Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.

SPLK-4001 Practice Questions
SPLK-4001 Practice Test
SPLK-4001 Practice Exam
SPLK-4001 Exam Questions
SPLK-4001 Study Guide

killexams.com

Splunk

SPLK-4001

Splunk O11y Cloud Certified Metrics User

ORDER FULL VERSION

https://killexams.com/pass4sure/exam-detail/SPLK-4001

Question: 887

An engineer uses data('jvm.memory.used').sum(by='app_name'). If the metric also has a host_id dimension, what happens to that dimension in the output?

A. It is replaced by the app_name dimension.
B. It is converted into a tag.
C. It is preserved and used as a secondary grouping.
D. It is dropped from the output metadata.

Answer: D
Explanation: When an aggregation is performed by a specific dimension, all other dimensions that are not part of the by clause are aggregated away and removed from the resulting time series metadata.

Question: 888

On the **Built-in "K8s Pods" Dashboard**, there is a chart for "CPU Throttling by Container." If you see a container with 0% CPU usage but 100% throttling, what is the most likely configuration error?

A. The CPU limit is set to a value lower than what is required for the process to even initialize (e.g., 1m).
B. The Liveness probe is failing.
C. The CPU request is too high for the node to handle.
D. The Pod is in a Succeeded state.

Answer: A
Explanation: Throttling is measured as the percentage of time periods where the container was ready to run but was prevented by the scheduler. If the limit is set extremely low (e.g., 1 millicore), the process may be throttled so aggressively that it cannot even register significant usage before the period ends.

Question: 889

Which of the following is true regarding "Data Resolution" in Splunk Observability Cloud?

A. It refers to the number of pixels used to render the charts in the browser
B. It represents the granularity of data, such as 1-second or 1-minute intervals
C. It is a type of metadata used to identify the source of a metric

D. It is the process of converting logs into metrics

Answer: B
Explanation: In the context of metrics, Resolution is synonymous with granularity. It defines how often a measurement is recorded and stored. High-resolution data (e.g., 1s) provides immense detail for real-time troubleshooting, while lower-resolution data (e.g., 1m or 1h) is used for long-term trending and performance optimization of the platform.

Question: 890

In a built-in chart for a rate metric derived from counters, the values appear lower than manual calculations at the same time range. The setup uses default rollups and no additional analytics. What explains this discrepancy?

A. The metric is treated as a gauge, applying Average regardless of origin.
B. The chart applies a rollup (e.g., Average or Sum) to the counter before any implicit rate computation or display, and resolution coarsening averages the rate over larger intervals.
C. Subscription to alerts forces a global min rollup on all rate-derived signals.
D. Built-in content disables rate calculations for counters to avoid negative values.

Answer: B
Explanation: Counters require explicit or implicit rate transformations for meaningful rate views; default rollups (Sum) followed by resolution-based aggregation can underrepresent peaks or averages in displayed rates compared to raw differencing at native scale. This is a common interpretation pitfall in built-in content. Content does not disable rates, and metric type handling follows defined rules rather than forcing gauge behavior or min rollups via subscriptions.

Question: 891

You are creating a single-instance dashboard for a "Customer ID." The number of customers is over 100,000. What is a potential issue with using a standard Dimension Variable for this?

A. Users must type the ID manually as the dropdown will be disabled.
B. The dropdown menu will only show the first 1,000 values.
C. The dashboard will crash due to high cardinality.
D. You cannot use variables for dimensions with more than 10,000 values.

Answer: B
Explanation: Dropdown menus for dashboard variables have a limit on the number of suggested values they can display (typically around 1,000). For extremely high-cardinality dimensions, users may need to start typing the value to filter the suggestions or use a "Text" type variable where they input the value

manually.

Question: 892

A user is configuring a notification for a detector using the "Email" integration. They want the email to include a direct link to the specific dashboard where the anomaly was first observed. Which feature should be used to ensure this link is dynamically included in the subscription?

