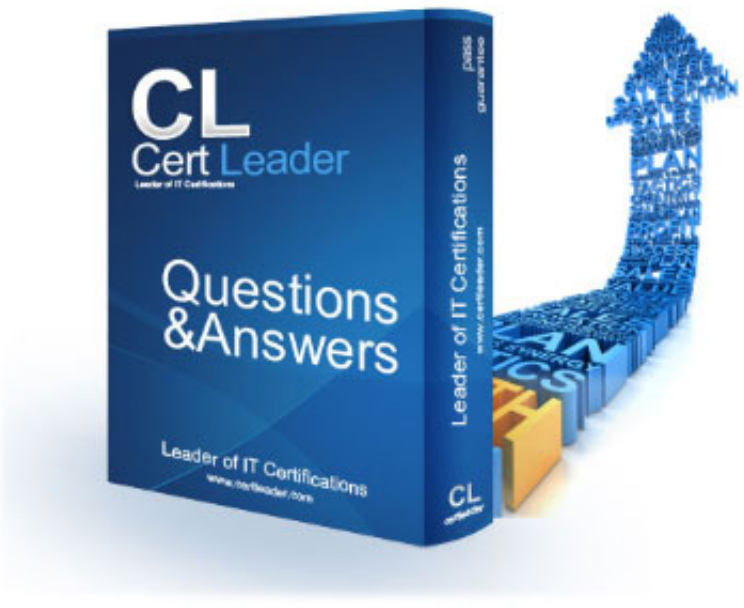

NEW QUESTION 1
Which analytic function can be used to discover peak page visits for a site over the last day?

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

Answer: A

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

NEW QUESTION 2
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)

Answer: D

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

NEW QUESTION 3
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/

Answer: B

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

NEW QUESTION 4
Clicking a metric name from the results in metric finder displays the metric in Chart Builder. What action needs to be taken in order to save the chart created in the UI?

A. Create a new dashboard and save the chart.
B. Save the chart to multiple dashboards.
C. Make sure that data is coming in for the metric then save the chart.
D. Save the chart to a dashboard.

Answer: D

Explanation:
According to the web search results, clicking a metric name from the results in metric finder displays the metric in Chart Builder1. Chart Builder is a tool that allows you to create and customize charts using metrics, dimensions, and analytics functions2. To save the chart created in the UI, you need to do the following steps:
? Click the Save button on the top right corner of the Chart Builder. This will open a
dialog box where you can enter the chart name and description, and choose the dashboard where you want to save the chart.
? Enter a name and a description for your chart. The name should be descriptive and unique, and the description should explain the purpose and meaning of the chart.
? Choose an existing dashboard from the drop-down menu, or create a new dashboard by clicking the + icon. A dashboard is a collection of charts that display metrics and events for your services or hosts3. You can organize and share dashboards with other users in your organization using dashboard groups3.
? Click Save. This will save your chart to the selected dashboard and redirect you to the dashboard view. You can also access your saved chart from the Dashboards menu on the left navigation bar.

NEW QUESTION 5

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

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

Answer: C

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

NEW QUESTION 9
How is it possible to create a dashboard group that no one else can edit?

A. Ask the admin to lock the dashboard group.
B. Restrict the write access on the dashboard group.
C. Link the dashboard group to the team.
D. Hide the edit menu on the dashboard group.

Answer: B

Explanation:
According to the web search results, dashboard groups are a feature of Splunk Observability Cloud that allows you to organize and share dashboards with other users in your organization1. You can set permissions for each dashboard group, such as who can view, edit, or manage the dashboards in the group1. To create a dashboard group that no one else can edit, you need to do the following steps:
? Create a dashboard group as usual, by selecting Dashboard Group from the
Create menu on the navigation bar, entering a name and description, and adding dashboards to the group1.
? Select Alert settings from the Dashboard actions menu () on the top right corner of the dashboard group. This will open a dialog box where you can configure the permissions for the dashboard group1.
? Under Write access, select Only me. This will restrict the write access to the
dashboard group to yourself only. No one else will be able to edit or delete the dashboards in the group1.
? Click Save. This will create a dashboard group that no one else can edit.

