



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.geekcert.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 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 2

A Software Engineer is troubleshooting an issue with memory utilization in their application. They released a new canary version to production and now want to determine if the average memory usage is lower for requests with the `\\canary\\` version dimension. They've already opened the graph of memory utilization for their service.

How does the engineer see if the new release lowered average memory utilization?

- A. On the chart for plot A, select Add Analytics, then select MeanTransformation. In the window that appears, select `\\version\\` from the Group By field.
- B. On the chart for plot A, scroll to the end and click Enter Function, then enter `\\A/B-\\`.



C. On the chart for plot A, select Add Analytics, then select Mean:Aggregation. In the window that appears, select `version` from the Group By field.

D. On the chart for plot A, click the Compare Means button. In the window that appears, type `version1`.

Correct Answer: C

The correct answer is C. On the chart for plot A, select Add Analytics, then select Mean:Aggregation. In the window that appears, select `version` from the Group By field.

This will create a new plot B that shows the average memory utilization for each version of the application. The engineer can then compare the values of plot B for the `canary` and `stable` versions to see if there is a significant difference. To

learn more about how to use analytics functions in Splunk Observability Cloud, you can refer to this documentation¹.

1: <https://docs.splunk.com/Observability/gdi/metrics/analytics.html>

QUESTION 3

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.

Correct Answer: AD

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.

QUESTION 4

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

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

Correct Answer: A

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 host, 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 changed To learn more about how to use metadata in Splunk Observability Cloud, you can refer to this documentation.

<https://docs.splunk.com/Observability/metrics-and-metadata/metrics.html#Dimensions>

<https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html>

QUESTION 5

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

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.