Splunk.SPLK-1003.v2026-01-06.q93

Exam Code:	SPLK-1003
Exam Name:	Splunk Enterprise Certified Admin
Certification Provider:	Splunk
Free Question Number:	93
Version:	v2026-01-06
# of views:	116
# of Questions views:	930
https://www.freepdfdumps.com/Splunk.SPLK-1003.v2026-01-06.q93.html	

NEW QUESTION: 1
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: 2
You update a props. conf file while Splunk is running. You do not restart Splunk and you run this command: splunk btoo1 props list -debug. What will the output be?
A. list of all the configurations on-disk that Splunk contains.
B. A verbose list of all configurations as they were when splunkd started.
C. A list of props. conf configurations as they are on-disk along with a file path from which the configuration is located
D. A list of the current running props, conf configurations along with a file path from which the configuration was made
Answer: C (LEAVE A REPLY)
https://docs.splunk.com/Documentation/Splunk/8.0.1/Troubleshooting/Usebtooltotroubleshootconfigurations
"The btool command simulates the merging process using the on-disk conf files and creates a report showing the merged settings."
"The report does not necessarily represent what's loaded in memory. If a conf file change is made that requires a service restart, the btool report shows the change even though that change isn't active."

NEW QUESTION: 3
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: 4
When Splunk is integrated with LDAP, which attribute can be changed in the Splunk UI for an LDAP user?
A. Default app
B. LDAP group
C. Password
D. Username
Answer: A (LEAVE A REPLY)
When Splunk is integrated with LDAP, most of the user attributes are managed by the LDAP server and cannot be changed in the Splunk UI. However, one exception is the default app attribute, which specifies which app a user sees when they log in to Splunk. This attribute can be changed in the Splunk UI by editing the user settings. Therefore, option A is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [Configure Splunk to use LDAP and map groups - Splunk Documentation]

NEW QUESTION: 5
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: 6
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: A (LEAVE A REPLY)
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 acknolwedgment comes in." Reference
https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/AboutHECIDXAck

NEW QUESTION: 7
Where can scripts for scripted inputs reside on the host file system? (select all that apply)
A. $SFLUNK_HOME/bin/scripts
B. $SPLUNK_HOME/etc/apps/bin
C. $SPLUNK_HOME/etc/system/bin
D. $S?LUNK_HOME/etc/apps/<your_app>/bin_

Answer: A,C,D (LEAVE A REPLY)
"Where to place the scripts for scripted inputs. The script that you refer to in $SCRIPT can reside in only one of the following places on the host file system:
$SPLUNK_HOME/etc/system/bin
$SPLUNK_HOME/etc/apps/<your_App>/bin
$SPLUNK_HOME/bin/scripts
As a best practice, put your script in the bin/ directory that is nearest to the inputs.conf file that calls your script on the host file system."

NEW QUESTION: 8
What are the required stanza attributes when configuring the transforms. conf to manipulate or remove events?
A. REGEX, DEST. FORMAT
B. REGEX. SRC_KEY, FORMAT
C. REGEX, DEST_KEY, FORMAT
D. REGEX, DEST_KEY FORMATTING
Answer: C (LEAVE A REPLY)
REGEX = <regular expression>
* Enter a regular expression to operate on your data.
FORMAT = <string>
* NOTE: This option is valid for both index-time and search-time field extraction. Index-time field extraction configuration require the FORMAT settings. The FORMAT settings is optional for search-time field extraction configurations.
* This setting specifies the format of the event, including any field names or values you want to add.
DEST_KEY = <key>
* NOTE: This setting is only valid for index-time field extractions.
* Specifies where SPLUNK software stores the expanded FORMAT results in accordance with the REGEX match.

NEW QUESTION: 9
Which of the following methods will connect a deployment client to a deployment server? (select all that apply)
A. Run $SPLUNK_ROME/bin/ splunk set deploy-poll : from the command line of the deployment client.
B. Create and edit a deploymentserver . conf file in SSPLVNE{ on the deployment server.
C. Create and edit a deploymentclient . conf file in SSPLTJNE(EOME/etc/ system/local on the deployment client.
D. Run $SPLUNK ROME/bin/spiunk set deploy-poi i : from the command line of the deployment server.
Answer: A,C (LEAVE A REPLY)
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: 10
The following stanza is active in indexes.conf:
[cat_facts]
maxHotSpanSecs = 3600
frozenTimePeriodInSecs = 2630000
maxTota1DataSizeMB = 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 maxTota1DataSizeMB, do not affect the retention time of the events. They only affect the size and duration of the buckets that store the events.
References:1:Set a retirement and archiving policy - Splunk Documentation

NEW QUESTION: 11
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: C (LEAVE A REPLY)

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/Configureforwardingwithoutputs.conf
Reference: https://docs.splunk.com/Documentation/Forwarder/8.0.5/Forwarder/Enableareceiver

NEW QUESTION: 12
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: B (LEAVE A REPLY)
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: 13
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: 14
What options are available when creating custom roles? (select all that apply)
A. Restrict search terms
B. Whitelist search terms
C. Limit the number of concurrent search jobs
D. Allow or restrict indexes that can be searched.
Answer: A,C,D (LEAVE A REPLY)
https://docs.splunk.com/Documentation/SplunkCloud/8.2.2106/Admin/ConcurrentLimits
"Set limits for concurrent scheduled searches. You must have the edit_search_concurrency_all and edit_search_concurrency_scheduled capabilities to configure these settings."

