ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

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

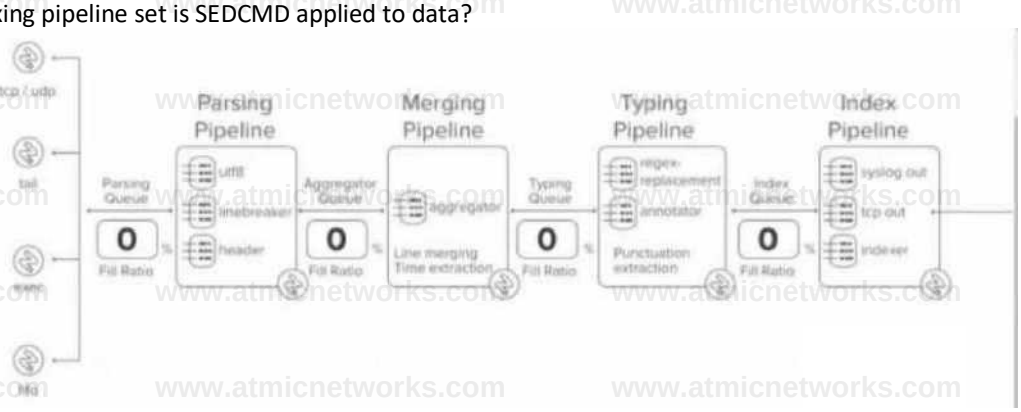

A. In the aggregator queue
B. In the parsing queue
C. In the exec pipeline
D. In the typing pipeline

Answer: D

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: 2

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

Answer: A

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: 3

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.

Answer: B

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: 4

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.

Answer: B

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. Noncertified 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: 5

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.

Answer: B

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: 6

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.

Answer: B

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: 7

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

lw,;;-|.w& my'Veb!:?™t,sswpXe 443 lei/Jef/K-L; .-r :-MS 6430) "GST / HTTP/1.1" *" ZOO Mi " • ■•HowiWb.i) (Windows NT 10.0;, Winf-1; X64) 4ppUMMtt.V537.36 (KHTHL. ilk* - k J Clr"rs/M.0.3112.113 SWfori/'-'".-'•••' v,;);i(i]7Q5

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.

Answer: D

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: 8

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.

Answer: A

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: 9

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:

2020-0.-31 C0:01;20 Qser—bob Sup*rSe*r*t^irrber—l231ʳt"'ᵉ90l2 operation-purchase

2020-0 -01 It: ʳ.:V Usoralicc SupwrSRcr^tNuitbar^l zlJSr.lfi'iC 2 Oprirat i on-purchase

A)

In props.conf.

.mures::/var/1 nq/pnr*hasns/transact?ons. 1 eg]

k£'?KX " (Super Jn-rrctNiutb-'-: •) \di 1 2)

DEST_KEY = _raw

FORMS! - ’ _;<-<xxxx>;xxxxx

B)

In props.conf:

.. AIJT; ; ' AL /Icy ‘'L'UXchdsea/LxdnadcLivris. luyl

TKMJ3F- RMS cleanup = rem’.'e ‘ensitive jota

In transforms.conf:

ram'vc sensitive-data|

■ . SUp : ' ’ i .1'12)

DEST_KEY ~ raw

FORMS! - -Ixxxxxxzxxxx

C)

In props.conf:

i* , '.’a ■/1-ı'ı ;»if haxe.-•‘rratitd ’ •■.', ;]

i IF'JkMS-cleanup ~ ramova sensitive lata

In transforms.conf:

.rente,■ _ ' lienal Live dalu]

REGEX - (SupeiSecietNumbei-i X'Jllli

DESTKEY - raw

FORMAT - SnperSberdtNuMbarx:91

D)

In props.conf:

; s .ur -^:; /var/1 g/pin ha.te- /t r ansaet e o? . It - i] TP-MISFOPHS-i'leanui ~ rcn>-ⁱv” Lens-t-'<_ taLa

In transforms.conf

. rent- , e ■.eabill. e data]

REGEX - (Stipeicv.ietNrrcl*!- ^ *

DEST KEY - raw

'-ill ■

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

Answer: B

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: 10

Which of the following are valid settings for file and directory monitor inputs?

A)

hri't, i ndrx, s'I-ITCH Tanqtk, T IF b. uiting, heist sngntairt

B)

hci.it, indexF sourefttype, TCP Routing host re^e*, host seqresnt

C)

b ^'., Lt. !■. ?:F I .::■'.,, L . ji.h, huiL ⅼ ri ⅼ

D)

Lost, index, a-xiiceiTpa,^UDFJxoiiting^ hoat_ieqe?t, hijst-Seiqmefit

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

Answer: B

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: 11

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

Answer: C

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/<app name>/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: 12

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.

Answer: C

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: 13

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.

Answer: A

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 s ourcetype.
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: 14

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.

Answer: D

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: 15

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.

Answer: D

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: 16

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

Answer: D

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: 17

Consider the following configurations:

'■.:< L'JHEJnEEE it.:.if,,. u:trx/. i. Ji[at. , :-.M

n 'i r i : ni i <• --sr-. . -i
AniTcwrYrr- ■■■■■■■ r "Trn.'d
; ndejc ~ sec ui i Ly

SSflWjiQHE ^#U / app» / # - 3 L I. /1 uv di / lap nit. c Jill

L..:LiL. - ' ;.n/iuu /jecait, I.M I
host = Icqsvri
.*1. lb '■ ”=itypa I itiU x JIR “l: re

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

Answer: C

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: 18

A monitor has been created in inputs. con: for a directory that contains a mix of file types.

