Number: SPLK-1005
Passing Score: 800
Time Limit: 120
File Version: 4.0

Exam Code: SPLK-1005

Exam Name: Splunk Cloud Certified Admin

Exam A

QUESTION 1
Which of the following statements regarding apps in Splunk Cloud is true?

A. Self-service install of premium apps is possible.
B. Only Cloud certified and vetted apps are supported.
C. Any app that can be deployed in an on-prem Splunk Enterprise environment is also supported on Splunk Cloud.
D. Self-service install is available for all apps on Splunkbase.

Correct Answer: B
Section:
Explanation:
In Splunk Cloud, only apps that have been certified and vetted by Splunk are supported. This is because Splunk Cloud is a managed service, and Splunk ensures that all apps meet specific security, performance, and compatibility requirements before they can be installed. This certification process guarantees that the apps won't negatively impact the overall environment, ensuring a stable and secure cloud service.
Self-service installation is available, but it is limited to apps that are certified for Splunk Cloud. Non-certified apps cannot be installed directly; they require a review and approval process by Splunk support.
Splunk Cloud
Reference: Refer to Splunk's documentation on app installation and the list of Cloud-vetted apps available on Splunkbase to understand which apps can be installed in Splunk Cloud.
Source:
Splunk Docs: About apps in Splunk Cloud
Splunkbase: Splunk Cloud Apps

QUESTION 2
When using Splunk Universal Forwarders, which of the following is true?

A. No more than six Universal Forwarders may connect directly to Splunk Cloud.
B. Any number of Universal Forwarders may connect directly to Splunk Cloud.
C. Universal Forwarders must send data to an Intermediate Forwarder.
D. There must be one Intermediate Forwarder for every three Universal Forwarders.

Correct Answer: B
Section:
Explanation:
Universal Forwarders can connect directly to Splunk Cloud, and there is no limit on the number of Universal Forwarders that may connect directly to it. This capability allows organizations to scale their data ingestion easily by deploying as many Universal Forwarders as needed without the requirement for intermediate forwarders unless additional data processing, filtering, or load balancing is required.
Splunk Documentation
Reference: Forwarding Data to Splunk Cloud

QUESTION 3
Which of the following is correct in regard to configuring a Universal Forwarder as an Intermediate Forwarder?

A. This can only be turned on using the Settings > Forwarding and Receiving menu in Splunk Web/UI.
B. The configuration changes can be made using Splunk Web. CU, directly in configuration files, or via a deployment app.
C. The configuration changes can be made using CU, directly in configuration files, or via a deployment app.
D. It is only possible to make this change directly in configuration files or via a deployment app.

Correct Answer: D
Section:
Explanation:
Configuring a Universal Forwarder (UF) as an Intermediate Forwarder involves making changes to its configuration to allow it to receive data from other forwarders before sending it to indexers.
D . It is only possible to make this change directly in configuration files or via a deployment app: This is the correct answer. Configuring a Universal Forwarder as an Intermediate Forwarder is done by editing the configuration files directly (like outputs.conf), or by deploying a pre-configured app via a deployment server. The Splunk Web UI (Management Console) does not provide an interface for configuring a Universal Forwarder as an Intermediate Forwarder.
A . This can only be turned on using the Settings > Forwarding and Receiving menu in Splunk Web/UI: Incorrect, as this applies to Heavy Forwarders, not Universal Forwarders.
B . The configuration changes can be made using Splunk Web, CLI, directly in configuration files, or via a deployment app: Incorrect, the Splunk Web UI is not used for configuring Universal Forwarders.
C . The configuration changes can be made using CLI, directly in configuration files, or via a deployment app: While CLI could be used for certain configurations, the specific Intermediate Forwarder setup is typically done via configuration files or deployment apps.
Splunk Documentation
Reference:
Universal Forwarder Configuration
Intermediate Forwarder Configuration

QUESTION 4
What does the followTail attribute do in inputs.conf?

A. Pauses a file monitor if the queue is full.
B. Only creates a tail checkpoint of the monitored file.
C. Ingests a file starting with new content and then reading older events.
D. Prevents pre-existing content in a file from being ingested.

Correct Answer: D
Section:
Explanation:
The followTail attribute in inputs.conf controls how Splunk processes existing content in a monitored file.
D . Prevents pre-existing content in a file from being ingested: This is the correct answer. When followTail = true is set, Splunk will ignore any pre-existing content in a file and only start monitoring from the end of the file, capturing new data as it is added. This is useful when you want to start monitoring a log file but do not want to index the historical data that might be present in the file.
A . Pauses a file monitor if the queue is full: Incorrect, this is not related to the followTail attribute.
B . Only creates a tail checkpoint of the monitored file: Incorrect, while a tailing checkpoint is created for state tracking, followTail specifically refers to skipping the existing content.
C . Ingests a file starting with new content and then reading older events: Incorrect, followTail does not read older events; it skips them.
Splunk Documentation
Reference:
followTail Attribute Documentation
Monitoring Files
These answers align with Splunk's best practices and available documentation on managing and configuring Splunk environments.

