PassQueen
a.k.a dumpsqueen.com

Splunk IT Service Intelligence Certified Admin Exam

Splunk SPLK-3002

Version Demo

Total Demo Questions: 11

Total Premium Questions: 118

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, ITSI Concepts	7
Topic 2, Service and KPI Management	49
Topic 3, Event Analytics	31
Topic 4, Glass Tables	7
Topic 5, ITSI Administration	24
Total	118

QUESTION NO: 1

Which of the following describe adaptive thresholds in ITSI? (Choose all that apply.)

A. They use historical KPI behavior to calculate threshold values.

B. They can reduce manual threshold maintenance for dynamic KPI behavior.

C. They are periodically updated to reflect newer KPI behavior.

D. They are configured only at the KPI base-search level.

ANSWER: A B C

Explanation:

Adaptive thresholds use historical KPI behavior to calculate threshold values that better reflect normal operating patterns. They are useful for KPIs whose expected values change over time and for which static threshold values would require frequent manual adjustment. ITSI retrains and updates adaptive thresholds regularly so newer KPI behavior can be incorporated into threshold calculations.

Adaptive thresholds are configured for KPIs, not at the base-search level. They also require historical KPI data because the threshold model must learn expected behavior from prior observations. See adaptive thresholds in ITSI.

QUESTION NO: 2

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: D

Explanation:

It is a minimum health indicator KPI. In Splunk IT Service Intelligence, KPI importance normally determines the relative influence that a KPI's health has when ITSI calculates the aggregate health score for its service. However, an importance value of 11 or higher has a special meaning: ITSI treats that KPI as a minimum health indicator. This establishes the KPI as a required health condition for the service rather than merely another weighted contribution to the aggregate score.

Consequently, the health state of a minimum health indicator constrains the overall service health. If the KPI is in a non-normal state, the service cannot present a healthier state than that KPI permits, even where the combined weighted condition of other KPIs might otherwise yield a better aggregate result. This is useful for KPIs representing essential dependencies or conditions that must remain healthy for the service to be considered healthy. The value 11 is therefore a threshold that changes how the KPI participates in service-health calculation, not a percentage, count, or automatic declaration that the entire service is critical.

See Splunk's documentation on KPI importance and service health and the current ITSI KPI importance documentation.

QUESTION NO: 3

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_rules. conf

B. By linking Entities to Service-Now configuration items.

C. By creating a notable event aggregation policy with a SNOW incident action.

D. By editing the associated correlation search and specifying an alert action.

ANSWER: C D

Explanation:

Creating a notable event aggregation policy with a SNOW incident action is a supported way to automate incident creation from ITSI. Multi-KPI alerts produce notable events, which ITSI can group into episodes according to aggregation-policy criteria. An aggregation policy can then apply an incident-management action to the resulting episode, allowing the relevant alert context to be sent to the configured ServiceNow instance. This approach is particularly useful when several related KPI conditions should result in one consolidated operational incident rather than separate tickets.

Editing the associated correlation search and specifying an alert action is also valid. A Multi-KPI alert is based on a correlation search, and alert actions define what Splunk does when that search generates its results. With the ServiceNow integration configured, the correlation search can invoke the ServiceNow incident alert action automatically when its trigger conditions are met. This enables incident creation directly at alert generation time and can include useful event fields in the ServiceNow record.

Both methods depend on having the applicable ServiceNow integration and credentials configured in Splunk, plus the required permissions for the integration account. Splunk's ITSI documentation describes configuring aggregation policies and notable-event actions, while the ServiceNow integration documentation describes using its alert action for incident creation. See Splunk ITSI aggregation policies and Splunk Add-on for ServiceNow alert-action configuration.

QUESTION NO: 4

ITSI Saved Search Scheduling is configured to use realtime_schedule = 0. Which statement is accurate about this configuration?

A. If this value is set to 0, the scheduler bases its determination of the next scheduled search execution time on the current time.

B. If this value is set to 0, the scheduler bases its determination of the next scheduled search on the last search execution time.

C. If this value is set to 0, the scheduler may skip scheduled execution periods.

D. If this value is set to 0, the scheduler might skip some execution periods to make sure that the scheduler is executing the searches running over the most recent time range.

