Splunk SPLK-2002 Practice Questions

Shared by Whitley on 18-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

When troubleshooting monitor inputs, which command checks the status of the tailed files?

Options:

A- splunk cmd btool inputs list | tail
B- splunk cmd btool check inputs layer
C- curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:FileStatus
D- curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:Tailstatus

Answer:

C

Explanation:

The curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:FileStatus command is used to check the status of the tailed files when troubleshooting monitor inputs. Monitor inputs are inputs that monitor files or directories for new data and send the data to Splunk for indexing. The TailingProcessor:FileStatus endpoint returns information about the files that are being monitored by the Tailing Processor, such as the file name, path, size, position, and status. The splunk cmd btool inputs list | tail command is used to list the inputs configurations from the inputs.conf file and pipe the output to the tail command. The splunk cmd btool check inputs layer command is used to check the inputs configurations for syntax errors and layering. The curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:Tailstatus command does not exist, and it is not a valid endpoint.

Question 2

Question Type: MultipleChoice

(Which option best is a minimum search head specification for a distributed Splunk environment?)

Options:

A- A 1Gb Ethernet NIC, optional 2nd NIC for a management network.
B- An x86 32-bit chip architecture.
C- 128 GB RAM.
D- Two physical CPU cores, or four vCPU at 2GHz or greater speed per core.

Answer:

D

Explanation:

According to the Splunk Enterprise Capacity Planning and Hardware Sizing Guidelines, a distributed Splunk environment's minimum search head specification must ensure that the system can efficiently manage search parsing, ad-hoc query execution, and knowledge object replication. Splunk officially recommends using a 64-bit x86 architecture system with a minimum of two physical CPU cores (or four vCPUs) running at 2 GHz or higher per core for acceptable performance.

Search heads are CPU-intensive components, primarily constrained by processor speed and the number of concurrent searches they must handle. Memory and disk space should scale with user concurrency and search load, but CPU capability remains the baseline requirement. While 128 GB RAM (Option C) is suitable for high-demand or Enterprise Security (ES) deployments, it exceeds the minimum hardware specification for general distributed search environments.

Splunk no longer supports 32-bit architectures (Option B). While a 1Gb Ethernet NIC (Option A) is common, it is not part of the minimum computational specification required by Splunk for search heads. The critical specification is processor capability --- two physical cores or equivalent.

Reference (Splunk Enterprise Documentation):

* Splunk Enterprise Capacity Planning Manual -- Hardware and Performance Guidelines

* Search Head Sizing and System Requirements

* Distributed Deployment Manual -- Recommended System Specifications

* Splunk Hardware and Performance Tuning Guide

Question 3

Question Type: MultipleChoice

How does the average run time of all searches relate to the available CPU cores on the indexers?

Options:

A- Average run time is independent of the number of CPU cores on the indexers.
B- Average run time decreases as the number of CPU cores on the indexers decreases.
C- Average run time increases as the number of CPU cores on the indexers decreases.
D- Average run time increases as the number of CPU cores on the indexers increases.

Answer:

C

Explanation:

The average run time of all searches increases as the number of CPU cores on the indexers decreases. The CPU cores are the processing units that execute the instructions and calculations for the data. The number of CPU cores on the indexers affects the search performance, because the indexers are responsible for retrieving and filtering the data from the indexes. The more CPU cores the indexers have, the faster they can process the data and return the results. The less CPU cores the indexers have, the slower they can process the data and return the results. Therefore, the average run time of all searches is inversely proportional to the number of CPU cores on the indexers. The average run time of all searches is not independent of the number of CPU cores on the indexers, because the CPU cores are an important factor for the search performance. The average run time of all searches does not decrease as the number of CPU cores on the indexers decreases, because this would imply that the search performance improves with less CPU cores, which is not true.The average run time of all searches does not increase as the number of CPU cores on the indexers increases, because this would imply that the search performance worsens with more CPU cores, which is not true

Question 4

Question Type: MultipleChoice

Which props.conf setting has the least impact on indexing performance?

Options:

A- SHOULD_LINEMERGE
B- TRUNCATE
C- CHARSET
D- TIME_PREFIX

Answer:

C

Explanation:

According to the Splunk documentation1, the CHARSET setting in props.conf specifies the character set encoding of the source data. This setting has the least impact on indexing performance, as it only affects how Splunk interprets the bytes of the data, not how it processes or transforms the data. The other options are false because:

The SHOULD_LINEMERGE setting in props.conf determines whether Splunk breaks events based on timestamps or newlines.This setting has a significant impact on indexing performance, as it affects how Splunk parses the data and identifies the boundaries of the events2.

The TRUNCATE setting in props.conf specifies the maximum number of characters that Splunk indexes from a single line of a file.This setting has a moderate impact on indexing performance, as it affects how much data Splunk reads and writes to the index3.

The TIME_PREFIX setting in props.conf specifies the prefix that directly precedes the timestamp in the event data.This setting has a moderate impact on indexing performance, as it affects how Splunk extracts the timestamp and assigns it to the event

Question 5

Question Type: MultipleChoice

What is the recommended order of activities in the Splunk deployment process?