How would a Cloud Admin fine-tune assigned sourcetypes for different files in the directory during the input phase?

A. On the Indexer parsing the data, leave sourcetype as automatic for the directory monitor. Then create a props.conf that assigns a specific sourcetype by source stanza.

B. On the forwarder collecting the data, leave sourcetype as automatic for the directory monitor.
Then create a props. conf that assigns a specific sourcetype by source stanza.

C. On the Indexer parsing the data, set multiple sourcetype_source attributes for the directory monitor collecting the files. Then create a props, com that filters out unwanted files.

D. On the forwarder collecting the data, set multiple 3ourcotype_sourc« attributes for the directory monitor collecting the files. Then create a props. conf that filters out unwanted files.

Answer: B

Explanation:

When dealing with a directory containing a mix of file types, it's essential to fine-tune the sourcetypes for different files to ensure accurate data parsing and indexing.

B . On the forwarder collecting the data, leave sourcetype as automatic for the directory monitor. Then create a props.conf that assigns a specific sourcetype by source stanza: This is the correct answer. In this approach, the Universal Forwarder is set up with a directory monitor where the sourcetype is initially left as automatic. Then, a props.conf file is configured to specify different sourcetypes based on the source (filename or path). This ensures that as the data is collected, it is appropriately categorized by

sourcetype according to the file type.
Splunk Documentation Reference:
Configuring Inputs and Sourcetypes
Fine-tuning sourcetypes

Question: 19

Windows Input types are collected in Splunk via a script which is configurable using the GUI. What is this type of input called?

A. Batch
B. Scripted
C. Modular
D. Front-end

Answer: C

Explanation:
Windows inputs in Splunk, particularly those that involve more advanced data collection capabilities beyond simple file monitoring, can utilize scripts or custom inputs. These are typically referred to as Modular Inputs.
C . Modular: This is the correct answer. Modular Inputs are designed to be configurable via the Splunk Web UI and can collect data using custom or predefined scripts, handling more complex data collection tasks. This is the type of input that is used for collecting Windows-specific data such as Event Logs, Performance Monitoring, and other similar inputs.
Splunk Documentation Reference:
Modular Inputs
Windows Data Collection

Question: 20

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

Answer: D

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: 21

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

Answer: B

Explanation:
The props.conf file is a crucial configuration file in Splunk that is used for both search-time and indextime 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: 22

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

Answer: C

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: 23

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

Answer: B

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: 24

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.

Answer: B

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: 25

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.

Answer: C

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

Question: 26

Which of the following is true when using Intermediate Forwarders?

A. Intermediate Forwarders may be a mix of Universal and Heavy Forwarders.
B. All Intermediate Forwarders must be Heavy Forwarders.
C. Intermediate Forwarders may be Universal Forwarders or Heavy Forwarders, but may not be mixed.
D. All Intermediate Forwarders must be Universal Forwarders.

Answer: B

Explanation:

Intermediate Forwarders are special types of forwarders that sit between Universal Forwarders and indexers to perform additional processing tasks such as routing, filtering, or load balancing data before it reaches the indexers.

B . All Intermediate Forwarders must be Heavy Forwarders is the correct answer. Heavy Forwarders are the only type of forwarder that can perform the necessary tasks required of an Intermediate Forwarder, such as parsing data, applying transformations, and routing based on specific rules. Universal Forwarders are lightweight and cannot perform these complex tasks, thus cannot serve as Intermediate Forwarders.

Splunk Documentation Reference:

Intermediate Forwarders

Question: 27

When should Splunk Cloud Support be contacted?

A. For scripted input troubleshooting.
B. For all configuration changes.
C. When unable to resolve issues or perform problem isolation.
D. For resizing, license changes, or any purchases.

Answer: C

Explanation:

Splunk Cloud Support should be contacted when issues arise that cannot be resolved internally or when problem isolation has been unsuccessful.

C . When unable to resolve issues or perform problem isolation is the correct answer. Splunk Cloud Support is typically involved when internal troubleshooting has been exhausted, and the issue requires expert assistance or deeper investigation. While scripted input troubleshooting might be handled by internal teams, contacting support for unresolved issues is the appropriate step.

Splunk Documentation Reference:

When to Contact Splunk Support

Question: 28

Which of the following is a valid stanza in props. conf?

A. [sourcetype::linux_secure]
B. [host=nyc25]
C. [host::nyc*]
D. [host:nyc*]

Answer: A

Explanation:
In props.conf, valid stanzas can include source types, hosts, and source specifications. The correct syntax uses colons for specific types, such as source types and hosts, but follows a particular format: A . [sourcetype::linux_secure] is the correct answer. This is a valid stanza format for a source type in props.conf. It indicates that the following configurations apply specifically to the linux_secure source type.

B . [host=nyc25]: Incorrect, the correct format for a host-based stanza uses double colons, not an equal sign.
C . [host::nyc]:* Incorrect, wildcards are not used in this manner within props.conf.
D . [host
]:* Incorrect, the correct format requires double colons for host stanzas.

