Getting Started with Splunk IT Service Intelligence (ITSI)

Table of Contents

Introduction

Hey there! Thanks for checking out this Getting Started Guide. You may be exploring Splunk IT Service Intelligence (ITSI), or already have it but are unsure how best to utilize it. Well this guide is here to help!

We put this guide together to answer all of your questions and help you feel confident navigating around the platform. But most of all, we know ITSI has an immense amount of value to offer you and your organization so we hope this guide leaves you feeling empowered when using ITSI.

Let's get started!

Product Overview

ITSI consists of two primary components: **Service Insights** and **Event Analytics**.

Here is a high-level view of Splunk ITSI and its key capabilities:

Service Insights allows you to visualize the health of your business and IT services as well as the health of related services. This top-down mapping approach helps correlate business services to technical processes in the underlying infrastructure layer, enabling you to quickly identify and triage problems and identify root causes directly from the service layer.

Event Analytics uses events from across your IT landscape and various monitoring tools to provide a unified operational console of all your events and service-impacting issues. By integrating event analytics with incident management tools and help desk applications, you can accelerate incident investigation and automate remedial actions. Event analytics uses machine learning for real-time clustering and automation prioritization to reduce event noise — detecting atypical conditions and only enabling alerts when behavior strays from the predetermined norm.

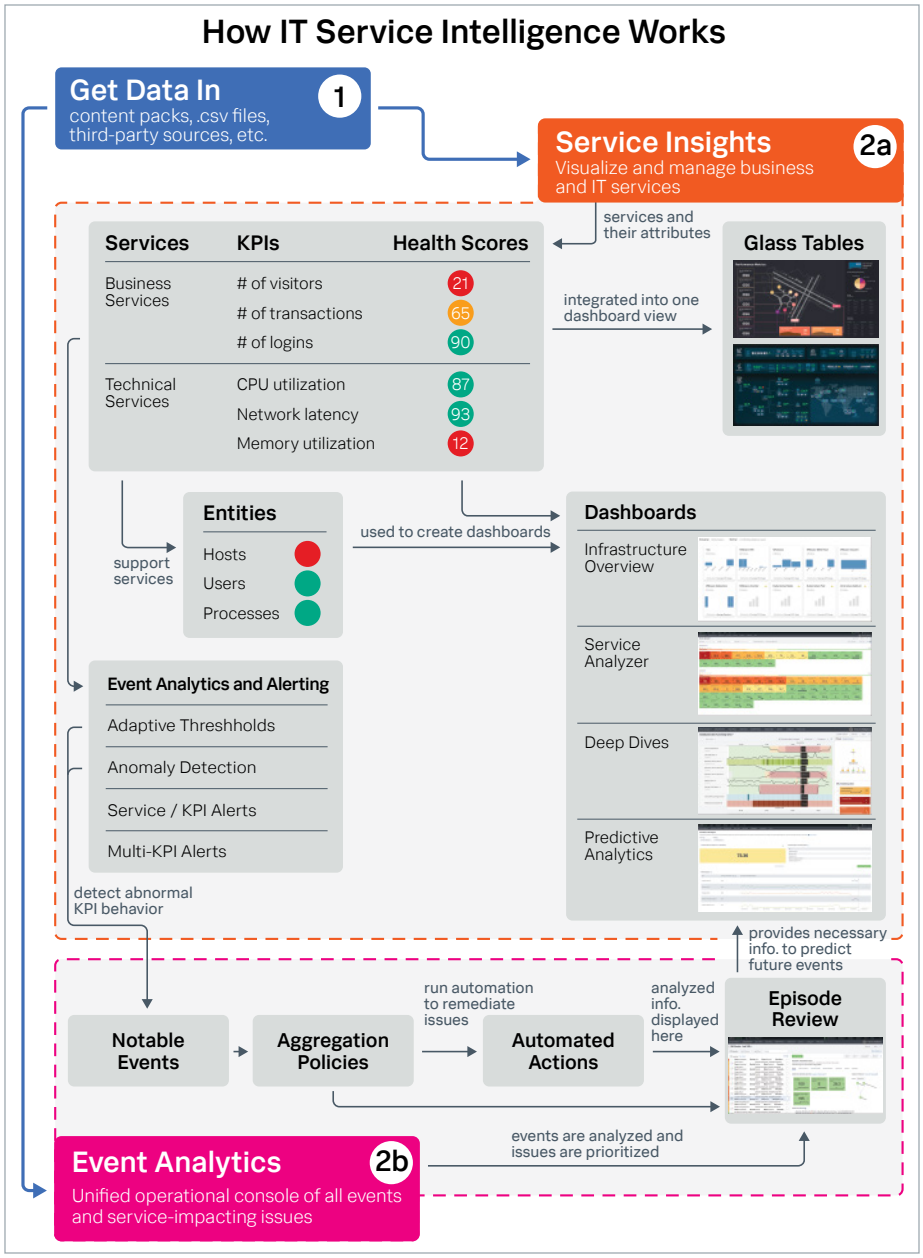

How ITSI Can Make Your Job Easier

There are four primary challenges IT and service teams face as they support digital transformation initiatives with complex environments:

- Lack of visibility into the health of the business
- Poor service performance
- High-severity outages and slowdowns
- Slow incident response and resolution times

These operational issues put your revenue, customer experience, employee effectiveness and innovation at risk. The unfortunate reality is that legacy IT tools just aren't equipped to handle the way businesses operate today — customer-focused, service-centric, and increasingly hybrid and interconnected digital businesses.

Splunk ITSI addresses these challenges by applying machine learning to data for 360° service monitoring, predictive analytics and streamlined incident management.

Here are some ways that ITSI can make your job easier:

Business and Technical Service Level Monitoring	• ITSI is customizable and monitors performance in the way that the business operates • Protect business service-level agreements with dashboards to: - Monitor service health - Troubleshoot alerts - Perform root cause analysis • Achieve end-to-end visibility across your entire IT environment
Machine Learning and Predictive Analytics	• ITSI uses advanced analytics like: - Anomaly detection - Adaptive thresholding - Predictive health scores (see Service Health Scores under Features and Definitions > Service Insights Terms) • Monitors key performance indicators (KPI) data and prevents issues 30 minutes in advance
Intelligent Incident Management	• ITSI accelerates mean-time-to-resolution (MTTR) using: - Event correlation and clustering using machine learning - Automated incident prioritization - Integrations with IT service management (ITSM) tools

Get Started with ITSI

Step 1: Getting Data In

There are two ways to get data into ITSI: entities and content packs.

Entities & Entity Integration

Entities and entity integrations are used to collect and aggregate data into Splunk ITSI. Data is collected into what we call Entities – you could define entities any way that fits your needs, but this usually includes data from servers, DNS groups, firewalls or other devices. Data can be metrics, logs, traces — anything that helps you gain better visibility into the health of the services you are responsible for. Data is streamed and collected from native systems or management/monitoring tools like Splunk Infrastructure Monitoring.

All entities exist in the Global team and can be created in a few ways:

Manually create a single entity in ITSI

Create a single entity in ITSI to associate events your Splunk platform deployment receives.

Prerequisites: you have to log in as a user with the itoa_admin or itoa_team_admin ITSI role.

For more information, see Documentation - create a single entity in ITSI.

Manually import entities from a Splunk search in ITSI

Create entities from ITSI module searches, saved searches or ad hoc searches using indexed data coming into your Splunk platform deployment.

ITSI uses the `itsiimportobjects` command to import entities from searches.

Warning: You can import a maximum of 50,000 entities at a time in ITSI. If you attempt to import more than 50,000 entities, only the first 50,000 are imported.

Prerequisites:

- ITSI role: You have to log in as a user with the itoa_admin or itoa_team_admin ITSI role and access to the Global team.
- Indexed data: You must have already indexed data you want to associate with entities.

For more information, see Documentation - import entities from a search in ITSI.

Manually import entities from a CSV file in ITSI

Importing entities from CSV files is an efficient way to define multiple entities. You can dump data from a change management database (CMDB) or asset inventory database into a CSV file and automate the import for ongoing updates.

ITSI uses the `itsiimportobjects` command to import entities from a CSV file. All events your Splunk platform deployment indexes from a manual entity import from a CSV file is stored in the `itsi_import_objects` index, and each event has the `itsi_import_objects:csv` source type.

