



Free Questions for *SPLK-1003* by *dumpssheet*

Shared by *Dunn* on *29-01-2024*

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

When using a directory monitor input, specific source types can be selectively overridden using which configuration file?

Options:

- A- sourcetypes . conf
- B- trans forms . conf
- C- outputs . conf
- D- props . conf

Answer:

D

Explanation:

When using a directory monitor input, specific source types can be selectively overridden using the props.conf file. According to the Splunk documentation¹, "You can specify a source type for data based on its input and source. Specify source type for an input. You can assign the source type for data coming from a specific input, such as /var/log/. If you use Splunk Cloud Platform, use Splunk Web to

define source types. If you use Splunk Enterprise, define source types in Splunk Web or by editing the inputs.conf configuration file." However, this method is not very granular and assigns the same source type to all data from an input. To override the source type on a per-event basis, you need to use the props.conf file and the transforms.conf file². The props.conf file contains settings that determine how the Splunk platform processes incoming data, such as how to segment events, extract fields, and assign source types². The transforms.conf file contains settings that modify or filter event data during indexing or search time². You can use these files to create rules that match specific patterns in the event data and assign different source types accordingly². For example, you can create a rule that assigns a source type of apache_error to any event that contains the word "error" in the first line².

Question 2

Question Type: MultipleChoice

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

Options:

- A- services/collector
- B- data/collector
- C- services/inputs?raw

D- services/data/collector

Answer:

A

Explanation:

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

`https://mysplunkserver.example.com:8088/services/collector`

Question 3

Question Type: MultipleChoice

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

Options:

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

Answer:

B, C

Explanation:

The indexing pipeline and the parsing pipeline are the two pipelines that are responsible for transforming the raw data into events and preparing them for indexing. The indexing pipeline applies index-time settings, such as timestamp extraction, line breaking, host extraction, and source type recognition. The parsing pipeline applies parsing settings, such as field extraction, event segmentation, and event annotation.

Question 4

Question Type: MultipleChoice

A user recently installed an application to index NCINX access logs. After configuring the application, they realize that no data is being ingested. Which configuration file do they need to edit to ingest the access logs to ensure it remains unaffected after upgrade?

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

Options:

- A-** Option A
- B-** Option B
- C-** Option C
- D-** Option D

Answer:

A

Explanation:

This option corresponds to the file path "\$SPLUNK_HOME/etc/apps/splunk_TA_nginx/local/inputs.conf". This is the configuration file that the user needs to edit to ingest the NGINX access logs to ensure it remains unaffected after upgrade. This is explained in the Splunk documentation, which states:

The local directory is where you place your customized configuration files. The local directory is empty when you install Splunk Enterprise. You create it when you need to override or add to the default settings in a configuration file. The local directory is never overwritten during an upgrade.

Question 5

Question Type: MultipleChoice

Load balancing on a Universal Forwarder is not scaling correctly. The forwarder's outputs. and the tcpout stanza are setup correctly. What else could be the cause of this scaling issue? (select all that apply)

Options:

- A-** The receiving port is not properly setup to listen on the right port.
- B-** The inputs . conf'S _SYSZOG_ROVTING is not setup to use the right group names.
- C-** The DNS record used is not setup with a valid list of IP addresses.
- D-** The indexAndForward value is not set properly.

Answer:

A, C

Explanation:

The possible causes of the load balancing issue on the Universal Forwarder are A and C. The receiving port and the DNS record are both factors that affect the ability of the Universal Forwarder to distribute data across multiple receivers. If the receiving port is not properly set up to listen on the right port, or if the DNS record used is not set up with a valid list of IP addresses, the Universal Forwarder might fail to connect to some or all of the receivers, resulting in poor load balancing.

Question 6

Question Type: MultipleChoice

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

Options:

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

Answer:

D

Explanation:

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

Question 7

Question Type: MultipleChoice

When working with an indexer cluster, what changes with the global precedence when comparing to a standalone deployment?

Options:

- A- Nothing changes.
- B- The peer-apps local directory becomes the highest priority.
- C- The app local directories move to second in the priority list.
- D- The system default directory' becomes the highest priority.

Answer:

C

Explanation:

The app local directories move to second in the priority list. This is explained in the Splunk documentation, which states:

In a clustered environment, the precedence of configuration files changes slightly from that of a standalone deployment. The app local directories move to second in the priority list, after the peer-apps local directory. This means that any configuration files in the app local directories on the individual peers are overridden by configuration files of the same name and type in the peer-apps local directory on the master node.

Question 8

Question Type: MultipleChoice

The following stanzas in inputs.conf are currently being used by a deployment client:

```
[udp: //145.175.118.177:1001
```

```
Connection_host = dns
```

```
sourcetype = syslog
```

Which of the following statements is true of data that is received via this input?

Options:

- A-** If Splunk is restarted, data will be queued and then sent when Splunk has restarted.
- B-** Local firewall ports do not need to be opened on the deployment client since the port is defined in inputs.conf.
- C-** The host value associated with data received will be the IP address that sent the data.
- D-** If Splunk is restarted, data may be lost.

Answer:

D

Explanation:

This is because the input type is UDP, which is an unreliable protocol that does not guarantee delivery, order, or integrity of the data packets. UDP does not have any mechanism to resend or acknowledge the data packets, so if Splunk is restarted, any data that was in transit or in the buffer may be dropped and not indexed.

Question 9

Question Type: MultipleChoice

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

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

Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

B

Explanation:

`curl "https://mysplunkserver.example.com:8088/services/collector" \ -H "Authorization: Splunk DF4S7ZE4-3GS1-8SFS-E777-0284GG91PF67" \ -d '{"event": "Hello World"}, {"event": "Hola Mundo"}, {"event": "Hallo Welt"}'`. This is the correct curl command to send multiple events through HTTP Event Collector (HEC), which is a token-based API that allows you to send data to Splunk Enterprise from any application that can make an HTTP request. The command has the following components:

The URL of the HEC endpoint, which consists of the protocol (https), the hostname or IP address of the Splunk server (mysplunkserver.example.com), the port number (8088), and the service name (services/collector).

The header that contains the authorization token, which is a unique identifier that grants access to the HEC endpoint. The token is prefixed with Splunk and enclosed in quotation marks. The token value (DF4S7ZE4-3GS1-8SFS-E777-0284GG91PF67) is an example and should be replaced with your own token value.

The data payload that contains the events to be sent, which are JSON objects enclosed in curly braces and separated by commas. Each event object has a mandatory field called event, which contains the raw data to be indexed. The event value can be a string, a number, a boolean, an array, or another JSON object. In this case, the event values are strings that say hello in different languages.

Question 10

Question Type: MultipleChoice

A non-clustered Splunk environment has three indexers (A,B,C) and two search heads (X, Y). During a search executed on search head X, indexer A crashes. What is Splunk's response?

Options:

- A-** Update the user in Splunk web informing them that the results of their search may be incomplete.
- B-** Repeat the search request on indexer B without informing the user.
- C-** Update the user in Splunk web that their results may be incomplete and that Splunk will try to re-execute the search.
- D-** Inform the user in Splunk web that their results may be incomplete and have them attempt the search from search head Y.

Answer:

A

Explanation:

[This is explained in the Splunk documentation1, which states:](#)

If an indexer goes down during a search, the search head notifies you that the results might be incomplete. The search head does not attempt to re-run the search on another indexer.

To Get Premium Files for SPLK-1003 Visit

<https://www.p2pexams.com/products/splk-1003>

For More Free Questions Visit

<https://www.p2pexams.com/splunk/pdf/splk-1003>