NEW QUESTION: 15
What will the following inputs. conf stanza do?
[script://myscript . sh]
Interval=0
A. The script will run at the default interval of 60 seconds.
B. The script will not be run.
C. The script will be run only once for each time Splunk is restarted.
D. The script will be run. As soon as the script exits, Splunk restarts it.
Answer: C (LEAVE A REPLY)
* The inputs.conf file is used to configure inputs, distributed inputs such as forwarders, and file system monitoring in Splunk1.
* The [script://myscript.sh] stanza specifies a script input, which means that Splunk runs the script and indexes its output1.
* The interval setting determines how often Splunk runs the script. If the interval is set to 0, the script runs only once when Splunk starts up1. If the interval is omitted, the script runs at the default interval of 60 seconds2.
* Therefore, option C is correct, and the other options are incorrect.

NEW QUESTION: 16
Which authentication methods are natively supported within Splunk Enterprise? (select all that apply)
A. LDAP
B. SAML
C. RADIUS
D. Duo Multifactor Authentication
Answer: (SHOW ANSWER)
Reference:https://docs.splunk.com/Documentation/Splunk/7.3.1/Security/SetupuserauthenticationwithSplunk
Splunk authentication: Provides Admin, Power and User by default, and you can define your own roles using a list of capabilities. If you have an Enterprise license, Splunk authentication is enabled by default. See Set up user authentication with Splunk's built-in system for more information. LDAP: Splunk Enterprise supports authentication with its internal authentication services or your existing LDAP server. See Set up user authentication with LDAP for more information. Scripted authentication API: Use scripted authentication to integrate Splunk authentication with an external authentication system, such as RADIUS or PAM. See Set up user authentication with external systems for more information. Note: Authentication, including native authentication, LDAP, and scripted authentication, is not available in Splunk Free.

NEW QUESTION: 17
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: 18
Which Splunk component consolidates the individual results and prepares reports in a distributed environment?
A. Indexers
B. Forwarder
C. Search head
D. Search peers
Answer: (SHOW ANSWER)
https://docs.splunk.com/Documentation/Splunk/7.3.1/DistSearch/Howuserscancontroldistributedsearches
"From the user standpoint, specifying and running a distributed search is essentially the same as running any other search. Behind the scenes, the search head distributes the query to its search peers, and consolidates the results when presenting them to the user."

NEW QUESTION: 19
What is the difference between the two wildcards ... and - for the monitor stanza in inputs, conf?
A. ... is not supported in monitor stanzas
B. There is no difference, they are interchangable and match anything beyond directory boundaries.
C. * matches anything in that specific directory path segment, whereas ... recurses through subdirectories as well.
D. ... matches anything in that specific directory path segment, whereas - recurses through subdirectories as well.
Answer: C (LEAVE A REPLY)

https://docs.splunk.com/Documentation/Splunk/7.3.0/Data/Specifyinputpathswithwildcards The ellipsis wildcard searches recursively through directories and any number of levels of subdirectories to find matches.
If you specify a folder separator (for example, //var/log/.../file), it does not match the first folder level, only subfolders.
* The asterisk wildcard matches anything in that specific folder path segment.
Unlike ..., * does not recurse through subfolders.

NEW QUESTION: 20
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: C (LEAVE A REPLY)
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: 21
When deploying apps, which attribute in the forwarder management interface determines the apps that clients install?
A. App Class
B. Client Class
C. Server Class
D. Forwarder Class
Answer: C (LEAVE A REPLY)
<https://docs.splunk.com/Documentation/Splunk/8.0.6/Updating/Deploymentserverarchitecture> https://docs.splunk.com/Splexicon:Serverclass

NEW QUESTION: 22
Which of the following is accurate regarding the input phase?
A. Breaks data into events with timestamps.
B. Applies event-level transformations.
C. Fine-tunes metadata.
D. Performs character encoding.
Answer: D (LEAVE A REPLY)

https://docs.splunk.com/Documentation/Splunk/latest/Deploy/Datapipeline "The data pipeline segments in depth. INPUT - In the input segment, Splunk software consumes data. It acquires the raw data stream from its source, breaks it into 64K blocks, and annotates each block with some metadata keys. The keys can also include values that are used internally, such as the character encoding of the data stream, and values that control later processing of the data, such as the index into which the events should be stored. PARSING Annotating individual events with metadata copied from the source-wide keys. Transforming event data and metadata according to regex transform rules."

