Version: 6.1

Question: 1

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.

Answer: C

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.

Answer: A

Question: 3

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

Answer: D

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/Customizehealthcheck

Question: 4

What should be considered when running the following CLI commands with a goal of accelerating an index cluster migration to new hardware?

```
$SPLUNK_HOME/bin/splunk edit cluster-config -max_peer_build_load 3

$SPLUNK_HOME/bin/splunk edit cluster-config -max_peer_rep_load 6
```

server.conf

```
[clustering]

max_peer_build_load = 2

max_peer_rep_load = 5
```

A. Data ingestion rate
B. Network latency and storage IOPS
C. Distance and location
D. SSL data encryption

Answer: B

Question: 5

Which statement is true about subsearches?

A. Subsearches are faster than other types of searches.
B. Subsearches work best for joining two large result sets.
C. Subsearches run at the same time as their outer search.
D. Subsearches work best for small result sets.

Answer: D

Reference: https://community.splunk.com/t5/Archive/Looking-for-way-to-explain-why-subsearches-are-so- slow/m-p/479133

Question: 6

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: 7

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.

Answer: D

Question: 8

What is the primary driver behind implementing indexer clustering in a customer's environment?

A. To improve resiliency as the search load increases.
B. To reduce indexing latency.
C. To scale out a Splunk environment to offer higher performance capability.
D. To provide higher availability for buckets of data.

Answer: D

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