Splunk.SPLK-2002.v2022-08-22.q74

Exam Code:	SPLK-2002
Exam Name:	Splunk Enterprise Certified Architect
Certification Provider:	Splunk
Free Question Number:	74
Version:	v2022-08-22
# of views:	2561
# of Questions views:	740
https://www.freeqas.com/qa/Splunk/SPLK-2002/Splunk.SPLK-2002.v2022-08-22.q74.html	

NEW QUESTION: 1
The KV store forms its own cluster within a SHC. What is the maximum number of SHC members KV store will form?
A. 25
B. 100
C. Unlimited
D. 50
Answer: C (LEAVE A REPLY)

NEW QUESTION: 2
When Splunk indexes data in a non clustered environment, what kind of files does it create by default?
A. Index and .tsidx files.
B. Rawdata and index files.
C. Compressed and .tsidx files.
D. Compressed and meta data files.
Answer: B (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Indexer/Aboutindexesandindexers

NEW QUESTION: 3
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.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 4
What is the logical first step when starting a deployment plan?
A. Determine what apps and use cases will be implemented.
B. Gather statistics on the expected adoption of Splunk for sizing.
C. Inventory the currently deployed logging infrastructure.
D. Collect the initial requirements for the deployment from all stakeholders.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 5
A multi-site indexer cluster can be configured using which of the following? (Select all that apply.)
A. Run a splunk edit cluster-config command from the CLI.
B. Directly edit SPLUNK_HOME/etc/system/local/server.conf
C. Directly edit SPLUNK_HOME/etc/system/default/server.conf
D. Via Splunk Web.
Answer: (SHOW ANSWER)

NEW QUESTION: 6
A Splunk instance has the following settings in SPLUNK_HOME/etc/system/local/server.conf:
[clustering]
mode = master
replication_factor = 2
pass4SymmKey = password123
Which of the following statements describe this Splunk instance? (Select all that apply.)
A. This is a multi-site cluster.
B. This cluster's search factor is 2.
C. This Splunk instance needs to be restarted.
D. This instance is missing the master_uri attribute.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 7
What is the algorithm used to determine captaincy in a Splunk search head cluster?
A. Raft distributed consensus.
B. Rapt distributed consensus.
C. Rift distributed consensus.
D. Round-robin distribution consensus.
Answer: A (LEAVE A REPLY)

Explanation/Reference: https://answers.splunk.com/answers/664102/need-to-know-about-raft-directory-on-search-head- c.html

NEW QUESTION: 8
Which of the following should be done when installing Enterprise Security on a Search Head Cluster? (Select all that apply.)
A. Copy the Enterprise Security configurations to the deployer.
B. Use the deployer to deploy Enterprise Security to the cluster members.
C. Install Enterprise Security on a staging instance.
D. Install Enterprise Security on the deployer.
Answer: B,D (LEAVE A REPLY)

NEW QUESTION: 9
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 sourcetype. Further investigation reveals that not all web logs 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 for this issue?
A. The indexers may have different configurations than the heavy forwarders.
B. The forwarders managed by the other department are an older version than the rest.
C. The search head may have different configurations than the indexers.
D. The data inputs are not properly configured across all the forwarders.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 10
The frequency in which a deployment client contacts the deployment server is controlled by what?
A. phoneHomeIntervalInSecs attribute in deploymentclient.conf
B. polling_interval attribute in outputs.conf
C. phoneHomeIntervalInSecs attribute in outputs.conf
D. polling_interval attribute in deploymentclient.conf
Answer: A (LEAVE A REPLY)

NEW QUESTION: 11
Which of the following is true regarding Splunk Enterprise performance? (Select all that apply.)
A. Adding search peers increases the maximum size of search results.
B. Adding RAM to an existing search heads provides additional search capacity.
C. Adding search peers increases the search throughput as search load increases.
D. Adding search heads provides additional CPU cores to run more concurrent searches.
Answer: B,D (LEAVE A REPLY)

Explanation/Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Capacity/HowsavedsearchesaffectSplunkEnterpriseperformance