ANSWER: B

Explanation:

The accurate statement is: "If this value is set to 0, the scheduler bases its determination of the next scheduled search on the last search execution time." In Splunk saved-search scheduling, `realtime_schedule` controls how the scheduler calculates the next execution time for a scheduled search. A value of `0` enables continuous scheduling, meaning that each subsequent run is calculated from the time the search actually last ran. This behavior maintains the intended interval between completed executions, rather than anchoring the next run to the wall-clock time.

This setting is relevant when searches are delayed because of scheduler workload, resource contention, or other execution delays. By using the prior execution as the reference point, continuous scheduling makes the schedule progress from actual search activity. Administrators can configure this behavior in the saved-search stanza in `savedsearches.conf`, including for searches that support ITSI services, KPI calculations, and related operational workloads. Splunk documents `realtime_schedule` as the setting that selects this continuous scheduling behavior when set to `0`. See the savedsearches.conf specification and Splunk's scheduled alert configuration documentation.

QUESTION NO: 5

Which of the following statements about ITSI team access are accurate? (Choose all that apply.)

A. Teams can restrict service content displayed in ITSI views.

B. Services can be associated with one or more teams.

C. Users need underlying Splunk permissions to search supporting data.

D. Team membership automatically grants access to every Splunk index.

ANSWER: A B C

Explanation:

Teams provide service-level access control in ITSI user interface views. Administrators can associate services and service templates with teams, then assign users to those teams through the applicable ITSI roles. The `itoa_team_analyst` role is intended for analysts who need access to content associated with their assigned teams.

Team membership does not grant Splunk index permissions. Users still require the underlying Splunk platform permissions necessary to search the data that supports their ITSI workflows. See Configure teams in ITSI and Configure users and roles in ITSI.

QUESTION NO: 6

In Episode Review, what is the result of clicking an episode's Acknowledge button?

A. Assign the current user as owner.

B. Change status from New to Acknowledged.

C. Change status from New to In Progress and assign the current user as owner.

D. Change status from New to Acknowledged and assign the current user as owner.

ANSWER: C

Explanation:

Clicking *Change status from New to In Progress and assign the current user as owner.* records that an analyst has taken responsibility for investigating the episode. In IT Service Intelligence Episode Review, a newly created episode begins in the *New* state. Acknowledging it transitions the episode to *In Progress*, making its active-investigation status visible to other users. The acknowledge action also assigns the episode to the user who clicked the button, establishing a clear owner for follow-up, triage, and resolution work.

This combined status-and-ownership update supports operational workflows by preventing uncertainty over whether an episode is being handled and by identifying the analyst responsible for it. The owner can subsequently manage the episode's lifecycle, including reviewing contributing notable events, adding comments, assigning it to another user if appropriate, and ultimately resolving or closing it according to the organization's process. Splunk's Episode Review documentation specifically describes acknowledgment as moving an episode from *New* to *In Progress* and assigning it to the acknowledging user. See the Splunk ITSI Episode Review overview and the current ITSI Episode Review documentation.

QUESTION NO: 7

Which of the following items describe ITSI teams? (select all that apply)

A. Teams should have itoa admin roles added with read-only permissions for services and entities.

B. Services should be assigned to the 'global' team if all users need access to it.

C. By default, all services are owned by the built-in 'global' team and administered by the 'itoa_admin' role.

D. A new team admin role should be created for each team. The new role should inherit the 'itoa_team_admin' role.

ANSWER: B C D

Explanation:

Services should be assigned to the 'global' team if all users need access to it. The built-in global team is intended for ITSI content that must be broadly available, rather than restricted to a specific operational group. Team assignment is a central part of ITSI's object-access model, so associating a service with the appropriate team makes the service available to the roles assigned to that team.

By default, all services are owned by the built-in 'global' team and administered by the 'itoa_admin' role. This provides an initial administrative ownership model for services and allows ITSI administrators to manage services before responsibility is delegated to more specific teams. Teams can subsequently be used to organize service ownership and limit administrative scope.

A new team admin role should be created for each team. The new role should inherit the 'itoa_team_admin' role. Inheriting from the built-in team-administrator role supplies the baseline ITSI capabilities required for team-level administration, while a dedicated role lets administrators assign users and permissions according to the responsibilities of that particular team. This supports delegated management without requiring every team administrator to receive unrestricted platform-wide administration. See Splunk's documentation on configuring ITSI teams and configuring services.