Warning: You can import a maximum of 50,000 entities at a time in ITSI. If you attempt to import more than 50,000 entities, only the first 50,000 are imported.

Prerequisites:

- **ITSI role:** You have to log in as a user with the itoa_ admin ITSI role.
- **CSV file:** You must have a CSV file that contains entity definitions. Specify column names in the first row. In each subsequent row, specify an entity title and entity type, as well as one or more entity aliases, and one or more entity information fields. To associate an entity with a service, provide a column with the name of the service. Importing from a CSV file has a limit of one service and one entity per row. There is no limit on the number of dependent services, entity aliases or entity rule values per row. A CSV file can contain multiple rows. Importing from a CSV file supports five different separators: comma (,), semicolon (;), pipe (|), tab (\t) and caret (^).

In this example you want to create two entities called appserver-04 and appserver-05, and associate appserver-04 with the Web A service and associate appserver-05 with the Web B service. The Web A service already exists in ITSI but the Web B service does not. The following image shows the CSV file to import:

ITService	service_desc	host	vendor	IP	itsi_role
Web A	This is Web A	appserver-04	Linux	10.2.1.133	web_server
Web B	This is Web B	appserver-05	Microsoft	10.2.1.134	web_server

For more information, see Documentation - import entities from a CSV file in ITSI.

After you import entities either by creating single entities or from a Splunk search, you can configure recurring imports to update existing entities and create new entities. However, you can't set up a recurring entity import from a CSV file. To configure recurring entity imports from data that's stored in a CSV file, you have to configure a universal forwarder to monitor the CSV file and send data to your Splunk platform deployment, run an entity import from a Splunk search, and configure a recurring import from the Splunk search.

For more information, see Set up a recurring entity import from a CSV file.

You can also automatically create entities and collect data on a recurring basis with ITSI entity integrations. The integrations that are available are:

- Unix and Linux entity integration in ITSI
- Windows entity integration in ITSI
- VMware vSphere entity integration in ITSI
- Splunk Infrastructure Monitoring entity integration in ITSI

To learn more, see Overview of entity integrations in ITSI.

Content Packs and Splunk App for Content Packs

If Content packs are individual preconfigured packs that provide capabilities for a specific use case. They can be installed directly within ITSI. Many content packs include service templates, so you can easily link one of your existing services to predetermined key performance indicators (KPIs), allowing for efficient setup and easier integrations.

Splunk App for Content Packs is a free Splunkbase app for ITSI (version 4.9 and later) that acts as a one-stop shop for content packs, and out-of-the-box searches and dashboards for common IT infrastructure monitoring sources. With this app, you no longer need to use the backup/restore functionality to install content packs. Instead, the app contains a library of readily updated content packs and is used to update all of them, rather than individually updating each content pack.

For guidance on how to onboard data for use with the desired Content Pack, see the details included in the Content Pack documentation for how to get data in.

- **Prerequisite:** You must have command-line access and Splunk admin access to an ITSI v4.9 or later instance
- **Step 1:** Download the Splunk App for Content Packs on Splunkbase.
- **Step 2:** Install the app per the instructions on the Splunk Docs page.
- **Step 3:** Go to **Configuration > Data Integrations** to see the available content packs. Data Integrations is the top-level GDI guidance we give for common data sources like Unix and Linux, Windows, and Splunk Infrastructure Monitoring.

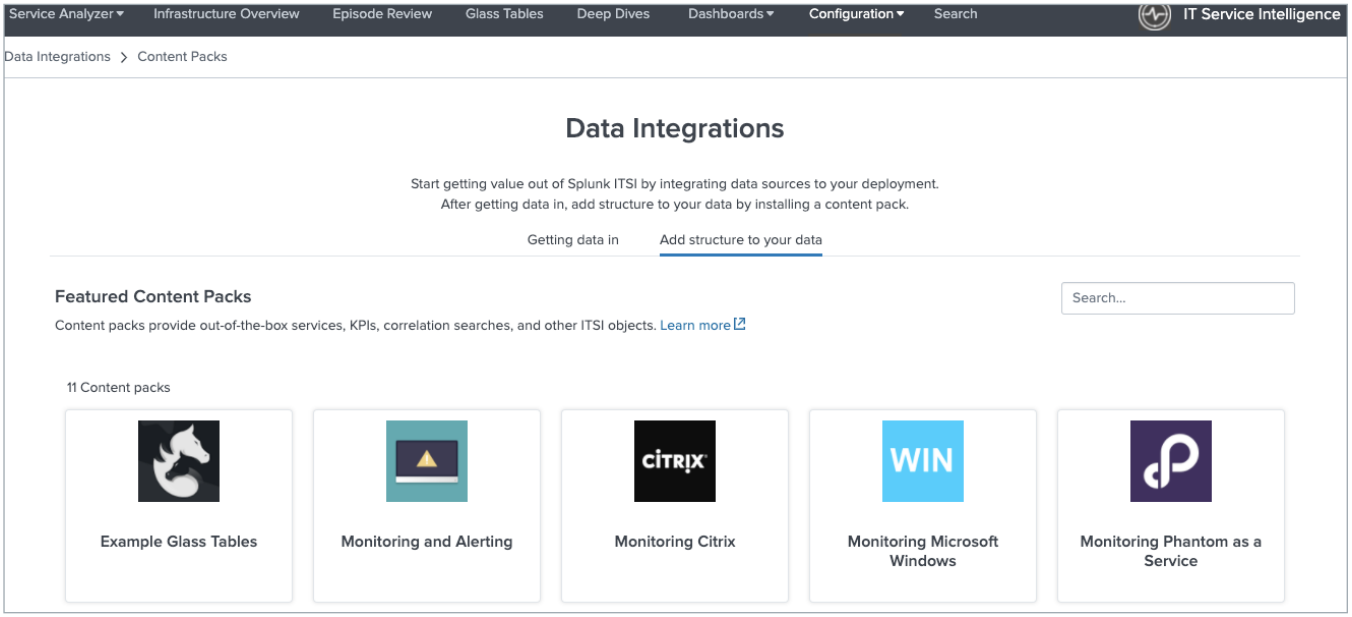

Make sure to install the associated Add-On for the Content Pack you downloaded! For example, there is a corresponding Unix and Linux Add-On that works with the Monitoring Unix and Linux content pack.

For more information regarding: how to install the Splunk App for Content Packs on a Splunk Cloud Platform or on-premises environments, how to install content packs for ITSI version 4.8.x and below, and to see a list of available content packs, see the Splunk Content Packs Manual.

Resources

- Blog post: Introducing Splunk App for Content Packs
- Blog post: Content Pack for Microsoft Exchange

Step 2a: Services and Service Insights

A service is a set of interconnected applications and hosts that are configured to offer a specific service to the organization. These services can be internal — an organization's email system — or external — an organization's website.

You can create business and technical services that model those within your environment. Some services might have dependencies on other services. Services contain key performance indicators (KPIs), which make it possible to monitor service health via service health scores, perform root cause analysis, receive alerts, and ensure that your IT operations are in compliance with business service-level agreements (SLAs).

ITSI Service Insights allow you to create glass tables to help you monitor real-time interrelationships and dependencies via KPIs and service health scores across your IT and business services in one view. Glass tables also feature a drawing canvas where you can add visualizations in the form of KPIs and service health scores, upload images and icons, and add charts.

Here are two examples of glass tables:

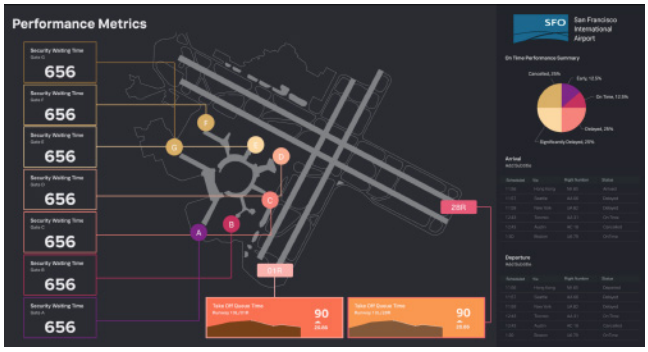

A glass table used to evaluate performance metrics at San Francisco International Airport.