QUESTION 5
In case of a Change Request, which of the following should submit a support case for Splunk Support?

A. The party requesting the change.
B. Certified Splunk Cloud administrator.
C. Splunk infrastructure owner.
D. Any person with the appropriate entitlement

Correct Answer: D
Section:

Explanation:
In Splunk Cloud, when there is a need for a change request that might involve modifying settings, upgrading, or other actions requiring Splunk Support, the process typically requires submitting a support case.
D . Any person with the appropriate entitlement: This is the correct answer. Any individual who has the necessary permissions or entitlements within the Splunk environment can submit a support case. This includes administrators or users who have been granted the ability to engage with Splunk Support. The request does not necessarily have to come from a Certified Splunk Cloud Administrator or the infrastructure owner; rather, it can be submitted by anyone with the correct level of access.
Splunk Documentation
Reference:
Submitting a Splunk Support Case
Managing User Roles and Entitlements

QUESTION 6
Consider the following configurations:

```
$SPLUNK_HOME/etc/apps/unix/local/inputs.conf

    [monitor:///var/log/secure.log]
    sourcetype = access_combined
    index = security


$SPLUNK_HOME/etc/apps/search/local/inputs.conf

    [monitor:///var/log/secure.log]
    host = logsvr1
    sourcetype = linux_secure
```

What is the value of the sourcetype property for this stanza based on Splunk's configuration file precedence?

A. NULL, or unset, due to configuration conflict
B. access_corabined
C. linux aacurs
D. linux_secure, access_combined

Correct Answer: C
Section:
Explanation:
When there are conflicting configurations in Splunk, the platform resolves them based on the configuration file precedence rules. These rules dictate which settings are applied based on the hierarchy of the configuration files.
In the provided configurations:
The first configuration in $SPLUNK_HOME/etc/apps/unix/local/inputs.conf sets the sourcetype to access_combined.
The second configuration in $SPLUNK_HOME/etc/apps/search/local/inputs.conf sets the sourcetype to linux_secure.
Configuration File Precedence:
In Splunk, configurations in local directories take precedence over those in default.
If two configurations are in local directories of different apps, the alphabetical order of the app names determines the precedence.
Since 'search' comes after 'unix' alphabetically, the configuration in $SPLUNK_HOME/etc/apps/search/local/inputs.conf will take precedence.
Therefore, the value of the sourcetype property for this stanza is linux_secure.
Splunk Documentation
Reference:
Configuration File Precedence
Resolving Conflicts in Splunk Configurations
This confirms that the correct answer is C. linux_secure.

QUESTION 7
In which of the following situations should Splunk Support be contacted?

A. When a custom search needs tuning due to not performing as expected.
B. When an app on Splunkbase indicates Request Install.
C. Before using the delete command.
D. When a new role that mirrors sc_admin is required.

Correct Answer: B
Section:
Explanation:
In Splunk Cloud, when an app on Splunkbase indicates 'Request Install,' it means that the app is not available for direct self-service installation and requires intervention from Splunk Support. This could be because the app needs to undergo an additional review for compatibility with the managed cloud environment or because it requires special installation procedures.
In these cases, customers need to contact Splunk Support to request the installation of the app. Support will ensure that the app is properly vetted and compatible with Splunk Cloud before proceeding with the installation.
Splunk Cloud
Reference: For further details, consult Splunk's guidelines on requesting app installations in Splunk Cloud and the processes involved in reviewing and approving apps for use in the cloud environment.
Source:
Splunk Docs: Install apps in Splunk Cloud Platform
Splunkbase: App request procedures for Splunk Cloud

QUESTION 8
The following Apache access log is being ingested into Splunk via a monitor input:

```
192.2.14.109 my.webserver.example - - 443 [09/Sep/2020:06:35:25 -0400] "GET / HTTP/1.1" "" 200 409 "-" "Mozilla/5.0
(Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/60.0.3112.113 Safari/537.36" 502 3110 1705
```

How does Splunk determine the time zone for this event?

A. The value of the TZ attribute in props. cont for the a :ces3_ccwbined sourcetype.
B. The value of the TZ attribute in props, conf for the my.webserver.example host.
C. The time zone of the Heavy/Intermediate Forwarder with the monitor input.
D. The time zone indicator in the raw event data.

