



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-2002 Practice Questions
SPLK-2002 Practice Test
SPLK-2002 Practice Exam
SPLK-2002 Exam Questions
SPLK-2002 Study Guide



killexams.com

Splunk

SPLK-2002

Splunk Enterprise Certified Architect

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/SPLK-2002>



Question: 1083

You are troubleshooting a Splunk deployment where events from a heavy forwarder are not searchable. The props.conf file defines a custom source type with SHOULD_LINEMERGE = true and a custom LINE_BREAKER. However, events are merged incorrectly, causing search issues. Which configuration change would most effectively resolve this issue?

- A. Set SHOULD_LINEMERGE = false and verify LINE_BREAKER in props.conf
- B. Increase max_events in limits.conf to handle larger events
- C. Adjust TIME_FORMAT in props.conf to improve timestamp parsing
- D. Enable data integrity checking in inputs.conf

Answer: A

Explanation: Incorrect event merging is often caused by SHOULD_LINEMERGE = true when the LINE_BREAKER is sufficient to split events. Setting SHOULD_LINEMERGE = false and verifying the LINE_BREAKER regex in props.conf ensures events are split correctly without unnecessary merging. max_events in limits.conf affects event size, not merging. TIME_FORMAT impacts timestamp parsing, not event boundaries. Data integrity checking in inputs.conf does not address merging issues.

Question: 1084

A single-site indexer cluster with a replication factor of 3 and a search factor of 2 experiences a bucket freeze. What does the cluster master do when a bucket is frozen?

- A. Ensures another copy is made on other peers
- B. Deletes all copies of the bucket
- C. Stops fix-up activities for the bucket
- D. Rolls all copies to frozen immediately

Answer: C

Explanation: When a bucket is frozen in an indexer cluster (e.g., due to retention

policies), the cluster master stops performing fix-up activities for that bucket, such as ensuring replication or search factor compliance. The bucket is no longer actively managed, and its copies age out per retention settings. The cluster master does not create new copies, delete copies, or roll them to frozen immediately.

Question: 1085

In a scenario where you have multiple search heads configured in a clustered environment using the Raft consensus algorithm, how does the algorithm enhance the reliability of search operations? .

- A. It allows for automatic failover to a standby search head if the primary fails
- B. It ensures that all search heads have a synchronized view of the data
- C. It enables the direct indexing of search results to the primary search head
- D. It maintains a log of decisions made by the search heads for auditing purposes

Answer: A, B

Explanation: The Raft consensus algorithm enhances reliability by allowing automatic failover and ensuring that all search heads maintain a synchronized view of the data, which is crucial for consistent search results.

Question: 1086

When implementing search head clustering, which configuration option is essential to ensure that search load is distributed evenly across the available search heads?

- A. Enable load balancing through a search head dispatcher
- B. Use a single search head to avoid confusion
- C. Set up dedicated search heads for each data type
- D. Ensure all search heads have the same hardware specifications

Answer: A

Explanation: Enabling load balancing through a search head dispatcher ensures that search queries are evenly distributed among the search heads, optimizing the performance and efficiency of search operations.

Question: 1087

A Splunk deployment ingests 1.5 TB/day of data from various sources, including HTTP Event Collector (HEC) inputs. The architect needs to ensure that HEC events are indexed with a custom source type based on the client application. Which configuration should be applied?

- A. inputs.conf: [http://hec_input] sourcetype = custom_app
- B. props.conf: [http] TRANSFORMS-sourcetype = set_custom_sourcetype
- C. transforms.conf: [set_custom_sourcetype] REGEX = app_name=client1 DEST_KEY = MetaData:Sourcetype FORMAT = custom_app
- D. inputs.conf: [http://hec_input] token =

Answer: B, C

Explanation: To dynamically set a custom source type for HEC events, props.conf uses TRANSFORMS-sourcetype = set_custom_sourcetype to reference a transform. In transforms.conf, the [set_custom_sourcetype] stanza uses REGEX to match the app_name=client1 field and sets DEST_KEY = MetaData:Sourcetype to assign the custom_app source type. Static sourcetype assignment in inputs.conf is not dynamic. The token setting in inputs.conf is unrelated to source type assignment.

Question: 1088

A telecommunications provider is deploying Splunk Enterprise to monitor its network infrastructure. The Splunk architect is tasked with integrating Splunk with a third-party incident management system that supports REST API calls for ticket creation. The integration requires Splunk to send a POST request with a JSON payload containing network event details whenever a critical issue is detected. The Splunk environment includes a search head cluster and an indexer cluster with a search factor of 3. Which of the following configurations are necessary for this integration?

