



Splunk

Exam Questions SPLK-2002

Splunk Enterprise Certified Architect

NEW QUESTION 1

Stakeholders have identified high availability for searchable data as their top priority. Which of the following best addresses this requirement?

- A. Increasing the search factor in the cluster.
- B. Increasing the replication factor in the cluster.
- C. Increasing the number of search heads in the cluster.
- D. Increasing the number of CPUs on the indexers in the cluster.

Answer: B

NEW QUESTION 2

Search dashboards in the Monitoring Console indicate that the distributed deployment is approaching its capacity. Which of the following options will provide the most search performance improvement?

- A. Replace the indexer storage to solid state drives (SSD).
- B. Add more search heads and redistribute users based on the search type.
- C. Look for slow searches and reschedule them to run during an off-peak time.
- D. Add more search peers and make sure forwarders distribute data evenly across all indexers.

Answer: C

NEW QUESTION 3

A Splunk architect has inherited the Splunk deployment at Buttercup Games and end users are complaining that the events are inconsistently formatted for a web sourcetype. Further investigation reveals that not all web logs flow through the same infrastructure: some of the data goes through heavy forwarders and some of the forwarders are managed by another department. Which of the following items might be the cause for this issue?

- A. The search head may have different configurations than the indexers.
- B. The data inputs are not properly configured across all the forwarders.
- C. The indexers may have different configurations than the heavy forwarders.
- D. The forwarders managed by the other department are an older version than the rest.

Answer: D

NEW QUESTION 4

Which of the following should be included in a deployment plan?

- A. Business continuity and disaster recovery plans.
- B. Current logging details and data source inventory.
- C. Current and future topology diagrams of the IT environment.
- D. A comprehensive list of stakeholders, either direct or indirect.

Answer: D

NEW QUESTION 5

Which of the following commands is used to clear the KV store?

- A. splunk clean kvstore
- B. splunk clear kvstore
- C. splunk delete kvstore
- D. splunk reinitialize kvstore

Answer: A

NEW QUESTION 6

Which of the following clarification steps should be taken if apps are not appearing on a deployment client? (Select all that apply.)

- A. Check serverclass.conf of the deployment server.
- B. Check deploymentclient.conf of the deployment client.
- C. Check the content of SPLUNK_HOME/etc/apps of the deployment server.
- D. Search for relevant events in splunkd.log of the deployment server.

Answer: ABC

NEW QUESTION 7

Which command will permanently decommission a peer node operating in an indexer cluster?

- A. splunk stop -f
- B. splunk offline -f
- C. splunk offline --enforce-counts
- D. splunk decommission --enforce counts

Answer: C

NEW QUESTION 8

Which of the following is a good practice for a search head cluster deployer?

- A. The deployer only distributes configurations to search head cluster members when they “phone home”.
- B. The deployer must be used to distribute non-replicable configurations to search head cluster members.
- C. The deployer must distribute configurations to search head cluster members to be valid configurations.
- D. The deployer only distributes configurations to search head cluster members with splunk apply shcluster-bundle.

Answer: A

NEW QUESTION 9

A new Splunk customer is using syslog to collect data from their network devices on port 514. What is the best practice for ingesting this data into Splunk?

- A. Configure syslog to send the data to multiple Splunk indexers.
- B. Use a Splunk indexer to collect a network input on port 514 directly.
- C. Use a Splunk forwarder to collect the input on port 514 and forward the data.
- D. Configure syslog to write logs and use a Splunk forwarder to collect the logs.

Answer: C

NEW QUESTION 10

Before users can use a KV store, an admin must create a collection. Where is a collection is defined?

- A. kvstore.conf
- B. collection.conf
- C. collections.conf
- D. kvcollections.conf

Answer: C

NEW QUESTION 10

Which search head cluster component is responsible for pushing knowledge bundles to search peers, replicating configuration changes to search head cluster members, and scheduling jobs across the search head cluster?

- A. Master
- B. Captain
- C. Deployer
- D. Deployment server

Answer: B

NEW QUESTION 12

Which of the following describe migration from single-site to multisite index replication?

