Actual exam question from Splunk's SPLK-3003

Question #: 1

Topic #: 1

[All SPLK-3003 Questions]

How does Monitoring Console (MC) initially identify the server role(s) of a new Splunk Instance?

A. The MC uses a REST endpoint to query the server.

B. Roles are manually assigned within the MC.

C. Roles are read from distsearch.conf.

D. The MC assigns all possible roles by default.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003
Question #: 2
Topic #: 1
[All SPLK-3003 Questions]

A customer has asked for a five-node search head cluster (SHC), but does not have the storage budget to use a replication factor greater than 2. They would like to understand what might happen in terms of the users' ability to view historic scheduled search results if they log onto a search head which doesn't contain one of the 2 copies of a given search artifact.
Which of the following statements best describes what would happen in this scenario?

A. The search head that the user has logged onto will proxy the required artifact over to itself from a search head that currently holds a copy. A copy will also be replicated from that search head permanently, so it is available for future use.

B. Because the dispatch folder containing the search results is not present on the search head, the user will not be able to view the search results.

C. The user will not be able to see the results of the search until one of the search heads is restarted, forcing synchronization of all dispatched artifacts across all search heads.

D. The user will not be able to see the results of the search until the Splunk administrator issues the apply shcluster-bundle command on the search head deployer, forcing synchronization of all dispatched artifacts across all search heads.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 3

Topic #: 1

[All SPLK-3003 Questions]

Monitoring Console (MC) health check configuration items are stored in which configuration file?

A. healthcheck.conf

B. alert_actions.conf

C. distsearch.conf

D. checklist.conf

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 4

Topic #: 1

[All SPLK-3003 Questions]

What should be considered when running the following CLI commands with a goal of accelerating an index cluster migration to new hardware?

```
$SPLUNK_HOME/bin/splunk edit cluster-config -max_peer_build_load 3

$SPLUNK_HOME/bin/splunk edit cluster-config -max_peer_rep_load 6
```

server.conf

```
[clustering]

max_peer_build_load = 2

max_peer_rep_load = 5
```

A. Data ingestion rate

B. Network latency and storage IOPS

C. Distance and location

D. SSL data encryption

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 5

Topic #: 1

[All SPLK-3003 Questions]

Which statement is true about subsearches?

A. Subsearches are faster than other types of searches.

B. Subsearches work best for joining two large result sets.

C. Subsearches run at the same time as their outer search.

D. Subsearches work best for small result sets.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003
Question #: 6
Topic #: 1
[All SPLK-3003 Questions]

A customer has been using Splunk for one year, utilizing a single/all-in-one instance. This single Splunk server is now struggling to cope with the daily ingest rate. Also, Splunk has become a vital system in day-to-day operations making high availability a consideration for the Splunk service. The customer is unsure how to design the new environment topology in order to provide this.
Which resource would help the customer gather the requirements for their new architecture?

A. Direct the customer to the docs.splunk.com and tell them that all the information to help them select the right design is documented there.

B. Ask the customer to engage with the sales team immediately as they probably need a larger license.

C. Refer the customer to answers.splunk.com as someone else has probably already designed a system that meets their requirements.

D. Refer the customer to the Splunk Validated Architectures document in order to guide them through which approved architectures could meet their requirements.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 7

Topic #: 1

[All SPLK-3003 Questions]

The customer has an indexer cluster supporting a wide variety of search needs, including scheduled search, data model acceleration, and summary indexing.

Here is an excerpt from the cluster mater's server.conf:

```
[clustering]
replication_factor=2
search_factor=1
summary_replication-false
```

Which strategy represents the minimum and least disruptive change necessary to protect the searchability of the indexer cluster in case of indexer failure?

A. Enable maintenance mode on the CM to prevent excessive fix-up and bring the failed indexer back online.

B. Leave replication_factor=2, increase search_factor=2 and enable summary_replication.

C. Convert the cluster to multi-site and modify the server.conf to be site_replication_factor=2, site_search_factor=2.

D. Increase replication_factor=3, search_factor=2 to protect the data, and allow there to always be a searchable copy.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 8

Topic #: 1

[All SPLK-3003 Questions]

What is the primary driver behind implementing indexer clustering in a customer's environment?