Splunk Documentation Reference:
props.conf Specification

Question: 29

Where does the regex replacement processor run?

A. Merging pipeline
B. Typing pipeline
C. Index pipeline
D. Parsing pipeline

Answer: D

Explanation:

The regex replacement processor is part of the parsing stage in Splunk's data ingestion pipeline. This stage is responsible for handling data transformations, which include applying regex replacements. D . Parsing pipeline is the correct answer. The parsing pipeline is where initial data transformations, including regex replacement, occur before the data is indexed. This stage processes events as they are parsed from raw data, including applying any regex-based modifications.
Splunk Documentation Reference:
Data Processing Pipelines in Splunk

Question: 30

What is the correct syntax to monitor /apache/too/logo, /apache/bor/logs, and /apache/bar/l/logo?

A)

, IiiuJi_ L._ L : / / / dptiuhtj / ' /1 L*ya ■

B)

I ’ I : -'i - ' f ■ .''-•■i ■ . /n; - h. /i -; i • ’. • -, -| - ■ /bn: '1 ■

C)

i; L : uf^- J.? u.. i

D)

"rt -. T ti. .-'nr-irr.--'f" /Iri^s^ ^j.r Ti^'b-! i ■'11^.% -,*1 ipn I..- ?i-r/1 .■' vp:]

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

Answer: B

Explanation:

In the context of Splunk, when configuring data inputs to monitor specific directories, the correct syntax must match the directory paths accurately and adhere to the format recognized by Splunk. Option A: [monitor:///apache/*/logs] - This syntax would attempt to monitor all directories under /apache/ that contain the word logs, which is not what the question is asking. It is incorrect for the paths given in the question.

Option B: [monitor:///apache/foo/logs, /apache/bar/logs, /apache/bar/1/logs] - This syntax correctly lists the specific paths /apache/foo/logs, /apache/bar/logs, and /apache/bar/1/logs separately. This is the correct answer as it precisely matches the paths given in the question.

Option C: [monitor:///apache/.../logs] - The triple dots syntax (...) is used to match any subdirectories under /apache/. This would monitor all logs directories within any subdirectory structure under /apache/, which again, does not specifically match the paths given in the question.

Option D: [monitor:///apache/foo/logs, /apache/bar/logs, and /apache/bar/1/logs] - This syntax includes the word "and", which is not valid in the Splunk monitor stanza. The syntax should list the paths separated by commas, without additional words.

Thus, Option B is the correct syntax to monitor the specified paths in Splunk.

For additional reference, you can check the official Splunk documentation on monitoring inputs which provides guidelines on how to configure monitoring of files and directories.

Question: 31

In Splunk terminology, what is an index?

A. A data repository that contains raw, compressed data along with psidx files.
B. A data repository that contains raw, compressed data along with tsidx files.
C. A data repository that contains raw, uncompressed data along with psidx files.
D. A data repository that contains raw, uncompressed data along with tsidx files.

Answer: B

Explanation:

In Splunk, an index is a data repository that stores both raw data and associated indexing information. Specifically, the raw data is stored in a compressed format, and the indexing information is stored in tsidx files (time series index files). These tsidx files enable fast searching and retrieval of data based on time. The correct terminology and structure make option B accurate. Splunk Documentation Reference: Splunk Indexes

Question: 32

When monitoring network inputs, there will be times when the forwarder is unable to send data to the indexers. Splunk uses a memory queue and a disk queue. Which setting is used for the disk queue?

A. queueSize
B. maxQeueSize
C. diskQiioiioiiizo
D. persistentQueueSize

Answer: D

Explanation:

When a forwarder is unable to send data to indexers, it queues the data in memory and optionally on disk. The setting used for the disk queue is persistentQueueSize. This configuration defines the size of the disk queue that stores data temporarily on the forwarder when it cannot immediately forward the data to an indexer.

Splunk Documentation Reference: Configure forwarding and receiving in Splunk

Question: 33

Which of the following takes place during the input phase?

A. Splunk annotates data with only 3 metadata keys: host, source, and sourcetype.
B. Splunk sets the character encoding of the data.
C. Splunk looks at the contents of the data to apply the correct source.
D. Splunk breaks data into individual lines.

Answer: B

Explanation:

During the input phase in Splunk, the system processes incoming data by first setting the character encoding of the data. This step ensures that the data is correctly interpreted by Splunk, allowing it to be parsed and processed properly later in the pipeline. Other options describe actions that occur during later phases, such as parsing and indexing.

Splunk Documentation Reference: How data moves through the data pipeline

Question: 34

Which of the following stanzas would enable a TCP input on port 1025, allowing traffic from all IP addresses except 10.5.5.1?

A)

;tqpi// I ' . ■ . I : P ■ i

B)

>nₚr//|^r'
acr npthY.-.rr - [1] 11 , 5 . '.. I

C)

tcp1/710251

D)

• Lup://10Z5]

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

Answer: B

Explanation:

In Splunk, to configure a TCP input on a specific port and restrict traffic from certain IP addresses, you can use the acceptFrom setting. The correct stanza that enables a TCP input on port 1025 and allows traffic from all IP addresses except 10.5.5.1 would look like this: [tcp://1025]

acceptFrom = !10.5.5.1

Here, !10.5.5.1 denotes that traffic from this IP should be denied, while all other IP addresses are allowed. Therefore, Option B is correct.

