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

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 2

A customer has asked for a five-node search head cluster (SHC), but does not have the storage budget to use a replication factor greater than 2. They would like to understand what might happen in terms of the users\\' ability to view historic scheduled search results if they log onto a search head which doesn\\'t contain one of the 2 copies of a given search artifact.

Which of the following statements best describes what would happen in this scenario?

A. The search head that the user has logged onto will proxy the required artifact over to itself from a search head that currently holds a copy. A copy will also be replicated from that search head permanently, so it is available for future use.

B. Because the dispatch folder containing the search results is not present on the search head, the user will not be able to view the search results.

C. The user will not be able to see the results of the search until one of the search heads is restarted, forcing synchronization of all dispatched artifacts across all search heads.

D. The user will not be able to see the results of the search until the Splunk administrator issues the apply shcluster-bundle command on the search head deployer, forcing synchronization of all dispatched artifacts across all search heads.

Correct Answer: A

QUESTION 3

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.

Correct Answer: C

QUESTION 4

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

Correct Answer: B

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

QUESTION 5

Consider the search shown below.

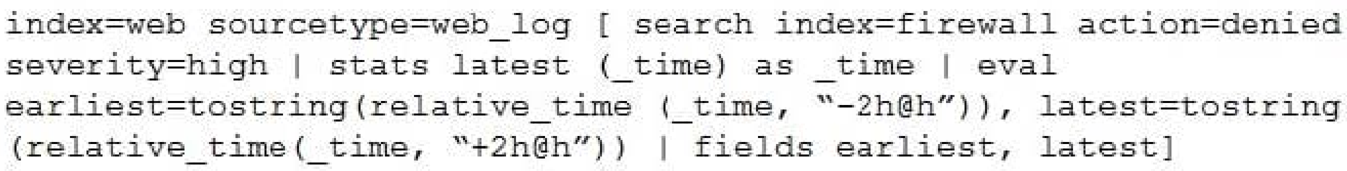

```
index=web sourcetype=web_log [ search index=firewall action=denied
severity=high | stats latest (_time) as _time | eval
earliest=tostring(relative_time (_time, “-2h@h”)), latest=tostring
(relative_time(_time, “+2h@h”)) | fields earliest, latest]
```

What is this search\\'s intended function?

A. To return all the web_log events from the web index that occur two hours before and after the most recent high severity, denied event found in the firewall index.

B. To find all the denied, high severity events in the firewall index, and use those events to further search for lateral movement within the web index.

C. To return all the web_log events from the web index that occur two hours before and after all high severity, denied events found in the firewall index.

D. To search the firewall index for web logs that have been denied and are of high severity.

Correct Answer: C

QUESTION 6

The universal forwarder (UF) should be used whenever possible, as it is smaller and more efficient. In which of the following scenarios would a heavy forwarder (HF) be a more appropriate choice?

A. When a predictable version of Python is required.

B. When filtering 10% - 5% of incoming events.

C. When monitoring a log file.

D. When running a script.

Correct Answer: B

Reference: https://www.splunk.com/en_us/blog/tips-and-tricks/universal-or-heavy-that-is-the-question.html

QUESTION 7

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

Correct Answer: D

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

QUESTION 8

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 9

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 10

What happens when an index cluster peer freezes a bucket?

A. All indexers with a copy of the bucket will delete it.

B. The cluster master will ensure another copy of the bucket is made on the other peers to meet the replication settings.

C. The cluster master will no longer perform fix-up activities for the bucket.

D. All indexers with a copy of the bucket will immediately roll it to frozen.

Correct Answer: C

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

QUESTION 11

In which directory should base config app(s) be placed to initialize an indexer?

A. $SPLUNK_HOME/etc/

B. $SPLUNK_HOME/etc/apps

C. $SPLUNK_HOME/etc/system/local

D. $SPLUNK_HOME/etc/slave-apps

Correct Answer: B

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

QUESTION 12

In addition to the normal responsibilities of a search head cluster captain, which of the following is a default behavior?

A. The captain is not a cluster member and does not perform normal search activities.

B. The captain is a cluster member who performs normal search activities.

C. The captain is not a cluster member but does perform normal search activities.

D. The captain is a cluster member but does not perform normal search activities.

Correct Answer: B

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/
SHCarchitecture#Search_head_cluster_captain

QUESTION 13

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

Correct Answer: D

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

QUESTION 14

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 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