NEW QUESTION: 12
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.confof the deployment server.
B. Check deploymentclient.confof the deployment client.
C. Check the content of SPLUNK_HOME/etc/appsof the deployment server.
D. Search for relevant events in splunkd.logof the deployment server.
Answer: A,B,C (LEAVE A REPLY)
Explanation/Reference: https://answers.splunk.com/answers/177021/why-is-deployment-client-not-picking-up-changes-
to.html

NEW QUESTION: 13
What does setting site=site0 on all Search Head Cluster members do in a multi-site indexer cluster?
A. Enables multisite search artifact replication.
B. Enables automatic search site affinity discovery.
C. Disables search site affinity.
D. Sets all members to dynamic captaincy.
Answer: (SHOW ANSWER)

NEW QUESTION: 14
Which of the following is an indexer clustering requirement?
A. Must use shared storage.
B. Must reside on a dedicated rack.
C. Must have at least three members.
D. Must share the same license pool.
Answer: D (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Admin/Distdeploylicenses

NEW QUESTION: 15
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/spool/searchpeers
B. SPLUNK_HOME/var/log/searchpeers
C. SPLUNK_HOME/var/run/searchpeers

D. SPLUNK_HOME/var/lib/searchpeers
Answer: C (LEAVE A REPLY)

NEW QUESTION: 16
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: (SHOW ANSWER)
Explanation/Reference:
https://docs.splunk.com/Documentation/SplunkCloud/8.0.0/Data/Monitornetworkports

NEW QUESTION: 17
Stakeholders have identified high availability for searchable data as their top priority. Which of the following best addresses this requirement?
A. Increasing the replication factor in the cluster.
B. Increasing the number of search heads in the cluster.
C. Increasing the number of CPUs on the indexers in the cluster.
D. Increasing the search factor in the cluster.
Answer: (SHOW ANSWER)

NEW QUESTION: 18
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: (SHOW ANSWER)

NEW QUESTION: 19
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 search is locked out?
A. 500GB. After this limit, search is locked out.
B. 300GB. After this limit, search is locked out.
C. Search is not locked out. Violations are still recorded.
D. 800GB. After this limit, search is locked out.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 20
Consider a use case involving firewall data. 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.)
A. Verify if Technical Add-On needs to be installed onto both a search head or indexer.
B. Identify the maximum number of forwarders Technical Add-On can support.
C. Identify number of scheduled or real-time searches.
D. Validate if this Technical Add-On enables event data for a data model.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 21
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. Add another search head to the cluster.
B. Change limits.conf value for max_searches_per_cpu to a higher value.
C. server.conf captain_is_adhoc_searchhead = true.
D. Create a job server on the cluster.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 22
Which search head cluster component is responsible for pushing knowledge bundles to search peers, replicating configuration changes to search head cluster members, and scheduling jobs across the search head cluster?
A. Master
B. Captain
C. Deployer
D. Deployment server
Answer: B (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture

NEW QUESTION: 23
In which phase of the Splunk Enterprise data pipeline are indexed extraction configurations processed?
A. Input
B. Search
C. Parsing
D. Indexing
Answer: C (LEAVE A REPLY)
Explanation/Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Admin/Configurationparametersandthedatapipeline

NEW QUESTION: 24
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.
Answer: B (LEAVE A REPLY)
Explanation/Reference: https://answers.splunk.com/answers/760348/search-head-version-compatibility.html

NEW QUESTION: 25
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-bundlecommand from the deployer.
C. Run the splunk resync shcluster-replicated-configcommand on this member.
D. Run the clean raftcommand on all members of the search head cluster.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 26
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.
Answer: D (LEAVE A REPLY)

Explanation/Reference:
https://docs.splunk.com/Documentation/CoE/ssf/Handbook/StakeholderReg

NEW QUESTION: 27
Which of the following statements describe search head clustering? (Select all that apply.)
A. A deployer is required.
B. The deployer must have sufficient CPU and network resources to process service requests and push configurations.
C. At least three search heads are needed.
D. Search heads must meet the high-performance reference server requirements.
Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 28
Which of the following are client filters available in serverclass.conf? (Select all that apply.)
A. Splunk server role.
B. IP address.
C. DNS name.
D. Platform (machine type).
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 29
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 (LEAVE A REPLY)
Explanation/Reference:

