Number: SPLK-2002
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: SPLK-2002

Exam Name: Splunk Enterprise Certified Architect

Exam A

QUESTION 1
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.

Correct Answer: A, B, D
Section:
Explanation:
Search head clustering is a Splunk feature that allows a group of search heads to share configurations, apps, and knowledge objects, and to provide high availability and scalability for searching. Search head clustering has the following characteristics:
A deployer is required. A deployer is a Splunk instance that distributes the configurations and apps to the members of the search head cluster. The deployer is not a member of the cluster, but a separate instance that communicates with the cluster master.
At least three search heads are needed. A search head cluster must have at least three search heads to form a quorum and to ensure high availability. If the cluster has less than three search heads, it cannot function properly and will enter a degraded mode.
The deployer must have sufficient CPU and network resources to process service requests and push configurations. The deployer is responsible for handling the requests from the cluster master and the cluster members, and for pushing the configurations and apps to the cluster members. Therefore, the deployer must have enough CPU and network resources to perform these tasks efficiently and reliably.
Search heads do not need to meet the high-performance reference server requirements, as this is not a mandatory condition for search head clustering. The high-performance reference server requirements are only recommended for optimal performance and scalability of Splunk deployments, but they are not enforced by Splunk.

QUESTION 2
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.

Correct Answer: A, C, D
Section:
Explanation:
When building a deployment plan, the architect should perform the following tasks:
Use case checklist. A use case checklist is a document that lists the use cases that the deployment will support, along with the data sources, the data volume, the data retention, the data model, the dashboards, the reports, the alerts, and the roles and permissions for each use case.A use case checklist helps to define the scope and the functionality of the deployment, and to identify the dependencies and the requirements for each use case1
Inventory data sources. An inventory of data sources is a document that lists the data sources that the deployment will ingest, along with the data type, the data format, the data location, the data collection method, the data volume, the data frequency, and the data owner for each data source.An inventory of data sources helps to determine the data ingestion strategy, the data parsing and enrichment, the data storage and retention, and the data security and compliance for the deployment1
Review network topology. A review of network topology is a process that examines the network infrastructure and the network connectivity of the deployment, along with the network bandwidth, the network latency, the network security, and the network monitoring for the deployment.A review of network topology helps to optimize the network performance and reliability, and to identify the network risks and mitigations for the deployment1
Installing Splunk apps is not a task that the architect should perform when building a deployment plan, as it is a task that the administrator should perform when implementing the deployment plan.Installing Splunk apps is a technical activity that requires access to the Splunk instances and the Splunk configurations, which are not available at the planning stage

QUESTION 3
Because Splunk indexing is read/write intensive, it is important to select the appropriate disk storage solution for each deployment. Which of the following statements is accurate about disk storage?

A. High performance SAN should never be used.
B. Enable NFS for storing hot and warm buckets.
C. The recommended RAID setup is RAID 10 (1 + 0).
D. Virtualized environments are usually preferred over bare metal for Splunk indexers.

Correct Answer: C
Section:
Explanation:
Splunk indexing is read/write intensive, as it involves reading data from various sources, writing data to disk, and reading data from disk for searching and reporting. Therefore, it is important to select the appropriate disk storage solution for each deployment, based on the performance, reliability, and cost requirements. The recommended RAID setup for Splunk indexers is RAID 10 (1 + 0), as it provides the best balance of performance and reliability. RAID 10 combines the advantages of RAID 1 (mirroring) and RAID 0 (striping), which means that it offers both data redundancy and data distribution.RAID 10 can tolerate multiple disk failures, as long as they are not in the same mirrored pair, and it can improve the read and write speed, as it can access multiple disks in parallel2
High performance SAN (Storage Area Network) can be used for Splunk indexers, but it is not recommended, as it is more expensive and complex than local disks. SAN also introduces additional network latency and dependency, which can affect the performance and availability of Splunk indexers.SAN is more suitable for Splunk search heads, as they are less read/write intensive and more CPU intensive2
NFS (Network File System) should not be used for storing hot and warm buckets, as it can cause data corruption, data loss, and performance degradation. NFS is a network-based file system that allows multiple clients to access the same files on a remote server. NFS is not compatible with Splunk index replication and search head clustering, as it can cause conflicts and inconsistencies among the Splunk instances. NFS is also slower and less reliable than local disks, as it depends on the network bandwidth and availability.NFS can be used for storing cold and frozen buckets, as they are less frequently accessed and less critical for Splunk operations2
Virtualized environments are not usually preferred over bare metal for Splunk indexers, as they can introduce additional overhead and complexity. Virtualized environments can affect the performance and reliability of Splunk indexers, as they share the physical resources and the network with other virtual machines. Virtualized environments can also complicate the monitoring and troubleshooting of Splunk indexers, as they add another layer of abstraction and configuration.Virtualized environments can be used for Splunk indexers, but they require careful planning and tuning to ensure optimal performance and availability2

QUESTION 4
Which of the following are possible causes of a crash in Splunk? (select all that apply)

A. Incorrect ulimit settings.
B. Insufficient disk IOPS.
C. Insufficient memory.
D. Running out of disk space.

Correct Answer: A, B, C, D
Section:
Explanation:
All of the options are possible causes of a crash in Splunk.According to the Splunk documentation1, incorrect ulimit settings can lead to file descriptor exhaustion, which can cause Splunk to crash or hang.Insufficient disk IOPS can also cause Splunk to crash or become unresponsive, as Splunk relies heavily on disk performance2.Insufficient memory can cause Splunk to run out of memory and crash, especially when running complex searches or handling large volumes of data3.Running out of disk space can cause Splunk to stop indexing data and crash, as Splunk needs enough disk space to store its data and logs4.
1: Configure ulimit settings for Splunk Enterprise2: Troubleshoot Splunk performance issues3: Troubleshoot memory usage4: Troubleshoot disk space issues

QUESTION 5
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.

Correct Answer: D
Section:
Explanation:
Adding more search peers and making sure forwarders distribute data evenly across all indexers will provide the most search performance improvement when the distributed deployment is approaching its capacity. Adding more search peers will increase the search concurrency and reduce the load on each indexer. Distributing data evenly across all indexers will ensure that the search workload is balanced and no indexer becomes a bottleneck. Replacing the indexer storage to SSD will improve the search performance, but it is a costly and time-consuming option. Adding more search heads will not improve the search performance if the indexers are the bottleneck. Rescheduling slow searches to run during an off-peak time will reduce the search contention, but it will not improve the search performance for each individual search. For more information, see [Scale your indexer cluster] and [Distribute data across your indexers] in the Splunk documentation.