Splunk Documentation Reference: Inputs.conf - acceptFrom

Question: 35

Which of the following is not considered a best practice for the deployment server?

A. Create small, single-purpose deployment apps.
B. Dedicate a Splunk instance as the deployment server.
C. Use a Linux server as the deployment server.
D. Create large, multi-purpose deployment apps.

Answer: D

Explanation:

In Splunk, it's considered best practice to create small, single-purpose deployment apps rather than large, multi-purpose ones. This approach ensures better manageability, easier updates, and clearer version control. Option D, which suggests creating large, multi-purpose deployment apps, is not a best practice.

Splunk Documentation Reference: Deployment Server Best Practices

Question: 36

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.

Answer: D

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: 37

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 dat

a. 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.

Answer: B

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: 38

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.

Answer: A

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: 39

In which file can the SH0ULD_LINEMERCE setting be modified?

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

Answer: C

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: 40

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

Answer: B

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: 41

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.

Answer: A

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: 42

When is data deleted from a Splunk Cloud index?

A. When buckets roll to frozen, without a defined archive.
B. When data is deleted via the Splunk Cloud Admin GUI.
C. When TA_Delete is downloaded and enabled from SplunkBase.
D. When the daleteindex command is executed from the CLI.

Answer: A

Explanation:
In Splunk Cloud, data is deleted from an index when the buckets roll to the frozen stage and no archive is defined. When data in a bucket reaches the frozen stage, it is deleted unless a frozen-to- archival script is configured to move the data elsewhere. This process is part of the index lifecycle management in Splunk.
Splunk Documentation Reference: Managing Indexes

Question: 43

What is the recommended method to test the onboarding of a new data source before putting it in production?

A. Send test data to a test index.
B. Send data to the associated production index.
C. Replicate Splunk deployment in a test environment.
D. Send data to the chance index.

Answer: A

Explanation:
The recommended method to test the onboarding of a new data source before putting it into production is to send test data to a test index. This approach allows you to validate data parsing, field extractions, and indexing behavior without affecting the production environment or data.
Splunk Documentation Reference: Onboarding New Data Sources

Question: 44

Which of the following is an accurate statement about the delete command?

A. The delete command removes events from disk.
B. By default, only admins can run the delete command.
C. Events are virtually deleted by marking them as deleted.
D. Deleting events reclaims disk space.

Answer: C

Explanation:
The delete command in Splunk does not remove events from disk but rather marks them as "deleted" in the index. This means the events are not accessible via searches, but they still occupy space on disk. Only users with the can_delete capability (typically admins) can use the delete
command.
Splunk Documentation Reference: Delete Command

Question: 45

What can be used in a Splunk Cloud environment to create new sourcetypes?

A. Data Preview
B. props. conf can be edited directly from the GUI
C. Splunk's CLI
D. Deployment Server

Answer: A

Explanation:
In a Splunk Cloud environment, the Data Preview feature is used to create and test new sourcetypes. This feature allows you to upload sample data, configure parsing settings, and define sourcetypes interactively without directly editing configuration files like props.conf or using the CLI.
Splunk Documentation Reference: Data Preview

Question: 46

Which of the following tasks is the responsibility of a Splunk Cloud administrator?

A. Configuring deployer
B. Configuring cluster master
C. Configuring indexers
D. Configuring indexes

Answer: D

Explanation:
In Splunk Cloud, configuring indexes is one of the primary responsibilities of a Splunk Cloud administrator. This task includes setting up new indexes, managing retention policies, and configuring index settings as required by the organization's data retention and compliance policies. Other tasks like configuring deployer, cluster master, or indexers are typically handled by Splunk Enterprise administrators, not Splunk Cloud administrators.
Splunk Documentation Reference: Splunk Cloud Administrator Guide

Question: 47

Which statement is true about monitor inputs?

A. Monitor inputs are configured in the monitor, conf file.
B. The ignoreOlderThan option allows files to be ignored based on the file modification time.
C. The crSalt setting is required.
D. Monitor inputs can ignore a file's existing content, indexing new data as it arrives, by configuring the tailProcessor option.

Answer: B

Explanation:
The statement about monitor inputs that is true is that the ignoreOlderThan option allows files to be ignored based on their file modification time. This setting helps prevent Splunk from indexing older data that is not relevant or needed.
Splunk Documentation Reference: Monitor files and directories

Question: 48

Where is the recommended place to deploy input apps that are not permitted on Splunk Cloud?

A. Universal Forwarder or Heavy Forwarder.
B. Heavy Forwarder only.
C. Universal Forwarder only.
D. Apps cannot be installed on on-prem instances.

Answer: A

Explanation:
For input apps that are not permitted on Splunk Cloud, the recommended place to deploy them is on a Universal Forwarder or Heavy Forwarder. These forwarders handle data collection and preprocessing before sending the data to Splunk Cloud. This setup allows organizations to leverage apps and configurations that are not supported directly in the cloud environment.
Splunk Documentation Reference: Forwarding Data to Splunk Cloud

Question: 49

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

■ .....9^10 Ui.^uo astion=new_txan3actiori pi;_n.irr—!55672jlB6‘l635n value=2,55 ccy=®?

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:
2920 ■ lb 11:1 : ■ • : 11 •r:»r.«" crin;i<-: inn /zjjiT^I¹";' :.<xxms:2i va_::e-2.5-. "-.f.i

A)

_ ..::... I i ji : . r^ . . _ I
SELQE it ....>. nun - .-I?'.-, nuti- ., I IJJI.L _.....

