DUMPSARENA

Splunk O11y Cloud Certified Metrics User Exam

Splunk SPLK-4001

Version Demo

Total Demo Questions: 7

Total Premium Questions: 78

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Splunk Observability Cloud Overview	7
Topic 2, Metrics	37
Topic 3, Dashboards	11
Topic 4, Detectors	23
Total	78

QUESTION NO: 1

Which of the following chart visualization types are unaffected by changing the time picker on a dashboard? (select all that apply)

A. Single Value

B. Heatmap

C. Line

D. List

ANSWER: A D

Explanation:

Single Value and **List** are unaffected by a dashboard time picker change because these visualizations are designed to present the most recent value available for the selected metric or computed expression, rather than redraw a time-series view over the dashboard's selected interval. A Single Value chart surfaces one current aggregate or latest data value, using its configured rollup and resolution behavior. Likewise, a List chart presents the latest values for metric time series, organized by their dimensions, so it is useful for seeing the current state of many entities such as hosts, services, or containers.

This behavior makes both chart types appropriate for at-a-glance operational status panels where the immediate value matters more than historical movement. Their displayed result can still reflect the chart's signal, filters, aggregation, resolution, and rollup configuration; however, changing the dashboard-level time range does not alter these visualization types in the way it does charts intended to render data across time. Splunk's chart reference documents the intended use and behavior of available visualization types, including Single Value and List charts. See the Splunk Observability Cloud chart types reference and the dashboard time range documentation.

QUESTION NO: 2

An SRE wants one detector to provide an early warning when API latency exceeds 400 milliseconds and a more urgent alert when it exceeds 700 milliseconds. Which two actions should the SRE take? (select two)

A. Create one alert rule with a static threshold of 400 milliseconds.

B. Create separate alert rules for the 400 millisecond and 700 millisecond thresholds.

C. Assign Warning severity to the lower-threshold rule and Critical severity to the higher-threshold rule.

D. Add additional notification recipients to the lower-threshold rule.

ANSWER: B C

Explanation:

A detector can contain multiple alert rules, and each rule can have its own threshold and severity. Creating separate rules for the 400 ms and 700 ms conditions allows the detector to represent escalating service impact without requiring separate detectors.

Assigning Warning severity to the lower threshold and Critical severity to the higher threshold communicates the intended urgency to responders. A single rule cannot represent both threshold levels as distinct alert severities, and changing notification recipients does not create escalation logic.

QUESTION NO: 3

For a high-resolution metric, what is the highest possible native resolution of the metric?

A. 2 seconds

B. 15 seconds

C. 1 second

D. 5 seconds

ANSWER: C

Explanation:

1 second is the highest possible native resolution for a high-resolution metric in Splunk Observability Cloud. Native resolution describes the interval at which the metric's original data points are received and stored before time-based rollups are used for longer-term querying and visualization. A one-second interval provides the finest granularity available, allowing users to inspect rapid changes in system and application behavior, such as brief latency spikes, bursty request rates, or short-lived resource saturation. This level of detail is particularly valuable when troubleshooting infrastructure and services whose important behavior can occur between broader collection intervals.

Splunk Observability Cloud subsequently generates rollups at coarser intervals so that historical metric data remains efficient to query over longer time ranges. Those rollups do not change the fact that the finest available native data for a high-resolution metric is captured once per second. When selecting chart time ranges and assessing the detail available for an investigation, users should distinguish the original one-second data from rolled-up representations selected for broader time windows. Splunk's metric documentation describes metric resolution and rollups, including the role of high-resolution data: Splunk Observability Cloud: Metric resolution and rollups. See also Splunk Help: Metric resolution.

QUESTION NO: 4

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 URL.

B. Export to CSV.

C. Send an SMS message.

D. Send to email addresses.

ANSWER: A C D

Explanation:

Invoke a webhook URL, **Send an SMS message**, and **Send to email addresses** are valid ways to configure delivery of detector alert notifications. A webhook recipient sends alert event information to an HTTP endpoint, allowing an organization to connect Observability Cloud alerts with an external automation, incident-management, ticketing, or messaging workflow. The receiving service can use the payload to create or update an incident, enrich an event, or initiate remediation.