Another glass table used to evaluate business, operational and SLA performances, along with infrastructure status.

Service Insights also enables you to create and use four kinds of dashboards: infrastructure overview, service analyzer, deep dives and predictive analytics.

Infrastructure Overview Dashboards provide a consolidated view of all of your entities grouped by type with drill downs to entity-specific dashboards for operating systems, virtual infrastructures, containers and cloud services.

Here is an example of an infrastructure overview dashboard:

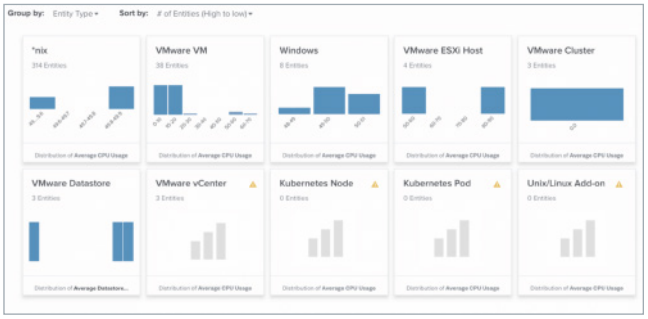

Service Analyzer Dashboards provide instant visibility into the health of your services and KPIs as either a tile view (top 50 worst performing Services and KPIs) or tree view (all selected services and their dependencies). You can visually correlate services to underlying infrastructure with a tile or tree view.

Here is a tile view example of a service analyzer dashboard:

Deep Dives Dashboards are an investigative tool to help you identify and analyze issues in your IT environment.

Here's what a deep dive dashboard looks like:

Deep dives display a swim lane view of KPIs and service health scores over time to help you zoom in on metric and log data and visually correlate root cause.

Use swim lane displays of multiple KPIs and correlate metrics over time to identify root causes.

Predictive Analytics Dashboards predict future incidents 30 minutes in advance using machine learning algorithms and historical service health scores.

Here is an example of a predictive analytics dashboard:

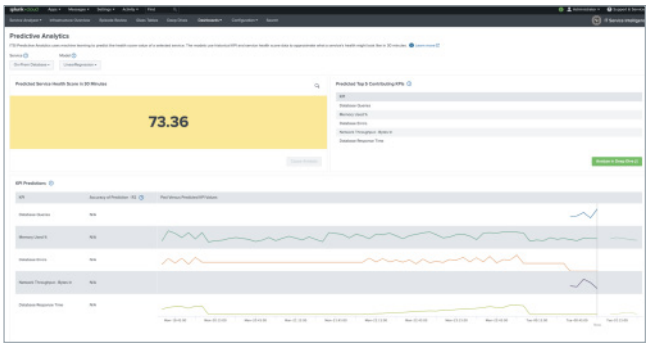

Top five contributing service metrics are displayed to guide troubleshooting.

To learn more about these data models, see the Splunk ITSI interactive demo.

The following are included in the demo: Glass Tables, Predictive Analytics Dashboard, Service Analyzer Dashboard and Deep Dive Dashboard.

Service Modeling and Service Decomposition

Before you are ready to set up your dashboards and services in Splunk ITSI, it's important to identify what services will provide the most value.

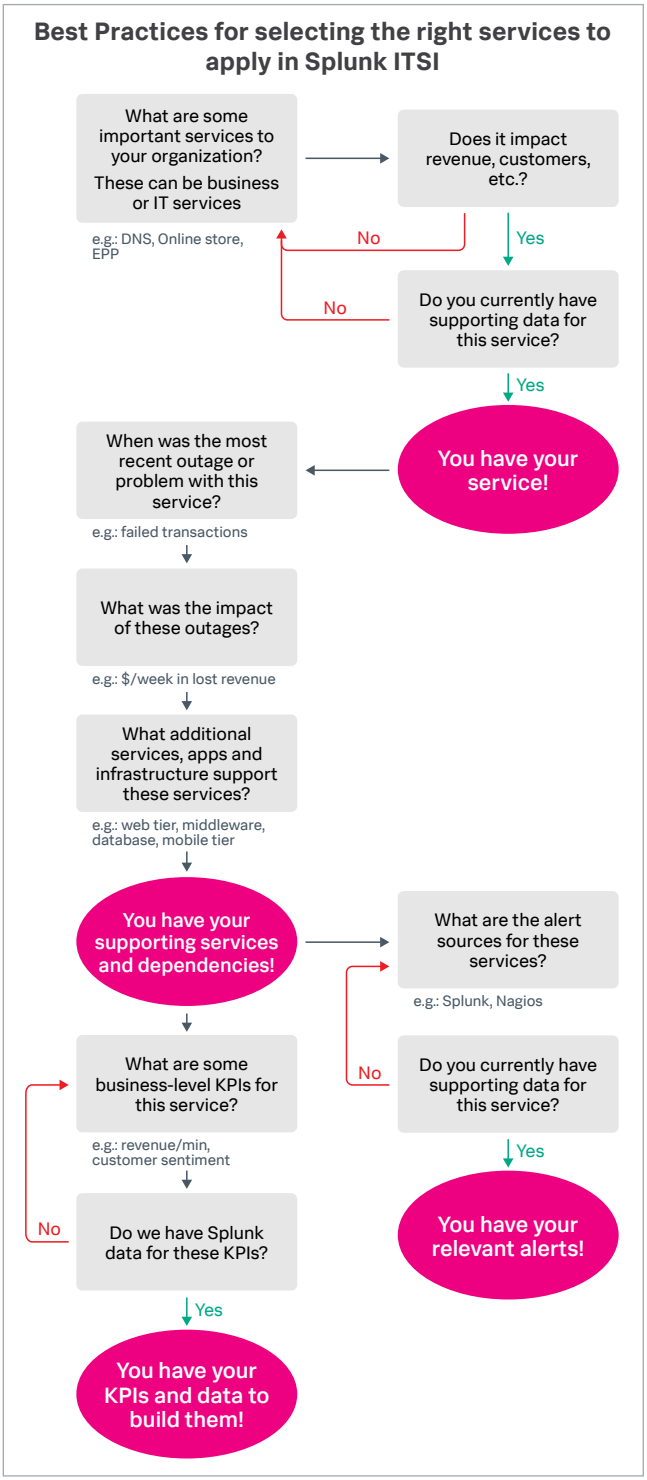

To learn more, see Tech Talk: Service Decomposition, ITSI Service Discovery Workshop, and ITSI Implementation Success offering.

Create KPIs for Your Services

ITSI KPIs are Key Performance Indicators that are helpful in determining the health of the service they belong to. KPIs are recurring saved searches that return the value of an IT performance metric. They are created within a specific service and define everything needed to generate searches to understand the underlying data, including how to access, aggregate, and qualify with thresholds. There are two types of KPIs: business and technical.

Doing pre-work with service decomposition to correctly identify what services are most valuable to the organization is a good first step to identifying appropriate KPIs to map to these services. Please schedule time with your account team or leverage OnDemand Services to be guided through a service discovery and decomposition workshop.

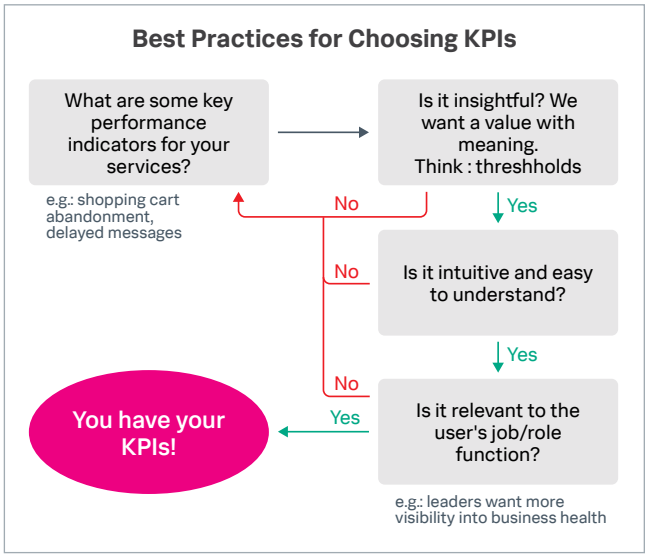

Good KPIs have the following characteristics:

- Simple
- Use base searches
- Provide data regularly
- Self normalizing data
- Data with deltas, not counters

You can also use Content Packs for preconfigured services and KPIs. Here are some KPIS available in the Microsoft 365 Content Pack:

Availability KPIs	• Extended recovery • False positive • Investigating • Restoring service • Normal service
Performance KPIs	• Added delegation entry • Added service principal • Set company information • Set password policy
Group Administration Activities KPIs	• Added group • Added member to group • Deleted group • Updated group
Login Activity KPIs	• Authentication methods • Distinct user sign-ins • Logins by region • Logon errors • User agents • User types
Office 365 Security & Compliance Center	• Mail flow • Elevation of exchange admin privilege • Unusual external user file activity • Potentially malicious URL click was detected

After you have your services, entity rules, KPIs, and service dependencies planned out, you can finally create services in ITSI! There are three ways to do so:

- Create a single service
- Import services from a CSV file
- Import services from a search

For more information regarding creating services, see the Service Insights Manual.

