Splunk

Exam SPLK-3003

Splunk Core Certified Consultant

Version: 3.0

[Total Questions: 85]

Question No : 1

The Splunk Validated Architectures (SVAs) document provides a series of approved Splunk topologies. Which statement accurately describes how it should be used by a customer?

A. Customer should look at the category tables, pick the highest number that their budget permits, then select this design topology as the chosen design.
B. Customers should identify their requirements, provisionally choose an approved design that meets them, then consider design principles and best practices to come to an informed design decision.
C. Using the guided requirements gathering in the SVAs document, choose a topology that suits requirements, and be sure not to deviate from the specified design.
D. Choose an SVA topology code that includes Search Head and Indexer Clustering because it offers the highest level of resilience.

Answer: B
Reference: https://www.splunk.com/en_us/blog/tips-and-tricks/splunk-validated-architectures.html

Question No : 2

In which of the following scenarios is a subsearch the most appropriate?

A. When joining results from multiple indexes.
B. When dynamically filtering hosts.
C. When filtering indexed fields.
D. When joining multiple large datasets.

Answer: A

Question No : 3

Report acceleration has been enabled for a specific use case. In which bucket location is the corresponding CSV file located?

A. thawedPath
B. summaryHomePath

C. tstatsHomePath
D. homePath, coldPath

Answer: B
Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Knowledge/Manageacceleratedsearchsummaries

Question No : 4

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

Answer: B

Question No : 5

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.

Answer: B

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

Question No : 6

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.

Answer: B

Question No : 7

Which of the following statements applies to indexer discovery?

A. The Cluster Master (CM) can automatically discover new indexers added to the cluster.
B. Forwarders can automatically discover new indexers added to the cluster.
C. Deployment servers can automatically configure new indexers added to the cluster.
D. Search heads can automatically discover new indexers added to the cluster.

Answer: D

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

Question No : 8

A customer wants to understand how Splunk bucket types (hot, warm, cold) impact search performance within their environment. Their indexers have a single storage device for all data. What is the proper message to communicate to the customer?

A. The bucket types (hot, warm, or cold) have the same search performance characteristics within the customer's environment.
B. While hot, warm, and cold buckets have the same search performance characteristics within the customers environment, due to their optimized structure, the thawed buckets are the most performant.
C. Searching hot and warm buckets result in best performance because by default the cold buckets are miniaturized by removing TSIDX files to save on storage cost.
D. Because the cold buckets are written to a cheaper/slower storage volume, they will be slower to search compared to hot and warm buckets which are written to Solid State Disk (SSD).

Answer: D

Question No : 9

A customer has downloaded the Splunk App for AWS from Splunkbase and installed it in a search head cluster following the instructions using the deployer. A power user modifies a dashboard in the app on one of the search head cluster members. The app containing an updated dashboard is upgraded to the latest version by following the instructions via the deployer.

What happens?

A. The updated dashboard will not be deployed globally to all users, due to the conflict with the power user's modified version of the dashboard.
B. Applying the search head cluster bundle will fail due to the conflict.
C. The updated dashboard will be available to the power user.
D. The updated dashboard will not be available to the power user; they will see their modified version.

Answer: A

Question No : 10

A customer has a new set of hardware to replace their aging indexers. What method would reduce the amount of bucket replication operations during the migration process?

A. Disable the indexing ports on the old indexers.
B. Disable replication ports on the old indexers.
C. Put the old indexers into manual detention.
D. Put the old indexers into automatic detention.

Answer: D

Question No : 11

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.

Answer: A

Question No : 12

A customer has been using Splunk for one year, utilizing a single/all-in-one instance. This single Splunk server is now struggling to cope with the daily ingest rate. Also, Splunk has become a vital system in day-to-day operations making high availability a consideration for the Splunk service. The customer is unsure how to design the new environment topology in order to provide this.

Which resource would help the customer gather the requirements for their new architecture?

