Splunk Core Certified Consultant

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(s): C

---

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(s): A

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(s): D

---

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(s): B

---

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(s): A

---

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(s): D

---

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(s): D

---

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(s): D

---

9. In a single indexer cluster, where should the Monitoring Console (MC) be installed?

A. Deployer sharing with master cluster.

B. License master that has 50 clients or more.

C. Cluster master node

D. Production Search Head

Answer(s): C

10. 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(s): A

---

11. A customer's deployment server is overwhelmed with forwarder connections after adding an additional 1000 clients. The default phone home interval is set to 60 seconds. To reduce the number of connection failures to the DS what is recommended?

A. Create a tiered deployment server topology.

B. Reduce the phone home interval to 6 seconds.

C. Leave the phone home interval at 60 seconds.

D. Increase the phone home interval to 600 seconds.

Answer(s): A

---

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

A.

B.

C.

D.

Answer(s): C

---

13. What is the Splunk PS recommendation when using the deployment server and building deployment apps?

A. Carefully design smaller apps with specific configuration that can be reused.

B. Only deploy Splunk PS base configurations via the deployment server.

C. Use $SPLUNK_HOME/etc/system/local configurations on forwarders and only deploy TAs via the deployment server.

D. Carefully design bigger apps containing multiple configs.

Answer(s): B

---

14. Which of the following processor occur in the indexing pipeline?

A. tcp out, syslog out

B. Regex replacement, annotator

C. Aggregator

D. UTF-8, linebreaker, header

Answer(s): D

---

15. Which configuration item should be set to false to significantly improve data ingestion performance?

A. AUTO_KV_JSON

B. BREAK_ONLY_BEFORE_DATE

C. SHOULD_LINEMERGE

D. ANNOTATE_PUNCT

Answer(s): C

---

16. 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(s): D

---

17. 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(s): B

---

18. 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.

Answer(s): D

---

19. 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(s): C

---

20. 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.

Answer(s): D
