

Exam Questions SPLK-2002

Splunk Enterprise Certified Architect

<https://www.2passeasy.com/dumps/SPLK-2002/>



NEW QUESTION 1

Search dashboards in the Monitoring Console indicate that the distributed deployment is approaching its capacity. Which of the following options will provide the most search performance improvement?

- A. Replace the indexer storage to solid state drives (SSD).
- B. Add more search heads and redistribute users based on the search type.
- C. Look for slow searches and reschedule them to run during an off-peak time.
- D. Add more search peers and make sure forwarders distribute data evenly across all indexers.

Answer: C

NEW QUESTION 2

A Splunk architect has inherited the Splunk deployment at Buttercup Games and end users are complaining that the events are inconsistently formatted for a web sourcetype. Further investigation reveals that not all web logs flow through the same infrastructure: some of the data goes through heavy forwarders and some of the forwarders are managed by another department. Which of the following items might be the cause for this issue?

- A. The search head may have different configurations than the indexers.
- B. The data inputs are not properly configured across all the forwarders.
- C. The indexers may have different configurations than the heavy forwarders.
- D. The forwarders managed by the other department are an older version than the rest.

Answer: D

NEW QUESTION 3

A customer has installed a 500GB Enterprise license. They also purchased and installed a 300GB, no enforcement license on the same license master. How much data can the customer ingest before search is locked out?

- A. 300G
- B. After this limit, search is locked ou
- C. B.500G
- D. After this limit, search is locked out.
- E. 800G
- F. After this limit, search is locked out.
- G. Search is not locked ou
- H. Violations are still recorded.

Answer: D

NEW QUESTION 4

A multi-site indexer cluster can be configured using which of the following? (Select all that apply.)

- A. Via Splunk Web.
- B. Directly edit SPLUNK_HOME/etc/system/local/server.conf
- C. Run a splunk edit cluster-config command from the CLI.
- D. Directly edit SPLUNK_HOME/etc/system/default/server.conf

Answer: AB

NEW QUESTION 5

Which Splunk tool offers a health check for administrators to evaluate the health of their Splunk deployment?

- A. btool
- B. DiagGen
- C. SPL Clinic
- D. Monitoring Console

Answer: D

NEW QUESTION 6

Which Splunk Enterprise offering has its own license?

- A. Splunk Cloud Forwarder
- B. Splunk Heavy Forwarder
- C. Splunk Universal Forwarder
- D. Splunk Forwarder Management

Answer: C

NEW QUESTION 7

Which of the following clarification steps should be taken if apps are not appearing on a deployment client? (Select all that apply.)

- A. Check serverclass.conf of the deployment server.
- B. Check deploymentclient.conf of the deployment client.

- C. Check the content of SPLUNK_HOME/etc/apps of the deployment server.
- D. Search for relevant events in splunkd.log of the deployment server.

Answer: ABC

NEW QUESTION 8

Which of the following security options must be explicitly configured (i.e. which options are not enabled by default)?

- A. Data encryption between Splunk Web and splunkd.
- B. Certificate authentication between forwarders and indexers.
- C. Certificate authentication between Splunk Web and search head.
- D. Data encryption for distributed search between search heads and indexers.

Answer: B

NEW QUESTION 9

Splunk Enterprise platform instrumentation refers to data that the Splunk Enterprise deployment logs in the _introspection index. Which of the following logs are included in this index? (Select all that apply.)

- A. audit.log
- B. metrics.log
- C. disk_objects.log
- D. resource_usage.log

Answer: CD

NEW QUESTION 10

Which of the following are true statements about Splunk indexer clustering?

- A. All peer nodes must run exactly the same Splunk version.
- B. The master node must run the same or a later Splunk version than search heads.
- C. The peer nodes must run the same or a later Splunk version than the master node.
- D. The search head must run the same or a later Splunk version than the peer nodes.

Answer: B

NEW QUESTION 10

A customer plans to ingest 600 GB of data per day into Splunk. They will have six concurrent users, and they also want high data availability and high search performance. The customer is concerned about cost and wants to spend the minimum amount on the hardware for Splunk. How many indexers are recommended for this deployment?

- A. Two indexers not in a cluster, assuming users run many long searches.
- B. Three indexers not in a cluster, assuming a long data retention period.
- C. Two indexers clustered, assuming high availability is the greatest priority.
- D. Two indexers clustered, assuming a high volume of saved/scheduled searches.

Answer: D

NEW QUESTION 15

At which default interval does metrics.log generate a periodic report regarding license utilization?

- A. 10 seconds
- B. 30 seconds
- C. 60 seconds
- D. 300 seconds

Answer: B

NEW QUESTION 17

Which of the following is a good practice for a search head cluster deployer?

- A. The deployer only distributes configurations to search head cluster members when they “phone home”.
- B. The deployer must be used to distribute non-replicable configurations to search head cluster members.
- C. The deployer must distribute configurations to search head cluster members to be valid configurations.
- D. The deployer only distributes configurations to search head cluster members with splunk apply shcluster-bundle.