A. Direct the customer to the docs.splunk.com and tell them that all the information to help them select the right design is documented there.
B. Ask the customer to engage with the sales team immediately as they probably need a larger license.
C. Refer the customer to answers.splunk.com as someone else has probably already designed a system that meets their requirements.
D. Refer the customer to the Splunk Validated Architectures document in order to guide them through which approved architectures could meet their requirements.

Answer: D
Reference: https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf

Question No : 13

As data enters the indexer, it proceeds through a pipeline where event processing occurs. In which pipeline does line breaking occur?

A. Indexing
B. Typing
C. Merging
D. Parsing

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

Question No : 14

Which of the following server.conf stanzas indicates the Indexer Discovery feature has not been fully configured (restart pending) on the Master Node?

A.
```
[indexer_discovery]
pass4SymmKey = $7$XcXl1lu3820Jbui14oVe324+mvx6gCKKv6kf2zEaVB6Ie4DcZ647CnLVlfW
```

B.
```
[clustering]
mode = master
pass4SymmKey = $7$tYTXzke+1r+3DULTHHDUTmYOXdtZJPxm21XwMARrJE20jsmicp9C3ni0
```

C.
```
[indexer_discovery]
pass4SymmKey = idxdiscovery
```

D.
```
[clustering]
mode = forwarder
pass4SymmKey = $7$PU9SBXww63Vz3UJdDYGIN0UrdscRh83ssC2pEpwE6P3gn50iNF094g==
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C
Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/indexerdiscovery

Question No : 15

A customer is using both internal Splunk authentication and LDAP for user management.

If a username exists in both $SPLUNK_HOME/etc/passwd and LDAP, which of the following statements is accurate?

A. The internal Splunk authentication will take precedence.
B. Authentication will only succeed if the password is the same in both systems.
C. The LDAP user account will take precedence.
D. Splunk will error as it does not support overlapping usernames

Answer: A

Question No : 16

A customer is having issues with truncated events greater than 64K. What configuration should be deployed to a universal forwarder (UF) to fix the issue?

A. None. Splunk default configurations will process the events as needed; the UF is not causing truncation.
B. Configure the best practice magic 6 or great 8 props.conf settings.
C. EVENT_BREAKER_ENABLE and EVENT_BREAKER regular expression settings per sourcetype.
D. Global EVENT_BREAKER_ENABLE and EVENT_BREAKER regular expression settings.

Answer: C
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/Resolvedataqualityissues

Question No : 17

A customer wants to implement LDAP because managing local Splunk users is becoming too much of an overhead. What configuration details are needed from the customer to implement LDAP authentication?

A. API: Python script with PAM/RADIUS details.
B. LDAP server: port, bind user credentials, path/to/groups, path/to/user.
C. LDAP server: port, bind user credentials, base DN for groups, base DN for users.
D. LDAP REST details, base DN for groups, base DN for users.

Answer: C
Reference: https://www.learnsplunk.com/splunk-ldap-authentication-configuration.html

Question No : 18

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

Answer: C

Question No : 19

A Splunk Index cluster is being installed and the indexers need to be configured with a license master. After the customer provides the name of the license master, what is the next step?

A. Enter the license master configuration via Splunk web on each indexer before disabling Splunk web.
B. Update /opt/splunk/etc/master-apps/_cluster/default/server.conf on the cluster master and apply a cluster bundle.
C. Update the Splunk PS base config license app and copy to each indexer.
D. Update the Splunk PS base config license app and deploy via the cluster master.

Answer: C

Question No : 20

When using SAML, where does user authentication occur?

A. Splunk generates a SAML assertion that authenticates the user.
B. The Service Provider (SP) decodes the SAML request and authenticates the user.
C. The Identity Provider (IDP) decodes the SAML request and authenticates the user.
D. The Service Provider (SP) generates a SAML assertion that authenticates the user.

Answer: A

Question No : 21

When setting up a multisite search head and indexer cluster, which nodes are required to declare site membership?

A. Search head cluster members, deployer, indexers, cluster master
B. Search head cluster members, deployment server, deployer, indexers, cluster master
C. All splunk nodes, including forwarders, must declare site membership
D. Search head cluster members, indexers, cluster master

Answer: D
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/SHCandindexercluster