Splunk.SPLK-1003.v2023-11-20.q78

Exam Code:	SPLK-1003
Exam Name:	Splunk Enterprise Certified Admin
Certification Provider:	Splunk
Free Question Number:	78
Version:	v2023-11-20
# of views:	2935
# of Questions views:	780
https://www.freepdfdumps.com/Splunk.SPLK-1003.v2023-11-20.q78.html	

NEW QUESTION: 1
Which Splunk component would one use to perform line breaking prior to indexing?
A. Heavy Forwarder
B. Universal Forwarder
C. Search head
D. This can only be done at the indexing layer.
Answer: A (LEAVE A REPLY)
Explanation
According to the Splunk documentation1, a heavy forwarder is a Splunk Enterprise instance that can parse and filter data before forwarding it to an indexer. A heavy forwarder can perform line breaking, which is the process of splitting incoming data into individual events based on a set of rules2. A heavy forwarder can also apply other transformations to the data, such as field extractions, event type matching, or masking sensitive data3.

NEW QUESTION: 2
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)
Explanation
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: 3

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)
Explanation
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: 4
Which Splunk component requires a Forwarder license?
A. Heavy forwarder
B. Universal forwarder
C. Heaviest forwarder
D. Search head
Answer: A (LEAVE A REPLY)

NEW QUESTION: 5
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)
Explanation
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."

NEW QUESTION: 6
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)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.0.1/Forwarding/Aboutforwardingandreceivingdata
https://docs.splunk.com/Documentation/Forwarder/8.1.1/Forwarder/Configureforwardingwithoutputs.conf#:~:tex

NEW QUESTION: 7
Which Splunk forwarder has a built-in license?
A. Cloud forwarder
B. Light forwarder
C. Universal forwarder
D. Heavy forwarder
Answer: C (LEAVE A REPLY)

NEW QUESTION: 8
Immediately after installation, what will a Universal Forwarder do first?
A. Automatically detect any indexers in its subnet and begin routing data.
B. Begin generating internal Splunk logs.
C. Begin reading local files on its server.
D. Send an email to the operator that the installation process has completed.
Answer: B (LEAVE A REPLY)
Explanation
Immediately after installation, a universal forwarder will start generating internal Splunk logs that contain information about its own operation, such as configuration changes, data inputs, and forwarding activities1. These logs are stored in the $SPLUNK_HOME/var/log/splunk directory on the universal forwarder machine1. The universal forwarder will not automatically detect any indexers in its subnet and begin routing data, as it needs to be configured with the IP address and port number of the indexer or the deployment server2. The universal forwarder will not begin reading local files on its server, as it needs to be configured with the data inputs that specify which files or directories to monitor2. The universal forwarder will not send an email to the operator that the installation process has completed, as this is not a default behavior of the universal forwarder and would require additional configuration3.

NEW QUESTION: 9
The priority of layered Splunk configuration files depends on the file's:
A. Owner
B. Weight
C. Context
D. Creation time
Answer: C (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/7.3.0/Admin/Wheretofindtheconfigurationfiles

"To determine the order of directories for evaluating configuration file precendence, Splunk software considers each file's context. Configuration files operate in either a global context or in the context of the current app and user"

NEW QUESTION: 10
The LINE_BREAKER attribute is configured in which configuration file?
A. inpucs.conf
B. transforms.conf
C. indexes.conf
D. props.conf
Answer: D (LEAVE A REPLY)

NEW QUESTION: 11
Which of the methods listed below supports muti-factor authentication?
A. Lightweight Directory Access Protocol (LDAP)
B. Security Assertion Markup Language (SAML)
C. Single Sign-on (SSO)
D. OpenID
Answer: B (LEAVE A REPLY)
Explanation
SAML is an open standard for exchanging authentication and authorization data between parties, especially between an identity provider and a service provider1. SAML supports multi-factor authentication by allowing the identity provider to require the user to present two or more factors of evidence to prove their identity2. For example, the user may need to enter a password and a one-time code sent to their phone, or scan their fingerprint and face.

