Splunk

Exam Questions SPLK-3002

Splunk IT Service Intelligence Certified Admin Exam

NEW QUESTION 1
Which scenario would benefit most by implementing ITSI?

A. Monitoring of business services functionality.
B. Monitoring of system hardware.
C. Monitoring of system process statuses
D. Monitoring of retail sales metrics.

Answer: A

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/AboutSI
Splunk IT Service Intelligence (ITSI) is a monitoring and analytics solution that uses artificial intelligence and machine learning to provide insights into the health and performance of IT services. ITSI lets you create services that represent the critical components of your IT infrastructure, such as applications, databases, servers, networks, and so on. You can then monitor the status and performance of these services using key performance indicators (KPIs), which are metrics that measure aspects of service health, such as availability, latency, error rate, and so on. ITSI also provides tools for visualizing, investigating, and alerting on service issues, such as service analyzers, glass tables, deep dives, episode review, and so on. The scenario that would benefit most by implementing ITSI is monitoring of business service functionality, because ITSI enables you to measure and improve the quality and reliability of your IT services and align them with your business objectives. References: What is Splunk IT Service Intelligence?

NEW QUESTION 2
How can admins manually control groupings of notable events?

A. Correlation searches.
B. Multi-KPI alerts.
C. notable_event_grouping.conf
D. Aggregation policies.

Answer: D

Explanation:
In Splunk IT Service Intelligence (ITSI), administrators can manually control the grouping of notable events using aggregation policies. Aggregation policies allow for the definition of criteria based on which notable events are grouped together. This includes configuring rules based on event fields, severity, source, or other event attributes. Through these policies, administrators can tailor the event grouping logic to meet the specific needs of their environment, ensuring that related events are grouped in a manner that facilitates efficient analysis and response. This feature is crucial for managing the volume of events and focusing on the most critical issues by effectively organizing related events into manageable groups.

NEW QUESTION 3
Which of the following items describe ITSI Deep Dive capabilities? (Choose all that apply.)

A. Comparing a service??s notable events over a time period.
B. Visualizing one or more Service KPIs values by time.
C. Examining and comparing alert levels for KPIs in a service over time.
D. Comparing swim lane values for a slice of time.

Answer: BCD

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/DeepDives
A deep dive is a dashboard that allows you to analyze the historical trends and anomalies of your KPIs and metrics in ITSI. A deep dive displays a timeline of events and swim lanes of data that you can customize and filter to investigate issues and perform root cause analysis. Some of the capabilities of deep dives are:
* B. Visualizing one or more service KPIs values by time. This is true because you can add KPI swim lanes to a deep dive to show the values and severity levels of one or more KPIs over time. You can also compare KPIs from different services or entities using service swapping or entity splitting.
* C. Examining and comparing alert levels for KPIs in a service over time. This is true because you can add alert swim lanes to a deep dive to show the alert levels and counts for one or more KPIs over time. You can also drill down into the alert details and view the notable events associated with each alert.
* D. Comparing swim lane values for a slice of time. This is true because you can use the time range selector to zoom in or out of a specific time range in a deep dive. You can also use the time brush to select a slice of time and compare the swim lane values for that time period.
The other option is not a capability of deep dives because:
A. Comparing a service??s notable events over a time period. This is not true because deep dives do not display notable events, which are alerts generated byITSI based on certain conditions or correlations. Notable events are displayed in other dashboards, such as episode review or glass tables.
References: [Overview of deep dives in ITSI], [Add swim lanes to a deep dive in ITSI]

NEW QUESTION 4
Within a correlation search, dynamic field values can be specified with what syntax?

A. fieldname
B. <fieldname /fieldname>
C. %fieldname%
D. eval(fieldname)

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/Splunk/8.2.2/Search/Searchindexes
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. References: [Use dynamic field values in correlation searches in ITSI]

NEW QUESTION 5
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.