- A. Develop a custom alert action using a Python script to format the JSON payload and send it to the incident management system's REST API
- B. Configure a webhook in Splunk's alert settings to send event data directly to the incident management system

- C. Install a third-party add-on on the search head cluster to handle authentication and communication with the incident management system
- D. Update the outputs.conf file on the indexers to forward event data to the incident management system's REST API

Answer: A, C

Explanation: A custom alert action with a Python script enables precise JSON payload formatting and secure API calls to the incident management system. A third-party add-on can simplify authentication and communication, if available. Using a webhook without customization is insufficient for complex payload requirements. Updating outputs.conf on indexers is incorrect, as alert actions are managed at the search head level.

Question: 1089

When ingesting network data from different geographical locations, which configuration aspect must be addressed to ensure low-latency data processing and accurate event timestamping?

- A. Utilize edge devices to preprocess data before ingestion
- B. Configure local indexes at each geographical site
- C. Set up a centralized index with global timestamp settings
- D. Adjust the maxLatency parameter to accommodate network delays

Answer: A, B

Explanation: Using edge devices helps preprocess data to minimize latency, and configuring local indexes ensures that data is stored and processed closer to its source.

Question: 1090

You are using the btool command to troubleshoot an issue with a Splunk app configuration. Which command would you use to see a merged view of all configuration files used by the app, including inherited settings from other apps?

- A. splunk btool app list
- B. splunk btool --debug list
- C. splunk btool list --app

D. splunk btool show config

Answer: B

Explanation: Using --debug with the btool command provides a detailed merged view of all configuration files, including inherited settings, which is crucial for troubleshooting.

Question: 1091

A Splunk architect is troubleshooting slow searches on a virtual index that queries HDFS data for a logistics dashboard. The configuration is:

```
[logistics]
vix.provider = hdfs
vix.fs.default.name = hdfs://namenode:8021
vix.splunk.search.splitter = 1500
```

The dashboard search is:

```
index=logistics sourcetype=shipment_logs | timechart span=1h count by status
```

Which of the following will improve search performance?

- A. Reduce vix.splunk.search.splitter to lower MapReduce overhead
- B. Enable vix.splunk.search.cache.enabled = true in indexes.conf
- C. Rewrite the search to use stats instead of timechart for aggregation
- D. Increase vix.splunk.search.mr.maxsplits to allow more parallel tasks

Answer: A, B

Explanation: Reducing vix.splunk.search.splitter decreases the number of MapReduce splits, reducing overhead and improving search performance. Enabling vix.splunk.search.cache.enabled = true caches results, speeding up dashboard refreshes. Rewriting the search to use stats instead of timechart does not significantly improve performance for HDFS virtual indexes, as both commands involve similar processing. Increasing vix.splunk.search.mr.maxsplits creates more splits, potentially increasing overhead and slowing searches.

Question: 1092

In a search head cluster with a deployer, an architect needs to distribute a new app to all members. The app contains non-replicable configurations in server.conf. Which command should be executed on the deployer to propagate these changes?

- A. splunk resync shcluster-replicated-config
- B. splunk apply shcluster-bundle
- C. splunk transfer shcluster-captain
- D. splunk clean raft

Answer: B

Explanation: To distribute a new app with non-replicable configurations (such as server.conf) to search head cluster members, the splunk apply shcluster-bundle command is executed on the deployer. This pushes the configuration bundle to all members, ensuring consistency. The splunk resync shcluster-replicated-config command is for member synchronization, not app distribution. The other options are unrelated to configuration deployment.

Question: 1093

You are tasked with ingesting data from an application that generates XML logs. Which configuration parameter is essential for ensuring that the XML data is parsed correctly and maintains its structure?

- A. Set the sourcetype to a predefined XML format
- B. Adjust the linebreaking setting to accommodate XML tags
- C. Enable auto_sourcetype to simplify the configuration process
- D. Configure the timestamp extraction settings to match XML date formats

Answer: A, B

Explanation: Defining the sourcetype as XML helps with proper parsing rules, while adjusting linebreaking settings ensures that XML tags are correctly handled during ingestion.

Question: 1094

When developing a custom app in Splunk that relies on complex searches and dashboards, which knowledge objects should be prioritized for reuse to enhance maintainability and consistency across the application?