NEW QUESTION: 12
In this source definition the MAX_TIMESTAMP_LOOKHEAD is missing. Which value would fit best?
Event example:
A. MAX_TIMESTAMP_L0CKAHEAD = 5
B. MAX_TIMESTAMP_LOOKAHEAD - 10
C. MAX_TIMESTAMF_LOOKHEAD = 20
D. MAX TIMESTAMP LOOKAHEAD - 30
Answer: (SHOW ANSWER)
Explanation
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: 13
What is the name of the object that stores events inside of an index?

A. Container
B. Bucket
C. Data layer
D. Indexer
Answer: B (LEAVE A REPLY)
Explanation
A bucket is the object that stores events inside of an index. According to the Splunk documentation1, "An index is a collection of directories, also called buckets, that contain index files. Each bucket represents a specific time range." A bucket can be in one of several states, such as hot, warm, cold, frozen, or thawed1. Buckets are managed by indexers or clusters of indexers1.

NEW QUESTION: 14
Which of the following is an appropriate description of a deployment server in a non-cluster environment?
A. Allows management of remote Splunk instances, requires Enterprise license, handles job of sending configurations, can automatically restart remote Splunk instances.
B. Allows management of remote Splunk instances, requires no license, handles job of sending configurations, can automatically restart remote Splunk instances.
C. Allows management of remote Splunk instances, requires Enterprise license, handles job of sending configurations, can manually restart remote Splunk instances.
D. Allows management of local Splunk instances, requires Enterprise license, handles job of sending configurations packaged as apps. can automatically restart remote Splunk instances.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 15
Which of the following is valid distribute search group?
A)
B)
C)
D)
A. option A
B. Option D
C. Option C
D. Option B
Answer: (SHOW ANSWER)

NEW QUESTION: 16
What conf file needs to be edited to set up distributed search groups?
A. props.conf
B. search.conf
C. distsearch.conf
D. distibutedsearch.conf

Answer: C (LEAVE A REPLY)
Explanation
"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"

NEW QUESTION: 17
Which layers are involved in Splunk configuration file layering? (select all that apply)
A. App context
B. User context
C. Global context
D. Forwarder context
Answer: A,B,C (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/latest/Admin/Wheretofindtheconfigurationfiles To determine the order of directories for evaluating configuration file precedence, Splunk software considers each file's context. Configuration files operate in either a global context or in the context of the current app and user: Global. Activities like indexing take place in a global context. They are independent of any app or user.
For example, configuration files that determine monitoring or indexing behavior occur outside of the app and user context and are global in nature. App/user. Some activities, like searching, take place in an app or user context. The app and user context is vital to search-time processing, where certain knowledge objects or actions might be valid only for specific users in specific apps.

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: C (LEAVE A REPLY)
Explanation
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
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: C (LEAVE A REPLY)
Explanation
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: 20
How do you remove missing forwarders from the Monitoring Console?
A. By restarting Splunk.
B. By rescanning active forwarders.
C. By rebuilding the forwarder asset table.
D. By reloading the deployment server.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 21
When indexing a data source, which fields are considered metadata?
A. host, raw, sourcetype
B. sourcetype, source, host
C. source, host, time
D. time, sourcetype, source
Answer: B (LEAVE A REPLY)

NEW QUESTION: 22
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: C (LEAVE A REPLY)
Explanation
http://www.splunk.com/view/SP-CAAAGPR

NEW QUESTION: 23
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)
Explanation
https://docs.splunk.com/Documentation/Splunk/latest/Data/Monitornetworkports

NEW QUESTION: 24
Which of the following monitor inputs stanza headers would match all of the following files?
/var/log/www1/secure.log
/var/log/www/secure.l
/var/log/www/logs/secure.logs
/var/log/www2/secure.log
A. [monitor:///var/log/.../secure.*
B. [monitor:///var/log/www1/secure.log]
C. [monitor:///var/log/www*/secure.*]
D. [monitor:///var/log/www1/secure.*]
Answer: (SHOW ANSWER)

