

Splunk.SPLK-1003.v2026-09-08.q91

Exam Code:	SPLK-1003
Exam Name:	Splunk Enterprise Certified Admin
Certification Provider:	Splunk
Free Question Number:	91
Version:	v2026-09-08
# of views:	111
# of Questions views:	1107
https://www.freecram.net/torrent/Splunk.SPLK-1003.v2026-09-08.q91.html	

NEW QUESTION: 1

Which of the following are supported options when configuring optional network inputs?

- A. Metadata override, sender filtering options, network input queues (quantum queues)
- B. Metadata override, sender filtering options, network input queues (memory/persistent queues)
- C. Filename override, sender filtering options, network output queues (memory/persistent queues)
- D. Metadata override, receiver filtering options, network input queues (memory/persistent queues)

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/latest/Data/Monitornetworkports>

NEW QUESTION: 2

There is a file with a vast amount of old data. Which of the following inputs.conf attributes would allow an admin to monitor the file for updates without indexing the pre-existing data?

- A. IgnoreOlderThan
- B. allowList
- C. monitor
- D. followTail

Answer: (SHOW ANSWER)

* IgnoreOlderThan: This setting filters files for indexing based on their age. It does not prevent indexing of old data already in the file.

* allowList: This setting allows specifying patterns to include files for monitoring, but it does not control indexing of pre-existing data.

* monitor: This is the default method for monitoring files but does not address indexing pre-existing data.

* followTail: This attribute, when set in inputs.conf, ensures that Splunk starts reading a file from the end (tail) and does not index existing old data. It is ideal for scenarios with large files where only new updates are relevant.

References:

- * Splunk Docs: Monitor text files
- * Splunk Docs: Configure followTail in inputs.conf

NEW QUESTION: 3

After how many warnings within a rolling 30-day period will a license violation occur with an enforced Enterprise license?

- A. 1
- B. 3
- C. 4
- D. 5

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Admin/Aboutlicenseviolations>

"Enterprise Trial license. If you get five or more warnings in a rolling 30 days period, you are in violation of your license.

Dev/Test license. If you generate five or more warnings in a rolling 30-day period, you are in violation of your license.

Developer license. If you generate five or more warnings in a rolling 30-day period, you are in violation of your license. BUT for Free license. If you get three or more warnings in a rolling 30 days period, you are in violation of your license."

Reference: <https://docs.splunk.com/Documentation/Splunk/8.0.5/Admin/Aboutlicenseviolations>

NEW QUESTION: 4

What event-processing pipelines are used to process data for indexing? (select all that apply)

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

Answer: ([SHOW ANSWER](#))

The indexing pipeline and the parsing pipeline are the two pipelines that are responsible for transforming the raw data into events and preparing them for indexing. The indexing pipeline applies index-time settings, such as timestamp extraction, line breaking, host extraction, and source type recognition. The parsing pipeline applies parsing settings, such as field extraction, event segmentation, and event annotation.

NEW QUESTION: 5

The following stanza is active in indexes.conf:

```
[cat_facts]
```

```
maxHotSpanSecs = 3600
```

```
frozenTimePeriodInSecs = 2630000
```

```
maxTotalDataSizeMB = 650000
```

All other related indexes.conf settings are default values.

If the event timestamp was 3739283 seconds ago, will it be searchable?

- A. Yes, only if the bucket is still hot.
- B. No, because the index will have exceeded its maximum size.
- C. Yes, only if the index size is also below 650000 MB.
- D. No, because the event time is greater than the retention time.

Answer: ([SHOW ANSWER](#))

The correct answer is D. No, because the event time is greater than the retention time.

According to the Splunk documentation¹, the `frozenTimePeriodInSecs` setting in `indexes.conf` determines how long Splunk software retains indexed data before deleting it or archiving it to a remote storage. The default value is 188697600 seconds, which is equivalent to six years. The setting can be overridden on a per-index basis.

In this case, the `cat_facts` index has a `frozenTimePeriodInSecs` setting of 2630000 seconds, which is equivalent to about 30 days. This means that any event that is older than 30 days from the current time will be removed from the index and will not be searchable.

The event timestamp was 3739283 seconds ago, which is equivalent to about 43 days. This means that the event is older than the retention time of the `cat_facts` index and will not be searchable.

The other settings in the stanza, such as `maxHotSpanSecs` and `maxTotalDataSizeMB`, do not affect the retention time of the events. They only affect the size and duration of the buckets that store the events.

References: ¹Set a retirement and archiving policy - Splunk Documentation

NEW QUESTION: 6

Which feature in Splunk allows Event Breaking, Timestamp extractions, and any advanced configurations found in `props.conf` to be validated all through the UI?

- A. Apps
- B. Search
- C. Data preview
- D. Forwarder inputs

Answer: ([SHOW ANSWER](#))

<http://www.splunk.com/view/SP-CAAAGPR>

NEW QUESTION: 7

In this source definition the `MAX_TIMESTAMP_LOOKHEAD` is missing. Which value would fit best?

```
[sshd_syslog]
TIME_PREFIX = ^
TIME_FORMAT = %Y-%m-%d %H:%M:%S.%3N %z
TIME_BREAKER = > \d{4}-\d{2}-\d{2} \d{2}:\d{2}:\d{2}
SHOULD_LINEMERGE = false
TRUNCATE = 0
```

Event example:

- A. `MAX_TIMESTAMP_LOOKHEAD = 5`
- B. `MAX_TIMESTAMP_LOOKHEAD = 10`
- C. `MAX_TIMESTAMP_LOOKHEAD = 20`
- D. `MAX_TIMESTAMP_LOOKHEAD = 30`

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/6.2.0/Data/Configuretimestamprecognition>

"Specify how far (how many characters) into an event Splunk software should look for a timestamp." since `TIME_PREFIX = ^` and timestamp is from 0-29 position, so `D=30` will pick up the WHOLE timestamp correctly.

NEW QUESTION: 8

Which Splunk component(s) would break a stream of syslog inputs into individual events? (select all that apply)

A. Universal Forwarder

B. Search head

C. Heavy Forwarder

D. Indexer

Answer: ([SHOW ANSWER](#))

The correct answer is C and D. A heavy forwarder and an indexer are the Splunk components that can break a stream of syslog inputs into individual events.

A universal forwarder is a lightweight agent that can forward data to a Splunk deployment, but it does not perform any parsing or indexing on the data. A search head is a Splunk component that handles search requests and distributes them to indexers, but it does not process incoming data.

A heavy forwarder is a Splunk component that can perform parsing, filtering, routing, and aggregation on the data before forwarding it to indexers or other destinations. A heavy forwarder can break a stream of syslog inputs into individual events based on the line breaker and should `linemerge` settings in the `inputs.conf` file¹.

An indexer is a Splunk component that stores and indexes data, making it searchable. An indexer can also break a stream of syslog inputs into individual events based on the `props.conf` file settings, such as `TIME_FORMAT`, `MAX_TIMESTAMP_LOOKAHEAD`, and `line_breaker`².

A Splunk component is a software process that performs a specific function in a Splunk deployment, such as data collection, data processing, data storage, data search, or data visualization.

Syslog is a standard protocol for logging messages from network devices, such as routers, switches, firewalls, or servers.

Syslog messages are typically sent over UDP or TCP to a central syslog server or a Splunk instance.

Breaking a stream of syslog inputs into individual events means separating the data into discrete records that can be indexed and searched by Splunk. Each event should have a timestamp, a host, a source, and a sourcetype, which are the default fields that Splunk assigns to the data.

References:

1: Configure inputs using Splunk Connect for Syslog - Splunk Documentation

2: `inputs.conf` - Splunk Documentation

3: How to configure `props.conf` for proper line breaking ... - Splunk Community

4: Reliable syslog/tcp input - splunk bundle style | Splunk

5: Configure inputs using Splunk Connect for Syslog - Splunk Documentation

6: About configuration files - Splunk Documentation

[7]: Configure your OSSEC server to send data to the Splunk Add-on for OSSEC - Splunk Documentation

[8]: Splunk components - Splunk Documentation

[9]: Syslog - Wikipedia

[10]: About default fields - Splunk Documentation

NEW QUESTION: 9

Which of the following applies only to Splunk index data integrity check?

A. Raw data in the index

B. Lookup table

C. Summary Index

D. Data model acceleration

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 10

Where are deployment server apps mapped to clients?

- A. Apps tab in forwarder management interface or clientapps.conf.
- B. Clients tab in forwarder management interface or deploymentclient.conf.
- C. Server Classes tab in forwarder management interface or serverclass.conf.
- D. Client Applications tab in forwarder management interface or clientapps.conf.

Answer: ([SHOW ANSWER](#))

Reference: [https://docs.splunk.com/Documentation/Splunk/8.0.5/Updating/](https://docs.splunk.com/Documentation/Splunk/8.0.5/Updating/Updateconfigurations#2._Reload_the_deployment_server)

[Updateconfigurations#2._Reload_the_deployment_server](https://docs.splunk.com/Documentation/Splunk/8.0.5/Updating/Updateconfigurations#2._Reload_the_deployment_server)

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Updating/Useserverclass.conf>

" Use serverclass.conf to define server classes " " The most important settings define the set of deployment clients and the set of apps for each server class. "

NEW QUESTION: 11

In a distributed environment, which Splunk component is used to distribute apps and configurations to the other Splunk instances?

- A. Indexer
- B. Deployer
- C. Forwarder
- D. Deployment server

Answer: D ([LEAVE A REPLY](#))

The deployer is a Splunk Enterprise instance that you use to distribute apps and certain other configuration updates to search head cluster members. The set of updates that the deployer distributes is called the configuration

bundle.<https://docs.splunk.com/Documentation/Splunk/8.1.3/DistSearch>

[/PropagateSHCconfigurationchanges#:~:text=The%20deployer%20is%20a%20Splunk,is%20called%20the%20configuration%20bundle.](https://docs.splunk.com/Documentation/Splunk/8.1.3/DistSearch/PropagateSHCconfigurationchanges#:~:text=The%20deployer%20is%20a%20Splunk,is%20called%20the%20configuration%20bundle.)