NEW QUESTION: 23
When configuring monitor inputs with whitelists or blacklists, what is the supported method of filtering the lists?
A. Slash notation
B. Regular expression
C. Irregular expression
D. Wildcard-only expression
Answer: B (LEAVE A REPLY)
https://docs.splunk.com/Documentation/Splunk/latest/Data
/Whitelistorblacklistspecificincomingdata#Include_or_exclude_specific_incoming_data

NEW QUESTION: 24
Which setting in indexes. conf allows data retention to be controlled by time?
A. maxDaysToKeep
B. moveToFrozenAfter
C. maxDataRetentionTime
D. frozenTimePeriodInSecs
Answer: D (LEAVE A REPLY)
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Setaretirementandarchivingpolicy

NEW QUESTION: 25
When are knowledge bundles distributed to search peers?
A. After a user logs in.
B. When Splunk is restarted.
C. When adding a new search peer.
D. When a distributed search is initiated.
Answer: D (LEAVE A REPLY)
"The search head replicates the knowledge bundle periodically in the background or when initiating a search.
" "As part of the distributed search process, the search head replicates and distributes its knowledge objects to its search peers, or indexers. Knowledge objects include saved searches, event types, and other entities used in searching accorss indexes. The search head needs to distribute this material to its search

peers so that they can properly execute queries on its behalf." Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/DistSearch/Whatsearchheadssend

NEW QUESTION: 26
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 incomple 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: A (LEAVE A REPLY)
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: 27
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: 28
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: A,C (LEAVE A REPLY)
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: 29
A Universal Forwarder has the following active stanza in inputs . conf:
[monitor: //var/log]
disabled = O
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: 30
What are the minimum required settings when creating a network input in Splunk?
A. Protocol, port number
B. Protocol, port, location
C. Protocol, username, port
D. Protocol, IP. port number
Answer: A (LEAVE A REPLY)
https://docs.splunk.com/Documentation/Splunk/8.0.5/Admin/Inputsconf
[tcp://<remote server>:<port>]
*Configures the input to listen on a specific TCP network port.
*If a <remote server> makes a connection to this instance, the input uses this stanza to configure itself.
*If you do not specify <remote server>, this stanza matches all connections on the specified port.
*Generates events with source set to "tcp:<port>", for example: tcp:514
*If you do not specify a sourcetype, generates events with sourcetype set to "tcp-raw"

NEW QUESTION: 31
Consider a company with a Splunk distributed environment in production. The Compliance Department wants to start using Splunk; however, they want to ensure that no one can see their reports or any other knowledge objects. Which Splunk Component can be added to implement this policy for the new team?
A. Indexer
B. Search head
C. Deployment server
D. Universal forwarder
Answer: B (LEAVE A REPLY)

NEW QUESTION: 32
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 typicallyrequired 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: 33
When working with an indexer cluster, what changes with the global precedence when comparing to a standalone deployment?
A. Nothing changes.
B. The peer-apps local directory becomes the highest priority.
C. The app local directories move to second in the priority list.
D. The system default directory' becomes the highest priority.
Answer: (SHOW ANSWER)
The app local directories move to second in the priority list. This is explained in the Splunk documentation, which states:
In a clustered environment, the precedence of configuration files changes slightly from that of a standalone deployment. The app local directories move to second in the priority list, after the peer-apps local directory. This means that any configuration files in the app local directories on the individual peers are overridden by configuration files of the same name and type in the peer-apps local directory on the master node.

NEW QUESTION: 34
How do you remove missing forwarders from the Monitoring Console?
A. By rescanning active forwarders.
B. By restarting Splunk.
C. By reloading the deployment server.
D. By rebuilding the forwarder asset table.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 35
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: B (LEAVE A REPLY)
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: 36
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: 37
Which forwarder is recommended by Splunk to use in a production environment?
A. Heavy forwarder
B. SSL forwarder
C. Lightweight forwarder
D. Universal forwarder
Answer: D (LEAVE A REPLY)

Reference:https://community.splunk.com/t5/Getting-Data-In/Splunk-forwarder/m-p/18009 The forwarder that is recommended by Splunk to use in a production environment is the universal forwarder.
The universal forwarder is a lightweight Splunk agent that forwards data to indexers or other forwarders. The universal forwarder has a small footprint and consumes minimal system resources. It also supports secure and reliable data forwarding with encryption and acknowledgement features. Therefore, option D is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [About forwarding and receiving data - Splunk Documentation]

NEW QUESTION: 38
Which Splunk component requires a Forwarder license?
A. Heavy forwarder
B. Search head
C. Heaviest forwarder
D. Universal forwarder
Answer: A (LEAVE A REPLY)

NEW QUESTION: 39
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: A,B,C (LEAVE A REPLY)
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 configuartion file on Splunk Enterprise indexer and heavy forwarder instances.

NEW QUESTION: 40
Given a forwarder with the following outputs.conf configuration:
[tcpout : mypartner]
Server = 145.188.183.184:9097
[tcpout : hfbank]
server = inputsl . mysplunkhfs . corp : 9997 , inputs2 . mysplunkhfs . corp : 9997 Which of the following is a true statement?
A. Data will continue to flow to hfbank if 145.1 g a) 183.184 : 9097 is unreachable.
B. Data is not encrypted to mypartner because 145.188 .183.184 : 9097 is specified by IP.
C. Data is encrypted to mypartner because 145.183.184 : 9097 is specified by IP.
D. Data will eventually stop flowing everywhere if 145.188.183.184 : 9097 is unreachable.
Answer: A (LEAVE A REPLY)

The outputs.conf file defines how forwarders send data to receivers1. You can specify some output configurations at installation time (Windows universal forwarders only) or the CLI, but most advanced configuration settings require that you edit outputs.conf1.
The [tcpout:...] stanza specifies a group of forwarding targets that receive data over TCP2. You can define multiple groups with different names and settings2.
The server setting lists one or more receiving hosts for the group, separated by commas2. If you specify multiple hosts, the forwarder load balances the data across them2.
Therefore, option A is correct, because the forwarder will send data to both inputsl.mysplunkhfs.corp:9997 and inputs2.mysplunkhfs.corp:9997, even if 145.188.183.184:9097 is unreachable.