Correct Answer: D
Section:
Explanation:
In Splunk, when ingesting logs such as an Apache access log, the time zone for each event is typically determined by the time zone indicator present in the raw event data itself. In the log snippet you provided, the time zone is indicated by -0400, which specifies that the event's timestamp is 4 hours behind UTC (Coordinated Universal Time).
Splunk uses this information directly from the event to properly parse the timestamp and apply the correct time zone. This ensures that the event's time is accurately reflected regardless of the time zone in which the Splunk instance or forwarder is located.
Splunk Cloud
Reference: For further details, you can review Splunk documentation on timestamp recognition and time zone handling, especially in relation to log files and data ingestion configurations.
Source:
Splunk Docs: How Splunk software handles timestamps
Splunk Docs: Configure event timestamp recognition

QUESTION 9
What syntax is required in inputs.conf to ingest data from files or directories?

A. A monitor stanza, sourcetype, and Index is required to ingest data.
B. A monitor stanza, sourcetype, index, and host is required to ingest data.
C. A monitor stanza and sourcetype is required to ingest data.

D. Only the monitor stanza is required to ingest data.

Correct Answer: A
Section:
Explanation:
In Splunk, to ingest data from files or directories, the basic configuration in inputs.conf requires at least the following elements:
monitor stanza: Specifies the file or directory to be monitored.
sourcetype: Identifies the format or type of the incoming data, which helps Splunk to correctly parse it.
index: Determines where the data will be stored within Splunk.
The host attribute is optional, as Splunk can auto-assign a host value, but specifying it can be useful in certain scenarios. However, it is not mandatory for data ingestion.
Splunk Cloud
Reference: For more details, you can consult the Splunk documentation on inputs.conf file configuration and best practices.
Source:
Splunk Docs: Monitor files and directories
Splunk Docs: Inputs.conf examples

QUESTION 10
The following sample log event shows evidence of credit card numbers being present in the transactions. loc file.

```
2020-09-10 11:19:00 action=new_transaction cc_num=4556723486763517 value=2.55 ccy=GBP
```

Which of these SEDCM3 settings will mask this and other suspected credit card numbers with an Y character for each character being masked? The indexed event should be formatted as follows:

```
Masked version:
2020-09-10 11:19:00 action=new_transaction cc_num=4556723xxxxxxxxx value=2.55 ccy=GBP
```

A)

```
[source::.../transactions.log]
SEDCMD-mask_num = s/(?cc_num=\d{7})\d{9}/\1xxxxxxxxx/g
```

B)

```
[source::.../transactions.log]
SEDCMD-mask_num = s/(?cc_num=\d{7})486763517/\1xxxxxxxxx/g
```

C)

```
[source::.../transactions.log]
SEDCMD-mask_num = s/cc_num=(?\d{7})486763517/cc_num=xxxxxxxxx/g
```

D)

```
[source::.../transactions.log]
SEDCMD-mask_num = s/(?cc_num=\d{7})\d{9}//g
```

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

Correct Answer: A
Section:
Explanation:
The correct SEDCMD setting to mask the credit card numbers, ensuring that the masked version replaces each digit with an 'x' character, is Option A.
The SEDCMD syntax works as follows:
s/ starts the substitute command.
(?cc_num=\d{7})\d{9}/ matches the specific pattern of the credit card number in the logs.
\1xxxxxxxxx replaces the matched portion with the first captured group (the first 7 digits of the cc_num), followed by 9 'x' characters to mask the remaining digits.

/g ensures that the substitution is applied globally, throughout the string.
Thus, Option A correctly implements this requirement.
Splunk Documentation
Reference: SEDCMD for Masking Data

QUESTION 11
Which of the following is a correct statement about Universal Forwarders?

A. The Universal Forwarder must be able to contact the license master.
B. A Universal Forwarder must connect to Splunk Cloud via a Heavy Forwarder.
C. A Universal Forwarder can be an Intermediate Forwarder.
D. The default output bandwidth is 500KBps.

Correct Answer: C
Section:
Explanation:
A Universal Forwarder (UF) can indeed be configured as an Intermediate Forwarder. This means that the UF can receive data from other forwarders and then forward that data on to indexers or Splunk Cloud, effectively acting as a relay point in the data forwarding chain.
Option A is incorrect because a Universal Forwarder does not need to contact the license master; only indexers and search heads require this.
Option B is incorrect as Universal Forwarders can connect directly to Splunk Cloud or via other forwarders.
Option D is also incorrect because the default output bandwidth limit for a UF is typically much higher than 500KBps (default is 256KBps per pipeline, but can be configured).
Splunk Documentation
Reference: Universal Forwarder

QUESTION 12
Which of the following is true when integrating LDAP authentication?

A. Splunk stores LDAP end user names and passwords on search heads.
B. The mapping of LDAP groups to Splunk roles happens automatically.
C. Splunk Cloud only supports Active Directory LDAP servers.
D. New user data is cached the first time a user logs in.

