Splunk

Exam Questions SPLK-2002

Splunk Enterprise Certified Architect

NEW QUESTION 1
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 2
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 3
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 4
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 5
Which of the following is true regarding Splunk Enterprise performance? (Select all that apply.)

A. Adding search peers increases the maximum size of search results.
B. Adding RAM to an existing search heads provides additional search capacity.
C. Adding search peers increases the search throughput as search load increases.
D. Adding search heads provides additional CPU cores to run more concurrent searches.

Answer: BD

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
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
The guidance Splunk gives for estimating size on for syslog data is 50% of original data size. How does this divide between files in the index?

A. rawdata is: 10%, tsidx is: 40%
B. rawdata is: 15%, tsidx is: 35%
C. rawdata is: 35%, tsidx is: 15%
D. rawdata is: 40%, tsidx is: 10%

Answer: B

NEW QUESTION 9
The frequency in which a deployment client contacts the deployment server is controlled by what?

A. polling_interval attribute in outputs.conf
B. phoneHomeIntervalInSecs attribute in outputs.conf
C. polling_interval attribute in deploymentclient.conf
D. phoneHomeIntervalInSecs attribute in deploymentclient.conf

Answer: D

NEW QUESTION 10
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 10
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 11
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 14
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 16
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 19
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 21
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 22
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 25
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 27
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 30
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 31
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 33

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 35
Which of the following is a best practice to maximize indexing performance?

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

Answer: D

NEW QUESTION 39
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 43
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 46
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 48
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 51
What is a Splunk Job? (Select all that apply.)

A. A user-defined Splunk capability.
B. Searches that are subjected to some usage quota.
C. A search process kicked off via a report or an alert.
D. A child OS process manifested from the splunkd process.

Answer: A

NEW QUESTION 54
When Splunk is installed, where are the internal indexes
stored by default?

A. SPLUNK_HOME/bin
B. SPLUNK_HOME/var/lib
C. SPLUNK_HOME/var/run
D. SPLUNK_HOME/etc/system/default

Answer: B

NEW QUESTION 58
What is the logical first step when starting a deployment plan?

A. Inventory the currently deployed logging infrastructure.
B. Determine what apps and use cases will be implemented.
C. Gather statistics on the expected adoption of Splunk for sizing.
D. Collect the initial requirements for the deployment from all stakeholders.

Answer: D

NEW QUESTION 62
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 65
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 67
......

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](https://www.surepassexam.com/SPLK-2002-exam-dumps.html)