

SPLK-3003.VCEplus.premium.exam.85q

Number: SPLK-3003
Passing Score: 800
Time Limit: 120 min
File Version: 1.0



Website: <https://vceplus.com>

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

SPLK-3003

Splunk Core Certified Consultant

Version 1.0

Exam A

QUESTION 1

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 3 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`

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/Customizehealthcheck>

QUESTION 4

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 5 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.



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://community.splunk.com/t5/Archive/Looking-for-way-to-explain-why-subsearches-are-so-slow/m-p/479133>

QUESTION 6

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf>

QUESTION 7

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 master'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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 8 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.

Correct Answer: D

Section: (none)

Explanation



Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howclusteredsearchworks>

QUESTION 9 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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/WheretohostDMC>

QUESTION 10

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 11

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 12 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 = $7$XcX11lu3820Jbui14oVe324+mvx6gCKKv6kf2zEaVB6Te4DcZ647CnLV1fW

[clustering]
mode = master
pass4SymmKey = $7$tYTXzke+1r+3DULTHHDUTmYOXdTzJpXm21XwMARrJE20jsmicp9C3ni0

[indexer_discovery]
pass4SymmKey = idxdiscovery

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



B.

C.

D.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/indexerdiscovery>

QUESTION 13

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: https://www.splunk.com/en_us/blog/platform/adding-a-deployment-server-forwarder-management-to-a-new-or-existing-splunk-cloud-or-splunk-enterprise-deployment.html

QUESTION 14 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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howindexingworks#Event_processing_and_the_data_pipeline

QUESTION 15 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`



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.0.6/Data/Configureeventlinebreaking>

QUESTION 16

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 17 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Bucketsandclusters>

QUESTION 18

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 19 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.learnsplunk.com/splunk-ldap-authentication-configuration.html>

QUESTION 20

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 21

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

- A. UDP stream
- B. TCP stream
- C. Temporary file
- D. STDOUT/STDERR

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/latest/Admin/inputsconf>

QUESTION 22

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 customer's 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).

Correct Answer: D

Section: (none)

Explanation



Explanation/Reference:

QUESTION 23

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`

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 24 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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.learnsplunk.com/splunk-interview-questions.html>

QUESTION 25

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 one 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:



QUESTION 26

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 = securitiy
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`

