Verified SPLK-4001 Exam Dumps PDF [2024 Access using BraindumpsIT [Q28-Q46

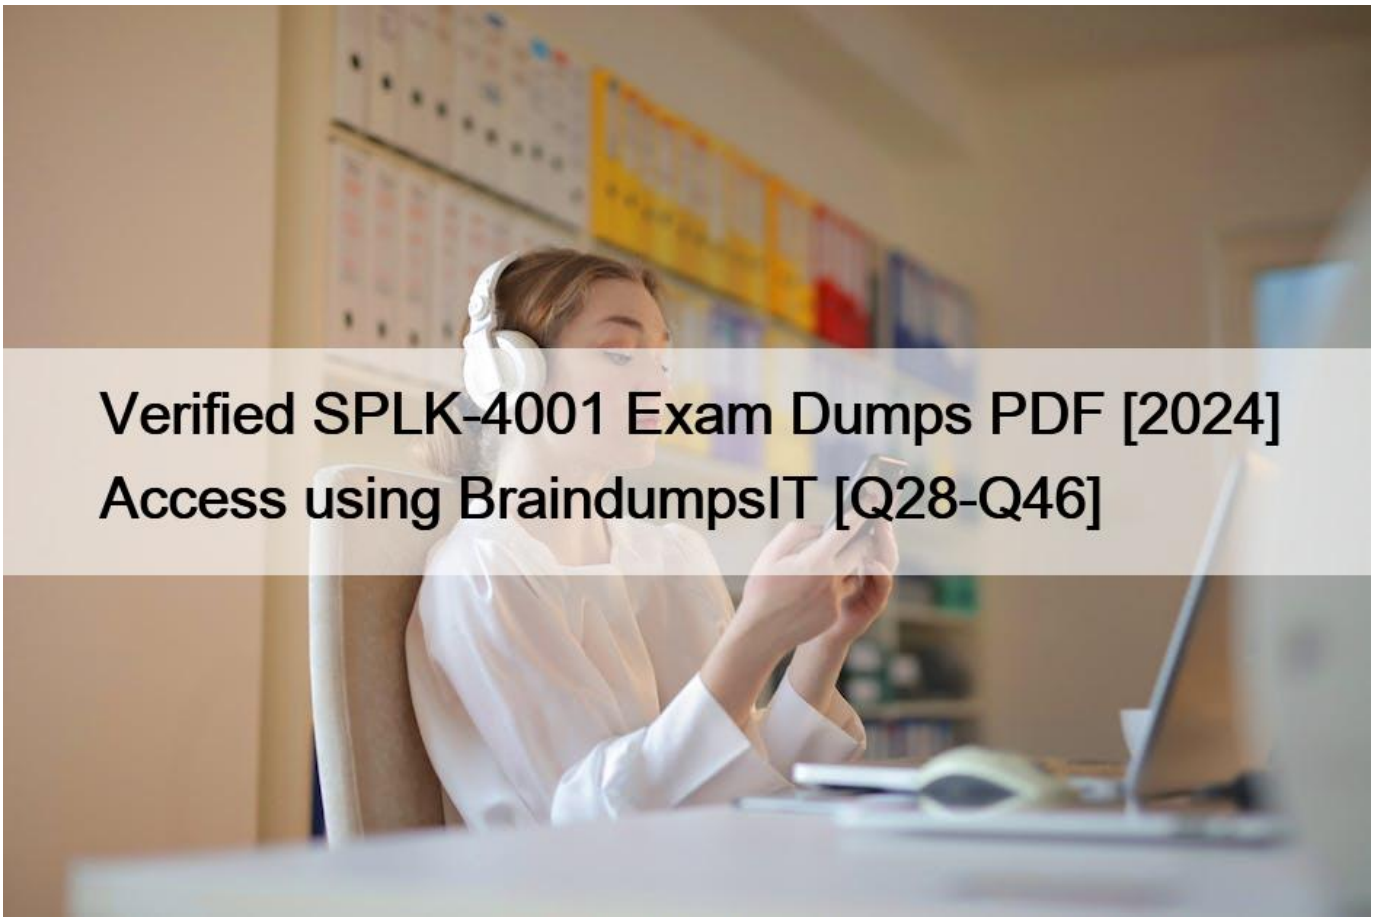

Verified SPLK-4001 Exam Dumps PDF [2024] Access using BraindumpsIT
Try Best SPLK-4001 Exam Questions from Training Expert BraindumpsIT

