Number: SPLK-1003
Passing Score: 800
Time Limit: 120
File Version: 5.0

Exam Code: SPLK-1003

Exam Name: Splunk Enterprise Certified Admin

Exam A

QUESTION 1
Consider the following stanza in inputs.conf:

```
[script:///opt/splunk/etc/apps/search/bin/lister.sh
disabled = 0
interval = 60.0
sourcetype = lister
```

What will the value of the source filed be for events generated by this scripts input?

A. /opt/splunk/ecc/apps/search/bin/liscer.sh
B. unknown
C. liscer
D. liscer.sh

Correct Answer: A
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.2.2/Admin/Inputsconf
-Scroll down to source = <string>
*Default: the input file path

QUESTION 2
Which of the following applies only to Splunk index data integrity check?

A. Lookup table
B. Summary Index
C. Raw data in the index
D. Data model acceleration

Correct Answer: C
Section:

QUESTION 3
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.

Correct Answer: D
Section:
Explanation:
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.

QUESTION 4
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.

Correct Answer: C
Section:
Explanation:
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.

QUESTION 5
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

Correct Answer: D
Section:
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.

QUESTION 6
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.

Correct Answer: C
Section:
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.

QUESTION 7
Which of the following types of data count against the license daily quota?

A. Replicated data
B. splunkd logs
C. Summary index data
D. Windows internal logs

Correct Answer: D
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.0.3/Admin/Distdeploylicenses#Clustered_deployments_and_licensing_issues
Reference: https://community.splunk.com/t5/Deployment-Architecture/License-usage-in-Indexer-Cluster/m-p/493548

QUESTION 8
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

Correct Answer: D
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/9.0.0/DistSearch/Distributedsearchgroups

QUESTION 9
Where are license files stored?

A. $SPLUNK_HOME/etc/secure
B. $SPLUNK_HOME/etc/system
C. $SPLUNK_HOME/etc/licenses
D. $SPLUNK_HOME/etc/apps/licenses

Correct Answer: C

Section:

QUESTION 10
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 configuration files have not been tampered with for auditing and/or legal purposes
C. To ensure that user passwords 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

Correct Answer: D
Section:

QUESTION 11
Which is a valid stanza for a network input?

A. [udp://172.16.10.1:9997]
connection = dns
sourcetype = dns
B. [any://172.16.10.1:10001]
connection_host = ip
sourcetype = web
C. [tcp://172.16.10.1:9997]
connection_host = web
sourcetype = web
D. [tcp://172.16.10.1:10001]
connection_host = dns
sourcetype = dns

Correct Answer: D
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.1.1/Data/Monitornetworkports
Reference: https://docs.splunk.com/Documentation/SplunkCloud/8.0.2006/Data/Bypassautomaticsourcetypeassignment

QUESTION 12
Which additional component is required for a search head cluster?

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

Correct Answer: A
Section:
Explanation:
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.

QUESTION 13
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.

Correct Answer: D
Section:
Explanation:
"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

QUESTION 14
Assume a file is being monitored and the data was incorrectly indexed to an exclusive index. The index is cleaned and now the data must be reindexed. What other index must be cleaned to reset the input checkpoint information for that file?

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

Correct Answer: D
Section:
Explanation:
--reset Reset the fishbucket for the given key or file in the btree. Resetting the checkpoint for an active monitor input reindexes data, resulting in increased license use.
https://docs.splunk.com/Documentation/Splunk/8.1.1/Troubleshooting/CommandlinetoolsforusewithSupport
Reference: http://docshare02.docshare.tips/files/4773/47733589.pdf

QUESTION 15
After configuring a universal forwarder to communicate with an indexer, which index can be checked via the Splunk Web UI for a successful connection?

A. index=main
B. index=test
C. index=summary
D. index=_internal

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/Validateyourconfiguration

QUESTION 16
Which of the following are available input methods when adding a file input in Splunk Web? (Choose all that apply.)

A. Index once.
B. Monitor interval.
C. On-demand monitor.
D. Continuously monitor.

Correct Answer: A, D
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.2.2/Data/Howdoyouwanttoadddata
The fastest way to add data to your Splunk Cloud instance or Splunk Enterprise deployment is to use Splunk Web. After you access the Add Data page, choose one of three options for getting data into your Splunk platform deployment with Splunk Web: (1) Upload, (2) Monitor, (3) Forward The Upload option lets you upload a file or archive of files for indexing. When you choose Upload option, Splunk Web opens the upload process page. Monitor. For Splunk Enterprise installations, the Monitor option lets you monitor one or more files, directories, network streams, scripts, Event Logs (on Windows hosts only), performance metrics, or any other type of machine data that the Splunk Enterprise instance has access to.

