Vendor: Splunk

Exam Code:SPLK-2002

Exam Name: Splunk Enterprise Certified Architect

Version: Demo

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.

Correct Answer: B

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture

---

QUESTION 2

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.

Correct Answer: C

Reference: https://wiki.splunk.com/Community:BestPracticeForConfiguringSyslogInput

---

QUESTION 3

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

Correct Answer: C

QUESTION 4

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

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.1/Data/
Troubleshoottheinputprocess#Troubleshoot_your_tailed_files

QUESTION 5

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.

Correct Answer: A

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/DeploymultisiteSHC

QUESTION 6

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

Correct Answer: A

Reference: https://answers.splunk.com/answers/237859/can-i-delete-all-data-from-a-kv-store-at-once.html

---

QUESTION 7

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).

Correct Answer: C

QUESTION 8

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.

Correct Answer: D

Reference: https://answers.splunk.com/answers/617608/can-we-increase-parallelingestionpipelines-in-ahe.html

---

QUESTION 9

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.

Correct Answer: B

QUESTION 10

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.

Correct Answer: AC

QUESTION 11

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.

Correct Answer: ABD

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Rebalancethecluster

QUESTION 12

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.

Correct Answer: B