[2024 Use Valid SPLK-2002 Exam - Actual Exam Question & Answer [Q20-Q36

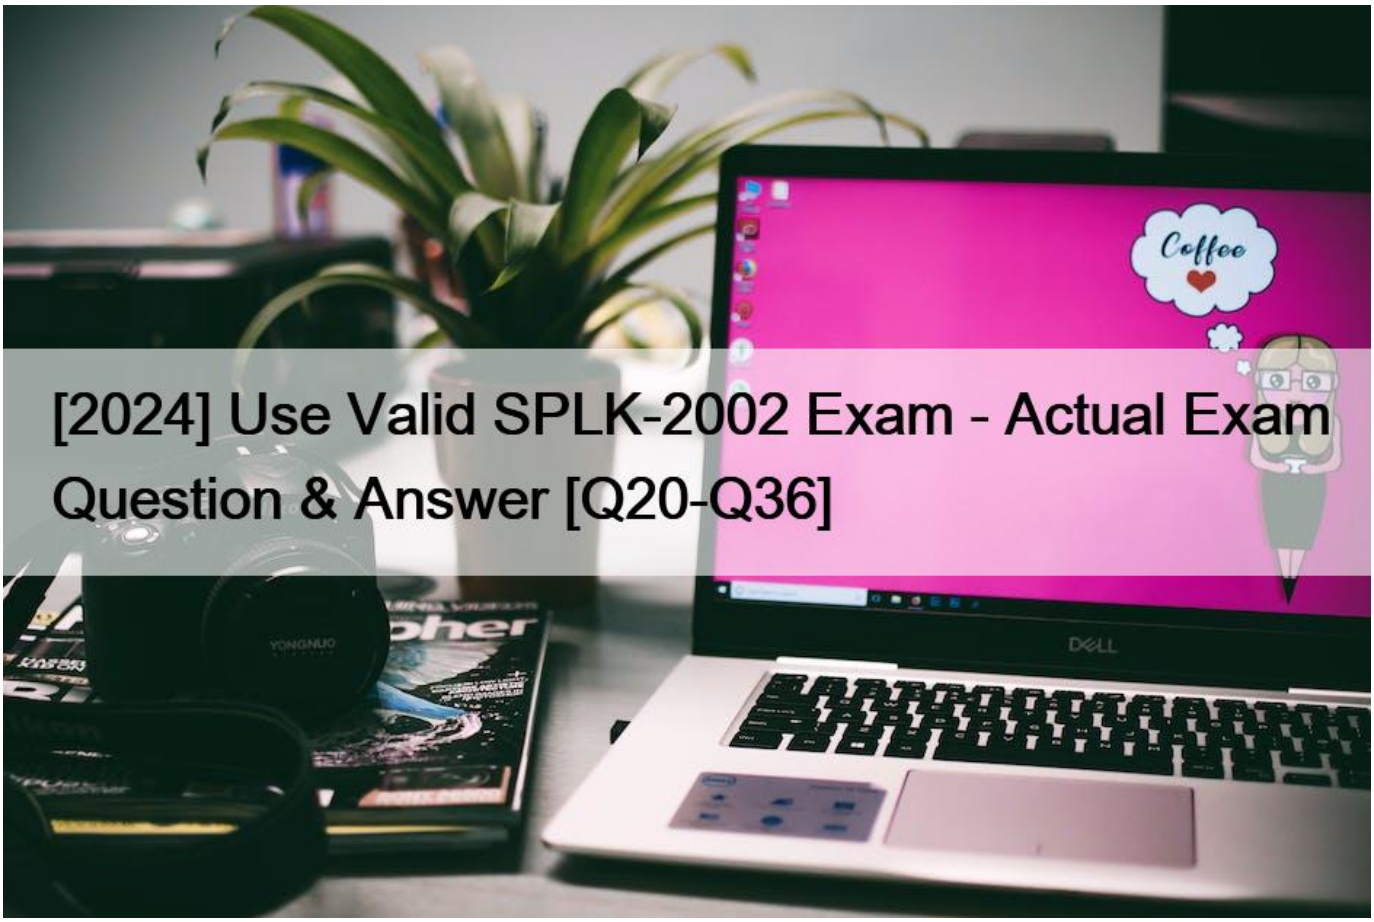

[2024] Use Valid SPLK-2002 Exam - Actual Exam Question & Answer
Test Engine to Practice SPLK-2002 Test Questions

Earning the Splunk SPLK-2002 certification is a significant achievement and can open up new career opportunities in the field of data analytics and cybersecurity. Certified architects are in high demand by organizations seeking to leverage the power of Splunk to gain insights into their data and improve their security posture. With the SPLK-2002 certification, professionals can demonstrate their expertise in designing and managing complex Splunk environments, making them valuable assets to any organization that relies on this powerful platform.