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Updating/Updateconfigurations> First line says it all: "The deployment server distributes deployment apps to clients." Reference:

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Updating/Updateconfigurations>

NEW QUESTION: 12

When enabling data integrity control, where does Splunk Enterprise store the hash files for each bucket?

- A. Splunk Enterprise stores hash files in the logdata directory of the corresponding bucket.
- B. Splunk Enterprise stores hash files in the rawdata directory of the corresponding bucket.
- C. Splunk Enterprise stores hash files in the hashdata directory of the corresponding bucket.
- D. Splunk Enterprise stores hash files in the metadata directory of the corresponding bucket.

Answer: ([SHOW ANSWER](#))

* Data integrity controls in Splunk ensure that indexed data has not been tampered with.

* When enabled, Splunk calculates hashes for each bucket and stores these hash files in the rawdata directory of the corresponding bucket.

* Incorrect Options:

* A, C, D: These directories do not store hash files.

References:

Splunk Docs: Configure data integrity controls

NEW QUESTION: 13

Search heads in a company's European offices need to be able to search data in their New York offices. They also need to restrict access to certain indexers. What should be configured to allow this type of action?

A. Indexer clustering

B. LDAP control

C. Distributed search

D. Search head clustering

Answer: (SHOW ANSWER)

The correct answer is C. Distributed search is the feature that allows search heads in a company's European offices to search data in their New York offices. Distributed search also enables restricting access to certain indexers by using the `splunk_server` field or the `server.conf` file¹.

Distributed search is a way to scale your Splunk deployment by separating the search management and presentation layer from the indexing and search retrieval layer. With distributed search, a Splunk instance called a search head sends search requests to a group of indexers, or search peers, which perform the actual searches on their indexes. The search head then merges the results back to the user².

Distributed search has several use cases, such as horizontal scaling, access control, and managing geo-dispersed data. For example, users in different offices can search data across the enterprise or only in their local area, depending on their needs and permissions².

The other options are incorrect because:

A) Indexer clustering is a feature that replicates data across a group of indexers to ensure data availability and recovery. Indexer clustering does not directly affect distributed search, although search heads can be configured to search across an indexer cluster³.

B) LDAP control is a feature that allows Splunk to integrate with an external LDAP directory service for user authentication and role mapping. LDAP control does not affect distributed search, although it can be used to manage user access to data and searches.

D) Search head clustering is a feature that distributes the search workload across a group of search heads that share resources, configurations, and jobs. Search head clustering does not affect distributed search, although the search heads in a cluster can search across the same set of indexers.

NEW QUESTION: 14

In this example, if `useACK` is set to true and the `maxQueueSize` is set to 7MB, what is the size of the wait queue on this universal forwarder?

A. 21MB

B. 28MB

C. 14MB

D. 7MB

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/latest/Forwarding/Protectagainstlossofin-flightdata#:~:text=The%20default%20for%20the%20maxQueueSize,wait%20queue%20size%20is%2021MB.>

<https://docs.splunk.com/Documentation/Splunk/latest/Forwarding/Protectagainstlossofin-flightdata>

NEW QUESTION: 15

A request has been made to restrict lookup files up to 500 megabytes for replication. Anything larger should not be replicated. Which of the following parameters provides the correct control for this scenario?

A. maxBundleSize

B. maxMemoryBundleSize

C. excludeReplicatedLookupSize

D. includeReplicatedLookupSize

Answer: ([SHOW ANSWER](#))

In Splunk Enterprise, when knowledge bundles (which include lookup files, configurations, and other knowledge objects) are replicated between search heads and indexers, administrators can control the maximum size of lookup files that are eligible for replication.

The correct parameter to use is `excludeReplicatedLookupSize`, defined in `distsearch.conf`. This parameter specifies a maximum file size (in megabytes) beyond which lookup files are excluded from bundle replication. By setting this to 500, any lookup file larger than 500 MB will not be replicated to search peers.

This is especially important for performance optimization and preventing unnecessary network load during search head to indexer communication.

Example configuration (`distsearch.conf`):

[replicationSettings]

`excludeReplicatedLookupSize = 500`

Reference (Splunk Documentation):

* `distsearch.conf.spec` and example # `excludeReplicatedLookupSize`

* Splunk Enterprise Distributed Search Manual # "Control knowledge bundle replication between search heads and indexers"

* Splunk Admin Manual # "Prevent large lookup files from being replicated"

NEW QUESTION: 16

Which of the following must be done to define user permissions when integrating Splunk with LDAP?

A. Map Users

B. Map Groups

C. Map LDAP Inheritance

D. Map LDAP to Active Directory

Answer: B ([LEAVE A REPLY](#))

<https://docs.splunk.com/Documentation/Splunk/8.1.3/Security/ConfigureLDAPwithSplunkWeb>

"You can map either users or groups, but not both. If you are using groups, all users must be members of an appropriate group. Groups inherit capabilities from the highest level role they're a member of." "If your LDAP environment does not have group entries, you can treat each user as its own group." Reference:

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/ConfigureLDAPwithSplunkWeb>

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:

<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 17

In addition to single, non-clustered Splunk instances, what else can the deployment server push apps to?

- A. Universal forwarders
- B. Splunk Cloud
- C. Linux package managers
- D. Windows using WMI

Answer: (SHOW ANSWER)

Reference: <https://community.splunk.com/t5/Deployment-Architecture/Push-apps-from-deployment-server-automatically-to-universal/m-p/328191> The deployment server is a Splunk component that distributes apps and other configurations to deployment clients, which are Splunk instances that receive updates from the deployment server. The deployment server can push apps to single, non-clustered Splunk instances, as well as universal forwarders, which are lightweight Splunk agents that forward data to indexers. Therefore, option A is the correct answer.

References: Splunk Enterprise Certified Admin | Splunk, [About deployment server and forwarder management - Splunk Documentation]

NEW QUESTION: 18

Windows can prevent a Splunk forwarder from reading open files. If files need to be read while they are being written to, what type of input stanza needs to be created?

- A. Tail Reader
- B. Upload
- C. MonitorNoHandle
- D. Monitor

Answer: (SHOW ANSWER)

The correct answer is C. MonitorNoHandle.

MonitorNoHandle is a type of input stanza that allows a Splunk forwarder to read files on Windows systems as Windows writes to them. It does this by using a kernel-mode filter driver to capture raw data as it gets written to the file1. This input stanza is useful for files that get locked open for writing, such as the Windows DNS server log file2.

The other options are incorrect because:

A: Tail Reader is not a valid input stanza in Splunk. It is a component of the Tailing Processor, which is responsible for monitoring files and directories for new data³.

B: Upload is a type of input stanza that allows Splunk to index a single file from a local or network file system. It is not suitable for files that are constantly being updated, as it only indexes the file once and does not monitor it for changes⁴.

D: Monitor is a type of input stanza that allows Splunk to monitor files and directories for new data. However, it may not work for files that Windows prevents Splunk from reading while they are open. In such cases, MonitorNoHandle is a better option².

A Splunk forwarder is a lightweight agent that can forward data to a Splunk deployment. There are two types of forwarders: universal and heavy. A universal forwarder can only forward data, while a heavy forwarder can also perform parsing, filtering, routing, and aggregation on the data before forwarding it⁵.

An input stanza is a section in the inputs.conf configuration file that defines the settings for a specific type of input, such as files, directories, network ports, scripts, or Windows event logs. An input stanza starts with a square bracket, followed by the input type and the input path or name. For example, [monitor:///var/log] is an input stanza for monitoring the /var/log directory.

References:

1: Monitor files and directories - Splunk Documentation

2: How to configure props.conf for proper line breaking ... - Splunk Community

3: How Splunk Enterprise monitors files and directories - Splunk Documentation

4: Upload a file - Splunk Documentation

5: Use forwarders to get data into Splunk Enterprise - Splunk Documentation

6: inputs.conf - Splunk Documentation

NEW QUESTION: 19

What are the values for `host` and `index` for `[stanza1]` used by Splunk during index time, given the following configuration files?

```
SPLUNK_HOME/etc/system/local/inputs.conf:
[stanza1]
host=server1

SPLUNK_HOME/etc/apps/search/local/inputs.conf:
[stanza1]
host=searchsvr1
index=searchinfo

SPLUNK_HOME/etc/apps/search/local/inputs.conf:
[stanza1]
host=unixsvr1
index=unixinfo
```

A. `host=server1``index=unixinfo`

B. `host=server1``index=searchinfo`

C. `host=searchsvr1``index=searchinfo`

D. `host=unixsvr1``index=unixinfo`

Answer: (SHOW ANSWER)

- etc/system/local/ has better precedence at index time - for identical settings in the same file, the last one overwrite others, see :<https://community.splunk.com/t5/Getting-Data-In/What-is-the-precedence-for-identical-stanzas-within-a-single-m-p/283566>

NEW QUESTION: 20

Which artifact is required in the request header when creating an HTTP event?

- A. ackID
- B. Token
- C. Manifest
- D. Host name

Answer: (SHOW ANSWER)

Reference:<https://docs.splunk.com/Documentation/Splunk/8.2.3/Data/FormateventsforHTTPEventCollector> When creating an HTTP event, the request header must include a token that identifies the HTTP Event Collector (HEC) endpoint. The token is a 32-character hexadecimal string that is generated when the HEC endpoint is created. The token is used to authenticate the request and route the event data to the correct index.

