Exam Questions SPLK-3002

Splunk IT Service Intelligence Certified Admin Exam

https://www.2passeasy.com/dumps/SPLK-3002/

NEW QUESTION 1
When must a service define entity rules?

A. If the intention is for the KPIs in the service to filter to only entities assigned to the service.
B. To enable entity cohesion anomaly detection.
C. If some or all of the KPIs in the service will be split by entity.
D. If the intention is for the KPIs in the service to have different aggregate v
E. entity KPI values.

Answer: A

Explanation:
Provide a value to filter the service to a specific set of entities. These entity rule values are meant to be custom for each service.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/EntityRules
A is the correct answer because a service must define entity rules if the intention is for the KPIs in the service to filter to only entities assigned to the service. Entity rules are filters that match entities to services based on entity aliases or entity metadata. If you enable the Filter to Entities in Service option for a KPI, you need to define entity rules for the service to ensure that the KPI search results only include the relevant entities for the service. Otherwise, the KPI search results might include entities that are not part of the service or exclude entities that are part of the service. References: [Define entities for a service in ITSI], [Configure KPI settings in ITSI]

NEW QUESTION 2
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 processld' or 'by host'.

Answer: A

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.

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
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.

Answer: C

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/DeepDives
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. References: [Create default deep dives for services in ITSI]

NEW QUESTION 6
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.

Answer: D

Explanation:
When an episode warrants investigation, the analyst acknowledges the episode, which moves the status from New to In Progress.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/EpisodeOverview
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.
References: Overview of Episode Review in ITSI, [Episode actions in Episode Review]

NEW QUESTION 7
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.

Answer: ABD

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.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/KPIthresholds 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. References: Configure KPI thresholds in
ITSI, Apply anomaly detection to a KPI in ITSI, Generate events with correlation searches in ITSI

NEW QUESTION 8
Which of the following describes a realistic troubleshooting workflow in ITSI?

A. Correlation Search –> Deep Dive –> Notable Event
B. Service Analyzer –> Notable Event Review –> Deep Dive
C. Service Analyzer –> Aggregation Policy –> Deep Dive
D. Correlation search –> KPI –> Aggregation Policy

Answer: B

Explanation:
A realistic troubleshooting workflow in ITSI is:
? B. Service Analyzer –> Notable Event Review –> Deep Dive
This workflow involves using the Service Analyzer dashboard to monitor the health and performance of your services and KPIs, using the Notable Event Review dashboard to investigate and manage the notable events generated by ITSI, and using the Deep Dive dashboard to analyze the historical trends and anomalies of your KPIs and metrics.
The other workflows are not realistic because they involve components that are not part of the troubleshooting process, such as correlation search, aggregation policy, and KPI.These components are used to create and configure the alerts and episodes that ITSI generates, not to investigate and resolve them. References: [Service Analyzer dashboard in
ITSI], Overview of Episode Review in ITSI, [Overview of deep dives in ITSI]

NEW QUESTION 9

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

Answer: A

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 orattributes. 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.

NEW QUESTION 10
Which index is used to store KPI values?

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

Answer: A

Explanation:
The IT Service Intelligence (ITSI) metrics summary index, itsi_summary_metrics, is a metrics-based summary index that stores KPI data.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/MetricsIndexRef
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. References: Overview of ITSI indexes

NEW QUESTION 10
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.

Answer: A

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.

NEW QUESTION 12
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 14
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 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 19
Which index will contain useful error messages when troubleshooting ITSI issues?

A. _introspection
B. _internal
C. itsi_summary
D. itsi_notable_audit

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/TroubleshootRE The index that will contain useful error messages when troubleshooting ITSI issues is:
* B. _internal. This is true because the _internal index contains logs and metrics generated by Splunk processes, such as splunkd and metrics.log. These logs can help you diagnose problems with your Splunk environment, including ITSI components and features.
The other indexes will not contain useful error messages because:
* A. _introspection. This is not true because the _introspection index contains data about Splunk resource usage, such as CPU, memory, disk space, and so on. These data can help you monitor the performance and health of your Splunk environment, but not the error messages.
* C. itsi_summary. This is not true because the itsi_summary index contains summarized data for your KPIs and services, such as health scores, severity levels, threshold values, and so on. These data can help you analyze the trends and anomalies of your IT services, but not the error messages.
* D. itsi_notable_audit. This is not true because the itsi_notable_audit index contains audit data for your notable events and episodes, such as creation time, owner

NEW QUESTION 23
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.

Answer: D

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 isbeneficial in complex environments where the interplay between different performance metrics needs to be considered to accurately detect and diagnose issues.

NEW QUESTION 28
What is an episode?