QUESTION 6
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 source. Further investigation reveals that not all weblogs 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 of 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: C
Section:
Explanation:
The indexers may have different configurations than the heavy forwarders, which might cause the issue of inconsistently formatted events for a web sourcetype. The heavy forwarders perform parsing and indexing on the data before sending it to the indexers. If the indexers have different configurations than the heavy forwarders, such as different props.conf or transforms.conf settings, the data may be parsed or indexed differently on the indexers, resulting in inconsistent events. The search head configurations do not affect the event formatting, as the search head does not parse or index the data. The data inputs configurations on the forwarders do not affect the event formatting, as the data inputs only determine what data to collect and how to monitor it. The forwarder version does not affect the event formatting, as long as the forwarder is compatible with the indexer. For more information, see [Heavy forwarder versus indexer] and [Configure event processing] in the Splunk documentation.

QUESTION 7
A customer has installed a 500GB Enterprise license. They also purchased and installed a 300GB, no enforcement license on the same license master. How much data can the customer ingest before the search is locked out?

A. 300GB. After this limit, the search is locked out.
B. 500GB. After this limit, the search is locked out.
C. 800GB. After this limit, the search is locked out.
D. Search is not locked out. Violations are still recorded.

Correct Answer: D
Section:
Explanation:
Search is not locked out when a customer has installed a 500GB Enterprise license and a 300GB, no enforcement license on the same license master. The no enforcement license allows the customer to exceed the license quota without locking search, but violations are still recorded. The customer can ingest up to 800GB of data per day without violating the license, but if they ingest more than that, they will incur a violation. However, the violation will not lock search, as the no enforcement license overrides the enforcement policy of the Enterprise license. For more information, see [No enforcement licenses] and [License violations] in the Splunk documentation.

QUESTION 8
What does the deployer do in a Search Head Cluster (SHC)? (Select all that apply.)

A. Distributes apps to SHC members.

B. Bootstraps a clean Splunk install for a SHC.
C. Distributes non-search-related and manual configuration file changes.
D. Distributes runtime knowledge object changes made by users across the SHC.

Correct Answer: A, C
Section:
Explanation:
The deployer distributes apps and non-search related and manual configuration file changes to the search head cluster members. The deployer does not bootstrap a clean Splunk install for a search head cluster, as this is done by the captain. The deployer also does not distribute runtime knowledge object changes made by users across the search head cluster, as this is done by the replication factor. For more information, seeUse the deployer to distribute apps and configuration updatesin the Splunk documentation.

QUESTION 9
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

Correct Answer: D
Section:
Explanation:
When using the props.conf LINE_BREAKER attribute to delimit multi-line events, the SHOULD_LINEMERGE attribute should be set to false. This tells Splunk not to merge events that have been broken by the LINE_BREAKER. Setting the SHOULD_LINEMERGE attribute to true, auto, or none will cause Splunk to ignore the LINE_BREAKER and merge events based on other criteria. For more information, seeConfigure event line breakingin the Splunk documentation.

QUESTION 10
Which of the following should be included in a deployment plan?

A. Business continuity and disaster recovery plans.
B. Current logging details and data source inventory.
C. Current and future topology diagrams of the IT environment.
D. A comprehensive list of stakeholders, either direct or indirect.

Correct Answer: A, B, C
Section:
Explanation:
A deployment plan should include business continuity and disaster recovery plans, current logging details and data source inventory, and current and future topology diagrams of the IT environment. These elements are essential for planning, designing, and implementing a Splunk deployment that meets the business and technical requirements. A comprehensive list of stakeholders, either direct or indirect, is not part of the deployment plan, but rather part of the project charter. For more information, seeDeployment planningin the Splunk documentation.

QUESTION 11
A multi-site indexer cluster can be configured using which of the following? (Select all that apply.)

A. Via Splunk Web.
B. Directly edit SPLUNK_HOME/etc./system/local/server.conf
C. Run a Splunk edit cluster-config command from the CLI.

D. Directly edit SPLUNK_HOME/etc/system/default/server.conf

Correct Answer: B, C
Section:
Explanation:
A multi-site indexer cluster can be configured by directly editing SPLUNK_HOME/etc/system/local/server.conf or running a splunk edit cluster-config command from the CLI. These methods allow the administrator to specify the site attribute for each indexer node and the site_replication_factor and site_search_factor for the cluster. Configuring a multi-site indexer cluster via Splunk Web or directly editing SPLUNK_HOME/etc/system/default/server.conf are not supported methods. For more information, seeConfigure the indexer cluster with server.confin the Splunk documentation.

QUESTION 12
Which index-time props.conf attributes impact indexing performance? (Select all that apply.)

A. REPORT
B. LINE_BREAKER
C. ANNOTATE_PUNCT
D. SHOULD_LINEMERGE

Correct Answer: B, D
Section:
Explanation:
The index-time props.conf attributes that impact indexing performance are LINE_BREAKER and SHOULD_LINEMERGE. These attributes determine how Splunk breaks the incoming data into events and whether it merges multiple events into one. These operations can affect the indexing speed and the disk space consumption. The REPORT attribute does not impact indexing performance, as it is used to apply transforms at search time. The ANNOTATE_PUNCT attribute does not impact indexing performance, as it is used to add punctuation metadata to events at search time. For more information, see [About props.conf and transforms.conf] in the Splunk documentation.

QUESTION 13
Which of the following are client filters available in serverclass.conf? (Select all that apply.)

A. DNS name.
B. IP address.
C. Splunk server role.
D. Platform (machine type).

Correct Answer: A, B, D
Section:
Explanation:
The client filters available in serverclass.conf are DNS name, IP address, and platform (machine type). These filters allow the administrator to specify which forwarders belong to a server class and receive the apps and configurations from the deployment server. The Splunk server role is not a valid client filter in serverclass.conf, as it is not a property of the forwarder. For more information, see [Use forwarder management filters] in the Splunk documentation.

QUESTION 14
What log file would you search to verify if you suspect there is a problem interpreting a regular expression in a monitor stanza?

A. btool.log
B. metrics.log
C. splunkd.log
D. tailing_processor.log

Correct Answer: D
Section:
Explanation:
The tailing_processor.log file would be the best place to search if you suspect there is a problem interpreting a regular expression in a monitor stanza. This log file contains information about how Splunk monitors files and directories, including any errors or warnings related to parsing the monitor stanza. The splunkd.log file contains general information about the Splunk daemon, but it may not have the specific details about the monitor stanza. The btool.log file contains information about the configuration files, but it does not log the runtime behavior of the monitor stanza. The metrics.log file contains information about the performance metrics of Splunk, but it does not log the event breaking issues. For more information, seeAbout Splunk Enterprise loggingin the Splunk documentation.

