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.passapply.com/splk-3003.html

100% Passing Guarantee
100% Money Back Assurance

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

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

What is the default push mode for a search head cluster deployer app configuration bundle?

A. full

B. merge_to_default

C. default_only

D. local_only

Correct Answer: B

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/PropagateSHCconfigurationchanges#:~:text=The%20deployer%20push%20mode%20determines,default%20push%20mode%20is%20merge_to_default%20

QUESTION 2

Which of the following server roles should be configured for a host which indexes its internal logs locally?

A. Cluster master

B. Indexer

C. Monitoring Console (MC)

D. Search head

Correct Answer: B

Reference: https://community.splunk.com/t5/Deployment-Architecture/How-to-identify-Splunk-Instancerole-by-internal-logs/m-p/365555

QUESTION 3

A customer has a network device that transmits logs directly with UDP or TCP over SSL. Using PS best practices, which ingestion method should be used?

A. Open a TCP port with SSL on a heavy forwarder to parse and transmit the data to the indexing tier.

B. Open a UDP port on a universal forwarder to parse and transmit the data to the indexing tier.

C. Use a syslog server to aggregate the data to files and use a heavy forwarder to read and transmit the data to the indexing tier.

D. Use a syslog server to aggregate the data to files and use a universal forwarder to read and transmit the data to the indexing tier.

Correct Answer: D

QUESTION 4

A site from a multi-site indexer cluster needs to be decommissioned. Which of the following actions must be taken?

A. Nothing. Decommissioning a site is not possible.

B. Create an alias for where the new data should be sent.

C. Remove the site from the list of available sites.

D. Remove the site from the list of available sites and create an alias for where the new data should be sent.

Correct Answer: D

QUESTION 5

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 6

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 7

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 8

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 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?

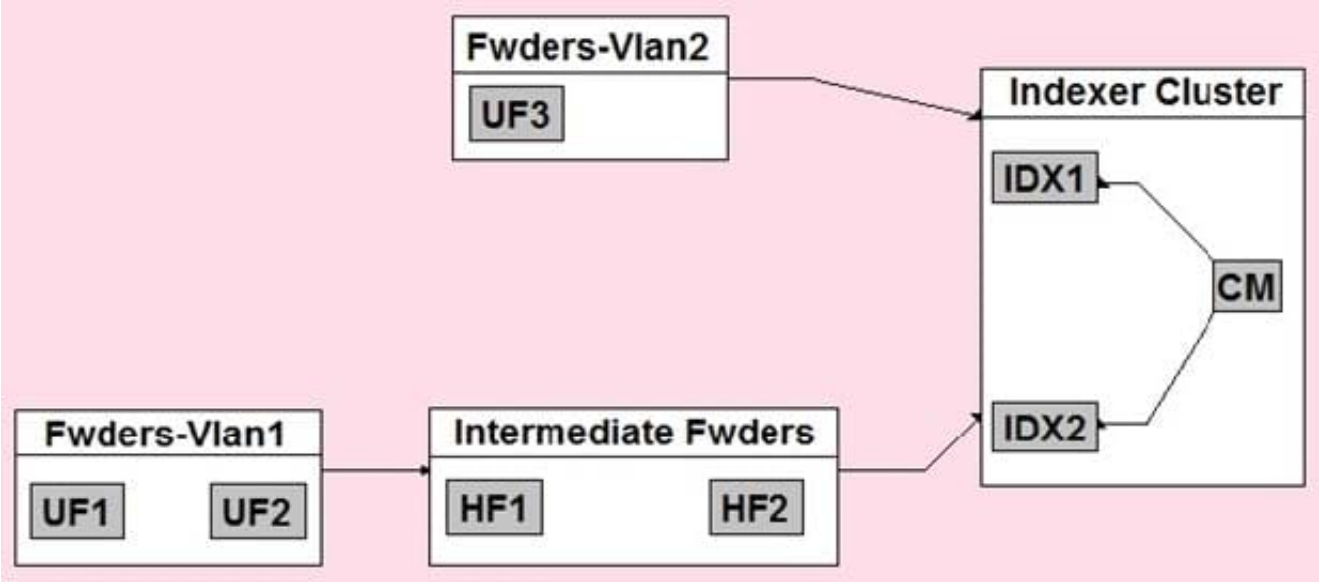

A. All universal forwarders.

B. Only the indexers.

C. All heavy forwarders.

D. On all parsing Splunk instances.

Correct Answer: D

QUESTION 10

A 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).

Correct Answer: D

QUESTION 11

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 12

A customer has a Universal Forwarder (UF) with an inputs.conf monitoring its splunkd.log. The data is sent through a heavy forwarder to an indexer. Where does the Index time parsing occur?

A. Indexer

B. Universal forwarder

C. Search head

D. Heavy forwarder

Correct Answer: D

Reference: https://www.learnsplunk.com/splunk-interview-questions.html

QUESTION 13

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 14

A customer has a number of inefficient regex replacement transforms being applied. When under heavy load the indexers are struggling to maintain the expected indexing rate. In a worst case scenario, which queue(s) would be

expected to fill up?

A. Typing, merging, parsing, input

B. Parsing

C. Typing

D. Indexing, typing, merging, parsing, input

Correct Answer: B

QUESTION 15

Which command is most efficient in finding the pass4SymmKey of an index cluster?

A. find / -name server.conf -print | grep pass4SymKey

B. $SPLUNK_HOME/bin/splunk search | rest splunk_server=local /servicesNS/-/unhash_app/storage/ passwords

C. $SPLUNK_HOME/bin/splunk btool server list clustering | grep pass4SymmKey

D. $SPLUNK_HOME/bin/splunk btool clustering list clustering --debug | grep pass4SymmKey

Correct Answer: D

Reference: https://community.splunk.com/t5/Deployment-Architecture/Which-instance-or-configuration-filein-my-Splunk-environment/m-p/241486