Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.

KILL EXAMS

SPLK-4001 Dumps
SPLK-4001 Braindumps
SPLK-4001 Real Questions
SPLK-4001 Practice Test
SPLK-4001 Actual Questions

PASS

Splunk

SPLK-4001

Splunk O11y Cloud Certified Metrics User

Question: 171

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.

Answer: A,B,C,D

Explanation:

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 issues1

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 aggregation2

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 behavior3

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.

1: https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors 2: https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Best-practices-for-detectors 3: 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: 172

An SRE came across an existing detector that is a good starting point for a detector they want to create. They clone the detector, update the metric, and add multiple new signals.

As a result of the cloned detector, which of the following is true?

A. The new signals will be reflected in the original detector.
B. The new signals will be reflected in the original chart.
C. You can only monitor one of the new signals.
D. The new signals will not be added to the original detector.

Answer: D

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document1, cloning a detector creates a copy of the detector that you can modify without affecting the original detector. You can change the metric, filter, and signal settings of the cloned detector. However, the new signals that you add to the cloned detector will not be reflected in the original detector, nor in the original chart that the detector was based on. Therefore, option D is correct.

Option A is incorrect because the new signals will not be reflected in the original detector. Option B is incorrect because the new signals will not be reflected in the original chart. Option C is incorrect because you can monitor all of the new signals that you add to the cloned detector.

Question: 173

Which of the following are supported rollup functions in Splunk Observability Cloud?

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: 174

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 MeanrTransformation. 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-l'.
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.

Answer: C

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 âversionâ 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 âcanaryâ and âstableâ 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: 175

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?

A. Single-instance dashboard
B. Machine dashboard
C. Multiple-service dashboard
D. Server dashboard

Answer: A

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.

Question: 176

To refine a search for a metric a customer types host: test-*.

What does this filter return?

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: 177

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

Answer: A

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document1, 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(âcache.hitsâ), counters(âcache.missesâ))

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(âcache.hitsâ), counters(âcache.missesâ))

Question: 178

Which of the following are correct ports for the specified components in the OpenTelemetry Collector?

A. gRPC (4000), SignalFx (9943), Fluentd (6060)
B. gRPC (6831), SignalFx (4317), Fluentd (9080)
C. gRPC (4459), SignalFx (9166), Fluentd (8956)
D. gRPC (4317), SignalFx (9080), Fluentd (8006)

Answer: D

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

Question: 179

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?

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-metrics2:
https://docs.splunk.com/Observability/gdi/metrics/detectors.html#Plot-cap 3:
https://docs.splunk.com/Observability/gdi/metrics/search.html

Question: 180

An SRE creates a new detector to receive an alert when server latency is higher than 260 milliseconds. Latency below 260 milliseconds is healthy for their service. The SRE creates a New Detector with a Custom Metrics Alert Rule for latency and sets a Static Threshold alert condition at 260ms.

How can the number of alerts be reduced?

A. Adjust the threshold.
B. Adjust the Trigger sensitivity. Duration set to 1 minute.
C. Adjust the notification sensitivity. Duration set to 1 minute.
D. Choose another signal.

Answer: B

Explanation:

According to the Splunk O11y Cloud Certified Metrics User Track document1, trigger sensitivity is a setting that determines how long a signal must remain above or below a threshold before an alert is triggered. By default, trigger sensitivity is set to Immediate, which means that an alert is triggered as soon as the signal crosses the threshold. This can result in a lot of alerts, especially if the signal fluctuates frequently around the threshold value. To reduce the number of alerts, you can adjust the trigger sensitivity to a longer duration, such as 1 minute, 5 minutes, or 15 minutes. This means that an alert is only triggered if the signal stays above or below the threshold for the specified duration. This can help filter out noise and focus on more persistent issues.

Question: 181

Where does the Splunk distribution of the OpenTelemetry Collector store the configuration files on Linux machines by default?

A. /opt/splunk/
B. /etc/otel/collector/
C. /etc/opentelemetry/
D. /etc/system/default/

Answer: B

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

Question: 182

Which of the following rollups will display the time delta between a datapoint being sent and a datapoint being received?

A. Jitter
B. Delay
C. Lag
D. Latency

Answer: C

Explanation:

According to the Splunk Observability Cloud documentation1, lag is a rollup function that returns the difference between the most recent and the previous data point values seen in the metric time series reporting interval. This can be used to measure the time delta between a data point being sent and a data point being received, as long as the data points have timestamps that reflect their send and receive times. For example, if a data point is sent at 10:00:00 and received at 10:00:05, the lag value for that data point is 5 seconds.