Therefore, option B is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [About HTTP Event Collector - Splunk Documentation]

NEW QUESTION: 21

Load balancing on a Universal Forwarder is not scaling correctly. The forwarder's outputs. and the tcpout stanza are setup correctly. What else could be the cause of this scaling issue? (select all that apply)

- A. The receiving port is not properly setup to listen on the right port.
- B. The inputs . conf'S _SYSZOG _ROVTING is not setup to use the right group names.
- C. The DNS record used is not setup with a valid list of IP addresses.
- D. The indexAndForward value is not set properly.

Answer: (SHOW ANSWER)

The possible causes of the load balancing issue on the Universal Forwarder are A and C. The receiving port and the DNS record are both factors that affect the ability of the Universal Forwarder to distribute data across multiple receivers. If the receiving port is not properly set up to listen on the right port, or if the DNS record used is not set up with a valid list of IP addresses, the Universal Forwarder might fail to connect to some or all of the receivers, resulting in poor load balancing.

NEW QUESTION: 22

Which of the following are supported configuration methods to add inputs on a forwarder? (select all that apply)

- A. CLI
- B. Edit inputs . conf
- C. Edit forwarder.conf
- D. Forwarder Management

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Forwarder/8.2.1/Forwarder/HowtoforwarddatatoSplunkEnterprise>

"You can collect data on the universal forwarder using several methods. Define inputs on the universal forwarder with the CLI. You can use the CLI to define inputs on the universal forwarder. After you define the inputs, the universal forwarder

collects data based on those definitions as long as it has access to the data that you want to monitor. Define inputs on the universal forwarder with configuration files. If the input you want to configure does not have a CLI argument for it, you can configure inputs with configuration files. Create an inputs.conf file in the directory, \$SPLUNK_HOME/etc/system/local

NEW QUESTION: 23

Which additional component is required for a search head cluster?

- A. Deployer
- B. Cluster Master
- C. Monitoring Console
- D. Management Console

Answer: (SHOW ANSWER)

Reference:<https://docs.splunk.com/Documentation/Splunk/8.0.5/DistSearch/SHCdeploymentoverview> The deployer. This is a Splunk Enterprise instance that distributes apps and other configurations to the cluster members. It stands outside the cluster and cannot run on the same instance as a cluster member. It can, however, under some circumstances, reside on the same instance as other Splunk Enterprise components, such as a deployment server or an indexer cluster master node.

NEW QUESTION: 24

A Universal Forwarder is monitoring a very active syslog stream and as a result is unable to switch between destinations. How would an admin safely remediate this issue?

- A. Configure and enable the LINE_BREAKER on the forwarder.
- B. Configure useAck on the forwarder.
- C. Configure forceTimebasedAutoLB on the forwarder.
- D. Configure and enable the FVFNT BREAKER on the forwarder.

Answer: (SHOW ANSWER)

The Universal Forwarder (UF) handles data forwarding to indexers. When monitoring a continuous and high-volume syslog stream over TCP, the UF may not detect an end-of-file (EOF) condition, which is typically required to trigger load balancing between multiple indexers. This can result in the UF continuously sending data to a single indexer, potentially leading to uneven load distribution.

To address this, Splunk provides the forceTimebasedAutoLB setting in the outputs.conf configuration file.

Enabling this setting allows the UF to switch between indexers at regular time intervals, regardless of EOF detection. This ensures a more balanced distribution of data across multiple indexers, even in scenarios with continuous data streams like syslog.

Reference:

Configure forwarding with outputs.conf - Splunk Documentation

NEW QUESTION: 25

Which data pipeline phase is the last opportunity for defining event boundaries?

- A. Input phase
- B. Indexing phase
- C. Parsing phase
- D. Search phase

Answer: (SHOW ANSWER)

Reference: <https://docs.splunk.com/Documentation/Splunk/8.2.3/Admin/Configurationparametersandthedatapipeline>

The parsing phase is the process of extracting fields and values from raw data. The parsing phase respects `LINE_BREAKER`, `SHOULD_LINEMERGE`, `BREAK_ONLY_BEFORE_DATE`, and all other line merging settings in `props.conf`. These settings determine how Splunk breaks the data into events based on certain criteria, such as timestamps or regular expressions. The event boundaries are defined by the `props.conf` file, which can be modified by the administrator. Therefore, the parsing phase is the last opportunity for defining event boundaries.

NEW QUESTION: 26

What is the valid option for a `[monitor]` stanza in `inputs.conf`?

- A. enabled
- B. datasource
- C. server_name
- D. ignoreOlderThan

Answer: (SHOW ANSWER)

Setting: `ignoreOlderThan = <time_window>` Description: "Causes the input to stop checking files for updates if the file modification time has passed the `<time_window>` threshold." Default: 0 (disabled) Reference: <https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/Monitorfilesanddirectorieswithininputs.conf>

NEW QUESTION: 27

Which default Splunk role could be assigned to provide users with the following capabilities?

Create saved searches

Edit shared objects and alerts

Not allowed to create custom roles

- A. admin
- B. power
- C. user
- D. splunk-system-role

Answer: (SHOW ANSWER)

Reference: <https://docs.splunk.com/Documentation/Splunk/8.2.3/Admin/Aboutusersandroles> The power role is a default Splunk role that grants users the ability to create saved searches, edit shared objects and alerts, and access advanced search commands. However, the power role does not allow users to create custom roles, which is a privilege reserved for the admin role. Therefore, option B is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [About configuring role-based user access - Splunk Documentation]

NEW QUESTION: 28

The `LINE_BREAKER` attribute is configured in which configuration file?

- A. `props.conf`
- B. `indexes.conf`
- C. `inputs.conf`

D. transforms.conf

Answer: (SHOW ANSWER)

Reference: <https://docs.splunk.com/Documentation/SplunkCloud/8.2.2105/Data/Configureeventlinebreaking>

NEW QUESTION: 29

Which file will be matched for the following monitor stanza in inputs.conf?

[monitor: ///var/log/*/bar/*.txt]

A. /var/log/host_460352847/temp/bar/file/csv/foo.txt

B. /var/log/host_460352847/bar/foo.txt

C. /var/log/host_460352847/bar/file/foo.txt

D. /var/ log/ host_460352847/temp/bar/file/foo.txt

Answer: (SHOW ANSWER)

The correct answer is C. /var/log/host_460352847/bar/file/foo.txt.

The monitor stanza in inputs.conf is used to configure Splunk to monitor files and directories for new data.

The monitor stanza has the following syntax1:

[monitor://<input path>]

The input path can be a file or a directory, and it can include wildcards (*) and regular expressions. The wildcards match any number of characters, including none, while the regular expressions match patterns of characters. The input path is case-sensitive and must be enclosed in double quotes if it contains spaces1.

In this case, the input path is /var/log//bar/.txt, which means Splunk will monitor any file with the .txt extension that is located in a subdirectory named bar under the /var/log directory. The subdirectory bar can be at any level under the /var/log directory, and the * wildcard will match any characters before or after the bar and .txt parts1.

Therefore, the file /var/log/host_460352847/bar/file/foo.txt will be matched by the monitor stanza, as it meets the criteria. The other files will not be matched, because:

A: /var/log/host_460352847/temp/bar/file/csv/foo.txt has a .csv extension, not a .txt extension.

B: /var/log/host_460352847/bar/foo.txt is not located in a subdirectory under the bar directory, but directly in the bar directory.

D: /var/log/host_460352847/temp/bar/file/foo.txt is located in a subdirectory named file under the bar directory, not directly in the bar directory.

NEW QUESTION: 30

Amanda is tasked with hiding the first 5 digits of the account number in the following log and replacing them with xxxxx.

Example events:

[22/Oct/2014:00:46:27] VendorID=9112 Code=B AcctID=4902636940

[22/Oct/2014:00:48:40] VendorID=1004 Code=J AcctID=4236256056

[22/Oct/2014:00:50:02] VendorID=5034 Code=H AcctID=0462999288

Which props.conf configuration would achieve this goal?

A. [source:....../vendor_sales.log]

TRANSFORMS-acct = s/AcctID=\d{5}(\d{5})/AcctID=xxxxx\1/g

B. [source:....../vendor_sales.log]

REPLACE-acct = s/AcctID=\d{5}(\d{5})/AcctID=xxxxx\1/g

C. [source::.../vendor_sales.log]

SEDCMD-acct = s/AcctID=\d{5}(\d{5})/AcctID=xxxxx\1/g

D. [source::.../vendor_sales.log]

SED-acct = s/AcctID=\d{5}(\d{5})/AcctID=xxxxx\1/g

Answer: C (LEAVE A REPLY)

The correct props.conf setting is:

[source::.../vendor_sales.log]

SEDCMD-acct = s/AcctID=\d{5}(\d{5})/AcctID=xxxxx\1/g

SEDCMD is the correct props.conf attribute for performing sed-style substitutions against raw event data during indexing.

The regular expression matches AcctID= followed by the first five digits, captures the last five digits, and replaces the first five digits with xxxxx.

Example:

AcctID=4902636940

becomes:

AcctID=xxxxx36940

Why the other options are incorrect:

A). TRANSFORMS-acct is used to call a transform stanza from transforms.conf, not to directly run a sed substitution.

B). REPLACE-acct is not the valid props.conf setting for this task.

D). SED-acct is not the valid Splunk setting; the correct setting name is SEDCMD- < class > .

Reference: Splunk Enterprise Admin Manual - props.conf specification, SEDCMD- < class > setting; Splunk Enterprise Getting Data In Manual - anonymize data with SEDCMD.

NEW QUESTION: 31

Syslog files are being monitored on a Heavy Forwarder.

Where would the appropriate TRANSFORMS setting be deployed to reroute logs based on the event message?

A. Heavy Forwarder

B. Indexer

C. Search head

D. Deployment server

Answer: (SHOW ANSWER)

A Heavy Forwarder is a Splunk instance that can parse and filter data before forwarding it to another Splunk instance, such as an indexer¹. A Heavy Forwarder can also perform index-time field extractions using the TRANSFORMS setting².