A. A workflow task.
B. A deep dive.
C. A notable event group.
D. A notable event.

Answer: C

Explanation:
It's a deduplicated group of notable events occurring as part of a larger sequence, or an incident or period considered in isolation.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/EpisodeOverview
An episode is a deduplicated group of notable events occurring as part of a larger sequence, or an incident or period considered in isolation. An episode helps you reduce alert noise and focus on the most important issues affecting your IT services. An episode is created by an aggregation policy, which is a set of rules that determines how to group notable events based on certain criteria, such as severity, source, title, and so on. You can use episode review to view, manage, and resolve episodes in ITSI. The statement that defines an episode is:
* C. A notable event group. This is true because an episode is composed of one or more notable events that are related by some common factor.
The other options are not definitions of an episode because:
* A. A workflow task. This is not true because a workflow task is an action that you can perform on an episode, such as assigning an owner, changing the status, adding comments, and so on.
* B. A deep dive. This is not true because a deep dive is a dashboard that allows you to analyze the historical trends and anomalies of your KPIs and metrics in ITSI.
* D. A notable event. This is not true because a notable event is an alert generated by ITSI based on certain conditions or correlations, not a group of alerts.
References: [Overview of Episode Review in ITSI], [Overview of aggregation policies in ITSI]

NEW QUESTION 30
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 33
When installing ITSI to support a Distributed Search Architecture, which of the following items apply? (Choose all that apply.)

A. Copy SA-IndexCreation to all indexers.
B. Copy SA-IndexCreation to the etc/apps directory on the index cluster master node.
C. Extract installer package into etc/apps directory of the cluster deployer node.
D. Extract ITSI app package into etc/apps directory of search head.

Answer: A

Explanation:
Copy SA-IndexCreation to $SPLUNK_HOME/etc/apps/ on all individual indexers in your environment.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Install/InstallSHC
A is the correct answer because when installing ITSI to support a distributed search architecture, you need to copy SA-IndexCreation to all indexers. SA-IndexCreation is an app that contains the definitions of the ITSI indexes, such as itsi_summary, itsi_tracked_alerts, itsi_grouped_alerts, etc. You need to copy this app to all indexers to ensure that they can store and search the ITSI data. B is not a correct answer because you do not need to copy SA-IndexCreation to the etc/apps directory on the index cluster master node. The index cluster master node does not store or search data, it only manages the replication and availability of data across the index cluster peers. C is not a correct answer because you do not need to extract the installer package into etc/apps directory of the cluster deployer node. The cluster deployer node is used to distribute apps and configuration updates to the search head cluster members. You need to extract the installer package into etc/shcluster/apps directory of the cluster deployer node instead. D is not a correct answer because you do not need to extract the ITSI app package into etc/apps directory of search head. You need to extract the ITSI app package into etc/shcluster/apps directory of the cluster deployer node and use the deployer to push the app to all search head cluster members. References: [Install Splunk IT ServiceIntelligence on a search head cluster], [Install Splunk IT Service Intelligence on an indexer cluster]

NEW QUESTION 38
How can Service Now incidents be created automatically when a Multi-KPI alert triggers? (select all that apply)

A. By creating a custom etc/apps/SA-ITOA/workflow_rule
B. conf
C. By linking Entities to Service-Now configuration items.
D. By creating a notable event aggregation policy with a SNOW incident action.
E. By editing the associated correlation search and specifying an alert action.

Answer: CD

Explanation:
To automatically create ServiceNow incidents when a Multi-KPI alert triggers in Splunk IT Service Intelligence (ITSI), the following approaches can be used:
* C.By creating a notable event aggregation policy with a ServiceNow (SNOW) incident action:ITSI allows the creation of notable event aggregation policies that can specify actions to be taken when certain conditions are met. One of these actions can be the creation of an incident in ServiceNow, directly linking the alerting mechanism in ITSI with incident management in ServiceNow.
* D.By editing the associated correlation search and specifying an alert action: Correlation searches in ITSI are used to identify patterns or conditions that signify notable events. These searches can be configured to include alert actions, such as creating a ServiceNow incident, whenever the search conditions are met. This direct integration ensures that incidents are automatically generated in ServiceNow, based on the specific criteria defined in the correlation search.
Options A and B are not standard practices for integrating ITSI with ServiceNow for automatic incident creation. The configuration typically involves setting up actionable alert mechanisms within ITSI that are specifically designed to integrate with external systems like ServiceNow.

NEW QUESTION 40
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.

Answer: D

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/Visualizations#collapseDesktop8
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 hocsearches, 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. References: Overview of the glass table editor in ITSI, [Configure service swapping on glass tables]

NEW QUESTION 41
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.

Answer: D

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. References: Create and update entities using a CSV file in ITSI, Delete entities in bulk in ITSI, Delete entities using the | deleteentity command in ITSI

NEW QUESTION 43
Which of the following is an advantage of using adaptive time thresholds?

A. Automatically update thresholds daily to manage dynamic changes to KPI values.
B. Automatically adjust KPI calculation to manage dynamic event data.
C. Automatically adjust aggregation policy grouping to manage escalating severity.
D. Automatically adjust correlation search thresholds to adjust sensitivity over time.