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 27

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"
```

```
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
```

```
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
```

```
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
```



B. C.

D.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 28 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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 29 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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/latest/Admin/Wheretofindtheconfigurationfiles>



QUESTION 30 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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/Howdoyouwanttoadddata>

QUESTION 31 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.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/Connectclustersearchheadstosearchpeers>

QUESTION 32 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`

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Setaretirementandarchivingpolicy>

QUESTION 33

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.

Correct Answer: C

Section: (none)

Explanation

**Explanation/Reference:**

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Bucketsandclusters>

QUESTION 34

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 35 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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howindexingworks#:~:text=Splunk%20Enterprise%20can%20index%20any,events%20indexes%20and%20metrics%20indexes>

QUESTION 36 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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/PropagateSHCconfigurationchanges#:~:text=The%20deployer%20push%20mode%20determines,default%20push%20mode%20is%20merge_to_default%20

QUESTION 37 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.



Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 38

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 39

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://community.splunk.com/t5/Getting-Data-In/Is-it-possible-to-have-a-Splunk-universal-forwarder-read-a-td-p/172752>

QUESTION 40 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Manageappdeployment>

QUESTION 41 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)



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 42

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 43 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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Security/ConfigureLDAPwithconfigurationfiles>

QUESTION 44

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 45

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 46

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/PropagateSHCconfigurationchanges>