NEW QUESTION: 30
Which index-time props.conf attributes impact indexing performance? (Select all that apply.)
A. LINE_BREAKER
B. REPORT
C. ANNOTATE_PUNCT
D. SHOULD_LINEMERGE
Answer: A,D (LEAVE A REPLY)

NEW QUESTION: 31
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
Answer: C (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Indexer/Configurethepeerindexes

NEW QUESTION: 32
To reduce the captain's work load in a search head cluster, what setting will prevent scheduled searches from running on the captain?
A. adhoc_searchhead = true(on all members)
B. adhoc_searchhead = true(on the current captain)
C. captain_is_adhoc_searchhead = true(on all members)
D. captain_is_adhoc_searchhead = true(on the current captain)
Answer: D (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/DistSearch/Adhocclustermember

NEW QUESTION: 33
Which Splunk server role regulates the functioning of indexer cluster?
A. Indexer
B. Deployer
C. Master Node
D. Monitoring Console
Answer: C (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Deploy/Indexercluster

NEW QUESTION: 34
Which Splunk server role regulates the functioning of indexer cluster?
A. Indexer

B. Deployer
C. Master Node
D. Monitoring Console
Answer: (SHOW ANSWER)
Explanation
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Deploy/Indexercluster

NEW QUESTION: 35
When planning a search head cluster, which of the following is true?
A. The search head captain must be assigned to the largest search head in the cluster.
B. All search heads must use the same operating system.
C. All search heads must be members of the cluster (no standalone search heads).
D. All indexers must belong to the underlying indexer cluster (no standalone indexers).
Answer: (SHOW ANSWER)

NEW QUESTION: 36
Which of the following artifacts are included in a Splunk diagfile? (Select all that apply.)
A. OS settings.
B. Internal logs.
C. Customer data.
D. Configuration files.
Answer: B,D (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Troubleshooting/Generateadiag

NEW QUESTION: 37
Before users can use a KV store, an admin must create a collection. Where is a collection is defined?
A. kvstore.conf
B. collection.conf
C. collections.conf
D. kvcollections.conf
Answer: C (LEAVE A REPLY)
Explanation
Explanation/Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Knowledge/DefineaKVStorelookupinSplunkWeb

NEW QUESTION: 38
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
Answer: (SHOW ANSWER)
Explanation/Reference: https://answers.splunk.com/answers/237859/can-i-delete-all-data-from-a-kv-store-at-once.html

NEW QUESTION: 39
Which of the following is true regarding Splunk Enterprise performance? (Select all that apply.)
A. Adding RAM to an existing search heads provides additional search capacity.
B. Adding search peers increases the search throughput as search load increases.
C. Adding search heads provides additional CPU cores to run more concurrent searches.
D. Adding search peers increases the maximum size of search results.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 40
When configuring a Splunk indexer cluster, what are the default values for replication and search factor?
replication_factor = 2
A. search factor = 3
B. search_factor = 2
replication_factor = 3
C. search factor = 3
replication_factor = 3
D. search_factor = 2
replication_factor = 2
Answer: (SHOW ANSWER)

NEW QUESTION: 41
Which of the following statements describe licensing in a clustered Splunk deployment? (Select all that apply.)
A. Each cluster member requires its own clustering license.
B. Cluster members must share the same license pool and license master.
C. Replicated data does not count against licensing.
D. Free licenses do not support clustering.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 42
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: (SHOW ANSWER)

NEW QUESTION: 43
Which Splunk Enterprise offering has its own license?
A. Splunk Heavy Forwarder
B. Splunk Universal Forwarder
C. Splunk Cloud Forwarder
D. Splunk Forwarder Management
Answer: B (LEAVE A REPLY)

NEW QUESTION: 44
Which CLI command converts a Splunk instance to a license slave?
A. splunk add licenses
B. splunk list licenser-slaves
C. splunk list licenser-localslave
D. splunk edit licenser-localslave
Answer: D (LEAVE A REPLY)

NEW QUESTION: 45
How does the average run time of all searches relate to the available CPU cores on the indexers?
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 (LEAVE A REPLY)
Explanation/Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Capacity/Accommodatemanysimultaneoussearches

NEW QUESTION: 46
A Splunk user successfully extracted an ip address into a field called src_ip. Their colleague cannot see that field in their search results with events known to have src_ip. Which of the following may explain the problem? (Select all that apply.)
A. The field was extracted as a private knowledge object.
B. The events are tagged as communicate, but are missing the network tag.
C. The Typing Queue, which does regular expression replacements, is blocked.
D. The colleague did not explicitly use the field in the search and the search was set to Fast Mode.

