SPLK-3003 Q&As

Splunk Core Certified Consultant

Pass Splunk SPLK-3003 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/splk-3003.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Splunk Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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

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

QUESTION 2

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

QUESTION 3

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

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

QUESTION 4

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

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

QUESTION 5

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

QUESTION 6

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

QUESTION 7

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

QUESTION 8

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

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

QUESTION 9

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

QUESTION 10

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 = security
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

QUESTION 11

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

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

QUESTION 12

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

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

QUESTION 13

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

QUESTION 14

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

QUESTION 15

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

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