QUESTION 17
What is the valid option for a [monitor] stanza in inputs.conf?

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

Correct Answer: D
Section:
Explanation:
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

QUESTION 18
Which of the following is a benefit of distributed search?

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

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

QUESTION 19
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

Correct Answer: C
Section:
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
Reference: https://docs.splunk.com/Documentation/Forwarder/8.0.5/Forwarder/Enableareceiver

QUESTION 20
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

Correct Answer: D
Section:
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

QUESTION 21
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>###-##-$2
KEY = _raw
B. props.conf
[mask-SSN]
REGEX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"
FORMAT = $1<SSN>###-##-$2
DEST_KEY = _raw
C. transforms.conf
[mask-SSN]
REX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"
FORMAT = $1<SSN>###-##-$2
DEST_KEY = _raw

D. transforms.conf
[mask-SSN]
REGEX = (?ms)^(.)\<[SSN>\d{3}-?\d{2}-?(\d{4}.*)$"
FORMAT = $1<SSN>###-##-$2
DEST_KEY = _raw

Correct Answer: D
Section:
Explanation:
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

QUESTION 22
If an update is made to an attribute in inputs.conf on a universal forwarder, on which Splunk component would the fishbucket need to be reset in order to reindex the data?

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

Correct Answer: A
Section:
Explanation:
https://www.splunk.com/en_us/blog/tips-and-tricks/what-is-this-fishbucket-thing.html
"Every Splunk instance has a fishbucket index, except the lightest of hand-tuned lightweight forwarders, and if you index a lot of files it can get quite large. As any other index, you can change the retention policy to control the size via indexes.conf" Reference https://community.splunk.com/t5/Archive/How-to-reindex-data-from-a-forwarder/tdp/93310

QUESTION 23
How can native authentication be disabled in Splunk?

A. Remove the $SPLUNK_HOME/etc/passwd file
B. Create an empty $SPLUNK_HOME/etc/passwd file
C. Set SPLUNK_AUTHENTICATION=false in splunk-launch.conf
D. Set nativeAuthentication=false in authentication.conf

Correct Answer: B
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/Secureyouradminaccount

QUESTION 24
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

Correct Answer: C
Section:
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.

QUESTION 25
Which of the following configuration files are used with a universal forwarder? (Choose all that apply.)

A. inputs.conf
B. monitor.conf
C. outputs.conf
D. forwarder.conf

Correct Answer: A, C
Section:
Explanation:
https://docs.splunk.com/Documentation/Forwarder/8.0.5/Forwarder/Configuretheuniversalforwarder
--Key configuration files are: inputs.conf controls how the forwarder collects data. outputs.conf controls how the forwarder sends data to an indexer or other forwarder server.conf for connection and performance tuning deploymentclient.conf for connecting to a deployment server Reference: https://docs.splunk.com/Documentation/Forwarder/8.0.5/Forwarder/ Configuretheuniversalforwarder

QUESTION 26
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.

Correct Answer: A
Section:
Explanation:
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-ANDblacklist-forthesame/td-p/390910

QUESTION 27
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

Correct Answer: C
Section:
Explanation:
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-toforwardinternal-data-to-the/td-p/111658

QUESTION 28
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

Correct Answer: A
Section:
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

QUESTION 29
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

Correct Answer: B
Section:
Explanation:
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..."

QUESTION 30
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

Correct Answer: C
Section:
Explanation:
<https://docs.splunk.com/Documentation/Splunk/8.0.6/Updating/Deploymentserverarchitecture>
https://docs.splunk.com/Splexicon:Serverclass

QUESTION 31
In this source definition the MAX_TIMESTAMP_LOOKHEAD is missing. Which value would fit best?

```
[sshd_syslog]
TIME_PREFIX = ^
TIME_FORMAT = %Y-%m-%d %H:%M:%S.%3N %z
LINE_BREAKER = ([\r\n]+)\d{4}-\d{2}-\d{2} \d{2}:\d{2}:\d{2}
SHOULD_LINEMERGE = false
TRUNCATE = 0
```

Event example:

```
2018-04-13 13:42:41.214 -0500 server sshd[26219]: Connection from 172.0.2.60 port 47366
```

A. MAX_TIMESTAMP_L0CKAHEAD = 5
B. MAX_TIMESTAMP_LOOKAHEAD - 10
C. MAX_TIMESTAMF_LOOKHEAD = 20
D. MAX TIMESTAMP LOOKAHEAD - 30

Correct Answer: D
Section:
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.

