DUMPSQUEEN

Splunk IT Service Intelligence Certified Admin Exam

Splunk SPLK-3002

Version Demo

Total Demo Questions: 10

Total Premium Questions: 53

Buy Premium PDF

https://dumpsqueen.com

support@dumpsqueen.com

dumpsqueen.com

QUESTION NO: 1

When in maintenance mode, which of the following is accurate?

A. Once the window is over, KPIs and notable events will begin to be generated again.

B. KPIs are shown in blue while in maintenance mode.

C. Maintenance mode slots are scheduled on a per hour basis.

D. Service health scores and KPI events are deleted until the window is over.

ANSWER: A

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/REBestPractice

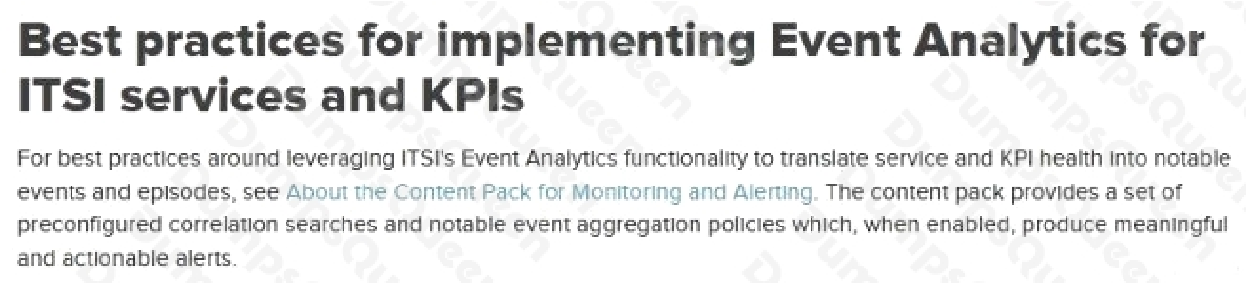

Best practices for implementing Event Analytics for ITSI services and KPIs

For best practices around leveraging ITSI's Event Analytics functionality to translate service and KPI health into notable events and episodes, see About the Content Pack for Monitoring and Alerting. The content pack provides a set of preconfigured correlation searches and notable event aggregation policies which, when enabled, produce meaningful and actionable alerts.

QUESTION NO: 2

Which of the following is a best practice when configuring maintenance windows?

A. Disable any glass tables that reference a KPI that is part of an open maintenance window.

B. Develop a strategy for configuring a service’s notable event generation when the service’s maintenance window is open.

C. Give the maintenance window a buffer, for example, 15 minutes before and after actual maintenance work.

D. Change the color of services and entities that are part of an open maintenance window in the service analyzer.

ANSWER: C

Explanation:

It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work.

