Splunk

Exam Questions SPLK-3002

Splunk IT Service Intelligence Certified Admin Exam

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
What happens when an anomaly is detected?

A. A separate correlation search needs to be created in order to see it.
B. A SNMP trap will be sent.
C. An anomaly alert will appear in core splunk, in index=main.
D. An anomaly alert will appear as a notable event in Episode Review.

Answer: D

Explanation:
When an anomaly is detected in Splunk IT Service Intelligence (ITSI), it typically generates a notable event that can be reviewed and managed in the Episode Review dashboard. The Episode Review is part of ITSI's Event Analytics framework and serves as a centralized location for reviewing, annotating, and managing notable events, including those generated by anomaly detection. This process enables IT operators and analysts to efficiently identify, prioritize, and respond to potential issues highlighted by the
anomaly alerts. The integration of anomaly alerts into the Episode Review dashboard streamlines the workflow for managing and investigating these alerts within the broader context of IT service management and operational intelligence.

NEW QUESTION 3
Which of the following describes default deep dives?

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

Answer: C

Explanation:
In Splunk IT Service Intelligence (ITSI), default deep dives are auto- generated and can be accessed via the Service Analyzer. Deep dives are an essential feature of ITSI that provide an in-depth, granular view into the health and performance of services and their associated KPIs. These default deep dives are automatically created for each service, allowing users to quickly drill down into the detailed operational metrics and performance data of their services. By accessing these deep dives through the Service Analyzer, ITSI users can efficiently investigate issues, understand service dependencies, and make informed decisions to maintain optimal service health. The auto-generated nature of these default deep dives simplifies the monitoring and analysis process, providing immediate insights into service performance without the need for manual setup or configuration.

NEW QUESTION 4
Which of the following is a best practice when configuring maintenance windows?

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

Answer: C

Explanation:
It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/AboutMW
A maintenance window is a period of time when a service or entity is undergoing maintenance operations or does not require active monitoring. It is a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work. This gives the system an opportunity to catch up with the maintenance state and reduces the chances of ITSI generating false positives during maintenance operations. For example, if a server will be shut down for maintenance at 1:00PM and restarted at 5:00PM, the ideal maintenance window is 12:30PM to 5:30PM. The 15- to 30-minute time buffer is a rough estimate based on 15 minutes being the time period over which most KPIs are configured to search data and identify alert triggers. References: Overview of maintenance windows in ITSI

NEW QUESTION 5

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 6
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 7
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 8
After a notable event has been closed, how long will the meta data for that event remain in the KV Store by default?

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

Answer: A

Explanation:
By default, notable event metadata is archived after six months to keep the KV store from growing too large.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/TrimNECollections

NEW QUESTION 9
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 10
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 10
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 13
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 18
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 21
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 22
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%20considers%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 27
Which of the following accurately describes base searches used for KPIs in a service?

A. Base searches can be used for multiple services.
B. A base search can only be used by its service and all dependent services.
C. All the metrics in a base search are used by one service.
D. All the KPIs in a service use the same base search.

Answer: A

Explanation:
KPI base searches let you share a search definition across multiple KPIs in IT Service Intelligence (ITSI). Create base searches to consolidate multiple similar KPIs, reduce search load, and improve search performance.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/BaseSearch
A base search is a search definition that can be shared across multiple KPIs that use the same data source. Base searches can improve search performance and reduce search load by consolidating multiple similar KPIs. The statement that accurately describes base searches used for KPIs in a service is:
A. Base searches can be used for multiple services. This means that you can create a base search for a service and use it for other services that have similar data sources and KPIs. For example, if you have multiple services that monitor web server performance, you can create a base search that queries the web server logs and use it for all the services that need to calculate KPIs based on those logs.

NEW QUESTION 29
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 31
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 34
Which of the following items describe ITSI teams? (select all that apply)

A. Teams should have itoa admin roles added with read-only permissions for services and entities.
B. Services should be assigned to the 'global' team if all users need access to it.
C. By default, all services are owned by the built-in 'global' team and administered by the 'itoa_admin' role.
D. A new team admin role should be created for each tea
E. The new role should inherit the 'itoa_team_admin' role.

Answer: BCD

Explanation:
In Splunk IT Service Intelligence (ITSI), teams are used to organize services, KPIs, and other objects within ITSI to facilitate access control and management:
B.Services should be assigned to the 'global' team if all users need access to it:The 'global' team in ITSI is a built-in concept that denotes universal accessibility. Assigning services to the 'global' team makes them accessible to all ITSI users, irrespective of their specific team memberships. This is useful for services that are relevant across the entire organization.
* C.By default, all services are owned by the built-in 'global' team and administered by the 'itoa_admin' role:This default setting ensures that upon creation, services are accessible to administrators and can be further re-assigned or refined for access by specific teams as needed.
* D.A new team admin role should be created for each team. The new role should inherit the 'itoa_team_admin' role:This best practice allows for granular access control and management within teams. Each team can have its own administrators with the appropriate level of access and permissions tailored to the needs of that team, derived from the capabilities of the 'itoa_team_admin' role.
The concept of adding 'itoa admin roles' with read-only permissions contradicts the typical use case for administrative roles, which usually require more than read-only access to manage services and entities effectively.