A. Notification Variable Substitution
B. Tip Sheets
C. Dashboard Deep-linking
D. Alert Correlation

Answer: B
Explanation: "Tip Sheets" or custom message templates in Splunk O11y Cloud allow users to add context to alert notifications. By using specific variables or markdown, users can include links to dashboards, runbooks, or specific troubleshooting steps that are sent along with the alert notification to the subscribed user or team.

Question: 893

Scenario:
Your team is monitoring a Kubernetes cluster with 500 ephemeral pods where CPU utilization metrics arrive irregularly due to autoscaling. A static threshold detector on cpu.utilization frequently flaps, triggering alerts during brief spikes that resolve quickly, and misses persistent issues in subgroups. The detector uses a 1-minute resolution with no duration configured and aggregates across all pods without grouping.

A. Increase detector resolution to 5 minutes and apply a 10-minute duration threshold to require sustained violations before alerting.
B. Switch to Outlier Detection condition without population grouping or smoothing, as it inherently handles ephemeral sources better than static thresholds.
C. Use a static threshold with a rolling mean transformation over a 15-minute window and group by service dimension to monitor per-subgroup 95th percentile.
D. Configure the detector to ignore late datapoints via extrapolation policy set to "last value" and remove all duration settings to capture all transient events in ephemeral infrastructure.

Answer: C
Explanation: For monitoring populations in ephemeral infrastructure with cyclic or bursty patterns, best practices recommend using aggregation functions like percentile to track subgroups (e.g., by service

dimension) rather than raw individual MTS, combined with smoothing transformations such as rolling mean to reduce noise and flapping. A duration or percent-of-duration setting further prevents alerts on transient violations, while static thresholds can be appropriate when combined with analytics. Outlier Detection is useful for deviation from population norms but requires proper configuration for large-scale sources. Ignoring late datapoints or shortening duration without smoothing often exacerbates flapping or misses real issues, and high resolution without duration increases noise in aperiodic data common to ephemeral pods. Extrapolation policies help with delayed data but do not address root causes like lack of aggregation or transformation.

Question: 894

You need to collect metrics from a Linux host that is running in a Docker container. To get accurate host-level metrics (not just container-level), which volumes must be mapped to the OTel Collector container?

A. /proc, /sys, and / (root)
B. /bin and /usr/lib
C. /dev/log and /var/log
D. /var/run/docker.sock and /etc/hostname

Answer: A
Explanation: For the hostmetrics receiver to see the physical host's stats from inside a container, it must have access to the host's /proc and /sys filesystems. Typically, these are mounted into the container (e.g., as /hostfs/proc and /hostfs/sys), and the receiver is configured to look at those specific paths using the root_path setting.

Question: 895

In **Kubernetes Navigator**, you use the "Group By" feature to organize the map by kubernetes_standard_label_app_kubernetes_io_name. Why is this considered a best practice for investigating large-scale outages?

A. It bypasses the need for SignalFlow permissions.
B. It automatically restarts Pods that are grouped together.
C. It reduces the amount of data the Splunk O11y collector needs to send.
D. It organizes the visualization logically by application or service rather than by physical host.

Answer: D
Explanation: By default, the Navigator might group by Node. However, for an SRE, grouping by the app label allows them to see the health of an entire logical service (comprised of multiple Pods) across the entire cluster, making it easier to see if an issue is localized to one node or widespread across the service.

Question: 896

For a dashboard tracking infrastructure, the mean disk usage is needed for the top 25% busiest servers over a rolling 30-minute period. What function setup achieves this?

A. Bottom N on usage
B. Full mean with calendar week
C. Ratio of disk to total capacity
D. top(25%, disk.usage) then mean:transformation with 30m moving window

Answer: D
Explanation: applying analytics functions to a subset of MTS in a signal uses top() with percentage, then applies the mean over a moving time window to focus on the busiest servers' smoothed utilization.

Question: 897

A monitoring team is setting up an alert for disk.utilization. They want the alert to trigger only if the average utilization stays above 90% for 5 minutes. Which concept defines the frequency at which the alert evaluates the incoming MTS?