QUESTION 32
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

Correct Answer: A, B, D
Section:
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

QUESTION 33
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

Correct Answer: A, C, D
Section:
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

QUESTION 34
Which setting in indexes. conf allows data retention to be controlled by time?

A. maxDaysToKeep
B. moveToFrozenAfter
C. maxDataRetentionTime
D. frozenTimePeriodlnSecs

Correct Answer: D
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Setaretirementandarchivingpolicy

QUESTION 35
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

Correct Answer: B, D
Section:
Explanation:
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.

QUESTION 36
In case of a conflict between a whitelist and a blacklist input setting, which one is used?

A. Blacklist
B. Whitelist

C. They cancel each other out.
D. Whichever is entered into the configuration first.

Correct Answer: A
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.0.4/Data/Whitelistorblacklistspecificincomingdata "
It is not necessary to define both an allow list and a deny list in a configuration stanza. The settings are independent. If you do define both filters and a file matches them both, Splunk Enterprise does not index that file, as the blacklist filter overrides the whitelist filter." Source:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/Whitelistorblacklistspecificincomingdata

QUESTION 37
In which Splunk configuration is the SEDCMD used?

A. props, conf
B. inputs.conf
C. indexes.conf
D. transforms.conf

Correct Answer: A
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.0.5/Forwarding/Forwarddatatothirdpartysystemsd
"You can specify a SEDCMD configuration in props.conf to address data that contains characters that the third-party server cannot process. "

QUESTION 38
Which of the following are supported configuration methods to add inputs on a forwarder? (select all that apply)

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

Correct Answer: A, B, D
Section:
Explanation:
https://docs.splunk.com/Documentation/Forwarder/8.2.1/Forwarder/HowtoforwarddatatoSplunkEnterprise
"You can collect data on the universal forwarder using several methods. Define inputs on the universal forwarder with the CLI. You can use the CLI to define inputs on the universal forwarder.
After you define the inputs, the universal forwarder collects data based on those definitions as long as it has access to the data that you want to monitor. Define inputs on the universal forwarder with configuration files. If the input you want to configure does not have a CLI argument for it, you can configure inputs with configuration files. Create an inputs.conf file in the directory, $SPLUNK_HOME/etc/system/local

QUESTION 39
Which parent directory contains the configuration files in Splunk?

A. SSFLUNK_HOME/etc
B. SSPLUNK_HOME/var
C. SSPLUNK_HOME/conf
D. SSPLUNK_HOME/default

Correct Answer: A
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Admin/Configurationfiledirectories
Section titled, Configuration file directories, states "A detailed list of settings for each configuration file is provided in the .spec file names for that configuration file. You can find the latest version of the .spec and .example files in the $SPLUNK_HOME/etc system/README folder of your Splunk Enterprise installation..."

QUESTION 40
Which forwarder type can parse data prior to forwarding?

A. Universal forwarder
B. Heaviest forwarder
C. Hyper forwarder
D. Heavy forwarder

Correct Answer: D
Section:
Explanation:
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."

QUESTION 41
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

Correct Answer: B
Section:
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.

QUESTION 42
Which Splunk component does a search head primarily communicate with?

A. Indexer
B. Forwarder
C. Cluster master
D. Deployment server

Correct Answer: A
Section:

QUESTION 43

How is data handled by Splunk during the input phase of the data ingestion process?

A. Data is treated as streams.
B. Data is broken up into events.
C. Data is initially written to disk.
D. Data is measured by the license meter.

Correct Answer: A
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.0.5/Deploy/Datapipeline
"In the input segment, Splunk software consumes data. It acquires the raw data stream from its source, breaks in into 64K blocks, and annotates each block with some metadata keys."
Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/Deploy/Datapipeline

QUESTION 44
Which option on the Add Data menu is most useful for testing data ingestion without creating inputs.conf?

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

Correct Answer: A
Section:

QUESTION 45
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.

Correct Answer: B
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/ConsiderationsfordecidinghowtomonitorWindowsdata
"The Splunk platform collects remote Windows data for indexing in one of two ways: From Splunk forwarders, Using Windows Management Instrumentation (WMI). For Splunk Cloud deployments, you must use the Splunk Universal Forwarder on a Windows machines to montior remote Windows data."

QUESTION 46
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

Correct Answer: B
Section:
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."
Reference:
https://docs.splunk.com/Documentation/Splunk/8.0.5/Security/ConfigureLDAPwithSplunkWeb

QUESTION 47
In which phase do indexed extractions in props.conf occur?

