Vendor: Splunk

Exam Code: SPLK-1005

Exam Name: Splunk Cloud Certified Admin

Version: DEMO

QUESTION 1
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.
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.

QUESTION 2
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.
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.

QUESTION 3
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:
[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.

QUESTION 4
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. 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.

QUESTION 5
What is the correct syntax to monitor /apache/too/logo, /apache/bor/logs, and /apache/bar/l/logo?

A. (monitor:///apache/*/logs]
B. (monitor:///apache/foo/logs, /apache/bar/logs, /apache/bar/1/logs]
C. [monitor:///apache/.../logs]
D. [monitor:///apache/foo/logs, /apache/bar/logs, and /apache/bar/1/logs]

Answer: C
Explanation:
Splunk's monitor:// input supports wildcards for efficient log monitoring:
* (single-level wildcard) matches only one directory level.
... (multi-level wildcard) matches multiple directory levels.

QUESTION 6
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.

QUESTION 7
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.

QUESTION 8
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.

QUESTION 9
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. [tcp://10.5.5.1:1025]
B. [tcp://1025]
acceptFrom = !10.5.5.1
C. [tcp://1025]
D. [tcp://1025]
acceptFrom = 10.5.5.1

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.

QUESTION 10
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.

QUESTION 11
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.

QUESTION 12
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.

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.

QUESTION 13
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.

QUESTION 14
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.

QUESTION 15
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.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**