The SPLK-2002 certification exam covers a broad range of topics, including Splunk deployment, data management, security, performance optimization, and troubleshooting. SPLK-2002 exam consists of 100 multiple-choice questions that test the candidate's ability to design, deploy, and manage Splunk Enterprise environments. Additionally, the exam evaluates the candidate's understanding of Splunk best practices, troubleshooting techniques, and deployment methodologies. Passing the SPLK-2002 exam indicates that the candidate has a comprehensive understanding of Splunk Enterprise architecture and is capable of deploying and managing Splunk Enterprise environments to meet the needs of their organization.

Q20. Stakeholders have identified high availability for searchable data as their top priority. Which of the following best addresses this requirement?
* Increasing the search factor in the cluster.
* Increasing the replication factor in the cluster.
* Increasing the number of search heads in the cluster.
* Increasing the number of CPUs on the indexers in the cluster.

Explanation/Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture

Q21. Which of the following tasks should the architect perform when building a deployment plan? (Select all that apply.)
* Use case checklist.
* Install Splunk apps.
* Inventory data sources.
* Review network topology.

Explanation

When building a deployment plan, the architect should perform the following tasks:

* Use case checklist. A use case checklist is a document that lists the use cases that the deployment will support, along with the data sources, the data volume, the data retention, the data model, the dashboards, the reports, the alerts, and the roles and permissions for each use case. A use case checklist helps to define the scope and the functionality of the deployment, and to identify the dependencies and the requirements for each use case1

* Inventory data sources. An inventory of data sources is a document that lists the data sources that the deployment will ingest, along with the data type, the data format, the data location, the data collection method, the data volume, the data frequency, and the data owner for each data source. An inventory of data sources helps to determine the data ingestion strategy, the data parsing and enrichment, the data storage and retention, and the data security and compliance for the deployment1

* Review network topology. A review of network topology is a process that examines the network infrastructure and the network connectivity of the deployment, along with the network bandwidth, the network latency, the network security, and the network monitoring for the deployment. A review of network topology helps to optimize the network performance and reliability, and to identify the network risks and mitigations for the deployment1 Installing Splunk apps is not a task that the architect should perform when building a deployment plan, as it is a task that the administrator should perform when implementing the deployment plan. Installing Splunk apps is a technical activity that requires access to the Splunk instances and the Splunk configurations, which are not available at the planning stage

Q22. When Splunk indexes data in a non-clustered environment, what kind of files does it create by default?
* Index and .tsidx files.
* Rawdata and index files.
* Compressed and .tsidx files.
* Compressed and meta data files.

Explanation

When Splunk indexes data in a non-clustered environment, it creates index and .tsidx files by default. The index files contain the raw data that Splunk has ingested, compressed and encrypted. The .tsidx files contain the time-series index that maps the timestamps and event IDs of the raw data. The rawdata and index files are not the correct terms for the files that Splunk creates. The compressed and .tsidx files are partially correct, but compressed is not the proper name for the index files. The compressed and meta data files are also partially correct, but meta data is not the proper name for the .tsidx files.

Q23. Which of the following is a valid use case that a search head cluster addresses?

* Provide redundancy in the event a search peer fails.
* Search affinity.
* Knowledge Object replication.
* Increased Search Factor (SF).

The correct answer is C. Knowledge Object replication. This is a valid use case that a search head cluster addresses, as it ensures that all the search heads in the cluster have the same set of knowledge objects, such as saved searches, dashboards, reports, and alerts1. The search head cluster replicates the knowledge objects across the cluster members, and synchronizes any changes or updates1. This provides a consistent user experience and avoids data inconsistency or duplication1. The other options are not valid use cases that a search head cluster addresses. Option A, providing redundancy in the event a search peer fails, is not a use case for a search head cluster, but for an indexer cluster, which maintains multiple copies of the indexed data and can recover from indexer failures2. Option B, search affinity, is not a use case for a search head cluster, but for a multisite indexer cluster, which allows the search heads to preferentially search the data on the local site, rather than on a remote site3. Option D, increased Search Factor (SF), is not a use case for a search head cluster, but for an indexer cluster, which determines how many searchable copies of each bucket are maintained across the indexers4. Therefore, option C is the correct answer, and options A, B, and D are incorrect.

