

SPLK-2002 Dumps

Splunk Enterprise Certified Architect

<https://www.certleader.com/SPLK-2002-dumps.html>



NEW QUESTION 1

Which of the following will cause the greatest reduction in disk size requirements for a cluster of N indexers running Splunk Enterprise Security?

- A. Setting the cluster search factor to N-1.
- B. Increasing the number of buckets per index.
- C. Decreasing the data model acceleration range.
- D. Setting the cluster replication factor to N-1.

Answer: D

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

What does the deployer do in a Search Head Cluster (SHC)? (Select all that apply.)

- A. Distributes apps to SHC members.
- B. Bootstraps a clean Splunk install for a SHC.
- C. Distributes non-search related and manual configuration file changes.
- D. Distributes runtime knowledge object changes made by users across the SHC.

Answer: A

NEW QUESTION 4

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 5

Which Splunk Enterprise offering has its own license?

- A. Splunk Cloud Forwarder
- B. Splunk Heavy Forwarder
- C. Splunk Universal Forwarder
- D. Splunk Forwarder Management

Answer: C

NEW QUESTION 6

Which component in the splunkd.log will log information related to bad event breaking?

- A. Audittrail
- B. EventBreaking
- C. IndexingPipeline
- D. AggregatorMiningProcessor

Answer: D

NEW QUESTION 7

Which Splunk server role regulates the functioning of indexer cluster?

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

Answer: C

NEW QUESTION 8

When adding or rejoining a member to a search head cluster, the following error is displayed:

Error pulling configurations from the search head cluster captain; consider performing a destructive configuration resync on this search head cluster member.

What corrective action should be taken?

- A. Restart the search head.
- B. Run the splunk apply shcluster-bundle command from the deployer.
- C. Run the clean raft command on all members of the search head cluster.
- D. Run the splunk resync shcluster-replicated-config command on this member.

Answer: B

NEW QUESTION 9

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 10

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 10

Which CLI command converts a Splunk instance to a license slave?

- A. splunk add licenses
- B. splunk list licenser-slaves
- C. splunk edit licenser-localslave
- D. splunk list licenser-localslave

Answer: C

NEW QUESTION 15

Which of the following can a Splunk diag contain?

- A. Search history, Splunk users and their roles, running processes, indexed data
- B. Server specs, current open connections, internal Splunk log files, index listings
- C. KV store listings, internal Splunk log files, search peer bundles listings, indexed data
- D. Splunk platform configuration details, Splunk users and their roles, current open connections, index listings

Answer: B

NEW QUESTION 17

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 21

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 26

To optimize the distribution of primary buckets; when does primary rebalancing automatically occur? (Select all that apply.)

- A. Rolling restart completes.
- B. Master node rejoins the cluster.
- C. Captain joins or rejoins cluster.
- D. A peer node joins or rejoins the cluster.

Answer: ABD

NEW QUESTION 27

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 31

Configurations from the deployer are merged into which location on the search head cluster member?

- A. SPLUNK_HOME/etc/system/local
- B. SPLUNK_HOME/etc/apps/APP_HOME/local
- C. SPLUNK_HOME/etc/apps/search/default
- D. SPLUNK_HOME/etc/apps/APP_HOME/default

Answer: A

NEW QUESTION 33

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 37

How does IT Service Intelligence (ITSI) impact the planning of a Splunk deployment?

- A. ITSI requires a dedicated deployment server.
- B. The amount of users using ITSI will not impact performance.
- C. ITSI in a Splunk deployment does not require additional hardware resources.
- D. Depending on the Key Performance Indicators that are being tracked, additional infrastructure may be needed.

Answer: D

NEW QUESTION 40

In the deployment planning process, when should a person identify who gets to see network data?

- A. Deployment schedule
- B. Topology diagramming
- C. Data source inventory
- D. Data policy definition

Answer: C

NEW QUESTION 41

A Splunk instance has the following settings in SPLUNK_HOME/etc/system/local/server.conf:

[clustering] mode = master

replication_factor = 2

pass4SymmKey = password123

Which of the following statements

describe this Splunk instance?

(Select all that apply.)

- A. This is a multi-site cluster.
- B. This cluster's search factor is 2.
- C. This Splunk instance needs to be restarted.
- D. This instance is missing the master_uri attribute.

Answer: AC

NEW QUESTION 43

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 48

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 51

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 52

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

- 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

NEW QUESTION 53

Which of the following is a best practice to maximize indexing performance?

- A. Use automatic sourcetype.
- B. Use the Splunk default settings.
- C. Not use pre-trained source types.
- D. Minimize configuration generality.

Answer: D

NEW QUESTION 54

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 56

Splunk configuration parameter settings can differ between multiple .conf files of the same name contained within different apps. Which of the following directories has the highest precedence?

- A. System local directory.
- B. System default directory.
- C. App local directories, in ASCII order.
- D. App default directories, in ASCII order.

Answer: A

NEW QUESTION 59

Which of the following statements about integrating with third-party systems is true? (Select all that apply.)

- A. A Hadoop application can search data in Splunk.
- B. Splunk can search data in the Hadoop File System (HDFS).

- C. You can use Splunk alerts to provision actions on a third-party system.
- D. You can forward data from Splunk forwarder to a third-party system without indexing it first.

Answer: CD

NEW QUESTION 61

In a distributed environment, knowledge object bundles are replicated from the search head to which location on the search peer(s)?

- A. SPLUNK_HOME/var/lib/searchpeers
- B. SPLUNK_HOME/var/log/searchpeers
- C. SPLUNK_HOME/var/run/searchpeers
- D. SPLUNK_HOME/var/spool/searchpeers

Answer: C

NEW QUESTION 66

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 70

When configuring a Splunk indexer cluster, what are the default values for replication and search factor?

- A. replication_factor = 2 search_factor = 2
- B. replication_factor = 2 searchfactor = 3
- C. replication_factor = 3 search_factor = 2
- D. replication_factor = 3 searchfactor = 3

Answer: A

NEW QUESTION 73

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 75

What is the default log size for Splunk internal logs?

- A. 10MB
- B. 20 MB
- C. 25MB
- D. 30MB

Answer: C

NEW QUESTION 77

Which of the following options can improve reliability of syslog delivery to Splunk? (Select all that apply.)

- A. Use TCP syslog.
- B. Configure UDP inputs on each Splunk indexer to receive data directly.
- C. Use a network load balancer to direct syslog traffic to active backend syslog listeners.
- D. Use one or more syslog servers to persist data with a Universal Forwarder to send the data to Splunk indexers.

Answer: CD

NEW QUESTION 78

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 81

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SPLK-2002 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SPLK-2002-dumps.html>