QUESTION 15
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

Correct Answer: D
Section:
Explanation:
The Monitoring Console is the Splunk tool that offers a health check for administrators to evaluate the health of their Splunk deployment. The Monitoring Console provides dashboards and alerts that show the status and performance of various Splunk components, such as indexers, search heads, forwarders, license usage, and search activity. The Monitoring Console can also run health checks on the deployment and identify any issues or recommendations. The btool is a command-line tool that shows the effective settings of the configuration files, but it does not offer a health check. The DiagGen is a tool that generates diagnostic snapshots of the Splunk environment, but it does not offer a health check. The SPL Clinic is a tool that analyzes and optimizes SPL queries, but it does not offer a health check. For more information, seeAbout the Monitoring Consolein the Splunk documentation.

QUESTION 16
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: B
Section:
Explanation:
In a four site indexer cluster, the configuration that stores two searchable copies at the origin site, one searchable copy at site2, and a total of four searchable copies is site_search_factor = origin:2, site2:1, total:4. This configuration tells the cluster to maintain two copies of searchable data at the site where the data originates, one copy of searchable data at site2, and a total of four copies of searchable data across all sites. The site_search_factor determines how many copies of searchable data are maintained by the cluster for each site. The site_replication_factor determines how many copies of raw data are maintained by the cluster for each site. For more information, seeConfigure multisite indexer clusters with server.confin the Splunk documentation.

QUESTION 17
Which of the following is true regarding Splunk Enterprise's performance? (Select all that apply.)

A. Adding search peers increases the maximum size of search results.
B. Adding RAM to existing search heads provides additional search capacity.
C. Adding search peers increases the search throughput as the search load increases.
D. Adding search heads provides additional CPU cores to run more concurrent searches.

Correct Answer: C, D
Section:
Explanation:
The following statements are true regarding Splunk Enterprise performance:
Adding search peers increases the search throughput as search load increases. This is because adding more search peers distributes the search workload across more indexers, which reduces the load on each indexer and improves the search speed and concurrency.
Adding search heads provides additional CPU cores to run more concurrent searches. This is because adding more search heads increases the number of search processes that can run in parallel, which improves the search performance and scalability. The following statements are false regarding Splunk Enterprise performance:
Adding search peers does not increase the maximum size of search results. The maximum size of search results is determined by the maxresultrows setting in the limits.conf file, which is independent of the number of search peers.
Adding RAM to an existing search head does not provide additional search capacity. The search capacity of a search head is determined by the number of CPU cores, not the amount of RAM. Adding RAM to a search head may improve the search performance, but not the search capacity. For more information, seeSplunk Enterprise performancein the Splunk documentation.

QUESTION 18
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
Section:
Explanation:
The Splunk Universal Forwarder is the only Splunk Enterprise offering that has its own license. The Splunk Universal Forwarder license allows the forwarder to send data to any Splunk Enterprise or Splunk Cloud instance without consuming any license quota. The Splunk Heavy Forwarder does not have its own license, but rather consumes the license quota of the Splunk Enterprise or Splunk Cloud instance that it sends data to. The Splunk Cloud Forwarder and the Splunk Forwarder Management are not separate Splunk Enterprise offerings, but rather features of the Splunk Cloud service. For more information, see [About forwarder licensing] in the Splunk documentation.

QUESTION 19
Which component in the splunkd.log will log information related to bad event breaking?

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

Correct Answer: D
Section:
Explanation:
The AggregatorMiningProcessor component in the splunkd.log file will log information related to bad event breaking. The AggregatorMiningProcessor is responsible for breaking the incoming data into events and applying the props.conf settings. If there is a problem with the event breaking, such as incorrect timestamps, missing events, or merged events, the AggregatorMiningProcessor will log the error or warning messages in the splunkd.log file. The Audittrail component logs information about the audit events, such as user actions, configuration changes, and search activity. The EventBreaking component logs information about the event breaking rules, such as the LINE_BREAKER and SHOULD_LINEMERGE settings. The IndexingPipeline component logs information about the indexing pipeline, such as the parsing, routing, and indexing phases. For more information, seeAbout Splunk Enterprise loggingand [Configure event line breaking] in the Splunk documentation.

QUESTION 20
Which Splunk server role regulates the functioning of indexer cluster?

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

Correct Answer: C
Section:
Explanation:
The master node is the Splunk server role that regulates the functioning of the indexer cluster. The master node coordinates the activities of the peer nodes, such as data replication, data searchability, and data recovery. The master node also manages the cluster configuration bundle and distributes it to the peer nodes. The indexer is the Splunk server role that indexes the incoming data and makes it searchable. The deployer is the Splunk server role that distributes apps and configuration updates to the search head cluster members. The monitoring console is the Splunk server role that monitors the health and performance of the Splunk deployment. For more information, seeAbout indexer clusters and index replicationin the Splunk documentation.

QUESTION 21
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: D
Section:
Explanation:
When adding or rejoining a member to a search head cluster, and 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.
The corrective action that should be taken is to run the splunk resync shcluster-replicated-config command on this member. This command will delete the existing configuration files on this member and replace them with the latest configuration files from the captain. This will ensure that the member has the same configuration as the rest of the cluster. Restarting the search head, running the splunk apply shcluster-bundle command from the deployer, or running the clean raft command on all members of the search head cluster are not the correct actions to take in this scenario. For more information, seeResolve configuration inconsistencies across cluster membersin the Splunk documentation.

QUESTION 22
Which of the following commands is used to clear the KV store?

A. splunk clean kvstore
B. splunk clear kvstore
C. splunk delete kvstore
D. splunk reinitialize kvstore

Correct Answer: A
Section:
Explanation:
The splunk clean kvstore command is used to clear the KV store. This command will delete all the collections and documents in the KV store and reset it to an empty state. This command can be useful for troubleshooting KV store issues or resetting the KV store data. The splunk clear kvstore, splunk delete kvstore, and splunk reinitialize kvstore commands are not valid Splunk commands. For more information, seeUse the CLI to manage the KV storein the Splunk documentation.

QUESTION 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

Correct Answer: B
Section:
Explanation:
Increasing the number of parallel ingestion pipelines in server.conf is most likely to improve indexing performance when indexing is slow and real-time search results are delayed in a Splunk environment with two indexers and one search head. The parallel ingestion pipelines allow Splunk to process multiple data streams simultaneously, which increases the indexing throughput and reduces the indexing latency. Increasing the maximum number of hot buckets in indexes.conf will not improve indexing performance, but rather increase the disk space consumption and the bucket rolling time. Decreasing the maximum size of the search pipelines in limits.conf will not improve indexing performance, but rather reduce the search performance and the search concurrency. Decreasing the maximum concurrent scheduled searches in limits.conf will not improve indexing performance, but rather reduce the search capacity and the search availability. For more information, seeConfigure parallel ingestion pipelinesin the Splunk documentation.