The TRANSFORMS setting is used to configure data transformations in the transforms.conf file³. The transforms.conf file contains settings and values that you can use to configure host and source type overrides, anonymize sensitive data, route events to different indexes, create index-time and search-time field extractions, and set up lookup tables³.

The TRANSFORMS setting can be deployed to the Heavy Forwarder where the syslog files are being monitored, so that the logs can be rerouted based on the event message before they are forwarded to the indexer². This can improve the performance and efficiency of data processing and indexing².

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:
<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 32

An index stores its data in buckets. Which default directories does Splunk use to store buckets? (Choose all that apply.)

- A. bucketdb
- B. frozendb
- C. colddb
- D. db

Answer: (SHOW ANSWER)

Reference: <https://wiki.splunk.com/Deploy:BucketRotationAndRetention>

NEW QUESTION: 33

Event processing occurs at which phase of the data pipeline?

- A. Search
- B. Indexing
- C. Parsing
- D. Input

Answer: (SHOW ANSWER)

According to the Splunk documentation¹, event processing occurs at the parsing phase of the data pipeline. The parsing phase is where Splunk software processes incoming data into individual events, extracts timestamp information, assigns source types, and performs other tasks to make the data searchable¹. The parsing phase can also apply field extractions, event type matching, and other transformations to the events².

NEW QUESTION: 34

An organization wants to collect Windows performance data from a set of clients, however, installing Splunk software on these clients is not allowed. What option is available to collect this data in Splunk Enterprise?

- A. Use Local Windows host monitoring.
- B. Use Windows Remote Inputs with WMI.
- C. Use Local Windows network monitoring.
- D. Use an index with an Index Data Type of Metrics.

Answer: (SHOW ANSWER)

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

"The Splunk platform collects remote Windows data for indexing in one of two ways: From Splunk forwarders, Using Windows Management Instrumentation (WMI). For Splunk Cloud deployments, you must use the Splunk Universal Forwarder on a Windows machines to montior remote Windows data."

NEW QUESTION: 35

When configuring HTTP Event Collector (HEC) input, how would one ensure the events have been indexed?

- A. Enable indexer acknowledgment.
- B. Enable forwarder acknowledgment.
- C. `splunk check-integrity -index <index name>`
- D. `index=_internal component=ACK | stats count by host`

Answer: (SHOW ANSWER)

Per the provided Splunk reference URL

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/AboutHECIDXAck>

"While HEC has precautions in place to prevent data loss, it's impossible to completely prevent such an occurrence, especially in the event of a network failure or hardware crash. This is where indexer acknowledgment comes in." Reference <https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/AboutHECIDXAck>

NEW QUESTION: 36

In which Splunk configuration is the SEDCMD used?

- A. `props.conf`
- B. `inputs.conf`
- C. `indexes.conf`
- D. `transforms.conf`

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Forwarding/Forwarddatatothird-partysystems>

"You can specify a SEDCMD configuration in `props.conf` to address data that contains characters that the third-party server cannot process. "

NEW QUESTION: 37

What is the correct curl to send multiple events through HTTP Event Collector?

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: (SHOW ANSWER)

`curl "https://mysplunkserver.example.com:8088/services/collector" \ -H "Authorization: Splunk DF4S7ZE4-3GS1-8SFS-E777-0284GG91PF67" \ -d '{"event": "Hello World"}, {"event": "Hola Mundo"}, {"event":`

`"Hallo Welt"}'`. This is the correct curl command to send multiple events through HTTP Event Collector (HEC), which is a token-based API that allows you to send data to Splunk Enterprise from any application that can make an HTTP request. The command has the following components:

The URL of the HEC endpoint, which consists of the protocol (`https`), the hostname or IP address of the Splunk server (`mysplunkserver.example.com`), the port number (`8088`), and the service name (`services/collector`).

The header that contains the authorization token, which is a unique identifier that grants access to the HEC endpoint. The token is prefixed with Splunk and enclosed in quotation marks. The token value (DF4S7ZE4-3GS1-8SFS-E777-0284GG91PF67) is an example and should be replaced with your own token value.

The data payload that contains the events to be sent, which are JSON objects enclosed in curly braces and separated by commas. Each event object has a mandatory field called event, which contains the raw data to be indexed. The event value can be a string, a number, a boolean, an array, or another JSON object. In this case, the event values are strings that say hello in different languages.

NEW QUESTION: 38

This file has been manually created on a universal forwarder

```
/opt/splunkforwarder/etc/apps/my_TA/local/inputs.conf>

[monitor:///var/log/messages]
sourcetype=syslog
index=syslog
```

A new Splunk admin comes in and connects the universal forwarders to a deployment server and deploys the same app with a new

```
inputs.conf file>

/opt/splunk/etc/deployment-apps/my_TA/local/inputs.conf

[monitor:///var/log/maillog]
sourcetype=maillog
index=syslog
```

Which file is now monitored?

- A. /var/log/maillog
- B. /var/log/maillog and /var/log/messages
- C. /var/log/messages
- D. none of the above

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 39

What conf file needs to be edited to set up distributed search groups?

- A. props.conf
- B. search.conf
- C. distsearch.conf
- D. distributedsearch.conf

Answer: ([SHOW ANSWER](#))

" You can group your search peers to facilitate searching on a subset of them. Groups of search peers are known as " distributed search groups. " You specify distributed search groups in the distsearch.conf file " Reference:

<https://docs.splunk.com/Documentation/Splunk/8.0.5/DistSearch/Distributedsearchgroups>

NEW QUESTION: 40

Which of the following are methods for adding inputs in Splunk? (select all that apply)

- A. CLI
- B. Splunk Web
- C. Editing inputs.conf
- D. Editing monitor.conf

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/8.2.2/Data/Configureyourinputs> Add your data to Splunk Enterprise. With Splunk Enterprise, you can add data using Splunk Web or Splunk Apps. In addition to these methods, you also can use the following methods. -The Splunk Command Line Interface (CLI) -The inputs.conf configuration file. When you specify your inputs with Splunk Web or the CLI, the details are saved in a configuration file on Splunk Enterprise indexer and heavy forwarder instances.

NEW QUESTION: 41

After automatic load balancing is enabled on a forwarder, the time interval for switching indexers can be updated by using which of the following attributes?

- A. channelTTL
- B. connectionTimeout
- C. autoLBFrequency
- D. secsInFailureInterval

Answer: ([SHOW ANSWER](#))

Reference:<https://docs.splunk.com/Documentation/Forwarder/8.2.1/Forwarder/Configureloadbalancing>

NEW QUESTION: 42

An add-on has configured field aliases for source IP address and destination IP address fields. A specific user prefers not to have those fields present in their user context. Based on the defaultprops.conf below, which `SPLUNK_HOME/etc/users/buttercup/myTA/local/props.conf` stanza can be added to the user's local context to disable the field aliases?

```
SPLUNK_HOME/etc/apps/myTA/default/props.conf
[mySourcetype]
FIELDALIAS-cim-src_ip = sourceIPAddress as src_ip
FIELDALIAS-cim-dest-ip = destinationIPAddress as dest_ip
```

```
[mySourcetype]
disable FIELDALIAS-cim-src_ip
disable FIELDALIAS-cim-dest-ip
```

```
[mySourcetype]
FIELDALIAS-cim-src_ip =
FIELDALIAS-cim-dest-ip =
```

```
[mySourcetype]
unset FIELDALIAS-cim-src_ip
unset FIELDALIAS-cim-dest-ip
```

```
[mySourcetype]
#FIELDALIAS-cim-src_ip = sourceIPAddress as src_ip
#FIELDALIAS-cim-dest-ip = destinationIPAddress as dest_ip
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/latest/Admin/Howtoeditaconfigurationfile#Clear%20a%20setting>

NEW QUESTION: 43

A non-clustered Splunk environment has three indexers (A,B,C) and two search heads (X, Y). During a search executed on search head X, indexer A crashes. What is Splunk's response?

- A. Update the user in Splunk web informing them that the results of their search may be incomplete.
- B. Repeat the search request on indexer B without informing the user.
- C. Update the user in Splunk web that their results may be incomplete and that Splunk will try to re-execute the search.
- D. Inform the user in Splunk web that their results may be incomplete and have them attempt the search from search head Y.

Answer: ([SHOW ANSWER](#))

This is explained in the Splunk documentation¹, which states:

If an indexer goes down during a search, the search head notifies you that the results might be incomplete. The search head does not attempt to re-run the search on another indexer.

NEW QUESTION: 44

Which of the following apply to how distributed search works? (select all that apply)

- A. The search head dispatches searches to the peers
- B. The search peers pull the data from the forwarders.
- C. Peers run searches in parallel and return their portion of results.

D. The search head consolidates the individual results and prepares reports

Answer: ([SHOW ANSWER](#))

Users log on to the search head and run reports: - The search head dispatches searches to the peers - Peers run searches in parallel and return their portion of results - The search head consolidates the individual results and prepares reports

NEW QUESTION: 45

Which option on the Add Data menu is most useful for testing data ingestion without creating inputs.conf?

- A. Monitor option
- B. Download option
- C. Upload option
- D. Forward option

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 46

Which configuration files are used to transform raw data ingested by Splunk? (Choose all that apply.)

- A. props.conf
- B. inputs.conf
- C. rawdata.conf
- D. transforms.conf

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/8.1.1/Knowledge/Configureadvancedextractionswithfieldtransforms>

use transformations with props.conf and transforms.conf to:

- Mask or delete raw data as it is being indexed
- Override sourcetype or host based upon event values
- Route events to specific indexes based on event content
- Prevent unwanted events from being indexed

Reference: <https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/Configurtimestamprecognition>

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:

<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 47

What is the correct example to redact a plain-text password from raw events?

- A. in props.conf:[identity]REGEX-redact_pw = s/password=(\[^\|/s\]+)/ #####REACTED#####/g
- B. in props.conf:[identity]SEDCMD-redact_pw = s/password=(\[^\|/s\]+)/ #####REACTED#####/g

C. in transforms.conf:[identity]SEDCMD-redact_pw = s/password=(^[^,|/s]+)/ #####REACTED#####/g

D. in transforms.conf:[identity]REGEX-redact_pw = s/password=(^[^,|/s]+)/ #####REACTED#####/g

Answer: (SHOW ANSWER)

The correct answer is B. in props.conf:

[identity]

SEDCMD-redact_pw = s/password=(^[^,|/s]+)/ #####REACTED#####/g

According to the Splunk documentation¹, to redact sensitive data from raw events, you need to use the SEDCMD attribute in the props.conf file. The SEDCMD attribute applies a sed expression to the raw data before indexing. The sed expression can use the s command to replace a pattern with a substitution string. For example, the following sed expression replaces any occurrence of password= followed by any characters until a comma, whitespace, or slash with #####REACTED#####:

s/password=(^[^,|/s]+)/ #####REACTED#####/g

The g flag at the end means that the replacement is applied globally, not just to the first match.

Option A is incorrect because it uses the REGEX attribute instead of the SEDCMD attribute. The REGEX attribute is used to extract fields from events, not to modify them.

Option C is incorrect because it uses the transforms.conf file instead of the props.conf file. The transforms.conf file is used to define transformations that can be applied to fields or events, such as lookups, evaluations, or replacements. However, these transformations are applied after indexing, not before.

Option D is incorrect because it uses both the wrong attribute and the wrong file. There is no REGEX-redact_pw attribute in the transforms.conf file.

References:1:Redact data from events - Splunk Documentation

NEW QUESTION: 48

In which phase of the index time process does the license metering occur?

A. input phase

B. Parsing phase

C. Indexing phase

D. Licensing phase

Answer: (SHOW ANSWER)

"When ingesting event data, the measured data volume is based on the new raw data that is placed into the indexing pipeline. Because the data is measured at the indexing pipeline, data that is filtered and dropped prior to indexing does not count against the license volume quota."

<https://docs.splunk.com/Documentation/Splunk/8.0.6/Admin/HowSplunklicensingworks>

NEW QUESTION: 49

The CLI command splunk add forward-server indexer: < receiving-port > will create stanza(s) in which configuration file?

A. inputs.conf

B. indexes.conf

C. outputs.conf

D. servers.conf

Answer: (SHOW ANSWER)

The CLI command " Splunk add forward-server indexer: < receiving-port > " is used to define the indexer and the listening port on forwards. The command creates this kind of entry " [tcpout-server:// < ip address > : < port >] " in the outputs.conf file.

<https://docs.splunk.com/Documentation/Forwarder/8.2.2/Forwarder/Configureforwardingwithoutoutputs.conf> Reference:

<https://docs.splunk.com/Documentation/Forwarder/8.0.5/Forwarder/Enableareceiver>

NEW QUESTION: 50

A Universal Forwarder has the following active stanza in inputs . conf:

