

DUMPSBOSS.

**Splunk Certified Cybersecurity Defense
Engineer**

Splunk SPLK-5002

Version Demo

Total Demo Questions: 11

Total Premium Questions: 113

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

**support@dumpsboss.co
dumpsboss.co**

QUESTION NO: 1

What are the key components of Splunk's indexing process?(Choosethree)

- A. Parsing
- B. Searching
- C. Indexing
- D. Alerting
- E. Input phase

ANSWER: A C E

Explanation:

Splunk's data-processing pipeline begins with the input phase, in which Splunk receives machine data through configured inputs such as files and directories, network ports, scripted inputs, HTTP Event Collector, or forwarders. This stage establishes how data enters the Splunk environment and provides initial source context. Splunk then performs parsing, where it processes the incoming data stream into discrete events. Parsing includes important event-processing activities such as identifying event boundaries, extracting or assigning timestamps, and determining default metadata including host, source, and sourcetype. These steps make the incoming data consistently structured for subsequent storage and search operations.

After parsing, indexing writes the processed events to an index. During this process, Splunk stores both the raw event data and index files that support efficient retrieval at search time. The resulting indexed data can then be searched using SPL, with the assigned metadata available for filtering, routing, and analysis. Therefore, **Input phase, Parsing, and Indexing** describe the core flow by which data is received, prepared as events, and stored for search. Splunk's overview of this pipeline is available in [How Splunk works](#), and its ingestion guidance is documented in [What Splunk can index](#).

QUESTION NO: 2

What Splunk process ensures that duplicate data is not indexed?

- A. Data deduplication
- B. Metadata tagging
- C. Indexer clustering
- D. Event parsing

ANSWER: D

Explanation:

Event parsing is correct because Splunk processes incoming data before committing it to the index, identifying event boundaries and applying parsing-related processing that makes events suitable for indexing. For file-based inputs, this ingestion workflow also works with Splunk's file-tracking mechanism, commonly called the fishbucket. Splunk records information about files it has read, including a checksum-based identifier and the read position, so that it can recognize previously processed content and resume from the appropriate point rather than ingesting the same data again. This behavior is particularly important for monitored log files and file rotation scenarios, where a file may be renamed, revisited, or replaced while Splunk is collecting it. The checksum is based on an initial portion of the file, and Splunk uses that tracking information to determine whether the file is new or has already been handled. Thus, event parsing as part of the data-ingestion pipeline is the applicable process in this question: it prepares event data for indexing while the associated file-tracking state prevents duplicate ingestion. See Splunk's documentation on [how Splunk handles log file rotation and CRCs](#) and the [Splunk Enterprise data-processing pipeline](#).

QUESTION NO: 3

A security team is implementing risk-based alerting for suspicious behavior associated with user accounts. Individual detections are low confidence, but several related detections involving the same user should become actionable when their combined risk is high.

Which practices support an effective risk-based alerting design? (Choose three)

- A. Use a consistent user value as the risk object.
- B. Create a notable event for every low-confidence contributing detection.
- C. Assign risk scores that reflect detection confidence and organizational context.
- D. Assign all risk events to a single generic shared object.
- E. Enrich the user with identity and asset context.

ANSWER: A C E

Explanation:

Using a consistent user value as the risk object, assigning risk scores that reflect detection confidence and organizational context, and enriching the user with identity and asset context support effective risk-based alerting. A consistent risk object allows related risk events to accumulate against the same entity. Calibrated scoring prevents a low-value signal from contributing the same weight as strong evidence or activity involving a privileged identity. Identity and asset context helps analysts understand the significance of the accumulated risk and investigate it efficiently.

Creating a notable event for every low-confidence contributing detection recreates the alert-volume problem that risk-based alerting is intended to reduce. Assigning every risk event to a generic shared object prevents meaningful accumulation and prioritization for the affected user.

QUESTION NO: 4

A security engineer is creating a recurring report for control owners and auditors. The report must demonstrate that security monitoring controls are operating effectively and that exceptions are being addressed.

Which elements should the report include? (Choose three)

- A. Data-source coverage and timeliness information.
- B. An unfiltered export of all raw events.
- C. Measurable control or detection results.
- D. A list of dashboard color selections.
- E. Documented exceptions and remediation status.

ANSWER: A C E

Explanation:

Data-source coverage and timeliness information, measurable control or detection results, and documented exceptions and remediation status provide useful evidence of control operation. Coverage and timeliness show whether the expected telemetry is available for monitoring. Measurable results demonstrate the activity and outcomes of the relevant controls or detections. Exceptions and remediation status show that identified deficiencies are tracked and managed rather than ignored.

