Free Questions for SPLK-2002

Shared by Banks on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

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

Answer:

D

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 2

Question Type: MultipleChoice

When converting from a single-site to a multi-site cluster, what happens to existing single-site clustered buckets?

Options:

A- They will continue to replicate within the origin site and age out based on existing policies.
B- They will maintain replication as required according to the single-site policies, but never age out.
C- They will be replicated across all peers in the multi-site cluster and age out based on existing policies.
D- They will stop replicating within the single-site and remain on the indexer they reside on and age out according to existing policies.

Answer:

D

Explanation:

When converting from a single-site to a multi-site cluster, existing single-site clustered buckets will maintain replication as required according to the single-site policies, but never age out. Single-site clustered buckets are buckets that were created before the conversion to a multi-site cluster. These buckets will continue to follow the single-site replication and search factors, meaning that they will have the same number of copies and searchable copies across the cluster, regardless of the site. These buckets will never age out, meaning that they will never be frozen or deleted, unless they are manually converted to multi-site buckets. Single-site clustered buckets will not continue to replicate within the origin site, because they will be distributed across the cluster according to the single-site policies. Single-site clustered buckets will not be replicated across all peers in the multi-site cluster, because they will follow the single-site replication factor, which may be lower than the multi-site total replication factor.Single-site clustered buckets will not stop replicating within the single-site and remain on the indexer they reside on, because they will still be subject to the replication and availability rules of the cluster

Question 3

Question Type: MultipleChoice

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?

Options:

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.

Answer:

D

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 4

Question Type: MultipleChoice

How does the average run time of all searches relate to the available CPU cores on the indexers?

Options:

A- Average run time is independent of the number of CPU cores on the indexers.
B- Average run time decreases as the number of CPU cores on the indexers decreases.
C- Average run time increases as the number of CPU cores on the indexers decreases.
D- Average run time increases as the number of CPU cores on the indexers increases.

Answer:

C

Explanation:

The average run time of all searches increases as the number of CPU cores on the indexers decreases. The CPU cores are the processing units that execute the instructions and calculations

for the data. The number of CPU cores on the indexers affects the search performance, because the indexers are responsible for retrieving and filtering the data from the indexes. The more CPU cores the indexers have, the faster they can process the data and return the results. The less CPU cores the indexers have, the slower they can process the data and return the results. Therefore, the average run time of all searches is inversely proportional to the number of CPU cores on the indexers. The average run time of all searches is not independent of the number of CPU cores on the indexers, because the CPU cores are an important factor for the search performance. The average run time of all searches does not decrease as the number of CPU cores on the indexers decreases, because this would imply that the search performance improves with less CPU cores, which is not true.The average run time of all searches does not increase as the number of CPU cores on the indexers increases, because this would imply that the search performance worsens with more CPU cores, which is not true

Question 5

Question Type: MultipleChoice

Which command is used for thawing the archive bucket?

Options:

A- Splunk collect
B- Splunk convert
C- Splunk rebuild
D- Splunk dbinspect

Answer:

C

Explanation:

The splunk rebuild command is used for thawing the archive bucket. Thawing is the process of restoring frozen data back to Splunk for searching. Frozen data is data that has been archived or deleted from Splunk after reaching the end of its retention period. To thaw a bucket, the user needs to copy the bucket from the archive location to the thaweddb directory under SPLUNK_HOME/var/lib/splunk and run the splunk rebuild command to rebuild the .tsidx files for the bucket. The splunk collect command is used for collecting diagnostic data from a Splunk instance. The splunk convert command is used for converting configuration files from one format to another. The splunk dbinspect command is used for inspecting the status and properties of the buckets in an index.

Question 6

Question Type: MultipleChoice

Consider a use case involving firewall dat

a. There is no Splunk-supported Technical Add-On, but the vendor has built one. What are the items that must be evaluated before installing the add-on? (Select all that apply.)

Options:

A- Identify number of scheduled or real-time searches.
B- Validate if this Technical Add-On enables event data for a data model.
C- Identify the maximum number of forwarders Technical Add-On can support.
D- Verify if Technical Add-On needs to be installed onto both a search head or indexer.

Answer:

A, B

Explanation:

A Technical Add-On (TA) is a Splunk app that contains configurations for data collection, parsing, and enrichment. It can also enable event data for a data model, which is useful for creating dashboards and reports. Therefore, before installing a TA, it is important to identify the number of scheduled or real-time searches that will use the data model, and to validate if the TA enables event data for a data model. The number of forwarders that the TA can support is not relevant, as the TA is installed on the indexer or search head, not on the forwarder.The installation location of the TA depends on the type of data and the use case, so it is not a fixed requirement

Question 7

Question Type: MultipleChoice

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

Options:

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

Answer:

A

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 8

Question Type: MultipleChoice

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?

Options:

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:

B

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.

P2P exams

Question 9

Question Type: MultipleChoice

When should multiple search pipelines be enabled?

Options:

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

P2P exams

Explanation:

Multiple search pipelines should be enabled only if CPU and memory resources are significantly under-utilized. Search pipelines are the processes that execute search commands and return results. Multiple search pipelines can improve the search performance by running concurrent searches in parallel. However, multiple search pipelines also consume more CPU and memory resources, which can affect the overall system performance. Therefore, multiple search pipelines should be enabled only if there are enough CPU and memory resources available, and if the system is not bottlenecked by disk I/O or network bandwidth.The number of concurrent users, the disk IOPS, and the Splunk Enterprise version are not relevant factors for enabling multiple search pipelines

Question 10

Question Type: MultipleChoice

Which Splunk server role regulates the functioning of indexer cluster?

Options:

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

Answer:

C

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 11

Question Type: MultipleChoice

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

Options:

A- telnet

B- tcpdump
C- splunk btool
D- splunk btprobe

Answer:

A, B

Explanation:

The telnet and tcpdump tools can be leveraged to diagnose connection problems between an indexer and forwarder. The telnet tool can be used to test the connectivity and port availability between the indexer and forwarder. The tcpdump tool can be used to capture and analyze the network traffic between the indexer and forwarder. The splunk btool command can be used to check the configuration files of the indexer and forwarder, but it cannot diagnose the connection problems. The splunk btprobe command does not exist, and it is not a valid tool.

Question 12

Question Type: MultipleChoice

Which of the following describe migration from single-site to multisite index replication?

Options:

A- A master node is required at each site.
B- Multisite policies apply to new data only.
C- Single-site buckets instantly receive the multisite policies.
D- Multisite total values should not exceed any single-site factors.

Answer:

B

Explanation:

Migration from single-site to multisite index replication only affects new data, not existing data. Multisite policies apply to new data only, meaning that data that is ingested after the migration will follow the multisite replication and search factors. Existing data, or data that was ingested before the migration, will retain the single-site policies, unless they are manually converted to

multisite buckets. Single-site buckets do not instantly receive the multisite policies, nor do they automatically convert to multisite buckets. Multisite total values can exceed any single-site factors, as long as they do not exceed the number of peer nodes in the cluster.A master node is not required at each site, only one master node is needed for the entire cluster

To Get Premium Files for SPLK-2002 Visit
https://www.p2pexams.com/products/splk-2002

For More Free Questions Visit
https://www.p2pexams.com/splunk/pdf/splk-2002

