Splunk.SPLK-3002.vMar-2024.by.Huynoyo.40q

Number: SPLK-3002
Passing Score: 800
Time Limit: 120
File Version: 4.0

Exam Code: SPLK-3002

Exam Name: Splunk IT Service Intelligence Certified Admin

Exam A

QUESTION 1
Which of the following is a characteristic of base searches?

A. Search expression, entity splitting rules, and thresholds are configured at the base search level.
B. It is possible to filter to entities assigned to the service for calculating the metrics for the service's KPIs.
C. The fewer KPIs that share a common base search, the more efficiency a base search provides, and anomaly detection is more efficient.
D. The base search will execute whether or not a KPI needs it.

Correct Answer: B
Section:
Explanation:
A base search is a search definition that can be shared across multiple KPIs that use the same data source. Base searches can improve search performance and reduce search load by consolidating multiple similar KPIs. One of the characteristics of base searches is that it is possible to filter to entities assigned to the service for calculating the metrics for the service's KPIs. This means that you can use entity filtering rules to specify which entities are relevant for each KPI based on the base search results.
Reference:Create KPI base searches in ITSI, [Filter entities for KPIs based on base searches]

QUESTION 2
What are valid ITSI Glass Table editor capabilities? (Choose all that apply.)

A. Creating glass tables.
B. Correlation search creation.
C. Service swapping configuration.
D. Adding KPI metric lanes to glass tables.

Correct Answer: A, C, D
Section:
Explanation:
Create a glass table to visualize and monitor the interrelationships and dependencies across your IT and business services.
The service swapping settings are saved and apply the next time you open the glass table.
You can add metrics like KPIs, ad hoc searches, and service health scores that update in real time against a background that you design. Glass tables show real-time data generated by KPIs and services.
The glass table editor is a tool that allows you to create and edit glass tables in ITSI. Some of the capabilities of the glass table editor are:
Creating glass tables from scratch or from existing templates.
Configuring service swapping on widgets to toggle displaying metrics from different services.
Adding KPI metric lanes to glass tables to show historical trends of KPI values.
The glass table editor does not support correlation search creation, which is a separate feature in ITSI that allows you to create searches that look for relationships between data points and generate notable events.
Reference:Overview of the glass table editor in ITSI, [Configure service swapping on glass tables], [Add KPI metric lanes to glass tables], [Overview of correlation searches in ITSI]

QUESTION 3
Which of the following is an advantage of an adaptive time threshold?

A. Automatically alerting when KPI value patterns change over time.
B. Automatically adjusting thresholds as normal KPI values change over time.
C. Automatically adjusting to holiday schedules.

D. Automatically predicting future degradation of KPI values over time.

Correct Answer: B
Section:
Explanation:
An adaptive time threshold in the context of Splunk IT Service Intelligence (ITSI) refers to the capability of dynamically adjusting threshold values for Key Performance Indicators (KPIs) based on historical data trends and patterns. This feature allows thresholds to evolve as the 'normal' behavior of KPIs changes over time, ensuring that alerts remain relevant and reduce the likelihood of false positives or negatives. The advantage of this approach is that it accommodates for natural fluctuations in KPI values that may occur due to changes in business operations, seasonality, or other factors, without requiring manual threshold adjustments. This makes the monitoring system more resilient and responsive to actual conditions, improving the overall effectiveness of IT operations management.

QUESTION 4
Which of the following best describes an ITSI Glass Table?

A. A view which displays a system topology overlaid with KPI metrics.
B. A view which describes a topology.
C. A dashboard which displays a system topology.
D. A view showing KPI values in a variety of visual styles.

Correct Answer: A
Section:
Explanation:
An ITSI Glass Table provides a customizable, high-level view that can display a system's topology overlaid with real-time Key Performance Indicator (KPI) metrics and service health scores. This visualization tool allows users to create a visual representation of their IT infrastructure, applications, and services, integrating live data to monitor the health and performance of each component in context. The ability to overlay KPI metrics on the system topology enables IT and business stakeholders to quickly understand the operational status and health of various elements within their environment, facilitating more informed decision-making and rapid response to issues.

QUESTION 5
Which of the following statements describe default glass tables in ITSI?

A. The Service Health Score default glass table.
B. There is one default glass table per service.
C. There is one service template default glass table.
D. There are no default glass tables.

Correct Answer: D
Section:
Explanation:
In Splunk IT Service Intelligence (ITSI), glass tables are fully customizable dashboards that provide a visual representation of an organization's IT environment, along with the health and status of services and KPIs. Unlike some pre-configured views or dashboards that might come with default setups in various platforms, ITSI does not provide default glass tables out of the box. Instead, users are encouraged to create their own glass tables tailored to their specific monitoring needs and operational views. This approach ensures that each organization can design glass tables that best represent their unique infrastructure, applications, and service landscapes, providing a more personalized and relevant operational overview.

QUESTION 6
Which of the following is part of setting up a new aggregation policy?

A. Filtering criteria
B. Policy version
C. Review order
D. Module rules

Correct Answer: A
Section:
Explanation:
When setting up a new aggregation policy in Splunk IT Service Intelligence (ITSI), one of the crucial components is defining the filtering criteria. This aspect of the aggregation policy determines which events should be included in the aggregation based on specific conditions or attributes. The filtering criteria can be based on various event fields such as severity, source, event type, and other custom fields relevant to the organization's monitoring strategy. By specifying the filtering criteria, ITSI administrators can ensure that the aggregation policy is applied only to the pertinent events, thus facilitating more targeted and effective event management and reducing noise in the operational environment. This helps in organizing and prioritizing events more efficiently, enhancing the overall incident management process within ITSI.

QUESTION 7
Which of the following is a recommended best practice for ITSI installation?

A. ITSI should not be installed on search heads that have Enterprise Security installed.
B. Before installing ITSI, make sure the Common Information Model (CIM) is installed.
C. Install the Machine Learning Toolkit app if anomaly detection must be configured.
D. Install ITSI on one search head in a search head cluster and migrate the configuration bundle to other search heads.

Correct Answer: A
Section:
Explanation:
One of the recommended best practices for Splunk IT Service Intelligence (ITSI) installation is to avoid installing ITSI on search heads that already have Splunk Enterprise Security (ES) installed. This recommendation stems from potential resource conflicts and performance issues that can arise when both resource-intensive applications are deployed on the same instance. Both ITSI and ES are complex applications that require significant system resources to function effectively, and running them concurrently on the same search head can lead to degraded performance, conflicts in resource allocation, and potential stability issues. It's generally advised to segregate these applications onto separate Splunk instances to ensure optimal performance and stability for both platforms.

QUESTION 8
Which views would help an analyst identify that a memory usage KPI is going critical? (select all that apply)

A. Memory KPI in a glass table.
B. Memory panel of the OS Host Details view in the Operating System module.
C. Memory swim lane in a Deep Dive.
D. Service & KPI tiles in the Service Analyzer.

Correct Answer: A, B, C, D
Section:
Explanation:
To identify that a memory usage KPI is going critical, an analyst can leverage multiple views within Splunk IT Service Intelligence (ITSI), each offering a different perspective or level of detail:
A) Memory KPI in a glass table: A glass table can display the current status of the memory usage KPI, along with other related KPIs and services, providing a high-level overview of system health.
B) Memory panel of the OS Host Details view in the Operating System module: This specific panel within the OS Host Details view offers detailed metrics and trends related to memory usage, allowing for in-depth analysis.
C) Memory swim lane in a Deep Dive: Deep Dives allow analysts to visually track the performance and status of KPIs over time. A swim lane dedicated to memory usage can highlight periods where the KPI goes critical, along with the context of other related KPIs.
D) Service & KPI tiles in the Service Analyzer: The Service Analyzer provides a comprehensive overview of all services and their KPIs. The tiles related to memory usage can quickly alert analysts to critical conditions through color-coded indicators.
Each of these views contributes to a comprehensive monitoring strategy, enabling analysts to detect and respond to critical memory usage conditions from various analytical perspectives.

QUESTION 9
How should entities be handled during the data audit phase of requirements gathering?

A. Entity meta-data for info and aliases should be identified and recorded as requirements.

B. Entities should be noted based upon Service KPI requirements such as 'by host' or 'by product line'.

C. Entities must be identified for every Service KPI defined and recorded in requirements.

D. Entities identified should be included in the entity filtering requirements, such as 'by processId' or 'by host'.

Correct Answer: A
Section:
Explanation:
During the data audit phase of requirements gathering for Splunk IT Service Intelligence (ITSI), it's crucial to identify and record the meta-data for entities, focusing on information (info) and aliases. This step involves understanding and documenting the key attributes and identifiers that describe each entity, such as host names, IP addresses, device types, or other relevant characteristics. These attributes are used to categorize and uniquely identify entities within ITSI, enabling more effective mapping of data to services and KPIs. By meticulously recording this meta-data, organizations ensure that their ITSI implementation is aligned with their specific monitoring needs and infrastructure, facilitating accurate service modeling and event management. This practice is foundational for setting up ITSI to reflect the actual IT environment, enhancing the relevance and effectiveness of the monitoring and analysis capabilities.

QUESTION 10
Which of the following is the best use case for configuring a Multi-KPI Alert?

A. Comparing content between two notable events.

B. Using machine learning to evaluate when data falls outside of an expected pattern.

C. Comparing anomaly detection between two KPIs.

D. Raising an alert when one or more KPIs indicate an outage is occurring.

Correct Answer: D
Section:
Explanation:
A multi-KPI alert is a type of correlation search that is based on defined trigger conditions for two or more KPIs. When trigger conditions occur simultaneously for each KPI, the search generates a notable event. For example, you might create a multi-KPI alert based on two common KPIs: CPU load percent and web requests. A sudden simultaneous spike in both CPU load percent and web request KPIs might indicate a DDOS (Distributed Denial of Service) attack. Multi-KPI alerts can bring such trending behaviors to your attention early, so that you can take action to minimize any impact on performance. Multi-KPI alerts are useful for correlating the status of multiple KPIs across multiple services. They help you identify causal relationships, investigate root cause, and provide insights into behaviors across your infrastructure. The best use case for configuring a multi-KPI alert is to raise an alert when one or more KPIs indicate an outage is occurring, such as when the service health score drops below a certain threshold or when multiple KPIs have critical severity levels.
Reference:Create multi-KPI alerts in ITSI

QUESTION 11
In distributed search, which components need to be installed on instances other than the search head?

A. SA-IndexCreation and SA-ITSI-Licensechecker on indexers.

B. SA-IndexCreation and SA-ITOA on indexers; SA-ITSI-Licensechecker and SA-UserAccess on the license master.

C. SA-IndexCreation on idexers; SA-ITSI-Licensechecker and SA-UserAccess on the license master.

D. SA-ITSI-Licensechecker on indexers.

Correct Answer: A
Section:
Explanation:
SA-IndexCreationis required on all indexers. For non-clustered, distributed environments, copySA-IndexCreationto$SPLUNK_HOME/etc/apps/on individual indexers.
In distributed search, the components that need to be installed on instances other than the search head are SA-IndexCreation and SA-ITSI-Licensechecker on indexers. SA-IndexCreation is an add-on that creates the indexes required by ITSI, such as itsi_summary and itsi_tracked_alerts. SA-ITSI-Licensechecker is an add-on that monitors the license usage of ITSI and generates alerts when the license limit is exceeded or about to expire. These components need to be installed on indexers because they handle the data ingestion and storage functions for ITSI. The other components, such as ITSI app and SA-ITOA, need to be installed on the search head(s) because they handle the search management and presentation functions for ITSI.
Reference:Install IT Service Intelligence in a distributed environment

QUESTION 12
Anomaly detection can be enabled on which one of the following?

A. KPI
B. Multi-KPI alert
C. Entity
D. Service

Correct Answer: A
Section:
Explanation:
A is the correct answer because anomaly detection can be enabled on a KPI level in ITSI. Anomaly detection allows you to identify trends and outliers in KPI search results that might indicate an issue with your system. You can enable anomaly detection for a KPI by selecting one of the two anomaly detection algorithms in the KPI configuration panel.
Reference:Apply anomaly detection to a KPI in ITSI

QUESTION 13
After a notable event has been closed, how long will the meta data for that event remain in the KV Store by default?

A. 6 months.
B. 9 months.
C. 1 year.
D. 3 months.

Correct Answer: A
Section:
Explanation:
By default, notable event metadata is archived after six months to keep the KV store from growing too large.

QUESTION 14
Which of the following is a best practice for identifying the most effective services with which to start an iterative ITSI deployment?