QUESTION 47 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.

Correct Answer: C

Section: (none)

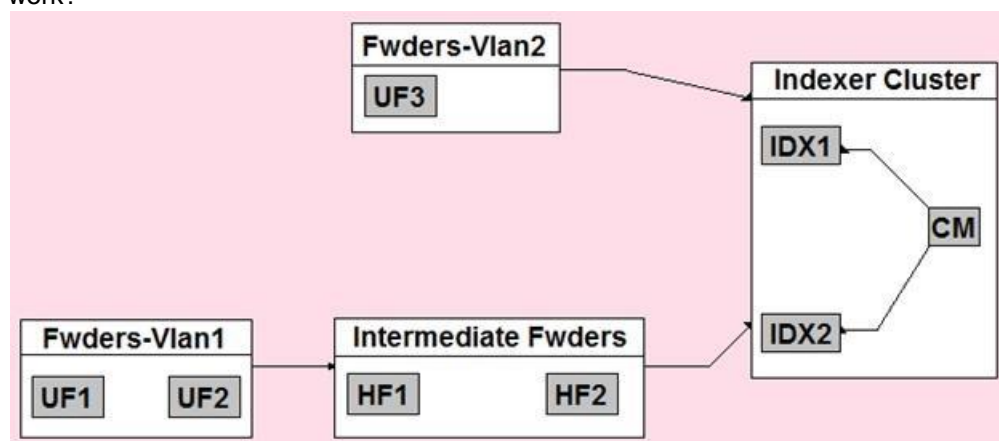
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/UseTheHTTPEventCollector>

QUESTION 48

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 49

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Knowledge/Manageacceleratedsearchsummaries>

QUESTION 50 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`

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://community.splunk.com/t5/Deployment-Architecture/Which-instance-or-configuration-file-in-my-Splunk-environment/m-p/241486>

QUESTION 51 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`

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 52

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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/Resolvedataqualityissues>

QUESTION 53

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 54 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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howindexingworks#Event_processing_and_the_data_pipeline

QUESTION 55

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 56 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`

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 57

A customer with a large distributed environment has blacklisted a large `lookup` from the search bundle to decrease the bundle size using `distsearch.conf`. After this change, when running searches utilizing the `lookup` that was blacklisted they see error messages in the Splunk Search UI stating the `lookup` file does not exist.

What can the customer do to resolve the issue?

- A. The search needs to be modified to ensure the `lookup` command specifies parameter `local=true`.
- B. The blacklisted `lookup` definition stanza needs to be modified to specify setting `allow_caching=true`.
- C. The search needs to be modified to ensure the `lookup` command specified parameter `blacklist=false`.
- D. The `lookup` cannot be blacklisted; the change must be reverted.

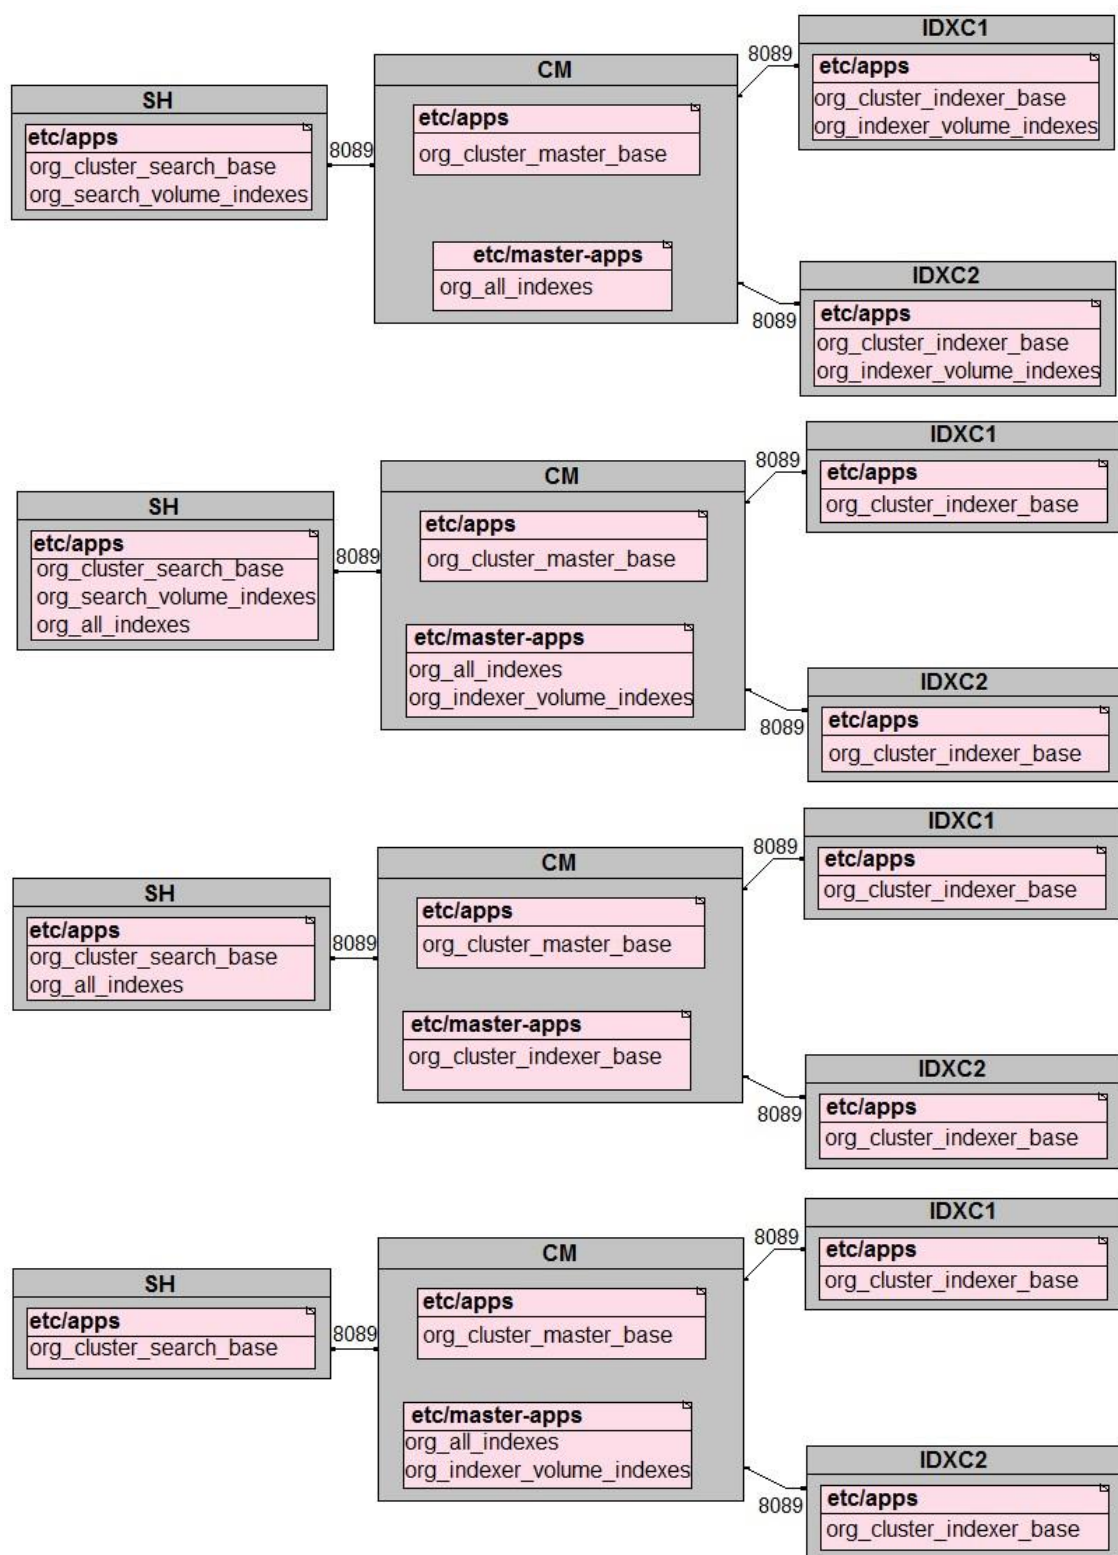
Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 58 In preparation for the deployment of a new environment for a customer, which of the following mappings are correct per PS best practices? A.



B. C.

D.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 59 Which of the following statements is true, as it pertains to search head clustering (SHC)?

- A. SHC is supported on AIX, Linux, and Windows operating systems.
- B. Maximum number of nodes for a SHC is 10.
- C. SHC members must run on the same hardware specifications.
- D. Minimum number of nodes for a SHC is 5.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



QUESTION 60 Where are Splunk Data Model Acceleration (DMA) summaries stored?

