



Splunk SPLK-4001 Practice Questions

Shared by Kinney on 18-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

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

Options:

- 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>

Question 2

Question Type: MultipleChoice

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?

Options:

- 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>

Question 3

Question Type: MultipleChoice

Interpreting data in charts can be affected by which of the following? (select all that apply)

Options:

- A- Tags
- B- Chart resolution

- C- Rollups
- D- Analytics functions

Answer:

B, C, D

Question 4

Question Type: MultipleChoice

What constitutes a single metrics time series (MTS)?

Options:

- 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.

Answer:

B

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

Question 5

Question Type: MultipleChoice

What information is needed to create a detector?

Options:

- 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

Answer:

C

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 6

Question Type: MultipleChoice

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

Options:

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

P2P
exams

Answer:

A

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.

P2P
exams

Question 7

Question Type: MultipleChoice

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?

Options:

- 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:

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

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:

```
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>

Question 8

Question Type: MultipleChoice

Which option best are supported rollup functions in Splunk Observability Cloud?

Options:

- A- average, latest, lag, min, max, sum, rate
- B- std_dev, mean, median, mode, min, max
- C- sigma, epsilon, pi, omega, beta, tau
- D- 1min, 5min, 10min, 15min, 30min

Answer:

A

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document¹, Observability Cloud has the following rollup functions: Sum: (default for counter metrics): Returns the sum of all data points in the MTS reporting interval. Average (default for gauge metrics): Returns the average value of all data points in the MTS reporting interval. Min: Returns the minimum data point value seen in the MTS reporting interval. Max: Returns the maximum data point value seen in the MTS reporting interval. Latest: Returns the most recent data point value seen in the MTS reporting interval. Lag: Returns the difference between the most recent and the previous data point values seen in the MTS reporting interval. Rate: Returns the rate of change of data points in the MTS reporting interval. Therefore, option A is correct.

Question 9

Question Type: MultipleChoice

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?

Options:

- 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.

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.

Question 10

Question Type: MultipleChoice

What Pod conditions does the Analyzer panel in Kubernetes Navigator monitor? (Select all that apply)

Options:

- A- Not Scheduled
- B- Unknown
- C- Failed
- D- Pending

P2P
exams

Answer:

A, B, C, D

Explanation:

The Pod conditions that the Analyzer panel in Kubernetes Navigator monitors are:

Not Scheduled: This condition indicates that the Pod has not been assigned to a Node yet. This could be due to insufficient resources, node affinity, or other scheduling constraints¹

Unknown: This condition indicates that the Pod status could not be obtained or is not known by the system. This could be due to communication errors, node failures, or other unexpected situations¹

Failed: This condition indicates that the Pod has terminated in a failure state. This could be due to errors in the application code, container configuration, or external factors¹

Pending: This condition indicates that the Pod has been accepted by the system, but one or more of its containers has not been created or started yet. This could be due to image pulling, volume mounting, or network issues¹

Therefore, the correct answer is A, B, C, and D.

To learn more about how to use the Analyzer panel in Kubernetes Navigator, you can refer to this documentation².

1: <https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#pod-phase> 2: <https://docs.splunk.com/observability/infrastructure/monitor/k8s-nav.html#Analyzer-panel>

To Get Premium Files for SPLK-4001 Visit

<https://www.p2pexams.com/products/splk-4001>

For More Free Questions Visit

<https://www.p2pexams.com/splunk/pdf/splk-4001>

20%
DISCOUNT

P2P
exams