A. Only include KPIs if they will be used in multiple services.
B. Analyze the business to determine the most critical services.
C. Focus on low-level services.
D. Define a large number of key services early.

Correct Answer: B
Section:
Explanation:
A best practice for identifying the most effective services with which to start an iterative ITSI deployment is to analyze the business to determine the most critical services that have the most impact on revenue, customer satisfaction, or other key performance indicators. You can use the Service Analyzer to prioritize and monitor these services.
Reference:Service Analyzer

QUESTION 15
When creating a custom deep dive, what color are services/KPIs in maintenance mode within the topology view?

A. Gray

B. Purple
C. Gear Icon
D. Blue

Correct Answer: A
Section:
Explanation:
When creating a custom deep dive, services or KPIs that are in maintenance mode are shown in gray color in the topology view. This indicates that they are not actively monitored and do not generate alerts or notable events.
Reference:Deep Dives

QUESTION 16
Which deep dive swim lane type does not require writing SPL?

A. Event lane.
B. Automatic lane.
C. Metric lane.
D. KPI lane.

Correct Answer: D
Section:
Explanation:
A KPI lane is a type of deep dive swim lane that does not require writing SPL. You can simply select a service and a KPI from a drop-down list and ITSI will automatically populate the lane with the corresponding data. You can also adjust the threshold settings and time range for the KPI lane.
Reference: [KPI Lanes]

QUESTION 17
Which of the following items apply to anomaly detection? (Choose all that apply.)

A. Use AD on KPIs that have an unestablished baseline of data points. This allows the ML pattern to perform it's magic.
B. A minimum of 24 hours of data is needed for anomaly detection, and a minimum of 4 entities for cohesive analysis.
C. Anomaly detection automatically generates notable events when KPI data diverges from the pattern.
D. There are 3 types of anomaly detection supported in ITSI: adhoc, trending, and cohesive.

Correct Answer: B, C
Section:
Explanation:
Anomaly detection is a feature of ITSI that uses machine learning to detect when KPI data deviates from a normal pattern. The following items apply to anomaly detection:
B) A minimum of 24 hours of data is needed for anomaly detection, and a minimum of 4 entities for cohesive analysis. This ensures that there is enough data to establish a baseline pattern and compare different entities within a service.
C) Anomaly detection automatically generates notable events when KPI data diverges from the pattern. You can configure the sensitivity and severity of the anomaly detection alerts and assign them to episodes or teams.
Reference: [Anomaly Detection]

QUESTION 18
Which of the following is a best practice when configuring maintenance windows?

A. Disable any glass tables that reference a KPI that is part of an open maintenance window.
B. Develop a strategy for configuring a service's notable event generation when the service's maintenance window is open.
C. Give the maintenance window a buffer, for example, 15 minutes before and after actual maintenance work.

D. Change the color of services and entities that are part of an open maintenance window in the service analyzer.

Correct Answer: C
Section:
Explanation:
It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work.
A maintenance window is a period of time when a service or entity is undergoing maintenance operations or does not require active monitoring. It is a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work. This gives the system an opportunity to catch up with the maintenance state and reduces the chances of ITSI generating false positives during maintenance operations. For example, if a server will be shut down for maintenance at 1:00PM and restarted at 5:00PM, the ideal maintenance window is 12:30PM to 5:30PM. The 15- to 30-minute time buffer is a rough estimate based on 15 minutes being the time period over which most KPIs are configured to search data and identify alert triggers.
Reference:Overview of maintenance windows in ITSI

QUESTION 19
In Episode Review, what is the result of clicking an episode's Acknowledge button?

A. Assign the current user as owner.
B. Change status from New to Acknowledged.
C. Change status from New to In Progress and assign the current user as owner.
D. Change status from New to Acknowledged and assign the current user as owner.

Correct Answer: D
Section:
Explanation:
When an episode warrants investigation, the analyst acknowledges the episode, which moves the status fromNewtoIn Progress.
An episode represents a disruption of service operation causing impact to business operations. It is a deduplicated group of notable events occurring as part of a larger sequence, or an incident or period considered in isolation. In Episode Review, you can manage the episodes and their statuses using various actions. One of the actions is Acknowledge, which changes the status of an episode from New to Acknowledged and assigns the current user as the owner. This action indicates that someone is working on resolving the episode and prevents duplicate efforts from other users.
Reference:Overview of Episode Review in ITSI, [Episode actions in Episode Review]

QUESTION 20
Which glass table feature can be used to toggle displaying KPI values from more than one service on a single widget?