A. Inputs phase
B. Parsing phase
C. Indexing phase
D. Searching phase

Correct Answer: B
Section:
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
Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/Admin/Configurationparametersandthedatapipeline

QUESTION 48
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.

Correct Answer: A
Section:
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

QUESTION 49
What action is required to enable forwarder management in Splunk Web?

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

Correct Answer: C
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/Splunk/8.2.1/Updating/Forwardermanagementoverview
https://docs.splunk.com/Documentation/MSApp/2.0.3/MSInfra/Setupadeploymentserver
"To activate deployment server, you must place at least one app into%SPLUNK_HOME%\etc\deployment-apps on the host you want to act as deployment server. In this case, the app is the "send to indexer" app you created earlier, and the host is the indexer you set up initially.

QUESTION 50
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.

Correct Answer: D
Section:
Explanation:
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."

QUESTION 51
When indexing a data source, which fields are considered metadata?

A. source, host, time
B. time, sourcetype, source
C. host, raw, sourcetype
D. sourcetype, source, host

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/SplunkCloud/8.2.2105/SearchReference/Metadata

QUESTION 52
What is the default value of LINE_BREAKER?

A. \r\n
B. ([\r\n]+)
C. \r+\n+
D. (\r\n+)

Correct Answer: B
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/SplunkCloud/8.2.2105/Data/Configureeventlinebreaking
Line breaking, which uses the LINE_BREAKER setting to split the incoming stream of data into separate lines. By default, the LINE_BREAKER value is any sequence of newlines and carriage returns.
In regular expression format, this is represented as the following string: ([\r\n]+). You don't normally need to adjust this setting, but in cases where it's necessary, you must configure it in the props.conf configuration file on the forwarder that sends the data to Splunk Cloud Platform or a Splunk Enterprise indexer. The LINE_BREAKER setting expects a value in regular expression format.

QUESTION 53
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.*]
C. [monitor:///var/log/www1/secure.log]
D. [monitor:///var/log/www*/secure.*]

Correct Answer: C
Section:
Explanation:

Reference:
https://docs.splunk.com/Documentation/Splunk/8.2.1/Data/Monitorfilesanddirectorieswithinputs.conf

QUESTION 54
What are the values for host and index for [stanza1] used by Splunk during index time, given the following configuration files?

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

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

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

A. host=server1
   index=unixinfo
B. host=server1
   index=searchinfo
C. host=searchsvr1
   index=searchinfo
D. host=unixsvr1
   index=unixinfo

Correct Answer: A
Section:
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-theprecedence-for-identical-stanzas-within-a-single/m-p/283566

QUESTION 55
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

Correct Answer: C, D
Section:
Explanation:
Reference: https://wiki.splunk.com/Deploy:BucketRotationAndRetention

QUESTION 56
The LINE_BREAKER attribute is configured in which configuration file?

A. props.conf
B. indexes.conf
C. inpucs.conf
D. transforms.conf

Correct Answer: A
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/SplunkCloud/8.2.2105/Data/Configureeventlinebreaking

QUESTION 57
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

Correct Answer: C
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/Forwarder/8.2.1/Forwarder/Configureloadbalancing

QUESTION 58
A log file contains 193 days worth of timestamped events. Which monitor stanza would be used to collect data 45 days old and newer from that log file?

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

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.splunk.com/Documentation/Splunk/8.2.1/Data/Configuretimestamprecognition

QUESTION 59
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. 90 days
B. 60 days

C. 7 days
D. 14 days

Correct Answer: B
Section:
Explanation:
Reference: https://docs.splunk.com/Documentation/Splunk/8.2.1/Admin/MoreaboutSplunkFree
https://docs.splunk.com/Documentation/Splunk/8.2.3/Admin/TypesofSplunklicenses

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

Correct Answer: D
Section:

QUESTION 61
Which of the following is an appropriate description of a deployment server in a non-cluster environment?

A. Allows management of local Splunk instances, requires Enterprise license, handles job of sending configurations packaged as apps. can automatically restart remote Splunk instances.
B. Allows management of remote Splunk instances, requires Enterprise license, handles job of sending configurations, can automatically restart remote Splunk instances.
C. Allows management of remote Splunk instances, requires no license, handles job of sending configurations, can automatically restart remote Splunk instances.
D. Allows management of remote Splunk instances, requires Enterprise license, handles job of sending configurations, can manually restart remote Splunk instances.