Correct Answer: D
Section:
Explanation:
When integrating LDAP authentication with Splunk, new user data is cached the first time a user logs in. This means that Splunk does not store LDAP usernames and passwords; instead, it relies on the LDAP server for authentication. The mapping of LDAP groups to Splunk roles must be configured manually; it does not happen automatically. Additionally, Splunk Cloud supports various LDAP servers, not just Active Directory.
Splunk Documentation
Reference: LDAP Authentication

QUESTION 13
A Splunk Cloud administrator is looking to allow a new group of Splunk users in the marketing department to access the Splunk environment and view a dashboard with relevant data. These users need to access marketing data (stored in the marketing_data index), but shouldn't be able to access other data, such as events related to security or operations.
Which approach would be the best way to accomplish these requirements?

A. Create a new user with access to the marketing_data index assigned.
B. Create a new role that inherits the user role and remove the capability to search indexes other than marketing_data.
C. Create a new role that inherits the admin rote and assign access to the marketing_dat.a index.

D. Create a new role that does not inherit from any other role, turn on the same capabilities as the user role, and assign access to the marketing_data index.

Correct Answer: B
Section:
Explanation:
The best approach to meet the requirements of the marketing department is to create a new role that inherits the user role but with restricted access to only the marketing_data index. This setup allows users to perform searches and view dashboards while ensuring they cannot access other indexes such as those containing security or operations data.
Splunk Documentation
Reference: Splunk Role-based Access Control

QUESTION 14
Files from multiple systems are being stored on a centralized log server. The files are organized into directories based on the original server they came from. Which of the following is a recommended approach for correctly setting the host values based on their origin?

A. Use the host segment, setting.
B. Set host = * in the monitor stanza.
C. The host value cannot be dynamically set.
D. Manually create a separate monitor stanza for each host, with the nose = value set.

Correct Answer: A
Section:
Explanation:
The recommended approach for setting the host values based on their origin when files from multiple systems are stored on a centralized log server is to use the host_segment setting. This setting allows you to dynamically set the host value based on a specific segment of the file path, which can be particularly useful when organizing logs from different servers into directories.
Splunk Documentation
Reference: Inputs.conf - host_segment

QUESTION 15
In which file can the SH0ULD_LINEMERCE setting be modified?

A. transforms.conf
B. inputs.conf
C. props.conf
D. outputs.conf

Correct Answer: C
Section:
Explanation:
The SHOULD_LINEMERGE setting is used in Splunk to control whether or not multiple lines of an event should be combined into a single event. This setting is configured in the props.conf file, where Splunk handles data parsing and field extraction. Setting SHOULD_LINEMERGE = true merges lines together based on specific rules.
Splunk Documentation
Reference: props.conf - SHOULD_LINEMERGE

QUESTION 16
What is the recommended approach to collect data from network devices?

A. TCP/UDP Feed > Heavy Forwarder > Intermediate Forwarder > Splunk Cloud
B. TCP/UDP Feed > Syslog Server with Universal Forwarder > Splunk Cloud

C. TCP/UDP Feed > Universal Forwarder > Intermediate Forwarder > Splunk Cloud
D. TCP/UDP Feed > Intermediate Forwarder > Heavy Forwarder > Splunk Cloud

Correct Answer: B
Section:
Explanation:
The recommended approach to collect data from network devices is to use a Syslog server with a Universal Forwarder (UF) installed. The network devices send data to the Syslog server, which then forwards the data to Splunk Cloud using the Universal Forwarder. This method ensures reliable data ingestion and processing while maintaining flexibility in handling different types of network device data.
Splunk Documentation
Reference: Best practices for getting data in

QUESTION 17
When a forwarder phones home to a Deployment Server it compares the check-sum value of the forwarder's app to the Deployment Server's app. What happens to the app If the check-sum values do not match?

A. The app on the forwarder is always deleted and re-downloaded from the Deployment Server.
B. The app on the forwarder is only deleted and re-downloaded from the Deployment Server if the forwarder's app has a smaller check-sum value.
C. The app is downloaded from the Deployment Server and the changes are merged.
D. A warning is generated on the Deployment Server stating the apps are out of sync. An Admin will need to confirm which version of the app should be used.

Correct Answer: A
Section:
Explanation:
When a forwarder phones home to a Deployment Server, it compares the checksum of its apps with those on the Deployment Server. If the checksums do not match, the app on the forwarder is always deleted and re-downloaded from the Deployment Server. This ensures that the forwarder has the most current and correct version of the app as dictated by the Deployment Server.
Splunk Documentation
Reference: Deployment Server Overview

QUESTION 18
Which of the following app installation scenarios can be achieved without involving Splunk Support?

A. Deploy premium apps.
B. Install apps via the Request Install button.
C. Install apps via self-service.
D. Install apps that have not gone through the vetting process.