Answer: D (LEAVE A REPLY)
Explanation/Reference: https://answers.splunk.com/answers/657187/map-command-field-not-being-evaluated.html

NEW QUESTION: 47
In the deployment planning process, when should a person identify who gets to see network data?
A. Deployment schedule
B. Topology diagramming
C. Data source inventory
D. Data policy definition
Answer: (SHOW ANSWER)
Explanation/Reference:

NEW QUESTION: 48
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.
Answer: (SHOW ANSWER)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCdeploymentoverview

NEW QUESTION: 49
Of the following types of files within an index bucket, which file type may consume the most disk?
A. Metadata (.data)
B. Rawdata
C. Bloom filter
D. Inverted index (.tsidx)
Answer: C (LEAVE A REPLY)

NEW QUESTION: 50
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.
Answer: C (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/Distsearchsystemrequirements

NEW QUESTION: 51
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 (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/Whatsearchheadssend

NEW QUESTION: 52
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 sourcetype. Further investigation reveals that not all web logs 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 for this issue?
A. The forwarders managed by the other department are an older version than the rest.
B. The indexers may have different configurations than the heavy forwarders.
C. The search head may have different configurations than the indexers.
D. The data inputs are not properly configured across all the forwarders.
Answer: (SHOW ANSWER)

NEW QUESTION: 53
Which component in the splunkd.log will log information related to bad event breaking?

A. Audittrail
B. IndexingPipeline
C. AggregatorMiningProcessor
D. EventBreaking
Answer: C (LEAVE A REPLY)

NEW QUESTION: 54
To optimize the distribution of primary buckets; when does primary rebalancing automatically occur? (Select all
that apply.)
A. Rolling restart completes.
B. Master node rejoins the cluster.
C. Captain joins or rejoins cluster.
D. A peer node joins or rejoins the cluster.
Answer: (SHOW ANSWER)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Rebalancethecluster

NEW QUESTION: 55
Which of the following will cause the greatest reduction in disk size requirements for a cluster of N indexers
running Splunk Enterprise Security?
A. Setting the cluster search factor to N-1.
B. Increasing the number of buckets per index.
C. Decreasing the data model acceleration range.
D. Setting the cluster replication factor to N-1.
Answer: D (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Systemrequirements

NEW QUESTION: 56
Which of the following is a way to exclude search artifacts when creating a diag?
A. SPLUNK_HOME/bin/splunk diag --exclude
B. SPLUNK_HOME/bin/splunk diag --debug --refresh
C. SPLUNK_HOME/bin/splunk diag --disable=dispatch
D. SPLUNK_HOME/bin/splunk diag --filter-searchstrings
Answer: A (LEAVE A REPLY)
Explanation
Explanation/Reference: https://splunkonbigdata.com/2018/10/01/splunk-diag/

NEW QUESTION: 57

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: (SHOW ANSWER)
Explanation/Reference: https://answers.splunk.com/answers/617608/can-we-increase-parallelingestionpipelines-in-a-
he.html

NEW QUESTION: 58
As a best practice, where should the internal licensing logs be stored?
A. License server.
B. Deployment layer.
C. Search head layer.
D. Indexing layer.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 59
What does setting site=site0on all Search Head Cluster members do in a multi-site indexer cluster?
A. Disables search site affinity.
B. Sets all members to dynamic captaincy.
C. Enables multisite search artifact replication.
D. Enables automatic search site affinity discovery.
Answer: A (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/DeploymultisiteSHC

NEW QUESTION: 60
Which of the following should be included in a deployment plan?
A. A comprehensive list of stakeholders, either direct or indirect.
B. Current and future topology diagrams of the IT environment.
C. Current logging details and data source inventory.
D. Business continuity and disaster recovery plans.
Answer: (SHOW ANSWER)

NEW QUESTION: 61
In search head clustering, which of the following methods can you use to transfer captaincy to a different
member? (Select all that apply.)