Question: 183

Which of the following is optional, but highly recommended to include in a datapoint?

A. Metric name
B. Timestamp
C. Value
D. Metric type

Answer: D

Explanation:

The correct answer is D. Metric type.

A metric type is an optional, but highly recommended field that specifies the kind of measurement that a datapoint represents. For example, a metric type can be gauge, counter, cumulative counter, or histogram. A metric type helps Splunk Observability Cloud to interpret and display the data correctly1

To learn more about how to send metrics to Splunk Observability Cloud, you can refer to this documentation2.

1: https://docs.splunk.com/Observability/gdi/metrics/metrics.html#Metric-types 2: https://docs.splunk.com/Observability/gdi/metrics/metrics.html

Question: 184

Which analytic function can be used to discover peak page visits for a site over the last day?

A. Maximum: Transformation (24h)
B. Maximum: Aggregation (Id)
C. Lag: (24h)
D. Count: (Id)

Answer: A

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(âpage.visitsâ))

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.

Question: 185

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?

A. Adding debug into the metrics receiver pipeline:

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

B. Adding logging into the metrics receiver pipeline:

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

C. Adding logging into the metrics exporter pipeline:

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

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: 186

What information is needed to create a detector?

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: 187

A customer has a large population of servers. They want to identify the servers where utilization has increased the most since last week.

Which analytics function is needed to achieve this?

A. Rate
B. Sum transformation
C. Tlmeshift
D. Standard deviation

Answer: C

Explanation:

The correct answer is

C. Timeshift.

According to the Splunk Observability Cloud documentation1, timeshift is an analytic function that allows you to compare the current value of a metric with its value at a previous time interval, such as an hour ago or a week ago. You can use the timeshift function to measure the change in a metric over time and identify trends, anomalies, or patterns. For example, to identify the servers where utilization has increased the most since last week, you can use the following SignalFlow code: timeshift(1w, counters(âserver.utilizationâ))

This will return the value of the server.utilization counter metric for each server one week ago. You can then subtract this value from the current value of the same metric to get the difference in utilization. You can also use a chart to visualize the results and sort them by the highest difference in utilization.

Question: 188

The alert recipients tab specifies where notification messages should be sent when alerts are triggered or cleared.

Which of the below options can be used? (select all that apply)

A. Invoke a webhook UR
B. Export to CS
C. Send an SMS message.
D. Send to email addresses.

Answer: A,B

Explanation:

The alert recipients tab specifies where notification messages should be sent when alerts are triggered or cleared.

The options that can be used are:

Invoke a webhook URL. This option allows you to send a HTTP POST request to a custom URL that can perform various actions based on the alert information. For example, you can use a webhook to create a ticket in a service desk

system, post a message to a chat channel, or trigger another workflow1

Send an SMS message. This option allows you to send a text message to one or more phone numbers when an alert is triggered or cleared. You can customize the message content and format using variables and templates2

Send to email addresses. This option allows you to send an email notification to one or more recipients when an alert is triggered or cleared. You can customize the email subject, body, and attachments using variables and templates. You can also include information from search results, the search job, and alert triggering in the email3

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

1: https://docs.splunk.com/Documentation/Splunk/latest/Alert/Webhooks 2: https://docs.splunk.com/Documentation/Splunk/latest/Alert/SMSnotification 3: https://docs.splunk.com/Documentation/Splunk/latest/Alert/Emailnotification

Question: 189

With exceptions for transformations or timeshifts, at what resolution do detectors operate?

A. 10 seconds
B. The resolution of the chart
C. The resolution of the dashboard
D. Native resolution

Answer: D

Explanation:

According to the Splunk Observability Cloud documentation1, detectors operate at the native resolution of the metric or dimension that they monitor, with some exceptions for transformations or timeshifts. The native resolution is the frequency at which the data points are reported by the source. For example, if a metric is reported every 10 seconds, the detector will evaluate the metric every 10 seconds. The native resolution ensures that the detector uses the most granular and accurate data available for alerting.

Question: 190

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

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 organizationâs usage of the platform.

They include metrics such as:

Ingest metrics: Measure the data youâre sending to Infrastructure Monitoring, such as the number of data points youâve 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) youâve 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

SAMPLE QUESTIONS

These questions are for demo purpose only. ***Full version*** *is up to date and contains actual questions and answers.*

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:

Actual Exam Questions*: Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

Exam Dumps*: Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

Practice Tests*: Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

Guaranteed Success*: Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

Updated Content: *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

Technical Support*: Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*