Get Started with Service Insights

Service Insights within Splunk ITSI consists of various dashboard views, alerts and metrics so that you can effectively monitor and map services within your organization. Here are some ways to get better acquainted with the various available features and views.

Tasks to tackle

- ☑ Navigate Service Analyzer
  - Explore the Tile View of Services
  - Filter to the Shared Infrastructure Service and Show Dependencies
  - How many Services are in Shared Infrastructure?
  - Navigate to the Tree View
- ☑ Navigate KPIs and Health Scores
- ☑ Navigate Entities
- ☑ Use the Content Pack for Monitoring and Alerting to create notable events and episodes
- ☑ Configure a service health score alert

White Paper: Modern IT Management with AIOps.

E-Book: A Guide to Modern IT Service Management with AIOps

Additional Resources

- Tech Talk: Service Modeling
- .conf presentation: Top KPIs to consider in Splunk ITSI
- Webinar: Getting Started with ITSI Part 1
- Tech Talk: Top 10 Glass Table Dashboard Design Principles
- How Service Health Scores are calculated
- Blog post: Best Practices for Thresholding KPIs
- Blog post: Best Practices for Alert Configuration
- Blog post: Best Practices for Adaptive Thresholding

Step 2b: Event Analytics

Event Analytics in Splunk ITSI is where you can streamline your incident management workflows, from alert management to incident response triggers.

Get Started with Event Analytics

Tasks to tackle

Ingest events through correlational searches. The data itself comes from Splunk indexes, but ITSI only focuses on a subset of all Splunk Enterprise data. This subset is generated by correlation searches. A correlation search is a specific type of saved search that generates notable events from the search results.

See Overview of correlation searches in ITSI.

Configure aggregation policies to group events into episodes. Once notable events start coming in, they need to be organized so you can start gaining value from them. Configure an aggregation policy to define which notable events are related to each other and group them into episodes. An episode contains a chronological sequence of events that tells the story of a problem or issue. In the backend, a component called the Rules Engine executes the aggregation policies you configure.

For more information, see Overview of aggregation policies in ITSI.

Set up automated actions to take on episodes.

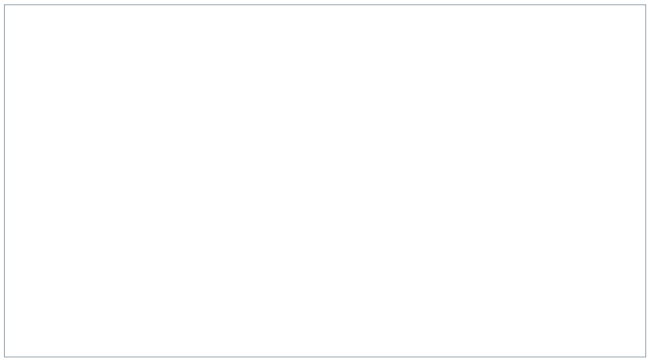

For more information, see Configure episode action rules in ITSI.

You can run actions on episodes either automatically using aggregation policies or manually in Episode Review. Some actions, like sending an email or pinging a host, are shipped with ITSI. You can also create tickets in external ticketing systems like ServiceNow or Remedy. Finally, actions can also be modular alerts that are shipped with Splunk add-ons or apps, or custom actions that you configure.

To learn more about event analytics, see the documentation and Event Analytics section (step 7 / 8) on the Splunk ITSI interactive demo.

Event Analytics Best Practices for Third-Party Data Sources

To avoid duplicate events, use the same frequency and time range in correlation searches.

When configuring a correlation search, consider using the same value for the search frequency and time range to avoid duplicate events. For example, a search might run every five minutes and also look back every five minutes.

If there's latency in your data and you need to look for events you might have missed, consider expanding the time range. For example, the search could run every minute but look back 5 minutes.

To reduce load on your system, don't use a time range greater than 5 minutes.

Exceeding a calculation window of 5 minutes can put a lot of load on your system, especially if you have a lot of events coming in. If you want to avoid putting extra load on your system, consider reducing the time range to 5 minutes or less.

One exception is if your data is coming in more sporadically. For example, if your data comes in every 15 minutes, consider using a 15-minute time range.

Normalize all the important fields in your third-party events.

When you're creating correlation searches, don't only normalize on obvious fields that exist in a lot of data sources, like host, severity, event type, message, and so on. It's also important to normalize fields that you know are important in your events. For example, when you're looking at Windows event logs, what do you look at to know if something is good or bad? Normalize those fields as well and use them to build out a common information model.

Perform this normalization process for every data source you have so you can easily identify important fields when creating aggregation policies.

Create one correlation search per data source.

For every third-party data source you're bringing into ITSI, create a single correlation search to normalize those fields and generate notable events. For example, one for SCOM, one for SolarWinds, and so on.

Don't create too many aggregation policies.

Limit the number of aggregation policies you enable in your environment. Too many aggregation policies create too many groups, which produces an overly granular view of your IT environment. By limiting the number of policies, you create more end-to-end visibility and avoid creating silos of collaboration between groups in your organization. Make sure to group events according to how those events are related, not based on how people work to resolve those issues.

Only select 5-10 fields for Smart Mode analysis.

Additional Resources

- Webinar: Getting Started with ITSI part 2
- Blog post: Event Storms with ITSI

Frequently Asked Questions

What is the deployment process and how long does it take?

For customers who already have Splunk Enterprise or Splunk Cloud, Splunk ITSI can be deployed and configured in a matter of days. ITSI is also included in IT Cloud Plus.

What are the use cases associated with Splunk ITSI?

Service Insights and Monitoring - Gain access to a service and application view between both your business and IT services, at both a high level, so you can understand the health, performance quickly across silos, while also being able to dive into deeper investigations to find the root cause of an incident faster. Use AI powered by machine learning to predict imminent outages 30-40 minutes in advance based on your service health.

Event Analytics and Management - Apply machine learning to your events to reduce event noise, decrease your MTTR, and the number of IT incidents, to go from a reactive to proactive IT.

What is Service Intelligence for SAP?

Service Intelligence for SAP is generally available as of March 31, 2021 and will help ITSI customers understand and monitor their complex SAP environments. For more information on Service Intelligence for SAP, refer to documentation.

What is the Splunk App for Content Packs?

Splunk App for Content Packs is a free app available to ITSI users. It contains a content library of readily updated Content Packs to quickly begin monitoring common IT infrastructure and use cases. For more information on Splunk App for Content Packs, see Features and Definitions.

Is there a demo available?

Yes! An interactive guided tour of ITSI is available and includes information about:

- Glass tables
- Predictive Analytics Dashboard
- Adaptive Thresholding
- Unified Views and Data-to-Everything
- Service Analyzer Dashboard
- Deep Dive Dashboard
- Event Analytics

Where can I see information about the latest release?

You can refer to Splunk ITSI Release Documentation for all new features, enhancements and capabilities.

How can I upgrade my version of ITSI?

Please reach out to your support account team for upgrades.

Features and Definitions

Content Packs and Splunk App for Content Packs

Content packs are individual preconfigured packs that provide capabilities for a specific use case. They can be installed directly within ITSI. Many content packs include service templates, so you can easily link one of your existing services to predetermined KPIs, allowing for efficient setup and easier integrations.