Answer: A

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/REBestPractice
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. References: [Overview of maintenance windows in ITSI]

NEW QUESTION 6
Where are KPI search results stored?

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

Answer: D

Explanation:
Search results are processed, created, and written to the itsi_summary index via an alert action.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/BaseSearch
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. References: Overview of ITSI indexes

NEW QUESTION 7
Which of the following can generate notable events?

A. Through ad-hoc search results which get processed by adaptive thresholds.
B. When two entity aliases have a matching value.
C. Through scheduled correlation searches which link to their respective services.
D. Manually selected using the Notable Event Review panel.

Answer: C

Explanation:
Notable events in Splunk IT Service Intelligence (ITSI) are primarily generated through scheduled correlation searches. These searches are designed to monitor data for specific conditions or patterns defined by the ITSI administrator, and when these conditions are met, a notable event is created. These correlation searches are often linked to specific services or groups of services, allowing for targeted monitoring and alerting based on the operational needs of those services. This mechanism enables ITSI to provide timely and relevant alerts that can be further investigated and managed through the Episode Review dashboard, facilitating efficient incident response and management within the IT environment.

NEW QUESTION 8
Which of the following items describe ITSI Backup and Restore functionality? (Choose all that apply.)

A. A pre-configured default ITSI backup job is provided that can be modified, but not deleted.
B. ITSI backup is inclusive of KV Store, ITSI Configurations, and index dependencies.
C. kvstore_to_json.py can be used in scripts or command line to backup ITSI for full or partial backups.
D. ITSI backups are stored as a collection of JSON formatted files.

Answer: CD

Explanation:
ITSI provides a kvstore_to_json.py script that lets you backup/restore ITSI configuration data, perform bulk service KPI operations, apply time zone offsets for ITSI objects, and regenerate KPI search schedules.
When you run a backup job, ITSI saves your data to a set of JSON files compressed into a single ZIP file.
Reference:
https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/kvstorejson
https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/BackupandRestoreITSIconfig
C and D are correct answers because ITSI backup and restore functionality uses
kvstore_to_json.py as a command line script or as part of custom scripts to backup ITSI datafor full or partial backups. ITSI backups are also stored as a collection of JSON formatted files that contain KV store objects such as services, KPIs, glass tables, etc. A is not a correct answer because there is no pre-configured default ITSI backup job provided. You can create your own backup jobs or use the command line script or custom scripts to backup ITSI data. B is not a correct answer because ITSI backup is not inclusive of index dependencies. ITSI backup only includes KV store objects and optionally some .conf files. You need to use other methods to backup index data. References: [Overview of backing up and restoring ITSI KV store data], [Create a full backup of ITSI], [Create a partial backup of ITSI]

NEW QUESTION 9
Which ITSI components are required before a module can be created?

A. One or more entity import saved searches.
B. One or more services with KPIs and their associated base searches.
C. One or more datamodels.
D. One or more correlation searches and their associated entities.

Answer: C

Explanation:
Before a module can be created in Splunk IT Service Intelligence (ITSI), it is essential to have one or more datamodels established. Datamodels in Splunk provide a structured format for organizing and interpreting data, which is crucial for modules within ITSI. Modules often rely on datamodels to extract, transform, and present data in a meaningful way, especially when dealing with complex datasets across various sources. Datamodels serve as the foundation for the module's ability to categorize and analyze data efficiently, enabling the creation of KPIs, services, and visualizations that are aligned with the specific needs of the module. Having these datamodels in place ensures that the module can function correctly and provide valuable insights into the monitored IT environments.

NEW QUESTION 10
Which of the following is a recommended best practice for service and glass table design?

A. Plan and implement services first, then build detailed glass tables.
B. Always use the standard icons for glass table widgets to improve portability.
C. Start with base searches, then services, and then glass tables.
D. Design glass tables first to discover which KPIs are important.