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/AboutMW

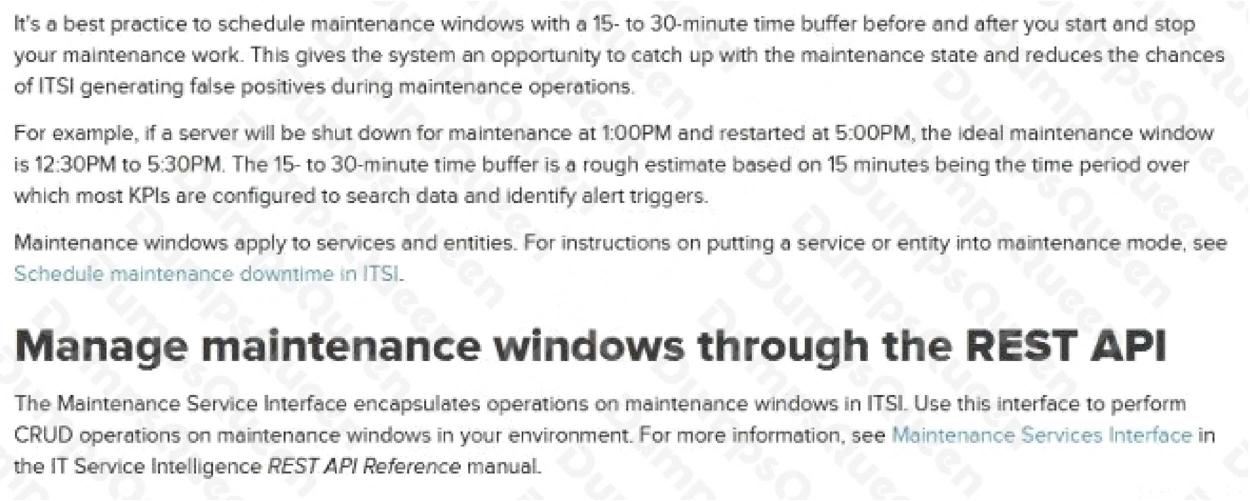

It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work. This gives the system an opportunity to catch up with the maintenance state and reduces the chances of ITSI generating false positives during maintenance operations.

For example, if a server will be shut down for maintenance at 1:00PM and restarted at 5:00PM, the ideal maintenance window is 12:30PM to 5:30PM. The 15- to 30-minute time buffer is a rough estimate based on 15 minutes being the time period over which most KPIs are configured to search data and identify alert triggers.

Maintenance windows apply to services and entities. For instructions on putting a service or entity into maintenance mode, see Schedule maintenance downtime in ITSI.

Manage maintenance windows through the REST API

The Maintenance Service Interface encapsulates operations on maintenance windows in ITSI. Use this interface to perform CRUD operations on maintenance windows in your environment. For more information, see Maintenance Services Interface in the IT Service Intelligence *REST API Reference* manual.

QUESTION NO: 3

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: A B D

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

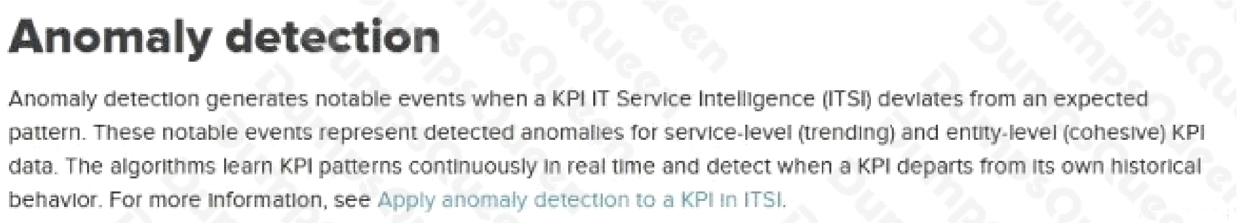

Anomaly detection

Anomaly detection generates notable events when a KPI IT Service Intelligence (ITSI) deviates from an expected pattern. These notable events represent detected anomalies for service-level (trending) and entity-level (cohesive) KPI data. The algorithms learn KPI patterns continuously in real time and detect when a KPI departs from its own historical behavior. For more information, see Apply anomaly detection to a KPI in ITSI.

https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/NotableEvents

QUESTION NO: 4

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:

When an episode warrants investigation, the analyst acknowledges the episode, which moves the status from New to In Progress. Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/EA/EpisodeOverview

QUESTION NO: 5

Which of the following items apply to anomaly detection? (Choose all that apply.)

A. Use AD on KPIs that have an unestablished baseline of data points. This allows the ML pattern to perform it’s magic.

B. A minimum of 24 hours of data is needed for anomaly detection, and a minimum of 4 entities for cohesive analysis.

C. Anomaly detection automatically generates notable events when KPI data diverges from the pattern.

D. There are 3 types of anomaly detection supported in ITSI: adhoc, trending, and cohesive.

ANSWER: B C

Explanation:

The KPI must be split by entity, and a minimum of four entities is required.