- A. In `tstatsHomePath`
- B. In the `.tsidx` files.
- C. In `summaryHomePath`
- D. In `journal.gz`

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Knowledge/Acceleratedatamodels#:~:text=Splunk%20software%20creates%20ad%20hoc,your%20indexes%20alongside%20index%20buckets>

QUESTION 61 When can the Search Job Inspector be used to debug searches?

- A. If the search has not expired.
- B. If the search is currently running.
- C. If the search has been queued.
- D. If the search has expired.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Search/ViewsearchjobpropertieswiththeJobInspector>

QUESTION 62

A Splunk Index cluster is being installed and the indexers need to be configured with a license master. After the customer provides the name of the license master, what is the next step?

- A. Enter the license master configuration via Splunk web on each indexer before disabling Splunk web.
- B. Update `/opt/splunk/etc/master-apps/_cluster/default/server.conf` on the cluster master and apply a cluster bundle.
- C. Update the Splunk PS base config license app and copy to each indexer.
- D. Update the Splunk PS base config license app and deploy via the cluster master.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 63

A customer has three users and is planning to ingest 250GB of data per day. They are concerned with search uptime, can tolerate up to a two-hour downtime for the search tier, and want advice on single search head versus a search head cluster. (SHC).

Which recommendation is the most appropriate?

- A. The customer should deploy two active search heads behind a load balancer to support HA.
- B. The customer should deploy a SHC with a single member for HA; more members can be added later.
- C. The customer should deploy a SHC, because it will be required to support the high volume of data.
- D. The customer should deploy a single search head with a warm standby search head and an rsync process to synchronize configurations.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 64 Which of the following is the most efficient search?

- A. `index=www status=200 uri=/cart/checkout | append [search index = sales] | stats count, sum(revenue) as total_revenue by session_id | table total_revenue session_id`
- B. `(index=www status=200 uri=/cart/checkout) OR (index=sales) | stats count, sum(revenue) as total_revenue by session_id | table total_revenue session_id`
- C. `index=www | append [search index = sales] | stats count, sum(revenue) as total_revenue by session_id | table total_revenue session_id`
- D. `(index=www) OR (index=sales) | search (index=www status=200 uri=/cart/checkout) OR (index=sales) | stats count, sum(revenue) as total_revenue by session_id | table total_revenue session_id`

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 65

Consider the search shown below.