Answer: A

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/GTOverview
A is the correct answer because it is recommended to plan and implement services first, then build detailed glass tables that reflect the service hierarchy and dependencies. This way, you can ensure that your glass tables provide accurate and meaningful service-level insights. Building glass tables first might lead to unnecessary or irrelevant KPIs that do not align with your service goals. References: Splunk IT Service Intelligence Service Design Best Practices

NEW QUESTION 10
What effects does the KPI importance weight of 11 have on the overall health score of a service?

A. At least 10% of the KPIs will go critical.
B. Importance weight is unused for health scoring.
C. The service will go critical.
D. It is a minimum health indicator KPI.

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/KPIImportance#:~:text=ITSI%20con
siders%20KPIs%20that%20have,other%20KPIs%20in%20the%20service
The KPI importance weight is a value that indicates how much a KPI contributes to the overall health score of a service. The importance weight can range from 1 (lowest) to 10 (highest). The statement that applies when configuring a KPI importance weight of 11 is:
* B. Importance weight is unused for health scoring. This is true because an importance weight of 11 is invalid and cannot be used for health scoring. The maximum value for importance weight is 10.
The other statements do not apply because:
* A. At least 10% of the KPIs will go critical. This is not true because an importance weight of 11 does not affect the severity level of any KPIs.
* C. The service will go critical. This is not true because an importance weight of 11 does not
affect the health score or status of any service.
* D. It is a minimum health indicator KPI. This is not true because an importance weight of 11 does not indicate anything about the minimum health level of a KPI.
References: Set KPI importance values in ITSI

NEW QUESTION 15
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.

Answer: B

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.

NEW QUESTION 17
What can a KPI widget on a glass table drill down into?

A. Another glass table.
B. A Splunk dashboard.
C. A custom deep dive.
D. Any of the above.

Answer: D

Explanation:
In Splunk IT Service Intelligence (ITSI), a KPI widget on a glass table can be configured to drill down into a variety of destinations based on the needs of the user and the design of the glass table. This flexibility allows users to dive deeper into the data or analysis represented by the KPI widget, providing context and additional insights. The destinations for drill-downs from a KPI widget can include:
* A. Another glass table, offering a different perspective or more detailed view related to the KPI. B. A Splunk dashboard that provides broader analysis or incorporates data frommultiple sources. C. A custom deep dive for in-depth, time-series analysis of the KPI and related metrics.
This versatility makes KPI widgets powerful tools for navigating through the wealth of operational data and insights available in ITSI, facilitating effective monitoring and decision- making.

NEW QUESTION 21
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.

Answer: D

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/MKA
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 twocommon 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. References: Create multi-KPI alerts in ITSI

NEW QUESTION 25
After ITSI is initially deployed for the operations department at a large company, another department would like to use ITSI but wants to keep their information private from the operations group. How can this be achieved?

A. Create service templates for each group and create the services from the templates.
B. Create teams for each department and assign KPIs to each team.
C. Create services for each group and set the permissions of the services to restrict them to each group.
D. Create teams for each department and assign services to the teams.

Answer: D

Explanation:
In Splunk IT Service Intelligence (ITSI), creating teams for each department and assigning services to those teams is an effective way to segregate data and ensure that information remains private between different groups within an organization. Teams in ITSI provide a mechanism for role-based access control, allowing administrators to define which users or groups have access to specific services, KPIs, and dashboards. By setting up teams corresponding to each department and then assigning services to these teams, ITSI canaccommodate multi-departmental use within the same instance while maintaining strict access controls. This ensures that each department can only view and interact with the data and services relevant to their operations, preserving confidentiality and data integrity across the organization.

NEW QUESTION 27
Which step is required to install ITSI on a single Search Head?

A. Untar the ITSI package in <splunk home>/etc/apps
B. Run splunk_apply shcluster-bundle
C. Use the Splunk -> Manage Apps Dashboard to download and install.
D. All of the above.