A. Service templates.
B. Service dependencies.
C. Ad-hoc search.
D. Service swapping.

Correct Answer: D
Section:
Explanation:
A glass table is a visualization tool that allows you to monitor the interrelationships and dependencies across your IT and business services. You can add metrics like KPIs, ad hoc searches, and service health scores that update in real time against a background that you design. One of the features of glass tables is service swapping, which enables you to toggle displaying KPI values from more than one service on a single widget. You can use service swapping to compare metrics across different services without creating multiple glass tables or widgets.
Reference:Overview of the glass table editor in ITSI, [Configure service swapping on glass tables]

QUESTION 21
Which index is used to store KPI values?

A. itsi_summary_metrics
B. itsi_metrics
C. itsi_service_health
D. itsi_summary

Correct Answer: A
Section:
Explanation:
The IT Service Intelligence (ITSI) metrics summary index,itsi_summary_metrics, is a metrics-based summary index that stores KPI data.
A is the correct answer because the itsi_summary_metrics index is used to store KPI values in ITSI. This index improves the performance of the searches dispatched by ITSI, particularly for very large environments. Every KPI is summarized in both the itsi_summary events index and the itsi_summary_metrics metrics index.
Reference:Overview of ITSI indexes

QUESTION 22
Where are KPI search results stored?

A. The default index.
B. KV Store.
C. Output to a CSV lookup.
D. The itsi_summary index.

Correct Answer: D
Section:
Explanation:
Search results are processed, created, and written to the itsi_summary index via an alert action.
D is the correct answer because KPI search results are stored in the itsi_summary index in ITSI. This index is an events index that stores the results of scheduled KPI searches. Summary indexing lets you run fast searches over large data sets by spreading out the cost of a computationally expensive report over time.
Reference:Overview of ITSI indexes

QUESTION 23
Which ITSI functions generate notable events? (Choose all that apply.)

A. KPI threshold breaches.
B. KPI anomaly detection.
C. Multi-KPI alert.
D. Correlation search.

Correct Answer: A, B, D
Section:
Explanation:
After you configure KPI thresholds, you can set up alerts to notify you when aggregate KPI severities change. ITSI generates notable events in Episode Review based on the alerting rules you configure.
Anomaly detection generates notable events when a KPI IT Service Intelligence (ITSI) deviates from an expected pattern.
Notable events are typically generated by a correlation search.
https://docs.splunk.com/Documentation/ITSI/4.10.1/SI/AboutSI
A, B, and D are correct answers because ITSI can generate notable events when a KPI breaches a threshold, when a KPI detects an anomaly, or when a correlation search matches a defined pattern. These are the main ways that ITSI can alert you to potential issues or incidents in your IT environment.
Reference:Configure KPI thresholds in ITSI,Apply anomaly detection to a KPI in ITSI,Generate events with correlation searches in ITSI

QUESTION 24
Which of the following describes a way to delete multiple duplicate entities in ITSI?

A. Via c CSV upload.
B. Via the entity lister page.
C. Via a search using the | deleteentity command.
D. All of the above.

Correct Answer: D
Section:
Explanation:
D is the correct answer because ITSI provides multiple ways to delete multiple duplicate entities. You can use a CSV upload to overwrite existing entities with new or updated information, or delete them by setting the action field to delete. You can also use the entity lister page to select multiple entities and delete them in bulk. Alternatively, you can use a search command called | deleteentity to delete entities that match certain criteria.
Reference:Create and update entities using a CSV file in ITSI,Delete entities in bulk in ITSI,Delete entities using the | deleteentity command in ITSI

QUESTION 25
Which capabilities are enabled through ''teams''?

A. Teams allow searches against the itsi_summary index.
B. Teams restrict notable event alert actions.
C. Teams restrict searches against the itsi_notable_audit index.
D. Teams allow restrictions to service content in UI views.

Correct Answer: D
Section:
Explanation:
D is the correct answer because teams allow you to restrict access to service content in UI views such as service analyzers, glass tables, deep dives, and episode review. Teams also control access to services and KPIs for editing and viewing purposes. Teams do not affect the ability to search against the itsi_summary index, restrict notable event alert actions, or restrict searches against the itsi_notable_audit index.
Reference:Overview of teams in ITSI

QUESTION 26
Besides creating notable events, what are the default alert actions a correlation search can execute? (Choose all that apply.)

A. Ping a host.
B. Send email.
C. Include in RSS feed.
D. Run a script.

Correct Answer: B, C, D
Section:
Explanation:
Throttling applies to any correlation search alert type, including notable events and actions (RSS feed, email, run script, and ticketing).
B, C, and D are correct answers because they are the default alert actions that a correlation search can execute besides creating notable events. You can configure a correlation search to send an email, include the results in an RSS feed, or run a custom script when the search matches a defined pattern. Ping a host is not a default alert action for correlation searches.
Reference:Configure correlation search settings in ITSI

QUESTION 27
Within a correlation search, dynamic field values can be specified with what syntax?

A. fieldname

B. <fieldname /fieldname>

C. %fieldname%

D. eval(fieldname)

Correct Answer: B
Section:
Explanation:
B is the correct answer because dynamic field values can be specified with <fieldname /fieldname> syntax within a correlation search. This syntax allows you to insert values from fields returned by the correlation search into alert actions such as email subject or body. For example, <host /host> inserts the value of the host field into the email.
Reference: [Use dynamic field values in correlation searches in ITSI]

QUESTION 28
In maintenance mode, which features of KPIs still function?

A. KPI searches will execute but will be buffered until the maintenance window is over.

B. KPI searches still run during maintenance mode, but results go to itsi_maintenance_summary index.

C. New KPIs can be created, but existing KPIs are locked.

D. KPI calculations and threshold settings can be modified.

Correct Answer: A
Section:
Explanation:
It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work. This gives the system an opportunity to catch up with the maintenance state and reduces the chances of ITSI generating false positives during maintenance operations.
A is the correct answer because KPI searches still run during maintenance mode, but the results are buffered until the maintenance window is over. This means that no alerts are triggered during maintenance mode, but once it ends, the buffered results are processed and alerts are generated if necessary. You cannot create new KPIs or modify existing KPIs during maintenance mode.
Reference: [Overview of maintenance windows in ITSI]

QUESTION 29
Which of the following describes enabling smart mode for an aggregation policy?

A. Configure --> Policies --> Smart Mode --> Enable, select ''fields'', click ''Save''

B. Enable grouping in Notable Event Review, select ''Smart Mode'', select ''fields'', and click ''Save''

C. Edit the aggregation policy, enable smart mode, select fields to analyze, click ''Save''

D. Edit the notable event view, enable smart mode, select ''fields'', and click ''Save''

Correct Answer: C
Section:
Explanation:
1. From the ITSI main menu, clickConfiguration>Notable Event Aggregation Policies.
2. Select a custom policy or the Default Policy.
3. Under Smart Mode grouping, enableSmart Mode.
4. ClickSelect fields. A dialog displays the fields found in your notable events from the last 24 hours.
C is the correct answer because smart mode is a feature of aggregation policies that allows ITSI to automatically group notable events based on the fields that have the most impact on the event occurrence. You can enable smart mode for an aggregation policy by editing the policy, selecting the smart mode option, and choosing the fields to analyze. You can also specify a minimum number of events to trigger smart mode and a maximum number of groups to create.

Reference:Configure smart mode for aggregation policies in ITSI

QUESTION 30
Which of the following are the default ports that must be configured on Splunk to use ITSI?

A. SplunkWeb (8405), SplunkD (8519), and HTTP Collector (8628)
B. SplunkWeb (8089), SplunkD (8088), and HTTP Collector (8000)
C. SplunkWeb (8000), SplunkD (8089), and HTTP Collector (8088)
D. SplunkWeb (8088), SplunkD (8089), and HTTP Collector (8000)

Correct Answer: C
Section:
Explanation:
C is the correct answer because ITSI uses the default ports of Splunk Enterprise for its communication and data collection. SplunkWeb uses port 8000, SplunkD uses port 8089, and HTTP Event Collector uses port 8088. These ports can be changed if needed, but they must match the configuration of Splunk Enterprise.
Reference:Ports used by ITSI

QUESTION 31
Which of the following is a good use case regarding defining entities for a service?