Correct Answer: C
Section:
Explanation:
In Splunk Cloud, you can install apps via self-service, which allows you to install certain approved apps without involving Splunk Support. This self-service capability is provided for apps that have already been vetted and approved for use in the Splunk Cloud environment.
Option A typically requires support involvement because premium apps often need licensing or other special considerations.
Option B might involve the Request Install button, but some apps might still require vetting or support approval.
Option D is incorrect because apps that have not gone through the vetting process cannot be installed via self-service and would require Splunk Support for evaluation and approval.
Splunk Documentation
Reference: Install apps on Splunk Cloud

QUESTION 19
Which file or folder below is not a required part of a deployment app?

A. app.conf (in default or local)
B. local.meta
C. metadata folder
D. props.conf

Correct Answer: D
Section:
Explanation:
When creating a deployment app in Splunk, certain files and folders are considered essential to ensure proper configuration and operation:
app.conf (in default or local): This is required as it defines the app's metadata and behaviors.
local.meta: This file is important for defining access permissions for the app and is often included.
metadata folder: The metadata folder contains files like local.meta and default.meta and is typically required for defining permissions and other metadata-related settings.
props.conf: While props.conf is essential for many Splunk apps, it is not mandatory unless you need to define specific data parsing or transformation rules.
D . props.conf is the correct answer because, although it is commonly used, it is not a mandatory part of every deployment app. An app may not need data parsing configurations, and thus, props.conf might not be present in some apps.
Splunk Documentation
Reference:
Building Splunk Apps
Deployment Apps
This confirms that props.conf is not a required part of a deployment app, making it the correct answer.

QUESTION 20
Which of the following files is used for both search-time and index-time configuration?

A. inputs.conf
B. props.conf
C. macros.conf
D. savesearch.conf

Correct Answer: B
Section:
Explanation:
The props.conf file is a crucial configuration file in Splunk that is used for both search-time and index-time configurations.
At index-time, props.conf is used to define how data should be parsed and indexed, such as timestamp recognition, line breaking, and data transformations.
At search-time, props.conf is used to configure how data should be searched and interpreted, such as field extractions, lookups, and sourcetypes.
B . props.conf is the correct answer because it is the only file listed that serves both index-time and search-time purposes.
Splunk Documentation
Reference:
props.conf - configuration for search-time and index-time

QUESTION 21
What Splunk command will allow an administrator to view the runtime configuration instructions for a monitored file in Inputs. cont on the forwarders?

A. ./splunk _internal call /services/data/input.3/filemonitor
B. ./splunk show config inputs.conf
C. ./splunk _internal rest /services/data/inputs/monitor
D. ./splunk show config inputs

Correct Answer: C
Section:
Explanation:
To view the runtime configuration instructions for a monitored file in inputs.conf on the forwarder, the correct command to use involves accessing the internal REST API that provides details on data inputs.
C . ./splunk _internal rest /services/data/inputs/monitor is the correct answer. This command uses Splunk's internal REST endpoint to retrieve information about monitored files, including their runtime configurations as defined in inputs.conf.
Splunk Documentation
Reference:
Splunk REST API - Data Inputs

QUESTION 22
Which of the following lists all parameters supported by the acceptFrom argument?

A. IPv4, IPv6, CIDRs, DNS names, Wildcards
B. IPv4, IPv6, CIDRs, DNS names
C. CIDRs, DNS names, Wildcards
D. IPv4. CIDRs, DNS names. Wildcards

Correct Answer: B
Section:
Explanation:
The acceptFrom parameter is used in Splunk to specify which IP addresses or DNS names are allowed to send data to a Splunk instance. The supported formats include IPv4, IPv6, CIDR notation, and DNS names.
B . IPv4, IPv6, CIDRs, DNS names is the correct answer. These are the valid formats that can be used with the acceptFrom argument. Wildcards are not supported in acceptFrom parameters for security reasons, as they would allow overly broad access.
Splunk Documentation
Reference:
acceptFrom Parameter Usage

QUESTION 23
At what point in the indexing pipeline set is SEDCMD applied to data?