1: About search head clusters 2: About indexer clusters and index replication 3: Configure search affinity 4:

Configure the search factor

Q24. Which of the following statements describe search head clustering? (Select all that apply.)
* A deployer is required.
* At least three search heads are needed.
* Search heads must meet the high-performance reference server requirements.
* The deployer must have sufficient CPU and network resources to process service requests and push configurations.

Q25. The frequency in which a deployment client contacts the deployment server is controlled by what?
* polling_interval attribute in outputs.conf
* phoneHomeIntervalInSecs attribute in outputs.conf
* polling_interval attribute in deploymentclient.conf
* phoneHomeIntervalInSecs attribute in deploymentclient.conf

Explanation

The frequency in which a deployment client contacts the deployment server is controlled by the phoneHomeIntervalInSecs attribute in deploymentclient.conf. This attribute specifies how often the deployment client checks in with the deployment server to get updates on the apps and configurations that it should receive. The polling_interval attribute in outputs.conf controls how often the forwarder sends data to the indexer or another forwarder. The polling_interval attribute in deploymentclient.conf and the phoneHomeIntervalInSecs attribute in outputs.conf are not valid Splunk attributes. For more information, see Configure deployment clients and Configure forwarders with outputs.conf in the Splunk documentation.

Q26. Which of the following is true for indexer cluster knowledge bundles?
* Only app-name/local is pushed.
* app-name/default and app-name/local are merged before pushing.
* Only app-name/default is pushed.
* app-name/default and app-name/local are pushed without change.

According to the Splunk documentation1, indexer cluster knowledge bundles are the configuration files that the cluster master distributes to the peer nodes as part of the cluster configuration bundle. The knowledge bundles contain the knowledge objects, such as event types, tags, lookups, and so on, that are relevant for indexing and searching the data. The cluster master creates the knowledge bundles by merging the app-name/default and app-name/local directories from the apps that reside on the master node. The cluster master then pushes the knowledge bundles to the peer nodes, where they reside under the

$SPLUNK_HOME/var/run directory2. The other options are false because:

* Only app-name/local is pushed. This is false because the cluster master pushes both the app-name/default and app-name/local directories, after merging them, to the peer nodes. The app-name/local directory contains the local customizations of the app configuration, while the app-name/default directory contains the default app configuration3.

* Only app-name/default is pushed. This is false because the cluster master pushes both the app-name/default and app-name/local directories, after merging them, to the peer nodes. The app-name/default directory contains the default app configuration, while the app-name/local directory contains the local customizations of the app configuration3.

* app-name/default and app-name/local are pushed without change. This is false because the cluster

* master merges the app-name/default and app-name/local directories before pushing them to the peer nodes. This ensures that the peer nodes have the latest and consistent configuration of the apps3.

Q27. Which index-time props.conf attributes impact indexing performance? (Select all that apply.)
* REPORT
* LINE_BREAKER
* ANNOTATE_PUNCT
* SHOULD_LINEMERGE
The index-time props.conf attributes that impact indexing performance are LINE_BREAKER and SHOULD_LINEMERGE. These attributes determine how Splunk breaks the incoming data into events and whether it merges multiple events into one. These operations can affect the indexing speed and the disk space consumption. The REPORT attribute does not impact indexing performance, as it is used to apply transforms at search time. The ANNOTATE_PUNCT attribute does not impact indexing performance, as it is used to add punctuation metadata to events at search time. For more information, see [About props.conf and transforms.conf] in the Splunk documentation.

Q28. Which tool(s) can be leveraged to diagnose connection problems between an indexer and forwarder? (Select all that apply.)
* telnet
* tcpdump
* splunk btool
* splunk btprobe

Q29. Which of the following describe migration from single-site to multisite index replication?
* A master node is required at each site.
* Multisite policies apply to new data only.
* Single-site buckets instantly receive the multisite policies.
* Multisite total values should not exceed any single-site factors.
Explanation/Reference: https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Migratetomultisite

Q30. Which command is used for thawing the archive bucket?
* Splunk collect
* Splunk convert
* Splunk rebuild
* Splunk dbinspect

Q31. In a four site indexer cluster, which configuration stores two searchable copies at the origin site, one searchable copy at site2, and a total of four searchable copies?
* site_search_factor = origin:2, site1:2, total:4
* site_search_factor = origin:2, site2:1, total:4

* site_replication_factor = origin:2, site1:2, total:4
* site_replication_factor = origin:2, site2:1, total:4

Explanation

In a four site indexer cluster, the configuration that stores two searchable copies at the origin site, one searchable copy at site2, and a total of four searchable copies is site_search_factor = origin:2, site2:1, total:4.