NEW QUESTION 38
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 43
What is the default importance value for dependent services?? health scores?

A. 11
B. 1
C. Unassigned
D. 10

Answer: D

Explanation:
By default, impacting service health scores have an importance value of 11.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/Dependencies
A service template is a predefined set of KPIs and entity rules that you can apply to a service or a group of services. A service template helps you standardize the configuration and monitoring of similar services across your IT environment. A service template can also include dependent services, which are services that are required for another service to function properly. For example, a web server service might depend on a database service and a network service. The default importance value for dependent services?? health scores is:
* D. 10. This is true because the importance value indicates how much a dependent service contributes to the health score of the parent service. The default value is 10, which means that the dependent service has the highest impact on the parent service??s healthscore. You can change the importance value of a dependent service in the service template settings.
The other options are not correct because:
* A. 11. This is not true because 11 is an invalid value for importance. The valid range is from 1 (lowest) to 10 (highest).
* B. 1. This is not true because 1 is the lowest value for importance, not the default value. A value of 1 means that the dependent service has the lowest impact on the parent service??s health score.
* C. Unassigned. This is not true because every dependent service has an assigned importance value, which defaults to 10.
References: Create and manage service templates in ITSI, Set KPI importance values in ITSI

NEW QUESTION 47
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 50
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 54
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 56
In maintenance mode, which features of KPIs still function?

A. KPI searches will execute but will be buffered until the maintenance window is over.
B. KPI searches still run during maintenance mode, but results go to itsi_maintenance_summary index.
C. New KPIs can be created, but existing KPIs are locked.
D. KPI calculations and threshold settings can be modified.

Answer: A

Explanation:
It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work. This gives the system an opportunity to catch up with the maintenance state and reduces the chances of ITSI generating false positives during maintenance operations.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/AboutMW
A is the correct answer because KPI searches still run during maintenance mode, but the results are buffered until the maintenance window is over. This means that no alerts are triggered during maintenance mode, but once it ends, the buffered results are processed and alerts are generated if necessary. You cannot create new KPIs or modify existing KPIs during maintenance mode. References: [Overview of maintenance windows in ITSI]

NEW QUESTION 60
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 65
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 67
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 72
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 75
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 76
When a KPI's aggregate value is calculated, which function is called?

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

Answer: B

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

NEW QUESTION 81
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 83
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 85
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 90
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 95
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 99

Which of the following is a characteristic of base searches?

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

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/BaseSearch
A base search is a search definition that can be shared across multiple KPIs that use the same data source. Base searches can improve search performance and reduce search load by consolidating multiple similar KPIs. One of the characteristics of base searches is that it is possible to filter to entities assigned to the service for calculating the metrics for the service??s KPIs. This means that you can use entity filtering rules to specify which entities are relevant for each KPI based on the base search results. References: Create KPI base searches in ITSI, [Filter entities for KPIs based on base searches]

NEW QUESTION 103
When changing a service template, which of the following will be added to linked services by default?

A. Thresholds.
B. Entity Rules.
C. New KPIs.
D. Health score.

Answer: C

Explanation:
? C. New KPIs. This is true because when you add new KPIs to a service template, they will be automatically added to all the services that are linked to that template. This helps you keep your services consistent and up-to-date with the latest KPI definitions.
The other options will not be added to linked services by default because:
? A. Thresholds. This is not true because when you change thresholds in a service template, they will not affect the existing thresholds in the linked services. You need to manually apply the threshold changes to each linked service if you want them to inherit the new thresholds from the template.
? B. Entity rules. This is not true because when you change entity rules in a service
template, they will not affect the existing entity rules in the linked services. You need to manually apply the entity rule changes to each linked service if you want them to inherit the new entity rules from the template.
? D. Health score. This is not true because when you change health score settings
in a service template, they will not affect the existing health score settings in the linked services. You need to manually apply the health score changes to each linked service if you want them to inherit the new health score settings from the template.
References: Create and manage service templates in ITSI, [Apply service template changes to linked services in ITSI]

NEW QUESTION 104
Which of the following is a valid type of Multi-KPI Alert?

A. Score over composite.
B. Value over time.
C. Status over time.
D. Rise over run.

Answer: B

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/MKA
B is the correct answer because value over time is a valid type of Multi-KPI Alert in ITSI. A Multi-KPI Alert is a type of alert that triggers when multiple KPIs from one or more services meet certain conditions within a specified time range. Value over time is a condition that compares the current value of a KPI to its previous values over a specified time range. For example, you can create a Multi-KPI Alert that triggers when the CPU usage and memory usage of a service are both higher than their average values in the last 24 hours. References: [Create Multi-KPI alerts in ITSI], [Multi-KPI alert conditions in ITSI]

NEW QUESTION 109
......