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

A. Gray
B. Purple
C. Gear Icon
D. Blue

Answer: A

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

NEW QUESTION 3
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 4
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 5
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 6
What is the minimum number of entities a KPI must be split by in order to use Entity Cohesion anomaly detection?

A. 3
B. 4
C. 5
D. 2

Answer: D

Explanation:
For Entity Cohesion anomaly detection in Splunk IT Service Intelligence (ITSI), the minimum number of entities a KPI must be split by is 2. Entity Cohesion as a method of anomaly detection focuses on identifying anomalies based on the deviation of an entity's behavior in comparison to other entities within the same group or cohort. By requiring a minimum of only two entities, ITSI allows for the comparison of entities to detect significant deviations in one entity's performance or behavior, which could indicate potential issues. This method leverages the idea that entities performing similar functions or within the same service should exhibit similar patterns of behavior, and significant deviations could be indicative of anomalies. The low minimum requirement of two entities ensures that this powerful anomaly detection feature can be utilized even in smaller environments.

NEW QUESTION 7
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 8
Which step is required to install ITSI on a single Search Head?

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

Answer: C

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

NEW QUESTION 9
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 10
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 10
What are valid considerations when designing an ITSI Service? (Choose all that apply.)

A. Service access control requirements for ITSI Team Access should be considered, and appropriate teams provisioned prior to creating the ITSI Service.
B. Entities, entity meta-data, and entity rules should be planned carefully to support the service design and configuration.
C. Services, entities, and saved searches are stored in the ITSI app, while events created by KPI execution are stored in the itsi_summary index.
D. Backfill of a KPI should always be selected so historical data points can be used immediately and alerts based on that data can occur.

Answer: ABC

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/ImplementPerms
A, B, and C are correct answers because service access control requirements for ITSI Team Access should be considered before creating the ITSI Service, as different teams may have different permissions and views of the service data. Entities, entity meta-data, and entity rules should also be planned carefully to support the service design and configuration, as they determine how ITSI maps data sources to services and KPIs. Services, entities, and saved searches are stored in the ITSI app, while events created by KPI execution are stored in the itsi_summary index for faster retrieval and analysis. References: ITSI service design best practices, Overview of ITSI indexes

NEW QUESTION 14
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 17
What should be considered when onboarding data into a Splunk index, assuming that ITSI will need to use this data?

A. Use | stats functions in custom fields to prepare the data for KPI calculations.
B. Check if the data could leverage pre-built KPIs from modules, then use the correct TA to onboard the data.
C. Make sure that all fields conform to CIM, then use the corresponding module to import related services.
D. Plan to build as many data models as possible for ITSI to leverage

Answer: B

Explanation:
Reference: https://newoutlook.it/download/book/splunk/advanced-splunk.pdf
When onboarding data into a Splunk index, assuming that ITSI will need to use this data, you should consider the following:
* B. Check if the data could leverage pre-built KPIs from modules, then use the correct TA to onboard the data. This is true because modules are pre-packaged sets of services, KPIs, and dashboards that are designed for specific types of data sources, such as operating systems, databases, web servers, and so on. Modules help you quickly set up and monitor your IT services using best practices and industry standards. To use modules, you need to install and configure the correct technical add-ons (TAs) that extract and normalize the data fields required by the modules.
The other options are not things you should consider because:
* A. Use | stats functions in custom fields to prepare the data for KPI calculations. This is not true because using | stats functions in custom fields can cause performance issues and inaccurate results when calculating KPIs. You should use | stats functions only in base searches or ad hoc searches, not in custom fields.
* C. Make sure that all fields conform to CIM, then use the corresponding module to import related services. This is not true because not all modules require CIM-compliant data sources. Some modules have their own data models and field extractions that are specific to their data sources. You should check the documentation of each module to see what data requirements and dependencies they have.
* D. Plan to build as many data models as possible for ITSI to leverage. This is not true because building too many data models can cause performance issues and resource consumption in your Splunk environment. You should only build data models that are necessary and relevant for your ITSI use cases.
References: Overview of modules in ITSI, [Install technical add-ons for ITSI modules]

NEW QUESTION 18
Which capabilities are enabled through ??teams???

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

Answer: D

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

NEW QUESTION 20
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 21
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 24
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 27
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 32
Which of the following actions can be performed with a deep dive?

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

Answer: A

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

NEW QUESTION 35
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 38
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

SPLK-3002 Practice Exam Features:

* SPLK-3002 Questions and Answers Updated Frequently

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

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

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