A. To improve resiliency as the search load increases.

B. To reduce indexing latency.

C. To scale out a Splunk environment to offer higher performance capability.

D. To provide higher availability for buckets of data.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 9

Topic #: 1

[All SPLK-3003 Questions]

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

A. Deployer sharing with master cluster.

B. License master that has 50 clients or more.

C. Cluster master node

D. Production Search Head

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003
Question #: 10
Topic #: 1
[All SPLK-3003 Questions]

A customer has downloaded the Splunk App for AWS from Splunkbase and installed it in a search head cluster following the instructions using the deployer. A power user modifies a dashboard in the app on one of the search head cluster members. The app containing an updated dashboard is upgraded to the latest version by following the instructions via the deployer.
What happens?

A. The updated dashboard will not be deployed globally to all users, due to the conflict with the power user's modified version of the dashboard.

B. Applying the search head cluster bundle will fail due to the conflict.

C. The updated dashboard will be available to the power user.

D. The updated dashboard will not be available to the power user; they will see their modified version.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 11

Topic #: 1

[All SPLK-3003 Questions]

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

A. Create a tiered deployment server topology.

B. Reduce the phone home interval to 6 seconds.

C. Leave the phone home interval at 60 seconds.

D. Increase the phone home interval to 600 seconds.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 12

Topic #: 1

[All SPLK-3003 Questions]

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 = $7$XcXl1lu3820Jbui14oVe324+mvx6gCKKv6kf2zEaVB6Ie4DcZ647CnLVlfW
```

B.

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

C.

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

D.

