SPLK-3003 Q&As

Splunk Core Certified Consultant

Pass Splunk SPLK-3003 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/splk-3003.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Splunk Official Exam Center

QUESTION 1

A customer has three users and is planning to ingest 250GB of data per day. They are concerned with search uptime, can tolerate up to a two-hour downtime for the search tier, and want advice on single search head versus a search head cluster. (SHC).

Which recommendation is the most appropriate?

A. The customer should deploy two active search heads behind a load balancer to support HA.

B. The customer should deploy a SHC with a single member for HA; more members can be added later.

C. The customer should deploy a SHC, because it will be required to support the high volume of data.

D. The customer should deploy a single search head with a warm standby search head and an rsync process to synchronize configurations.

Correct Answer: D

QUESTION 2

When a bucket rolls from cold to frozen on a clustered indexer, which of the following scenarios occurs?

A. All replicated copies will be rolled to frozen; original copies will remain.

B. Replicated copies of the bucket will remain on all other indexers and the Cluster Master (CM) assigns a new primary bucket.

C. The bucket rolls to frozen on all clustered indexers simultaneously.

D. Nothing. Replicated copies of the bucket will remain on all other indexers until a local retention rule causes it to roll.

Correct Answer: B

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Bucketsandclusters

QUESTION 3

Monitoring Console (MC) health check configuration items are stored in which configuration file?

A. healthcheck.conf

B. alert_actions.conf

C. distsearch.conf

D. checklist.conf

Correct Answer: D

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/Customizehealthcheck

QUESTION 4

In an environment that has Indexer Clustering, the Monitoring Console (MC) provides dashboards to monitor environment health. As the environment grows over time and new indexers are added, which steps would ensure the MC is aware of the additional indexers?

A. No changes are necessary, the Monitoring Console has self-configuration capabilities.

B. Using the MC setup UI, review and apply the changes.

C. Remove and re-add the cluster master from the indexer clustering UI page to add new peers, then apply the changes under the MC setup UI.

D. Each new indexer needs to be added using the distributed search UI, then settings must be saved under the MC setup UI.

Correct Answer: B

QUESTION 5

A customer wants to migrate from using Splunk local accounts to use Active Directory with LDAP for their Splunk user accounts instead. Which configuration files must be modified to connect to an Active Directory LDAP provider?

A. authentication.conf, authorize.conf, ldap.conf

B. authentication.conf, ldap.conf

C. authentication.conf

D. authorize.conf, authentication.conf

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Security/ ConfigureLDAPwithconfigurationfiles

QUESTION 6

Where does the bloomfilter reside?

A. $SPLUNK_HOME/var/lib/splunk/indexfoo/db/db_1553504858_1553504507_8

B. $SPLUNK_HOME/var/lib/splunk/indexfoo/db/db_1553504858_1553504507_8/*.tsidx

C. $SPLUNK_HOME/var/lib/splunk/fishbucket

D. $SPLUNK_HOME/var/lib/splunk/indexfoo/db/db_1553504858_1553504507_8/rawdata

Correct Answer: C

QUESTION 7

A customer has the following Splunk instances within their environment: An indexer cluster consisting of a cluster master/master node and five clustered indexers, two search heads (no search head clustering), a deployment server, and a license master. The deployment server and license master are running on their own single-purpose instances. The customer would like to start using the Monitoring Console (MC) to monitor the whole environment.

On the MC instance, which instances will need to be configured as distributed search peers by specifying them via the UI using the settings menu?

A. Just the cluster master/master node.

B. Indexers, search heads, deployment server, license master, cluster master/master node.

C. Search heads, deployment server, license master, cluster master/master node

D. Deployment server, license master

Correct Answer: C

QUESTION 8

An index receives approximately 50GB of data per day per indexer at an even and consistent rate. The customer would like to keep this data searchable for a minimum of 30 days. In addition, they have hourly scheduled searches that process a week\\'s worth of data and are quite sensitive to search performance.

Given ideal conditions (no restarts, nor drops/bursts in data volume), and following PS best practices, which of the following sets of indexes.conf settings can be leveraged to meet the requirements?

A. frozenTimePeriodInSecs, maxDataSize, maxVolumeDataSizeMB, maxHotBuckets

B. maxDataSize, maxTotalDataSizeMB, maxHotBuckets, maxGlobalDataSizeMB

C. maxDataSize, frozenTimePeriodInSecs, maxVolumeDataSizeMB

D. frozenTimePeriodInSecs, maxWarmDBCount, homePath.maxDataSizeMB, maxHotSpanSecs

Correct Answer: B

QUESTION 9

In the diagrammed environment shown below, the customer would like the data read by the universal forwarders to set an indexed field containing the UF\\'s host name. Where would the parsing configurations need to be installed for this to work?