Splunk App for Content Packs is a free application for ITSI (version 4.9 and later) that acts as a one-stop shop for content packs, and out-of-the-box searches and dashboards for common IT infrastructure monitoring sources. With this app, you no longer need to use the backup/restore functionality to install content packs. Instead, the app contains a library of readily updated content packs and is used to update all of them, rather than individually updating each content pack. Getting started with Splunk for IT operations use cases has never been easier!

Check out the release documentation for content packs for more information.

Service Insights Terms

Services

An ITSI service is a set of interconnected applications and hosts that are configured to offer a specific service to the organization. These services can be internal — an organization's email system — or external — an organization's website.

You can create business and technical services that model those within your environment. Some services might have dependencies on other services. Services contain key performance indicators (KPIs), which make it possible to monitor service health via service health scores, perform root cause analysis, receive alerts, and ensure that your IT operations are in compliance with business service-level agreements (SLAs).

For more information about creating services, see Overview of Creating Services in ITSI and ITSI Thresholding Basics.

Entities

An entity is an IT infrastructure component that is managed to support IT/business services.

Each entity is unique: it can be identified based on its specific attributes and relationships to other IT processes.

ITSI entities can be any of the following components:

- Physical, virtual, or cloud resources
- Network devices (switches, routers)
- Users (AD/LDAP)
- Operating systems or processes
- Software application (db, web server, business app)
- Application process instances (ex: 2 instances of the same web server application is 2 separate entities)

Entities contain information ITSI uses to associate services with information found in Splunk searches, imports and integrations. You can use this entity information to filter items according to the entity definition.

An entity is not a service. You must 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).

See the Entity Integrations Manual for more information about importing, defining, and managing entities.

Key Performance Indicators (KPIs)

Key performance indicators are recurring saved searches that return the value of an IT performance metric.

KPIs are created within a specific service and define everything needed to generate searches to understand the underlying data, including how to access, aggregate and qualify with thresholds.

There are two types of KPIs: business and technical. Some examples of business KPIs are number of site visitors, number of transactions and number of logins. On the other hand, a few examples of technical KPIs are CPU load percentage, memory used percentage and response time.

You can use these metrics to measure and ensure that performance remains within acceptable parameters.

For more information, see Overview of creating KPIs in ITSI.

Service Health Scores

A service's health score is the weighted average of the severity levels of a service's KPIs and dependencies. ITSI reflects service health scores with a severity level and color in the Service Analyzer and in the glass tables.

ITSI uses a combination of individual KPI health scores and their importance settings to calculate the overall service health score. After creating a KPI, you may optionally change the importance of the KPI in order to increase or reduce its impact on the service health score.

ITSI considers KPIs that have an importance value of 11 as a special case that represents a "minimum health indicator" for the service. When a KPI with an importance value of 11 reaches the critical state, the overall health score for the service turns critical, regardless of the status of other KPIs in the service.

For more information, see How service health scores are calculated.

Dashboards

Infrastructure Overview Dashboard

A consolidated view of all your data integrations and investigation tools for operating systems, virtual infrastructures, containers, and cloud services.

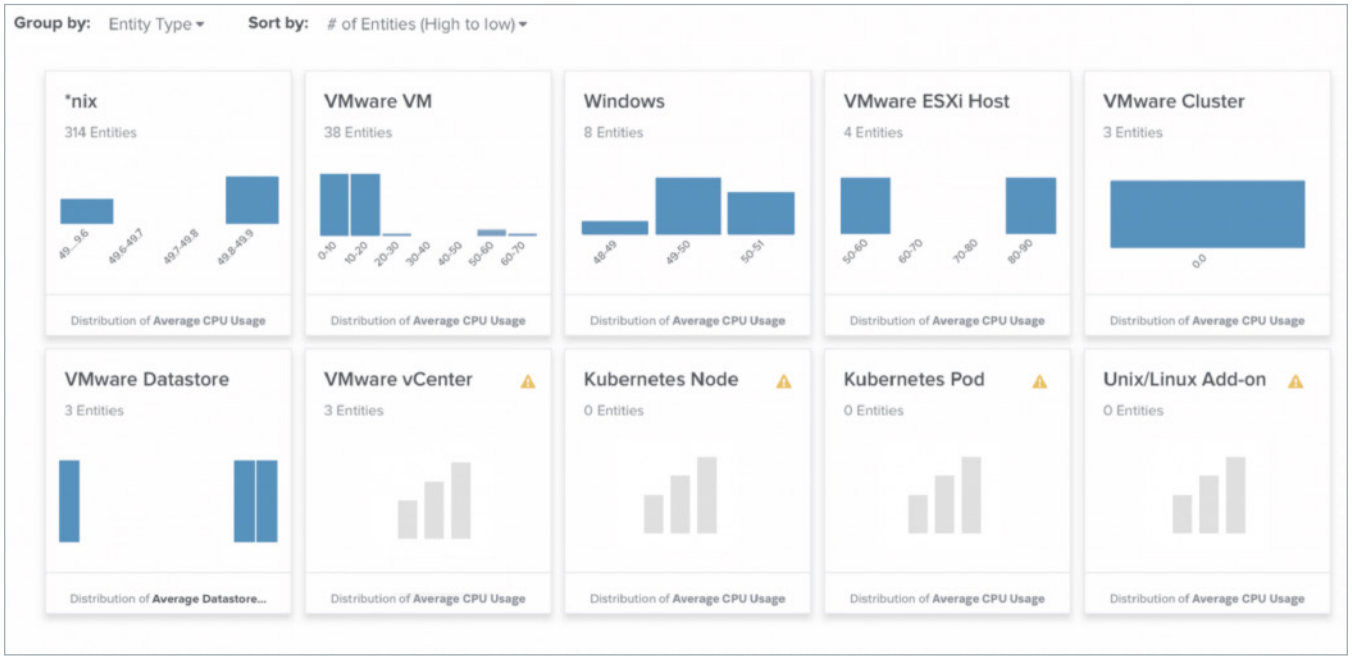

Figure 1: An example of an infrastructure overview dashboard.

Service Analyzer Dashboard

Service Analyzer helps you map dependencies based on a connection between devices and applications in a tile or topology view.

Visually correlate services to underlying infrastructure with a tile or tree view. Drill down to code level and identify root causes directly from service monitoring dashboards.

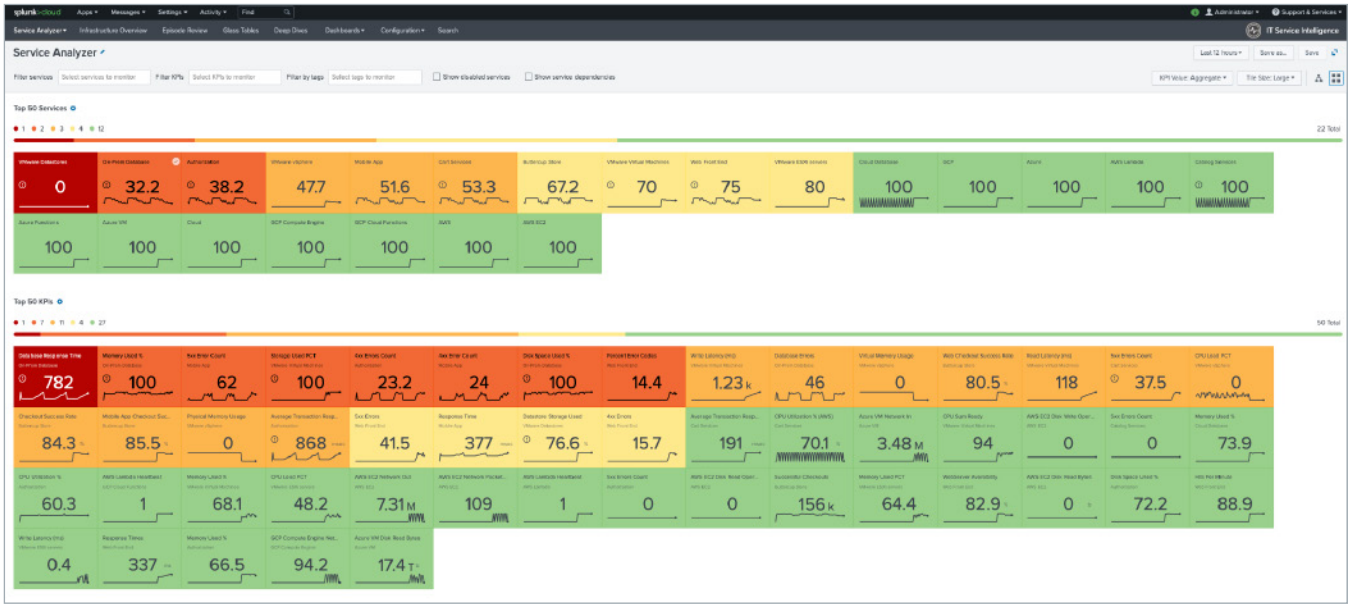

Figure 2: Tile View.