NEW QUESTION: 41
Which of the following authentication types requires scripting in Splunk?
A. ADFS
B. LDAP
C. SAML
D. RADIUS
Answer: D (LEAVE A REPLY)
https://answers.splunk.com/answers/131127/scripted-authentication.html
Scripted Authentication: An option for Splunk Enterprise authentication. You can use an authentication system that you have in place (such as PAM or RADIUS) by configuring authentication.conf to use a script instead of using LDAP or Splunk Enterprise default authentication.

NEW QUESTION: 42
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: B,C (LEAVE A REPLY)
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: 43
What action could be taken to prevent a license warning with an ingest-based license?
A. Add a new license before midnight on the indexer(s).
B. Delete the data before midnight on the indexer(s).
C. Add a new license before midnight on the license manager.
D. Delete the data before midnight on the license manager.
Answer: C (LEAVE A REPLY)

In Splunk Enterprise, license warnings occur when the daily indexing volume exceeds the licensed quota. These warnings are tracked from midnight to midnight based on the system clock of the license manager. If the number of warnings surpasses the allowed threshold within a specified period, a license violation ensues, potentially restricting search capabilities.
To prevent a license warning from escalating to a violation, administrators have until midnight to address the issue. The recommended action is toadd a new licenseto the license manager before midnight. This increases the daily indexing volume quota, ensuring that the current day's data ingestion falls within the permissible limits.
It's important to note that deleting data from indexers or the license manager does not retroactively reduce the recorded license usage for the day. Once data is indexed, it contributes to the day's license volume, and its removal does not negate that contribution.
Reference:
About license violations - Splunk Documentation

NEW QUESTION: 44
In which scenario would a Splunk Administrator want to enable data integrity check when creating an index?
A. To ensure that data has not been tampered with for auditing and/or legal purposes
B. To ensure that user passwords have not been tampered with for auditing and/or legal purposes.
C. To ensure that configuration files have not been tampered with for auditing and/or legal purposes
D. To ensure that hot buckets are still open for writes and have not been forced to roll to a cold state
Answer: A (LEAVE A REPLY)

NEW QUESTION: 45
Which of the following are required when defining an index in indexes. conf? (select all that apply)
A. coldPath
B. homePath
C. frozenPath
D. thawedPath
Answer: (SHOW ANSWER)
homePath = $SPLUNK_DB/hatchdb/db
coldPath = $SPLUNK_DB/hatchdb/colddb
thawedPath = $SPLUNK_DB/hatchdb/thaweddb
https://docs.splunk.com/Documentation/Splunk/latest/Admin/Indexesconf
https://docs.splunk.com/Documentation/Splunk/7.3.1/Admin/Indexesconf#PER_INDEX_OPTIONS

NEW QUESTION: 46
Social Security Numbers (PII) data is found in log events, which is against company policy. SSN format is as follows: 123-44-5678.
Which configuration file and stanza pair will mask possible SSNs in the log events?
A. props.conf[mask-SSN]REX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"FORMAT = $1<SSN>###-

##-$2KEY = _raw
B. props.conf[mask-SSN]REGEX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"FORMAT = $1<SSN>###-##-$2DEST_KEY = _raw
C. transforms.conf[mask-SSN]REX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"FORMAT = $1<SSN>###-##-$2DEST_KEY = _raw
D. transforms.conf[mask-SSN]REGEX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"FORMAT = $1<SSN>###-##-$2DEST_KEY = _raw
Answer: D (LEAVE A REPLY)
because transforms.conf is the right configuration file to state the regex expression.https://docs.splunk.com/Documentation/Splunk/8.1.0/Admin/Transformsconf
Reference: https://community.splunk.com/t5/Archive/How-to-mask-SSN-into-our-logs-going-into-Splunk/tdp/
433035

NEW QUESTION: 47
In a customer managed Splunk Enterprise environment, what is the endpoint URI used to collect data?
A. services/ collector
B. services/ inputs ? raw
C. services/ data/ collector
D. data/ collector
Answer: (SHOW ANSWER)
The answer to your question is C. services/data/collector. This is the endpoint URI used to collect data in a customer managed Splunk Enterprise environment.According to the Splunk documentation1, "The HTTP Event Collector REST API endpoint is /services/data/collector.You can use this endpoint to send events to HTTP Event Collector on a Splunk Enterprise or Splunk Cloud Platform deployment." You can also use this endpoint to send events to a specific token or index1. For example, you can use the following curl command to send an event with the token 578254cc-05f5-46b5-957b-910d1400341a and the index main: curl -k https://localhost:8088/services/data/collector -H'Authorization: Splunk 578254cc-05f5-46b5-957b-910d1400341a'-d'{"index":"main","event":"Hello, world!"}'

NEW QUESTION: 48
Using the CLI on the forwarder, how could the current forwarder to indexer configuration be viewed?
A. splunk btool server list --debug

B. splunk list forward-indexer
C. splunk list forward-server
D. splunk btool indexes list --debug
Answer: C (LEAVE A REPLY)
Reference:https://community.splunk.com/t5/All-Apps-and-Add-ons/How-do-I-configure-a-Splunk-Forwarder- on-Linux/m-p/72078 The CLI command to view the current forwarder to indexer configuration is splunk list forward-server. This command displays the hostnames and port numbers of the indexers that the forwarder sends datato. Therefore, option C is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [Use CLI commands to manage your forwarders - Splunk Documentation]