This configuration tells the cluster to maintain two copies of searchable data at the site where the data originates, one copy of searchable data at site2, and a total of four copies of searchable data across all sites.

The site_search_factor determines how many copies of searchable data are maintained by the cluster for each site. The site_replication_factor determines how many copies of raw data are maintained by the cluster for each site. For more information, see Configure multisite indexer clusters with server.conf in the Splunk documentation.

Q32. Which of the following are true statements about Splunk indexer clustering?

* All peer nodes must run exactly the same Splunk version.
* The master node must run the same or a later Splunk version than search heads.
* The peer nodes must run the same or a later Splunk version than the master node.
* The search head must run the same or a later Splunk version than the peer nodes.

Explanation

The following statements are true about Splunk indexer clustering:

* All peer nodes must run exactly the same Splunk version. This is a requirement for indexer clustering, as different Splunk versions may have different data formats or features that are incompatible with each other. All peer nodes must run the same Splunk version as the master node and the search heads that connect to the cluster.

* The search head must run the same or a later Splunk version than the peer nodes. This is a recommendation for indexer clustering, as a newer Splunk version may have new features or bug fixes that improve the search functionality or performance. The search head should not run an older Splunk version than the peer nodes, as this may cause search errors or failures. The following statements are false about Splunk indexer clustering:

* The master node must run the same or a later Splunk version than the search heads. This is not a requirement or a recommendation for indexer clustering, as the master node does not participate in the search process. The master node should run the same Splunk version as the peer nodes, as this ensures the cluster compatibility and functionality.

* The peer nodes must run the same or a later Splunk version than the master node. This is not a requirement or a recommendation for indexer clustering, as the peer nodes do not coordinate the cluster activities. The peer nodes should run the same Splunk version as the master node, as this ensures the cluster compatibility and functionality. For more information, see [About indexer clusters and index replication] and [Upgrade an indexer cluster] in the Splunk documentation.

Q33. Which of the following is a best practice to maximize indexing performance?

* Use automatic source typing.
* Use the Splunk default settings.
* Not use pre-trained source types.
* Minimize configuration generality.

A best practice to maximize indexing performance is to minimize configuration generality. Configuration generality refers to the use of generic or default settings for data inputs, such as source type, host, index, and timestamp. Minimizing configuration generality means using specific and accurate settings for each data input, which can reduce the processing overhead and improve the indexing throughput. Using automatic source typing, using the Splunk default settings, and not using pre-trained source types are examples of

configuration generality, which can negatively affect the indexing performance

Q34. A three-node search head cluster is skipping a large number of searches across time. What should be done to increase scheduled search capacity on the search head cluster?
* Create a job server on the cluster.
* Add another search head to the cluster.
* server.conf captain_is_adhoc_searchhead = true.
* Change limits.conf value for max_searches_per_cpu to a higher value.

Q35. What log file would you search to verify if you suspect there is a problem interpreting a regular expression in a monitor stanza?
* btool.log
* metrics.log
* splunkd.log
* tailing_processor.log

The tailing_processor.log file would be the best place to search if you suspect there is a problem interpreting a regular expression in a monitor stanza. This log file contains information about how Splunk monitors files and directories, including any errors or warnings related to parsing the monitor stanza. The splunkd.log file contains general information about the Splunk daemon, but it may not have the specific details about the monitor stanza. The btool.log file contains information about the configuration files, but it does not log the runtime behavior of the monitor stanza. The metrics.log file contains information about the performance metrics of Splunk, but it does not log the event breaking issues. For more information, see About Splunk Enterprise logging in the Splunk documentation.

Q36. Stakeholders have identified high availability for searchable data as their top priority. Which of the following best addresses this requirement?
* Increasing the search factor in the cluster.
* Increasing the replication factor in the cluster.
* Increasing the number of search heads in the cluster.
* Increasing the number of CPUs on the indexers in the cluster.

Explanation

https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture