Splunk

Exam Questions SPLK-1004

Splunk Core Certified Advanced Power User

NEW QUESTION 1
Which element attribute is required for event annotation?

A. <search type="event_annotation">
B. <search style="annotation">
C. <search type=$annotation$>
D. <search type="annotation">

Answer: D

Explanation:
In Splunk dashboards, event annotations are used to add informative overlays on timeline visualizations to mark significant events. The required element attribute to define an event annotation within a dashboard panel is <search type="annotation"> (Option D). This attribute specifies that the search within this element is intended to generate annotations, which are then overlaid on the timeline based on the time and information provided by the search results.

NEW QUESTION 2
If a nested macro expands to a search string that begins with a generating command, what additional syntax is needed?

A. Double tick marks around the nested macro.
B. A comma before the nested macro.
C. Square brackets around the nested macro.
D. A pipe character before the nested macro.

Answer: C

Explanation:
When a nested macro in Splunk expands to a search string that begins with a generating command, square brackets (Option C) are needed around the nested macro. This syntax ensures that the expanded macro is correctly interpreted as part of the overall search command structure. Generating commands in Splunk are those that can start a search pipeline and do not require input from a preceding command, such as search, inputlookup, and datamodel. Encapsulating the nested macro in square brackets allows Splunk to process it as an independent subsearch or command within the larger search query. The other options, including double tick marks, a comma, and a pipe character, do not provide the correct syntax for this purpose.

NEW QUESTION 3
Which is a regex best practice?

A. Use complex expressions rather than simple ones.
B. Avoid backtracking.
C. Use greedy operators (. *) instead of non-greedy operators (. *?).
D. Use * rather than +.

Answer: B

Explanation:
In regex (regular expressions), one of the best practices is to avoid backtracking when possible. Backtracking occurs when the regex engine revisits previous parts of the input string to attempt different permutations of the pattern, which can significantly degrade performance, especially with complex patterns on large inputs. Designing regex patterns to minimize or avoid backtracking can lead to more efficient and faster evaluations.

NEW QUESTION 4
If a search contains a subsearch, what is the order of execution?

A. The order of execution depends on whether either search uses a stats command.
B. The inner search executes first.
C. The otter search executes first.
D. The two searches are executed in parallel.

Answer: B

Explanation:
In a Splunk search containing a subsearch, the inner subsearch executes first (Option B). The result of the subsearch is then passed to the outer search. This is because the outer search often depends on the results of the inner subsearch to complete its execution. For example, a subsearch might be used to identify a list of relevant terms or values which are then used by the outer search to filter or manipulate the main dataset.

NEW QUESTION 5
Which of the following is not a common default time field?

A. date_zone
B. date minute
C. date_year
D. date_day

Answer: A

Explanation:
In Splunk, common default time fields include date_minute, date_year, and date_day, which represent the minute, year, and day parts of event timestamps, respectively. date_zone (Option A) is not recognized as a common default time field in Splunk. The platform typically uses fields like _time and various date_* fields for time-related information but does not use date_zone as a standard time field.

NEW QUESTION 6
Why is the transaction command slow in large splunk deployments?

A. It forces the search to run in fast mode.
B. transaction or runs on each Indexer in parallel.
C. It forces all event data to be returned to the search head.
D. transaction runs a hidden eval to format fields.

Answer: C

Explanation:
The transaction command can be slow in large Splunk deployments because it requires all event data relevant to the transaction to be returned to the search head (Option C). This process can be resource-intensive, especially for transactions that span a large volume of data or time, as it involves aggregating and sorting events across potentially many indexers before the transaction logic can be applied.

NEW QUESTION 7
What default Splunk role can use the Log Event alert action?

A. Power
B. User
C. can_delete
D. Admin

Answer: D

Explanation:
In Splunk, the Admin role (Option D) has the capability to use the Log Event alert action among many other administrative privileges. The Log Event alert action allows Splunk to create an event in an index based on the triggering of an alert, providing a way to log and track alert occurrences over time. The Admin role typically encompasses a wide range of permissions, including the ability to configure and manage alert actions.