The Splunk SPLK-4001 exam consists of 65 questions that are to be completed in 90 minutes. The questions are in a multiple-choice format with a passing score of 70% or higher. Splunk O11y Cloud Certified Metrics User certification is valid for two years and must be renewed after the expiration date. There is a registration fee to take the exam, and the exam can be taken at any Pearson VUE test center or online through remote proctoring. Splunk O11y Cloud Certified Metrics User certification exam requires thorough preparation, including self-study, practice tests, and attending Splunk training courses.

NO.28 What is the limit on the number of properties that an MTS can have?
* 64
* 36
* No limit
* 50
Explanation

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 QA1 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 documentation2.

1:

https://docs.splunk.com/Observability/metrics-and-metadata/metrics-dimensions-mts.html#Custom-properties

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

NO.29 A user wants to add a link to an existing dashboard from an alert. When they click the dimension value in the alert message, they are taken to the dashboard keeping the context. How can this be accomplished? (select all that apply)
* Build a global data link.
* Add a link to the Runbook URL.
* Add a link to the field.
* Add the link to the alert message body.

Explanation

The possible ways to add a link to an existing dashboard from an alert are:

Build a global data link. A global data link is a feature that allows you to create a link from any dimension value in any chart or table to a dashboard of your choice. You can specify the source and target dashboards, the dimension name and value, and the query parameters to pass along. When you click on the dimension value in the alert message, you will be taken to the dashboard with the context preserved1 Add a link to the field. A field link is a feature that allows you to create a link from any field value in any search result or alert message to a dashboard of your choice. You can specify the field name and value, the dashboard name and ID, and the query parameters to pass along. When you click on the field value in the alert message, you will be taken to the dashboard with the context preserved2 Therefore, the correct answer is A and C.

To learn more about how to use global data links and field links in Splunk Observability Cloud, you can refer to these documentations12.

1: https://docs.splunk.com/Observability/gdi/metrics/charts.html#Global-data-links 2:

https://docs.splunk.com/Observability/gdi/metrics/search.html#Field-links

NO.30 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?
* On the chart for plot A, select Add Analytics, then select MeanrTransformation. In the window that appears, select ‘version’ from the Group By field.
* On the chart for plot A, scroll to the end and click Enter Function, then enter ‘A/B-l’.
* 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.
* On the chart for plot A, click the Compare Means button. In the window that appears, type ‘version1.
Explanation

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

NO.31 Changes to which type of metadata result in a new metric time series?
* Dimensions
* Properties
* Sources
* Tags
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

NO.32 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?
* The detector has an incorrect alert rule.
* The detector has an incorrect signal,
* The detector is disabled.
* The detector has a muting rule.
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.

NO.33 A DevOps engineer wants to determine if the latency their application experiences is growing fester after a new software release a week ago. They have already created two plot lines, A and B, that represent the current latency and the latency a week ago, respectively. How can the engineer use these two plot lines to determine the rate of change in latency?
* Create a temporary plot by dragging items A and B into the Analytics Explorer window.
* Create a plot C using the formula (A-B) and add a scale:percent function to express the rate of change as a percentage.
* Create a plot C using the formula (A/B-l) and add a scale: 100 function to express the rate of change as a percentage.
* Create a temporary plot by clicking the Change% button in the upper-right corner of the plot showing lines A and B.

Explanation

The correct answer is C. Create a plot C using the formula (A/B-l) and add a scale: 100 function to express the rate of change as a percentage.

To calculate the rate of change in latency, you need to compare the current latency (plot A) with the latency a week ago (plot B). One way to do this is to use the formula (A/B-l), which gives you the ratio of the current latency to the previous latency minus one. This ratio represents how much the current latency has increased or decreased relative to the previous latency. For example, if the current latency is 200 ms and the previous latency is 100 ms, then the ratio is (200/100-l) = 1, which means the current latency is 100% higher than the previous latency1 To express the rate of change as a percentage, you need to multiply the ratio by 100. You can do this by adding a scale: 100 function to the formula. This function scales the values of the plot by a factor of 100. For example, if the ratio is 1, then the scaled value is 100%2 To create a plot C using the formula (A/B-l) and add a scale: 100 function, you need to follow these steps:

Select plot A and plot B from the Metric Finder.