QUESTION 24
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
Section:
Explanation:
The guidance Splunk gives for estimating size on for syslog data is 50% of original data size. This divides between files in the index as follows: rawdata is 15%, tsidx is 35%. The rawdata is the compressed version of the original data, which typically takes about 15% of the original data size. The tsidx is the index file that contains the time-series metadata and the inverted index, which typically takes about 35% of the original data size. The total size of the rawdata and the tsidx is about 50% of the original data size. For more information, see [Estimate your storage requirements] in the Splunk documentation.

QUESTION 25
In an existing Splunk environment, the new index buckets that are created each day are about half the size of the incoming data. Within each bucket, about 30% of the space is used for rawdata and about 70% for index files. What additional information is needed to calculate the daily disk consumption, per indexer, if indexer clustering is implemented?

A. Total daily indexing volume, number of peer nodes, and number of accelerated searches.
B. Total daily indexing volume, number of peer nodes, replication factor, and search factor.
C. Total daily indexing volume, replication factor, search factor, and number of search heads.
D. Replication factor, search factor, number of accelerated searches, and total disk size across cluster.

Correct Answer: B
Section:
Explanation:
The additional information that is needed to calculate the daily disk consumption, per indexer, if indexer clustering is implemented, is the total daily indexing volume, the number of peer nodes, the replication factor, and the search factor. These information are required to estimate how much data is ingested, how many copies of raw data and searchable data are maintained, and how many indexers are involved in the cluster. The number of accelerated searches, the number of search heads, and the total disk size across the cluster are not relevant for calculating the daily disk consumption, per indexer. For more information, see [Estimate your storage

requirements] in the Splunk documentation.

QUESTION 26
A three-node search head cluster is skipping a large number of searches across time. What should be done to increase scheduled search capacity on the search head cluster?

A. Create a job server on the cluster.
B. Add another search head to the cluster.
C. server.conf captain_is_adhoc_searchhead = true.
D. Change limits.conf value for max_searches_per_cpu to a higher value.

Correct Answer: D
Section:
Explanation:
Changing the limits.conf value for max_searches_per_cpu to a higher value is the best option to increase scheduled search capacity on the search head cluster when a large number of searches are skipped across time. This value determines how many concurrent scheduled searches can run on each CPU core of the search head. Increasing this value will allow more scheduled searches to run at the same time, which will reduce the number of skipped searches. Creating a job server on the cluster, running the server.conf captain_is_adhoc_searchhead = true command, or adding another search head to the cluster are not the best options to increase scheduled search capacity on the search head cluster. For more information, see [Configure limits.conf] in the Splunk documentation.

QUESTION 27
The frequency in which a deployment client contacts the deployment server is controlled by what?

A. polling_interval attribute in outputs.conf
B. phoneHomeIntervalInSecs attribute in outputs.conf
C. polling_interval attribute in deploymentclient.conf
D. phoneHomeIntervalInSecs attribute in deploymentclient.conf

Correct Answer: D
Section:
Explanation:
The frequency in which a deployment client contacts the deployment server is controlled by the phoneHomeIntervalInSecs attribute in deploymentclient.conf. This attribute specifies how often the deployment client checks in with the deployment server to get updates on the apps and configurations that it should receive. The polling_interval attribute in outputs.conf controls how often the forwarder sends data to the indexer or another forwarder. The polling_interval attribute in deploymentclient.conf and the phoneHomeIntervalInSecs attribute in outputs.conf are not valid Splunk attributes. For more information, seeConfigure deployment clientsandConfigure forwarders with outputs.confin the Splunk documentation.

QUESTION 28
To activate replication for an index in an indexer cluster, what attribute must be configured in indexes.conf on all peer nodes?

A. repFactor = 0
B. replicate = 0
C. repFactor = auto
D. replicate = auto

Correct Answer: C
Section:
Explanation:
To activate replication for an index in an indexer cluster, the repFactor attribute must be configured in indexes.conf on all peer nodes. This attribute specifies the replication factor for the index, which determines how many copies of raw data are maintained by the cluster. Setting the repFactor attribute to auto will enable replication for the index. The replicate attribute in indexes.conf is not a valid Splunk attribute. The repFactor attribute in outputs.conf and the replicate attribute in deploymentclient.conf are not related to replication for an index in an indexer cluster. For more information, seeConfigure indexes for indexer clustersin the Splunk documentation.

QUESTION 29
Which of the following clarification steps should be taken if apps are not appearing on a deployment client? (Select all that apply.)

A. Check serverclass.conf of the deployment server.
B. Check deploymentclient.conf of the deployment client.
C. Check the content of SPLUNK_HOME/etc/apps of the deployment server.
D. Search for relevant events in splunkd.log of the deployment server.

Correct Answer: A, B, D
Section:
Explanation:
The following clarification steps should be taken if apps are not appearing on a deployment client:
Check serverclass.conf of the deployment server. This file defines the server classes and the apps and configurations that they should receive from the deployment server. Make sure that the deployment client belongs to the correct server class and that the server class has the desired apps and configurations.
Check deploymentclient.conf of the deployment client. This file specifies the deployment server that the deployment client contacts and the client name that it uses. Make sure that the deployment client is pointing to the correct deployment server and that the client name matches the server class criteria.
Search for relevant events in splunkd.log of the deployment server. This file contains information about the deployment server activities, such as sending apps and configurations to the deployment clients, detecting client check-ins, and logging any errors or warnings. Look for any events that indicate a problem with the deployment server or the deployment client.
Checking the content of SPLUNK_HOME/etc/apps of the deployment server is not a necessary clarification step, as this directory does not contain the apps and configurations that are distributed to the deployment clients. The apps and configurations for the deployment server are stored in SPLUNK_HOME/etc/deployment-apps. For more information, seeConfigure deployment server and clientsin the Splunk documentation.

QUESTION 30
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
Section:
Explanation:
The minimum reference server specification for a Splunk indexer is 12 CPU cores, 12GB RAM, and 800 IOPS. This specification is based on the assumption that the indexer will handle an average indexing volume of 100GB per day, with a peak of 300GB per day, and a typical search load of 1 concurrent search per 1GB of indexing volume. The other specifications are either higher or lower than the minimum requirement. For more information, see [Reference hardware] in the Splunk documentation.

QUESTION 31
Which of the following security options must be explicitly configured (i.e. which options are not enabled by default)?