A. Use the Monitoring Console.
B. Use the Search Head Clustering settings menu from Splunk Web on any member.
C. Run the splunk transfer shcluster-captaincommand from the current captain.
D. Run the splunk transfer shcluster-captaincommand from the member you would like to become the captain.
Answer: B,D (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/DistSearch/Transfercaptain

NEW QUESTION: 62
A Splunk user successfully extracted an ip address into a field called src_ip. Their colleague cannot see that field in their search results with events known to have src_ip. Which of the following may explain the problem? (Select all that apply.)
A. The field was extracted as a private knowledge object.
B. The colleague did not explicitly use the field in the search and the search was set to Fast Mode.
C. The events are tagged as communicate, but are missing the network tag.
D. The Typing Queue, which does regular expression replacements, is blocked.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 63
Which of the following tasks should the architect perform when building a deployment plan? (Select all that apply.)
A. Review network topology.
B. Inventory data sources.
C. Install Splunk apps.
D. Use case checklist.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 64
What is the algorithm used to determine captaincy in a Splunk search head cluster?
A. Rift distributed consensus.

B. Raft distributed consensus.
C. Rapt distributed consensus.
D. Round-robin distribution consensus.
Answer: (SHOW ANSWER)

NEW QUESTION: 65
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
Answer: D (LEAVE A REPLY)
Explanation
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Sitereplicationfactor

NEW QUESTION: 66
Which of the following will cause the greatest reduction in disk size requirements for a cluster of N indexers running Splunk Enterprise Security?
A. Setting the cluster search factor to N-1.
B. Increasing the number of buckets per index.
C. Decreasing the data model acceleration range.
D. Setting the cluster replication factor to N-1.
Answer: (SHOW ANSWER)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Systemrequirements

NEW QUESTION: 67
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 clustered, assuming a high volume of saved/scheduled searches.
B. Two indexers not in a cluster, assuming users run many long searches.
C. Three indexers not in a cluster, assuming a long data retention period.
D. Two indexers clustered, assuming high availability is the greatest priority.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 68

Splunk configuration parameter settings can differ between multiple .conf files of the same name contained within different apps. Which of the following directories has the highest precedence?

A. System local directory.

B. System default directory.

C. App local directories, in ASCII order.

D. App default directories, in ASCII order.

Answer: (SHOW ANSWER)

Explanation/Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.2/Admin/Wheretofindtheconfigurationfiles

NEW QUESTION: 69

Which of the following is a way to exclude search artifacts when creating a diag?

A. SPLUNK_HOME/bin/splunk diag --debug --refresh

B. SPLUNK_HOME/bin/splunk diag --exclude

C. SPLUNK_HOME/bin/splunk diag --filter-searchstrings

D. SPLUNK_HOME/bin/splunk diag --disable=dispatch

Answer: B (LEAVE A REPLY)

NEW QUESTION: 70

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

A. site_mappings

B. available_sites

C. site_search_factor

D. site_replication_factor

Answer: A (LEAVE A REPLY)

Explanation/Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Decommissionasite

NEW QUESTION: 71

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

A. Single-site buckets instantly receive the multisite policies.

B. Multisite total values should not exceed any single-site factors.

C. A master node is required at each site.

D. Multisite policies apply to new data only.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 72

Which command is used for thawing the archive bucket?

A. Splunk dbinspect

B. Splunk rebuild
C. Splunk convert
D. Splunk collect
Answer: B (LEAVE A REPLY)

NEW QUESTION: 73
When configuring a Splunk indexer cluster, what are the default values for replication and search factor?
replication_factor = 2
A. search_factor = 2
replication_factor = 2
B. search factor = 3
replication_factor = 3
C. search_factor = 2
replication_factor = 3
D. search factor = 3
Answer: A (LEAVE A REPLY)
Explanation/Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Thesearchfactor

NEW QUESTION: 74
Which server.conf attribute should be added to the master node's server.conf file when decommissioning a site in an indexer cluster?
A. available_sites
B. site_mappings
C. site_search_factor
D. site_replication_factor
Answer: B (LEAVE A REPLY)