



Free Questions for [SPLK-4001](#) by [dumpshq](#)

Shared by [Acosta](#) on [29-01-2024](#)

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

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 2

Question Type: MultipleChoice

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

Options:

- 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>

Question 3

Question Type: MultipleChoice

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

Options:

A- Dimensions

B- Properties

C- Sources

D- Tags

Answer:

A

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 4

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.

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.

Question 5

Question Type: MultipleChoice

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?

Options:

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>

Question 6

Question Type: MultipleChoice

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?

Options:

- 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.

Question 7

Question Type: MultipleChoice

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

Options:

- 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:

A, D

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.

Question 8

Question Type: MultipleChoice

Which of the following aggregate analytic functions will allow a user to see the highest or lowest n values of a metric?

Options:

- A- Maximum / Minimum
- B- Best/Worst
- C- Exclude / Include
- D- Top / Bottom

Answer:

D

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 9

Question Type: MultipleChoice

Which of the following are true about organization metrics? (select all that apply)

Options:

- 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:

A, C, D

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>

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>