Answer: C

Explanation:
To install Splunk IT Service Intelligence (ITSI) on a single Search Head, one of the straightforward methods is to use the Splunk Web interface, specifically the "Manage Apps" dashboard, to download and install ITSI. This method is user-friendly and does not require manual file handling or command-line operations. By navigating to "Manage Apps" in the Splunk Web interface, users can find ITSI in the app repository or upload the ITSI installation package if it has been downloaded previously. From there, the installation process is initiated through the Splunk Web interface, simplifying the setup process. This approach ensures that the installation follows Splunk's standard app installation procedures, helping to avoid common installation errors and ensuring that ITSI is correctly integrated into the Splunk environment.

NEW QUESTION 32
When deploying ITSI on a distributed Splunk installation, which component must be installed on the search head(s)?

A. SA-ITOA
B. ITSI app
C. All ITSI components
D. SA-ITSI-Licensechecker

Answer: B

Explanation:
Install SA-ITSI-Licensechecker and SA-UserAccess on any license master in a distributed or search head cluster environment. If a search head in your environment is also a license
master, the license master components are installed when you install ITSI on the search heads.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Install/InstallDD
When deploying ITSI on a distributed Splunk installation, the component that must be installed on the search head(s) is the ITSI app. The ITSI app contains the main features and functionality of ITSI, such as service creation and management, KPI configuration, glass table creation and editing, episode review, deep dives, and so on. The ITSI app also contains some add-ons that provide additional functionality, such as SA-ITOA (IT Operations Analytics), SA-UserAccess (User Access Management), and SA-Utils (Utility Functions). The ITSI app must be installed on the search head(s) because it handles the search management and presentation functions for ITSI. References: Install IT Service Intelligence in a distributed environment

NEW QUESTION 36
Which of the following services often has KPIs but no entities?

A. Security Service.
B. Network Service.
C. Business Service.
D. Technical Service.

Answer: C

Explanation:
In the context of Splunk IT Service Intelligence (ITSI), a Business Service often has Key Performance Indicators (KPIs) but might not have directly associated entities. Business Services represent high-level aggregations of organizational functions or processes and are typically measured by KPIs that reflect the performance of underlying technical services or components rather than direct infrastructure entities. For example, a Business Service might monitor overall transaction completion times or customer satisfaction scores, which are abstracted from the specific technical entities that underlie these metrics. This abstraction allows Business Services to provide a business-centric view of IT health and performance, focusing on outcomes rather than specific technical components.

NEW QUESTION 39
Which index contains ITSI Episodes?

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

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/IndexOverview
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. References: [Overview of episodes in ITSI]

NEW QUESTION 44
Which anomaly detection algorithm is included within ITSI?

A. Entity cohesion
B. Standard deviation
C. Linear regression
D. Infantile regression

Answer: A

Explanation:
Among the anomaly detection algorithms included within Splunk IT Service Intelligence (ITSI), "Entity Cohesion" is a notable option. The Entity Cohesion algorithm is designed to detect anomalies by comparing the behavior of one entity against the collective behavior of a group of similar entities. This approach is particularly useful in scenarios where entities are expected to exhibit similar patterns of behavior under normal conditions. Anomalies are identified when an entity's metrics deviate significantly from the group norm, suggesting a potential issue with thatspecific entity. This method leverages the concept of cohesion among similar entities to enhance the accuracy and relevance of anomaly detection within ITSI environments.

NEW QUESTION 49
Which deep dive swim lane type does not require writing SPL?

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

Answer: D

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. References: [KPI Lanes]

NEW QUESTION 52
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.

Answer: ABCD

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.

NEW QUESTION 53
Which of the following are characteristics of ITSI service dependencies? (select all that apply)

A. If a primary service has a dependent service KPI and the KPI's importance level is changed, the dependency is broken.
B. It is best practice to use the dependent service's built-in 'ServiceHealthScore' KPI to reflect impact to the primary service.
C. Setting the dependent service KPI importance level will be treated as any other KPI in the primary service's health score.
D. Impactful dependent services should only be configured to one primary service to avoid false negatives in Multi KPI Alerts.

Answer: BC

Explanation:
In the context of Splunk IT Service Intelligence (ITSI), service dependencies allow for the modeling of relationships between services, where the health of one service (dependent) can affect the health of another (primary).
* B.It is best practice to use the dependent service's built-in 'ServiceHealthScore' KPI to reflect impact to the primary service:Utilizing the 'ServiceHealthScore' KPI of a dependent service as part of the primary service's health calculation is a recommended practice. This approach ensures that changes in the health of the dependent service directly influence the primary service's overall health score, providing a more holistic view of service health within the IT environment.
* C.Setting the dependent service KPI importance level will be treated as any other KPI in the primary service's health score:When a dependent service's KPI is incorporated into a primary service, the importance level assigned to this KPI is factored into the primary service's overall health score calculation just like any other KPI. This means that the impact of the dependent service on the primary service can be weighted according to the business significance of the relationship between the services.
The other options are not accurate representations of ITSI service dependencies. Changes in KPI importance levels do not break dependencies, and there is no restriction on configuring impactful dependent services to only one primary service, as dependencies can be complex and multi-layered across various services.

NEW QUESTION 58
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)

Answer: C

Explanation:
Reference: https://splunk.github.io/docker-splunk/ARCHITECTURE.html
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.
References: Ports used by ITSI

NEW QUESTION 60
What is the range for a normal Service Health score category?

A. 20-40
B. 40-60
C. 60-80
D. 80-100

Answer: D

Explanation:
In Splunk IT Service Intelligence (ITSI), the Service Health Score is a metric that provides a quantifiable measure of the overall health and performance of a service. The score ranges from 0 to 100, with higher scores indicating better health. The range for a normal Service Health score category is typically from 80 to 100. Scores within this range suggest that the service is performing well, with no significant issues affecting its health. This categorization helps IT and business stakeholders quickly assess the operational status of their services, enabling them to focus on services that may require attention or intervention due to lower health scores.

NEW QUESTION 63
Which is the least permissive role required to modify default deep dives?

A. itoa_analyst
B. admin
C. power
D. itoa_admin

Answer: D

Explanation:
To modify default deep dives in Splunk IT Service Intelligence (ITSI), the least permissive role typically required is theitoa_adminrole. This role is specifically designed within ITSI to provide administrative capabilities, including the ability to configure and customize various aspects of ITSI, such as services, KPIs, and deep dives. The itoa_adminrole has the necessary permissions to edit and manage default deep dives, enabling users with this role to tailor the deep dives to meet specific operational requirements and preferences. Other roles likeitoa_analyst,admin, orpowermight not have sufficient privileges to modify default deep dives, as these roles are generally more restricted in terms of their ability to make broad changes within ITSI.

NEW QUESTION 66
There are two Smart Mode configuration settings that control how fields affect grouping. Which of these is correct?

A. Text deviation and category deviation.
B. Text similarity and category deviation.
C. Text similarity and category similarity.
D. Text deviation and category similarity.

Answer: C

Explanation:
In the context of Smart Mode configuration within Splunk IT Service Intelligence (ITSI), the two settings that control how fields affect grouping are "Text similarity" and "Category similarity." Smart Mode is a feature used in event grouping that leverages machine learning to automatically group related events. "Text similarity" refers to how closely the textual content of event fields must match for those events to be grouped together, taking into account commonalities in strings or narratives within the event data. "Category similarity," on the other hand, relates to the similarity in the categorical attributes of events, such as event types or source types, which helps in clustering events that are similar in nature or origin. Both of these settings are crucial in determining how events are grouped in ITSI, influencing the granularity and relevance of the event groupings based on textual and categorical similarities.

NEW QUESTION 70
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.