NEW QUESTION 8
Which field Is requited for an event annotation?

A. annotation_category
B. _time
C. eventype
D. annotation_label

Answer: B

Explanation:
For an event annotation in Splunk, the required field is time (Option B). The time field specifies the point or range in time that the annotation should be applied to in timeline visualizations, making it essential for correlating the annotation with the correct temporal context within the data.

NEW QUESTION 9
What is an example of the simple XML syntax for a base search and its post-srooess search?

A. <search id="myBaseSearch">, <search base="myBaseSearch">
B. <search globalsearch="myBaseSearch">, <search globalsearch>
C. <panel id="myBaseSearch">, <panel base="myBaseSearch">
D. <search id="myGlobalSearch">, <search base="myBaseSearch">

Answer: A

NEW QUESTION 10
Which statement about tsidx files is accurate?

A. Splunk updates tsidx files every 30 minutes.
B. Splunk removes outdated tsidx files every 5 minutes.
C. A tsidx file consists of a lexicon and a posting list.
D. Each bucket in each index may contain only one tsidx file.

Answer: C

Explanation:
A tsidx file in Splunk is an index file that contains indexed data, and it consists of two main parts: alexicon and a posting list (Option C). The lexicon is a list of unique terms found in the data, and the posting list is a list of references to the occurrences of these terms in the indexed data. This structure allows Splunk to efficiently search and retrieve data based on search terms.

NEW QUESTION 10
What does the query | makeresults generate?

A. A timestamp
B. A results field
C. An error message
D. The results of the previously run search.

Answer: B

Explanation:
The | makeresults command in Splunk generates a single event containing default fields, with theprimary purpose of creating sample data or a placeholder event for testing and development purposes. The most notable field it generates is _time, but it does not create a specific 'results' field per se. However, it's commonly used to create a base event for further manipulation with eval or other commands in search queries for demonstration, testing, or constructing specific scenarios.

NEW QUESTION 12
Which of the following has a schema or structure embedded in the data itself?

A. Dark data
B. Unstructured data
C. Embedded data
D. Self-describing data

Answer: D

Explanation:
Self-describing data (Option D) refers to data that includes information about its own structure or schema within the data itself. This characteristic makes it easier to understand and process the data because the structure and meaning of the data are embedded with the data, reducing the need for external definitions or mappings. Examples of self- describing data formats include JSON and XML, where elements and attributes describe the data they contain.

NEW QUESTION 13
What is a performance improvement technique unique to dashboards?

A. Using stats instead of transaction
B. Using global searches
C. Using report acceleration
D. Using datamodel acceleration

Answer: C

Explanation:
Using report acceleration (Option C) is a performance improvement technique unique to dashboards in Splunk. Report acceleration involves pre-computing the results of a report (which can be a saved search or a dashboard panel) and storing these results in a summary index, allowing dashboards to load faster by retrieving the pre-computed data instead of running the full search each time. This technique is especially useful for dashboards that rely on complex searches or searches over large datasets.

NEW QUESTION 14
Which commands should be used in place of a subsearch if possible?

A. untable and/or xyseries
B. stats and/or eval
C. mvexpand and/or where
D. bin and/or where

Answer: B

Explanation:
Using stats and/or eval commands in place of a subsearch is often recommended for performance optimization in Splunk searches. Subsearches can be resource-intensive and slow, especially when dealing with large datasets or complex search operations. The stats command is versatile and can be used for aggregation, summarization, and calculation of data, often achieving the same goals as a subsearch but more efficiently. The eval command is used for field calculations and conditional evaluations, allowing for the manipulation of search results without the need for a subsearch. These commands, when used effectively, can reduce the processing load and improve the speed of searches.

NEW QUESTION 18
What arguments are required when using the spath command?

A. input, output, index
B. input, output path
C. No arguments are required.
D. field, host, source

Answer: B

NEW QUESTION 23
When running a search, which Splunk component retrieves the individual results?