Correct Answer: B
Section:
Explanation:
Reference: https://docs.splunk.com/Documentation/Splunk/8.2.1/Admin/StartSplunk
https://docs.splunk.com/Documentation/Splunk/8.2.2/Updating/Deploymentserverarchitecture
"A deployment client is a Splunk instance remotely configured by a deployment server".

QUESTION 62
Which Splunk forwarder has a built-in license?

A. Light forwarder
B. Heavy forwarder
C. Universal forwarder
D. Cloud forwarder

Correct Answer: C
Section:
Explanation:
Reference: https://community.splunk.com/t5/Getting-Data-In/Do-we-need-a-license-for-Heavyforwarder/m-p/210451

QUESTION 63
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

Correct Answer: C
Section:
Explanation:
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.

QUESTION 64
Which default Splunk role could be assigned to provide users with the following capabilities?
Create saved searches
Edit shared objects and alerts
Not allowed to create custom roles

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

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

QUESTION 65
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

Correct Answer: A
Section:
Explanation:
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. Reference: Splunk Enterprise Certified Admin | Splunk, [Configure Splunk to use LDAP and map groups - Splunk Documentation]

QUESTION 66
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

Correct Answer: C
Section:
Explanation:
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 data to. Therefore, option C is the correct answer. Reference: Splunk Enterprise Certified Admin | Splunk, [Use CLI commands to manage your forwarders - Splunk Documentation]

QUESTION 67
Which artifact is required in the request header when creating an HTTP event?

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

Correct Answer: B
Section:
Explanation:
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. Reference: Splunk Enterprise Certified Admin | Splunk, [About HTTP Event Collector - Splunk Documentation]

QUESTION 68
All search-time field extractions should be specified on which Splunk component?

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

Correct Answer: D
Section:
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. Reference: Splunk Enterprise Certified Admin | Splunk, [About fields - Splunk Documentation]

QUESTION 69
In addition to single, non-clustered Splunk instances, what else can the deployment server push apps to?

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

Correct Answer: A
Section:
Explanation:
Reference: https://community.splunk.com/t5/Deployment-Architecture/Push-apps-fromdeployment-server-automatically-to-universal/m-p/328191
The deployment server is a Splunk component that distributes apps and other configurations to deployment clients, which are Splunk instances that receive updates from the deployment server.
The deployment server can push apps to single, non-clustered Splunk instances, as well as universal forwarders, which are lightweight Splunk agents that forward data to indexers. Therefore, option A is the correct answer. Reference: Splunk Enterprise Certified Admin | Splunk, [About deployment server and forwarder management - Splunk Documentation]

QUESTION 70
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>

Correct Answer: C
Section:
Explanation:
Reference: https://community.splunk.com/t5/Getting-Data-In/How-can-I-trigger-the-re-indexing-ofa-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. Reference: Splunk Enterprise Certified Admin | Splunk, [Use btprobe to troubleshoot file monitoring - Splunk Documentation]

QUESTION 71
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

Correct Answer: A
Section:
Explanation:
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.
Reference: Splunk Enterprise Certified Admin | Splunk, [Configure event line merging - Splunk Documentation]

QUESTION 72

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

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

QUESTION 73
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.

Correct Answer: A, C
Section:
Explanation:
Reference: https://community.splunk.com/t5/Getting-Data-In/Why-does-Splunk-have-multipleindexes/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.

QUESTION 74
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.

Correct Answer: C
Section:
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."

QUESTION 75
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

Correct Answer: D
Section:
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."
Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/Admin/Aboutlicenseviolations

QUESTION 76
Who provides the Application Secret, Integration, and Secret keys, as well as the API Hostname when setting up Duo for Multi-Factor Authentication in Splunk Enterprise?

A. Duo Administrator
B. LDAP Administrator
C. SAML Administrator
D. Trio Administrator

Correct Answer: A
Section:
Explanation:
Reference: https://duo.com/docs/splunk

QUESTION 77
When does a warm bucket roll over to a cold bucket?

A. When Splunk is restarted.
B. When the maximum warm bucket age has been reached.
C. When the maximum warm bucket size has been reached.
D. When the maximum number of warm buckets is reached.

Correct Answer: D
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.1.1/Indexer/HowSplunkstoresindexes
Once further conditions are met (for example, the index reaches some maximum number of warm buckets), the indexer begins to roll the warm buckets to cold, based on their age. It always selects the oldest warm bucket to roll to cold. Buckets continue to roll to cold as they age in this manner.
Cold buckets reside in a different location from hot and warm buckets. You can configure the location so that cold buckets reside on cheaper storage.
Reference: https://community.splunk.com/t5/Deployment-Architecture/Rolling-Hot-Data-to-to-Coldquicker/tdp/166653

QUESTION 78
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

Correct Answer: D
Section:
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=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

QUESTION 79
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.

