Exam Questions SPLK-3002

Splunk IT Service Intelligence Certified Admin Exam

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

NEW QUESTION 1
There are two departments using ITSI. Finance and Sales. Analysts in each department should not be allowed to see each other??s services. What are the role configuration steps required to accomplish this?

A. itoa_finance_admin, inherited from itoa_admin; itoa_sales_admin, inherited from itoa_team_admin; itoa_finance_analyst, inherited from itoa_analyst; itoa_sales_analyst, inherited from itoa_analyst.
B. itoa_finance_admin, inherited from itoa_admin; itoa_sales_admin, inherited fromitoa_team_admin; itoa_finance_analyst, inherited from itoa_team_analyst; itoa_sales_analyst, inherited from itoa_team_analyst.
C. itoa_finance_admin, inherited from itoa_admin; itoa_sales_admin, inherited from itoa_team_admin; itoa_finance_analyst, inherited from itoa_analyst; itoa_sales_analyst, inherited from itoa_team_analyst.
D. itoa_finance_admin, inherited from itoa_team_admin; itoa_sales_admin, inherited from itoa_team_admin; itoa_finance_analyst, inherited from itoa_analyst; itoa_sales_analyst, inherited from itoa_analyst.

Answer: C

Explanation:
C is the correct answer because teams are a feature of ITSI that allow you to restrict access to service content in UI views based on user roles. To create separate teams for finance and sales analysts, you need to create custom roles that inherit from the itoa_analyst role, which has read-only access to ITSI content. For example, you can create itoa_finance_analyst and itoa_sales_analyst roles that inherit from itoa_analyst. Then, you need to create custom teams that include these roles and assign them to the relevant services. For example, you can create a finance team that includes the itoa_finance_analyst role and assign it to the finance services. Similarly, you can create a sales team that includes the itoa_sales_analyst role and assign it to the sales services. This way, analysts in each department can only see their own services and not each other??s. References: Create teams in ITSI, Assign teams to services in ITSI

NEW QUESTION 2
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 3
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 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
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 7
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 8
Which of the following statements is accurate when using multiple policies?

A. New policies are applied after the default policy.
B. Policy processing is applied in a defined order.
C. An event can be processed by only a single policy.
D. New policies are applied before the default policy.

Answer: B

Explanation:
In Splunk IT Service Intelligence (ITSI), when using multiple event management policies, it is important to understand that policy processing is applied in a defined order. This order is crucial because it determines how events are processed and aggregated, and which rules are applied to events first. The order of policies can be customized, allowing administrators to prioritize certain policies over others based on the specific needs and operational logic of their IT environment. This feature provides flexibility in event management, enabling more precise control over event processing and ensuring that the most critical events are handled according to the desiredprecedence. This structured approach to policy processing helps in maintaining the efficiency and effectiveness of event management within ITSI.

NEW QUESTION 9
Which of the following can generate notable events?

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

Answer: C

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

facilitating efficient incident response and management within the IT environment.

NEW QUESTION 10
Which of the following is a best practice for identifying the most effective services with which to start an iterative ITSI deployment?

A. Only include KPIs if they will be used in multiple services.
B. Analyze the business to determine the most critical services.
C. Focus on low-level services.
D. Define a large number of key services early.

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/MKA
A best practice for identifying the most effective services with which to start an iterative ITSI deployment is to analyze the business to determine the most critical services that have the most impact on revenue, customer satisfaction, or other key performance indicators. You can use the Service Analyzer to prioritize and monitor these services. References: Service Analyzer

NEW QUESTION 10
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 11
For which ITSI function is it a best practice to use a 15-30 minute time buffer?

A. Correlation searches.
B. Adaptive thresholding.
C. Maintenance windows
D. Anomaly detection.

Answer: B

Explanation:
B is the correct answer because adaptive thresholding is a feature of ITSI that allows you to dynamically adjust KPI thresholds based on historical patterns and trends. Adaptive thresholding requires a time buffer of at least 15 minutes to calculate the thresholds based on the previous data points. The time buffer ensures that there is enough data to perform the calculations and avoid false positives or negatives. References: Configure adaptive thresholding for a KPI in ITSI

NEW QUESTION 15
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 19
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 24
When troubleshooting KPI search performance, which search names in job activity identify base searches?

A. Indicator - XXXX - Base Search
B. Indicator - Shared - xxxx - ITSI Search
C. Indicator - Base - xxxx - ITSI Search
D. Indicator - Base - XXXX - Shared Search

Answer: B

Explanation:
In the context of troubleshooting KPI search performance in Splunk IT Service Intelligence (ITSI), the search names in the job activity that identify base searches typically follow the pattern "Indicator - Shared - xxxx - ITSI Search." These base searches are fundamental components of the KPI calculation process, aggregating and preparing data for further analysis by KPIs. Identifying these base searches in the job activity is crucial for diagnosing performance issues, as these searches can be resource-intensive and impact overall system performance. Understanding the naming convention helps administrators and analysts quickly pinpoint the base searches related to specific KPIs, facilitating more effective troubleshooting and optimization of search performance within the ITSI environment.