SMS delivery provides a direct mobile notification channel for alert activity when a detector transitions into an alerting state or subsequently clears. Email recipients provide a standard notification channel for one or more specified addresses, enabling relevant responders and stakeholders to receive the alert message and its associated detector context. These recipient settings determine notification destinations; they do not change the detector's signal, condition, or threshold evaluation. Splunk's alerting documentation describes configuring notifications and recipient destinations for detector alerts, including integrations and message delivery behavior. See Splunk Observability Cloud: Alerts, detectors, and notifications and Splunk Observability Cloud: Users and teams.

QUESTION NO: 5

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 (9943), Fluentd (8006)

ANSWER: D

Explanation:

gRPC (4317), SignalFx (9943), Fluentd (8006) correctly identifies the default listening ports for these OpenTelemetry Collector receivers. The OpenTelemetry Protocol gRPC receiver conventionally listens on port 4317, which is the standard endpoint for applications and agents sending OTLP data over gRPC. The SignalFx receiver listens on port 9943 by default, allowing it to accept metrics sent using the SignalFx protocol. The Fluentd Forward receiver listens on port 8006, enabling Fluentd or Fluent Bit forward-protocol clients to send log data to the collector. These receiver endpoints are configurable, but the listed values are the distribution defaults and are important when configuring senders, network policies, service exposure, and firewall rules. Splunk documents the collector's exposed receiver endpoints, including the OTLP, SignalFx, and Fluentd endpoints, in its OpenTelemetry Collector exposed endpoints documentation. The upstream SignalFx receiver configuration also documents its default endpoint as `0.0.0.0:9943`: OpenTelemetry SignalFx receiver documentation.

QUESTION NO: 6

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

ANSWER: D

Explanation:

The built-in Kubernetes Navigator includes Map, Nodes, Workloads, Node Detail, Workload Detail, Pod Detail, and Container Detail. Kubernetes Navigator is Splunk Observability Cloud's purpose-built interface for exploring Kubernetes infrastructure and moving from a high-level cluster view to the health and performance details of individual resources. The Map view presents the relationships among Kubernetes entities, helping users understand the cluster's topology. Nodes and Workloads provide focused views for examining those resource categories across the cluster. From these views, users can select an individual entity and drill into its corresponding detail page: Node Detail for a node, Workload Detail for a workload, Pod Detail for a pod, and Container Detail for a container. This hierarchical navigation supports efficient troubleshooting by connecting broad infrastructure context with increasingly granular telemetry. Splunk documents Kubernetes Navigator as including these navigator views and entity-detail workflows, with the exact available content depending on the Kubernetes metrics collected through the Splunk Distribution of the OpenTelemetry Collector. See Splunk's Kubernetes Navigator documentation and its Infrastructure Navigators overview for the supported views and navigation behavior.

QUESTION NO: 7

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:

Only metrics with a dimension of host and a value beginning with test- is correct. In Splunk Observability Cloud metric search filters, the text before the colon identifies the dimension name, so `host` restricts the search to metric time series that include a `host` dimension. The value expression after the colon is evaluated against that dimension's value. Here, `test-*` uses the asterisk as a wildcard, matching values whose prefix is `test-`. For example, time series tagged with `host=test-01`, `host=test-api`, or `host=test-production` are returned, provided they otherwise match the metric search context. This type of dimension filtering is useful when hosts follow a naming convention and an analyst needs to focus quickly on a particular group of infrastructure without enumerating every individual host. The filter is specifically a dimension-name and dimension-value match, rather than a free-text match against metric names or arbitrary metadata. Splunk documents metric search and filtering behavior, including the use of dimensions to narrow metric results, in its Filter metrics documentation. Additional metric and metadata concepts are available in the Splunk Observability Cloud metrics documentation.