```
[monitor: //var/log]
```

```
disabled = 0
```

```
host = 460352847
```

An event from this input has a timestamp of 10:55. What timezone will Splunk add to the event as part of indexing?

- A. Universal Coordinated Time.
- B. The timezone of the search head.
- C. The timezone of the indexer that indexed the event.
- D. The timezone of the forwarder.

Answer: (SHOW ANSWER)

The correct answer is D. The timezone of the forwarder will be added to the event as part of indexing.

According to the Splunk documentation¹, Splunk software determines the time zone to assign to a timestamp using the following logic in order of precedence:

Use the time zone specified in raw event data (for example, PST, -0800), if present.

Use the TZ attribute set in props.conf, if the event matches the host, source, or source type that the stanza specifies.

If the forwarder and the receiving indexer are version 6.0 or higher, use the time zone that the forwarder provides.

Use the time zone of the host that indexes the event.

In this case, the event does not have a time zone specified in the raw data, nor does it have a TZ attribute set in props.conf. Therefore, the next rule applies, which is to use the time zone that the forwarder provides. A universal forwarder is a lightweight agent that can forward data to a Splunk deployment, and it knows its system time zone and sends that information along with the events to the indexer². The indexer then converts the event time to UTC and stores it in the _time field¹.

The other options are incorrect because:

A) Universal Coordinated Time (UTC) is not the time zone that Splunk adds to the event as part of indexing, but rather the time zone that Splunk uses to store the event time in the _time field. Splunk software converts the event time to UTC based on the time zone that it determines from the rules above¹.

B) The timezone of the search head is not relevant for indexing, as the search head is a Splunk component that handles search requests and distributes them to indexers, but it does not process incoming data³. The search head uses the user's timezone setting to determine the time range in UTC that should be searched and to display the timestamp of the results in the user's timezone².

C) The timezone of the indexer that indexed the event is only used as a last resort, if none of the other rules apply. In this case, the forwarder provides the time zone information, so the indexer does not use its own time zone¹.

NEW QUESTION: 51

Which of the following is the use case for the deployment server feature of Splunk?

- A. Managing distributed workloads in a Splunk environment.
- B. Automating upgrades of Splunk forwarder installations on endpoints.
- C. Orchestrating the operations and scale of a containerized Splunk deployment.
- D. Updating configuration and distributing apps to processing components, primarily forwarders.

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/8.2.2/Updating/Aboutdeploymentserver>

"The deployment server is the tool for distributing configurations, apps, and content updates to groups of Splunk Enterprise instances."

NEW QUESTION: 52

Which of the following methods will connect a deployment client to a deployment server? (select all that apply)

- A. Run `$SPLUNK_HOME/bin/splunk set deploy-poll` : from the command line of the deployment client.
- B. Create and edit a `deploymentserver.conf` file in `$SPLUNK_HOME/etc/system/local` on the deployment server.
- C. Create and edit a `deploymentclient.conf` file in `$SPLUNK_HOME/etc/system/local` on the deployment client.
- D. Run `$SPLUNK_HOME/bin/splunk set deploy-poll` : from the command line of the deployment server.

Answer: ([SHOW ANSWER](#))

The correct methods to connect a deployment client to a deployment server are A and C.

You can either run the command `splunk set deploy-poll < IP_address/hostname > : < management_port >` from the command line of the deployment client1 or create and edit a `deploymentclient.conf` file in `$SPLUNK_HOME/etc/system/local` on the deployment client2. Both methods require you to specify the IP address, hostname, and management port of the deployment server that you want the client to connect to.

NEW QUESTION: 53

A company moves to a distributed architecture to meet the growing demand for the use of Splunk. What parameter can be configured to enable automatic load balancing in the Universal Forwarder to send data to the indexers?

- A. Create one `outputs.conf` file for each of the server addresses in the indexing tier.
- B. Configure the `outputs.conf` file to point to any server in the indexing tier and Splunk will configure the data to be sent to all of the indexers.
- C. Splunk does not do load balancing and requires a hardware load balancer to balance traffic across the indexers.
- D. Set the stanza to have a `server` value equal to a comma-separated list of IP addresses and indexer ports for each of the indexers in the environment.

Answer: D ([LEAVE A REPLY](#))

Set the stanza to have a `server` value equal to a comma-separated list of IP addresses and indexer ports for each of the indexers in the environment. This is explained in the Splunk documentation1, which states:

To enable automatic load balancing, set the stanza to have a `server` value equal to a comma-separated list of IP addresses and indexer ports for each of the indexers in the environment. For example:

```
[tcpout] server=10.1.1.1:9997,10.1.1.2:9997,10.1.1.3:9997
```

The forwarder then distributes data across all of the indexers in the list.

NEW QUESTION: 54

All search-time field extractions should be specified on which Splunk component?

- A. Deployment server
- B. Universal forwarder
- C. Indexer
- D. Search head

Answer: ([SHOW ANSWER](#))

Search-time field extractions are the process of extracting fields from events after they are indexed. Search-time field extractions are specified on the search head, which is the Splunk component that handles searching and reporting. Search-time field extractions are configured in props.conf and transforms.conf files, which are located in the etc/system/local directory on the search head. Therefore, option D is the correct answer.

References: Splunk Enterprise Certified Admin | Splunk, [About fields - Splunk Documentation]

NEW QUESTION: 55

A log file contains 193 days worth of timestamped events. Which monitor stanza would be used to collect data 45 days old and newer from that log file?

- A. followTail = -45d
- B. ignore = 45d
- C. includeNewerThan = -35d
- D. ignoreOlderThan = 45d

Answer: ([SHOW ANSWER](#))

Reference: <https://docs.splunk.com/Documentation/Splunk/8.2.1/Data/Configuretimestamprecognition>

NEW QUESTION: 56

Which Splunk forwarder type allows parsing of data before forwarding to an indexer?

- A. Heavy forwarder
- B. Parsing forwarder
- C. Advanced forwarder
- D. Universal forwarder

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 57

Which Splunk component performs indexing and responds to search requests from the search head?

- A. Forwarder
- B. Search peer
- C. License master
- D. Search head cluster

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Splexicon:Searchpeer>

"A Splunk platform instance that responses to search requests from a search head. The term "Search peer" is usually synonymous with the indexer role in a distributed search topology..."

NEW QUESTION: 58

What is the correct attribute to set in inputs.conf in order to have data sent to a particular indexer group?

- A. `_SPLUNKDB_ROUTING`
- B. `_TCP_ROUTING`
- C. `_UDF_ROUTING`
- D. `_INGEST_ROUTING`

Answer: (SHOW ANSWER)

The correct attribute is:

`_TCP_ROUTING = < target_group >`

In inputs.conf, `_TCP_ROUTING` is used to route data from a specific input to a specific TCP output group.

The target group must be defined in outputs.conf.

Example:

inputs.conf

[monitor:///var/log/app.log]

index = app_index

`_TCP_ROUTING = indexer_group_1`

The corresponding group is defined in outputs.conf:

[tcpout:indexer_group_1]

server = indexer1.example.com:9997,indexer2.example.com:9997

This allows specific inputs to be sent to particular indexer groups.

Reference: Splunk Enterprise Forwarding Data Manual - route inputs to specific indexer groups; Splunk Enterprise Admin Manual - inputs.conf and outputs.conf specifications.

NEW QUESTION: 59

How often does Splunk recheck the LDAP server?

- A. Every 5 minutes
- B. Each time a user logs in
- C. Each time Splunk is restarted
- D. Varies based on `LDAP_refresh` setting.

Answer: B (LEAVE A REPLY)

<https://docs.splunk.com/Documentation/Splunk/8.0.6/Security/ManageSplunkuserroleswithLDAP>

NEW QUESTION: 60

Which of the following indexes come pre-configured with Splunk Enterprise? (select all that apply)

- A. `_license`
- B. `_Internal`
- C. `_external`
- D. `_thefishbucket`

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/8.0.5/Indexer/Howindexingworks>

NEW QUESTION: 61

A new forwarder has been installed with a manually created deploymentclient.conf.

What is the next step to enable the communication between the forwarder and the deployment server?

- A. Restart Splunk on the deployment server.
- B. Enable the deployment client in Splunk Web under Forwarder Management.
- C. Restart Splunk on the deployment client.
- D. Wait for up to the time set in the phoneHomeIntervalInSecs setting.

Answer: (SHOW ANSWER)

The next step to enable the communication between the forwarder and the deployment server after installing a new forwarder with a manually created deploymentclient.conf is to restart Splunk on the deployment client.

The deploymentclient.conf file contains the settings for the deployment client, which is a Splunk instance that receives updates from the deployment server. The file must include the targetUri attribute, which specifies the hostname and management port of the deployment server. To apply the changes in the deploymentclient.conf file, Splunk must be restarted on the deployment client. Therefore, option C is the correct answer.

References: Splunk Enterprise Certified Admin | Splunk, [Configure deployment clients - Splunk Documentation]

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:
<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 62

For single line event sourcetypes. it is most efficient to set SHOULD_linemerge to what value?

- A. True
- B. False
- C. <regex string>
- D. Newline Character

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/latest/Data/Configureeventlinebreaking> Attribute : SHOULD_LINEMERGE = [true|false] Description : When set to true, the Splunk platform combines several input lines into a single event, with configuration based on the settings described in the next section.

NEW QUESTION: 63

What action is required to enable forwarder management in Splunk Web?

- A. Navigate to Settings > Server Settings > General Settings, and set an App server port.
- B. Navigate to Settings > Forwarding and receiving, and click on Enable Forwarding.
- C. Create a server class and map it to a client in SPLUNK_HOME/etc/system/local/serverclass.conf.
- D. Place an app in the SPLUNK_HOME/etc/deployment-apps directory of the deployment server.

Answer: (SHOW ANSWER)

Reference:<https://docs.splunk.com/Documentation/Splunk/8.2.1/Updating/Forwardermanagementoverview>

<https://docs.splunk.com/Documentation/MSApp/2.0.3/MSInfra/Setupdeploymentserver>

"To activate deployment server, you must place at least one app into %SPLUNK_HOME%\etc\deployment- apps on the host you want to act as deployment server. In this case, the app is the "send to indexer" app you created earlier, and the host is the indexer you set up initially.

NEW QUESTION: 64

A security team needs to ingest a static file for a specific incident. The log file has not been collected previously and future updates to the file must not be indexed.

Which command would meet these needs?

- A. `splunk add one shot / opt/ incident [data .log -index incident`
- B. `splunk edit monitor /opt/incident/data.* -index incident`
- C. `splunk add monitor /opt/incident/data.log -index incident`
- D. `splunk edit oneshot [opt/ incident/data.* -index incident`

Answer: (SHOW ANSWER)

The correct answer is A. `splunk add one shot / opt/ incident [data . log -index incident` According to the Splunk documentation¹, the `splunk add one shot` command adds a single file or directory to the Splunk index and then stops monitoring it. This is useful for ingesting static files that do not change or update. The command takes the following syntax:
`splunk add one shot <file> -index <index_name>`

The file parameter specifies the path to the file or directory to be indexed. The index parameter specifies the name of the index where the data will be stored. If the index does not exist, Splunk will create it automatically.

Option B is incorrect because the `splunk edit monitor` command modifies an existing monitor input, which is used for ingesting files or directories that change or update over time. This command does not create a new monitor input, nor does it stop monitoring after indexing.

Option C is incorrect because the `splunk add monitor` command creates a new monitor input, which is also used for ingesting files or directories that change or update over time. This command does not stop monitoring after indexing.

Option D is incorrect because the `splunk edit oneshot` command does not exist. There is no such command in the Splunk CLI.

References:¹Monitor files and directories with inputs.conf - Splunk Documentation

NEW QUESTION: 65

What is the correct order of index time precedence?

(For each of the following, highest precedence is shown at the top and lowest precedence is shown at the bottom)

- A. `$SPLUNK_HOME/etc/system/default`
`$SPLUNK_HOME/etc/system/local`
`$SPLUNK_HOME/etc/users/<username>/<appname>/local`
`$SPLUNK_HOME/etc/apps/<appname>/default`
`$SPLUNK_HOME/etc/apps/<appname>/local`
- B. `$SPLUNK_HOME/etc/users/<username>/<appname>/local`
`$SPLUNK_HOME/etc/apps/<appname>/local`

\$SPLUNK_HOME/etc/apps/<appname>/default

\$SPLUNK_HOME/etc/system/local

\$SPLUNK_HOME/etc/system/default

C. \$SPLUNK_HOME/etc/users/<username>/<appname>/local (Highest)

\$SPLUNK_HOME/etc/system/default

\$SPLUNK_HOME/etc/apps/aaa/local

\$SPLUNK_HOME/etc/apps/zzz/default

\$SPLUNK_HOME/etc/system/local (Lowest of these listed)

D. \$SPLUNK_HOME/etc/apps/<appname>/default

\$SPLUNK_HOME/etc/apps/<appname>/local

\$SPLUNK_HOME/etc/users/<username>/<appname>/local

\$SPLUNK_HOME/etc/system/default

\$SPLUNK_HOME/etc/system/local

Answer: (SHOW ANSWER)

Splunk uses a layered configuration model where settings are loaded in a specific order.

This order determines which configuration takes precedence when multiple settings conflict.

At index time (the point when data is parsed and indexed), the configuration precedence is clearly defined in the official documentation.

From Splunk Docs (props.conf precedence):

Configuration file resolution order (highest to lowest precedence):

\$SPLUNK_HOME/etc/users/ < username > / < appname > /local

\$SPLUNK_HOME/etc/apps/ < appname > /local

\$SPLUNK_HOME/etc/apps/ < appname > /default

\$SPLUNK_HOME/etc/system/local

\$SPLUNK_HOME/etc/system/default

However, for index-time configurations, a slight difference applies:

system/local and users/local are often treated specially, but in practice and according to Splunk Docs, the system /local configs override apps/default, and so on.

In Option C, the correct precedence from highest to lowest is:

/etc/users/local (Highest)

/etc/system/default

/etc/apps/aaa/local

/etc/apps/zzz/default

/etc/system/local (Lowest of these listed)

Though system/local typically has high precedence, when users/local is involved, that is the ultimate override.

Splunk Docs confirms this in:

Configuration file precedence

Configuration layering reference

Therefore, Option C reflects the correct Splunk configuration file precedence order at index time.

NEW QUESTION: 66

Which of the following statements describes how distributed search works?

- A. Forwarders pull data from the search peers.
- B. Search heads store a portion of the searchable data.
- C. The search head dispatches searches to the search peers.
- D. Search results are replicated within the indexer cluster.

Answer: ([SHOW ANSWER](#))

URL <https://docs.splunk.com/Documentation/Splunk/8.2.2/DistSearch/Configuredistributedsearch>

"To activate distributed search, you add search peers, or indexers, to a Splunk Enterprise instance that you designate as a search head. You do this by specifying each search peer manually."

NEW QUESTION: 67

What is required when adding a native user to Splunk? (select all that apply)

- A. Password
- B. Username
- C. Full Name
- D. Default app

Answer: ([SHOW ANSWER](#))

According to the Splunk system admin course PDF, When adding native users, Username and Password ARE REQUIRED

NEW QUESTION: 68

When using a directory monitor input, specific source types can be selectively overridden using which configuration file?

- A. sourcetypes . conf
- B. trans forms . conf
- C. outputs . conf
- D. props . conf

Answer: ([SHOW ANSWER](#))

When using a directory monitor input, specific source types can be selectively overridden using the props.conf file. According to the Splunk documentation¹, "You can specify a source type for data based on its input and source. Specify source type for an input. You can assign the source type for data coming from a specific input, such as /var/log/. If you use Splunk Cloud Platform, use Splunk Web to define source types. If you use Splunk Enterprise, define source types in Splunk Web or by editing the inputs.conf configuration file." However, this method is not very granular and assigns the same source type to all data from an input. To override the source type on a per-event basis, you need to use the props.conf file and the transforms.conf file². The props.conf file contains settings that determine how the Splunk platform processes incoming data, such as how to segment events, extract fields, and assign source types². The transforms.conf file contains settings that modify or filter event data during indexing or search time². You can use these files to create rules that match specific patterns in the event data and assign different source types accordingly². For example, you can create a rule that assigns a source type of apache_error to any event that contains the word "error" in the first line².

NEW QUESTION: 69

Which Splunk configuration file is used to enable data integrity checking?

- A. props.conf

- B. global.conf
- C. indexes.conf
- D. data_integrity.conf

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/8.1.2/Security/Dataintegritycontrol#:~:text=When%20you%20enable%20data%20integrity%20control%2C%20Splunk%20Enterprise%20computes%20hashes,it%20to%20a%2011Hashes%20file.>

Reference: <https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/Dataintegritycontrol>

NEW QUESTION: 70

Which of the following is valid distribute search group?

- A)
- B)