Answer: A

NEW QUESTION 19

A new Splunk customer is using syslog to collect data from their network devices on port 514. What is the best practice for ingesting this data into Splunk?

- A. Configure syslog to send the data to multiple Splunk indexers.
- B. Use a Splunk indexer to collect a network input on port 514 directly.
- C. Use a Splunk forwarder to collect the input on port 514 and forward the data.
- D. Configure syslog to write logs and use a Splunk forwarder to collect the logs.

Answer: C

NEW QUESTION 23

How does IT Service Intelligence (ITSI) impact the planning of a Splunk deployment?

- A. ITSI requires a dedicated deployment server.
- B. The amount of users using ITSI will not impact performance.
- C. ITSI in a Splunk deployment does not require additional hardware resources.
- D. Depending on the Key Performance Indicators that are being tracked, additional infrastructure may be needed.

Answer: D

NEW QUESTION 25

In the deployment planning process, when should a person identify who gets to see network data?

- A. Deployment schedule
- B. Topology diagramming
- C. Data source inventory
- D. Data policy definition

Answer: C

NEW QUESTION 27

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

- A. Use the Monitoring Console.
- B. Use the Search Head Clustering settings menu from Splunk Web on any member.
- C. Run the splunk transfer shcluster-captain command from the current captain.
- D. Run the splunk transfer shcluster-captain command from the member you would like to become the captain.

Answer: BD

NEW QUESTION 31

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

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

Answer: D

NEW QUESTION 32

What does setting site=site0 on all Search Head Cluster members do in a multi-site indexer cluster?

- A. Disables search site affinity.
- B. Sets all members to dynamic captaincy.
- C. Enables multisite search artifact replication.
- D. Enables automatic search site affinity discovery.

Answer: A

NEW QUESTION 34

A search head has successfully joined a single site indexer cluster. Which command is used to configure the same search head to join another indexer cluster?

- A. splunk add cluster-config
- B. splunk add cluster-master
- C. splunk edit cluster-config
- D. splunk edit cluster-master

Answer: B

NEW QUESTION 35

When adding or decommissioning a member from a Search Head Cluster (SHC), what is the proper order of operations?

- A. 1. Delete Splunk Enterprise, if it exists.2. Install and initialize the instance.3. Join the SHC.
- B. 1. Install and initialize the instance.2. Delete Splunk Enterprise, if it exists.3. Join the SHC.
- C. 1. Initialize cluster rebalance operation.2. Remove master node from cluster.3. Trigger replication.
- D. 1. Trigger replication.2. Remove master node from cluster.3. Initialize cluster rebalance operation.

Answer: B

NEW QUESTION 37

When troubleshooting monitor inputs, which command checks the status of the tailed files?

- A. splunk cmd btool inputs list | tail
- B. splunk cmd btool check inputs layer
- C. curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:FileStatus
- D. curl https://serverhost:8089/services/admin/inputstatus/TailingProcessor:Tailstatus

Answer: C

NEW QUESTION 42

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

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

Answer: B

NEW QUESTION 44

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

- A. System local directory.
- B. System default directory.
- C. App local directories, in ASCII order.
- D. App default directories, in ASCII order.

Answer: A

NEW QUESTION 49

What is the algorithm used to determine captaincy in a Splunk search head cluster?

- A. Raft distributed consensus.
- B. Rapt distributed consensus.
- C. Rift distributed consensus.
- D. Round-robin distribution consensus.

Answer: A

NEW QUESTION 50

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

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

Answer: C

NEW QUESTION 51

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

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

Answer: AC

NEW QUESTION 54

Which of the following options can improve reliability of syslog delivery to Splunk? (Select all that apply.)

- A. Use TCP syslog.
- B. Configure UDP inputs on each Splunk indexer to receive data directly.
- C. Use a network load balancer to direct syslog traffic to active backend syslog listeners.
- D. Use one or more syslog servers to persist data with a Universal Forwarder to send the data to Splunk indexers.

Answer: CD

NEW QUESTION 58

Which of the following statements describe search head clustering? (Select all that apply.)

- A. A deployer is required.
- B. At least three search heads are needed.
- C. Search heads must meet the high-performance reference server requirements.

D. The deployer must have sufficient CPU and network resources to process service requests and push configurations.

Answer: AC

NEW QUESTION 59

Because Splunk indexing is read/write intensive, it is important to select the appropriate disk storage solution for each deployment. Which of the following statements is accurate about disk storage?

- A. High performance SAN should never be used.
- B. Enable NFS for storing hot and warm buckets.
- C. The recommended RAID setup is RAID 10 (1 + 0).
- D. Virtualized environments are usually preferred over bare metal for Splunk indexers.

Answer: C

NEW QUESTION 61

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SPLK-2002 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SPLK-2002 Product From:

<https://www.2passeasy.com/dumps/SPLK-2002/>

Money Back Guarantee

SPLK-2002 Practice Exam Features:

- * SPLK-2002 Questions and Answers Updated Frequently
- * SPLK-2002 Practice Questions Verified by Expert Senior Certified Staff
- * SPLK-2002 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SPLK-2002 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year