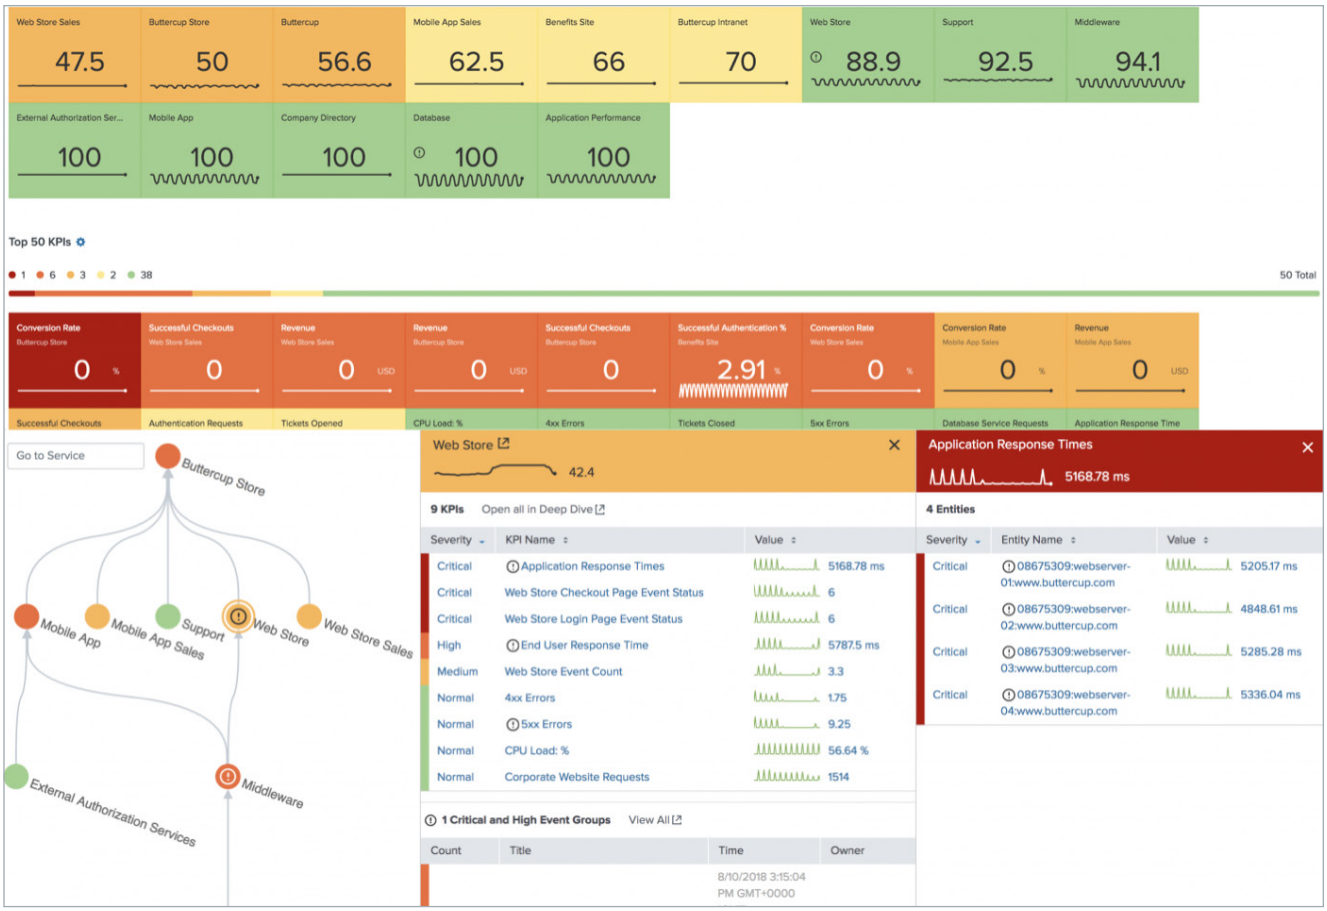

Figure 3: Tree View.

Deep Dives Dashboard

Deep dives are an investigative tool to help you identify and analyze issues in your IT environment.

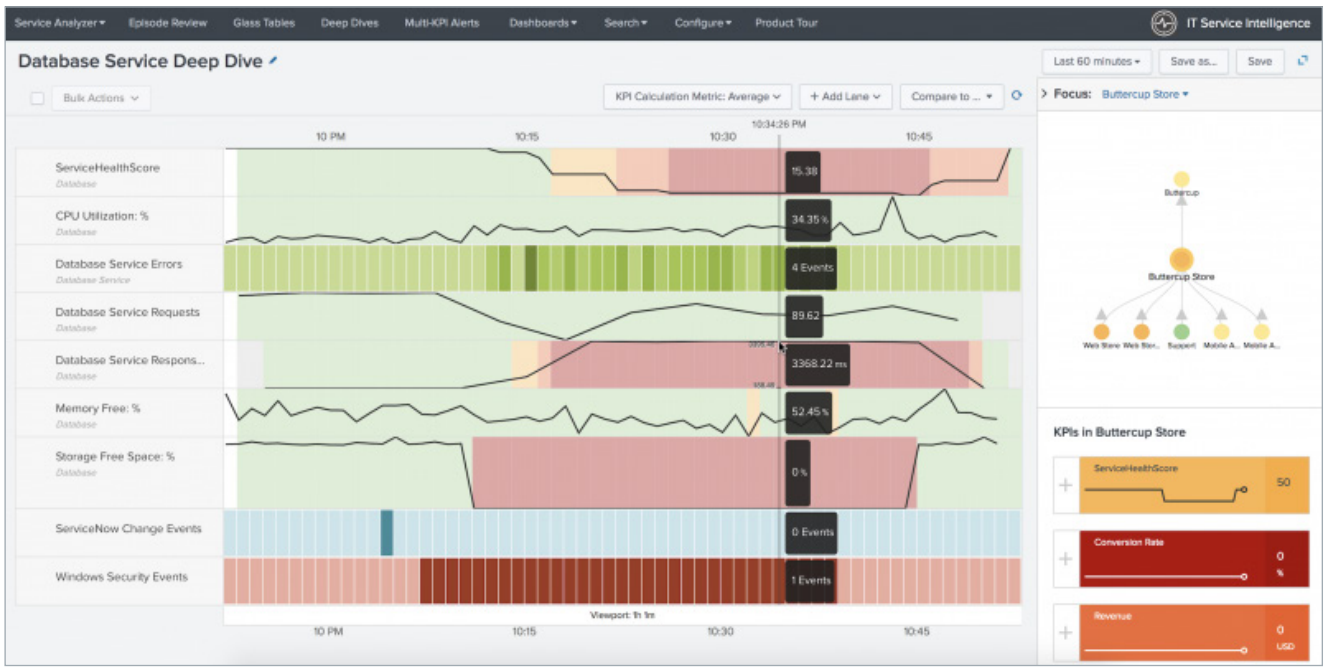

Figure 4: Deep dives display a side-by-side view of KPIs and service health scores over time to help you zoom in on metric and log data and visually correlate root cause.

Figure 5: Use side-by-side displays of multiple KPIs and correlate metrics over time to identify root causes.

For information, see Overview of deep dives in ITSI.

Predictive Analytics Dashboard

Predict future incidents 30 minutes in advance using machine learning algorithms and historical service health scores.