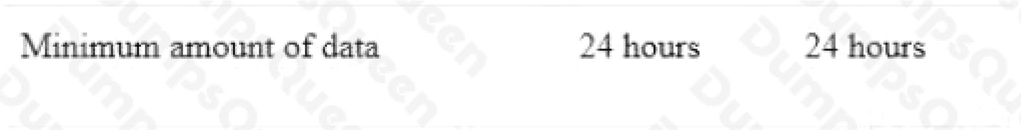

Minimum amount of data	24 hours	24 hours

If the KPI diverges from the normal pattern, ITSI creates a notable event in Episode Review.

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/AD

QUESTION NO: 6

What are valid ITSI Glass Table editor capabilities? (Choose all that apply.)

A. Creating glass tables.

B. Correlation search creation.

C. Service swapping configuration.

D. Adding KPI metric lanes to glass tables.

ANSWER: A C D

Explanation:

Create a glass table to visualize and monitor the interrelationships and dependencies across your IT and business services. The service swapping settings are saved and apply the next time you open the glass table.

You can add metrics like KPIs, ad hoc searches, and service health scores that update in real time against a background that you design. Glass tables show real-time data generated by KPIs and services. Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/GTOverview

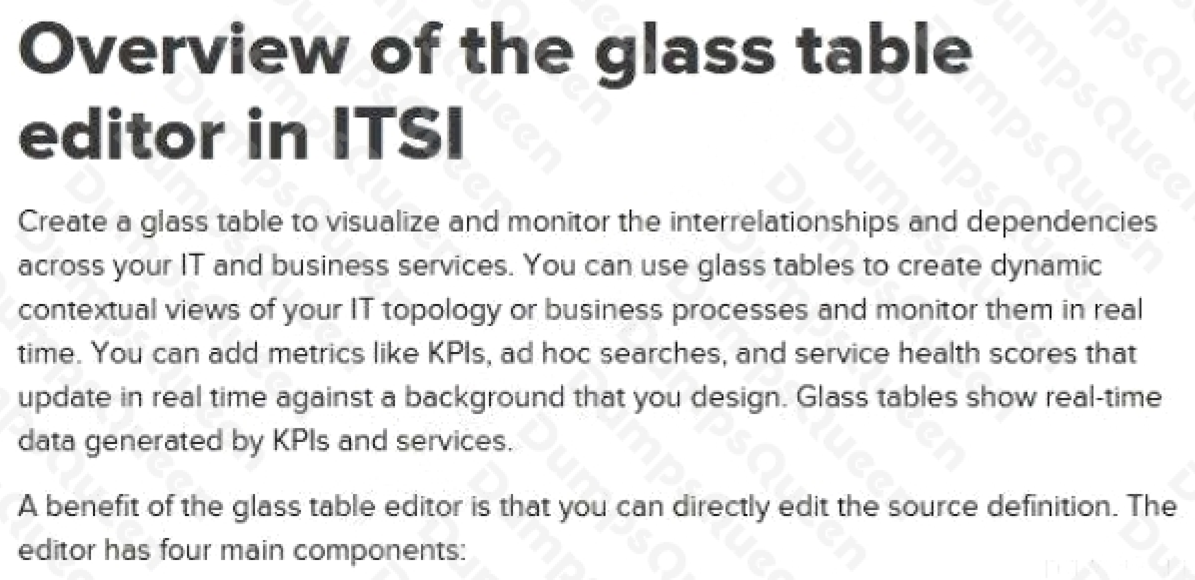

Overview of the glass table editor in ITSI

Create a glass table to visualize and monitor the interrelationships and dependencies across your IT and business services. You can use glass tables to create dynamic contextual views of your IT topology or business processes and monitor them in real time. You can add metrics like KPIs, ad hoc searches, and service health scores that update in real time against a background that you design. Glass tables show real-time data generated by KPIs and services.