Click on Add Analytics and choose Formula from the list of functions.

In the Formula window, enter (A/B-l) as the formula and click Apply.

Click on Add Analytics again and choose Scale from the list of functions.

In the Scale window, enter 100 as the factor and click Apply.

You should see a new plot C that shows the rate of change in latency as a percentage.

To learn more about how to use formulas and scale functions in Splunk Observability Cloud, you can refer to these documentations34.

1: https://www.mathsisfun.com/numbers/percentage-change.html 2:

https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Scale 3:

https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Formula 4:

https://docs.splunk.com/Observability/gdi/metrics/analytics.html#Scale

NO.34 Which analytic function can be used to discover peak page visits for a site over the last day?
* Maximum: Transformation (24h)
* Maximum: Aggregation (Id)
* Lag: (24h)
* Count: (Id)

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.

NO.35 What information is needed to create a detector?
* Alert Status, Alert Criteria, Alert Settings, Alert Message, Alert Recipients
* Alert Signal, Alert Criteria, Alert Settings, Alert Message, Alert Recipients
* Alert Signal, Alert Condition, Alert Settings, Alert Message, Alert Recipients
* Alert Status, Alert Condition, Alert Settings, Alert Meaning, Alert Recipients

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.

NO.36 One server in a customer’s data center is regularly restarting due to power supply issues. What type of dashboard could be used to view charts and create detectors for this server?
* Single-instance dashboard

* Machine dashboard
* Multiple-service dashboard
* Server dashboard

Explanation

According to the Splunk O11y Cloud Certified Metrics User Track document1, a single-instance dashboard is a type of dashboard that displays charts and information for a single instance of a service or host. You can use a single-instance dashboard to monitor the performance and health of a specific server, such as the one that is restarting due to power supply issues. You can also create detectors for the metrics that are relevant to the server, such as CPU usage, memory usage, disk usage, and uptime. Therefore, option A is correct.

NO.37 Which of the following are required in the configuration of a data point? (select all that apply)

* Metric Name
* Metric Type
* Timestamp
* Value

Explanation

The required components in the configuration of a data point are:

Metric Name: A metric name is a string that identifies the type of measurement that the data point represents, such as cpu.utilization, memory.usage, or response.time. A metric name is mandatory for every data point, and it must be unique within a Splunk Observability Cloud organization1 Timestamp: A timestamp is a numerical value that indicates the time at which the data point was collected or generated. A timestamp is mandatory for every data point, and it must be in epoch time format, which is the number of seconds since January 1, 1970 UTC1 Value: A value is a numerical value that indicates the magnitude or quantity of the measurement that the data point represents. A value is mandatory for every data point, and it must be compatible with the metric type of the data point1 Therefore, the correct answer is A, C, and D.

To learn more about how to configure data points in Splunk Observability Cloud, you can refer to this documentation1.

1: https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Data-points

NO.38 What constitutes a single metrics time series (MTS)?

* A series of timestamps that all reflect the same metric.
* A set of data points that all have the same metric name and list of dimensions.
* A set of data points that use different dimensions but the same metric name.
* A set of metrics that are ordered in series based on timestamp.

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

NO.39 The built-in Kubernetes Navigator includes which of the following?
* Map, Nodes, Workloads, Node Detail, Workload Detail, Group Detail, Container Detail
* Map, Nodes, Processors, Node Detail, Workload Detail, Pod Detail, Container Detail
* Map, Clusters, Workloads, Node Detail, Workload Detail, Pod Detail, Container Detail
* Map, Nodes, Workloads, Node Detail, Workload Detail, Pod Detail, Container Detail

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

NO.40 Which of the following are correct ports for the specified components in the OpenTelemetry Collector?
* gRPC (4000), SignalFx (9943), Fluentd (6060)
* gRPC (6831), SignalFx (4317), Fluentd (9080)
* gRPC (4459), SignalFx (9166), Fluentd (8956)
* gRPC (4317), SignalFx (9080), Fluentd (8006)

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

NO.41 Where does the Splunk distribution of the OpenTelemetry Collector store the configuration files on Linux machines by default?
* /opt/splunk/
* /etc/otel/collector/
* /etc/opentelemetry/
* /etc/system/default/

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

NO.42 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)
* Custom events that have been sent in from an external source.
* Events created when a detector clears an alert.
* Random alerts from active detectors.
* Events created when a detector triggers an alert.