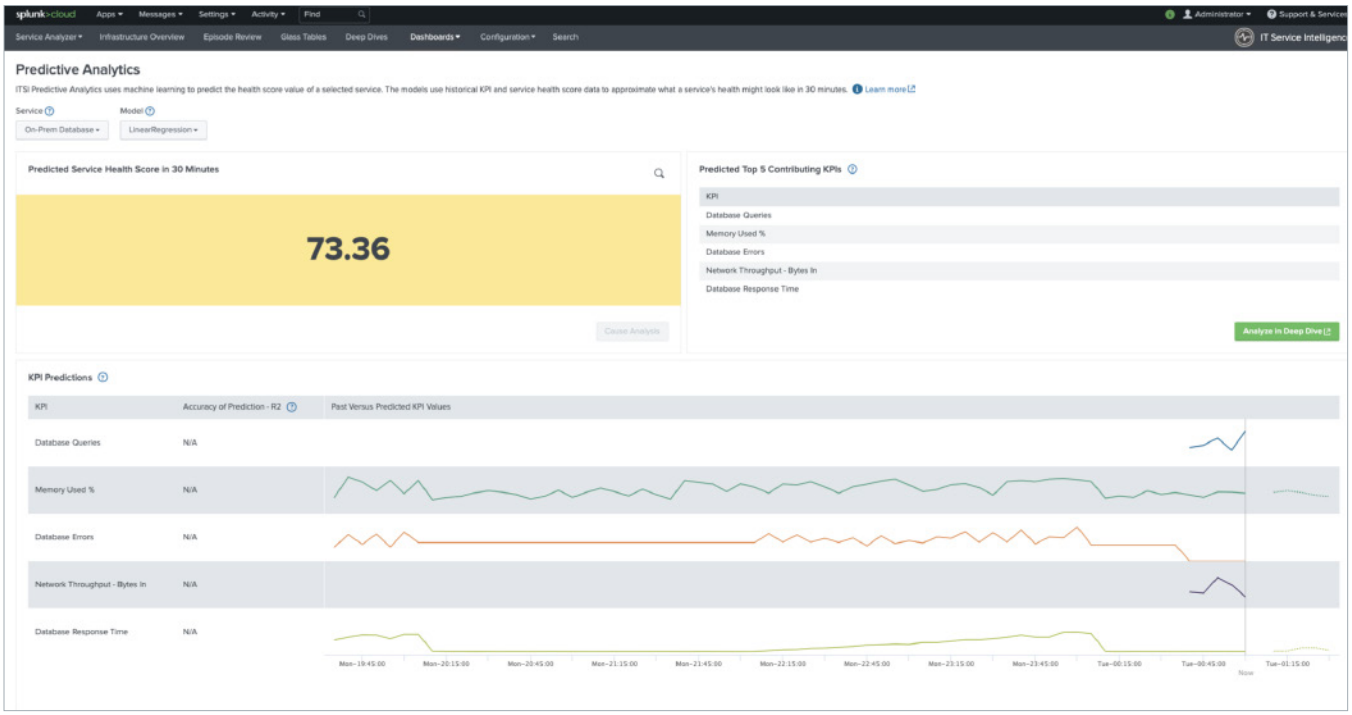

Figure 6: Top five contributing service metrics are displayed to guide troubleshooting.

Glass Tables

Glass tables are custom visualizations that help you monitor real-time interrelationships and dependencies via KPIs and service health scores across your IT and business services in one view.

Create glass tables to provide dynamic contextual views of your IT topology or business processes and monitor them in real time. Glass tables also feature a drawing canvas where you can add visualizations in the form of KPIs and service health scores, upload images and icons, and add charts.

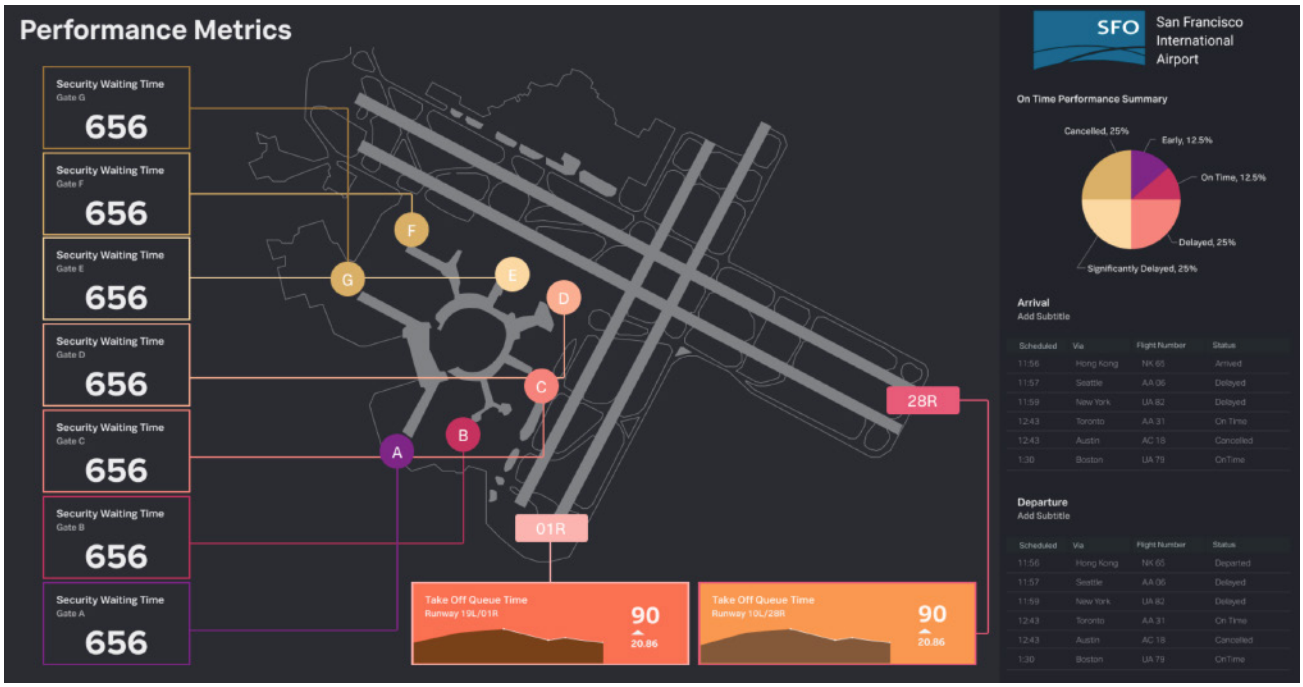

Figure 7: An example of a glass table used to evaluate performance metrics at San Francisco International Airport.

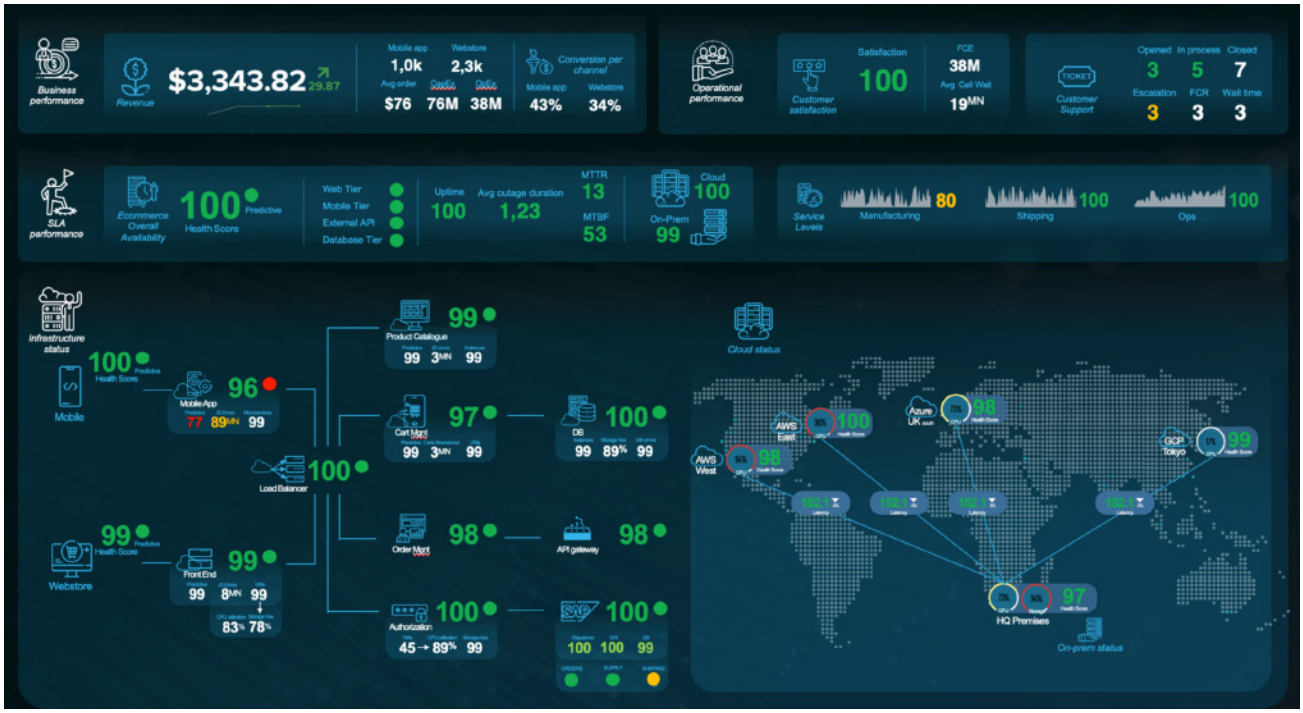

Figure 8: Another example of a glass table used to evaluate business, operational and SLA performances, along with infrastructure status.