- A. Event types to categorize logs according to specific criteria relevant to the application.
- B. Dashboards that can be dynamically updated based on user input and preferences.
- C. Macros that encapsulate complex search logic for simplified reuse.
- D. Field aliases that allow for standardization of field names across different datasets.

Answer: A, C, D

Explanation: Prioritizing event types, macros, and field aliases enhances maintainability and consistency within the app, allowing for easier updates and standardized data handling.

Question: 1095

At Buttercup Games, a Splunk architect is tasked with optimizing a complex search query that analyzes web access logs to identify users with high latency (response time > 500ms) across multiple data centers. The query must extract the client IP, calculate the average latency per user session, and filter sessions with more than 10 requests, while incorporating a custom field extraction for session_id using the regex pattern session=([a-z0-9]{32}). The dataset is massive, and performance is critical. Which of the following Search Processing Language (SPL) queries is the most efficient and accurate for this requirement?

- A. `sourcetype=web_access | rex field=_raw "session=([a-z0-9]{32})" | stats count, avg(response_time) as avg_latency by client_ip, session_id | where count > 10 AND avg_latency > 500`
- B. `sourcetype=web_access | extract session=([a-z0-9]{32}) | stats count, avg(response_time) as latency by session_id, client_ip | where latency > 500 AND count > 10`
- C. `sourcetype=web_access session=* | rex field=_raw "session=([a-z0-9]{32})" | eventstats avg(response_time) as avg_latency by client_ip, session_id | where avg_latency > 500 AND count > 10`
- D. `sourcetype=web_access | regex session=([a-z0-9]{32}) | stats count(response_time) as count, avg(response_time) as avg_latency by client_ip, session_id | where count > 10 AND avg_latency > 500`

Answer: A

Explanation: The query must efficiently extract the session_id using regex, calculate the average latency, and filter based on count and latency thresholds. The rex command is the correct choice for field extraction from _raw data, as extract is not a valid SPL command and regex filters events rather than extracting fields. The stats command is optimal for aggregating count and average latency by client_ip and session_id. Option A uses rex correctly, applies stats for aggregation, and filters with where, making it both accurate and efficient. Option C uses eventstats, which is less efficient for large datasets due to its event-level processing, and Option D incorrectly uses regex and count(response_time).

Question: 1096

A Splunk architect is managing a single-site indexer cluster with a replication factor of 3 and a search factor of 2. The cluster has four peer nodes, and the daily indexing volume is 400 GB. The architect needs to estimate the storage requirements for one year, assuming buckets are 50% of incoming data size. Which of the following factors are required for the calculation?

- A. Replication factor
- B. Search factor
- C. Number of peer nodes
- D. Daily indexing volume

Answer: A, B, D

Explanation: To estimate storage requirements for an indexer cluster, the replication factor (3) determines the number of bucket copies, the search factor (2) specifies the number of searchable copies (rawdata plus index files), and the daily indexing volume (400 GB, with buckets at 50% size) provides the base data size. The number of peer nodes affects distribution but not the total storage calculation, as storage is driven by replication and search factors.

Question: 1097

A Splunk architect is troubleshooting duplicate events in a deployment ingesting 600 GB/day of syslog data. The inputs.conf file includes:

```
[monitor:///logs/syslog/*.log]
index = syslog
sourcetype = syslog
crcSalt =
```

The architect suspects partial file ingestion due to network issues. Which configurations should the architect implement to prevent duplicates?

- A. Configure CHECK_METHOD = entire_md5
- B. Enable persistent queues on forwarders
- C. Increase replication factor to 3
- D. Add TIME_FORMAT in props.conf

Answer: A, B

Explanation: Configuring CHECK_METHOD = entire_md5 ensures Splunk verifies the entire file's hash, preventing partial ingestion duplicates. Enabling persistent queues buffers data during network issues, ensuring complete ingestion. Increasing the replication factor does not prevent duplicates. Adding TIME_FORMAT aids timestamp parsing but does not address duplicates.

Question: 1098

In a Splunk deployment ingesting 800 GB/day of data from scripted inputs, the architect notices that some events are indexed with incorrect timestamps due to varying time formats in the data. The scripted input generates JSON events with a "log_time" field in formats like "2026-04-21T12:00:00Z" or "04/21/2026 12:00:00". Which props.conf settings should be applied to ensure consistent timestamp extraction?