A complete unfiltered export of raw events is difficult for control owners and auditors to use as evidence without context or analysis. A list of dashboard color selections does not demonstrate that controls operated effectively.

QUESTION NO: 5

A security team is onboarding firewall data whose events contain the fields `src_ip` and `dest_ip`. Their Enterprise Security correlation searches expect the CIM Network Traffic fields `src` and `dest`.

What is the most appropriate way to map these existing fields without modifying the raw events?

- A. Configure search-time field aliases for the CIM fields.
- B. Configure index-time transformations for the CIM fields.
- C. Create dashboard-specific field extractions for the CIM fields.
- D. Write normalized results to a separate summary index.

ANSWER: A

Explanation:

Configure search-time field aliases for the CIM fields. is correct because field aliases map an existing extracted field to another field name at search time. In this case, aliases can make `src_ip` available as `src` and `dest_ip` available as `dest`, allowing the firewall events to satisfy the normalized field expectations of the applicable CIM data model and Enterprise Security content.

Index-time transformations would change data during ingestion and are unnecessary when the source fields already exist. A dashboard-specific extraction would not make the normalized fields consistently available to correlation searches. Creating a separate summary index would store derived results but would not normalize the original firewall events for general security use.

QUESTION NO: 6

A Splunk administrator is tasked with creating a weekly security report for executives.

What elements should they focus on?

- A. High-level summaries and actionable insights
- B. Detailed logs of every notable event
- C. Excluding compliance metrics to simplify reports
- D. Avoiding visuals to focus on raw data

ANSWER: A

Explanation:

High-level summaries and actionable insights are the appropriate focus because an executive security report must translate operational security data into information that supports leadership decisions. Rather than requiring executives to interpret individual events, the report should communicate the organization's overall security posture: material incidents, notable threat or risk trends, business impact, and whether key controls or compliance objectives need attention. In Splunk, this is commonly delivered through dashboards and scheduled reports that aggregate security-relevant data into understandable metrics and visualizations, such as incident volumes over time, critical detections, affected business services, or remediation status.

The actionable element is essential: leadership should be able to see what needs prioritization, resourcing, or escalation. For example, a report can identify a sustained increase in account attacks and pair that trend with a recommended control improvement or remediation initiative. This approach makes the report concise, outcome-oriented, and useful for governance while retaining links or drilldowns to supporting operational detail when needed. Splunk's security guidance emphasizes using dashboards to monitor and communicate security posture, and the NIST Cybersecurity Framework describes communicating cybersecurity risk information to stakeholders as part of effective governance. See [Splunk Security](#) and [NIST Cybersecurity Framework](#).

QUESTION NO: 7

What does Splunk's term "bucket" refer to in data indexing?

- A. A storage unit for archived data
- B. A collection of events with a specific retention policy
- C. A directory containing indexed data
- D. A database table for search results

ANSWER: C

Explanation:

In Splunk, a bucket is a directory on disk that contains a portion of an index's data and the associated index files. As events are ingested, Splunk writes them into buckets, which progress through lifecycle stages such as hot, warm, cold, frozen, and (when configured) thawed. A bucket contains the raw event data along with index structures that Splunk uses to locate and search that data efficiently. Each index is therefore made up of one or more buckets rather than a single monolithic data store. The bucket lifecycle supports efficient ingestion, searching, storage management, retention, and archival: data can move between storage tiers as it ages while remaining governed by the index configuration. Splunk documents that buckets are directories containing indexed data, and that indexers manage these buckets according to index settings and bucket lifecycle policies. See Splunk's [How Splunk stores indexes](#) and [About buckets and indexer clusters](#) documentation.

QUESTION NO: 8

Which features are crucial for validating integrations in Splunk SOAR? (Choose three)

- A. Testing API connectivity
- B. Monitoring data ingestion rates
- C. Verifying authentication methods
- D. Evaluating automated action performance
- E. Increasing indexer capacity

ANSWER: A C D

Explanation:

Testing API connectivity, verifying authentication methods, and evaluating automated action performance are crucial validation activities for Splunk SOAR integrations. An app integration depends on a configured asset being able to reach the external service or device through its defined endpoint, network path, and API interface. Testing connectivity therefore confirms that SOAR can establish the basic communication required before a playbook relies on the asset.

