Welcome we will get started shortly…

Splunk ITSI Workshop

Presenter:
Paul Winchester

16th June 2020

Note: This workshop is based on Splunk's own Splunk4Ninjas – ITSI sessions.

Also note: This is not an actual picture of me. But it's quite close...

Register and Create Your Environment

1. If you've not yet created your Splunk environment, use this link to register:

http://splunk4rookies.com/4320/self_register

2. Log into you instance:

http://<aws_instance>:8000/en-US/

username : admin

password : smartway

Please note that the instances take 30mins to initialise due to Machine Learning backfilling.

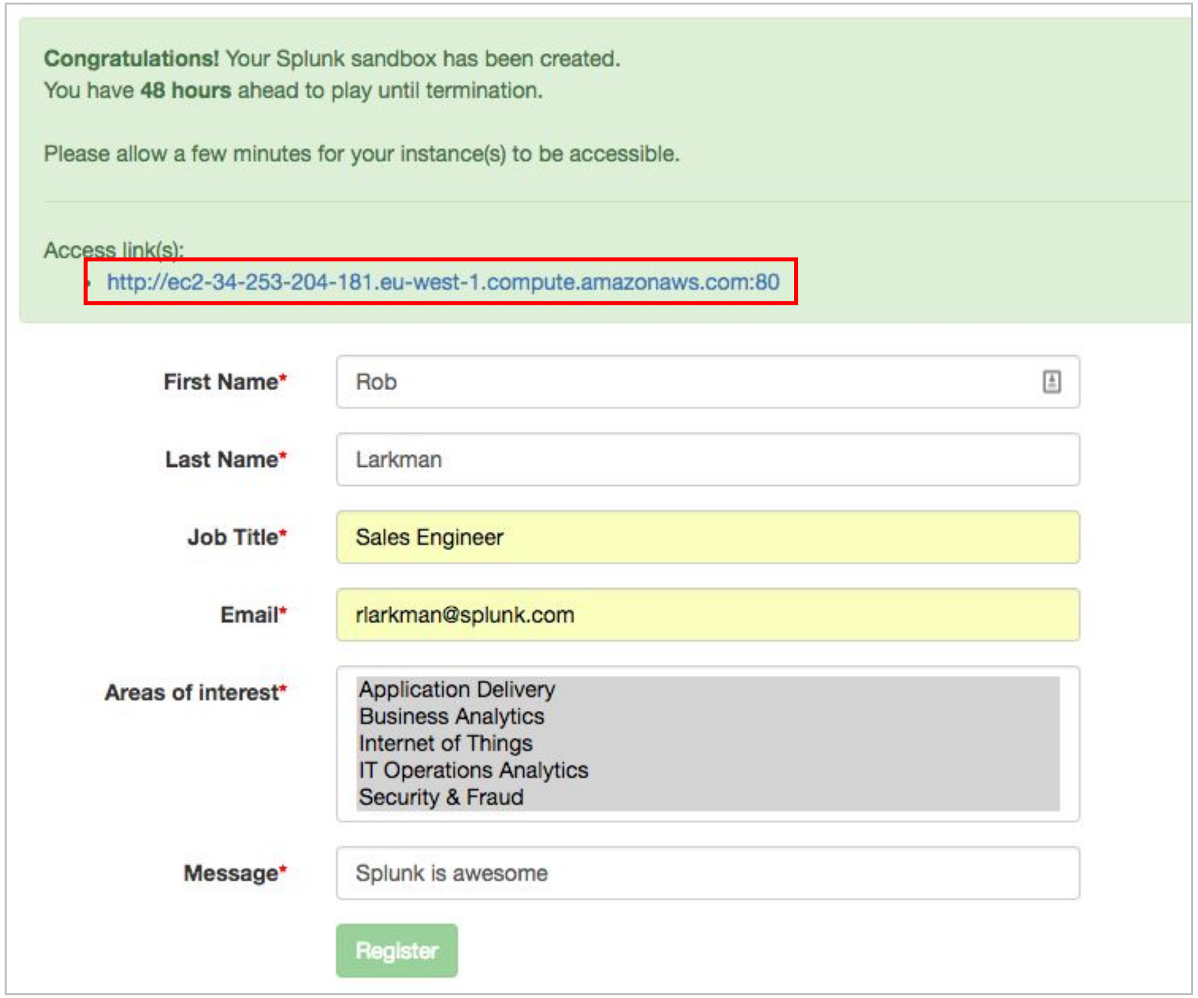

Agenda

- IT Service Intelligence Overview
- Create a Service
- Creation of two KPI’s
- Machine Learning
  - Apply Adaptive Thresholds
  - Predict Analytics
- Deep Dives
- Glass Tables
- Event Analytics Demystified
- Somerford’s approach to ITSI