A. Indexer
B. Search head
C. Universal forwarder
D. Master node

Answer: B

Explanation:
The Search head (Option B) in Splunk architecture is responsible for initiating and coordinating search activities across a distributed environment. When a search

is run, the search head parses the search query, distributes the search tasks to the appropriate indexers (which hold the actual data), and then consolidates the results retrieved by the indexers. The search head is the component that interacts with the user, presenting the final search results

NEW QUESTION 24
When possible, what is the best choice for summarizing data to improve search performance?

A. Us the fieldsummary command.
B. Data model acceleration
C. Report acceleration
D. Summary indexing

Answer: D

NEW QUESTION 26
What is the recommended way to create a field extraction that is both persistent and precise?

A. Use the rex command.
B. Use the Field Extractor and manually edit the generated regular expression.
C. Use the Field Extractor and let it automatically generate a regular expression.
D. Use the erex command.

Answer: B

NEW QUESTION 27
Which command processes a template for a set of related fields?

A. bin
B. xyseries
C. foreach
D. untable

Answer: C

Explanation:
The foreach command in Splunk is used to apply a processing step to each field in a set of related fields, making it ideal for performing repetitive tasks across multiple fields without having to specify each field individually. This command can process a template of commands or functions to apply to each specified field, thereby streamlining operations that need to be applied uniformly across multiple data points.

NEW QUESTION 32
When using the bin command, which argument sets the bin size?

A. mazDataSizeMB
B. max
C. volume
D. span

Answer: D

Explanation:
When using the bin command in Splunk, the span argument is used to set the size of each bin (Option D). The span argument determines the granularity or width of each bin when segmenting data over a time range or numerical field, which is essential for time series analysis, histogram generation, or other aggregated data visualizations.

NEW QUESTION 34
Which of the following best describes the process for tokenizing event data?

A. The event Cats is broken up by values in the punch field.
B. The event data is broken up by major breaker and then broken up further by minor breakers.
C. The event data is broken up by a series of user-defined regex patterns.
D. The event data has all punctuation stripped out and is then space delinked.

Answer: B

Explanation:
The process for tokenizing event data in Splunk is best described as breaking the event data up by major breakers and then further breaking it up by minor breakers (Option B). Major breakers typically identify the boundaries of events, while minor breakers further segment the event data intofields. This hierarchical approach to tokenization allows Splunk to efficiently parse and structure the incoming data for analysis.

NEW QUESTION 37
Which of the following are potential string results returned by the type of function?

A. True, False, Unknown
B. Number, Siring, Bool
C. Number, String, Null
D. Field, Value, Lookup

Answer: C

Explanation:
The typeof function in Splunk returns a string that represents the data type of the evaluated expression. The potential string results include "Number", "String", and "Null" (Option C). These indicate whether the evaluated expression is a numerical value, a string, or a null value, respectively, helping users understand the data types they are working with in their searches andscripts.

NEW QUESTION 40
What is returned when Splunk finds fewer than the minimum matches for each lookup value?

A. The default value NULL until the minimum match threshold is reached.
B. The default match value until the minimum match threshold Is reached.
C. The first match unless the time_field attribute is specified.
D. Only the first match.

Answer: A

Explanation:
When Splunk's lookup feature finds fewer than the minimum matches specified for each lookup value, it returns the default value NULL for those unmatched entries until the minimum match threshold is reached (Option A). This behavior ensures that lookups return consistent and expected results, even when the available data does not meet the specified criteria for a minimum number of matches.

NEW QUESTION 43
How is regex passed to the makemv command?

A. makemv be preceded by the erex command.
B. It is specified by the delim argument.
C. It Is specified by the tokenizer argument.
D. Makemv must be preceded by the rex command.

Answer: B

Explanation:
The regex is passed to the makemv command in Splunk using the delim argument (Option B). This argument specifies the delimiter used to split a single string field into multiple values, effectively creating a multivalue field from a field that contains delimited data.

NEW QUESTION 46
When and where do search debug messages appear to help with troubleshooting views?

A. In the Dashboard Editor, while the search is running.
B. In the Search Job Inspector, after the search completes.
C. In the Search Job Inspector, while the search is running.
D. In the Dashboard Editor, after the search completes.