Correct Answer: C
Section:
Explanation:
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 peerapps 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.

QUESTION 80
What happens when there are conflicting settings within two or more configuration files?

A. The setting is ignored until conflict is resolved.
B. The setting for both values will be used together.
C. The setting with the lowest precedence is used.
D. The setting with the highest precedence is used.

Correct Answer: D
Section:
Explanation:
When there are conflicting settings within two or more configuration files, the setting with the highest precedence is used. The precedence of configuration files is determined by a combination of the file type, the directory location, and the alphabetical order of the file names.

QUESTION 81
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.

Correct Answer: A, C
Section:
Explanation:
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.

QUESTION 82
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?

- ○ $SPLUNK_HOME/etc/apps/Splunk_TA_nginx/local/inputs.conf
- ○ $SPLUNK_HOME/etc/apps/Splunk_TA_nginx/default/inputs.conf
- ○ $SPLUNK_HOME/etc/system/default/Splunk_TA_nginx/local/inputs.conf
- ○ $SPLUNK_HOME/etc/users/admin/Splunk_TA_nginx/local/inputs.conf

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

Correct Answer: A
Section:
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.

QUESTION 83
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

Correct Answer: B, C
Section:
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.

QUESTION 84
In a customer managed Splunk Enterprise environment, what is the endpoint URI used to collect data?

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

Correct Answer: A
Section:
Explanation:
This is the endpoint URI used to collect data using the HTTP Event Collector (HEC), which is a tokenbased API that allows you to send data to Splunk Enterprise from any application that can make an HTTP request. The endpoint URI consists of the protocol (http or https), the hostname or IP address of the Splunk server, the port number (default is 8088), and the service name (services/collector). For example: https://mysplunkserver.example.com:8088/services/collector

QUESTION 85
Running this search in a distributed environment:

```
index=aws source=*/AWSLogs/314575187704/elasticloadbalancing/*
| lookup responsible_teams elb OUTPUT team
| eval team=coalesce(team,elb)
| stats sum(received_bytes) sum(sent_bytes) by team
| outputlookup current_prod_account_data
```


On what Splunk component does the eval command get executed?

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

Correct Answer: C
Section:
Explanation:
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.

QUESTION 86
When would the following command be used?

```
./splunk check-integrity -index [ index name ] [ -verbose ]
```

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.

Correct Answer: D
Section:
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.

QUESTION 87
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]

Correct Answer: B
Section:
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.

QUESTION 88
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

Correct Answer: B
Section:
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 onetime code sent to their phone, or scan their fingerprint and face.

QUESTION 89
A Splunk administrator has been tasked with developing a retention strategy to have frequently accessed data sets on SSD storage and to have older, less frequently accessed data on slower NAS storage. They have set a mount point for the NAS. Which parameter do they need to modify to set the path for the older, less frequently accessed data in indexes.conf?

A. homepath
B. thawedPath
C. summaryHomePath
D. colddeath

Correct Answer: D
Section:

Explanation:
The coldPath parameter defines the path for the cold buckets, which are the oldest and least frequently accessed data in an index1. By setting the coldPath to point to the NAS mount point, the Splunk administrator can achieve the retention strategy of having older data on slower NAS storage.

QUESTION 90
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.

Correct Answer: C
Section:
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.

QUESTION 91
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 reexecute 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.

Correct Answer: A
Section:
Explanation:
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.

QUESTION 92
What is the correct curl to send multiple events through HTTP Event Collector?

```
curl "https://mysplunkserver.example.com:8088/services/collector" \
-H "Authorization: Splunk DF457ZE4-3G51-85F5-H777-0284GG91PF67" \
-d "event": "Hello World", "Hola Mundo", "Hallo Welt"
```

```
curl "https://mysplunkserver.example.com:8088/services/collector" \
-H "Authorization: Splunk DF457ZE4-3G51-85F5-H777-0284GG91PF67" \
-d "event": "Hello World", "event": "Hola Mundo", "event": "Hallo Welt"
```

```
curl "https://mysplunkserver.example.com:8088/services/collector" \
-H "Authorization: Splunk DF457ZE4-3G51-85F5-H777-0284GG91PF67" \
-d '{"event": "Hello World"}{"event": "Hola Mundo"}{"event": "Hallo Welt", "nested": {"key1": "value1"}}'
```

```
curl "https://mysplunkserver.example.com:8088/services/collector" \
-H "Authorization: Splunk DF457ZE4-3G51-85F5-H777-0284GG91PF67" \
-d '{"event": "Hello World", "Hola Mundo", "Hallo Welt", "nested": {"key1": "value1"}}'
```

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

Correct Answer: B
Section:
Explanation:
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.