NEW QUESTION: 25
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)
Explanation
<https://docs.splunk.com/Documentation/Splunk/8.0.6/Updating/Deploymentserverarchitecture>
https://docs.splunk.com/Splexicon:Serverclass

NEW QUESTION: 26
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)
Explanation
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: 27
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: (SHOW ANSWER)

Explanation
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: 28
Which feature of Splunk's role configuration can be used to aggregate multiple roles intended for groups of users?
A. Linked roles
B. Grantable roles
C. Role federation
D. Role inheritance
Answer: D (LEAVE A REPLY)
Explanation
You can have a role inherit certain properties from one or more existing role
https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/Aboutusersandroles

NEW QUESTION: 29
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)
Explanation
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."

NEW QUESTION: 30
What are the values for host and index for [stanza1] used by Splunk during index time, given the following configuration files?
A. host=server1
index=unixinfo
B. host=server1
index=searchinfo
C. host=searchsvr1
index=searchinfo

D. host=unixsvr1
index=unixinfo
Answer: (SHOW ANSWER)
Explanation
- 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

NEW QUESTION: 31
Immediately after installation, what will a Universal Forwarder do first?
A. Automatically detect any indexers in its subnet and begin routing data.
B. Begin reading local files on its server.
C. Begin generating internal Splunk logs.
D. Send an email to the operator that the installation process has completed.
Answer: C (LEAVE A REPLY)
Explanation
Begin generating internal Splunk logs. Immediately after installation, a Universal Forwarder will start generating internal Splunk logs that contain information about its own operation, such as startup and shutdown events, configuration changes, data ingestion, and forwarding activities1. These logs are stored in the $SPLUNK_HOME/var/log/splunk directory on the Universal Forwarder machine2.

NEW QUESTION: 32
After an Enterprise Trial license expires, it will automatically convert to a Free license. How many days is an Enterprise Trial license valid before this conversion occurs?
A. 7 days
B. 60 days
C. 14 days
D. 90 days
Answer: B (LEAVE A REPLY)

NEW QUESTION: 33
Which of the following indexes come pre-configured with Splunk Enterprise? (select all that apply)

A. _license
B. _Internal
C. _external
D. _thefishbucket
Answer: B,D (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.0.5/Indexer/Howindexingworks

NEW QUESTION: 34
Which optional configuration setting in inputs .conf allows you to selectively forward the data to specific indexer(s)?
A. _TCP_ROUTING
B. _INDEXER_LIST
C. _INDEXER_GROUP
D. _INDEXER ROUTING
Answer: A (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/7.0.3/Forwarding/Routeandfilterdatad#Perform_selective_indexi
Specifies a comma-separated list of tcpout group names. Use this setting to selectively forward your data to specific indexers by specifying the tcpout groups that the forwarder should use when forwarding the data.
Define the tcpout group names in the outputs.conf file in [tcpout:<tcpout_group_name>] stanzas. The groups present in defaultGroup in [tcpout] stanza in the outputs.conf file.

NEW QUESTION: 35
What is the default value of LINE_BREAKER?
A. ([\r\n]+)
B. \r\n
C. \r+\n+
D. (\r\n+)
Answer: A (LEAVE A REPLY)

NEW QUESTION: 36
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)
Explanation

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

NEW QUESTION: 37
What is the command to reset the fishbucket for one source?
A. splunk clean eventdata -index _thefishbucket
B. splunk btool fishbucket reset <source>
C. splunk cmd btprobe -d SPLUNK_HOME/var/lib/splunk/fishbucket/splunk_private_db --file <source> --reset
D. rm -r ~/splunkforwarder/var/lib/splunk/fishbucket
Answer: C (LEAVE A REPLY)