```
index=web sourcetype=web_log [ search index=firewall action=denied
severity=high | stats latest (_time) as _time | eval
earliest=tostring(relative_time (_time, "-2h@h")), latest=tostring
(relative_time(_time, "+2h@h")) | fields earliest, latest]
```

What is this search's intended function?

- A. To return all the `web_log` events from the `web` index that occur two hours before and after the most recent high severity, denied event found in the `firewall` index.
- B. To find all the denied, high severity events in the `firewall` index, and use those events to further search for lateral movement within the `web` index.
- C. To return all the `web_log` events from the `web` index that occur two hours before and after all high severity, denied events found in the `firewall` index.
- D. To search the `firewall` index for web logs that have been denied and are of high severity.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 66

When setting up a multisite search head and indexer cluster, which nodes are required to declare site membership?

- A. Search head cluster members, deployer, indexers, cluster master
- B. Search head cluster members, deployment server, deployer, indexers, cluster master
- C. All splunk nodes, including forwarders, must declare site membership
- D. Search head cluster members, indexers, cluster master

Correct Answer: D

Section: (none)

Explanation



Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/SHCandindexercluster>

QUESTION 67 A customer is using both internal Splunk authentication and LDAP for user management.

If a username exists in both `$SPLUNK_HOME/etc/passwd` and LDAP, which of the following statements is accurate?

- A. The internal Splunk authentication will take precedence.
- B. Authentication will only succeed if the password is the same in both systems.
- C. The LDAP user account will take precedence.
- D. Splunk will error as it does not support overlapping usernames

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 68 When utilizing a subsearch within a Splunk SPL search query, which of the following statements is accurate?

- A. Subsearches have to be initiated with the `| subsearch` command.
- B. Subsearches can only be utilized with `| inputlookup` command.
- C. Subsearches have a default result output limit of 10000.

D. There are no specific limitations when using subsearches.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.0.6/Search/Aboutsubsearches#:~:text=By%20default%2C%20subsearches%20return%20a.will%20timeout%20before%20it%20completes>

QUESTION 69

A customer is migrating their existing Splunk Indexer from an old set of hardware to a new set of indexers. What is the earliest method to migrate the system?

- A. 1. Add new indexers to the cluster as peers, in the same site (if needed).
2. Ensure new indexers receive common configuration.
3. Decommission old indexers (one at a time) to allow time for CM to fix/migrate buckets to new hardware.
4. Remove all the old indexers from the CM's list.
- B. 1. Add new indexers to the cluster as peers, to a new site.
2. Ensure new indexers receive common configuration from the CM.
3. Decommission old indexers (one at a time) to allow time for CM to fix/migrate buckets to new hardware.
4. Remove all the old indexers from the CM's list.
- C. 1. Add new indexers to the cluster as peers, in the same site.
2. Update the replication factor by +1 to instruct the cluster to start replicating to new peers.
3. Allow time for CM to fix/migrate buckets to new hardware.
4. Remove all the old indexers from the CM's list.
- D. 1. Add new indexers to the cluster as new site.
2. Update cluster master (CM) `server.conf` to include the new available site.
3. Allow time for CM to fix/migrate buckets to new hardware.
4. Remove the old indexers from the CM's list.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



QUESTION 70 When using SAML, where does user authentication occur?

- A. Splunk generates a SAML assertion that authenticates the user.
- B. The Service Provider (SP) decodes the SAML request and authenticates the user.
- C. The Identity Provider (IDP) decodes the SAML request and authenticates the user.
- D. The Service Provider (SP) generates a SAML assertion that authenticates the user.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 71 Which of the following server roles should be configured for a host which indexes its internal logs locally?

- A. Cluster master
- B. Indexer
- C. Monitoring Console (MC)
- D. Search head

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://community.splunk.com/t5/Deployment-Architecture/How-to-identify-Splunk-Instance-role-by-internal-logs/m-p/365555>

QUESTION 72

The Splunk Validated Architectures (SVAs) document provides a series of approved Splunk topologies. Which statement accurately describes how it should be used by a customer?

- A. Customer should look at the category tables, pick the highest number that their budget permits, then select this design topology as the chosen design.
- B. Customers should identify their requirements, provisionally choose an approved design that meets them, then consider design principles and best practices to come to an informed design decision.
- C. Using the guided requirements gathering in the SVAs document, choose a topology that suits requirements, and be sure not to deviate from the specified design.
- D. Choose an SVA topology code that includes Search Head and Indexer Clustering because it offers the highest level of resilience.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: https://www.splunk.com/en_us/blog/tips-and-tricks/splunk-validated-architectures.html

QUESTION 73

In a large cloud customer environment with many (>100) dynamically created endpoint systems, each with a UF already deployed, what is the best approach for associating these systems with an appropriate serverclass on the deployment server?

- A. Work with the cloud orchestration team to create a common host-naming convention for these systems so a simple pattern can be used in the `serverclass.conf` `whitelist` attribute.
- B. Create a CSV lookup file for each severclass, manually keep track of the endpoints within this CSV file, and leverage the `whitelist.from_pathname` attribute in `serverclass.conf`.
- C. Work with the cloud orchestration team to dynamically insert an appropriate `clientName` setting into each endpoint's `local/deploymentclient.conf` which can be matched by `whitelist` in `serverclass.conf`.
- D. Using an installation bootstrap script run a CLI command to assign a `clientName` setting and permit `serverclass.conf` `whitelist` simplification.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 74 Which of the following is the most

efficient search? A.