A. Automatically associate entities to services using multiple entity aliases.
B. All of the entities have the same identifying field name.
C. Being able to split a CPU usage KPI by host name.
D. KPI total values are aggregated from multiple different category values in the source events.

Correct Answer: A
Section:
Explanation:
Define entities before creating services. When you configure a service, you can specify entity matching rules based on entity aliases that automatically add the entities to your service.
A is the correct answer because defining entities for a service allows you to automatically associate entities to services using multiple entity aliases. Entity aliases are alternative names or identifiers for an entity, such as host name, IP address, MAC address, or DNS name. ITSI matches entity aliases to fields in your data sources and assigns entities to services accordingly. This way, you can avoid manually adding entities to each service and ensure that your services reflect the latest changes in your environment.
Reference:Define entities for a service in ITSI

QUESTION 32
When in maintenance mode, which of the following is accurate?

A. Once the window is over, KPIs and notable events will begin to be generated again.
B. KPIs are shown in blue while in maintenance mode.
C. Maintenance mode slots are scheduled on a per hour basis.
D. Service health scores and KPI events are deleted until the window is over.

Correct Answer: A
Section:
Explanation:
A is the correct answer because when in maintenance mode, KPIs and notable events will begin to be generated again once the window is over. Maintenance mode is a feature of ITSI that allows you to temporarily suspend alerts and health score calculations for a service or an entity during planned maintenance or downtime. During maintenance mode, KPI searches still run, but the results are buffered until the window is over. Once the window is over, the buffered results are processed and alerts and health scores are generated if necessary.

Reference: [Overview of maintenance windows in ITSI]

QUESTION 33
In which index are active notable events stored?

A. itsi_notable_archive
B. itsi_notable_audit
C. itsi_tracked_alerts
D. itsi_tracked_groups

Correct Answer: C
Section:
Explanation:
In Splunk IT Service Intelligence (ITSI), notable events are created and managed within the context of its Event Analytics framework. These notable events are stored in the itsi_tracked_alerts index. This index is specifically designed to hold the active notable events that are generated by ITSI's correlation searches, which are based on the conditions defined for various services and their KPIs. Notable events are essentially alerts or issues that need to be investigated and resolved. The itsi_tracked_alerts index enables efficient storage, querying, and management of these events, facilitating the ITSI's event management and review process. The other options, such as itsi_notable_archive and itsi_notable_audit, serve different purposes, such as archiving resolved notable events and auditing changes to notable event configurations, respectively. Therefore, the correct answer for where active notable events are stored is the itsi_tracked_alerts index.

QUESTION 34
When a KPI's aggregate value is calculated, which function is called?

A. stats
B. tstats
C. fieldsummary
D. eval

Correct Answer: B
Section:
Explanation:
In Splunk IT Service Intelligence (ITSI), when a Key Performance Indicator (KPI) aggregate value is calculated, the tstats function is often called. The tstats function in Splunk is used for rapid statistical queries over large volumes of data, which is particularly useful in ITSI for efficiently calculating aggregate values of KPIs across potentially vast datasets. This function allows for quick aggregation and summarization of indexed data, which is essential for monitoring and analyzing the performance metrics that KPIs represent in ITSI. Unlike the stats command, which operates on already retrieved events, tstats works directly on indexed data, providing faster performance especially when dealing with high volumes of data typical in an IT environment. The tstats command is therefore fundamental in the backend processing of ITSI for calculating aggregate values of KPIs, enabling real-time and historical analysis of service health and performance.

QUESTION 35
Which of the following describes default deep dives?

A. Are manually generated and can be accessed via the Service Analyzer.
B. Include all KPIs of all services.
C. Are auto-generated and can be accessed via the Service Analyzer.
D. Include health scores of all services.

Correct Answer: C
Section:
Explanation:
In Splunk IT Service Intelligence (ITSI), default deep dives are auto-generated and can be accessed via the Service Analyzer. Deep dives are an essential feature of ITSI that provide an in-depth, granular view into the health and

performance of services and their associated KPIs. These default deep dives are automatically created for each service, allowing users to quickly drill down into the detailed operational metrics and performance data of their services. By accessing these deep dives through the Service Analyzer, ITSI users can efficiently investigate issues, understand service dependencies, and make informed decisions to maintain optimal service health. The auto-generated nature of these default deep dives simplifies the monitoring and analysis process, providing immediate insights into service performance without the need for manual setup or configuration.

