Splunk

Exam Questions SPLK-3002

Splunk IT Service Intelligence Certified Admin Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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
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 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
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 5
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 6
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 7
Which ITSI functions generate notable events? (Choose all that apply.)

A. KPI threshold breaches.
B. KPI anomaly detection.
C. Multi-KPI alert.
D. Correlation search.

Answer: ABD

Explanation:
After you configure KPI thresholds, you can set up alerts to notify you when aggregate KPI severities change. ITSI generates notable events in Episode Review based on the alerting rules you configure.
Anomaly detection generates notable events when a KPI IT Service Intelligence (ITSI) deviates from an expected pattern.
Notable events are typically generated by a correlation search.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/KPIthresholds https://docs.splunk.com/Documentation/ITSI/4.10.1/SI/AboutSI
A, B, and D are correct answers because ITSI can generate notable events when a KPI breaches a threshold, when a KPI detects an anomaly, or when a correlation search matches a defined pattern. These are the main ways that ITSI can alert you to potential issues or incidents in your IT environment. References: Configure KPI thresholds in
ITSI, Apply anomaly detection to a KPI in ITSI, Generate events with correlation searches in ITSI

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

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

Answer: B

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

NEW QUESTION 9
Which of the following 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 10
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 10
Which of the following best describes an ITSI Glass Table?

A. A view which displays a system topology overlaid with KPI metrics.
B. A view which describes a topology.
C. A dashboard which displays a system topology.
D. A view showing KPI values in a variety of visual styles.

Answer: A

Explanation:
An ITSI Glass Table provides a customizable, high-level view that can display a system's topology overlaid with real-time Key Performance Indicator (KPI) metrics and service health scores. This visualization tool allows users to create a visual representation of their IT infrastructure, applications, and services, integrating live data to monitor the health and performance of each component in context. The ability to overlay KPI metrics on the system topology enables IT and business stakeholders to quickly understand the operational status and health of various elements within their environment, facilitating more informed decision-making and rapid response to issues.

NEW QUESTION 11
Which of the following is a recommended best practice for service and glass table design?

A. Plan and implement services first, then build detailed glass tables.
B. Always use the standard icons for glass table widgets to improve portability.
C. Start with base searches, then services, and then glass tables.
D. Design glass tables first to discover which KPIs are important.

Answer: A

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/GTOverview
A is the correct answer because it is recommended to plan and implement services first, then build detailed glass tables that reflect the service hierarchy and dependencies. This way, you can ensure that your glass tables provide accurate and meaningful service-level insights. Building glass tables first might lead to unnecessary or irrelevant KPIs that do not align with your service goals. References: Splunk IT Service Intelligence Service Design Best Practices

NEW QUESTION 14
What effects does the KPI importance weight of 11 have on the overall health score of a service?

A. At least 10% of the KPIs will go critical.
B. Importance weight is unused for health scoring.
C. The service will go critical.
D. It is a minimum health indicator KPI.

Answer: B

Explanation:
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/KPIImportance#:~:text=ITSI%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 15
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 18
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 22
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 24
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 27
When installing ITSI to support a Distributed Search Architecture, which of the following items apply? (Choose all that apply.)

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

Answer: A

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

NEW QUESTION 32
Which glass table feature can be used to toggle displaying KPI values from more than one service on a single widget?

A. Service templates.
B. Service dependencies.
C. Ad-hoc search.
D. Service swapping.

Answer: D

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/Visualizations#collapseDesktop8
A glass table is a visualization tool that allows you to monitor the interrelationships and dependencies across your IT and business services. You can add metrics like KPIs, ad hocsearches, and service health scores that update in real time against a background that you design. One of the features of glass tables is service swapping, which enables you to toggle displaying KPI values from more than one service on a single widget. You can use service swapping to compare metrics across different services without creating multiple glass tables or widgets. References: Overview of the glass table editor in ITSI, [Configure service swapping on glass tables]

NEW QUESTION 33
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 36
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 41
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 43
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

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

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

Answer: B

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

NEW QUESTION 48
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 52
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 55
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 58
Which of the following is a problem requiring correction in ITSI?

A. Twoormore entitieswiththe same service ID.
B. Twoormore entitieswiththe same entity ID.
C. Twoormore entitieswiththe same value in a single alias field.
D. Twoormore entitieswiththe same entity key value inanyinfo field.

Answer: C

Explanation:
In Splunk IT Service Intelligence (ITSI), entities represent infrastructure components, applications, or other elements that are monitored. Each entity is uniquely identified by its entity ID, and entities can be associated with one or more services through the concept of aliases. A problem arises when two or more entities have the same value in a single alias field because aliases are used to match events to entities in ITSI. If multiple entities share the same alias value, ITSI might incorrectly associate data with the wrong entity, leading to inaccurate monitoring and analytics. This scenario requires correction to ensure that each alias uniquely identifies a single entity, thereby maintaining the integrity of the monitoring and analysis process within ITSI. The uniqueness of service IDs, entity IDs, and entity key values in info fields is also important but does not typically present the same level of issue as duplicate values in an alias field.

NEW QUESTION 63
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 67
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 71
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 72
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 77
Which of the following is a characteristic of custom deep dives?

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

Answer: C

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

NEW QUESTION 81
.....

Relate Links

100% Pass Your SPLK-3002 Exam with Exambible Prep Materials

https://www.exambible.com/SPLK-3002-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/