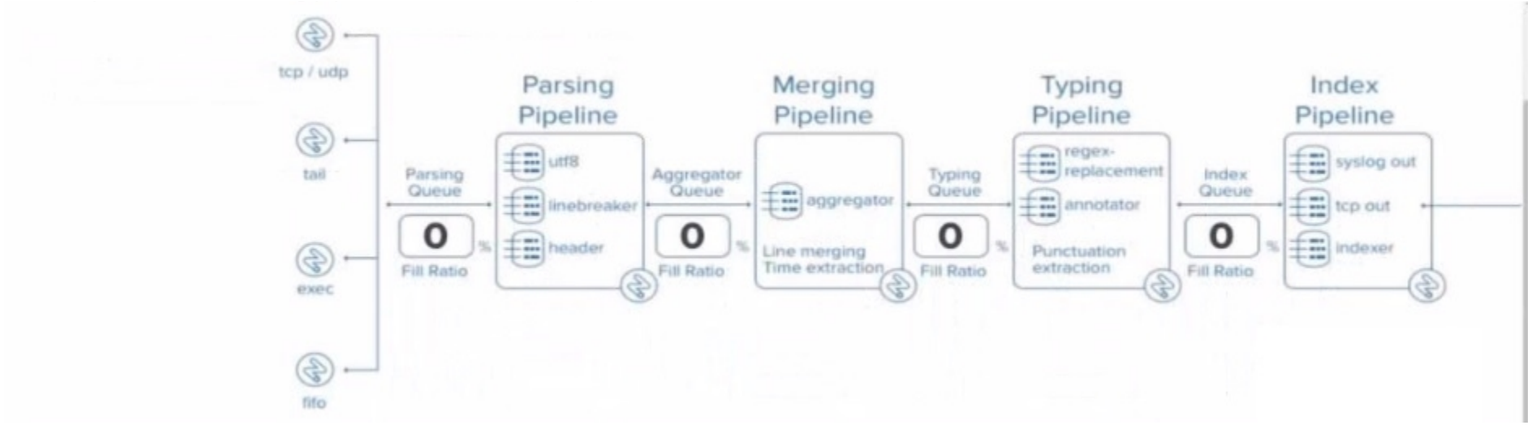

A. In the aggregator queue
B. In the parsing queue
C. In the exec pipeline
D. In the typing pipeline

Correct Answer: D
Section:
Explanation:
In Splunk, SEDCMD (Stream Editing Commands) is applied during the Typing Pipeline of the data indexing process. The Typing Pipeline is responsible for various tasks, such as applying regular expressions for field extractions, replacements, and data transformation operations that occur after the initial parsing and aggregation steps.
Here's how the indexing process works in more detail:
Parsing Pipeline: In this stage, Splunk breaks incoming data into events, identifies timestamps, and assigns metadata.
Merging Pipeline: This stage is responsible for merging events and handling time-based operations.
Typing Pipeline: The Typing Pipeline is where SEDCMD operations occur. It applies regular expressions and replacements, which is essential for modifying raw data before indexing. This pipeline is also responsible for field extraction and other similar operations.
Index Pipeline: Finally, the processed data is indexed and stored, where it becomes available for searching.
Splunk Cloud
Reference: To verify this information, you can refer to the official Splunk documentation on the data pipeline and indexing process, specifically focusing on the stages of the indexing pipeline and the roles they play. Splunk Docs often discuss the exact sequence of operations within the pipeline, highlighting when and where commands like SEDCMD are applied during data processing.
Source:
Splunk Docs: Managing Indexers and Clusters of Indexers
Splunk Answers: Community discussions and expert responses frequently clarify where specific operations occur within the pipeline.

QUESTION 24
When monitoring directories that contain mixed file types, which setting should be omitted from inputs, conf and instead be overridden in propo.conf?

A. sourcetype
B. host
C. source
D. index

Correct Answer: A
Section:
Explanation:
When monitoring directories containing mixed file types, the sourcetype should typically be overridden in props.conf rather than defined in inputs.conf. This is because sourcetype is meant to classify the type of data being ingested, and when dealing with mixed file types, setting a single sourcetype in inputs.conf would not be effective for accurate data classification. Instead, you can use props.conf to define rules that apply different sourcetypes based on the file path, file name patterns, or other criteria. This allows for more granular and accurate assignment of sourcetypes, ensuring the data is properly parsed and indexed according to its type.
Splunk Cloud
Reference: For further clarification, refer to Splunk's official documentation on configuring inputs and props, especially the sections discussing monitoring directories and configuring sourcetypes.
Source:
Splunk Docs: Monitor files and directories
Splunk Docs: Configure event line breaking and input settings with props.conf

QUESTION 25
How are HTTP Event Collector (HEC) tokens configured in a managed Splunk Cloud environment?

A. Any token will be accepted by HEC, the data may just end up in the wrong index.
B. A token is generated when configuring a HEC input, which should be provided to the application developers.
C. Obtain a token from the organization's application developers and apply it in Settings > Data Inputs > HTTP Event Collector > New Token.
D. Open a support case for each new data input and a token will be provided.

Correct Answer: B
Section:
Explanation:

In a managed Splunk Cloud environment, HTTP Event Collector (HEC) tokens are configured by an administrator through the Splunk Web interface. When setting up a new HEC input, a unique token is automatically generated. This token is then provided to application developers, who will use it to authenticate and send data to Splunk via the HEC endpoint.
This token ensures that the data is correctly ingested and associated with the appropriate inputs and indexes. Unlike the other options, which either involve external tokens or support cases, option B reflects the standard procedure for configuring HEC tokens in Splunk Cloud, where control over tokens remains within the Splunk environment itself.
Splunk Cloud
Reference: Splunk's documentation on HEC inputs provides detailed steps on creating and managing tokens within Splunk Cloud. This includes the process of generating tokens, configuring data inputs, and distributing these tokens to application developers.
Source:
Splunk Docs: HTTP Event Collector in Splunk Cloud Platform
Splunk Docs: Create and manage HEC tokens