NEW QUESTION: 49
Which Splunk forwarder type allows parsing of data before forwarding to an indexer?
A. Universal forwarder
B. Parsing forwarder
C. Heavy forwarder
D. Advanced forwarder
Answer: C (LEAVE A REPLY)

NEW QUESTION: 50
When using a directory monitor input, specific source type can be selectively overridden using which configuration file?
A. props.conf
B. sourcetypes.conf
C. transforms.conf
D. outputs.conf
Answer: A (LEAVE A REPLY)
Reference:https://docs.splunk.com/Documentation/SplunkCloud/latest/Data
/Bypassautomaticsourcetypeassignment
When using a directory monitor input, specific source types can be selectively overridden using props.conf. The props.conf file contains settings for parsing and indexing data, as well as search-time field extractions. The props.conf file can be used to assign or change source types for specific inputs using the sourcetype attribute. Therefore, option A is the correct answer. References: Splunk Enterprise Certified Admin | Splunk,
[Configure directory monitor inputs - Splunk Documentation]

NEW QUESTION: 51
Local user accounts created in Splunk store passwords in which file?
A. $ SFLUNK_HOME/etc/passwd
B. $ SFLUNK_HOME/etc/authentication
C. $ S?LUNK_HOME/etc/users/passwd.conf
D. $ SPLUNK HOME/etc/users/authentication.conf

Answer: (SHOW ANSWER)
Per the provided reference URL https://docs.splunk.com/Documentation/Splunk/7.3.1/Admin/User-seedconf
"To set the default username and password, place user-seed.conf in $SPLUNK_HOME/etc/system/local. You must restart Splunk to enable configurations. If the $SPLUNK_HOME/etc/passwd file is present, the settings in this file (user-seed.conf) are not used."

NEW QUESTION: 52
A configuration file in a deployed app needs to be directly edited. Which steps would ensure a successful deployment to clients?
A. Make the change in $SPLUNK HOME/etc/dep10yment apps/$appName/10ca1/ on the deployment server, and the change will be automatically sent to the deployment clients.
B. Make the change in $SPLUNK HOME /etc/apps/$appname/local/ on any of the deployment clients, and then run the command . / splunk reload deploy-server to push that change to the deployment server.
C. Make the change in $SPLUNK HOME/etc/dep10yment apps/$appName/10ca1/ on the deployment server, and then run $SPLUNK HOME/bin/sp1unk reload deploy-server.
D. Make the change in $SPLUNK HOME/etc/apps/$appName/defau1t on the deployment server, and it will be distributed down to the clients' own local versions.
Answer: (SHOW ANSWER)
According to the Splunk documentation1, to customize a configuration file, you need to create a new file with the same name in a local or app directory. Then, add the specific settings that you want to customize to the local configuration file. Never change or copy the configuration files in the default directory. The files in the default directory must remain intact and in their original location. The Splunk Enterprise upgrade process overwrites the default directory.
To deploy configuration files to deployment clients, you need to use the deployment server. The deployment server is a Splunk Enterprise instance that distributes content and updates to deployment clients2. The deployment server uses a directory called $SPLUNK_HOME/etc/deployment-apps to store the apps and configuration files that it deploys to clients2. To update the configuration files in this directory, you need to edit them manually and then run the command $SPLUNK_HOME/bin/sp1unk reload deploy-server to make the changes take effect2.
Therefore, option A is incorrect because it does not include the reload command. Option B is incorrect because it makes the change on a deployment client instead of the deployment server. Option D is incorrect because it changes the default directory instead of the local directory.
References: 1: How to edit a configuration file - Splunk Documentation 2: Deployment of configuration files - Splunk Community

NEW QUESTION: 53
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: B (LEAVE A REPLY)

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: 54

Which additional component is required for a search head cluster?

A. Deployer

B. Cluster Master

C. Monitoring Console

D. Management Console

Answer: A (LEAVE A REPLY)

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: 55

The following stanzas in inputs. conf are currently being used by a deployment client:

[udp: //145.175.118.177:1001

Connection_host = dns

sourcetype = syslog

Which of the following statements is true of data that is received via this input?

A. If Splunk is restarted, data will be queued and then sent when Splunk has restarted.

B. Local firewall ports do not need to be opened on the deployment client since the port is defined in inputs.conf.

C. The host value associated with data received will be the IP address that sent the data.

D. If Splunk is restarted, data may be lost.

Answer: D (LEAVE A REPLY)

This is because the input type is UDP, which is an unreliable protocol that does not guarantee delivery, order, or integrity of the data packets. UDP does not have any mechanism to resend or acknowledge the data packets, so if Splunk is restarted, any data that was in transit or in the buffer may be dropped and not indexed.

NEW QUESTION: 56

What type of Splunk license is pre-selected in a brand new Splunk installation?

A. Free license

B. Forwarder license

C. Enterprise trial license

D. Enterprise license

Answer: C (LEAVE A REPLY)

A Splunk Enterprise trial license gives you access to all the features of Splunk Enterprise for a limited period of time, usually 60 days1. After the trial period expires, you can either purchase a Splunk Enterprise license or switch to a Free license1.