```
[clustering]
mode = forwarder
pass4SymmKey = $7$PU9SBXww63Vz3UJdDYGIN0UrdscRh83ssC2pEpwE6P3gn50iNF094g==
```

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 13

Topic #: 1

[All SPLK-3003 Questions]

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

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

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

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

D. Carefully design bigger apps containing multiple configs.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 14

Topic #: 1

[All SPLK-3003 Questions]

Which of the following processor occur in the indexing pipeline?

A. tcp out, syslog out

B. Regex replacement, annotator

C. Aggregator

D. UTF-8, linebreaker, header

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 15

Topic #: 1

[All SPLK-3003 Questions]

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

A. AUTO_KV_JSON

B. BREAK_ONLY_BEFORE_DATE

C. SHOULD_LINEMERGE

D. ANNOTATE_PUNCT

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 16

Topic #: 1

[All SPLK-3003 Questions]

A customer has a new set of hardware to replace their aging indexers. What method would reduce the amount of bucket replication operations during the migration process?

A. Disable the indexing ports on the old indexers.

B. Disable replication ports on the old indexers.

C. Put the old indexers into manual detention.

D. Put the old indexers into automatic detention.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 17

Topic #: 1

[All SPLK-3003 Questions]

When a bucket rolls from cold to frozen on a clustered indexer, which of the following scenarios occurs?

A. All replicated copies will be rolled to frozen; original copies will remain.

B. Replicated copies of the bucket will remain on all other indexers and the Cluster Master (CM) assigns a new primary bucket.

C. The bucket rolls to frozen on all clustered indexers simultaneously.

D. Nothing. Replicated copies of the bucket will remain on all other indexers until a local retention rule causes it to roll.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 18

Topic #: 1

[All SPLK-3003 Questions]

A site from a multi-site indexer cluster needs to be decommissioned. Which of the following actions must be taken?

A. Nothing. Decommissioning a site is not possible.

B. Create an alias for where the new data should be sent.

C. Remove the site from the list of available sites.

D. Remove the site from the list of available sites and create an alias for where the new data should be sent.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 19

Topic #: 1

[All SPLK-3003 Questions]

A customer wants to implement LDAP because managing local Splunk users is becoming too much of an overhead. What configuration details are needed from the customer to implement LDAP authentication?

A. API: Python script with PAM/RADIUS details.

B. LDAP server: port, bind user credentials, path/to/groups, path/to/user.

C. LDAP server: port, bind user credentials, base DN for groups, base DN for users.

D. LDAP REST details, base DN for groups, base DN for users.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 20

Topic #: 1

[All SPLK-3003 Questions]

A customer has a search cluster (SHC) of six members split evenly between two data centers (DC). The customer is concerned with network connectivity between the two DCs due to frequent outages. Which of the following is true as it relates to SHC resiliency when a network outage occurs between the two DCs?

A. The SHC will function as expected as the SHC deployer will become the new captain until the network communication is restored.

B. The SHC will stop all scheduled search activity within the SHC.

C. The SHC will function as expected as the minimum required number of nodes for a SHC is 3.

D. The SHC will function as expected as the SHC captain will fall back to previous active captain in the remaining site.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 21

Topic #: 1

[All SPLK-3003 Questions]

A [script://] input sends data to a Splunk forwarder using which method?

A. UDP stream

B. TCP stream

C. Temporary file

D. STDOUT/STDERR

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 22

Topic #: 1

[All SPLK-3003 Questions]

A customer wants to understand how Splunk bucket types (hot, warm, cold) impact search performance within their environment. Their indexers have a single storage device for all data. What is the proper message to communicate to the customer?

A. The bucket types (hot, warm, or cold) have the same search performance characteristics within the customer's environment.

B. While hot, warm, and cold buckets have the same search performance characteristics within the customers environment, due to their optimized structure, the thawed buckets are the most performant.

C. Searching hot and warm buckets result in best performance because by default the cold buckets are miniaturized by removing TSIDX files to save on storage cost.

D. Because the cold buckets are written to a cheaper/slower storage volume, they will be slower to search compared to hot and warm buckets which are written to Solid State Disk (SSD).

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003
Question #: 23
Topic #: 1
[All SPLK-3003 Questions]

An index receives approximately 50GB of data per day per indexer at an even and consistent rate. The customer would like to keep this data searchable for a minimum of 30 days. In addition, they have hourly scheduled searches that process a week's worth of data and are quite sensitive to search performance.
Given ideal conditions (no restarts, nor drops/bursts in data volume), and following PS best practices, which of the following sets of indexes.conf settings can be leveraged to meet the requirements?

A. frozenTimePeriodInSecs, maxDataSize, maxVolumeDataSizeMB, maxHotBuckets

B. maxDataSize, maxTotalDataSizeMB, maxHotBuckets, maxGlobalDataSizeMB

C. maxDataSize, frozenTimePeriodInSecs, maxVolumeDataSizeMB

D. frozenTimePeriodInSecs, maxWarmDBCount, homePath.maxDataSizeMB, maxHotSpanSecs

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 24

Topic #: 1

[All SPLK-3003 Questions]

A customer has a Universal Forwarder (UF) with an inputs.conf monitoring its splunkd.log. The data is sent through a heavy forwarder to an indexer.

Where does the Index time parsing occur?

A. Indexer

B. Universal forwarder

C. Search head

D. Heavy forwarder

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 25

Topic #: 1

[All SPLK-3003 Questions]

The customer wants to migrate their current Splunk Index cluster to new hardware to improve indexing and search performance. What is the correct process and procedure for this task?

A. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the same configuration via the deployment server. 3. Decommission old peers one at a time. 4. Remove old peers from the CM's list. 5. Update forwarders to forward to the new peers.

B. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the cluster bundle and the same configuration as original peers. 3. Decommission old peers one at a time. 4. Remove old peers from the CM's list. 5. Update forwarders to forward to the new peers.

C. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the same configuration via the deployment server. 3. Update forwarders to forward to the new peers. 4. Decommission old peers on at a time. 5. Restart the cluster master (CM).

D. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the cluster bundle and the same configuration as original peers. 3. Update forwarders to forward to the new peers. 4. Decommission old peers one at a time. 5. Remove old peers from the CM's list.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 26

Topic #: 1

[All SPLK-3003 Questions]

Consider the scenario where the /var/log directory contains the files secure, messages, cron, audit. A customer has created the following inputs.conf stanzas in the same Splunk app in order to attempt to monitor the files secure and messages:

```
[monitor:///var/log]
sourcetype = syslog
index = secrutiy
disabled = false
whitelist = messages

[monitor:///var/log]
sourcetype  = syslog
index = security
disabled = false
whitelist = secure
```

Which file(s) will actually be actively monitored?

A. /var/log/secure

B. /var/log/messages

C. /var/log/messages, /var/log/cron, /var/log/audit, /var/log/secure

D. /var/log/secure, /var/log/messages

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 27

Topic #: 1

[All SPLK-3003 Questions]

A customer has written the following search:

```
sourcetype=purchase:orders
| table _time, customer, product, amount, order_id
| stats count sum(amount) AS amount latest (_time) AS _time by customer, order_id
| search customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| table _time, customer, order_id, amount, vip_status
| search vip_status= "true"
```

How can the search be rewritten to maximize efficiency?

A.

