SPLK-4001 Q&As

Splunk O11y Cloud Certified Metrics User

Pass Splunk SPLK-4001 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/splk-4001.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Splunk Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

The Sum Aggregation option for analytic functions does which of the following?

A. Calculates the number of MTS present in the plot.

B. Calculates 1/2 of the values present in the input time series.

C. Calculates the sum of values present in the input time series across the entire environment or per group.

D. Calculates the sum of values per time series across a period of time.

Correct Answer: C

According to the Splunk Test Blueprint - O11y Cloud Metrics User document1, one of the metrics concepts that is covered in the exam is analytic functions. Analytic functions are mathematical operations that can be applied to metrics to transform, aggregate, or analyze them. The Splunk O11y Cloud Certified Metrics User Track document2 states that one of the recommended courses for preparing for the exam is Introduction to Splunk Infrastructure Monitoring, which covers the basics of metrics monitoring and visualization. In the Introduction to Splunk Infrastructure Monitoring course, there is a section on Analytic Functions, which explains that analytic functions can be used to perform calculations on metrics, such as sum, average, min, max, count, etc. The document also provides examples of how to use analytic functions in charts and dashboards. One of the analytic functions that can be used is Sum Aggregation, which calculates the sum of values present in the input time series across the entire environment or per group. The document gives an example of how to use Sum Aggregation to calculate the total CPU usage across all hosts in a group by using the following syntax: sum(cpu.utilization) by hostgroup

QUESTION 2

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

QUESTION 3

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 4

Which of the following statements about adding properties to MTS are true? (select all that apply)

A. Properties can be set via the API.

B. Properties are sent in with datapoints.

C. Properties are applied to dimension key:value pairs and propagated to all MTS with that dimension

D. Properties can be set in the UI under Metric Metadata.

Correct Answer: AD

According to the web search results, properties are key-value pairs that you can assign to dimensions of existing metric time series (MTS) in Splunk Observability Cloud1. Properties provide additional context and information about the metrics, such as the environment, role, or owner of the dimension. For example, you can add the property use: QA to the host dimension of your metrics to indicate that the host that is sending the data is used for QA. To add properties to MTS, you can use either the API or the UI. The API allows you to programmatically create, update, delete, and list properties for dimensions using HTTP requests2. The UI allows you to interactively create, edit, and delete properties for dimensions using the Metric Metadata page under Settings3. Therefore, option A and D are correct.

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 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

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 down. 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 6

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.

Correct Answer: ABCD

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 issues 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 aggregation 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 behavior 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. https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best- practices-for-detectors 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

QUESTION 7

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

Correct Answer: A

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 data

To learn more about how to use Max Delay in Splunk Observability Cloud, you can refer to this documentation.

1: https://docs.splunk.com/observability/alerts-detectors-notifications/detector- options.html#Max-Delay

QUESTION 8

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)

Correct Answer: C

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.

QUESTION 9

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

A. 64

B. 36

C. No limit

D. 50

Correct Answer: A

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

dimensions-mts.html

QUESTION 10

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

According to the Splunk O11y Cloud Certified Metrics User Track document, 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"))