A Splunk Enterprise Free license allows you to index up to 500 MB of data per day, but some features are disabled, such as authentication, distributed search, and alerting2. You can switch to a Free license at any time during the trial period or after the trial period expires1.

A Splunk Enterprise Forwarder license is used with forwarders, which are Splunk instances that forward data to other Splunk instances. A Forwarder license does not allow indexing or searching of data3. You can install a Forwarder license on any Splunk instance that you want to use as a forwarder4.

A Splunk Enterprise commercial end-user license is a license that you purchase from Splunk based on either data volume or infrastructure. This license gives you access to all the features of Splunk Enterprise within a defined limit of indexed data per day (volume-based license) or vCPU count (infrastructure license). You can purchase and install this license after the trial period expires or at any time during the trial period1.

NEW QUESTION: 57

Which of the following accurately describes HTTP Event Collector indexer acknowledgement?

A. It requires a separate channel provided by the client.

B. It is configured the same as indexer acknowledgement used to protect in-flight data.

C. It can be enabled at the global setting level.

D. It stores status information on the Splunk server.

Answer: A (LEAVE A REPLY)

https://docs.splunk.com/Documentation/Splunk/8.2.2/Data/AboutHECIDXAck

- Section: About channels and sending data

Sending events to HEC with indexer acknowledgment active is similar to sending them with the setting off. There is one crucial difference: when you have indexer acknowledgment turned on, you must specify a channel when you send events. The concept of a channel was introduced in HEC primarily to prevent a fast client from impeding the performance of a slow client. When you assign one channel per client, because channels are treated equally on Splunk Enterprise, one client can't affect another. You must include a matching channel identifier both when sending data to HEC in an HTTP request and when requesting acknowledgment that events contained in the request have been indexed. If you don't, you will receive the error message, "Data channel is missing." Each request that includes a token for which indexer acknowledgment has been enabled must include a channel identifier, as shown in the following example cURL statement, where <data> represents the event data portion of the request

NEW QUESTION: 58

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: 59
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: C,D (LEAVE A REPLY)
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_breaker2.
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: 60
Which of the following statements accurately describes using SSL to secure the feed from a forwarder?
A. It does not encrypt the certificate password.
B. SSL automatically compresses the feed by default.
C. It requires that the forwarder be set to compressed=true.
D. It requires that the receiver be set to compression=true.
Answer: A (LEAVE A REPLY)
Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/AboutsecuringyourSplunkconfigurationwithSSL

NEW QUESTION: 61
When would the following command be used?
A. To verify' the integrity of a local index.
B. To verify the integrity of a SmartStore index.
C. To verify the integrity of a SmartStore bucket.
D. To verify the integrity of a local bucket.
Answer: (SHOW ANSWER)
To verify the integrity of a local bucket. The command ./splunk check-integrity -bucketPath [bucket path] [-verbose] is used to verify the integrity of a local bucket by comparing the hashes stored in the l1Hashes and l2Hash files with the actual data in the bucket1. This command can help detect any tampering or corruption of the data.

NEW QUESTION: 62

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: D (LEAVE A REPLY)

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: 63

The Splunk administrator wants to ensure data is distributed evenly amongst the indexers. To do this, he runs the following search over the last 24 hours:

index=*

What field can the administrator check to see the data distribution?

A. host

B. index

C. linecount

D. splunk_server

Answer: D (LEAVE A REPLY)

https://docs.splunk.com/Documentation/Splunk/8.2.2/Knowledge/Usedefaultfields splunk_server The splunk server field contains the name of the Splunk server containing the event. Useful in a distributed Splunk environment. Example: Restrict a search to the main index on a server named remote. splunk_server=remote index=main 404

NEW QUESTION: 64

What happens when the same username exists in Splunk as well as through LDAP?

A. Splunk user is automatically deleted from authentication.conf.

B. LDAP settings take precedence.

C. Splunk settings take precedence.

D. LDAP user is automatically deleted from authentication.conf

Answer: (SHOW ANSWER)

Reference:https://docs.splunk.com/Documentation/SplunkCloud/8.2.2105/Security /SetupuserauthenticationwithLDAP

Splunk platform attempts native authentication first. If authentication fails outside of a local account that doesn't exist, there is no attempt to use LDAP to log in. This is adapted from precedence of Splunk authentication schema.

NEW QUESTION: 65
How does the Monitoring Console monitor forwarders?
A. By pulling internal logs from forwarders.
B. By using the forwarder monitoring add-on
C. With internal logs forwarded by forwarders.
D. With internal logs forwarded by deployment server.
Answer: C (LEAVE A REPLY)
Quoting the following Splunk URL reference https://docs.splunk.com/Documentation/Splunk/8.2.2/DMC/DMCprerequisites "Monitoring Console setup prerequisites. Forward internal logs (both $SPLUNK_HOME/car/log/splunk and $SPLUNK_HOME/var/log/introspection) to indexers from all other components. Without this step, many dashboards will lack data."

NEW QUESTION: 66
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: D (LEAVE A REPLY)
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: 67
Which Splunk component does a search head primarily communicate with?
A. Cluster master
B. Indexer
C. Deployment server
D. Forwarder
Answer: (SHOW ANSWER)

NEW QUESTION: 68
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: (SHOW ANSWER)
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 form 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

NEW QUESTION: 69
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: A (LEAVE A REPLY)
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.