Answer: C

Explanation:
Search debug messages in Splunk appear in the Search Job Inspector while the search is running (Option C). The Search Job Inspector provides detailed information about a search job, including performance statistics, search job properties, and any messages or warnings generated during the search execution. This tool is invaluable for troubleshooting and optimizing searches, as it offers real-time insights into the search process and potential issues.

NEW QUESTION 51
When would a distributable streaming command be executed on an Indexer?

A. If any of the preceding search commands are executed on the search head.
B. If all preceding search commands are executed on me indexer, and a streamstatscommand is used.
C. If all preceding search commands are executed on the Indexer.
D. If some of the preceding search commands are executed on the indexer, and a Timerchart command is used.

Answer: C

Explanation:
A distributable streaming command would be executed on an indexer if all preceding search commands are executed on the indexer (Option C). Distributable streaming commands are designed to be executed where the data resides, reducing data transfer across the network and leveraging the processing capabilities of indexers. This enhances the overall efficiency and performance of Splunk searches, especially in distributed environments.

NEW QUESTION 53
How can the inspect button be disabled on a dashboard panel?

A. Set inspect.link.disabled to 1
B. Set link.inspect .visible to 0
C. Set link.inspectSearch.visible too
D. Set link.search.disabled to 1

Answer: B

Explanation:

To disable the inspect button on a dashboard panel in Splunk, you can set the link.inspect.visible attribute to 0 (Option B) in the panel's source code. This attribute controls the visibility of the inspect button, and setting it to 0 hides the button, preventing users from accessing the search inspector for that panel.

NEW QUESTION 54
Which of the following would exclude all entries contained in the lookup file baditems. csv from search results?

A. NOT [inputlookup baditems.csv]
B. NOT (lookup baditems.csv OUTPUT item)
C. WHERE item NOT IN (baditems.csv)
D. [NOT inputlookup baditems.csv]

Answer: A

Explanation:
The correct syntax to exclude all entries contained in the lookup file baditems.csv from search results is NOT [inputlookup baditems.csv]. This syntax uses a subsearch with the inputlookup command to retrieve the contents of the baditems.csv lookup file and then uses the NOT operator to exclude those results from the main search. This approach is efficient for filtering out unwanted data based on a predefined list of criteria stored in a lookup file.

NEW QUESTION 58
What are the four types of event actions?

A. stats, target, set, and unset
B. stats, target, change, and clear
C. eval, link, change, and clear
D. eval, link, set, and unset

Answer: C

Explanation:
The four types of event actions in Splunk are eval, link, change, and clear (Option C). These actions can be used in dashboard panel configurations to dynamically interact with or manipulate event data based on user inputs or other criteria. Eval is used for calculating fields, link for creating hyperlinks, change for modifying field values, and clear for removing field values or other data elements.

NEW QUESTION 60
When using a nested search macro, how can an argument value be passed to the inner macro?

A. The argument value may be passed to the outer macro.
B. An argument cannot be used with an inner nested macro.
C. An argument cannot be used with an outer nested macro.
D. The argument value must be specified in the outer macro.

Answer: A

Explanation:
When using a nested search macro in Splunk, an argument value can be passed to the inner macro by specifying the argument in the outer macro's invocation (Option A). This allows the outer macro to accept arguments from the user or another search command and then pass those arguments into the inner macro, enabling dynamic and flexible macro compositions that can adapt based on input parameters.

NEW QUESTION 61
Assuming a standard time zone across the environment, what syntax will always return ewnts from between 2:00am and 5:00am?

A. datehour>-2 AND date_hour<5
B. earliest=-2h@h AND latest=-5h@h
C. time_hour>-2 AND time_hour>-5
D. earliest=2h@ AND latest=5h3h

Answer: B

Explanation:
To always return events from between 2:00 AM and 5:00 AM, assuming a standard time zone across the environment, the correct Splunk search syntax is earliest=-2h@h AND latest=-5h@h (Option B). This syntax uses relative time modifiers to specify a range starting 2 hours ago from the current hour (-2h@h) and ending 5 hours ago from the current hour (-5h@h), effectively capturing the desired time window.

NEW QUESTION 66
......