Splunk

Exam SPLK-2002

Splunk Enterprise Certified Architect

Version: 7.0

[Total Questions: 90]

Question No : 1

Which of the following should be done when installing Enterprise Security on a Search Head Cluster? (Select all that apply.)

A. Install Enterprise Security on the deployer.
B. Install Enterprise Security on a staging instance.
C. Copy the Enterprise Security configurations to the deployer.
D. Use the deployer to deploy Enterprise Security to the cluster members.

Answer: A,D

Question No : 2

How does IT Service Intelligence (ITSI) impact the planning of a Splunk deployment?

A. ITSI requires a dedicated deployment server.
B. The amount of users using ITSI will not impact performance.
C. ITSI in a Splunk deployment does not require additional hardware resources.
D. Depending on the Key Performance Indicators that are being tracked, additional infrastructure may be needed.

Answer: D

Question No : 3

Which of the following options can improve reliability of syslog delivery to Splunk? (Select all that apply.)

A. Use TCP syslog.
B. Configure UDP inputs on each Splunk indexer to receive data directly.
C. Use a network load balancer to direct syslog traffic to active backend syslog listeners.
D. Use one or more syslog servers to persist data with a Universal Forwarder to send the data to Splunk indexers.

Answer: C,D

Question No : 4

Which of the following statements describe search head clustering? (Select all that apply.)

A. A deployer is required.
B. At least three search heads are needed.
C. Search heads must meet the high-performance reference server requirements.
D. The deployer must have sufficient CPU and network resources to process service requests and push configurations.

Answer: A,C

Question No : 5

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%

Answer: B

Question No : 6

In a distributed environment, knowledge object bundles are replicated from the search head to which location on the search peer(s)?

A. SPLUNK_HOME/var/lib/searchpeers
B. SPLUNK_HOME/var/log/searchpeers
C. SPLUNK_HOME/var/run/searchpeers
D. SPLUNK_HOME/var/spool/searchpeers

Answer: C

Question No : 7

When configuring a Splunk indexer cluster, what are the default values for replication and search factor?

A. replication_factor = 2search_factor = 2
B. replication_factor = 2search factor = 3
C. replication_factor = 3search_factor = 2
D. replication_factor = 3search factor = 3

Answer: C

Question No : 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).

Answer: C

Question No : 9

Which of the following tasks should the architect perform when building a deployment plan? (Select all that apply.)

A. Use case checklist.
B. Install Splunk apps.
C. Inventory data sources.
D. Review network topology.

Answer: D

Question No : 10

Which Splunk server role regulates the functioning of indexer cluster?

A. Indexer
B. Deployer
C. Master Node
D. Monitoring Console

Answer: C

Question No : 11

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.

Answer: D

Question No : 12

When troubleshooting monitor inputs, which command checks the status of the tailed files?

A. splunk cmd btool inputs list | tail
B. splunk cmd btool check inputs layer
C. curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:FileStatus
D. curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:Tailstatus

Answer: C

Question No : 13

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.

Answer: D

Question No : 14

In which phase of the Splunk Enterprise data pipeline are indexed extraction configurations processed?

A. Input
B. Search
C. Parsing
D. Indexing

Answer: C

Question No : 15

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.

Answer: D

Explanation: https://community.splunk.com/t5/Deployment-Architecture/How-to-resolve-error-quot-Error-pulling-configurations-from-the/m-p/354231

Question No : 16

Which server.conf attribute should be added to the master node's server.conf file when decommissioning a site in an indexer cluster?

A. site_mappings
B. available_sites
C. site_search_factor
D. site_replication_factor

Answer: A

Question No : 17

When using the props.conf LINE_BREAKER attribute to delimit multi-line events, the SHOULD_LINEMERGE attribute should be set to what?

A. Auto
B. None
C. True
D. False

Answer: C

Question No : 18

Of the following types of files within an index bucket, which file type may consume the most disk?

A. Rawdata
B. Bloom filter
C. Metadata (.data)
D. Inverted index (.tsidx)

Answer: B

Question No : 19

When should multiple search pipelines be enabled?

A. Only if disk IOPS is at 800 or better.
B. Only if there are fewer than twelve concurrent users.
C. Only if running Splunk Enterprise version 6.6 or later.
D. Only if CPU and memory resources are significantly under-utilized.

Answer: D

Question No : 20

Stakeholders have identified high availability for searchable data as their top priority. Which of the following best addresses this requirement?

A. Increasing the search factor in the cluster.
B. Increasing the replication factor in the cluster.
C. Increasing the number of search heads in the cluster.
D. Increasing the number of CPUs on the indexers in the cluster.

Answer: A

Explanation:

https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture

Question No : 21

Which tool(s) can be leveraged to diagnose connection problems between an indexer and forwarder? (Select all that apply.)

A. telnet
B. tcpdump
C. splunk btool
D. splunk btprobe

Answer: B,C

Question No : 22

Which Splunk tool offers a health check for administrators to evaluate the health of their Splunk deployment?

A. btool
B. DiagGen
C. SPL Clinic
D. Monitoring Console

Answer: D

Question No : 23

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

Answer: D

Question No : 24

Which Splunk Enterprise offering has its own license?

A. Splunk Cloud Forwarder
B. Splunk Heavy Forwarder
C. Splunk Universal Forwarder
D. Splunk Forwarder Management

Answer: C

Question No : 25

Search dashboards in the Monitoring Console indicate that the distributed deployment is approaching its capacity. Which of the following options will provide the most search performance improvement?

A. Replace the indexer storage to solid state drives (SSD).
B. Add more search heads and redistribute users based on the search type.
C. Look for slow searches and reschedule them to run during an off-peak time.
D. Add more search peers and make sure forwarders distribute data evenly across all indexers.

Answer: D

Question No : 26

Which component in the splunkd.log will log information related to bad event breaking?

A. Audittrail
B. EventBreaking
C. IndexingPipeline
D. AggregatorMiningProcessor