```
index=sales sourcetype=purchase:orders
| table _time, customer, product, amount, order_id
| stats count sum(amount) AS amount latest (_time) AS _time by customer, order_id
| search customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| table _time, customer, order_id, amount, vip_status
| search vip_status= "true"
```

B.

```
index=proxy source=proxy:data:syslog user= "timmy*"
| table _time, user, url, duration, category, action
| stats count sum(duration) AS duration last(url) AS url latest (_time) AS _time by user
| lookup user_status user OUTPUT status
| table _time, user, status
```

C.

```
index=sales sourcetype=purchase:orders customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| search vip_status= "true"
| stats sum(amount) AS amount latest (_time) AS _time by customer, order_id
| table _time, customer, order_id, amount
```

D.

```
index=sales sourcetype=purchase:orders customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| stats count sum(amount) AS amount latest (_time) AS _time by customer, order_id
| search vip_status= "true"
| table _time, customer, order_id, amount, vip_status
```

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 28

Topic #: 1

[All SPLK-3003 Questions]

How could a role in which all users must specify an index=clause in all searches be configured?

A. Set the authorize.conf setting: srchIndexesDefault to no value.

B. Set the authorize.conf setting: srchFilter to no value.

C. Set the authorize.conf setting: srchIndexesAllowed to no value.

D. Set the authorize.conf setting: srchJobsQuota to no value.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 29

Topic #: 1

[All SPLK-3003 Questions]

In which of the following scenarios should base configurations be used to provide consistent, repeatable, and supportable configurations?

A. For non-production environments to keep their configurations in sync.

B. To ensure every customer has exactly the same base settings.

C. To provide settings that do not need to be customized to meet customer requirements.

D. To provide settings that can be customized to meet customer requirements.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 30

Topic #: 1

[All SPLK-3003 Questions]

Data can be onboarded using apps, Splunk Web, or the CLI.

Which is the PS preferred method?

A. Create UDP input port 9997 on a UF.

B. Use the add data wizard in Splunk Web.

C. Use the inputs.conf file.

D. Use a scripted input to monitor a log file.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 31

Topic #: 1

[All SPLK-3003 Questions]

Which of the following statements applies to indexer discovery?

A. The Cluster Master (CM) can automatically discover new indexers added to the cluster.

B. Forwarders can automatically discover new indexers added to the cluster.

C. Deployment servers can automatically configure new indexers added to the cluster.

D. Search heads can automatically discover new indexers added to the cluster.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 32

Topic #: 1

[All SPLK-3003 Questions]

The data in Splunk is now subject to auditing and compliance controls. A customer would like to ensure that at least one year of logs are retained for both Windows and Firewall events. What data retention controls must be configured?

A. maxTotalDataSizeMB and frozenTimePeriodInSecs

B. coldToFrozenDir and coldToFrozenScript

C. Splunk Volume and maxTotalDataSizMB

D. Splunk Volume and frozenTimePeriodInSecs

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 33

Topic #: 1

[All SPLK-3003 Questions]

What happens when an index cluster peer freezes a bucket?

A. All indexers with a copy of the bucket will delete it.

B. The cluster master will ensure another copy of the bucket is made on the other peers to meet the replication settings.

C. The cluster master will no longer perform fix-up activities for the bucket.

D. All indexers with a copy of the bucket will immediately roll it to frozen.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003
Question #: 34
Topic #: 1
[All SPLK-3003 Questions]