Options:

A- Infrastructure Planning and Buildout
Splunk Deployment and Data Enrichment
User Planning and Rollout
B- User Planning and Rollout
Infrastructure Planning and Buildout
Splunk Deployment and Data Enrichment
C- Splunk Deployment and Data Enrichment
User Planning and Rollout
Infrastructure Planning and Buildout
D- Infrastructure Planning and Buildout
User Planning and Rollout
Splunk Deployment and Data Enrichment

Answer:

A

Explanation:

Splunk's official planning guidance explains that the deployment cycle begins with defining the system architecture, capacity planning, platform design, and topologies. Splunk states that before any component is installed, administrators must complete ''infrastructure planning and buildout,'' which includes determining indexer capacity, search head roles, clustering strategy, storage layout, and performance requirements. This foundational step ensures that all Splunk components have the proper hardware, network design, and scaling expectations.

After the environment is built, Splunk documentation states that the next stage is ''deployment and data onboarding,'' which includes configuring indexers, forwarders, parsing rules, event processing pipelines, data source validation, and enrichment steps such as field extractions, tagging, event types, and CIM alignment. Splunk describes this as the phase where you bring in data and confirm correctness, completeness, and normalization.

Only after the system is stable and populated with data does Splunk recommend ''user planning and rollout'', which includes developing dashboards, roles, knowledge objects, search best practices, and enabling user access. Splunk emphasizes that user onboarding should occur last, once infrastructure and data pipelines are fully validated.

Reference: Splunk Admin & Architect Study Guide; Splunk Deployment Planning Guidelines; Splunk Validated Architectures (Planning and Design Sections).

Question 6

Question Type: MultipleChoice

Which option best is a problem that could be investigated using the Search Job Inspector?

Options:

A- Error messages are appearing underneath the search bar in Splunk Web.
B- Dashboard panels are showing 'Waiting for queued job to start' on page load.
C- Different users are seeing different extracted fields from the same search.
D- Events are not being sorted in reverse chronological order.

Answer:

A

Explanation:

According to the Splunk documentation1, the Search Job Inspector is a tool that you can use to troubleshoot search performance and understand the behavior of knowledge objects, such as event types, tags, lookups, and so on, within the search. You can inspect search jobs that are currently running or that have finished recently. The Search Job Inspector can help you investigate error messages that appear underneath the search bar in Splunk Web, as it can show you the details of the search job, such as the search string, the search mode, the search timeline, the search log, the search profile, and the search properties.You can use this information to identify the cause of the error and fix it2. The other options are false because:

Dashboard panels showing ''Waiting for queued job to start'' on page load is not a problem that can be investigated using the Search Job Inspector, as it indicates that the search job has not started yet. This could be due to the search scheduler being busy or the search priority being low.You can use the Jobs page or the Monitoring Console to monitor the status of the search jobs and adjust the priority or concurrency settings if needed3.

Different users seeing different extracted fields from the same search is not a problem that can be investigated using the Search Job Inspector, as it is related to the user permissions and the knowledge object sharing settings.You can use the Access Controls page or the Knowledge Manager to manage the user roles and the knowledge object visibility4.

Events not being sorted in reverse chronological order is not a problem that can be investigated using the Search Job Inspector, as it is related to the search syntax and the sort command. You can use the Search Manual or the Search Reference to learn how to use the sort command and its options to sort the events by any field or criteria.

Question 7

Question Type: MultipleChoice

A customer currently has many deployment clients being managed by a single, dedicated deployment server. The customer plans to double the number of clients.

What could be done to minimize performance issues?

Options:

A- Modify deploymentclient. conf to change from a Pull to Push mechanism.
B- Reduce the number of apps in the Manager Node repository.
C- Increase the current deployment client phone home interval.
D- Decrease the current deployment client phone home interval.

Answer:

C

Explanation:

According to the Splunk documentation1, increasing the current deployment client phone home interval can minimize performance issues by reducing the frequency of communication between the clients and the deployment server. This can also reduce the network traffic and the load on the deployment server. The other options are false because:

Modifying deploymentclient.conf to change from a Pull to Push mechanism is not possible, as Splunk does not support a Push mechanism for deployment server2.

Reducing the number of apps in the Manager Node repository will not affect the performance of the deployment server, as the apps are only downloaded when there is a change in the configuration or a new app is added3.

Decreasing the current deployment client phone home interval will increase the performance issues, as it will increase the frequency of communication between the clients and the deployment server, resulting in more network traffic and load on the deployment server1.

Question 8

Question Type: MultipleChoice

As of Splunk 9.0, which index records changes to . conf files?

Options:

A- _configtracker
B- _introspection
C- _internal
D- _audit

Answer:

A

Explanation:

This is the index that records changes to .conf files as of Splunk 9.0.According to the Splunk

documentation1, the _configtracker index tracks the changes made to the configuration files on the Splunk platform, such as the files in the etc directory.The _configtracker index can help monitor and troubleshoot the configuration changes, and identify the source and time of the changes1. The other options are not indexes that record changes to .conf files.Option B, _introspection, is an index that records the performance metrics of the Splunk platform, such as CPU, memory, disk, and network usage2.Option C, _internal, is an index that records the internal logs and events of the Splunk platform, such as splunkd, metrics, and audit logs3.Option D, _audit, is an index that records the audit events of the Splunk platform, such as user authentication, authorization, and activity4. Therefore, option A is the correct answer, and options B, C, and D are incorrect.

1: About the _configtracker index2: About the _introspection index3: About the _internal index4: About the _audit index

P2P exams