- A. TIME_PREFIX = "log_time":
- B. TIME_FORMAT = %Y-%m-%dT%H:%M:%SZ
- C. TIME_FORMAT = %m/%d/%Y %H:%M:%S
- D. MAX_TIMESTAMP_LOOKAHEAD = 30

Answer: A, B, D

Explanation: To extract timestamps from the "log_time" field in JSON events, TIME_PREFIX = "log_time": specifies the start of the timestamp. The TIME_FORMAT

= %Y-%m-%dT%H:%M:%SZ handles the ISO8601 format (2026-04-21T12:00:00Z). The MAX_TIMESTAMP_LOOKAHEAD = 30 limits the number of characters Splunk searches for the timestamp, improving performance. The format %m/%d/%Y %H:%M:%S is not sufficient, as it does not cover the ISO8601 format.

Question: 1099

A Splunk architect is optimizing a deployment ingesting 900 GB/day of CSV logs with a 120-day retention period. The cluster has a replication factor of 2 and a search factor of 2. The indexes.conf file includes:

```
[main]
maxTotalDataSizeMB = 1200000
frozenTimePeriodInSecs = 10368000
```

What is the total storage requirement, and which adjustment would most reduce storage?

- A. 32.4 TB; Decrease search factor to 1
- B. 64.8 TB; Decrease replication factor to 1
- C. 32.4 TB; Enable summary indexing
- D. 64.8 TB; Increase maxHotBuckets

Answer: A

Explanation: Storage calculation: $(0.9 \text{ TB} \times 120 \text{ days} \times 0.5) \times (2 + 2 - 1) = 54 \text{ TB} \times 0.6 = 32.4 \text{ TB}$. Decreasing the search factor to 1 reduces tsidx copies, lowering storage significantly. Decreasing the replication factor compromises availability. Summary indexing does not reduce primary storage. Increasing maxHotBuckets affects memory, not storage.

Question: 1100

You are implementing role-based access control (RBAC) in a search head cluster. Which configurations are essential to ensure that users have appropriate access to knowledge objects? .

- A. Assigning roles that define specific permissions for knowledge objects
- B. Ensuring knowledge objects are shared at the app level rather than the user level
- C. Configuring user authentication methods that align with corporate policies
- D. Regularly auditing user access to knowledge objects to ensure compliance

Answer: A, C, D

Explanation: Defining roles and permissions ensures appropriate access control, while aligning authentication methods with policies is crucial for security. Regular audits help maintain compliance with access controls.

Question: 1101

You are troubleshooting a Splunk deployment where a universal forwarder is sending data to an indexer cluster, but events are not appearing in searches. The forwarder is configured to send data to a load-balanced indexer group via `outputs.conf`, and the `Splunkd.log` on the forwarder shows repeated "TcpOutputProc - Connection to indexer:9997 closed. Connection reset by peer" errors. Network connectivity tests confirm that port 9997 is open, and the indexer is receiving other data. Which step should you take to diagnose and resolve this issue?

- A. Run `tcpdump` on the indexer to capture packets on port 9997 and verify the connection handshake
- B. Increase the `maxQueueSize` in `inputs.conf` on the forwarder to buffer more events
- C. Check the indexer's `server.conf` for misconfigured SSL settings
- D. Adjust the forwarder's `limits.conf` to increase `maxKBps` for higher throughput

Answer: A

Explanation: The "Connection reset by peer" error in the forwarder's `Splunkd.log` indicates a network or configuration issue causing the indexer to terminate the connection. Running `tcpdump` on the indexer to capture packets on port 9997 is the most effective diagnostic step, as it allows you to verify the TCP handshake and identify potential issues like packet loss or firewall interference. Increasing `maxQueueSize` in `inputs.conf` addresses buffering but not connection issues. Checking SSL settings in `server.conf` is relevant only if SSL is enabled, which is not indicated. Adjusting `maxKBps` in `limits.conf` affects throughput but does not resolve connection resets.

Question: 1102

A Splunk architect is implementing a custom REST API endpoint to allow external systems to update knowledge objects in Splunk Enterprise. The endpoint is configured in `restmap.conf`:

```
[script:update_knowledge]
match = /update_knowledge
script = update_knowledge.py
requireAuthentication = true
```

The Python script fails to update knowledge objects due to insufficient permissions. Which of the following will resolve the issue?