NEW QUESTION: 38
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)
Explanation
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: 39
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)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.2.2/Admin/Distdeploylicenses#:~:text=License%20requiremen

"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: 40
What event-processing pipelines are used to process data for indexing? (select all that apply)
A. Parsing pipeline
B. Indexing pipeline
C. Typing pipeline
D. fifo pipeline
Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 41
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. Client Applications tab in forwarder management interface or clientapps.conf.
D. Server Classes tab in forwarder management interface or serverclass.conf.
Answer: (SHOW ANSWER)

NEW QUESTION: 42
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: A,B,D (LEAVE A REPLY)
Explanation
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: 43
In which scenario would a Splunk Administrator want to enable data integrity check when creating an index?
A. To ensure that hot buckets are still open for writes and have not been forced to roll to a cold state
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 data has not been tampered with for auditing and/or legal purposes
Answer: (SHOW ANSWER)

NEW QUESTION: 44
Which of the following describes a Splunk deployment server?
A. A Splunk Forwarder that deploys data to multiple indexers.
B. A Splunk app installed on a Splunk Enterprise server.
C. A Splunk Enterprise server that distributes apps.
D. A server that automates the deployment of Splunk Enterprise to remote servers.
Answer: C (LEAVE A REPLY)
Explanation
A Splunk deployment server is a system that distributes apps, configurations, and other assets to groups of Splunk Enterprise instances. You can use it to distribute updates to most types of Splunk Enterprise components: forwarders, non-clustered indexers, and search heads2.
A Splunk deployment server is available on every full Splunk Enterprise instance. To use it, you must activate it by placing at least one app into %SPLUNK_HOME%\etc\deployment-apps on the host you want to act as deployment server3.
A Splunk deployment server maintains the list of server classes and uses those server classes to determine what content to distribute to each client. A server class is a group of deployment clients that share one or more defined characteristics1.
A Splunk deployment client is a Splunk instance remotely configured by a deployment server. Deployment clients can be universal forwarders, heavy forwarders, indexers, or search heads. Each deployment client belongs to one or more server classes1.
A Splunk deployment app is a set of content (including configuration files) maintained on the deployment server and deployed as a unit to clients of a server class. A deployment app can be an existing Splunk Enterprise app or one developed solely to group some content for deployment purposes1.
Therefore, option C is correct, and the other options are incorrect.

NEW QUESTION: 45
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 ga. 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: (SHOW ANSWER)

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: 46
How can native authentication be disabled in Splunk?
A. Set SPLUNK_AUTHENTICATION=false in splunk-launch.conf
B. Create an empty $SPLUNK_HOME/etc/passwd file
C. Set nativeAuthentication=false in authentication.conf
D. Remove the $SPLUNK_HOME/etc/passwd file
Answer: B (LEAVE A REPLY)

NEW QUESTION: 47
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: D (LEAVE A REPLY)
Explanation
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: 48
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)
Explanation
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: 49
Which setting in indexes. conf allows data retention to be controlled by time?
A. maxDaysToKeep
B. moveToFrozenAfter
C. maxDataRetentionTime
D. frozenTimePeriodlnSecs
Answer: D (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Setaretirementandarchivingpolicy

NEW QUESTION: 50
In which Splunk configuration is the SEDCMD used?
A. props, conf
B. inputs.conf
C. indexes.conf
D. transforms.conf
Answer: (SHOW ANSWER)
Explanation
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: 51
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)
Explanation
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: 52
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)
Explanation
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

NEW QUESTION: 53
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)
Explanation
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: 54
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)
Explanation
https://docs.splunk.com/Documentation/Splunk/latest/Data/Whitelistorblacklistspecificincomingdata#Include_or

NEW QUESTION: 55
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)
Explanation
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: 56
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)
Explanation
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: 57
In inputs. conf, which stanza would mean Splunk was only reading one local file?
A. [read://opt/log/crashlog/Jan27crash.txt]
B. [monitor::/ opt/log/crashlog/Jan27crash.txt]
C. [monitor:/// opt/log/]
D. [monitor:/// opt/log/ crashlog/Jan27crash.txt]
Answer: B (LEAVE A REPLY)
Explanation

[monitor::/opt/log/crashlog/Jan27crash.txt]. This stanza means that Splunk is monitoring a single local file named Jan27crash.txt in the /opt/log/crashlog/ directory1. The monitor input type is used to monitor files and directories for changes and index any new data that is added2.