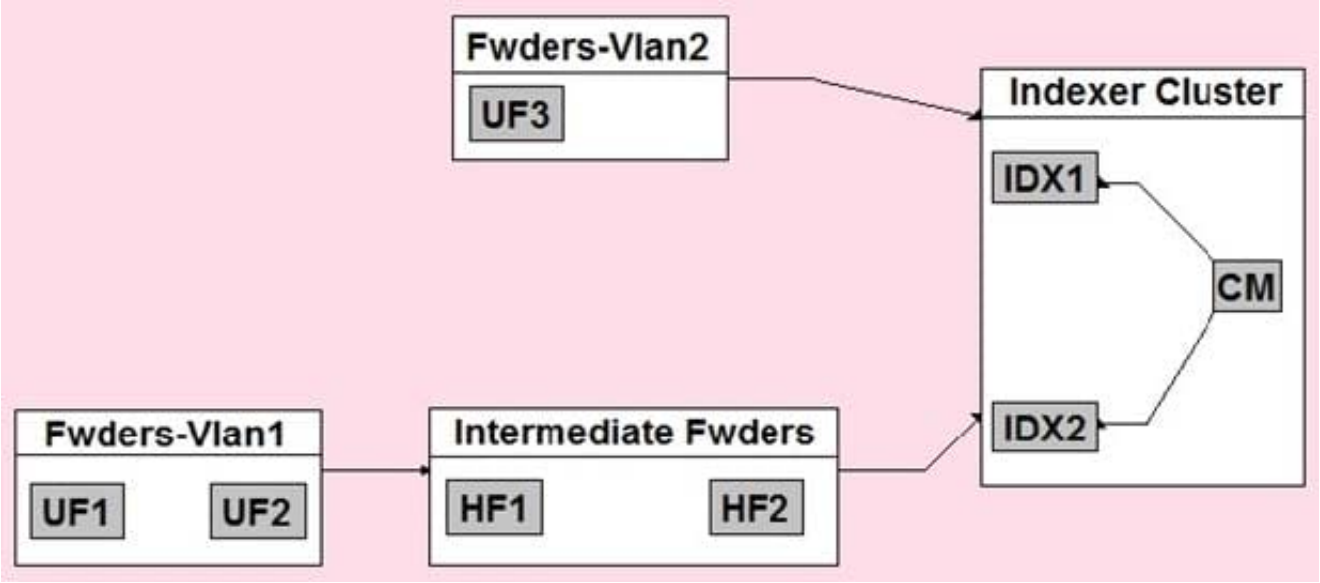

A. All universal forwarders.

B. Only the indexers.

C. All heavy forwarders.

D. On all parsing Splunk instances.

Correct Answer: D

QUESTION 10

A working search head cluster has been set up and used for 6 months with just the native/local Splunk user authentication method. In order to integrate the search heads with an external Active Directory server using LDAP, which of the following statements represents the most appropriate method to deploy the configuration to the servers?

A. Configure the integration in a base configuration app located in shcluster-apps directory on the search head deployer, then deploy the configuration to the search heads using the splunk apply shclusterbundle command.

B. Log onto each search using a command line utility. Modify the authentication.conf and authorize.conf files in a base configuration app to configure the integration.

C. Configure the LDAP integration on one Search Head using the Settings > Access Controls > Authentication Method and Settings > Access Controls > Roles Splunk UI menus. The configuration setting will replicate to the other nodes in the search head cluster eliminating the need to do this on the other search heads.

D. On each search head, login and configure the LDAP integration using the Settings > Access Controls > Authentication Method and Settings > Access Controls > Roles Splunk UI menus.

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Security/ConfigureLDAPwithSplunkWeb

QUESTION 11

A new search head cluster is being implemented. Which is the correct command to initialize the deployer node without restarting the search head cluster peers?

A. $SPLUNK_HOME/bin/splunk apply shcluster-bundle

B. $SPLUNK_HOME/bin/splunk apply cluster-bundle

C. $SPLUNK_HOME/bin/splunk apply shcluster-bundle -action stage D. $SPLUNK_HOME/bin/splunk apply cluster-bundle -action stage

Correct Answer: A

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/ PropagateSHCconfigurationchanges

QUESTION 12

What is required to setup the HTTP Event Collector (HEC)?

A. Each HEC input requires a unique name but token values can be shared.

B. Each HEC input requires an existing forwarder output group.

C. Each HEC input entry must contain a valid token.

D. Each HEC input requires a Source name field.

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/UsetheHTTPEventCollector

QUESTION 13

Where are Splunk Data Model Acceleration (DMA) summaries stored?

A. In tstatsHomePath

B. In the .tsidx files.

C. In summaryHomePath

D. In journal.gz

Correct Answer: A

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Knowledge/ Acceleratedatamodels#:~:text=Splunk%20software%20creates%20ad%20hoc,your%20indexes%20alongside%20index%20buckets

QUESTION 14

A customer with a large distributed environment has blacklisted a large lookup from the search bundle to decrease the bundle size using distsearch.conf. After this change, when running searches utilizing the lookup that was blacklisted they see error messages in the Splunk Search UI stating the lookup file does not exist.

What can the customer do to resolve the issue?

A. The search needs to be modified to ensure the lookup command specifies parameter local=true.

B. The blacklisted lookup definition stanza needs to be modified to specify setting allow_caching=true.

C. The search needs to be modified to ensure the lookup command specified parameter blacklist=false.

D. The lookup cannot be blacklisted; the change must be reverted.

Correct Answer: A

QUESTION 15

A customer is migrating their existing Splunk Indexer from an old set of hardware to a new set of indexers. What is the earliest method to migrate the system?

A. 1. Add new indexers to the cluster as peers, in the same site (if needed).

2.

Ensure new indexers receive common configuration.

3.

Decommission old indexers (one at a time) to allow time for CM to fix/migrate buckets to new

hardware.

4.

Remove all the old indexers from the CM\\'s list.

B. 1. Add new indexers to the cluster as peers, to a new site.

2.

Ensure new indexers receive common configuration from the CM.

3.

Decommission old indexers (one at a time) to allow time for CM to fix/migrate buckets to new

hardware.

4.

Remove all the old indexers from the CM\\'s list.

C. 1. Add new indexers to the cluster as peers, in the same site.

2.

Update the replication factor by +1 to Instruct the cluster to start replicating to new peers.

3.

Allow time for CM to fix/migrate buckets to new hardware.

4.

Remove all the old indexers from the CM\\'s list.

D. 1. Add new indexers to the cluster as new site.

2.

Update cluster master (CM) server.conf to include the new available site.

3.

Allow time for CM to fix/migrate buckets to new hardware.

4.

Remove the old indexers from the CM\\'s list.

Correct Answer: B