- A. Grant the `rest_properties_set` capability to the user's role in `authorize.conf`
- B. Ensure the script uses the Splunk SDK's `KnowledgeObjects` class
- C. Configure `allowRemoteAccess = true` in `server.conf`
- D. Set `capability::edit_objects` for the user's role in `authorize.conf`

Answer: A, B

Explanation: Granting the `rest_properties_set` capability in `authorize.conf` allows the user to modify knowledge objects via the REST API. Using the Splunk SDK's `KnowledgeObjects` class ensures the script correctly interacts with Splunk's knowledge object endpoints. The `allowRemoteAccess` setting in `server.conf` is unrelated to REST API permissions. The `edit_objects` capability does not exist in Splunk; knowledge object permissions are managed through REST-specific capabilities.

Question: 1103

A Splunk architect needs to ensure that sensitive information is only accessible to specific roles. What is the most effective method for achieving this through role capabilities?

- A. Create a new index specifically for sensitive data and restrict access.
- B. Use event-level permissions to hide sensitive information.
- C. Configure data masking for sensitive fields.
- D. Apply tags to events for controlled access.

Answer: A, B

Explanation: Creating a new index for sensitive data and applying event-level permissions are effective methods to ensure that sensitive information is only accessible to specific roles.

Question: 1104

In your Splunk environment, you want to create a dashboard that visualizes data trends over time for a specific application. You decide to use the timechart command. Which of the following SPL commands would best suit this purpose?

- A. `index=app_logs | timechart count by status`
- B. `index=app_logs | stats count by time`
- C. `index=app_logs | chart count over time by status`
- D. `index=app_logs | eval timestamp=strftime(_time, "%Y-%m-%d") | stats count by timestamp`

Answer: A

Explanation: The timechart command aggregates data over time and is specifically designed for visualizing trends, making it the best choice for this scenario.

Question: 1105

A Splunk architect is configuring a search pipeline for a dashboard that monitors network latency: `index=network sourcetype=ping_data | eval latency_status=if(latency > 100, "High", "Normal") | stats count by latency_status | sort -count`. The environment has 15 indexers, and the search is executed every 30 seconds, causing high search head load. Which configuration in `limits.conf` can reduce the load?

- A. `max_searches_per_cpu = 2`
- B. `max_events_per_search = 5000`
- C. `scheduler_max_searches = 10`
- D. `max_memtable_bytes = 10000000`

Answer: C

Explanation: The `scheduler_max_searches` parameter in `limits.conf` under the `[scheduler]` stanza limits the number of scheduled searches, reducing the search head load by throttling frequent executions. The `max_searches_per_cpu` parameter limits concurrent searches per CPU, not scheduled searches. The `max_events_per_search` parameter limits events processed, not execution frequency. The `max_memtable_bytes` parameter limits in-memory table sizes, which does not directly reduce load.

Question: 1106

A Splunk architect is configuring Splunk Web security for a deployment with 12 indexers and 5 search heads. The security policy requires TLS 1.3 and a 20-minute session timeout. The architect has a certificate (`web_cert.pem`) and private key (`web_privkey.pem`). Which of the following configurations in `web.conf` will meet these requirements?

- A. `[settings]`
`enableSplunkWebSSL = true`
`privKeyPath = /opt/splunk/etc/auth/web_privkey.pem`
`serverCert = /opt/splunk/etc/auth/web_cert.pem`
`sslVersions = tls1.3`
`sessionTimeout = 20m`
- B. `[settings]`
`enableSplunkWebSSL = true`
`privKeyPath = /opt/splunk/etc/auth/web_privkey.pem`
`certPath = /opt/splunk/etc/auth/web_cert.pem`
`sslVersions = tls1.3`
`sessionTimeout = 1200`
- C. `[settings]`
`enableSplunkWebSSL = true`
`privKeyPath = /opt/splunk/etc/auth/web_privkey.pem`
`serverCert = /opt/splunk/etc/auth/web_cert.pem`
`sslProtocol = tls1.3`
`sessionTimeout = 20`
- D. `[settings]`
`enableSplunkWebSSL = true`
`privKeyPath = /opt/splunk/etc/auth/web_privkey.pem`
`certPath = /opt/splunk/etc/auth/web_cert.pem`
`sslVersions = tls1.3`

sessionTimeout = 20m

Answer: A

Explanation: The [settings] stanza enables SSL (enableSplunkWebSSL = true), specifies the private key (privKeyPath) and certificate (serverCert), restricts to TLS 1.3 (sslVersions = tls1.3), and sets a 20-minute timeout (sessionTimeout = 20m). Incorrect options use certPath, sslProtocol, or incorrect timeout formats.



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 syllabus 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.