B)

. . iroj; 11.. ./ttiinaacti .-:iE-li .]
^C$WH#9 k_:ium = s/ (?eejuw>%# p JHW' * J^. 1.' /Xixxxxxxxxx/g

C)

~: i ■ :: . .. -'i l -1 ? .-■ t -i< .' i]
SECOCHia'sk_num =• s / cc_n 141^ (? \ d p 11 4^ £7 ^ J i 1 ■' / cc_n um^xxxxxxxxx / □

D)

’- ii"- ::... 'r.i -.i.-ia't ⌈ .- ' i]
SE'lCMzj-sraEk nun - ST/{?££ rrmrr^Xd{7J I \d[I//[1]]

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

Answer: A

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: 50

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.

Answer: C

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: 51

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.

Answer: C

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: 52

For the following data, what would be the correct attribute/value oair to use to successfully extract the correct timestamp from all the events?

^fi. 12 ■ '^: 11: ^ .j ¹ . . JiXnK ₜ i“, - ■ I.' i&ageHL [4Q&] !
itartir.^ uc-ii-e .11

EX 12 •’•;21:C; :i 1.1 l.t-xiv.;. I'?.. jm ■. 1:-. iX: .Xt. " . .' : Uf^lzeContr&l'ei: H£S3a£ tracing [ªpnw9v 3'nirc^" ■ ₃<j"=L3iL date" - "2018’09-21 20:10:3? HXiOC"r|

Sep 11 06ill:5tl hosT 1.Qxa5plf?(coin stnrpii'j^rit-[5;7 j ! .1. .-J.¹. . .1 1...I.J-.-UtiLA. L.ttl

A. TIMK_FORMAT = %b %d %H:%M:%S %z
B. DATETIME CONFIG = %Y-%m-%d %H:%M:%S %2
C. TIME_FORMAT = %b %d %H:%M:%S
D. DATETIKE CONFIG = Sb %d %H:%M:%S

Answer: C

Explanation:
The correct attribute/value pair to successfully extract the timestamp from the provided events is

TIME_FORMAT = %b %d %H:%M:%S. This format corresponds to the structure of the timestamps in the provided data:
%b represents the abbreviated month name (e.g., Sep).
%d represents the day of the month.
%H:%M:%S represents the time in hours, minutes, and seconds.
This format will correctly extract timestamps like "Sep 12 06:11:58".
Splunk Documentation Reference: Configure Timestamp Recognition

Question: 53

In what scenarios would transforms.conf be used?

A. Per-Event Index Routing, Applying Event Types, SEOCMD operations B. Per-Event Sourcetype, Per-Event Host Name, Per-Event Index Routing C. Per-Event Host Name, Per-Event Index Rooting, SEDCMD operations D. Per-Event Sourcetype, Per-Event Index Routing, Applying Event Types

Answer: B

Explanation:
transforms.conf is used for various advanced data processing tasks in Splunk, including: Per-Event Sourcetype: Dynamically assigning a sourcetype based on event content.
Per-Event Host Name: Dynamically setting the host field based on event content.
Per-Event Index Routing: Directing specific events to different indexes based on their content. Option B correctly identifies these common uses of transforms.conf.
Splunk Documentation Reference: transforms.conf - Configuration

Question: 54

Which monitor statement will retrieve only files that start with "access" in the directory /opt/log/ww2/?

.1 ^opt/lciij/wwi/

-rw-:--r-“ . :■ t : • •’ 3234 I QI>' May 7 2.2137 accea.**! .

iw _ i _ root i.ut J-k-J? Mu. / 22:37 _-._-_ij.__.

A. [monitor:///opt/lug/.../access]
B. [monitor:///opt/log/www2/access*]
C. [monitor:///opt/log/www2/]
D. [monitor:///opt/log/.../]

Answer: B

Explanation:
The correct monitor statement to retrieve only files that start with "access" in the directory /opt/log/www2/ is [monitor:///opt/log/www2/access*]. This configuration specifically targets files that begin with the name "access" and will match any such files within that directory, such as

"access.log".
Splunk Documentation Reference: Monitor files and directories

Question: 55

Li was asked to create a Splunk configuration to monitor syslog files stored on Linux servers at their organization. This configuration will be pushed out to multiple systems via a Splunk app using the on- prem deployment server.
The system administrators have provided Li with a directory listing for the logging locations on three syslog hosts, which are representative of the file structure for all systems collecting this dat a. An example from each system is shown below:

Host: '?-•-■•; -

He path: / . _ ı.ᴛ . v ii.z. i.-,_ ᴛ ..... ui..

Best; _

File path' /'.'si.'b'.i/i^'WitWsysl nQOX/lfmix ꜱᴘɪ''ɪᴛ*ɪ/«ꜱ'ı’«»".ɪᴠɪʟ)•■ ■•?