A. Data encryption between Splunk Web and splunkd.
B. Certificate authentication between forwarders and indexers.
C. Certificate authentication between Splunk Web and search head.
D. Data encryption for distributed search between search heads and indexers.

Correct Answer: B
Section:
Explanation:
The following security option must be explicitly configured, as it is not enabled by default:

Certificate authentication between forwarders and indexers. This option allows the forwarders and indexers to verify each other's identity using SSL certificates, which prevents unauthorized data transmission or spoofing attacks. This option is not enabled by default, as it requires the administrator to generate and distribute the certificates for the forwarders and indexers. For more information, see [Secure the communication between forwarders and indexers] in the Splunk documentation. The following security options are enabled by default:
Data encryption between Splunk Web and splunkd. This option encrypts the communication between the Splunk Web interface and the splunkd daemon using SSL, which prevents data interception or tampering. This option is enabled by default, as Splunk provides a self-signed certificate for this purpose. For more information, see [About securing Splunk Enterprise with SSL] in the Splunk documentation.
Certificate authentication between Splunk Web and search head. This option allows the Splunk Web interface and the search head to verify each other's identity using SSL certificates, which prevents unauthorized access or spoofing attacks. This option is enabled by default, as Splunk provides a self-signed certificate for this purpose. For more information, see [About securing Splunk Enterprise with SSL] in the Splunk documentation.
Data encryption for distributed search between search heads and indexers. This option encrypts the communication between the search heads and the indexers using SSL, which prevents data interception or tampering. This option is enabled by default, as Splunk provides a self-signed certificate for this purpose. For more information, see [Secure your distributed search environment] in the Splunk documentation.

QUESTION 32
Which of the following artifacts are included in a Splunk diag file? (Select all that apply.)

A. OS settings.
B. Internal logs.
C. Customer data.
D. Configuration files.

Correct Answer: B, D
Section:
Explanation:
The following artifacts are included in a Splunk diag file:
Internal logs. These are the log files that Splunk generates to record its own activities, such as splunkd.log, metrics.log, audit.log, and others. These logs can help troubleshoot Splunk issues and monitor Splunk performance.
Configuration files. These are the files that Splunk uses to configure various aspects of its operation, such as server.conf, indexes.conf, props.conf, transforms.conf, and others. These files can help understand Splunk settings and behavior. The following artifacts are not included in a Splunk diag file:
OS settings. These are the settings of the operating system that Splunk runs on, such as the kernel version, the memory size, the disk space, and others. These settings are not part of the Splunk diag file, but they can be collected separately using the diag --os option.
Customer data. These are the data that Splunk indexes and makes searchable, such as the rawdata and the tsidx files. These data are not part of the Splunk diag file, as they may contain sensitive or confidential information.
For more information, seeGenerate a diagnostic snapshot of your Splunk Enterprise deploymentin the Splunk documentation.

QUESTION 33
metrics. log is stored in which index?

A. main
B. _telemetry
C. _internal
D. _introspection

Correct Answer: C
Section:
Explanation:
According to the Splunk documentation1, metrics.log is a file that contains various metrics data for reviewing product behavior, such as pipeline, queue, thruput, and tcpout_connections.Metrics.log is stored in the _internal index by default2, which is a special index that contains internal logs and metrics for Splunk Enterprise. The other options are false because:
main is the default index for user data, not internal data3.
_telemetry is an index that contains data collected by the Splunk Telemetry feature, which sends anonymous usage and performance data to Splunk4.
_introspection is an index that contains data collected by the Splunk Monitoring Console, which monitors the health and performance of Splunk components.

QUESTION 34
A single-site indexer cluster has a replication factor of 3, and a search factor of 2. What is true about this cluster?

A. The cluster will ensure there are at least two copies of each bucket, and at least three copies of searchable metadata.
B. The cluster will ensure there are at most three copies of each bucket, and at most two copies of searchable metadata.
C. The cluster will ensure only two search heads are allowed to access the bucket at the same time.
D. The cluster will ensure there are at least three copies of each bucket, and at least two copies of searchable metadata.

Correct Answer: D
Section:
Explanation:
A single-site indexer cluster is a group of Splunk Enterprise instances that index and replicate data across the cluster1.A bucket is a directory that contains indexed data, along with metadata and other information2.A replication factor is the number of copies of each bucket that the cluster maintains1.A search factor is the number of searchable copies of each bucket that the cluster maintains1.A searchable copy is a copy that contains both the raw data and the index files3.A search head is a Splunk Enterprise instance that coordinates the search activities across the peer nodes1.
Option D is the correct answer because it reflects the definitions of replication factor and search factor. The cluster will ensure that there are at least three copies of each bucket, one on each peer node, to satisfy the replication factor of 3. The cluster will also ensure that there are at least two searchable copies of each bucket, one primary and one searchable, to satisfy the search factor of 2.The primary copy is the one that the search head uses to run searches, and the searchable copy is the one that can be promoted to primary if the original primary copy becomes unavailable3.
Option A is incorrect because it confuses the replication factor and the search factor. The cluster will ensure there are at least three copies of each bucket, not two, to meet the replication factor of 3. The cluster will ensure there are at least two copies of searchable metadata, not three, to meet the search factor of 2.
Option B is incorrect because it uses the wrong terms. The cluster will ensure there are at least, not at most, three copies of each bucket, to meet the replication factor of 3. The cluster will ensure there are at least, not at most, two copies of searchable metadata, to meet the search factor of 2.
Option C is incorrect because it has nothing to do with the replication factor or the search factor. The cluster does not limit the number of search heads that can access the bucket at the same time.The search head can search across multiple clusters, and the cluster can serve multiple search heads1.
1:The basics of indexer cluster architecture - Splunk Documentation2:About buckets - Splunk Documentation3:Search factor - Splunk Documentation

QUESTION 35
Which of the following configuration attributes must be set in server, conf on the cluster manager in a single-site indexer cluster?