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.

NO.43 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?

* Max Delay
* Duration
* Latency
* Extrapolation Policy

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

NO.44 Which of the following are accurate reasons to clone a detector? (select all that apply)

* To modify the rules without affecting the existing detector.
* To reduce the amount of billed TAPM for the detector.
* To add an additional recipient to the detector’s alerts.
* To explore how a detector was created without risk of changing it.

Explanation

The correct answers are A and D.

According to the Splunk Test Blueprint – O11y Cloud Metrics User document1, one of the alerting concepts that is covered in the exam is detectors and alerts. Detectors are the objects that define the conditions for generating alerts, and alerts are the notifications that are sent when those conditions are met.

The Splunk O11y Cloud Certified Metrics User Track document2 states that one of the recommended courses for preparing for the exam is Alerting with Detectors, which covers how to create, modify, and manage detectors and alerts.

In the Alerting with Detectors course, there is a section on Cloning Detectors, which explains that cloning a detector creates a copy of the detector with all its settings, rules, and alert recipients. The document also provides some reasons why you might want to clone a detector, such as:

To modify the rules without affecting the existing detector. This can be useful if you want to test different thresholds or conditions before applying them to the original detector.

To explore how a detector was created without risk of changing it. This can be helpful if you want to learn from an existing detector or use it as a template for creating a new one.

Therefore, based on these documents, we can conclude that A and D are accurate reasons to clone a detector.

B and C are not valid reasons because:

Cloning a detector does not reduce the amount of billed TAPM for the detector. TAPM stands for Tracked Active Problem Metric, which is a metric that has been alerted on by a detector. Cloning a detector does not change the number of TAPM that are generated

by the original detector or the clone.

Cloning a detector does not add an additional recipient to the detector’s alerts. Cloning a detector copies the alert recipients from the original detector, but it does not add any new ones. To add an additional recipient to a detector’s alerts, you need to edit the alert settings of the detector.

NO.45 Which of the following statements are true about local data links? (select all that apply)
* Anyone with write permission for a dashboard can add local data links that appear on that dashboard.
* Local data links can only have a Splunk Observability Cloud internal destination.
* Only Splunk Observability Cloud administrators can create local links.
* Local data links are available on only one dashboard.

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.

NO.46 Which of the following statements is true of detectors created from a chart on a custom dashboard?
* Changes made to the chart affect the detector.
* Changes made to the detector affect the chart.
* The alerts will show up in the team landing page.
* The detector is automatically linked to the chart.

Explanation

The correct answer is D. The detector is automatically linked to the chart.

When you create a detector from a chart on a custom dashboard, the detector is automatically linked to the chart. This means that you can see the detector status and alerts on the chart, and you can access the detector settings from the chart menu. You can also unlink the detector from the chart if you want to1 Changes made to the chart do not affect the detector, and changes made to the

detector do not affect the chart.

The detector and the chart are independent entities that have their own settings and parameters. However, if you change the metric or dimension of the chart, you might lose the link to the detector1 The alerts generated by the detector will show up in the Alerts page, where you can view, manage, and acknowledge them. You can also see them on the team landing page if you assign the detector to a team2 To learn more about how to create and link detectors from charts on custom dashboards, you can refer to this documentation1.

1: https://docs.splunk.com/observability/alerts-detectors-notifications/link-detectors-to-charts.html 2:

https://docs.splunk.com/observability/alerts-detectors-notifications/view-manage-alerts.html

Achieving the SPLK-4001 certification demonstrates that an individual has a thorough understanding of Splunk's Observability Cloud and is capable of using it to monitor and analyze data effectively. Splunk O11y Cloud Certified Metrics User certification is recognized by employers and can help professionals advance their careers. Additionally, SPLK-4001 certification holders are eligible to join the Splunk Trust, a community of top-performing Splunk professionals.

Latest 100% Passing Guarantee - Brilliant SPLK-4001 Exam Questions PDF:
https://www.braindumpsit.com/SPLK-4001_real-exam.html]