Host: u-. :.l ʜʟ:

He path: AaiV j..-'g/Lt< ^ .'i ʟA.: - ʙ^s![^^i/l|ijfu>^s$®isCT/^f^^«-l^t.302^9^® 1

A)

.ɪᴜᴊIL i; /7A-±r/lug/htLs ᴊɪ 1/sy lug ■ /lliiL*_j •.,.•. Lu?' ■ [

B)

tMSit^(y//^¥il^ift|^^fl?b^ mJ*/! i nuM_3flmnei/tyflTog.l&cjJ

C)

.n: ᴊɪ r; /’ /Zwiu/lug/netiwik/* * /lir/ja ^H^ury/jyslug.lug.*]

D)

■munirnr:/ Av.^r/lnVn-Tw^ '"I*/11jM^sii^jc^ 0$hi’"f |

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

Answer: A

Explanation:

The correct monitor statement that will capture all variations of the syslog file paths across different systems is [monitor:///var/log/network/syslog*/linux_secure/*].

This configuration works because:

syslog* matches directories that start with "syslog" (like syslog01, syslog02, etc.).

The wildcard * after linux_secure/ will capture all files within that directory, including different filenames like syslog.log and syslog.log.2020090801.

This setup will ensure that all the necessary files from the different syslog hosts are monitored.

Splunk Documentation Reference: Monitor files and directories

Question: 56

By default, which of the following capabilities are granted to the sc_admin role?

A. indexes_edit, edit token, admin_all_objects, delete_by_keyword
B. indexes_edit, fsh_manage, acs_conf, list_indexesdiscovert
C. indexes_edit, fsh_manage, admin_all_objects can_delete
D. indexes_edit, edit_token_http, admin _all objects, edit limits_conf

Answer: C

Explanation:

By default, the sc_admin role in Splunk Cloud is granted several important capabilities, including: indexes_edit: The ability to create, edit, and manage indexes.

fsh_manage: Manage full-stack monitoring integrations.

admin_all_objects: Full administrative control over all objects in Splunk.
can_delete: The ability to delete events using the delete command.
Option C correctly lists these default capabilities for the sc_admin role.
Splunk Documentation Reference: User roles and capabilities

Question: 57

When adding a directory monitor and specifying a sourcetype explicitly, it applies to all files in the directory and subdirectories. If automatic sourcetyping is used, a user can selectively override it in which file on the forwarder?

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

Answer: B

Explanation:
When a directory monitor is set up with automatic sourcetyping, a user can selectively override the sourcetype assignment by configuring the props.conf file on the forwarder. The props.conf file allows you to define how data should be parsed and processed, including assigning or overriding sourcetypes for specific data inputs.
Splunk Documentation Reference: props.conf configuration

Question: 58

Which of the following methods is valid for creating index-time field extractions?

A. Use the UI to create a sourcetype, specify the field name and corresponding regular expression with capture statement.
B. Create a configuration app with the index-time props.conf and/or transfoms. conf, and upload the

app via UI.
C. Use the CU app to define settings in fields.conf, and restart Splunk Cloud.
D. Use the rex command to extract the desired field, and then save as a calculated field.

Answer: B

Explanation:
The valid method for creating index-time field extractions is to create a configuration app that includes the necessary props.conf and/or transforms.conf configurations. This app can then be uploaded via the UI. Index-time field extractions must be defined in these configuration files to ensure that fields are extracted correctly during indexing.
Splunk Documentation Reference: Index-time field extractions

Question: 59

Which of the following is the default bandwidth limit in the Splunk Universal Forwarder credentials package?

A. 0KBps
B. 256 KBps

C. 512 KBps
D. 1024 KBps

Answer: B

Explanation:
The default bandwidth limit in the Splunk Universal Forwarder is set to 256 KBps. This setting is in place to prevent the forwarder from overwhelming network resources, and it can be adjusted as necessary based on the deployment's specific needs.
Splunk Documentation Reference: Universal Forwarder Configuration

Question: 60

A customer wants to mask unstructured data before sending it to Splunk Cloud. Where should SEBCMD be configured for this?

A. props. conf on a Splunk Cloud search head,
B. props.conf on a Heavy Forwarder.
C. transforms, cent on a Splunk Cloud indexer.
D. props. conf- on a Universal Forwarder.

Answer: B

Explanation:

To mask unstructured data before sending it to Splunk Cloud, the SEDCMD should be configured in the props.conf file on a Heavy Forwarder. The Heavy Forwarder is responsible for data parsing and

transformation before forwarding the data to Splunk Cloud. This ensures that sensitive data is masked before it reaches the indexing stage.
Splunk Documentation Reference: Using SEDCMD to Mask Data

Question: 61

How is the forwarder configuration app for Splunk Cloud obtained?

A. Use the wget URL presented when an sc_admin user logs in for the first time.
B. Download from the email sent to the person listed in the SHIP TO: field when the customer licensed Splunk Cloud.
C. Download from the Splunk Cloud UI under the Universal Forwarder app.
D. Download from Splunkbase using splunk.com credentials.

Answer: C

Explanation:
The forwarder configuration app can be accessed directly through the Splunk Cloud UI in the Universal Forwarder app, which simplifies the deployment process by allowing secure, direct download from the cloud instance. [Reference: Splunk Docs on forwarder setup for Splunk Cloud]

Question: 62

What is the name of the Splunk index that contains the most valuable information for troubleshooting a Splunk issue?

A. _internal
B. lastchanceindex
C. _monitoring
D. defaultdb

Answer: A

Explanation:
The _internal index stores logs that are valuable for troubleshooting, including information about system operations, indexers, and search head logs. This index provides insights necessary to diagnose many common issues. [Reference: Splunk Docs on indexes]

Question: 63

A log file is being ingested into Splunk, and a few events have no date stamp. How would Splunk first try to determine the missing date of the events?