A customer has the following Splunk instances within their environment: An indexer cluster consisting of a cluster master/master node and five clustered indexers, two search heads (no search head clustering), a deployment server, and a license master. The deployment server and license master are running on their own single-purpose instances. The customer would like to start using the Monitoring Console (MC) to monitor the whole environment.
On the MC instance, which instances will need to be configured as distributed search peers by specifying them via the UI using the settings menu?

A. Just the cluster master/master node.

B. Indexers, search heads, deployment server, license master, cluster master/master node.

C. Search heads, deployment server, license master, cluster master/master node

D. Deployment server, license master

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 35

Topic #: 1

[All SPLK-3003 Questions]

What does Splunk do when it indexes events?

A. Extracts the top 10 fields.

B. Extracts metadata fields such as host, source, sourcetype.

C. Performs parsing, merging, and typing processes on universal forwarders.

D. Create report acceleration summaries.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 36

Topic #: 1

[All SPLK-3003 Questions]

What is the default push mode for a search head cluster deployer app configuration bundle?

A. full

B. merge_to_default

C. default_only

D. local_only

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 37

Topic #: 1

[All SPLK-3003 Questions]

In which of the following scenarios is a subsearch the most appropriate?

A. When joining results from multiple indexes.

B. When dynamically filtering hosts.

C. When filtering indexed fields.

D. When joining multiple large datasets.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003
Question #: 38
Topic #: 1
[All SPLK-3003 Questions]

A customer has implemented their own Role Based Access Control (RBAC) model to attempt to give the Security team different data access than the Operations team by creating two new Splunk roles "" security and operations. In the srchIndexesAllowed setting of authorize.conf, they specified the network index under the security role and the operations index under the operations role. The new roles are set up to inherit the default user role.
If a new user is created and assigned to the operations role only, which indexes will the user have access to search?

A. operations, network, _internal, _audit

B. operations

C. No Indexes

D. operations, network

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 39

Topic #: 1

[All SPLK-3003 Questions]

A customer would like Splunk to delete files after they've been ingested. The Universal Forwarder has read/write access to the directory structure. Which input type would be most appropriate to use in order to ensure files are ingested and then deleted afterwards?

A. Script

B. Batch

C. Monitor

D. Fschange

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 40

Topic #: 1

[All SPLK-3003 Questions]

In which directory should base config app(s) be placed to initialize an indexer?

A. $SPLUNK_HOME/etc/<app_name>

B. $SPLUNK_HOME/etc/apps

C. $SPLUNK_HOME/etc/system/local

D. $SPLUNK_HOME/etc/slave-apps

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 41

Topic #: 1

[All SPLK-3003 Questions]

As a best practice which of the following should be used to ingest data on clustered indexers?

A. Monitoring (via a process), collecting data (modular inputs) from remote systems/applications

B. Modular inputs, HTTP Event Collector (HEC), inputs.conf monitor stanza

C. Actively listening on ports, monitoring (via a process), collecting data from remote systems/applications

D. splunktcp, splunktcp-ssl, HTTP Event Collector (HEC)

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 42

Topic #: 1

[All SPLK-3003 Questions]

When adding a new search head to a search head cluster (SHC), which of the following scenarios occurs?

A. The new search head connects to the captain and replays any recent configuration changes to bring it up to date.