Verifying authentication methods is equally essential because SOAR apps use asset configuration values such as API keys, tokens, OAuth details, certificates, or credentials to authorize requests. A connectivity test can expose invalid, expired, insufficiently privileged, or incorrectly configured authentication details, allowing them to be corrected before automated response activity is attempted.

Evaluating automated action performance validates that the integration can successfully execute the app actions that playbooks need, such as querying an indicator, creating a ticket, blocking an address, or retrieving endpoint information. Successful action execution should be assessed for accurate outputs, expected completion behavior, and usable results for downstream playbook logic. Splunk's app-development documentation describes connection testing and asset configuration, while its action documentation explains how apps expose actions for playbook use. See [Splunk SOAR app connection documentation](#) and [Splunk SOAR app actions documentation](#).

QUESTION NO: 9

Which features of Splunk are crucial for tuning correlation searches?(Choosethree)

- A. Using thresholds and conditions
- B. Reviewing notable event outcomes
- C. Enabling event sampling
- D. Disabling field extractions
- E. Optimizing search queries

ANSWER: A B E

Explanation:

Using thresholds and conditions is crucial because correlation searches need clear triggering criteria that distinguish meaningful security activity from normal operational behavior. Adjusting counts, time windows, and conditional logic lets defenders reduce alert noise while preserving detections that warrant investigation. Reviewing notable event outcomes provides the feedback loop required for effective tuning: analysts can assess event volume, severity, recurring patterns, and investigation disposition to identify where a detection needs refinement. Splunk Enterprise Security supports managing correlation searches and reviewing the notable events they generate, making operational outcomes a practical basis for improving detection fidelity.

Optimizing search queries is equally important because correlation searches run on a schedule and can materially affect search-head and indexer resources. Efficient SPL, appropriate time constraints, and accelerated data-model searches such as `tstats` where applicable help correlation searches complete reliably within their scheduled intervals. This supports timely notable-event generation as well as scalable detection engineering. Splunk's guidance on correlation searches and search performance reinforces the need to maintain efficient, well-scoped searches: [Splunk Enterprise Security correlation searches](#) and [Splunk search optimization quick tips](#).

QUESTION NO: 10

What is the purpose of using data models in building dashboards?

- A. To store raw data for compliance purposes
- B. To provide a consistent structure for dashboard queries
- C. To compress indexed data
- D. To reduce storage usage on Splunk instances

ANSWER: B

Explanation:

To provide a consistent structure for dashboard queries is correct because Splunk data models organize event data into a defined schema of datasets, fields, and relationships. This lets dashboard authors build searches against a predictable, reusable representation of data instead of repeatedly accounting for differences in raw sourcetypes and field extractions. A data model can also be accelerated, creating summaries that support faster reporting and visualization searches over the model's datasets. In security use cases, data models commonly align data from varied sources to the Common Information Model, enabling a dashboard to use the same normalized fields and dataset constraints regardless of whether events originated from endpoint, network, or authentication products. This improves dashboard consistency, maintainability, and—when acceleration is appropriate—performance. Data models are knowledge objects that describe and normalize data for search and reporting; they are not the primary mechanism for retaining raw events. Splunk documents the structure, use, and acceleration of data models in its [data model overview](#). For security-oriented normalization and common data-model usage, see Splunk's [Common Information Model documentation](#).

QUESTION NO: 11

An engineer observes a delay in data being indexed from a remote location. The universal forwarder is configured correctly. What should they check next?

- A. Review forwarder logs for queue blockages.
- B. Increase the indexer memory allocation.
- C. Optimize search head clustering.
- D. Reconfigure the props.conf file.

ANSWER: A

Explanation:

Review forwarder logs for queue blockages. is the appropriate next diagnostic step because a correctly configured universal forwarder can still be unable to send data promptly when its processing or output queues are blocked. Splunk forwarding uses queues to move events through input, parsing, and output stages. If a downstream indexer, network connection, or receiving pipeline cannot accept data at the required rate, backpressure can cause events to accumulate at the forwarder and produce indexing delay.

The engineer should inspect the universal forwarder's `splunkd.log` for blocked-queue, connection, throttling, or output-related messages. They can also use the forwarder's internal `metrics.log` queue measurements to identify sustained queue growth or blocked status. This establishes whether the delay originates in the forwarding path and provides evidence for the next remediation step, such as investigating receiver availability, throughput, or network behavior. Splunk documents forwarder troubleshooting and the use of internal metrics for diagnosing processing bottlenecks in its [forwarding-data troubleshooting guidance](#) and [metrics.log performance analysis documentation](#).