QUESTION 93
What event-processing pipelines are used to process data for indexing? (select all that apply)

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

Correct Answer: B, D
Section:

QUESTION 94
What is the correct example to redact a plain-text password from raw events?

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

Correct Answer: B
Section:
Explanation:
The correct answer is B. in props.conf:
[identity]
SEDCMD-redact_pw = s/password=([^,|/s]+)/ ####REACTED####/g
According to the Splunk documentation1, to redact sensitive data from raw events, you need to use the SEDCMD attribute in the props.conf file. The SEDCMD attribute applies a sed expression to the raw data before indexing. The sed expression can use the s command to replace a pattern with a substitution string. For example, the following sed expression replaces any occurrence of password= followed by any characters until a comma, whitespace, or slash with ####REACTED####:
s/password=([^,|/s]+)/ ####REACTED####/g
The g flag at the end means that the replacement is applied globally, not just to the first match.
Option A is incorrect because it uses the REGEX attribute instead of the SEDCMD attribute. The REGEX attribute is used to extract fields from events, not to modify them.
Option C is incorrect because it uses the transforms.conf file instead of the props.conf file. The transforms.conf file is used to define transformations that can be applied to fields or events, such as lookups, evaluations, or replacements. However, these transformations are applied after indexing, not before.
Option D is incorrect because it uses both the wrong attribute and the wrong file. There is no REGEX-redact_pw attribute in the transforms.conf file.

QUESTION 95
What is an example of a proper configuration for CHARSET within props.conf?

A. [host: : server. splunk. com] CHARSET = BIG5
B. [index: :main] CHARSET = BIG5
C. [sourcetype: : son] CHARSET = BIG5
D. [source: : /var/log/ splunk] CHARSET = BIG5

Correct Answer: A
Section:
Explanation:
According to the Splunk documentation1, to manually specify a character set for an input, you need to set the CHARSET key in the props.conf file. You can specify the character set by host, source, or sourcetype, but not by index.
https://docs.splunk.com/Documentation/Splunk/latest/Data/Configurecharactersetencoding

QUESTION 96
A security team needs to ingest a static file for a specific incident. The log file has not been collected previously and future updates to the file must not be indexed.
Which command would meet these needs?

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

C. splunk add monitor /opt/incident/data.log ---index incident
D. splunk edit oneshot [opt/ incident/data.* ---index incident

Correct Answer: A
Section:
Explanation:
The correct answer is A. splunk add one shot / opt/ incident [data . log ---index incident
According to the Splunk documentation1, the splunk add one shot command adds a single file or directory to the Splunk index and then stops monitoring it. This is useful for ingesting static files that do not change or update.
The command takes the following syntax:
splunk add one shot <file> -index <index_name>
The file parameter specifies the path to the file or directory to be indexed. The index parameter specifies the name of the index where the data will be stored. If the index does not exist, Splunk will create it automatically.
Option B is incorrect because the splunk edit monitor command modifies an existing monitor input, which is used for ingesting files or directories that change or update over time. This command does not create a new monitor input, nor does it stop monitoring after indexing.
Option C is incorrect because the splunk add monitor command creates a new monitor input, which is also used for ingesting files or directories that change or update over time. This command does not stop monitoring after indexing.
Option D is incorrect because the splunk edit oneshot command does not exist. There is no such command in the Splunk CLI.

QUESTION 97
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

Correct Answer: C
Section:
Explanation:
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!'}'

QUESTION 98
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.

Correct Answer: B
Section:
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.

QUESTION 99
A Universal Forwarder is collecting two separate sources of data (A,B). Source A is being routed through a Heavy Forwarder and then to an indexer. Source B is being routed directly to the indexer. Both sets of data require the masking of raw text strings before being written to disk. What does the administrator need to do t/1 /2nsure that the masking takes place successfully?

A. Make sure that props . conf and transforms . conf are both present on the in-dexer and the search head.
B. For source A, make sure that props . conf is in place on the indexer; and for source B, make sure transforms . conf is present on the Heavy Forwarder.
C. Make sure that props . conf and transforms . conf are both present on the Universal Forwarder.
D. Place both props . conf and transforms . conf on the Heavy Forwarder for source A, and place both props . conf and transforms . conf on the indexer for source B.