```
[searchGroup:Paris]
default = false
servers = server1:8089, server2:8089
```

- C)

```
[searchGroup:Paris]
default = false
servers = server1:9997, server2:9997
```

- D)

- A. option A
- B. Option B
- C. Option D
- D. Option C

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 71

Which of the following is true when authenticating users to Splunk using LDAP?

- A. LDAP group names must match the Splunk role name defined in authorize.conf.
- B. Splunk will search each LDAP strategy in the order in which they are listed in authentication.conf.
- C. Splunk only supports encrypted LDAP connections.
- D. LDAP will take precedence over local users with the same username as defined in etc/passwd.

Answer: ([SHOW ANSWER](#))

When configuring multiple LDAP strategies in Splunk, the order in which they are listed in the authentication.conf file determines the sequence in which Splunk searches them during authentication.

From the official Splunk documentation:

"To configure multiple LDAP strategies, set the authSettings setting to a comma-separated list of all strategies, in the order in which you want to query the strategies."

- Configure LDAP using configuration files - Splunk Documentation

Therefore, Splunk will search each LDAP strategy in the order specified in the authSettings parameter of the authentication.conf file.

Reference:

Configure LDAP using configuration files - Splunk Documentation

NEW QUESTION: 72

What hardware attribute would need to be changed to increase the number of simultaneous searches (ad-hoc and scheduled) on a single search head?

- A. Disk
- B. CPUs
- C. Memory
- D. Network interface cards

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/7.3.1/DistSearch/SHCarchitecture> Scroll down to section titled, How the cluster handles concurrent search quotas, "Overall search quota. This quota determines the maximum number of historical searches (combined scheduled and ad hoc) that the cluster can run concurrently. This quota is configured with max_Searches_per_cpu and related settings in limits.conf."

NEW QUESTION: 73

Which of the following enables compression for universal forwarders in outputs.conf ?

- A.

```
[udpout:mysplunk_indexer1]
compression=true
```
- B.

```
[tcpout]
defaultGroup=my_indexers
compressed=true
```
- C.

```
/opt/splunkforwarder/bin/splunk enable compression
```
- D.

```
[tcpout:my_indexers] server=mysplunk_indexer1:9997, mysplunk_indexer2:9997
decompression=false
```

Answer: (SHOW ANSWER)

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

Compression

#

This example sends compressed events to the remote indexer.

NOTE: Compression can be enabled TCP or SSL outputs only.

The receiver input port should also have compression enabled.

[tcpout]

server = splunkServer.example.com:4433

compressed = true

NEW QUESTION: 74

Which of the following is the recommended guideline for creating a new user role?

- A. Create a role that incorporates capabilities and index inheritance.
- B. Create a new unique role for each unique user.
- C. There are no recommended guidelines when creating new user roles.
- D. Create two roles based on capabilities and indexes, then utilize inheritance.

Answer: ([SHOW ANSWER](#))

When creating new user roles in Splunk, it's recommended to define roles that incorporate the necessary capabilities and index access permissions. This approach allows for granular control over user access and aligns with best practices for role-based access control.

"You can create custom roles and assign those roles to your users. Custom roles let you make granular adjustments to user access, including... Role inheritance... Capabilities... Allowed and default indexes..."

- About configuring role-based user access -

By creating roles that incorporate both capabilities and index access, administrators can efficiently manage user permissions and maintain a secure Splunk environment.

Reference:

About configuring role-based user access - Splunk Documentation

NEW QUESTION: 75

The volume of data from collecting log files from 50 Linux servers and 200 Windows servers will require multiple indexers. Following best practices, which types of Splunk component instances are needed?

- A. Indexers, search head, universal forwarders, license master
- B. Indexers, search head, deployment server, universal forwarders
- C. Indexers, search head, deployment server, license master, universal forwarder
- D. Indexers, search head, deployment server, license master, universal forwarder, heavy forwarder

Answer: ([SHOW ANSWER](#))

Indexers, search head, deployment server, license master, universal forwarder. This is the combination of Splunk component instances that are needed to handle the volume of data from collecting log files from 50 Linux servers and 200 Windows servers, following the best practices. The roles and functions of these components are:

Indexers: These are the Splunk instances that index the data and make it searchable. They also perform some data processing, such as timestamp extraction, line breaking, and field extraction. Multiple indexers can be clustered together to provide high availability, data replication, and load balancing.

Search head: This is the Splunk instance that coordinates the search across the indexers and merges the results from them. It also provides the user interface for searching, reporting, and dashboarding. A search head can also be clustered with other search heads to provide high availability, scalability, and load balancing.

Deployment server: This is the Splunk instance that manages the configuration and app deployment for the universal forwarders. It allows the administrator to centrally control the inputs.conf, outputs.conf, and other configuration files for the forwarders, as well as distribute apps and updates to them.

License master: This is the Splunk instance that manages the licensing for the entire Splunk deployment. It tracks the license usage of all the Splunk instances and enforces the license limits and violations. It also allows the administrator to add, remove, or change licenses.

Universal forwarder: These are the lightweight Splunk instances that collect data from various sources and forward it to the indexers or other forwarders. They do not index or parse the data, but only perform minimal processing, such as compression and encryption. They are installed on the Linux and Windows servers that generate the log files.

NEW QUESTION: 76

In a customer managed Splunk Enterprise environment, what is the endpoint URI used to collect data?

- A. services/collector
- B. data/collector
- C. services/inputs?raw
- D. services/data/collector

Answer: (SHOW ANSWER)

This is the endpoint URI used to collect data using the HTTP Event Collector (HEC), which is a token-based API that allows you to send data to Splunk Enterprise from any application that can make an HTTP request.

The endpoint URI consists of the protocol (http or https), the hostname or IP address of the Splunk server, the port number (default is 8088), and the service name (services/collector). For example:

`https://mysplunkserver.example.com:8088/services/collector`

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:

<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 77

An admin is running the latest version of Splunk with a 500 GB license. The current daily volume of new data is 300 GB per day. To minimize license issues, what is the best way to add 10 TB of historical data to the index?

- A. Buy a bigger Splunk license.
- B. Add 2.5 TB each day for the next 5 days.
- C. Add all 10 TB in a single 24 hour period.
- D. Add 200 GB of historical data each day for 50 days.

Answer: C (LEAVE A REPLY)

`https://docs.splunk.com/Documentation/Splunk/8.1.2/Admin/Aboutlicenseviolations`

"An Enterprise license stack with a license volume of 100 GB of data per day or more does not currently violate."

NEW QUESTION: 78

Which of the following is a valid method to create a Splunk user?

- A. Create a support ticket.
- B. Create a user on the host operating system.
- C. Splunk REST API.
- D. Add the username to users.conf.

Answer: ([SHOW ANSWER](#))

Splunk REST API is a valid method to create a Splunk user. The Splunk REST API allows administrators to create, edit, and delete users programmatically using HTTP requests. The other options are not valid methods to create a Splunk user. Creating a support ticket does not create a user, but requests assistance from Splunk support. Creating a user on the host operating system does not create a Splunk user, but a system user. Adding the username to users.conf does not create a user, but modifies the configuration file that stores user information. References = Configure users with the CLI, Configure users and roles

NEW QUESTION: 79

Which of the following is true regarding LDAP integration with Splunk Enterprise?

- A. Having the change authentication capability will not allow setup of the LDAP integration.
- B. Mappings can be changed at any time if the user has the power role.
- C. A user cannot log in via LDAP unless they have an associated Splunk role.
- D. LDAP integration will not function unless all groups are mapped to an LDAP group.

Answer: C ([LEAVE A REPLY](#))

In Splunk Enterprise, when integrating with an LDAP (Lightweight Directory Access Protocol) directory for authentication, user access is governed by the mapping between LDAP groups and Splunk roles. A user authenticated via LDAP must belong to at least one LDAP group that is mapped to a Splunk role. Without this mapping, the user can authenticate successfully against LDAP but will not be granted any role privileges inside Splunk, and therefore cannot log in to the Splunk web interface.

Splunk documentation explicitly states:

"When you integrate Splunk Enterprise with LDAP, a user must be assigned at least one Splunk role through an LDAP group mapping. If the user does not belong to a mapped group, they cannot log into Splunk." This ensures that user permissions are inherited from LDAP-to-role mappings and provides centralized management of authentication and authorization.

Reference (Splunk Documentation):

- * Splunk Enterprise Admin Manual # Securing Splunk Enterprise # Authenticate users with LDAP
- * authentication.conf.spec and example # LDAP configuration and role mapping
- * Splunk Docs: "Configure LDAP authentication"

NEW QUESTION: 80

Which of the following are reasons to create separate indexes? (Choose all that apply.)

- A. Different retention times.
- B. Increase number of users.
- C. Restrict user permissions.
- D. File organization.

Answer: ([SHOW ANSWER](#))

Reference:<https://community.splunk.com/t5/Getting-Data-In/Why-does-Splunk-have-multiple-indexes/m-p/12063>

Different retention times: You can set different retention policies for different indexes, depending on how long you want to keep the data. For example, you can have an index for security data that has a longer retention time than an index for performance data that has a shorter retention time.

Restrict user permissions: You can set different access permissions for different indexes, depending on who needs to see the data. For example, you can have an index for sensitive data that is only accessible by certain users or roles, and an index for public data that is accessible by everyone.

NEW QUESTION: 81

If an update is made to an attribute in inputs.conf on a universal forwarder, on which Splunk component would the fishbucket need to be reset in order to reindex the data?

- A. Indexer
- B. Forwarder
- C. Search head
- D. Deployment server

Answer: (SHOW ANSWER)

https://www.splunk.com/en_us/blog/tips-and-tricks/what-is-this-fishbucket-thing.html

"Every Splunk instance has a fishbucket index, except the lightest of hand-tuned lightweight forwarders, and if you index a lot of files it can get quite large. As any other index, you can change the retention policy to control the size via indexes.conf"

Reference <https://community.splunk.com/t5/Archive/How-to-reindex-data-from-a-forwarder/td-p/93310>

NEW QUESTION: 82

Running this search in a distributed environment:

On what Splunk component does the eval command get executed?

- A. Heavy Forwarders
- B. Universal Forwarders
- C. Search peers
- D. Search heads

Answer: (SHOW ANSWER)

The eval command is a distributable streaming command, which means that it can run on the search peers in a distributed environment¹. The search peers are the indexers that store the data and perform the initial steps of the search processing².

The eval command calculates an expression and puts the resulting value into a search results field¹. In your search, you are using the eval command to create a new field called

"responsible_team" based on the values in the "account" field.

NEW QUESTION: 83

Seven different network switches are sending traffic to a server hosting a Universal Forwarder. Three of the devices are sending TCP data and four of the devices are sending UDP data.

What is the minimum number of input stanzas that must be created on the Universal Forwarder to successfully capture data from all seven sources?

- A. One
- B. Seven
- C. Four
- D. Two

Answer: (SHOW ANSWER)

In Splunk Enterprise and Splunk Universal Forwarder, data inputs are configured using stanzas in inputs.conf

. Each stanza defines a listener for a particular input type (for example, TCP or UDP) and a specific port.

Splunk documentation states that a single TCP input stanza can receive data from multiple remote hosts sending to the same TCP port, and similarly, a single UDP input stanza can receive data from multiple devices sending to the same UDP port. Therefore, the number of sending devices does not determine the number of stanzas required; rather, the input protocol and port type do.

In this case:

- * All three TCP devices can send data to one TCP port (one stanza).

- * All four UDP devices can send data to one UDP port (one stanza).

Thus, the minimum number of input stanzas required is two - one for TCP and one for UDP.

Example configuration (inputs.conf):