NEW QUESTION: 58
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/messages
B. /var/log/maillog and /var/log/messages
C. /var/log/maillog
D. none of the above
Answer: C (LEAVE A REPLY)

NEW QUESTION: 59
Using SEDCMD in props.conf allows raw data to be modified. With the given event below, which option will mask the first three digits of the AcctID field resulting output: [22/Oct/2018:15:50:21] VendorID=1234 Code=B AcctID=xxx5309 Event:
[22/Oct/2018:15:50:21] VendorID=1234 Code=B AcctID=xxx5309
A. SEDCMD-1acct = s/VendorID=\d{3}(\d{4})/VendorID=xxx/g
B. SEDCMD-xxxAcct = s/AcctID=\d{3}(\d{4})/AcctID=xxx/g
C. SEDCMD-1acct = s/AcctID=\d{3}(\d{4})/AcctID=\1xxx/g
D. SEDCMD-1acct = s/AcctID=\d{3}(\d{4})/AcctID=xxx\1/g
Answer: (SHOW ANSWER)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.2.2/Data/Anonymizedata
Scrolling down to the section titled "Define the sed script in props.conf shows the correct syntax of an example which validates that the number/character /1 immediately preceded the /g

NEW QUESTION: 60
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. Search head
B. Universal forwarder
C. Deployment server
D. Indexer
Answer: A (LEAVE A REPLY)

NEW QUESTION: 61

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: (SHOW ANSWER)
Explanation
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: 62
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)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/ConsiderationsfordecidinghowtomonitorWindowsdat "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: 63
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: C (LEAVE A REPLY)
Explanation
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: 64
What is a role in Splunk? (select all that apply)
A. A classification that determines what capabilities a user has.
B. A classification that determines if a Splunk server can remotely control another Splunk server.
C. A classification that determines what functions a Splunk server controls.
D. A classification that determines what indexes a user can search.
Answer: A,D (LEAVE A REPLY)
Explanation
A role in Splunk is a classification that determines what capabilities and indexes a user has. A capability is a permission to perform a specific action or access a specific feature on the Splunk platform1. An index is a collection of data that Splunk software processes and stores2. By assigning roles to users, you can control what they can do and what data they can access on the Splunk platform.
Therefore, the correct answers are A and D. A role in Splunk determines what capabilities and indexes a user has. Option B is incorrect because Splunk servers do not use roles to remotely control each other. Option C is incorrect because Splunk servers use instances and components to determine what functions they control3.

References: 1: Define roles on the Splunk platform with capabilities - Splunk Documentation 2: About indexes and indexers - Splunk Documentation 3: Splunk Enterprise components - Splunk Documentation

NEW QUESTION: 65
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)
Explanation
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: 66
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: C (LEAVE A REPLY)
Explanation
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: 67
During search time, which directory of configuration files has the highest precedence?
A. $SFLUNK_KOME/etc/system/local
B. $SPLUNK_KCME/etc/system/default
C. $SPLUNK_HCME/etc/apps/app1/local
D. $SPLUNK HCME/etc/users/admin/local
Answer: D (LEAVE A REPLY)
Explanation

Adding further clarity and quoting same Splunk reference URL from @giubal"
"To keep configuration settings consistent across peer nodes, configuration files are managed from the cluster master, which pushes the files to the slave-app directories on the peer nodes. Files in the slave-app directories have the highest precedence in a cluster peer's configuration. Here is the expanded precedence order for cluster peers:
1.Slave-app local directories -- highest priority
2. System local directory
3. App local directories
4. Slave-app default directories
5. App default directories
6. System default directory --lowest priority