B. The new search head connects to the deployer and replays any recent configuration changes to bring it up to date.

C. The new search head connects to the captain and pulls the most recently deployed bundle. It then connects to the deployer and replays any recent configuration changes to bring it up to date.

D. The new search head connects to the deployer and pulls the most recently deployed bundle. It then connects to the captain and replays any recent configuration changes to bring it up to date.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 43

Topic #: 1

[All SPLK-3003 Questions]

A customer wants to migrate from using Splunk local accounts to use Active Directory with LDAP for their Splunk user accounts instead. Which configuration files must be modified to connect to an Active Directory LDAP provider?

A. authentication.conf, authorize.conf, ldap.conf

B. authentication.conf, ldap.conf

C. authentication.conf

D. authorize.conf, authentication.conf

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 44

Topic #: 1

[All SPLK-3003 Questions]

A customer has a number of inefficient regex replacement transforms being applied. When under heavy load the indexers are struggling to maintain the expected indexing rate. In a worst case scenario, which queue(s) would be expected to fill up?

A. Typing, merging, parsing, input

B. Parsing

C. Typing

D. Indexing, typing, merging, parsing, input

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 45

Topic #: 1

[All SPLK-3003 Questions]

A new single-site three indexer cluster is being stood up with replication_factor:2, search_factor:2. At which step would the Indexer Cluster be classed as "˜Indexing Ready' and be able to ingest new data?

Step 1: Install and configure Cluster Master (CM)/Master Node with base clustering stanza settings, restarting CM.

Step 2: Configure a base app in etc/master-apps on the CM to enable a splunktcp input on port 9997 and deploy index creation configurations.

Step 3: Install and configure Indexer 1 so that once restarted, it contacts the CM, download the latest config bundle.

Step 4: Indexer 1 restarts and has successfully joined the cluster.

Step 5: Install and configure Indexer 2 so that once restarted, it contacts the CM, downloads the latest config bundle

Step 6: Indexer 2 restarts and has successfully joined the cluster.

Step 7: Install and configure Indexer 3 so that once restarted, it contacts the CM, downloads the latest config bundle.

Step 8: Indexer 3 restarts and has successfully joined the cluster.

A. Step 2

B. Step 4

C. Step 6

D. Step 8

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 46

Topic #: 1

[All SPLK-3003 Questions]

A new search head cluster is being implemented. Which is the correct command to initialize the deployer node without restarting the search head cluster peers?

A. $SPLUNK_HOME/bin/splunk apply shcluster-bundle

B. $SPLUNK_HOME/bin/splunk apply cluster-bundle

C. $SPLUNK_HOME/bin/splunk apply shcluster-bundle ""action stage

D. $SPLUNK_HOME/bin/splunk apply cluster-bundle ""action stage

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 47

Topic #: 1

[All SPLK-3003 Questions]

What is required to setup the HTTP Event Collector (HEC)?

A. Each HEC input requires a unique name but token values can be shared.

B. Each HEC input requires an existing forwarder output group.

C. Each HEC input entry must contain a valid token.

D. Each HEC input requires a Source name field.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 48

Topic #: 1

[All SPLK-3003 Questions]

In the diagrammed environment shown below, the customer would like the data read by the universal forwarders to set an indexed field containing the UF's host name. Where would the parsing configurations need to be installed for this to work?

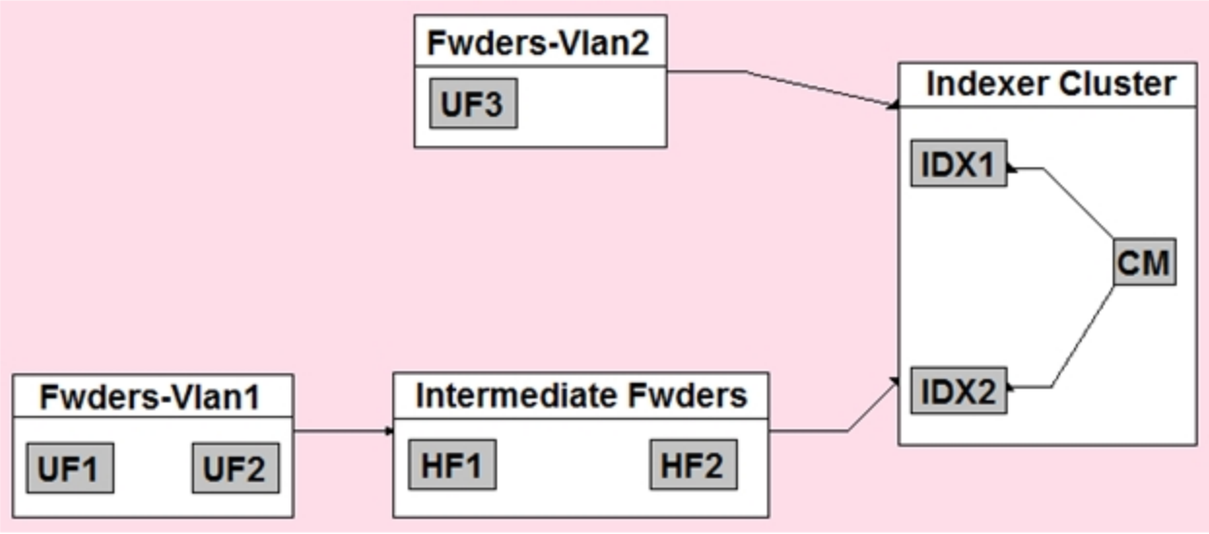

A. All universal forwarders.

B. Only the indexers.

C. All heavy forwarders.

D. On all parsing Splunk instances.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 49

Topic #: 1

[All SPLK-3003 Questions]

Report acceleration has been enabled for a specific use case. In which bucket location is the corresponding CSV file located?