QUESTION 36
Which of the following is a problem requiring correction in ITSI?

A. Two or more entities with the same service ID.
B. Two or more entities with the same entity ID.
C. Two or more entities with the same value in a single alias field.
D. Two or more entities with the same entity key value in any info field.

Correct Answer: C
Section:
Explanation:
In Splunk IT Service Intelligence (ITSI), entities represent infrastructure components, applications, or other elements that are monitored. Each entity is uniquely identified by its entity ID, and entities can be associated with one or more services through the concept of aliases. A problem arises when two or more entities have the same value in a single alias field because aliases are used to match events to entities in ITSI. If multiple entities share the same alias value, ITSI might incorrectly associate data with the wrong entity, leading to inaccurate monitoring and analytics. This scenario requires correction to ensure that each alias uniquely identifies a single entity, thereby maintaining the integrity of the monitoring and analysis process within ITSI. The uniqueness of service IDs, entity IDs, and entity key values in info fields is also important but does not typically present the same level of issue as duplicate values in an alias field.

QUESTION 37
Which index contains ITSI Episodes?

A. itsi_tracked_alerts
B. itsi_grouped_alerts
C. itsi_notable_archive
D. itsi_summary

Correct Answer: B
Section:
Explanation:
B is the correct answer because ITSI episodes are stored in the itsi_grouped_alerts index. This index contains notable events that have been grouped together based on predefined aggregation policies. Episodes help you reduce alert noise and focus on resolving incidents faster.
Reference: [Overview of episodes in ITSI]

QUESTION 38
Which of the following best describes a default deep dive?

A. It initially shows the health scores for all services.
B. It initially shows the highest importance KPIs.
C. It initially shows all of the KPIs for a selected service.
D. It initially shows all the entity swim lanes.

Correct Answer: C
Section:
Explanation:
C is the correct answer because a default deep dive initially shows all of the KPIs for a selected service. You can create a default deep dive by drilling down from another dashboard or by selecting a service from the deep dive

lister page. A default deep dive does not show health scores, importance scores, or entity swim lanes by default.
Reference: [Create default deep dives for services in ITSI]

QUESTION 39
Which of the following is a good use case for a Multi-KPI alert?

A. Alerting when the values of two or more KPIs go into maintenance mode.
B. Alerting when the trend of two or more KPIs indicates service failure is imminent.
C. Alerting when two or more KPIs are deviating from their typical pattern.
D. Alerting when comparing the values of two or more KPIs indicates an unusual condition is occurring.

Correct Answer: D
Section:
Explanation:
A Multi-KPI alert in Splunk IT Service Intelligence (ITSI) is designed to trigger based on the conditions of multiple Key Performance Indicators (KPIs). This type of alert is particularly useful when a single KPI's state is not sufficient to indicate an issue, but the correlation between multiple KPIs can provide a clearer picture of an emerging problem. The best use case for a Multi-KPI alert is therefore when comparing the values of two or more KPIs indicates an unusual condition is occurring. This allows for more nuanced and context-rich alerting mechanisms that can identify complex issues not detectable by monitoring individual KPIs. This approach is beneficial in complex environments where the interplay between different performance metrics needs to be considered to accurately detect and diagnose issues.

QUESTION 40
Which of the following actions can be performed with a deep dive?

A. Create a Multi-KPI alert from the deep dive's current state to warn of similar situations in the future.
B. Create a predictive analysis model from the deep dive to warn of future service degradation.
C. Create an anomaly detection alert to show when the same pattern begins in the future.
D. Create a custom service analyzer from selected deep dive lanes.

Correct Answer: A
Section:
Explanation:
Deep dives in Splunk IT Service Intelligence (ITSI) allow for an in-depth analysis of services and their KPIs over time, providing a detailed view of the operational health and performance trends. One of the powerful actions that can be performed with a deep dive is the creation of a Multi-KPI alert from the deep dive's current state. This functionality enables users to define alerts based on the complex conditions observed during the deep dive analysis, allowing for the early detection of similar situations in the future. By configuring a Multi-KPI alert directly from a deep dive, ITSI users can leverage their insights and observations to proactively monitor for patterns or conditions that may indicate potential service degradation or failure, enhancing the overall responsiveness and effectiveness of the IT monitoring strategy.