A. Lag
B. Dimensions
C. Resolution
D. Jitter

Answer: C
Explanation: Resolution is the frequency at which data is sampled or reported. If the resolution is 10 seconds, the alerting engine receives a new datapoint for the MTS every 10 seconds and can evaluate the condition. Higher resolution (frequent data) allows for faster alerting, while lower resolution (infrequent data) might delay the detection of a breach.

Question: 898

In the Kubernetes nodes navigator, a node appears unhealthy on the heatmap. What sequence of steps using the Navigator and Analyzer best isolates whether the problem is node-level (e.g., Kubelet issues) or pod-related? (Single-answer)

A. Apply a top analytic function on node metrics in a custom chart.

B. Filter the pods navigator globally by the node name and check only container images.
C. Drill into the node detail view, review node conditions and pod list, then access the Analyzer tab for AI insights on related patterns across pods/containers.
D. Use built-in dashboards exclusively to view aggregate node metrics without detail views.

Answer: C
Explanation: Node detail view shows conditions (e.g., Kubelet version/status) and pod lists for initial assessment, while the Analyzer tab delivers AI insights on patterns linking node issues to pods or containers, providing comprehensive isolation not available through global filters or dashboards alone.

Question: 899

While interacting with a built-in dashboard, you apply a "Global Filter" for env:production. You notice that one specific chart titled "Global System Health" does not change its data output. What is the most likely reason for this behavior?

A. The user lacks the "Filter Override" permission for that specific dashboard.
B. The chart is using a metric that does not have an env dimension associated with it.
C. The "Global System Health" chart is an embedded external image and not a native Splunk O11y chart.
D. The chart is hardcoded with a SignalFlow script that overrides the UI-level filters.

Answer: D
Explanation: In Splunk O11y Cloud, charts can be configured to "Listen to Global Filters" or to ignore them. If a chart's SignalFlow specifically defines filters or if the "Override Dashboard Filters" toggle is enabled in the chart settings, the dashboard-level global filters will not affect that specific visualization.

Question: 900

A metrics user builds a chart and wants to save it efficiently for reuse across multiple dashboard groups. What is the recommended practice? (Single-answer)

A. Use only single-instance dashboards to avoid sharing complexity.
B. Export the chart as an image and re-import manually into each dashboard.
C. Save the chart to a dashboard, then use dashboard groups or sharing features to organize and propagate access without duplicating the chart definition.
D. Create a new chart from scratch in each dashboard group.

Answer: C
Explanation: Saving once and leveraging dashboard groups or sharing avoids duplication while enabling organization. Export/import or recreation is inefficient. Single-instance is a style choice, not a sharing limitation.

Question: 901

A gauge metric memory.used is ingested with dimensions host and environment, and custom properties are later applied post-ingest to selected MTS values of the environment dimension. If a chart query filters on a custom property value that matches only a subset of MTS, explain the distinction in how this affects MTS identification versus the use of dimensions in the same query, and what happens to datapoint components during chart resolution adjustment to a coarser interval.

A. Both custom properties and dimensions equally redefine MTS identity upon application, causing new MTS creation, and coarser resolution drops the value component of datapoints without rollup.
B. Custom properties replace dimensions for MTS identification in filtered queries, and resolution adjustment interpolates missing timestamps rather than rolling up values.
C. Custom properties do not participate in MTS uniqueness and only enable post-ingest filtering/grouping, whereas dimensions define MTS identity; coarser chart resolution applies rollups to datapoint values while preserving the original timestamp and dimensions in the underlying MTS.
D. Dimensions are mutable post-ingest like custom properties, allowing dynamic MTS redefinition, while resolution changes affect only the metric name component of datapoints.