NEW QUESTION 10
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

Answer: D

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.

NEW 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.

Answer: C

Explanation:
One reason a user of Splunk Observability Cloud would want to subscribe to an alert is C. To receive an email notification when a detector is triggered.
A detector is a component of Splunk Observability Cloud that monitors metrics or events and triggers alerts when certain conditions are met. A user can create and configure detectors to suit their monitoring needs and goals1
A subscription is a way for a user to receive notifications when a detector triggers an alert. A user can subscribe to a detector by entering their email address in the Subscription tab of
the detector page. A user can also unsubscribe from a detector at any time2
When a user subscribes to an alert, they will receive an email notification that contains information about the alert, such as the detector name, the alert status, the alert severity, the alert time, and the alert message. The email notification also includes links to view the detector, acknowledge the alert, or unsubscribe from the detector2
To learn more about how to use detectors and subscriptions in Splunk Observability Cloud, you can refer to these documentations12.
1: https://docs.splunk.com/Observability/alerts-detectors-notifications/detectors.html 2: https://docs.splunk.com/Observability/alerts-detectors-notifications/subscribe-to- detectors.html

NEW QUESTION 15
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 20
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.

Answer: AD

Explanation:
According to the Splunk Lantern article Resolving 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.

NEW QUESTION 21
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.

Answer: B

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

NEW QUESTION 26
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 29
To refine a search for a metric a customer types host: test-*. What does this filter return?

A. Only metrics with a dimension of host and a value beginning with test-.
B. Error
C. Every metric except those with a dimension of host and a value equal to test.
D. Only metrics with a value of test- beginning with host.

Answer: A

Explanation:
The correct answer is A. Only metrics with a dimension of host and a value beginning with test-.
This filter returns the metrics that have a host dimension that matches the pattern test-. For example, test-01, test-abc, test-xyz, etc. The asterisk () is a wildcard character that can match any string of characters1
To learn more about how to filter metrics in Splunk Observability Cloud, you can refer to this documentation2.
1: https://docs.splunk.com/Observability/gdi/metrics/search.html#Filter-metrics 2: https://docs.splunk.com/Observability/gdi/metrics/search.html

NEW QUESTION 34
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 35
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

Answer: D

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

NEW QUESTION 38
Which of the following statements are true about local data links? (select all that apply)

A. Anyone with write permission for a dashboard can add local data links that appear on that dashboard.
B. Local data links can only have a Splunk Observability Cloud internal destination.
C. Only Splunk Observability Cloud administrators can create local links.
D. Local data links are available on only one dashboard.

Answer: AD

Explanation:
The correct answers are A and D.
According to the Get started with Splunk Observability Cloud document1, one of the topics that is covered in the Getting Data into Splunk Observability Cloud course is global and local data links. Data links are shortcuts that provide convenient access to related resources, such as Splunk Observability Cloud dashboards, Splunk Cloud Platform and Splunk Enterprise, custom URLs, and Kibana logs.
The document explains that there are two types of data links: global and local. Global data links are available on all dashboards and charts, while local data links are available on only one dashboard. The document also provides the following information about local data links:
? Anyone with write permission for a dashboard can add local data links that appear on that dashboard.
? Local data links can have either a Splunk Observability Cloud internal destination or an external destination, such as a custom URL or a Kibana log.
? Only Splunk Observability Cloud administrators can delete local data links. Therefore, based on this document, we can conclude that A and D are true statements about local data links. B and C are false statements because:
? B is false because local data links can have an external destination as well as an internal one.
? C is false because anyone with write permission for a dashboard can create local data links, not just administrators.

NEW QUESTION 40
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 41
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 45
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 48
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 53
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

Answer: C

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.

NEW QUESTION 57
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 62
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 63
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SPLK-4001 Exam with Our Prep Materials Via below:

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