QUESTION 26
A user has been asked to mask some sensitive data without tampering with the structure of the file /var/log/purchase/transactions. log that has the following format:

```
2020-01-01 00:01:20 User-bob SuperSecretNumber-123456789012 Operation-purchase
2020-01-01 16:15:32 User=alice SuperSecretNumber=123456789012 Operation=purchase
```

A)

In props.conf:
```
[source::/var/log/purchases/transactions.log]
REGEX = (SuperSecretNumber=)\d{12}
DEST_KEY = _raw
FORMAT = $1xxxxxxxxxxx
```

B)

In props.conf:
```
[source::/var/log/purchases/transactions.log]
TRANSFORMS-cleanup = remove_sensitive_data
```

In transforms.conf:
```
[remove_sensitive_data]
REGEX = (SuperSecretNumber=)\d{12}
DEST_KEY = _raw
FORMAT = $1xxxxxxxxxxx
```

C)

In props.conf:
```
[source::/var/log/purchases/transactions.log]
TRANSFORMS-cleanup = remove_sensitive_data
```

In transforms.conf:
```
[remove_sensitive_data]
REGEX = (SuperSecretNumber=)\d{12}
DEST_KEY = _raw
FORMAT = SuperSecretNumber::$1
```

D)

In props.conf:
```
[source::/var/log/purchases/transactions.log]
TRANSFORMS-cleanup = remove_sensitive_data
```

In transforms.conf:
```
[remove_sensitive_data]
REGEX = (SuperSecretNumber=\d{12})
DEST_KEY = _raw
FORMAT = xxxxxxxxxxx
```

A. Option A

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

Correct Answer: B
Section:
Explanation:
Option B is the correct approach because it properly uses a TRANSFORMS stanza in props.conf to reference the transforms.conf for removing sensitive data. The transforms stanza in transforms.conf uses a regular expression (REGEX) to locate the sensitive data (in this case, the SuperSecretNumber) and replaces it with a masked version using the FORMAT directive.
In detail:
props.conf refers to the transforms.conf stanza remove_sensitive_data by setting TRANSFORMS-cleanup = remove_sensitive_data.
transforms.conf defines the regular expression that matches the sensitive data and specifies how the sensitive data should be replaced in the FORMAT directive.
This approach ensures that sensitive information is masked before indexing without altering the structure of the log files.
Splunk Cloud
Reference: For further reference, you can look at Splunk's documentation regarding data masking and transformation through props.conf and transforms.conf.
Source:
Splunk Docs: Anonymize data
Splunk Docs: Props.conf and Transforms.conf

QUESTION 27
Which of the following are valid settings for file and directory monitor inputs?
A)

```
host, index, source_length, _TCP_Routing, host_segment
```

B)

```
host, index, sourcetype, _TCP_Routing, host_regex, host_segment
```

C)

```
host, index, directory, host_regex, host_segment
```

D)

```
host, index, sourcetype, _UDP_Routing, host_regex, host_segment
```

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

Correct Answer: B
Section:
Explanation:
In Splunk, when configuring file and directory monitor inputs, several settings are available that control how data is indexed and processed. These settings are defined in the inputs.conf file. Among the given options:
host: Specifies the hostname associated with the data. It can be set to a static value, or dynamically assigned using settings like host_regex or host_segment.
index: Specifies the index where the data will be stored.
sourcetype: Defines the data type, which helps Splunk to correctly parse and process the data.
TCP_Routing: Used to route data to specific indexers in a distributed environment based on TCP routing rules.
host_regex: Allows you to extract the host from the path or filename using a regular expression.
host_segment: Identifies the segment of the directory structure (path) to use as the host.
Given the options:

Option B is correct because it includes host, index, sourcetype, TCP_Routing, host_regex, and host_segment. These are all valid settings for file and directory monitor inputs in Splunk.
Splunk Documentation
Reference:
Monitor Inputs (inputs.conf)
Host Setting in Inputs
TCP Routing in Inputs
By referring to the Splunk documentation on configuring inputs, it's clear that Option B aligns with the valid settings used for file and directory monitoring, making it the correct choice.

QUESTION 28
Which of the following is not a path used by Splunk to execute scripts?

A. SPLUNK_HOME/etc/system/bin
B. SPLUNK HOME/etc/appa/<app name>/bin
C. SPLUNKHOMS/ctc/scripts/local
D. SPLUNK_HOME/bin/scripts

Correct Answer: C
Section:
Explanation:
Splunk executes scripts from specific directories that are structured within its installation paths. These directories typically include:
SPLUNK_HOME/etc/system/bin: This directory is used to store scripts that are part of the core Splunk system configuration.
SPLUNK_HOME/etc/apps//bin: Each Splunk app can have its own bin directory where scripts specific to that app are stored.
SPLUNK_HOME/bin/scripts: This is a standard directory for storing scripts that may be globally accessible within Splunk's environment.
However, C. SPLUNKHOMS/ctc/scripts/local is not a recognized or standard path used by Splunk for executing scripts. This path does not adhere to the typical directory structure within the SPLUNK_HOME environment, making it the correct answer as it does not correspond to a valid script execution path in Splunk.
Splunk Documentation
Reference:
Using Custom Scripts in Splunk
Directory Structure of SPLUNK_HOME

QUESTION 29
Which of the following are features of a managed Splunk Cloud environment?