QUESTION NO: 8

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:

itsi_summary_metrics is the correct index because Splunk IT Service Intelligence stores KPI values in this dedicated metrics summary index. ITSI calculates KPI measurements from each KPI's base search and writes the resulting time-series metric data to `itsi_summary_metrics`. This design lets ITSI efficiently retrieve KPI history for service health, glass tables, dashboards, threshold evaluation, and analytics without repeatedly running the original source searches over the full data set.

The index is metrics-based, meaning KPI values are stored as metrics with dimensions that identify the relevant ITSI objects and context. It is part of ITSI's internal data model and is created and maintained for the KPI summarization workflow. Administrators should recognize this index when validating KPI data retention, reviewing index storage, configuring related performance considerations, or investigating whether KPI values have been collected successfully. Splunk's ITSI documentation explicitly identifies `itsi_summary_metrics` as the metrics-based summary index used for KPI data. See the ITSI metrics index reference and the ITSI KPI search documentation.

QUESTION NO: 9

Which of the following can be configured for an ITSI correlation search? (Choose all that apply.)

A. Search schedule.

B. Trigger condition.

C. Notable event action.

D. KPI importance.

ANSWER: A B C

Explanation:

A correlation search is a scheduled saved search that can evaluate search results and generate notable events. Administrators can configure the search schedule, define the triggering criteria that determine when the search alerts, and configure a notable-event action. The notable-event action defines information used to create the event, such as title, severity, description, and other event fields.

KPI importance is configured on a KPI and affects the contribution of that KPI to a service health score. It is not a correlation-search setting. See Configure correlation searches in ITSI for additional details.

QUESTION NO: 10

Which anomaly detection algorithm fulfills the paired monitoring requirement?

A. Detection algorithm: Trending anomaly detectionMonitoring requirement: Produce an alert when an entity deviates from its historical behavior.

B. Detection algorithm: Entity cohesion anomaly detectionMonitoring requirement: Produce an alert when one entity in the KPI is not behaving similar to other entities in the KPI.

C. Detection algorithm: Trending anomaly detectionMonitoring requirement: Produce an alert when one entity in the KPI is not behaving similar to other entities in the KPI.

D. Detection algorithm: Entity cohesion anomaly detectionMonitoring requirement: Produce an alert when multiple KPIs in the service deviate from their historical behaviors.

ANSWER: B

Explanation:

Detection algorithm: Entity cohesion anomaly detection
Monitoring requirement: Produce an alert when one entity in the KPI is not behaving similar to other entities in the KPI.

Entity cohesion anomaly detection is designed for a KPI that is split by entity and whose entities are expected to exhibit comparable behavior. It evaluates the relationship among the entity-level time series rather than treating each entity solely as an independent historical baseline. This makes it appropriate for paired or peer-based monitoring: for example, identifying a host, region, store, or application instance whose KPI behavior diverges materially from the rest of its peer group.

In ITSI, the KPI must provide entity-level values so that the algorithm has a collection of entities to compare. Entity cohesion normalizes the participating entity series and identifies outlying behavior relative to the group. The resulting anomaly can then drive the KPI's alerting behavior, surfacing the specific entity that no longer exhibits cohesion with its peers. This directly implements a requirement to alert when one member of a KPI population behaves differently from the others, rather than merely identifying an unusual aggregate trend.

See Splunk's documentation on anomaly detection in IT Service Intelligence and configuring KPIs.

QUESTION NO: 11

Where is ITSI configuration data for objects such as services, entities, and glass tables primarily stored?

A. KV Store.

B. The _internal index.

C. The itsi_summary_metrics index.

D. A CSV lookup file.

ANSWER: A

Explanation:

KV Store. is correct. ITSI uses Splunk KV Store collections to persist configuration and operational data for many ITSI objects, including services, entities, service templates, glass tables, and event-management records. This allows ITSI to store structured object definitions and retrieve them for the ITSI user interface and supporting workflows.

KV Store configuration data is separate from indexed event data and KPI summary data. Administrators should include ITSI KV Store data in their backup and recovery planning. See Back up and restore ITSI configuration data.