Answer: C
Explanation: Dimensions are sent with data at ingest time and, together with the metric name and type, uniquely identify each MTS; changing or adding dimensions creates new MTS. Custom properties are applied after ingest to existing MTS (via dimension values) and provide additional context for filtering, grouping, or searching without altering MTS identity or creating new series. A datapoint consists of timestamp, metric name/type, value, and dimensions. When chart resolution is coarsened, Splunk Observability Cloud applies a rollup function to aggregate multiple datapoint values within the new interval into a single displayed datapoint per MTS, while the underlying stored datapoints and their components (including original timestamps and dimensions) remain unchanged.

Question: 902

While troubleshooting an alert that didn't fire during a known outage, an engineer notices the "Data Gap" icon on the chart. The metric's resolution is 1 minute, but data was missing for 3 minutes. The extrapolation policy was set to "Null." Why did the detector fail?

A. A "Null" value does not satisfy a comparison operator like > or <, so the condition could not be evaluated as true.
B. "Null" causes the detector to assume the last known state was "Healthy."
C. The detector uses "Mean" by default to fill gaps regardless of dashboard settings.
D. The detector automatically reboots when it encounters three consecutive nulls.

Answer: A
Explanation: Detectors require numerical data to evaluate conditions. If data is "Null" (due to late data or gaps with a Null policy), the detector cannot confirm that the threshold was crossed. Therefore, the "Duration" or "Percent of Time" requirements for the alert to fire will not be met.

Question: 903

The **Cluster Analyzer** identifies a "Hot Spot" in your cluster. When you drill down, you see a high value for container_cpu_usage_seconds_total but a very low kube_pod_container_resource_requests_cpu_cores. What is the primary troubleshooting conclusion?

A. The node's clock is out of sync with the Splunk O11y backend.
B. The Pod is "Guaranteed" and is performing as expected.
C. The Pod is "Burstable" and is currently using more than its fair share of the node.
D. The Kubelet is failing to report accurate usage metrics.

Answer: C
Explanation: A "Hot Spot" combined with high usage but low requests indicates that a Pod was scheduled with minimal requirements (Requests) but is now consuming a large amount of CPU. This can lead to the "Noisy Neighbor" effect where other Pods on the same node suffer.

Question: 904

A detector is built on a histogram MTS using a 10-second native but with detector resolution fixed at 1 minute. The histogram rollup buckets data, and a transformation extracts the 90th percentile from the distribution before thresholding. If late datapoints arrive irregularly and the extrapolation policy is "none" (drop gaps), derive the risk to detection reliability and specify which datapoint components remain reliable versus those impacted.

A. Gaps cause dropped intervals in the rolled-up distribution, leading to missing or incomplete percentile calculations for those minutes and potential missed alerts; the timestamp for non-gapped intervals, metric type (histogram), dimensions, and name remain reliable, while the value (distribution) component is absent or partial for gapped periods—native details are summarized but not lost for on-time data.
B. Histogram extraction makes percentile immune to gaps because buckets are cumulative like counters.
C. Fixed 1-minute resolution stores data at that granularity, eliminating late point issues and preserving every component uniformly.
D. "None" policy forces full MTS recreation on late arrival, duplicating all components and ensuring 100% reliability.

Answer: A

Explanation: Detector resolution drives rollup of native datapoints into 1-minute summaries (histogram bucket aggregation here). With "none" extrapolation, gaps result in no output value for the affected interval, so percentile extraction yields no result or a partial distribution, risking undetected threshold crossings. Reliable components include the MTS-defining dimensions, metric name/type, and timestamps for processed intervals. Impacted is primarily the value/distribution component for gapped periods. This setup requires careful policy selection aligned with data reliability expectations for histogram-based detectors.

Question: 905

In a high-traffic e-commerce platform, the team needs to visualize both the total revenue across all regions and the per-region breakdown simultaneously to identify contribution imbalances. Revenue is a gauge metric.

A. Create one plot with sum:aggregation for the global total, then add a second plot with the same metric using mean:aggregation grouped by region dimension.
B. Apply exclude analytics to remove certain regions from the total calculation plot.
C. Use a single plot with sum:transformation over a 1-hour moving window to derive both total and breakdown.
D. Configure two separate charts instead of combining plots in one view.