NEW QUESTION: 68
User role inheritance allows what to be inherited from the parent role? (select all that apply)
A. Parents
B. Capabilities
C. Index access
D. Search history
Answer: (SHOW ANSWER)
Explanation
https://docs.splunk.com/Documentation/Splunk/latest/Security/Aboutusersandroles#Role_inheritance
https://docs.splunk.com/Documentation/Splunk/7.3.1/Security/Aboutusersandroles#How_users_inherit_capabilit

NEW QUESTION: 69
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)
Explanation
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: 70
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: 71
An index stores its data in buckets. Which default directories does Splunk use to store buckets? (Choose all that apply.)
A. colddb
B. frozendb
C. db
D. bucketdb
Answer: (SHOW ANSWER)

NEW QUESTION: 72
In which phase do indexed extractions in props.conf occur?
A. Inputs phase
B. Parsing phase
C. Indexing phase
D. Searching phase
Answer: B (LEAVE A REPLY)
Explanation
The following items in the phases below are listed in the order Splunk applies them (ie LINE_BREAKER occurs before TRUNCATE).
Input phase
inputs.conf
props.conf
CHARSET
NO_BINARY_CHECK
CHECK_METHOD
CHECK_FOR_HEADER (deprecated)
PREFIX_SOURCETYPE
sourcetype

wmi.conf
regmon-filters.conf
Structured parsing phase
props.conf
INDEXED_EXTRACTIONS, and all other structured data header extractions
Parsing phase
props.conf
LINE_BREAKER, TRUNCATE, SHOULD_LINEMERGE, BREAK_ONLY_BEFORE_DATE, and all other line merging settings TIME_PREFIX, TIME_FORMAT, DATETIME_CONFIG (datetime.xml), TZ, and all other time extraction settings and rules TRANSFORMS which includes per-event queue filtering, per-event index assignment, per-event routing SEDCMD MORE_THAN, LESS_THAN transforms.conf stanzas referenced by a TRANSFORMS clause in props.conf LOOKAHEAD, DEST_KEY, WRITE_META, DEFAULT_VALUE, REPEAT_MATCH

NEW QUESTION: 73
How would you configure your distsearch conf to allow you to run the search below?
sourcetype=access_combined status=200 action=purchase splunk_setver_group=HOUSTON A)
B)
C)
D)
A. option A
B. Option B
C. Option C
D. Option D
Answer: C (LEAVE A REPLY)
Explanation
https://docs.splunk.com/Documentation/Splunk/8.0.3/DistSearch/Distributedsearchgroups

NEW QUESTION: 74
Using the CLI on the forwarder, how could the current forwarder to indexer configuration be viewed?
A. splunk list forward-server
B. splunk list forward-indexer
C. splunk btool indexes list --debug
D. splunk btool server list --debug
Answer: (SHOW ANSWER)

NEW QUESTION: 75
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: D (LEAVE A REPLY)
Explanation
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: 76
Which of the following Splunk components require a separate installation package?
A. Universal forwarder
B. Deployment server
C. License master
D. Heavy forwarder
Answer: (SHOW ANSWER)

NEW QUESTION: 77
Which of the following are reasons to create separate indexes? (Choose all that apply.)
A. File organization.
B. Increase number of users.
C. Restrict user permissions.
D. Different retention times.
Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 78
Which of the following is a valid distributed search group?
A. [distributedSearch:Paris] default = false servers = server1, server2
B. [searchGroup:Paris] default = false servers = server1:8089, server2:8089
C. [searchGroup:Paris] default = false servers = server1:9997, server2:9997
D. [distributedSearch:Paris] default = false servers = server1:8089; server2:8089
Answer: (SHOW ANSWER)
Explanation
https://docs.splunk.com/Documentation/Splunk/9.0.0/DistSearch/Distributedsearchgroups