A benefit of the glass table editor is that you can directly edit the source definition. The editor has four main components:

https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/ServiceSwap

QUESTION NO: 7

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

A. Service access control requirements for ITSI Team Access should be considered, and appropriate teams provisioned prior to creating the ITSI Service.

B. Entities, entity meta-data, and entity rules should be planned carefully to support the service design and configuration.

C. Services, entities, and saved searches are stored in the ITSI app, while events created by KPI execution are stored in the itsi_summary index.

D. Backfill of a KPI should always be selected so historical data points can be used immediately and alerts based on that data can occur.

ANSWER: A C

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/Configure/ImplementPerms

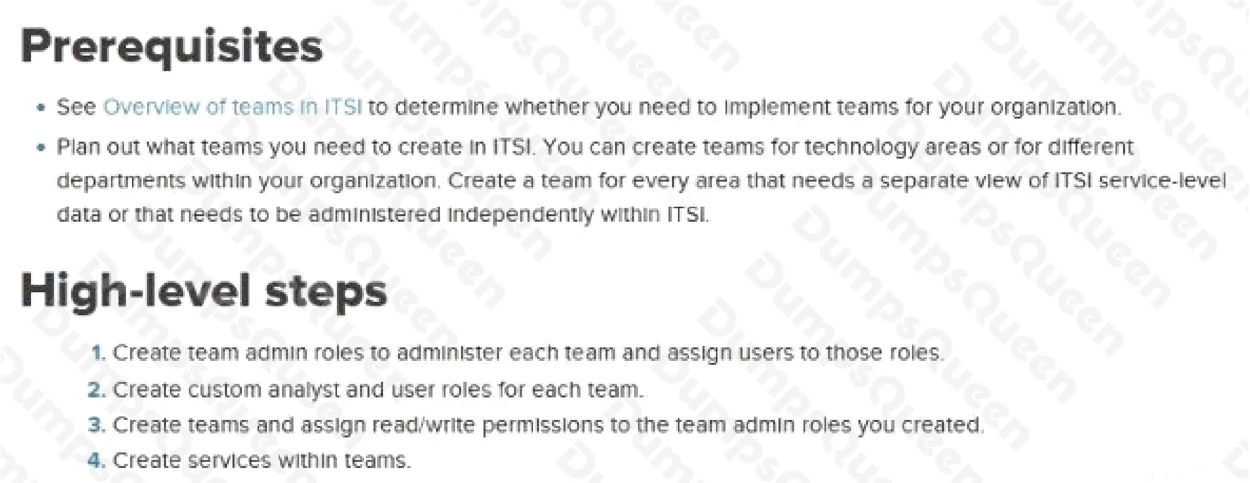

Prerequisites

- See Overview of teams in ITSI to determine whether you need to implement teams for your organization.
- Plan out what teams you need to create in ITSI. You can create teams for technology areas or for different departments within your organization. Create a team for every area that needs a separate view of ITSI service-level data or that needs to be administered independently within ITSI.

High-level steps

1. Create team admin roles to administer each team and assign users to those roles.
2. Create custom analyst and user roles for each team.
3. Create teams and assign read/write permissions to the team admin roles you created.
4. Create services within teams.

QUESTION NO: 8

Which of the following items describe ITSI Deep Dive capabilities? (Choose all that apply.)