```
index=foo sourcetype=bar | lookup local=t mylookup host OUTPUT  
host_flag | where host_flag= "true" | stats count by host
```

```
index=foo sourcetype=* | lookup mylookup host OUTPUT  
host_flag | where host_flag= "true" | stats count by host
```

```
index=foo sourcetype=bar | fields host | lookup mylookup host OUTPUT  
host_flag | where host_flag= "true" | stats count by host
```

```
index=foo sourcetype=bar | table host | lookup local=t mylookup host OUTPUT  
host_flag | where host_flag= "true" | stats count by host
```

- B.
- C.
- D.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 75

A customer has 30 indexers in an indexer cluster configuration and two search heads. They are working on writing SPL search for a particular use-case, but are concerned that it takes too long to run for short time durations.

How can the Search Job Inspector capabilities be used to help validate and understand the customer concerns?

- A. Search Job Inspector provides statistics to show how much time and the number of events each indexer has processed.
- B. Search Job Inspector provides a Search Health Check capability that provides an optimized SPL query the customer should try instead.
- C. Search Job Inspector cannot be used to help troubleshoot the slow performing search; customer should review `index=_introspection` instead.
- D. The customer is using the `transaction` SPL search command, which is known to be slow.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 76

A customer would like to remove the `output_file` capability from users with the default user role to stop them from filling up the disk on the search head with lookup files. What is the best way to remove this capability from users?

- A. Create a new role without the `output_file` capability that inherits the default user role and assign it to the users.
- B. Create a new role with the `output_file` capability that inherits the default user role and assign it to the users.
- C. Edit the default user role and remove the `output_file` capability.
- D. Clone the default user role, remove the `output_file` capability, and assign it to the users.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 77

A working search head cluster has been set up and used for 6 months with just the native/local Splunk user authentication method. In order to integrate the search heads with an external Active Directory server using LDAP, which of the following statements represents the most appropriate method to deploy the configuration to the servers?

- A. Configure the integration in a base configuration app located in `shcluster-apps` directory on the search head deployer, then deploy the configuration to the search heads using the `splunk apply shcluster-bundle` command.
- B. Log onto each search using a command line utility. Modify the `authentication.conf` and `authorize.conf` files in a base configuration app to configure the integration.
- C. Configure the LDAP integration on one Search Head using the `Settings > Access Controls > Authentication Method` and `Settings > Access Controls > Roles` Splunk UI menus. The configuration setting will replicate to the other nodes in the search head cluster eliminating the need to do this on the other search heads.
- D. On each search head, login and configure the LDAP integration using the `Settings > Access Controls > Authentication Method` and `Settings > Access Controls > Roles` Splunk UI menus.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.splunk.com/Documentation/Splunk/8.1.0/Security/ConfigureLDAPwithSplunkWeb>

QUESTION 78

In an environment that has Indexer Clustering, the Monitoring Console (MC) provides dashboards to monitor environment health. As the environment grows over time and new indexers are added, which steps would ensure the MC is aware of the additional indexers?

- A. No changes are necessary, the Monitoring Console has self-configuration capabilities.
- B. Using the MC setup UI, review and apply the changes.
- C. Remove and re-add the cluster master from the indexer clustering UI page to add new peers, then apply the changes under the MC setup UI.
- D. Each new indexer needs to be added using the distributed search UI, then settings must be saved under the MC setup UI.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 79

In addition to the normal responsibilities of a search head cluster captain, which of the following is a default behavior?

- A. The captain is not a cluster member and does not perform normal search activities.
- B. The captain is a cluster member who performs normal search activities.
- C. The captain is not a cluster member but does perform normal search activities.
- D. The captain is a cluster member but does not perform normal search activities.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/SHCarchitecture#Search_head_cluster_captain



QUESTION 80

What happens to the indexer cluster when the indexer Cluster Master (CM) runs out of disk space?

- A. A warm standby CM needs to be brought online as soon as possible before an indexer has an outage.
- B. The indexer cluster will continue to operate as long as no indexers fail.
- C. If the indexer cluster has site failover configured in the CM, the second cluster master will take over.
- D. The indexer cluster will continue to operate as long as a replacement CM is deployed within 24 hours.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 81

Which event processing pipeline contains the regex replacement processor that would be called upon to run event masking routines on events as they are ingested?

- A. Merging pipeline
- B. Indexing pipeline
- C. Typing pipeline
- D. Parsing pipeline

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 82 Which statement is correct?

- A. In general, `search` commands that can be distributed to the search peers should occur as early as possible in a well-tuned search.
- B. As a streaming command, `streamstats` performs better than `stats` since `stats` is just a reporting command.
- C. When trying to reduce a search result to unique elements, the `dedup` command is the only way to achieve this.
- D. Formatting commands such as `fieldformat` should occur as early as possible in the search to take full advantage of the often larger number of search peers.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 83

A non-ES customer has a concern about data availability during a disaster recovery event. Which of the following Splunk Validated Architectures (SVAs) would be recommended for that use case?

- A. Topology Category Code: M4
- B. Topology Category Code: M14
- C. Topology Category Code: C13
- D. Topology Category Code: C3

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf> (21)



QUESTION 84

The universal forwarder (UF) should be used whenever possible, as it is smaller and more efficient. In which of the following scenarios would a heavy forwarder (HF) be a more appropriate choice?

- A. When a predictable version of Python is required.
- B. When filtering 10%–15% of incoming events.
- C. When monitoring a log file.
- D. When running a script.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: https://www.splunk.com/en_us/blog/tips-and-tricks/universal-or-heavy-that-is-the-question.html

QUESTION 85

When monitoring and forwarding events collected from a file containing unstructured textual events, what is the difference in the Splunk2Splunk payload traffic sent between a universal forwarder (UF) and indexer compared to the Splunk2Splunk payload sent between a heavy forwarder (HF) and the indexer layer? (Assume that the file is being monitored locally on the forwarder.)

- A. The payload format sent from the UF versus the HF is exactly the same. The payload size is identical because they're both sending 64K chunks.
- B. The UF sends a stream of data containing one set of metadata fields to represent the entire stream, whereas the HF sends individual events, each with their own metadata fields attached, resulting in a larger payload.
- C. The UF will generally send the payload in the same format, but only when the sourcetype is specified in the `inputs.conf` and `EVENT_BREAKER_ENABLE` is set to `true`.
- D. The HF sends a stream of 64K TCP chunks with one set of metadata fields attached to represent the entire stream, whereas the UF sends individual events, each with their own metadata fields attached.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