A. Availability of premium apps, no IP address whitelisting or blacklisting, deployed in US East AWS region.
B. 20GB daily maximum data ingestion, no SSO integration, no availability of premium apps.
C. Availability of premium apps, SSO integration, IP address whitelisting and blacklisting.
D. Availability of premium apps, SSO integration, maximum concurrent search limit of 20.

Correct Answer: C
Section:
Explanation:
In a managed Splunk Cloud environment, several features are available to ensure that the platform is secure, scalable, and meets enterprise requirements. The key features include:
Availability of premium apps: Splunk Cloud supports the installation and use of premium apps such as Splunk Enterprise Security, IT Service Intelligence, etc.
SSO Integration: Single Sign-On (SSO) integration is supported, allowing organizations to leverage their existing identity providers for authentication.
IP address whitelisting and blacklisting: To enhance security, managed Splunk Cloud environments allow for IP address whitelisting and blacklisting to control access.
Given the options:
Option C correctly lists these features, making it the accurate choice.
Option A incorrectly states 'no IP address whitelisting or blacklisting,' which is indeed available.
Option B mentions 'no SSO integration' and 'no availability of premium apps,' both of which are inaccurate.

Option D talks about a 'maximum concurrent search limit of 20,' which does not represent the standard limit settings and may vary based on the subscription level.
Splunk Documentation
Reference:
Splunk Cloud Features and Capabilities
Single Sign-On (SSO) in Splunk Cloud
Security and Access Control in Splunk Cloud

QUESTION 30
Which of the following statements is true about data transformations using SEDCMD?

A. Can only be used to mask or truncate raw data.
B. Configured in props.conf and transform.conf.
C. Can be used to manipulate the sourcetype per event.
D. Operates on a REGEX pattern match of the source, sourcetype, or host of an event.

Correct Answer: A
Section:
Explanation:
SEDCMD is a directive used within the props.conf file in Splunk to perform inline data transformations. Specifically, it uses sed-like syntax to modify data as it is being processed.
A . Can only be used to mask or truncate raw data: This is the correct answer because SEDCMD is typically used to mask sensitive data, such as obscuring personally identifiable information (PII) or truncating parts of data to ensure privacy and compliance with security policies. It is not used for more complex transformations such as changing the sourcetype per event.
B . Configured in props.conf and transform.conf: Incorrect, SEDCMD is only configured in props.conf.
C . Can be used to manipulate the sourcetype per event: Incorrect, SEDCMD does not manipulate the sourcetype.
D . Operates on a REGEX pattern match of the source, sourcetype, or host of an event: Incorrect, while SEDCMD uses regex for matching patterns in the data, it does not operate on the source, sourcetype, or host specifically.
Splunk Documentation
Reference:
SEDCMD Usage
Mask Data with SEDCMD

QUESTION 31
Which of the following tasks is not managed by the Splunk Cloud administrator?

A. Forwarding events to Splunk Cloud.
B. Upgrading the indexer's Splunk software.
C. Managing knowledge objects.
D. Creating users and roles.

Correct Answer: B
Section:
Explanation:
In Splunk Cloud, several administrative tasks are managed by the Splunk Cloud administrator, but certain tasks related to the underlying infrastructure and core software management are handled by Splunk itself.
B . Upgrading the indexer's Splunk software is the correct answer. Upgrading Splunk software on indexers is a task that is managed by Splunk's operations team, not by the Splunk Cloud administrator. The Splunk Cloud administrator handles tasks like forwarding events, managing knowledge objects, and creating users and roles, but the underlying software upgrades and maintenance are managed by Splunk as part of the managed service.
Splunk Documentation
Reference:
Splunk Cloud Administration

QUESTION 32
What is a private app?

A. An app where only a specific role has read and write access.
B. An app that is only viewable by a specific user.
C. An app that is created and used only by a specific organization.
D. An app where only a specific role has read access.

Correct Answer: C
Section:
Explanation:
A private app in Splunk is one that is created and used within a specific organization, and is not publicly available in the Splunkbase app store.
C . An app that is created and used only by a specific organization is the correct answer. This type of app is developed internally and used by a particular organization, often tailored to meet specific internal needs. It is not shared with other organizations and remains private within that organization's Splunk environment.
Splunk Documentation
Reference:
Private Apps in Splunk