

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

A customer has installed a 500GB Enterprise license. They also purchased and installed a 300GB, no enforcement license on the same license master. How much data can the customer ingest before search is locked out?

- A. 300G
- B. After this limit, search is locked out
- C. B.500G
- D. After this limit, search is locked out.
- E. 800G
- F. After this limit, search is locked out.
- G. Search is not locked out
- H. Violations are still recorded.

Answer: D

NEW QUESTION 3

When using the props.conf LINE_BREAKER attribute to delimit multi-line events, the SHOULD_LINEMERGE attribute should be set to what?

- A. Auto
- B. None
- C. True
- D. False

Answer: C

NEW QUESTION 4

A multi-site indexer cluster can be configured using which of the following? (Select all that apply.)

- A. Via Splunk Web.
- B. Directly edit SPLUNK_HOME/etc/system/local/server.conf
- C. Run a splunk edit cluster-config command from the CLI.
- D. Directly edit SPLUNK_HOME/etc/system/default/server.conf

Answer: AB

NEW QUESTION 5

Which Splunk tool offers a health check for administrators to evaluate the health of their Splunk deployment?

- A. btool
- B. DiagGen
- C. SPL Clinic
- D. Monitoring Console

Answer: D

NEW QUESTION 6

In a four site indexer cluster, which configuration stores two searchable copies at the origin site, one searchable copy at site2, and a total of four searchable copies?

- A. site_search_factor = origin:2, site1:2, total:4
- B. site_search_factor = origin:2, site2:1, total:4
- C. site_replication_factor = origin:2, site1:2, total:4
- D. site_replication_factor = origin:2, site2:1, total:4

Answer: D

NEW QUESTION 7

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 8

Indexing is slow and real-time search results are delayed in a Splunk environment with two indexers and one search head. There is ample CPU and memory available on the indexers. Which of the following is most likely to improve indexing performance?

- A. Increase the maximum number of hot buckets in indexes.conf
- B. Increase the number of parallel ingestion pipelines in server.conf
- C. Decrease the maximum size of the search pipelines in limits.conf
- D. Decrease the maximum concurrent scheduled searches in limits.conf

Answer: D

NEW QUESTION 9

In an existing Splunk environment, the new index buckets that are created each day are about half the size of the incoming data. Within each bucket, about 30% of the space is used for rawdata and about 70% for index files.

What additional information is needed to calculate the daily disk consumption, per indexer, if indexer clustering is implemented?

- A. Total daily indexing volume, number of peer nodes, and number of accelerated searches.
- B. Total daily indexing volume, number of peer nodes, replication factor, and search factor.
- C. Total daily indexing volume, replication factor, search factor, and number of search heads.
- D. Replication factor, search factor, number of accelerated searches, and total disk size across cluster.

Answer: D

NEW QUESTION 10

A three-node search head cluster is skipping a large number of searches across time. What should be done to increase scheduled search capacity on the search head cluster?

- A. Create a job server on the cluster.
- B. Add another search head to the cluster.
- C. server.conf captain_is_adhoc_searchhead = true.
- D. Change limits.conf value for max_searches_per_cpu to a higher value.

Answer: D

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 13

What is the minimum reference server specification for a Splunk indexer?

- A. 12 CPU cores, 12GB RAM, 800 IOPS
- B. 16 CPU cores, 16GB RAM, 800 IOPS
- C. 24 CPU cores, 16GB RAM, 1200 IOPS
- D. 28 CPU cores, 32GB RAM, 1200 IOPS

Answer: A

NEW QUESTION 14

Which of the following artifacts are included in a Splunk diag file? (Select all that apply.)

- A. OS settings.
- B. Internal logs.
- C. Customer data.
- D. Configuration files.

Answer: BD

NEW QUESTION 16

Splunk Enterprise platform instrumentation refers to data that the Splunk Enterprise deployment logs in the _introspection index. Which of the following logs are included in this index? (Select all that apply.)

- A. audit.log
- B. metrics.log

- C. disk_objects.log
- D. resource_usage.log

Answer: CD

NEW QUESTION 19

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 20

Which of the following are true statements about Splunk indexer clustering?

- A. All peer nodes must run exactly the same Splunk version.
- B. The master node must run the same or a later Splunk version than search heads.
- C. The peer nodes must run the same or a later Splunk version than the master node.
- D. The search head must run the same or a later Splunk version than the peer nodes.

Answer: B

NEW QUESTION 24

At which default interval does metrics.log generate a periodic report regarding license utilization?

- A. 10 seconds
- B. 30 seconds
- C. 60 seconds
- D. 300 seconds

Answer: B

NEW QUESTION 27

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 32

Which Splunk internal index contains licenserelated events?

- A. _audit
- B. _license
- C. _internal
- D. _introspection

Answer: C

NEW QUESTION 33

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 37

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 42

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 46

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 51

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 53

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 54

The KV store forms its own cluster within a SHC. What is the maximum number of SHC members KV store will form?

- A. 25
- B. 50
- C. 100
- D. Unlimited

Answer: D

NEW QUESTION 59

Which command is used for thawing the archive bucket?

- A. Splunk collect
- B. Splunk convert
- C. Splunk rebuild
- D. Splunk dbinspect

Answer: C

NEW QUESTION 62

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 63

Which tool(s) can be leveraged to diagnose connection problems between an indexer and forwarder? (Select all that apply.)

- A. telnet

- B. tcpdump
- C. splunk btool
- D. splunk btprobe

Answer: BC

NEW QUESTION 64

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 67

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 70

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 74

When should multiple search pipelines be enabled?

- A. Only if disk IOPS is at 800 or better.
- B. Only if there are fewer than twelve concurrent users.
- C. Only if running Splunk Enterprise version 6.6 or later.
- D. Only if CPU and memory resources are significantly under-utilized.

Answer: D

NEW QUESTION 77

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 81

Which of the following should be done when installing Enterprise Security on a Search Head Cluster? (Select all that apply.)

- A. Install Enterprise Security on the deployer.
- B. Install Enterprise Security on a staging instance.
- C. Copy the Enterprise Security configurations to the deployer.
- D. Use the deployer to deploy Enterprise Security to the cluster members.

Answer: AD

NEW QUESTION 83

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 86

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 89

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 90

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

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

Answer: A

NEW QUESTION 95

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 99

Which of the following statements describe search head clustering? (Select all that apply.)

- A. A deployer is required.
- B. At least three search heads are needed.
- C. Search heads must meet the high-performance reference server requirements.
- D. The deployer must have sufficient CPU and network resources to process service requests and push configurations.

Answer: AC

NEW QUESTION 101

Which of the following tasks should the architect perform when building a deployment plan? (Select all that apply.)

- A. Use case checklist.
- B. Install Splunk apps.
- C. Inventory data sources.
- D. Review network topology.

Answer: D

NEW QUESTION 103

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 104

.....

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>