For more information, see Overview of the glass table editor in ITSI.

Advanced Analytics & Alerting

Adaptive Thresholds

Adaptive thresholds are a type of threshold configuration for a KPI that's populated based on a statistical distribution of KPI data. Adaptive thresholds should be configured so that alerts are only triggered when behavior strays from normal.

Machine learning algorithms are used to automatically update thresholds based on observed behaviors. This not only determines what should be considered normal in your IT environment but also prevents alerts from becoming stale. The thresholds automatically recalculate on a nightly basis to ensure that changes in behavior don't trigger false alerts.

To learn more about how to use adaptive thresholds, see Apply adaptive thresholds to a KPI in ITSI.

Anomaly Detection

Anomaly detection generates notable events when a KPI deviates from an expected pattern. These notable events represent detected abnormal behavior 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.

See Apply anomaly detection to a KPI in ITSI.

Service/KPI alerts

Enable alerting on a single KPI in ITSI so you can be alerted when aggregate KPI threshold values change. ITSI generates notable events in Episode Review based on the alerting rules you configure.

Use these alerts to investigate and take action on the severity changes of your individual KPIs before they negatively impact the service as a whole.

To do so, make sure you:

- Have written access to the service in order to enable KPI alerting.
- Create a KPI within a service and configure thresholds for it before you can enable alerting. For more information, see Step 7: Set Thresholds in the KPI configuration workflow.

To learn more about how to set up KPI alerts, see Receive alerts when KPI severity changes in ITSI.

Multi-KPI Alerts

Trigger alerts based on multiple service conditions. Define severity levels and trigger conditions, or assign weights to attribute relative importance.

Create a multi-KPI alert from a deep dive view when you see a correlation between two or more KPIs, and get notified next time a similar problem occurs.

To learn more about how to set up multi-KPI alerts, see Create multi-KPI alerts in ITSI

Event Analytics Terms

Intelligent Incident Management

Integrated event management and incident response. You can create a ticket, run a script or contact a team from one dashboard — resolving an incident without switching to another tool.

Episode Review

Incident management dashboard that prioritizes issues by severity and allows teams to trigger a response from the same view. Reduce "swivel-chair" operations and respond to the most urgent issue first.

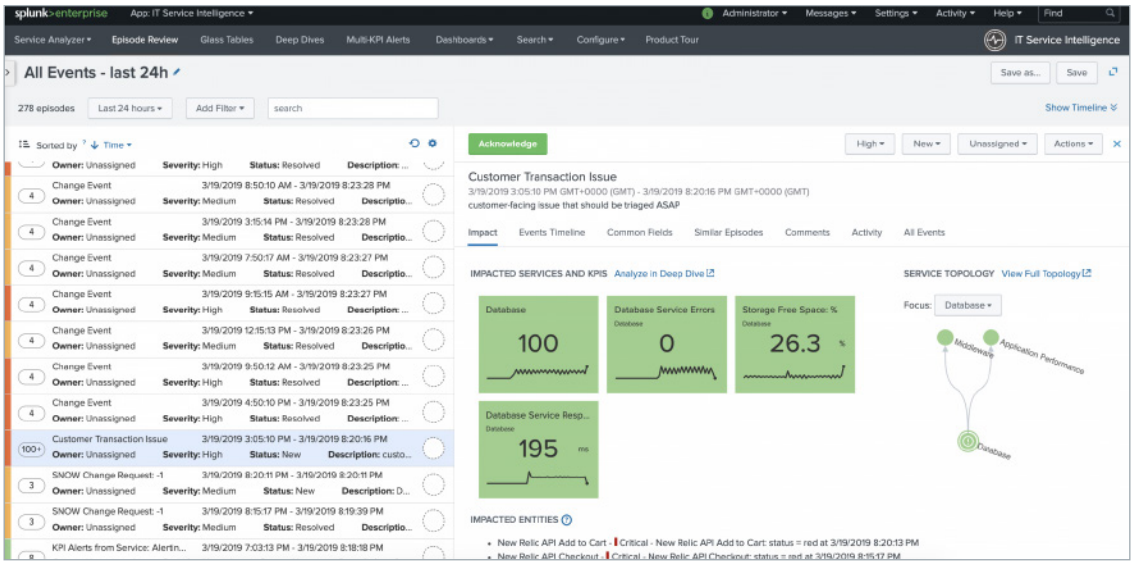

Figure 9: Snapshot of an episode review dashboard.

Mean Time to Detect (MTTD)

MTTD is a key performance indicator for IT incident management and measures how long an issue exists before the appropriate parties are aware of it.

Mean Time to Resolve (MTTR)

MTTR is the average time it takes to resolve an anomaly. This time includes detecting the anomaly (MTTD), diagnosing the problem, repairing the issue and ensuring that this anomaly won't happen again. By using ITSI's intelligent incident management capability, you can accelerate MTTR.

Rules Engine

The ITSI Rules Engine is a system for continuously processing notable events to allow for event grouping and deduplication, as well as automatic action execution, based on user-defined criteria. The system revolves around a continuously running indexed real-time search that streams all notable events into a custom search command.

The Rules Engine's functionality begins with correlation searches. Correlation searches generate notable events in ITSI, which are stored in the `itsi_tracked_alerts` index. The Rules Engine saved search accepts notable events into the `itsi_rules_engine` custom search command. The search command generates the internal structures required for aggregating events into episodes and executing actions.

The Rules Engine search periodically polls the configuration database for updates. If a policy indicates some action should be executed, the Rules Engine dispatches a REST request to the Event Management Interface to execute the action.

To see a diagram of the Rules Engine workflow or to learn more about the searches, see Overview of the ITSI Rules Engine.

Notable Events

Notable events represent detected abnormal behavior for service-level (trending) and entity-level (cohesive) KPI data. Notable events form the basis for Episodes which are an intelligently grouped set of related notable events.

Aggregation Policies

A notable event aggregation policy is the fundamental unit of event grouping in IT Service Intelligence (ITSI). Aggregation policies are the data structure the Rules Engine uses to group notable events into deduplicated episodes and organize them in Episode Review. These episodes have their own title, description, severity, status and assignee that are separate from the individual notable events within the episode. Aggregation policies are also the container for action rules that automate episode actions, such as sending an email or pinging a host.

To learn more about the 3 components of aggregation policies: 1) filtering, splitting, and breaking criteria, 2) episode information, and 3) action rules, see Overview of Aggregation Policies in ITSI.

Intelligent Event Correlation & Aggregation

Collect and enrich events from multiple sources into a single alerting framework. Real-time, automated event correlation triggers alerts as data enters the system, using out-of-the-box (OOTB) machine learning policies for immediate noise reduction.

Additional Resources

- On Demand Services Catalog
- Learning Path: ITSI Admins
- Learning Path: ITSI End Users
- Course: Implementing ITSI
- Course: Using ITSI
- ITSI Community
- .conf Sessions
- Events and Workshops
- Webinars

Use the resources and tools outlined in this Getting Started Guide to explore ITSI and all of its capabilities!
Please reach out to your account team if you have any questions or concerns.

21-20243-Splunk-Getting Started with ITSI-101-GSG