Correct Answer: D
Section:
Explanation:
The correct answer is D. Place both props . conf and transforms . conf on the Heavy Forwarder for source A, and place both props . conf and transforms . conf on the indexer for source B.
According to the Splunk documentation1, to mask sensitive data from raw events, you need to use the SEDCMD attribute in the props.conf file and the REGEX attribute in the transforms.conf file. The SEDCMD attribute applies a sed expression to the raw data before indexing, while the REGEX attribute defines a regular expression to match the data to be masked. You need to place these files on the Splunk instance that parses the data, which is usually the indexer or the heavy forwarder2. The universal forwarder does not parse the data, so it does not need these files.
For source A, the data is routed through a heavy forwarder, which can parse the data before sending it to the indexer. Therefore, you need to place both props.conf and transforms.conf on the heavy forwarder for source A, so that the masking takes place before indexing.
For source B, the data is routed directly to the indexer, which parses and indexes the data. Therefore, you need to place both props.conf and transforms.conf on the indexer for source B, so that the masking takes place before indexing.

QUESTION 100
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.

Correct Answer: D
Section:
Explanation:
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.

QUESTION 101

Event processing occurs at which phase of the data pipeline?

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

Correct Answer: C
Section:
Explanation:
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.

QUESTION 102
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.

Correct Answer: A
Section:
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.

QUESTION 103
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.

Correct Answer: A, D
Section:
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.

QUESTION 104
What is the name of the object that stores events inside of an index?

A. Container

B. Bucket
C. Data layer
D. Indexer

Correct Answer: B
Section:
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.

QUESTION 105
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.

Correct Answer: C
Section:
Explanation:
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.

QUESTION 106
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.

Correct Answer: C
Section:
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.

QUESTION 107
What type of Splunk license is pre-selected in a brand new Splunk installation?

A. Free license B. Forwarder license
B. Enterprise trial license
C. Enterprise license

Correct Answer: C
Section:
Explanation:
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.

QUESTION 108
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

Correct Answer: A, B, C
Section:
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.

QUESTION 109
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

Correct Answer: A, B, C
Section:

Explanation:
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.

QUESTION 110
Which of the following authentication types requires scripting in Splunk?

A. ADFS
B. LDAP
C. SAML
D. RADIUS

Correct Answer: D
Section:
Explanation:
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.

QUESTION 111
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.

Correct Answer: B
Section:
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."

QUESTION 112
How is a remote monitor input distributed to forwarders?

A. As an app.
B. As a forward.conf file.
C. As a monitor.conf file.
D. As a forwarder monitor profile.

Correct Answer: A
Section:
Explanation:
https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/Usingforwardingagents
Scroll down to the section Titled, How to configure forwarder inputs, and subsection Here are the main ways that you can configure data inputs on a forwarder Install the app or add-on that contains the inputs you wants
Reference: https://docs.splunk.com/Documentation/Splunk/8.0.5/Data/Usingforwardingagents

QUESTION 113
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.

Correct Answer: A
Section:
Explanation:
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.

QUESTION 114
Search heads in a company's European offices need to be able to search data in their New York offices. They also need to restrict access to certain indexers. What should be configured to allow this type of action?

A. Indexer clustering
B. LDAP control
C. Distributed search
D. Search head clustering

Correct Answer: C
Section:
Explanation:
The correct answer is C. Distributed search is the feature that allows search heads in a company's European offices to search data in their New York offices. Distributed search also enables restricting access to certain indexers by using the splunk_server field or the server.conf file1.
Distributed search is a way to scale your Splunk deployment by separating the search management and presentation layer from the indexing and search retrieval layer. With distributed search, a Splunk instance called a search head sends search requests to a group of indexers, or search peers, which perform the actual searches on their indexes. The search head then merges the results back to the user2.
Distributed search has several use cases, such as horizontal scaling, access control, and managing geo-dispersed data. For example, users in different offices can search data across the enterprise or only in their local area, depending on their needs and permissions2.
The other options are incorrect because:
A) Indexer clustering is a feature that replicates data across a group of indexers to ensure data availability and recovery. Indexer clustering does not directly affect distributed search, although search heads can be configured to search across an indexer cluster3.
B) LDAP control is a feature that allows Splunk to integrate with an external LDAP directory service for user authentication and role mapping. LDAP control does not affect distributed search, although it can be used to manage user access to data and searches.
D) Search head clustering is a feature that distributes the search workload across a group of search heads that share resources, configurations, and jobs. Search head clustering does not affect distributed search, although the search heads in a cluster can search across the same set of indexers.

QUESTION 115
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

Correct Answer: C
Section:
Explanation:
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.