Answer: BCD

Explanation:
Throttling applies to any correlation search alert type, including notable events and actions (RSS feed, email, run script, and ticketing).
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/ConfigCS
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. References: Configure correlation search settings in ITSI

NEW QUESTION 74
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 point
B. This allows the ML pattern to perform it??s magic.
C. A minimum of 24 hours of data is needed for anomaly detection, and a minimum of 4 entities for cohesive analysis.
D. Anomaly detection automatically generates notable events when KPI data diverges fromthe pattern.
E. There are 3 types of anomaly detection supported in ITSI: adhoc, trending, and cohesive.

Answer: BC

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/AD
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. References: [Anomaly Detection]

NEW QUESTION 79
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.

Answer: D

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.

NEW QUESTION 82
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.

Answer: D

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 alsocontrol 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. References: Overview of teams in ITSI

NEW QUESTION 83
What is the main purpose of the service analyzer?

A. Display a list of All Services and Entities.
B. Trigger external alerts based on threshold violations.
C. Allow Analysts to add comments to Alerts.
D. Monitor overall Service and KPI status.

Answer: D

Explanation:
Reference: https://docs.splunk.com/Documentation/MSExchange/4.0.3/Reference/ServiceAnalyzer
The service analyzer is a dashboard that allows you to monitor the overall service and KPI status in ITSI. The service analyzer displays a list of all services and their health scores, which indicate how well each service is performing based on its KPIs. You can also view the status and values of each KPI within a service, as well as drill down into deep dives or glass tables for further analysis. The service analyzer helps you identify issues affecting your services and prioritize them based on their impact and urgency. The main purpose of the service analyzer is:
* D. Monitor overall service and KPI status. This is true because the service analyzer provides a comprehensive view of the health and performance of your services and KPIs in real time.
The other options are not the main purpose of the service analyzer because:
* A. Display a list of all services and entities. This is not true because the service analyzer does not display entities, which are IT components that require management to deliver an IT service. Entities are displayed in other dashboards, such as entity management or entity health overview.
* B. Trigger external alerts based on threshold violations. This is not true because the service analyzer does not trigger alerts, which are notifications sent to external systems or users when certain conditions are met. Alerts are triggered by correlation searches or alert actions configured in ITSI.
* C. Allow analysts to add comments to alerts. This is not true because the service analyzer does not allow analysts to add comments to alerts, which are notifications sent to external systems or users

NEW QUESTION 85
Which of the following is a good use case for creating a custom module?

A. Modules are required to create entity and service import searches.
B. Modules are required to be able to create custom visualizations for deep dives.
C. Making it easy to migrate KPI base searches and related visualizations to other ITSI installations.
D. Creating a service template to make it easy to automatically create new services during service and entity import.

Answer: C

Explanation:
Creating a custom module in Splunk IT Service Intelligence (ITSI) is particularly beneficial for the purpose of migrating KPI base searches and related visualizations to other ITSI installations. Custom modules can encapsulate a set of configurations, searches, and visualizations that are tailored to specific monitoring needs or environments. By packaging these elements into a module, it becomes easier to transfer, deploy, and maintain consistency across different ITSI instances. This modularity supports the reuse of developed components, simplifying the process of scaling and replicating monitoring setups in diverse operational contexts. The ability to migrate these components seamlessly enhances operational efficiency and ensures that best practices and custom configurations can be shared across an organization's ITSI deployments.

NEW QUESTION 88
Which of the following is a problem requiring correction in ITSI?

A. Twoormore entitieswiththe same service ID.
B. Twoormore entitieswiththe same entity ID.
C. Twoormore entitieswiththe same value in a single alias field.
D. Twoormore entitieswiththe same entity key value inanyinfo field.

Answer: C

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.

NEW QUESTION 90
Which of the following are deployment recommendations for ITSI? (Choose all that apply.)

A. Deployments often require an increase of hardware resources above base Splunk requirements.
B. Deployments require a dedicated ITSI search head.
C. Deployments may increase the number of required indexers based on the number of KPI searches.
D. Deployments should use fastest possible disk arrays for indexers.