A. master_uri
B. site
C. replication_factor
D. site_replication_factor

Correct Answer: A
Section:
Explanation:
The correct configuration attribute to set in server.conf on the cluster manager in a single-site indexer cluster ismaster_uri.This attribute specifies the URI of the cluster manager, which is required for the peer nodes and search heads to communicate with it1. The other attributes are not required for a single-site indexer cluster, but they are used for a multisite indexer cluster.The site attribute defines the site name for each node in a multisite indexer cluster2.The replication_factor attribute defines the number of copies of each bucket to maintain across the entire multisite indexer cluster3.The site_replication_factor attribute defines the number of copies of each bucket to maintain across each site in a multisite indexer cluster4. Therefore, option A is the correct answer, and options B, C, and D are incorrect.
1: Configure the cluster manager2: Configure the site attribute3: Configure the replication factor4: Configure the site replication factor

QUESTION 36
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
Section:
Explanation:
The splunk offline --enforce-counts command will permanently decommission a peer node operating in an indexer cluster. This command will remove the peer node from the cluster and delete its data. This command should be used when the peer node is no longer needed or is being replaced by another node. The splunk stop -f command will stop the Splunk service on the peer node, but it will not decommission it from the cluster. The splunk offline -f command will take the peer node offline, but it will not delete its data or enforce the replication and search factors. The splunk decommission --enforce-counts command is not a valid Splunk command. For more information, seeRemove a peer node from an indexer clusterin the Splunk documentation.

QUESTION 37
Which CLI command converts a Splunk instance to a license slave?

A. splunk add licenses
B. splunk list licenser-slaves
C. splunk edit licenser-localslave
D. splunk list licenser-localslave

Correct Answer: C
Section:
Explanation:
The splunk edit licenser-localslave command is used to convert a Splunk instance to a license slave. This command will configure the Splunk instance to contact a license master and receive a license from it. This command should be used when the Splunk instance is part of a distributed deployment and needs to share a license pool with other instances. The splunk add licenses command is used to add a license to a Splunk instance, not to convert it to a license slave. The splunk list licenser-slaves command is used to list the license slaves that are connected to a license master, not to convert a Splunk instance to a license slave. The splunk list licenser-localslave command is used to list the license master that a license slave is connected to, not to convert a Splunk instance to a license slave. For more information, seeConfigure license slavesin the Splunk documentation.

QUESTION 38
Splunk Enterprise platform instrumentation refers to data that the Splunk Enterprise deployment logs in the _introspection index. Which of the following logs are included in this index? (Select all that apply.)

A. audit.log
B. metrics.log
C. disk_objects.log
D. resource_usage.log

Correct Answer: C, D
Section:
Explanation:
The following logs are included in the _introspection index, which contains data that the Splunk Enterprise deployment logs for platform instrumentation:
disk_objects.log. This log contains information about the disk objects that Splunk creates and manages, such as buckets, indexes, and files. This log can help monitor the disk space usage and the bucket lifecycle.
resource_usage.log. This log contains information about the resource usage of Splunk processes, such as CPU, memory, disk, and network. This log can help monitor the Splunk performance and identify any resource bottlenecks. The following logs are not included in the _introspection index, but rather in the _internal index, which contains data that Splunk generates for internal logging:
audit.log. This log contains information about the audit events that Splunk records, such as user actions, configuration changes, and search activity. This log can help audit the Splunk operations and security.
metrics.log. This log contains information about the performance metrics that Splunk collects, such as data throughput, data latency, search concurrency, and search duration. This log can help measure the Splunk performance and efficiency. For more information, seeAbout Splunk Enterprise loggingand [About the _introspection index] in the Splunk documentation.

QUESTION 39
Which of the following can a Splunk diag contain?

A. Search history, Splunk users and their roles, running processes, indexed data
B. Server specs, current open connections, internal Splunk log files, index listings

C. KV store listings, internal Splunk log files, search peer bundles listings, indexed data
D. Splunk platform configuration details, Splunk users and their roles, current open connections, index listings

Correct Answer: B
Section:
Explanation:
The following artifacts are included in a Splunk diag file:
Server specs. These are the specifications of the server that Splunk runs on, such as the CPU model, the memory size, the disk space, and the network interface. These specs can help understand the Splunk hardware requirements and performance.
Current open connections. These are the connections that Splunk has established with other Splunk instances or external sources, such as forwarders, indexers, search heads, license masters, deployment servers, and data inputs. These connections can help understand the Splunk network topology and communication.
Internal Splunk log files. These are the log files that Splunk generates to record its own activities, such as splunkd.log, metrics.log, audit.log, and others. These logs can help troubleshoot Splunk issues and monitor Splunk performance.
Index listings. These are the listings of the indexes that Splunk has created and configured, such as the index name, the index location, the index size, and the index attributes. These listings can help understand the Splunk data management and retention. The following artifacts are not included in a Splunk diag file:
Search history. This is the history of the searches that Splunk has executed, such as the search query, the search time, the search results, and the search user. This history is not part of the Splunk diag file, but it can be accessed from the Splunk Web interface or the audit.log file.
Splunk users and their roles. These are the users that Splunk has created and assigned roles to, such as the user name, the user password, the user role, and the user capabilities. These users and roles are not part of the Splunk diag file, but they can be accessed from the Splunk Web interface or the authentication.conf and authorize.conf files.
KV store listings. These are the listings of the KV store collections and documents that Splunk has created and stored, such as the collection name, the collection schema, the document ID, and the document fields. These listings are not part of the Splunk diag file, but they can be accessed from the Splunk Web interface or the mongod.log file.
Indexed data. These are the data that Splunk indexes and makes searchable, such as the rawdata and the tsidx files. These data are not part of the Splunk diag file, as they may contain sensitive or confidential information. For more information, seeGenerate a diagnostic snapshot of your Splunk Enterprise deploymentin the Splunk documentation.

QUESTION 40
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: A, D
Section:
Explanation:
The following statements are true about Splunk indexer clustering:
All peer nodes must run exactly the same Splunk version. This is a requirement for indexer clustering, as different Splunk versions may have different data formats or features that are incompatible with each other. All peer nodes must run the same Splunk version as the master node and the search heads that connect to the cluster.
The search head must run the same or a later Splunk version than the peer nodes. This is a recommendation for indexer clustering, as a newer Splunk version may have new features or bug fixes that improve the search functionality or performance. The search head should not run an older Splunk version than the peer nodes, as this may cause search errors or failures. The following statements are false about Splunk indexer clustering:
The master node must run the same or a later Splunk version than the search heads. This is not a requirement or a recommendation for indexer clustering, as the master node does not participate in the search process. The master node should run the same Splunk version as the peer nodes, as this ensures the cluster compatibility and functionality.
The peer nodes must run the same or a later Splunk version than the master node. This is not a requirement or a recommendation for indexer clustering, as the peer nodes do not coordinate the cluster activities. The peer nodes should run the same Splunk version as the master node, as this ensures the cluster compatibility and functionality. For more information, see [About indexer clusters and index replication] and [Upgrade an indexer cluster] in the Splunk documentation.

QUESTION 41
A customer plans to ingest 600 GB of data per day into Splunk. They will have six concurrent users, and they also want high data availability and high search performance. The customer is concerned about cost and wants to spend the minimum amount on the hardware for Splunk. How many indexers are recommended for this deployment?