- A. A master node is required at each site.
- B. Multisite policies apply to new data only.
- C. Single-site buckets instantly receive the multisite policies.
- D. Multisite total values should not exceed any single-site factors.

Answer: D

NEW QUESTION 16

Which of the following is a way to exclude search artifacts when creating a diag?

- A. `SPLUNK_HOME/bin/splunk diag --exclude`
- B. `SPLUNK_HOME/bin/splunk diag --debug --refresh`
- C. `SPLUNK_HOME/bin/splunk diag --disable=dispatch`
- D. `SPLUNK_HOME/bin/splunk diag --filter-searchstrings`

Answer: A

NEW QUESTION 21

Which of the following statements describe licensing in a clustered Splunk deployment? (Select all that apply.)

- A. Free licenses do not support clustering.
- B. Replicated data does not count against licensing.
- C. Each cluster member requires its own clustering license.
- D. Cluster members must share the same license pool and license master.

Answer: BD

NEW QUESTION 23

When planning a search head cluster, which of the following is true?

- A. All search heads must use the same operating system.
- B. All search heads must be members of the cluster (no standalone search heads).
- C. The search head captain must be assigned to the largest search head in the cluster.
- D. All indexers must belong to the underlying indexer cluster (no standalone indexers).

Answer: C

NEW QUESTION 27

A search head has successfully joined a single site indexer cluster. Which command is used to configure the same search head to join another indexer cluster?

- A. splunk add cluster-config
- B. splunk add cluster-master
- C. splunk edit cluster-config
- D. splunk edit cluster-master

Answer: B

NEW QUESTION 28

To improve Splunk performance, parallelIngestionPipelines setting can be adjusted on which of the following components in the Splunk architecture? (Select all that apply.)

- A. Indexers
- B. Forwarders
- C. Search head
- D. Cluster master

Answer: AB

NEW QUESTION 29

When adding or decommissioning a member from a Search Head Cluster (SHC), what is the proper order of operations?

- A. 1. Delete Splunk Enterprise, if it exists.2. Install and initialize the instance.3. Join the SHC.
- B. 1. Install and initialize the instance.2. Delete Splunk Enterprise, if it exists.3. Join the SHC.
- C. 1. Initialize cluster rebalance operation.2. Remove master node from cluster.3. Trigger replication.
- D. 1. Trigger replication.2. Remove master node from cluster.3. Initialize cluster rebalance operation.

Answer: B

NEW QUESTION 33

Of the following types of files within an index bucket, which file type may consume the most disk?

- A. Rawdata
- B. Bloom filter
- C. Metadata (.data)
- D. Inverted index (.tsidx)

Answer: B

NEW QUESTION 36

Which of the following is an indexer clustering requirement?

- A. Must use shared storage.
- B. Must reside on a dedicated rack.
- C. Must have at least three members.
- D. Must share the same license pool.

Answer: D

NEW QUESTION 41

Consider a use case involving firewall data. There is no Splunk-supported Technical Add-On, but the vendor has built one. What are the items that must be evaluated before installing the add-on? (Select all that apply.)

- A. Identify number of scheduled or real-time searches.
- B. Validate if this Technical Add-On enables event data for a data model.
- C. Identify the maximum number of forwarders Technical Add-On can support.
- D. Verify if Technical Add-On needs to be installed onto both a search head or indexer.

Answer: AC

NEW QUESTION 44

What is the default log size for Splunk internal logs?

- A. 10MB
- B. 20 MB
- C. 25MB

D. 30MB

Answer: C

NEW QUESTION 47

Because Splunk indexing is read/write intensive, it is important to select the appropriate disk storage solution for each deployment. Which of the following statements is accurate about disk storage?

- A. High performance SAN should never be used.
- B. Enable NFS for storing hot and warm buckets.
- C. The recommended RAID setup is RAID 10 (1 + 0).
- D. Virtualized environments are usually preferred over bare metal for Splunk indexers.

Answer: C

NEW QUESTION 52

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

SPLK-2002 Practice Exam Features:

- * SPLK-2002 Questions and Answers Updated Frequently
- * SPLK-2002 Practice Questions Verified by Expert Senior Certified Staff
- * SPLK-2002 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SPLK-2002 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SPLK-2002 Practice Test Here](#)