NEW QUESTION: 70
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: A (LEAVE A REPLY)
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- forthe-same/td-p/390910

NEW QUESTION: 71
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-partysystemsd
"You can specify a SEDCMD configuration in props.conf to address data that contains characters that the third-party server cannot process. "

NEW QUESTION: 72
This file has been manually created on a universal forwarder
A new Splunk admin comes in and connects the universal forwarders to a deployment server and deploys the same app with a new
Which file is now monitored?
A. /var/log/maillog
B. /var/log/messages
C. none of the above
D. /var/log/maillog and /var/log/messages
Answer: (SHOW ANSWER)

NEW QUESTION: 73
Which setting allows the configuration of Splunk to allow events to span over more than one line?
A. SHOULD_LINEMERGE = true
B. BREAK_ONLY_BEFORE_DATE = true
C. BREAK_ONLY_BEFORE = <REGEX pattern>
D. SHOULD_LINEMERGE = false
Answer: (SHOW ANSWER)
The setting that allows the configuration of Splunk to allow events to span over more than one line is SHOULD_LINEMERGE. This setting determines whether consecutive lines from a single source should be concatenated into a single event. If SHOULD_LINEMERGE is set to true, Splunk will attempt to merge multiple lines into one event based on certain criteria, such as timestamps or regular expressions. Therefore, option A is the correct answer. References: Splunk Enterprise Certified Admin | Splunk, [Configure event line merging - Splunk Documentation]

NEW QUESTION: 74
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: C (LEAVE A REPLY)
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: 75
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: B (LEAVE A REPLY)
* 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: 76
Event processing occurs at which phase of the data pipeline?
A. Search
B. Indexing
C. Parsing
D. Input
Answer: C (LEAVE A REPLY)
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: 77
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%20l1Hashes%20file.
Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/Dataintegritycontrol

NEW QUESTION: 78
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: B (LEAVE A REPLY)
https://docs.splunk.com/Documentation/Splunk/latest/Data/Monitornetworkports

NEW QUESTION: 79
The universal forwarder has which capabilities when sending data? (select all that apply)
A. Sending alerts
B. Compressing data
C. Obfuscating/hiding data
D. Indexer acknowledgement
Answer: B,D (LEAVE A REPLY)
https://docs.splunk.com/Documentation/Splunk/8.0.1/Forwarding/Aboutforwardingandreceivingdata
https://docs.splunk.com/Documentation/Forwarder/8.1.1/Forwarder/Configureforwardingwithoutputs.conf#:~:text=compressed%3Dtrue%20This%20tells%20the,the%20forwarder%20sends%20raw%20data.

NEW QUESTION: 80
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: 81
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: A (LEAVE A REPLY)
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: 82
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: A,C (LEAVE A REPLY)
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: 83
Which forwarder type can parse data prior to forwarding?
A. Universal forwarder
B. Heaviest forwarder
C. Hyper forwarder
D. Heavy forwarder
Answer: (SHOW ANSWER)
https://docs.splunk.com/Documentation/Splunk/latest/Forwarding/Typesofforwarders

"A heavy forwarder parses data before forwarding it and can route data based on criteria such as source or type of event."

NEW QUESTION: 84
Which of the following lists the three phases of the Splunk Indexing process in order?
A. Ingest phaseLicensing phaseParsing phase
B. Sourcetype phaseIndex phaseWrite-to-disk phase
C. Input phaseParsing phaseIndexing phase
D. Ingest phaseTransforming phaseIndexing phase
Answer: C (LEAVE A REPLY)
The Splunk indexing process consists of three main phases: Input, Parsing, and Indexing. Understanding these phases is crucial for configuring data inputs and managing data flow within Splunk.
* Input Phase: Splunk receives data from various sources, such as files, network ports, or scripted inputs.
* Parsing Phase: Splunk breaks the data into individual events, applies transformations, and extracts timestamps.
* Indexing Phase: Splunk writes the parsed events to disk and creates indexes for efficient searching.
From the official Splunk documentation:
"The data pipeline consists of three main phases: input, parsing, and indexing."
- How the Splunk platform indexes data - Splunk Documentation
Therefore, the correct order of the indexing process is: Input phase # Parsing phase # Indexing phase.
Reference:
How the Splunk platform indexes data - Splunk Documentation

NEW QUESTION: 85
When deploying apps on Universal Forwarders using the deployment server, what is the correct component and location of the app before it is deployed?
A. On Universal Forwarder, $SPLUNK_HOME/etc/apps
B. On Deployment Server, $SPLUNK_HOME/etc/apps
C. On Deployment Server, $SPLUNK_HOME/etc/deployment-apps
D. On Universal Forwarder, $SPLUNK_HOME/etc/deployment-apps
Answer: C (LEAVE A REPLY)
The correct answer is C. On Deployment Server, $SPLUNK_HOME/etc/deployment-apps.
A deployment server is a Splunk Enterprise instance that acts as a centralized configuration manager for any number of other instances, called "deployment clients". A deployment client can be a universal forwarder, a non-clustered indexer, or a search head1.
A deployment app is a directory that contains any content that you want to download to a set of deployment clients. The content can include a Splunk Enterprise app, a set of Splunk Enterprise configurations, or other content, such as scripts, images, and supporting files2.
You create a deployment app by creating a directory for it on the deployment server. The default location is