Answer: A

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/TimePolicies
Adaptive thresholds are thresholds calculated by machine learning algorithms that dynamically adapt and change based on the KPI??s observed behavior. Adaptive thresholds are useful for monitoring KPIs that have unpredictable or seasonal patterns that are difficult to capture with static thresholds. For example, you might use adaptive thresholds for a KPI that measures web traffic volume, which can vary depending on factors such as holidays, promotions, events, and so on.
The advantage of using adaptive thresholds is:
* A. Automatically update thresholds daily to manage dynamic changes to KPI values. This is true because adaptive thresholds use historical data from a training window to generate threshold values for each time block in a threshold template. Each night at midnight, ITSI recalculates adaptive threshold values for a KPI by organizing the data from the training window into distinct buckets and then analyzing each bucket separately.This way, the thresholds reflect the most recent changes in the KPI data and account for any anomalies or trends.
The other options are not advantages of using adaptive thresholds because:
* B. Automatically adjust KPI calculation to manage dynamic event data. This is not true because adaptive thresholds do not affect the KPI calculation, which is based on the base search and the aggregation method. Adaptive thresholds only affect the threshold values that are used to determine the KPI severity level.
* C. Automatically adjust aggregation policy grouping to manage escalating severity. This is not true because adaptive thresholds do not affect the aggregation policy, which is a set of rules that determines how to group notable events into episodes. Adaptive thresholds only affect the threshold values that are used to generate notable events based on KPI severity level.
* D. Automatically adjust correlation search thresholds to adjust sensitivity over time. This is not true because adaptive thresholds do not affect the correlation search, which is a search that looks for relationships between data points and generates notable events. Adaptive thresholds only affect the threshold values that are used by KPIs, which can be used as inputs for correlation searches.
References: Create adaptive KPI thresholds in ITSI

NEW QUESTION 47
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 52
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 53
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 58
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 62
To use Adaptive Threshholding, what is the minimum requirement for a set of KPI data?

A. 14 days old.
B. 7 days old.
C. 30 days old.
D. 10 days old.

Answer: B

Explanation:
To utilize Adaptive Thresholding in Splunk IT Service Intelligence (ITSI), the minimum requirement for a set of Key Performance Indicator (KPI) data is that it must be at least 7 days old. Adaptive Thresholding uses historical data to dynamically adjust thresholds based on observed patterns and trends. Having a minimum of 7 days worth of data allows the system to analyze a sufficient amount of information to identify normal ranges and variances in KPI behavior, thereby setting more accurate and contextually relevant thresholds. This requirementensures that the adaptive thresholds are based on a meaningful data set that reflects the typical operational conditions of the monitored services.

NEW QUESTION 66

In which index are active notable events stored?

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

Answer: C

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 theitsi_tracked_alertsindex. 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. Theitsi_tracked_alertsindex enables efficient storage, querying, and management of these events, facilitating the ITSI's event management and review process. The other options, such asitsi_notable_archiveand 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 theitsi_tracked_alertsindex.

NEW QUESTION 67
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 71
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 73
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 75
Which of the following are characteristics of service templates? (select all that apply)

A. Service templates can be modified after services are instantiated from it.
B. Service templates contain KPIs and KPI thresholds.
C. Service templates can contain specific or generic entity rules.
D. Service templates contain domain specific dashboards and deep dives.

Answer: BC

Explanation:
Service templates in Splunk IT Service Intelligence (ITSI) are designed to streamline the creation of services by providing pre-defined configurations:
* B.Service templates contain KPIs and KPI thresholds:This allows for the standardized deployment of services with predefined performance indicators and their associated thresholds, ensuring consistency across similar services.
* C.Service templates can contain specific or generic entity rules:These rules define how entities are associated with services created from the template, allowing for both broad and targeted applicability.
While service templates contain configurations for KPIs, thresholds, and entity rules, the ability to modify templates after services have been instantiated from them is limited. Changes to a template do not retroactively affect services already created from that template. Moreover, service templates do not inherently contain domain-specific dashboards or deep dives; these are created separately within ITSI.

NEW QUESTION 80
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 83
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??

Answer: C

Explanation:
* 1. From the ITSI main menu, click Configuration > Notable Event Aggregation Policies.
* 2. Select a custom policy or the Default Policy.
* 3. Under Smart Mode grouping, enable Smart Mode.
* 4. Click Select fields. A dialog displays the fields found in your notable events from the last 24 hours.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/SmartMode
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. References: Configure smart mode for aggregation policies in ITSI

NEW QUESTION 88
Which of the following is a characteristic of custom deep dives?

A. Allows itoa_analyst roles to add comments.
B. Requires at least 7 days' data to show anomalies.
C. Combines metric, event, KPI, and service health score lanes.
D. Uses drilldown to generate notable events via anomaly detection.

Answer: C

Explanation:
Custom deep dives in Splunk IT Service Intelligence (ITSI) are versatile and
highly customizable dashboards that allow users to analyze various types of data in a unified view. One of the key characteristics of custom deep dives is their ability to combine lanes of different data types, such as metrics, events, Key Performance Indicators (KPIs), and service health scores. This multifaceted approach provides a comprehensive and layered view of the IT environment, enabling analysts and operators to correlate different data types and gain deeper insights into the health and performance of services. By incorporating these diverse data lanes, custom deep dives facilitate a more holistic understanding of the operational landscape, aiding in more effective troubleshooting and decision-making.

NEW QUESTION 93

......