A. Comparing a service's notable events over a time period.

B. Visualizing one or more Service KPIs values by time.

C. Examining and comparing alert levels for KPIs in a service over time.

D. Comparing swim lane values for a slice of time.

ANSWER: B C D

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/DeepDives

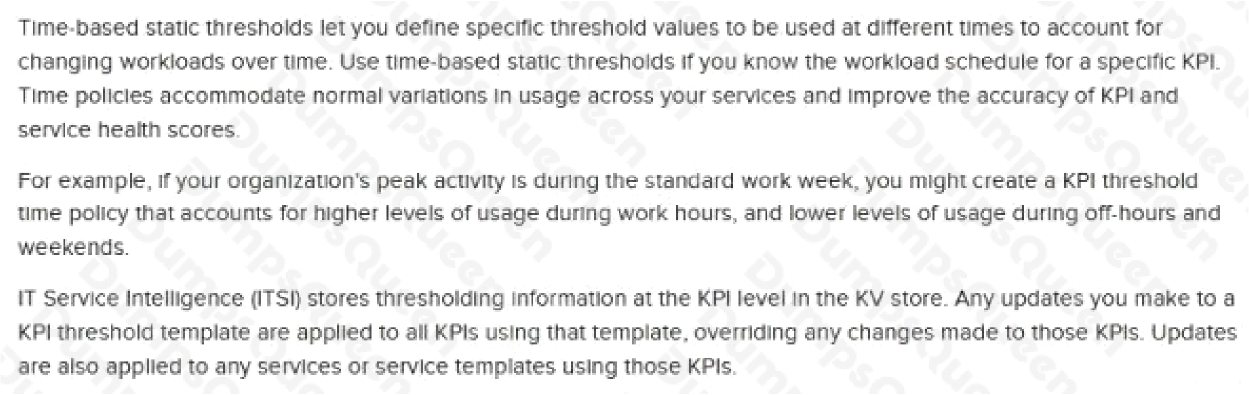

Time-based static thresholds let you define specific threshold values to be used at different times to account for changing workloads over time. Use time-based static thresholds if you know the workload schedule for a specific KPI. Time policies accommodate normal variations in usage across your services and improve the accuracy of KPI and service health scores.

For example, if your organization's peak activity is during the standard work week, you might create a KPI threshold time policy that accounts for higher levels of usage during work hours, and lower levels of usage during off-hours and weekends.

IT Service Intelligence (ITSI) stores thresholding information at the KPI level in the KV store. Any updates you make to a KPI threshold template are applied to all KPIs using that template, overriding any changes made to those KPIs. Updates are also applied to any services or service templates using those KPIs.

QUESTION NO: 9

Anomaly detection can be enabled on which one of the following?

A. KPI

B. Multi-KPI alert

C. Entity

D. Service

ANSWER: A

Explanation:

Enable anomaly detection to identify trends and outliers in KPI search results that might indicate an issue with your system.
Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/SI/AD

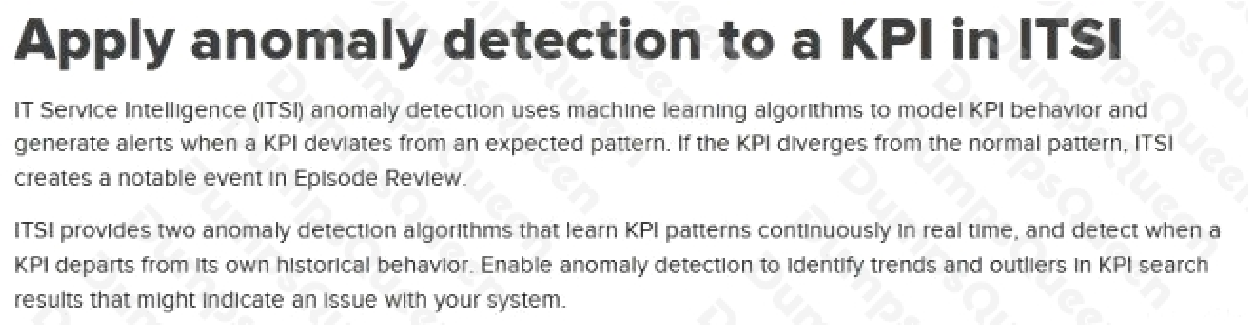

Apply anomaly detection to a KPI in ITSI

IT Service Intelligence (ITSI) anomaly detection uses machine learning algorithms to model KPI behavior and generate alerts when a KPI deviates from an expected pattern. If the KPI diverges from the normal pattern, ITSI creates a notable event in Episode Review.

