Exam : **SPLK-3002**

Title : Splunk IT Service Intelligence Certified Admin

Vendor : Splunk

Version : DEMO

NO.1 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 and analyzing 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.

NO.2 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
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

NO.3 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.

NO.4 Which of the following is a good use case regarding defining entities for a service?
A. Automatically associate entities to services using multiple entity aliases.
B. All of the entities have the same identifying field name.
C. Being able to split a CPU usage KPI by host name.
D. KPI total values are aggregated from multiple different category values in the source events.
Answer: A
Explanation:
Define entities before creating services. When you configure a service, you can specify entity matching rules based on entity aliases that automatically add the entities to your service.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Entity/About A is the correct answer because defining entities for a service allows you to automatically associate entities to services using multiple entity aliases. Entity aliases are alternative names or identifiers for an entity, such as host name, IP address, MAC address, or DNS name. ITSI matches entity aliases to fields in your data sources and assigns entities to services accordingly. This way, you can avoid manually adding entities to each service and ensure that your services reflect the latest changes in your environment.
References: Define entities for a service in ITSI

NO.5 Which of the following are characteristics of ITSI service dependencies? (select all that apply)
A. If a primary service has a dependent service KPI and the KPI's importance level is changed, the dependency is broken.
B. It is best practice to use the dependent service's built-in 'ServiceHealthScore' KPI to reflect impact to the primary service.
C. Setting the dependent service KPI importance level will be treated as any other KPI in the primary service's health score.
D. Impactful dependent services should only be configured to one primary service to avoid false negatives in Multi KPI Alerts.
Answer: B C
Explanation:

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

NO.6 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: A B C D

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.

NO.7 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]

NO.8 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

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]

NO.9 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 vs. 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]

NO.10 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.