



Splunk

Exam Questions SPLK-2002

Splunk Enterprise Certified Architect

NEW QUESTION 1

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 2

Which of the following are client filters available in serverclass.conf? (Select all that apply.)

- A. DNS name.
- B. IP address.
- C. Splunk server role.
- D. Platform (machine type).

Answer: AB

NEW QUESTION 3

Which Splunk server role regulates the functioning of indexer cluster?

- A. Indexer
- B. Deployer
- C. Master Node
- D. Monitoring Console

Answer: C

NEW QUESTION 4

Which of the following security options must be explicitly configured (i.e. which options are not enabled by default)?

- A. Data encryption between Splunk Web and splunkd.
- B. Certificate authentication between forwarders and indexers.
- C. Certificate authentication between Splunk Web and search head.
- D. Data encryption for distributed search between search heads and indexers.

Answer: B

NEW QUESTION 5

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 6

A customer plans to ingest 600 GB of data per day into Splunk. They will have six concurrent users, and they also want high data availability and high search performance. The customer is concerned about cost and wants to spend the minimum amount on the hardware for Splunk. How many indexers are recommended for this deployment?

- A. Two indexers not in a cluster, assuming users run many long searches.
- B. Three indexers not in a cluster, assuming a long data retention period.
- C. Two indexers clustered, assuming high availability is the greatest priority.
- D. Two indexers clustered, assuming a high volume of saved/scheduled searches.

Answer: D

NEW QUESTION 7

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 8

Which of the following statements describe a Search Head Cluster (SHC) captain? (Select all that apply.)

- A. Is the job scheduler for the entire SHC.
- B. Manages alert action suppressions (throttling).
- C. Synchronizes the member list with the KV store primary.
- D. Replicates the SHC's knowledge bundle to the search peers.

Answer: AD

NEW QUESTION 9

Which search will show all deployment client messages from the client (UF)?

- A. index=_audit component=DC* host=<ds> | stats count by message
- B. index=_audit component=DC* host=<uf> | stats count by message
- C. index=_internal component= DC* host=<uf> | stats count by message
- D. index=_internal component=DS* host=<ds> | stats count by message

Answer: D

NEW QUESTION 10

When Splunk indexes data in a non clustered environment, what kind of files does it create by default?

- A. Index and .tsidx files.
- B. Rawdata and index files.
- C. Compressed and .tsidx files.
- D. Compressed and meta data files.

Answer: B

NEW QUESTION 10

In search head clustering, which of the following methods can you use to transfer captaincy to a different member? (Select all that apply.)

- A. Use the Monitoring Console.
- B. Use the Search Head Clustering settings menu from Splunk Web on any member.
- C. Run the splunk transfer shcluster-captain command from the current captain.
- D. Run the splunk transfer shcluster-captain command from the member you would like to become the captain.

Answer: BD

NEW QUESTION 15

What does setting site=site0 on all Search Head Cluster members do in a multi-site indexer cluster?

- A. Disables search site affinity.
- B. Sets all members to dynamic captaincy.
- C. Enables multisite search artifact replication.
- D. Enables automatic search site affinity discovery.

Answer: A

NEW QUESTION 19

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

In which phase of the Splunk Enterprise data pipeline are indexed extraction configurations processed?

- A. Input
- B. Search
- C. Parsing
- D. Indexing

Answer: C

NEW QUESTION 27

When converting from a single-site to a multi-site cluster, what happens to existing single-site clustered buckets?

- A. They will continue to replicate within the origin site and age out based on existing policies.
- B. They will maintain replication as required according to the single-site policies, but never age out.

- C. They will be replicated across all peers in the multi-site cluster and age out based on existing policies.
- D. They will stop replicating within the single-site and remain on the indexer they reside on and age out according to existing policies.

Answer: B

NEW QUESTION 29

What is the algorithm used to determine captaincy in a Splunk search head cluster?

- A. Raft distributed consensus.
- B. Rapt distributed consensus.
- C. Rift distributed consensus.
- D. Round-robin distribution consensus.

Answer: A

NEW QUESTION 31

As a best practice, where should the internal licensing logs be stored?

- A. Indexing layer.
- B. License server.
- C. Deployment layer.
- D. Search head layer.

Answer: D

NEW QUESTION 36

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

- 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

NEW QUESTION 40

Which two sections can be expanded using the Search Job Inspector?

- A. Execution costs.
- B. Saved search history.
- C. Search job properties.
- D. Optimization suggestions.

Answer: BC

NEW QUESTION 43

.....

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](#)