A. Two indexers not in a cluster, assuming users run many long searches.
B. Three indexers not in a cluster, assuming a long data retention period.
C. Two indexers clustered, assuming high availability is the greatest priority.
D. Two indexers clustered, assuming a high volume of saved/scheduled searches.

Correct Answer: C
Section:
Explanation:
Two indexers clustered is the recommended deployment for a customer who plans to ingest 600 GB of data per day into Splunk, has six concurrent users, and wants high data availability and high search performance. This deployment will provide enough indexing capacity and search concurrency for the customer's needs, while also ensuring data replication and searchability across the cluster. The customer can also save on the hardware cost by using only two indexers. Two indexers not in a cluster will not provide high data availability, as there is no data replication or failover. Three indexers not in a cluster will provide more indexing capacity and search concurrency, but also more hardware cost and no data availability. The customer's data retention period, number of long searches, or volume of saved/scheduled searches are not relevant for determining the number of indexers. For more information, see [Reference hardware] and [About indexer clusters and index replication] in the Splunk documentation.

QUESTION 42
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.

Correct Answer: A, D
Section:
Explanation:
Syslog is a standard protocol for sending log messages from various devices and applications to a central server. Syslog can use either UDP or TCP as the transport layer protocol. UDP is faster but less reliable, as it does not guarantee delivery or order of the messages. TCP is slower but more reliable, as it ensures delivery and order of the messages. Therefore, to improve the reliability of syslog delivery to Splunk, it is recommended to use TCP syslog.
Another option to improve the reliability of syslog delivery to Splunk is to use one or more syslog servers to persist data with a Universal Forwarder to send the data to Splunk indexers. This way, the syslog servers can act as a buffer and store the data in case of network or Splunk outages. The Universal Forwarder can then forward the data to Splunk indexers when they are available.
Using a network load balancer to direct syslog traffic to active backend syslog listeners is not a reliable option, as it does not address the possibility of data loss or duplication due to network failures or Splunk outages.
Configuring UDP inputs on each Splunk indexer to receive data directly is also not a reliable option, as it exposes the indexers to the network and increases the risk of data loss or duplication due to UDP limitations.

QUESTION 43
What is the logical first step when starting a deployment plan?

A. Inventory the currently deployed logging infrastructure.
B. Determine what apps and use cases will be implemented.
C. Gather statistics on the expected adoption of Splunk for sizing.
D. Collect the initial requirements for the deployment from all stakeholders.

Correct Answer: D
Section:
Explanation:
The logical first step when starting a deployment plan is to collect the initial requirements for the deployment from all stakeholders. This includes identifying the business objectives, the data sources, the use cases, the security and compliance needs, the scalability and availability expectations, and the budget and timeline constraints. Collecting the initial requirements helps to define the scope and the goals of the deployment, and to align the expectations of all the parties involved.

Inventorying the currently deployed logging infrastructure, determining what apps and use cases will be implemented, and gathering statistics on the expected adoption of Splunk for sizing are all important steps in the deployment planning process, but they are not the logical first step. These steps can be done after collecting the initial requirements, as they depend on the information gathered from the stakeholders.

QUESTION 44
Which of the following strongly impacts storage sizing requirements for Enterprise Security?

A. The number of scheduled (correlation) searches.
B. The number of Splunk users configured.
C. The number of source types used in the environment.
D. The number of Data Models accelerated.

Correct Answer: D
Section:
Explanation:
Data Model acceleration is a feature that enables faster searches over large data sets by summarizing the raw data into a more efficient format. Data Model acceleration consumes additional disk space, as it stores both the raw data and the summarized data. The amount of disk space required depends on the size and complexity of the Data Model, the retention period of the summarized data, and the compression ratio of the data. According to the Splunk Enterprise Security Planning and Installation Manual, Data Model acceleration is one of the factors that strongly impacts storage sizing requirements for Enterprise Security. The other factors are the volume and type of data sources, the retention policy of the data, and the replication factor and search factor of the index cluster.The number of scheduled (correlation) searches, the number of Splunk users configured, and the number of source types used in the environment are not directly related to storage sizing requirements for Enterprise Security1
1: https://docs.splunk.com/Documentation/ES/6.6.0/Install/Plan#Storage_sizing_requirements

QUESTION 45
Which of the following is true regarding the migration of an index cluster from single-site to multi-site?

A. Multi-site policies will apply to all data in the indexer cluster.
B. All peer nodes must be running the same version of Splunk.
C. Existing single-site attributes must be removed.
D. Single-site buckets cannot be converted to multi-site buckets.

Correct Answer: C
Section:
Explanation:
According to the Splunk documentation1, when migrating an indexer cluster from single-site to multi-site, you must remove the existing single-site attributes from the server.conf file of each peer node. These attributes include replication_factor, search_factor, and cluster_label. You must also restart each peer node after removing the attributes. The other options are false because:
Multi-site policies will apply only to the data created after migration, unless you configure the manager node to convert legacy buckets to multi-site1.
All peer nodes do not need to run the same version of Splunk, as long as they are compatible with the manager node2.
Single-site buckets can be converted to multi-site buckets by changing the constrain_singlesite_buckets setting in the manager node's server.conf file to 'false'1.

QUESTION 46
What information is written to the __introspection log file?

A. File monitor input configurations.
B. File monitor checkpoint offset.
C. User activities and knowledge objects.
D. KV store performance.

Correct Answer: D
Section:

Explanation:
The __introspection log file contains data about the impact of the Splunk software on the host system, such as CPU, memory, disk, and network usage, as well as KV store performance1.This log file is monitored by default and the contents are sent to the _introspection index1. The other options are not related to the __introspection log file.File monitor input configurations are stored in inputs.conf2.File monitor checkpoint offset is stored in fishbucket3.User activities and knowledge objects are stored in the _audit and _internal indexes respectively4.

QUESTION 47
A customer has a four site indexer cluster. The customer has requirements to store five copies of searchable data, with one searchable copy of data at the origin site, and one searchable copy at the disaster recovery site (site4).
Which configuration meets these requirements?

