SPLK-3003$^{Q\&As}$

Splunk Core Certified Consultant

Pass Splunk SPLK-3003 Exam with 100% Guarantee

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

QUESTION 1

The customer wants to migrate their current Splunk Index cluster to new hardware to improve indexing and search performance. What is the correct process and procedure for this task?

A. 1. Install new indexers.

2.

Configure indexers into the cluster as peers; ensure they receive the same configuration via the deployment server.

3.

Decommission old peers one at a time.

4.

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

5.

Update forwarders to forward to the new peers.

B. 1. Install new indexers.

2.

Configure indexers into the cluster as peers; ensure they receive the cluster bundle and the same configuration as original peers.

3.

Decommission old peers one at a time.

4.

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

5.

Update forwarders to forward to the new peers.

C. 1. Install new indexers.

2.

Configure indexers into the cluster as peers; ensure they receive the same configuration via the deployment server.

3.

Update forwarders to forward to the new peers.

4.

Decommission old peers on at a time.

5.

Restart the cluster master (CM).

D. 1. Install new indexers.

2.

Configure indexers into the cluster as peers; ensure they receive the cluster bundle and the same configuration as original peers.

3.

Update forwarders to forward to the new peers.

4.

Decommission old peers one at a time.

5.

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

Correct Answer: C

QUESTION 2

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 3

When can the Search Job Inspector be used to debug searches?

A. If the search has not expired.

B. If the search is currently running.

C. If the search has been queued.

D. If the search has expired.

Correct Answer: A

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Search/ ViewsearchjobpropertieswiththeJobInspector

QUESTION 4

Which of the following processor occur in the indexing pipeline?

A. tcp out, syslog out

B. Regex replacement, annotator

C. Aggregator

D. UTF-8, linebreaker, header

Correct Answer: D

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/
Howindexingworks#Event_processing_and_the_data_pipeline

QUESTION 5

A customer has a multisite cluster (two sites, each site in its own data center) and users experiencing a slow response when searches are run on search heads located in either site. The Search Job Inspector shows the delay is being caused by search heads on either site waiting for results to be returned by indexers on the opposing site. The network team has confirmed that there is limited bandwidth available between the two data centers, which are in different geographic locations.

Which of the following would be the least expensive and easiest way to improve search performance?

A. Configure site_search_factor to ensure a searchable copy exists in the local site for each search head.

B. Move all indexers and search heads in one of the data centers into the same site.

C. Install a network pipe with more bandwidth between the two data centers.

D. Set the site setting on each indexer in the server.conf clustering stanza to be the same for all indexers regardless of site.

Correct Answer: A

QUESTION 6

How could a role in which all users must specify an index=clause in all searches be configured?

A. Set the authorize.conf setting: srchIndexesDefault to no value.

B. Set the authorize.conf setting: srchFilter to no value.

C. Set the authorize.conf setting: srchIndexesAllowed to no value.

D. Set the authorize.conf setting: srchJobsQuota to no value.

Correct Answer: B

QUESTION 7

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 8

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 9

When utilizing a subsearch within a Splunk SPL search query, which of the following statements is accurate?

A. Subsearches have to be initiated with the | subsearch command.

B. Subsearches can only be utilized with | inputlookup command.

C. Subsearches have a default result output limit of 10000.

D. There are no specific limitations when using subsearches.

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/8.0.6/Search/Aboutsubsearches#:~:text=By%20default%2C%20subsearches%20return%20a,will%20timeout%20before%20it%20completes

QUESTION 10

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 11

A customer has 30 indexers in an indexer cluster configuration and two search heads. They are working on writing SPL search for a particular use-case, but are concerned that it takes too long to run for short time durations.

How can the Search Job Inspector capabilities be used to help validate and understand the customer concerns?

A. Search Job Inspector provides statistics to show how much time and the number of events each indexer has processed.

B. Search Job Inspector provides a Search Health Check capability that provides an optimized SPL query the customer should try instead.

C. Search Job Inspector cannot be used to help troubleshoot the slow performing search; customer should review index=_introspection instead.

D. The customer is using the transaction SPL search command, which is known to be slow.

Correct Answer: A

QUESTION 12

The customer has an indexer cluster supporting a wide variety of search needs, including scheduled search, data model acceleration, and summary indexing. Here is an excerpt from the cluster mater\\'s server.conf:

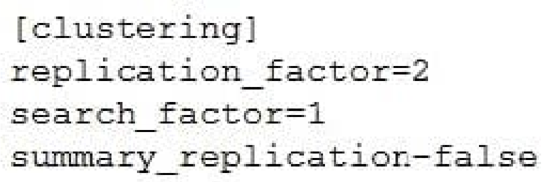

```
[clustering]
replication_factor=2
search_factor=1
summary_replication-false
```

Which strategy represents the minimum and least disruptive change necessary to protect the searchability of the indexer cluster in case of indexer failure?

A. Enable maintenance mode on the CM to prevent excessive fix-up and bring the failed indexer back online.

B. Leave replication_factor=2, increase search_factor=2 and enable summary_replication.

C. Convert the cluster to multi-site and modify the server.conf to be site_replication_factor=2, site_search_factor=2.

D. Increase replication_factor=3, search_factor=2 to protect the data, and allow there to always be a searchable copy.

Correct Answer: D

QUESTION 13

When monitoring and forwarding events collected from a file containing unstructured textual events, what is the difference in the Splunk2Splunk payload traffic sent between a universal forwarder (UF) and indexer compared to the Splunk2Splunk payload sent between a heavy forwarder (HF) and the indexer layer? (Assume that the file is being monitored locally on the forwarder.)

A. The payload format sent from the UF versus the HF is exactly the same. The payload size is identical because they\\'re both sending 64K chunks.

B. The UF sends a stream of data containing one set of medata fields to represent the entire stream, whereas the HF sends individual events, each with their own metadata fields attached, resulting in a lager payload.

C. The UF will generally send the payload in the same format, but only when the sourcetype is specified in the inputs.conf and EVENT_BREAKER_ENABLE is set to true.

D. The HF sends a stream of 64K TCP chunks with one set of metadata fields attached to represent the entire stream, whereas the UF sends individual events, each with their own metadata fields attached.

Correct Answer: B

QUESTION 14

A customer has a search cluster (SHC) of six members split evenly between two data centers (DC). The customer is concerned with network connectivity between the two DCs due to frequent outages. Which of the following is true as it relates to SHC resiliency when a network outage occurs between the two DCs?

A. The SHC will function as expected as the SHC deployer will become the new captain until the network communication is restored.

B. The SHC will stop all scheduled search activity within the SHC.

C. The SHC will function as expected as the minimum required number of nodes for a SHC is 3.

D. The SHC will function as expected as the SHC captain will fall back to previous active captain in the remaining site.

Correct Answer: D

QUESTION 15

How does Monitoring Console (MC) initially identify the server role(s) of a new Splunk Instance?

A. The MC uses a REST endpoint to query the server.

B. Roles are manually assigned within the MC.

C. Roles are read from distsearch.conf.

D. The MC assigns all possible roles by default.

Correct Answer: C