SPLK-2002^Q&As^

Splunk Enterprise Certified Architect

Pass Splunk SPLK-2002 Exam with 100% Guarantee

QUESTION 1

Indexing is slow and real-time search results are delayed in a Splunk environment with two indexers and one search head. There is ample CPU and memory available on the indexers. Which of the following is most likely to improve indexing performance?

A. Increase the maximum number of hot buckets in indexes.conf

B. Increase the number of parallel ingestion pipelines in server.conf

C. Decrease the maximum size of the search pipelines in limits.conf

D. Decrease the maximum concurrent scheduled searches in limits.conf

Correct Answer: D

QUESTION 2

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.

Correct Answer: D

QUESTION 3

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

Correct Answer: D

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Sitereplicationfactor

QUESTION 4

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.

Correct Answer: B

QUESTION 5

Splunk configuration parameter settings can differ between multiple .conf files of the same name contained within different apps. Which of the following directories has the highest precedence?

A. System local directory.

B. System default directory.

C. App local directories, in ASCII order.

D. App default directories, in ASCII order.

Correct Answer: A

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Admin/Wheretofindtheconfigurationfiles

QUESTION 6

Before users can use a KV store, an admin must create a collection. Where is a collection is defined?

A. kvstore.conf

B. collection.conf

C. collections.conf

D. kvcollections.conf

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Knowledge/ DefineaKVStorelookupinSplunkWeb

QUESTION 7

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.

Correct Answer: AC

QUESTION 8

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).

Correct Answer: C

QUESTION 9

Which Splunk Enterprise offering has its own license?

A. Splunk Cloud Forwarder

B. Splunk Heavy Forwarder

C. Splunk Universal Forwarder

D. Splunk Forwarder Management

Correct Answer: C

Reference: https://docs.splunk.com/Splexicon:Forwardinglicense

QUESTION 10

Which command will permanently decommission a peer node operating in an indexer cluster?

A. splunk stop -f

B. splunk offline -f

C. splunk offline --enforce-counts

D. splunk decommission --enforce counts

Correct Answer: C

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Takeapeeroffline

QUESTION 11

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%

Correct Answer: B

Reference: https://answers.splunk.com/answers/147951/what-is-the-compression-ratio-of-raw-data-insplunk.html

QUESTION 12

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.

Correct Answer: B

Reference: https://answers.splunk.com/answers/760348/search-head-version-compatibility.html

QUESTION 13

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

Correct Answer: A

Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Capacity/Referencehardware#Reference_host_specification

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.

Correct Answer: C

Reference: https://wiki.splunk.com/Community:BestPracticeForConfiguringSyslogInput

QUESTION 15

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.

Correct Answer: D