Answer: ABC

Explanation:
You might need to increase the hardware specifications of your own Enterprise Security
deployment above the minimum hardware requirements depending on your environment. Install Splunk Enterprise Security on a dedicated search head or search head cluster.
The Splunk platform uses indexers to scale horizontally. The number of indexers required in an Enterprise Security deployment varies based on the data volume, data type, retention requirements, search type, and search concurrency.
Reference: https://docs.splunk.com/Documentation/ES/latest/Install/DeploymentPlanning
A, B, and C are correct answers because ITSI deployments often require more hardware resources than base Splunk requirements due to the high volume of data ingestion and processing. ITSI deployments also require a dedicated search head that runs the ITSI app and handles all ITSI-related searches and dashboards. ITSI deployments may also increase the number of required indexers based on the number and frequency of KPI searches, which can generate a large amount of summary data. References: ITSI deployment overview, ITSI deployment planning

NEW QUESTION 92
Which of the following describes entities? (Choose all that apply.)

A. Entities must be IT devices, such as routers and switches, and must be identified by either IP value, host name, or mac address.
B. An abstract (pseudo/logical) entity can be used to split by for a KPI, although no entity rules or filtering can be used to limit data to a specific service.
C. Multiple entities can share the same alias value, but must have different role values.
D. To automatically restrict the KPI to only the entities in a particular service, select ??Filter to Entities in Service??.

Answer: BD

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/KPIfilter
Entities are IT components that require management to deliver an IT service. Each entity has specific attributes and relationships to other IT processes that uniquely identify it. Entities contain alias fields and informational fields that ITSI associates with indexed events. Some statements that describe entities are:
* B. An abstract (pseudo/logical) entity can be used to split by for a KPI, although no entity rules or filtering can be used to limit data to a specific service. An abstract entity is an entity that does not represent a physical host or device, but rather a logical grouping of data sources. For example, you can create an abstract entity for each business unit in your organization and use it to split by for a KPI that measures revenue or customer satisfaction. However, you cannot use entity rules or filtering to limit data to a specific service based on abstract entities, because they do not have alias fields that match indexed events.
* D. To automatically restrict the KPI to only the entities in a particular service, select ??Filter to Entities in Service??. This option allows you to filter the data sources for a KPI by the entities that are assigned to the service. For example, if you have a service for web servers and you want to monitor the CPU load percent for each web server entity, you can select this option to ensure that only the events from those entities are used for the KPI calculation.
References: Overview of entity integrations in ITSI, [Create KPI base searches in ITSI]

NEW QUESTION 94
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.

Answer: A

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.

NEW QUESTION 96
Which of the following is a characteristic of notable event groups?

A. Notable event groups combine independent notable events.
B. Notable event groups are created in the itsi_tracked_alerts index.
C. Notable event groups allow users to adjust threshold settings.
D. All of the above.

Answer: A

Explanation:
In Splunk IT Service Intelligence (ITSI), notable event groups are used to logically group related notable events, which enhances the manageability and analysis

of events:
A.Notable event groups combine independent notable events:This characteristic allows for the aggregation of related events into a single group, making it easier for users to manage and investigate related issues. By grouping events, users can focus on the broader context of an issue rather than getting lost in the details of individual events.
While notable event groups play a critical role in organizing and managing events in ITSI, they do not inherently allow users to adjust threshold settings, which is typically handled at the KPI or service level. Additionally, while notable event groups are utilized within the ITSI framework, the statement that they are created in the 'itsi_tracked_alerts' index might not fully capture the complexity of how event groups are managed and stored within the ITSI architecture.

NEW QUESTION 99
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

SPLK-3002 Practice Exam Features:

* SPLK-3002 Questions and Answers Updated Frequently

* SPLK-3002 Practice Questions Verified by Expert Senior Certified Staff

* SPLK-3002 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SPLK-3002 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year