ITSI provides two anomaly detection algorithms that learn KPI patterns continuously in real time, and detect when a KPI departs from its own historical behavior. Enable anomaly detection to identify trends and outliers in KPI search results that might indicate an issue with your system.

QUESTION NO: 10

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: A

Explanation:

Reference: https://docs.splunk.com/Documentation/ITSI/4.10.2/IModules/Troubleshootingmodules

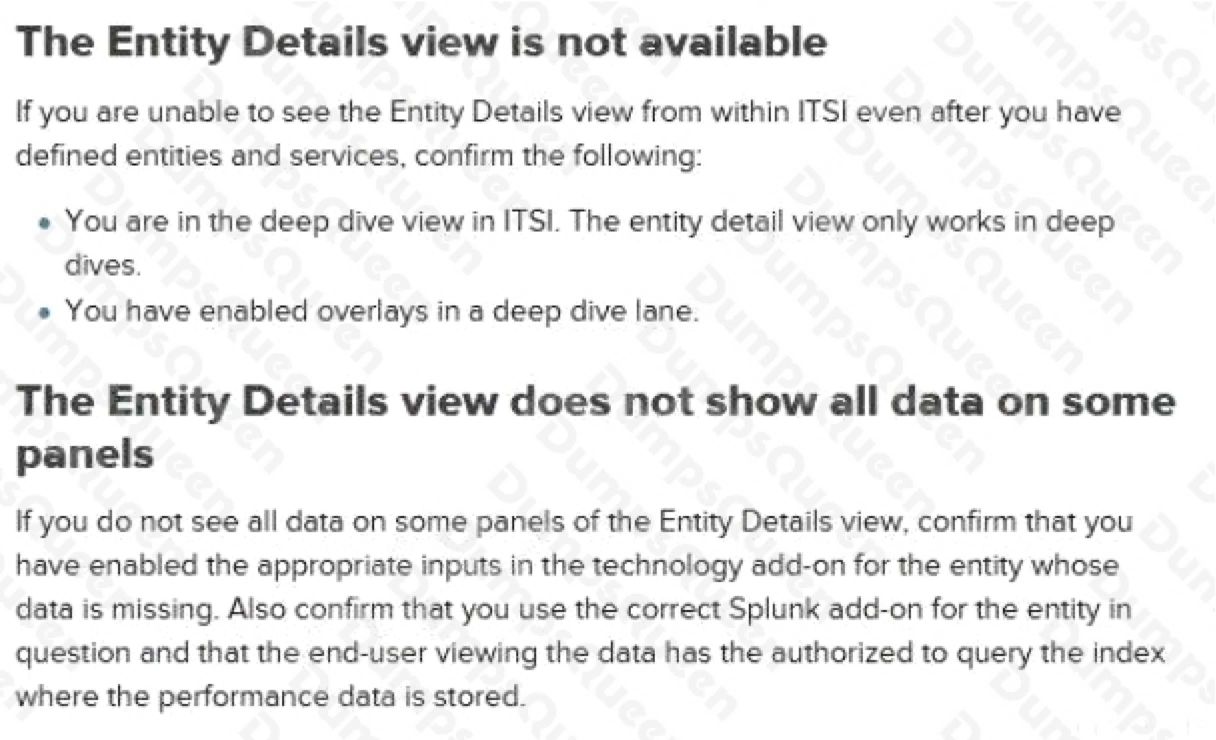

The Entity Details view is not available

If you are unable to see the Entity Details view from within ITSI even after you have defined entities and services, confirm the following:

- You are in the deep dive view in ITSI. The entity detail view only works in deep dives.
- You have enabled overlays in a deep dive lane.

The Entity Details view does not show all data on some panels

If you do not see all data on some panels of the Entity Details view, confirm that you have enabled the appropriate inputs in the technology add-on for the entity whose data is missing. Also confirm that you use the correct Splunk add-on for the entity in question and that the end-user viewing the data has the authorized to query the index where the performance data is stored.