```
# TCP input for three switches sending via TCP
```

```
[tcp://9997]
```

```
sourcetype = switch_logs
```

```
# UDP input for four switches sending via UDP
```

```
[udp://514]
```

```
sourcetype = switch_logs
```

This configuration ensures all seven devices' logs are collected without creating individual stanzas for each device.

Reference (Splunk Documentation):

- * Splunk Enterprise Admin Manual # Configure Data Inputs # "Listen for network data"

- * inputs.conf.spec and example # "You can configure a single TCP or UDP input to receive data from multiple remote hosts."

- * Splunk Universal Forwarder Manual # Configure Forwarding Inputs # "Universal Forwarders can listen on a single TCP or UDP port for multiple remote data sources."

NEW QUESTION: 84

What configuration file are remote Windows Management Instrumentation inputs defined in?

- A. wmi_inputs.conf
- B. inputs.conf
- C. None, the inputs are defined outside of Splunk.
- D. wmi.conf

Answer: (SHOW ANSWER)

In Splunk Enterprise, when collecting data from remote Windows machines using Windows Management Instrumentation (WMI), the configurations are defined in the wmi.conf file. This file specifies the parameters for connecting to WMI providers and defines the data inputs.

The wmi.conf file is located in the \$SPLUNK_HOME\etc\system\local\ directory. It contains stanzas that define global settings and input-specific configurations for WMI data collection. This setup allows Splunk to collect various types of data

from remote Windows systems, such as event logs and performance metrics, without requiring a forwarder on the remote machine.

Reference:

Monitor data through Windows Management Instrumentation (WMI) - Splunk Documentation wmi.conf - Splunk Documentation

NEW QUESTION: 85

Which of the following is a benefit of distributed search?

- A. Peers run search in sequence.
- B. Peers run search in parallel.
- C. Resilience from indexer failure.
- D. Resilience from search head failure.

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/8.2.2/DistSearch/Whatisdistributedsearch> Parallel reduce search processing If you struggle with extremely large high-cardinality searches, you might be able to apply parallel reduce processing to them to help them complete faster. You must have a distributed search environment to use parallel reduce search processing.

NEW QUESTION: 86

Assume a file is being monitored and the data was incorrectly indexed to an exclusive index. The index is cleaned and now the data must be reindexed. What other index must be cleaned to reset the input checkpoint information for that file?

- A. _audit
- B. _checkpoint
- C. _introspection
- D. _thefishbucket

Answer: (SHOW ANSWER)

--reset Reset the fishbucket for the given key or file in the btree. Resetting the checkpoint for an active monitor input reindexes data, resulting in increased license use.<https://docs.splunk.com/Documentation/Splunk/8.1.1/Troubleshooting/CommandlinetoolsforusewithSupport>

Reference:

<http://docshare02.docshare.tips/files/4773/47733589.pdf>

NEW QUESTION: 87

Which option accurately describes the purpose of the HTTP Event Collector (HEC)?

- A. A token-based HTTP input that is secure and scalable and that requires the use of forwarders
- B. A token-based HTTP input that is secure and scalable and that does not require the use of forwarders.
- C. An agent-based HTTP input that is secure and scalable and that does not require the use of forwarders.
- D. A token-based HTTP input that is insecure and non-scalable and that does not require the use of forwarders.

Answer: B (LEAVE A REPLY)

<https://docs.splunk.com/Documentation/Splunk/8.2.2/Data/UsetheHTTPEventCollector>

"The HTTP Event Collector (HEC) lets you send data and application events to a Splunk deployment over the HTTP and Secure HTTP (HTTPS) protocols. HEC uses a token-based authentication model. You can generate a token and then

configure a logging library or HTTP client with the token to send data to HEC in a specific format. This process eliminates the need for a Splunk forwarder when you send application events."

NEW QUESTION: 88

Which configuration file would be used to forward the Splunk internal logs from a search head to the indexer?

- A. props.conf
- B. inputs.conf
- C. outputs.conf
- D. collections.conf

Answer: (SHOW ANSWER)

<https://docs.splunk.com/Documentation/Splunk/8.1.1/DistSearch/Forwardsearchheaddata> Per the provided Splunk reference URL by @hwangho, scroll to section Forward search head data, subsection titled, 2. Configure the search head as a forwarder. "Create an outputs.conf file on the search head that configures the search head for load-balanced forwarding across the set of search peers (indexers)." Reference: <https://community.splunk.com/t5/Getting-Data-In/How-to-configure-search-head-to-forwardinternal-data-to-the/td-p/111658>

NEW QUESTION: 89

Which scenario is applicable given the stanzas in authentication.conf below?

[authentication]

externalTwoFactorAuthVendor = Duo

externalTwoFactorAuthSettings = duoMFA

[duoMFA]

integrationKey = aGFwcHliaXJ0aGRheU1pZGR5

secretKey = YXVzdHJhaWxpYW5Gb3JHcmVw

applicationKey = c3BsaW5raW5ndGhlcGx1bWJ1c3NpbmN1OTU

apiHostname = 466993018.duosecurity.com

failOpen = True

timeout = 60

- A. If Splunk cannot connect to the multifactor authentication provider, all logins will be denied.
- B. Multifactor authentication is required to log into the host operating system.
- C. The secretKey does not need to be protected since multifactor authentication is turned on.
- D. If Splunk cannot connect to the multifactor authentication provider, authentications will be successful without completing a multifactor challenge.

Answer: (SHOW ANSWER)

The failOpen setting in the [duoMFA] stanza determines how Splunk software handles authentication requests when it cannot connect to the Duo Security service. If failOpen is set to True, as in this example, Splunk software allows users to log in without completing a multifactor challenge. If failOpen is set to False, Splunk software denies all logins when it cannot connect to Duo Security. This setting is independent of the authentication type or the secretKey protection. References = Connect to Duo Security for multifactor authentication

NEW QUESTION: 90

What is the command to reset the fishbucket for one source?

- A. `rm -r ~/splunkforwarder/var/lib/splunk/fishbucket`
- B. `splunk clean eventdata -index _thefishbucket`
- C. `splunk cmd btprobe -d SPLUNK_HOME/var/lib/splunk/fishbucket/splunk_private_db --file < source > --reset`
- D. `splunk btool fishbucket reset < source >`

Answer: ([SHOW ANSWER](#))

Reference: <https://community.splunk.com/t5/Getting-Data-In/How-can-I-trigger-the-re-indexing-of-a-single-file/m-p/108568>

The fishbucket is a directory that stores information about the files that have been monitored and indexed by Splunk. The fishbucket helps Splunk avoid indexing duplicate data by keeping track of file signatures and offsets. To reset the fishbucket for one source, the command `splunk cmd btprobe` can be used with the `-reset` option and the name of the source file.

Therefore, option C is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [Use btprobe to troubleshoot file monitoring - Splunk Documentation]

NEW QUESTION: 91

On the deployment server, administrators can map clients to server classes using client filters. Which of the following statements is accurate?

- A. The blacklist takes precedence over the whitelist.
- B. The whitelist takes precedence over the blacklist.
- C. Wildcards are not supported in any client filters.
- D. Machine type filters are applied before the whitelist and blacklist.

Answer: ([SHOW ANSWER](#))

<https://docs.splunk.com/Documentation/Splunk/8.2.1/Updating/Filterclients> Reference:

<https://community.splunk.com/t5/Getting-Data-In/Can-I-use-both-the-whitelist-AND-blacklist-for-the-same/td-p/390910>

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:

<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

Valid SPLK-1003 Dumps shared by EduDump.com for Helping Passing SPLK-1003 Exam! EduDump.com now offer the **newest SPLK-1003 exam dumps**, the EduDump.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com SPLK-1003 dumps with Test Engine here:

<https://www.edudump.com/exams/Splunk/SPLK-1003/premium/> (232 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)