NEW QUESTION 29
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 33
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 36
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 41
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 46
Which of the following is the best use case for configuring a Multi-KPI Alert?

A. Comparing content between two notable events.
B. Using machine learning to evaluate when data falls outside of an expected pattern.
C. Comparing anomaly detection between two KPIs.
D. Raising an alert when one or more KPIs indicate an outage is occurring.

Answer: D

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/MKA
A multi-KPI alert is a type of correlation search that is based on defined trigger conditions for two or more KPIs. When trigger conditions occur simultaneously for each KPI, the search generates a notable event. For example, you might create a multi-KPI alert based on twocommon KPIs: CPU load percent and web requests. A sudden simultaneous spike in both CPU load percent and web request KPIs might indicate a DDOS (Distributed Denial of Service) attack. Multi-KPI alerts can bring such trending behaviors to your attention early, so that you can take action to minimize any impact on performance. Multi-KPI alerts are useful for correlating the status of multiple KPIs across multiple services. They help you identify causal relationships, investigate root cause, and provide insights into behaviors across your infrastructure. The best use case for configuring a multi-KPI alert is to raise an alert when one or more KPIs indicate an outage is occurring, such as when the service health score drops below a certain threshold or when multiple KPIs have critical severity
levels. References: Create multi-KPI alerts in ITSI

NEW QUESTION 50
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 53
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 56
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 58
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 60
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 65
Which of the following are characteristics of ITSI service dependencies? (select all that apply)

A. If a primary service has a dependent service KPI and the KPI's importance level is changed, the dependency is broken.
B. It is best practice to use the dependent service's built-in 'ServiceHealthScore' KPI to reflect impact to the primary service.
C. Setting the dependent service KPI importance level will be treated as any other KPI in the primary service's health score.

D. Impactful dependent services should only be configured to one primary service to avoid false negatives in Multi KPI Alerts.

Answer: BC

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

NEW QUESTION 67
Which of the following are the default ports that must be configured on Splunk to use ITSI?

A. SplunkWeb (8405), SplunkD (8519), and HTTP Collector (8628)
B. SplunkWeb (8089), SplunkD (8088), and HTTP Collector (8000)
C. SplunkWeb (8000), SplunkD (8089), and HTTP Collector (8088)
D. SplunkWeb (8088), SplunkD (8089), and HTTP Collector (8000)

Answer: C

Explanation:
Reference: https://splunk.github.io/docker-splunk/ARCHITECTURE.html
C is the correct answer because ITSI uses the default ports of Splunk Enterprise for its communication and data collection. SplunkWeb uses port 8000, SplunkD uses port 8089, and HTTP Event Collector uses port 8088. These ports can be changed if needed, but they must match the configuration of Splunk Enterprise.
References: Ports used by ITSI

NEW QUESTION 71
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 74
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 79
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 80
Which of the following statements describe default glass tables in ITSI?

A. The Service Health Score default glass table.
B. There is one default glass table per service.
C. There is one service template default glass table.
D. There are no default glass tables.

Answer: D

Explanation:
In Splunk IT Service Intelligence (ITSI), glass tables are fully customizable dashboards that provide a visual representation of an organization's IT environment, along with the health and status of services and KPIs. Unlike some pre-configured views or dashboards that might come with default setups in various platforms, ITSI does not provide default glass tables out of the box. Instead, users are encouraged to create their own glass tables tailored to their specific monitoring needs and operational views. This approach ensures that each organization can design glass tables that best represent their unique infrastructure, applications, and service landscapes, providing a more personalized and relevant operational overview.

NEW QUESTION 81
Which capabilities are enabled through ??teams???

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

Answer: D

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

NEW QUESTION 83
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
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 86
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 91
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 96
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 100
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 101
When working with a notable event group in the Notable Events Review dashboard, which of the following can be set at the individual or group level?

A. Service, status, owner.
B. Severity, status, owner.

C. Severity, comments, service.
D. Severity, status, service.

Answer: B

Explanation:
In the Notable Events Review dashboard within Splunk IT Service Intelligence (ITSI), when working with a notable event group, users can set or adjust certain attributes at the individual event level or at the group level. These attributes include:
? Severity:The importance or impact level of the notable event or group, which can be adjusted to reflect the current assessment of the situation.
? Status:The current state of the notable event or group, such as "New," "In Progress," or "Resolved," indicating the progress in addressing the event or group.
? Owner:The user or team responsible for managing and resolving the notable event or group.
These settings allow for effective management and tracking of notable events, ensuring that they are appropriately prioritized, acted upon, and resolved by the responsible parties.

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 107
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 112
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SPLK-3002 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SPLK-3002 Product From:

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

Money Back Guarantee

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