A. site_replication_factor = origin:2, site4:l, total:3
B. site_replication_factor = origin:l, site4:l, total:5
C. site_search_factor = origin:2, site4:l, total:3
D. site search factor = origin:1, site4:l, total:5

Correct Answer: B
Section:
Explanation:
The correct configuration to meet the customer's requirements issite_replication_factor = origin:1, site4:1, total:5. This means that each bucket will have one copy at the origin site, one copy at the disaster recovery site (site4), and three copies at any other sites. The total number of copies will be five, as required by the customer.The site_replication_factor determines how many copies of each bucket are stored across the sites in a multisite indexer cluster1.The site_search_factor determines how many copies of each bucket are searchable across the sites in a multisite indexer cluster2. Therefore, option B is the correct answer, and options A, C, and D are incorrect.
1: Configure the site replication factor2: Configure the site search factor

QUESTION 48
Which of the following server. conf stanzas indicates the Indexer Discovery feature has not been fully configured (restart pending) on the Master Node?
A)

```
[indexer_discovery]
 pass4SymmKey = $7$XcXl1lu4630Jbui14oVe295+mvx6gCKKv6kf2zEaVB6Ie4DcZ318nLVlfW
```

B)

```
[clustering]
 mode = forwarder
 pass4SymmKey = $7$PU9SBXww63Vz3UJdDYGIN0UrdscRh83ssC2pEpwE6P3gn50iNFO94g==
```

C)

```
[clustering]
 mode = master
 pass4SymmKey = $7$tYTXzke+1r+3DULTHHDUTmYOXdtZJPxm21XwMARrJE20jsmicp9C3ni0
```

D)

```
[indexer_discovery]
 pass4SymmKey = idxdiscovery
```

A. Option A
B. Option B
C. Option C
D. Option D

Correct Answer: A
Section:
Explanation:
The Indexer Discovery feature enables forwarders to dynamically connect to the available peer nodes in an indexer cluster. To use this feature, the manager node must be configured with the [indexer_discovery] stanza and a pass4SymmKey value. The forwarders must also be configured with the same pass4SymmKey value and the master_uri of the manager node. The pass4SymmKey value must be encrypted using the splunk _encrypt command. Therefore, option A indicates that the Indexer Discovery feature has not been fully configured on the manager node, because the pass4SymmKey value is not encrypted. The other options are not related to the Indexer Discovery feature. Option B shows the configuration of a forwarder that is part of an indexer cluster. Option C shows the configuration of a manager node that is part of an indexer cluster.Option D shows an invalid configuration of the [indexer_discovery] stanza, because the pass4SymmKey value is not encrypted and does not match the forwarders' pass4SymmKey value12
1: https://docs.splunk.com/Documentation/Splunk/9.1.2/Indexer/indexerdiscovery2:
https://docs.splunk.com/Documentation/Splunk/9.1.2/Security/Secureyourconfigurationfiles#Encrypt_the_pass4SymmKey_setting_in_server.conf

QUESTION 49
A customer currently has many deployment clients being managed by a single, dedicated deployment server. The customer plans to double the number of clients.
What could be done to minimize performance issues?

A. Modify deploymentclient. conf to change from a Pull to Push mechanism.
B. Reduce the number of apps in the Manager Node repository.
C. Increase the current deployment client phone home interval.
D. Decrease the current deployment client phone home interval.

Correct Answer: C
Section:
Explanation:
According to the Splunk documentation1, increasing the current deployment client phone home interval can minimize performance issues by reducing the frequency of communication between the clients and the deployment server. This can also reduce the network traffic and the load on the deployment server. The other options are false because:
Modifying deploymentclient.conf to change from a Pull to Push mechanism is not possible, as Splunk does not support a Push mechanism for deployment server2.
Reducing the number of apps in the Manager Node repository will not affect the performance of the deployment server, as the apps are only downloaded when there is a change in the configuration or a new app is added3.
Decreasing the current deployment client phone home interval will increase the performance issues, as it will increase the frequency of communication between the clients and the deployment server, resulting in more network traffic and load on the deployment server1.

QUESTION 50
Following Splunk recommendations, where could the Monitoring Console (MC) be installed in a distributed deployment with an indexer cluster, a search head cluster, and 1000 forwarders?

A. On a search peer in the cluster.
B. On the deployment server.
C. On the search head cluster deployer.
D. On a search head in the cluster.

Correct Answer: C
Section:
Explanation:
The Monitoring Console (MC) is the Splunk Enterprise monitoring tool that lets you view detailed topology and performance information about your Splunk Enterprise deployment1.The MC can be installed on any Splunk Enterprise instance that can access the data from all the instances in the deployment2.However, following the Splunk recommendations, the MC should be installed on the search head cluster deployer, which is a dedicated

instance that manages the configuration bundle for the search head cluster members3.This way, the MC can monitor the search head cluster as well as the indexer cluster and the forwarders, without affecting the performance or availability of the other instances4. The other options are not recommended because they either introduce additional load on the existing instances (such as A and D) or do not have access to the data from the search head cluster (such as B).
1:About the Monitoring Console - Splunk Documentation2:Add Splunk Enterprise instances to the Monitoring Console3:Configure the deployer - Splunk Documentation4: [Monitoring Console setup and use - Splunk Documentation]

QUESTION 51
A Splunk instance has crashed, but no crash log was generated. There is an attempt to determine what user activity caused the crash by running the following search:

```
index=_internal sourcetype=splunkd ("pipelines finished" OR "My GUID")
| transaction startswith="My GUID" endswith="pipelines finished" keepevicted=true keeporphans=true
| search closed_txn=0
| head 1
```

What does searching for closed_txn=0 do in this search?

A. Filters results to situations where Splunk was started and stopped multiple times.
B. Filters results to situations where Splunk was started and stopped once.
C. Filters results to situations where Splunk was stopped and then immediately restarted.
D. Filters results to situations where Splunk was started, but not stopped.

Correct Answer: D
Section:
Explanation:
Searching for closed_txn=0 in this search filters results to situations where Splunk was started, but not stopped. This means that the transaction was not completed, and Splunk crashed before it could finish the pipelines.The closed_txn field is added by the transaction command, and it indicates whether the transaction was closed by an event that matches the endswith condition1.A value of 0 means that the transaction was not closed, and a value of 1 means that the transaction was closed1. Therefore, option D is the correct answer, and options A, B, and C are incorrect.
1: transaction command overview

QUESTION 52
The master node distributes configuration bundles to peer nodes. Which directory peer nodes receive the bundles?

A. apps
B. deployment-apps
C. slave-apps
D. master-apps

Correct Answer: C
Section:
Explanation:
The master node distributes configuration bundles to peer nodes in theslave-appsdirectory under$SPLUNK_HOME/etc. The configuration bundle method is the only supported method for managing common configurations and app deployment across the set of peers.It ensures that all peers use the same versions of these files1.Bundles typically contain a subset of files (configuration files and assets) from$SPLUNK_HOME/etc/system,$SPLUNK_HOME/etc/apps, and$SPLUNK_HOME/etc/users2.The process of distributing knowledge bundles means that peers by default receive nearly the entire contents of the search head's apps3.