Answer: A
Explanation: Combining plots in one chart is efficient: one plot uses aggregation like sum without group by for the overall total, while the second applies the same metric with group by on the region dimension for the breakdown. This allows side-by-side or overlaid insights without multiple charts. Transformations or exclude would not separate total from segmented views properly.

Question: 906

Distinguishing chart types, when is a List chart particularly useful compared to a Heatmap?

A. For showing continuous time-series trends across dimensions.
B. When displaying the latest values for many MTS with optional sparklines and sorting, rather than color-based severity distribution.
C. When needing a matrix view of values over two dimensions.
D. For single latest value display with severity color only.

Answer: B
Explanation: List charts show tabular latest values with sparklines and sorting. Heatmap uses color for severity or values in a grid. Line is for trends. Single Value is for one metric.

Question: 907

To discriminate metadata types effectively in a large environment, an administrator reviews which fields can be used to filter without increasing MTS count. Which statements correctly describe the differences? (Select Two)

A. All metadata types sent with datapoints contribute equally to MTS identification.
B. Dimensions are ingest-time metadata that define MTS uniqueness and directly impact cardinality when varied.
C. Attributes are the only type that can be modified retroactively on historical datapoints.
D. Custom properties and tags are post-ingest metadata that add context for filtering and grouping but do not create new MTS or affect cardinality.

Answer: B,D
Explanation: Dimensions participate in MTS definition at ingest and therefore control cardinality. Custom properties and tags enrich MTS after ingestion for better search, dashboards, and detectors without multiplying the number of series. This distinction is fundamental for cost control and effective metadata strategy in Splunk Observability Cloud. Misunderstanding it often leads to unintended MTS growth or ineffective filters.

Question: 908

In troubleshooting with built-in Kubernetes dashboards, pod CPU charts with "rate" applied show spikes only when viewed at resolutions under 5 minutes. At longer ranges, spikes vanish. Explain the mechanic and recommended practice. (Single-answer)

A. Subscriptions prevent spike visualization in long-range views.
B. Coarser resolutions apply rollups that average the rate over larger intervals, diluting spikes; best practice is to use short ranges or native resolution for burst detection before alerting.
C. Rate function is limited to sub-5-minute resolutions in built-in content.
D. The dashboards enforce Min rollup for CPU regardless of selection.

Answer: B
Explanation: Rate on counters/gauges produces per-interval changes; larger resolution intervals average or sum these, smoothing spikes. Native or short-range viewing preserves burst visibility for accurate interpretation and detector setup in built-in content—a core best practice for rollups and resolution.

Question: 909

What are the three essential components that constitute a "Datapoint" before it is processed into the Splunk IM Metrics Store?

A. Org ID, Metric Name, and Tag set
B. Metric Name, Timestamp, and Value
C. MTS ID, Value, and Rollup Type
D. Metric Name, Dimensions, and Properties

Answer: B
Explanation: A datapoint is the most basic unit of data sent to Splunk IM. It must consist of a Metric Name (what is being measured), a Timestamp (when it was measured), and a Value (the measurement itself). While Dimensions are often included to define the MTS, the fundamental structure of the point itself relies on the name, time, and numerical value.

Question: 910

What does the "Impact of late datapoints" specifically refer to in the context of Splunk Observability Cloud?

A. The permanent loss of data if it arrives after the "Max Delay" window has closed for a detector.
B. The inability of the system to process any data that arrives more than 5 seconds after its timestamp.
C. The financial cost incurred by the organization for exceeding the 1-second resolution limit.
D. The delay between a metric reaching the ingest gateway and it appearing on a chart or being evaluated by a detector.

Answer: D
Explanation: Late datapoints refer to the latency in the data pipeline. This lag impacts real-time monitoring because detectors and charts must wait for this data. If the system is not configured to "wait" (via Max Delay), it evaluates time buckets as empty, leading to incorrect calculations, broken trends, and unreliable alerts.

KILLEXAMS.COM

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.

Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllbus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.