$SPLUNK_HOME/etc/deployment-apps, but this is configurable through the repositoryLocation attribute in serverclass.conf. Underneath this location, each app must have its own subdirectory. The name of the subdirectory serves as the app name in the forwarder management interface2.
The other options are incorrect because:
A: On Universal Forwarder, $SPLUNK_HOME/etc/apps. This is the location where the deployment app resides after it is downloaded from the deployment server to the universal forwarder. It is not the location of the app before it is deployed2.
B: On Deployment Server, $SPLUNK_HOME/etc/apps. This is the location where the apps that are specific to the deployment server itself reside. It is not the location where the deployment apps for the clients are stored2.
D: On Universal Forwarder, $SPLUNK_HOME/etc/deployment-apps. This is not a valid location for any app on a universal forwarder. The universal forwarder does not act as a deployment server and does not store deployment apps3.

NEW QUESTION: 86
A user is assigned two roles with the following search filters. What is the user's applied search filter?
A.
B.
C.
D.
Answer: A (LEAVE A REPLY)
When a user is assignedmultiple rolesin Splunk and each has a defined srchFilter, Splunk combines these filters using alogical ANDoperation. This ensures that the user can only search within the intersection of constraints imposed by each role.
From Splunk Docs:
"If a user has multiple roles assigned and multiple roles specify srchFilter, Splunk softwareANDs the filters together."
- Source: Splunk Documentation - authorize.conf
Let's break it down:
* role_A specifies: sourcetype!=json AND index=main
* role_B specifies: sourcetype=csv
To evaluate the effective search filter for the user, Splunk willANDthe two conditions:
(sourcetype=csv) AND (sourcetype!=json AND index=main)
This means the user's search is limited to events where:
* sourcetype=csv (from role_B)
* sourcetype!=json AND index=main (from role_A)
Combining them together logically:
srchFilter = ((sourcetype=csv) AND (sourcetype!=json AND index=main))
This is exactly what is shown inOption A.
Reference:
authorize.conf - Splunk Admin Manual

NEW QUESTION: 87
A user recently installed an application to index NCINX access logs. After configuring the application, they realize that no data is being ingested. Which configuration file do they need to edit to ingest the access logs to ensure it remains unaffected after upgrade?
A. Option A
B. Option B
C. Option C
D. Option D
Answer: A (LEAVE A REPLY)
This option corresponds to the file path "$SPLUNK_HOME/etc/apps/splunk_TA_nginx/local/inputs.conf". This is the configuration file that the user needs to edit to ingest the NGINX access logs to ensure it remains unaffected after upgrade. This is explained in the Splunk documentation, which states:
The local directory is where you place your customized configuration files. The local directory is empty when you install Splunk Enterprise. You create it when you need to override or add to the default settings in a configuration file. The local directory is never overwritten during an upgrade.

NEW QUESTION: 88
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/Monitorfilesanddirectorieswithinputs.conf

NEW QUESTION: 89
What is the correct order of steps in Duo Multifactor Authentication?
A. 1 Request Login2. Connect to SAML server3 Duo MFA4 Create User session5 Authentication Granted 6. Log into Splunk
B. 1. Request Login 2 Duo MFA3. Authentication Granted 4 Connect to SAML server5. Log into Splunk6. Create User session
C. 1 Request Login2 Check authentication / group mapping3 Authentication Granted4. Duo MFA5. Create User session6. Log into Splunk
D. 1 Request Login 2 Duo MFA3. Check authentication / group mapping4 Create User session5.Authentication Granted6 Log into Splunk
Answer: C (LEAVE A REPLY)
Using the provided DUO/Splunk reference URLhttps://duo.com/docs/splunk

Scroll down to the Network Diagram section and note the following 6 similar steps
1 - SPlunk connection initiated
2 - Primary authentication
3 - Splunk connection established to Duo Security over TCP port 443
4 - Secondary authentication via Duo Security's service
5 - Splunk receives authentication response
6 - Splunk session logged in.

NEW QUESTION: 90
Which of the following enables compression for universal forwarders in outputs. conf ?
A.
B.
C.
D.
Answer: B (LEAVE A REPLY)
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: 91
Which of the following statements describe deployment management? (select all that apply)
A. Requires an Enterprise license
B. Is responsible for sending apps to forwarders.
C. Once used, is the only way to manage forwarders
D. Can automatically restart the host OS running the forwarder.
Answer: (SHOW ANSWER)
https://docs.splunk.com/Documentation/Splunk/8.2.2/Admin/Distdeploylicenses#:~:text=License%20requirements,do%20not%20index%20external%20data.
"All Splunk Enterprise instances functioning as management components needs access to an Enterprise license. Management components include the deployment server, the indexer cluster manager node, the search head cluster deployer, and the monitoring console."
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: 92
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: B (LEAVE A REPLY)
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: 93
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/8.0.5/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

Valid SPLK-1003 Dumps shared by Actual4test.com for Helping Passing SPLK-1003 Exam! Actual4test.com now offer the **newest SPLK-1003 exam dumps**, the Actual4test.com SPLK-1003 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com SPLK-1003 dumps with Test Engine here: https://www.actual4test.com/SPLK-1003_examcollection.html (**203** Q&As Dumps, **30%OFF** **Special Discount:** **Freepdfdumps**)