A. Splunk will take the date of a previous event within the log file.
B. Splunk will use the current system time of the Indexer for the date.
C. Splunk will use the date of when the file monitor was created.
D. Splunk will take the date from the file modification time.

Answer: D

Explanation:
When events lack a timestamp, Splunk defaults to using the file modification time, which is accessible metadata for parsing time information if no timestamp is present in the log entry. [Reference: Splunk Docs on timestamp recognition]

Question: 64

Which of the following are default Splunk Cloud user roles?

A. must_delete, power, sc_admin
B. power, user, admin
C. apps, power, sc_admin
D. can delete, users, admin

Answer: B

Explanation:
Default Splunk Cloud roles include power, user, and admin, each with unique permissions suitable for common operational and administrative functions. [Reference: Splunk Docs on user roles in Splunk Cloud]

Question: 65

A customer has worked with their LDAP administrator to configure an LDAP strategy in Splunk. The configuration works, and user Mia can log into Splunk using her LDAP Account. After some time, the Splunk Cloud administrator needs to move Mia from the user role to the power role. How should they accomplish this?

A. Ask the LDAP administrator to move Mia's account to an appropriately mapped LDAP group.
B. Have Mia log into Splunk, then update her own role in user settings.
C. Create a role named Power in Splunk, then map Mia's account to that role.
D. Use the Cloud Monitoring Console app as an administrator to map Mia's account to the power role.

Answer: A

Explanation:
In Splunk Cloud, role-based access controls are managed by mapping LDAP groups to Splunk roles. Therefore, any change in roles should be managed by the LDAP administrator, who can adjust Mia's group to an LDAP group mapped to the power role.
[Reference: Splunk Docs on LDAP integration in Splunk Cloud]

Question: 66

Which configuration shown is used to enable a forwarder as a deployment client of the server 10.1.2.3?

A. [target-broker:deploymentServer] targetUri = 10.1.2.3:9997
B. [target-broker:deploymentserver] targetUri = 10.1.2.3:8089
C. [target-broker:deploymentserver] deploymentserver = 10.1.2.3:9997
D. [target-broker:deploymentserver] deploymentserver = 10.1.2.3:8089

Answer: B

Explanation:
For setting up a deployment client, the correct stanza syntax in inputs.conf includes specifying targetUri with the port 8089, which is the management port for Splunk instances, not the data port 9997. [Reference: Splunk Docs on deployment server configurations]

Question: 67

Which of the following would always require raising a support ticket?

A. Capacity or configuration changes in Splunk Cloud.
B. Search does not return expected results in Splunk Cloud.
C. A user is unable to log into Splunk Cloud.
D. Data is not indexed in Splunk Cloud.

Answer: A

Explanation:
Any modifications in capacity or configurations within Splunk Cloud require an official support ticket, as they are managed by Splunk Cloud support teams to ensure consistent and secure changes.
[Reference: Splunk Docs on Splunk Cloud support requests]

Question: 68

What information is identified during the input phase of the ingestion process?

A. Line breaking and timestamp.
B. A hash of the message payload.
C. Metadata fields like sourcetype and host.
D. SRC and DST IP addresses and ports.

Answer: C

Explanation:
During the input phase, Splunk assigns metadata fields such as sourcetype, host, and source, which are critical for data categorization and routing. [Reference: Splunk Docs on data ingestion stages]

Question: 69

Given the following set of files, which of the monitor stanzas below will result in Splunk monitoring all of the files ending with .log?
Files:
/var/log/www1/secure.log
/var/log/www1/access.log
/var/log/www2/logs/secure.log
/var/log/www2/access.log
/var/log/www2/access.log.1

A. [monitor:///var/log/*/*.log]
B. [monitor:///var/log/.../*.log]
C. [monitor:///var/log/*/*]
D. [monitor:///var/log/.../*]

Answer: B

Explanation:
The ellipsis (...) in [monitor:///var/log/.../*.log] allows Splunk to monitor files ending in .log in all nested directories under /var/log/. [Reference: Splunk Docs on monitor stanza syntax]

Question: 70

Which of the following is a valid method to test if a forwarder can successfully send data to Splunk Cloud?

A. Search the _audit index to confirm whether the forwarder ID was registered.
B. Use oneshot from the CLI on the forwarders, then check to see if those logs show up in the Splunk Cloud environment.
C. On Splunk Cloud UI, click Add Data and upload a test file, then search to see if the logs show up.
D. Ping the inputssl.example.splunkcloud.com to see if it returns the ping.

Answer: B

Explanation:
Using the oneshot command allows a direct check for data reception in the cloud environment. Logs can be verified in the cloud after the forwarder sends them. [Reference: Splunk Docs on testing forwarder data inputs]

Question: 71

Which of the following statements is true regarding sedcmd?

A. SEDCMD can be defined in either props.conf or transforms.conf.
B. SEDCMD does not work on Windows-based installations of Splunk.
C. SEDCMD uses the same syntax as Splunk's replace command.
D. SEDCMD provides search and replace functionality using regular expressions and substitutions.

Answer: D

Explanation:
SEDCMD in props.conf applies regular expressions to modify data as it is ingested. It is useful for transforming raw event data before indexing. [Reference: Splunk Docs on SEDCMD]

Question: 72

How is it possible to test a script from the Splunk perspective before using it within a scripted input?