A. thawedPath

B. summaryHomePath

C. tstatsHomePath

D. homePath, coldPath

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 50

Topic #: 1

[All SPLK-3003 Questions]

Which command is most efficient in finding the pass4SymmKey of an index cluster?

A. find / -name server.conf ""print | grep pass4SymKey

B. $SPLUNK_HOME/bin/splunk search | rest splunk_server=local /servicesNS/-/unhash_app/storage/passwords

C. $SPLUNK_HOME/bin/splunk btool server list clustering | grep pass4SymmKey

D. $SPLUNK_HOME/bin/splunk btool clustering list clustering --debug | grep pass4SymmKey

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 51

Topic #: 1

[All SPLK-3003 Questions]

Where does the bloomfilter reside?

A. $SPLUNK_HOME/var/lib/splunk/indexfoo/db/db_1553504858_1553504507_8

B. $SPLUNK_HOME/var/lib/splunk/indexfoo/db/db_1553504858_1553504507_8/*.tsidx

C. $SPLUNK_HOME/var/lib/splunk/fishbucket

D. $SPLUNK_HOME/var/lib/splunk/indexfoo/db/db_1553504858_1553504507_8/rawdata

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 52

Topic #: 1

[All SPLK-3003 Questions]

A customer is having issues with truncated events greater than 64K. What configuration should be deployed to a universal forwarder (UF) to fix the issue?

A. None. Splunk default configurations will process the events as needed; the UF is not causing truncation.

B. Configure the best practice magic 6 or great 8 props.conf settings.

C. EVENT_BREAKER_ENABLE and EVENT_BREAKER regular expression settings per sourcetype.

D. Global EVENT_BREAKER_ENABLE and EVENT_BREAKER regular expression settings.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 53

Topic #: 1

[All SPLK-3003 Questions]

A customer has a network device that transmits logs directly with UDP or TCP over SSL. Using PS best practices, which ingestion method should be used?

A. Open a TCP port with SSL on a heavy forwarder to parse and transmit the data to the indexing tier.

B. Open a UDP port on a universal forwarder to parse and transmit the data to the indexing tier.

C. Use a syslog server to aggregate the data to files and use a heavy forwarder to read and transmit the data to the indexing tier.

D. Use a syslog server to aggregate the data to files and use a universal forwarder to read and transmit the data to the indexing tier.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 54

Topic #: 1

[All SPLK-3003 Questions]

As data enters the indexer, it proceeds through a pipeline where event processing occurs. In which pipeline does line breaking occur?

A. Indexing

B. Typing

C. Merging

D. Parsing

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 55

Topic #: 1

[All SPLK-3003 Questions]

A customer has a multisite cluster (two sites, each site in its own data center) and users experiencing a slow response when searches are run on search heads located in either site. The Search Job Inspector shows the delay is being caused by search heads on either site waiting for results to be returned by indexers on the opposing site. The network team has confirmed that there is limited bandwidth available between the two data centers, which are in different geographic locations. Which of the following would be the least expensive and easiest way to improve search performance?

A. Configure site_search_factor to ensure a searchable copy exists in the local site for each search head.

B. Move all indexers and search heads in one of the data centers into the same site.

C. Install a network pipe with more bandwidth between the two data centers.

D. Set the site setting on each indexer in the server.conf clustering stanza to be the same for all indexers regardless of site.

Show Suggested Answer

Actual exam question from Splunk's SPLK-3003

Question #: 56

Topic #: 1

[All SPLK-3003 Questions]

A customer is using regex to whitelist access logs and secure logs from a web server, but only the access logs are being ingested. Which troubleshooting resource would provide insight into why the secure logs are not being ingested?

A. list monitor

B. oneshot

C. btprobe

D. tailingprocessor

Show Suggested Answer