A. splunk run <scriptname>
B. splunk script <scriptname>
C. ./$SPLUNK_HOME/etc/apps/<app>/bin/<scriptname>
D. splunk cmd <scriptname>

Answer: D

Explanation:
splunk cmd <scriptname> allows running scripts in Splunk's environment for testing purposes. This ensures the script behaves as expected within Splunk's CLI context. [Reference: Splunk Docs on scripted inputs]

Question: 73

What two files are used in the data transformation process?

A. parsing.conf and transforms.conf
B. props.conf and transforms.conf
C. transforms.conf and fields.conf
D. transforms.conf and sourcetypes.conf

Answer: B

Explanation:
props.conf and transforms.conf define data parsing, transformations, and routing rules, making them essential for data transformations. [Reference: Splunk Docs on props.conf and transforms.conf]

Question: 74

Where can an administrator download the Splunk Cloud Universal Forwarder credentials package?

A. Splunk Support.
B. Cloud Monitoring Console forwarder drop-down.
C. Universal Forwarder app in the Splunk Cloud search head.
D. Splunkbase.

Answer: C

Explanation:
The Universal Forwarder credentials package is available in the Splunk Cloud search head's Universal Forwarder app for secure, managed deployment. [Reference: Splunk Docs on Universal Forwarder credentials package]

Question: 75

When creating a new index, which of the following is true about archiving expired events?

A. Store expired events in private AWS-based storage.
B. Expired events cannot be archived.
C. Archive some expired events from an index and discard others.
D. Store expired events on-prem using your own storage systems.

Answer: D

Explanation:
In Splunk Cloud, expired events can be archived to customer-managed storage solutions, such as onpremises storage. This allows organizations to retain data beyond the standard retention period if needed. [Reference: Splunk Docs on data archiving in Splunk Cloud]

Question: 76

Due to internal security policies, a Splunk Cloud administrator cannot send data directly to Splunk Cloud from certain data sources. Additional parsing and API-based data sources also need to be sent to Splunk Cloud. What forwarder type should the Splunk Cloud administrator use to satisfy these requirements within their environment?

A. Syslog-ng server with a universal forwarder
B. Light forwarder as an intermediate forwarder
C. Heavy forwarder as an intermediate forwarder
D. Universal forwarder as an intermediate forwarder

Answer: C

Explanation:
A heavy forwarder is appropriate in this scenario because it can perform additional data parsing, filtering, and routing before forwarding data to Splunk Cloud. This is particularly useful for data that requires preprocessing or cannot be sent directly due to security policies. [Reference: Splunk Docs on forwarder types and capabilities]

Question: 77

Configuration folders named default contain configuration files/settings specified in the Splunk product or default settings specified in apps. Which of the following is recommended to override these settings?

A. It does not matter whether setting overrides are placed in default or local folders. Both are equally acceptable since Splunk will merge all the files together into one runtime model after each restart.
B. Any settings to be overridden should be modified in-place wherever the setting was found originally. For example, if overriding a setting originally found in system/default, it should be overridden there to ensure that the desired value is used by Splunk.
C. Overrides should be placed in a folder named local, ideally within a custom Splunk app. This ensures the overrides are preserved upon product or app upgrade and will also be easier to maintain/support.
D. Try to store all configuration overrides in system/local folder to keep all configurations in one place. This ensures the modification has the highest precedence over all other configuration entries.

Answer: C

Explanation:
Placing configuration overrides in the local folder within a custom app allows for easy maintenance and ensures that these overrides are preserved during upgrades, as files in default are overwritten. [Reference: Splunk Docs on configuration file precedence]

Question: 78

Which of the following is a valid monitor stanza for inputs.conf?

A. [monitor:///var/log/*.log] index = linux sourcetype = access_combined host = 489307057
B. [monitor:\\\var\log\httpd-[0-9].log] index = linux sourcetype = access_combined host = 489307057
C. [monitor:///var/log/httpd-[0-9].log] index = linux sourcetype = access_combined host = 489307057
D. [monitor:\\\var\log*.log] index = linux sourcetype = access_combined host = 489307057

Answer: C

Explanation:
[monitor:///var/log/httpd-[0-9].log] is a valid path and syntax for inputs.conf to monitor files ending in .log under /var/log, with other correct index, sourcetype, and host settings specified. [Reference: Splunk Docs on monitor stanzas]

Question: 79

What is the default port for sending data via HTTP Event Collector to Splunk Cloud?

A. 443
B. 8088

C. 9997
D. 8000

Answer: B

Explanation:
The default port for HTTP Event Collector (HEC) in Splunk Cloud is 8088, which is used for data ingestion via HEC. [Reference: Splunk Docs on HTTP Event Collector settings]

Question: 80

In Splunk Cloud, which of the following statements regarding REST API is true?

A. REST API and Splunk HEC are on the same port.
B. All REST API endpoints are open and available by default.
C. REST API is not available in Splunk Cloud.
D. A subset of REST API endpoints are enabled for customers to manage Splunk.

Answer: D

Explanation:
Splunk Cloud enables only a subset of REST API endpoints for customer use to ensure security and control over the environment, allowing essential functionality while maintaining a secure setup. [Reference: Splunk Docs on REST API access in Splunk Cloud]