IT Service Intelligence

Complexity obscures the data you need

IT Ops teams continue to struggle to monitor, investigate, analyze & act

DEV/APPS

CLOUD

OFFICE

DR

REMOTE

SECURITY

STORAGE

NETWORK

SERVERS

BUILDING

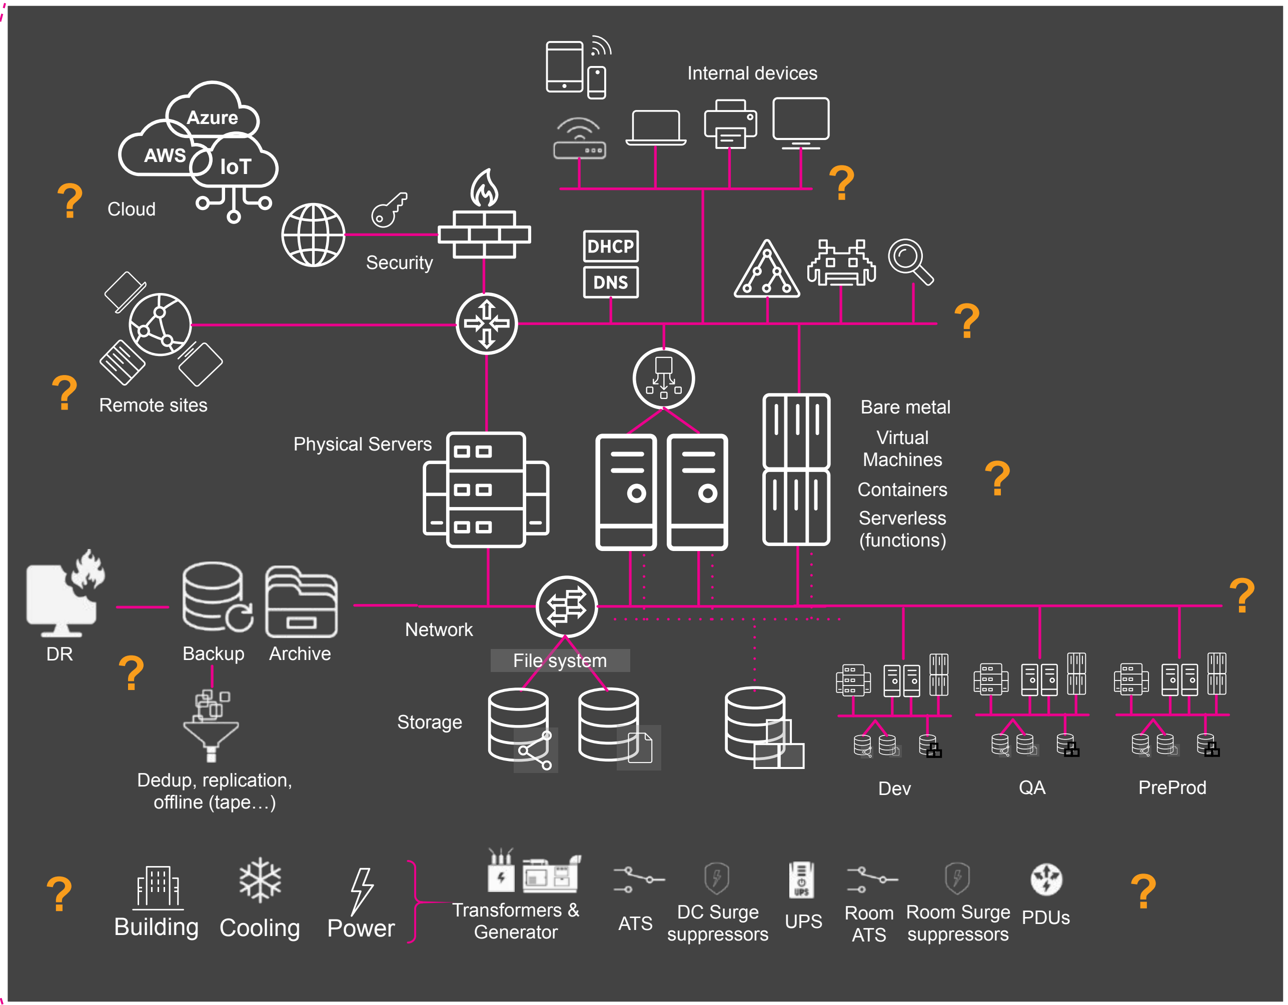

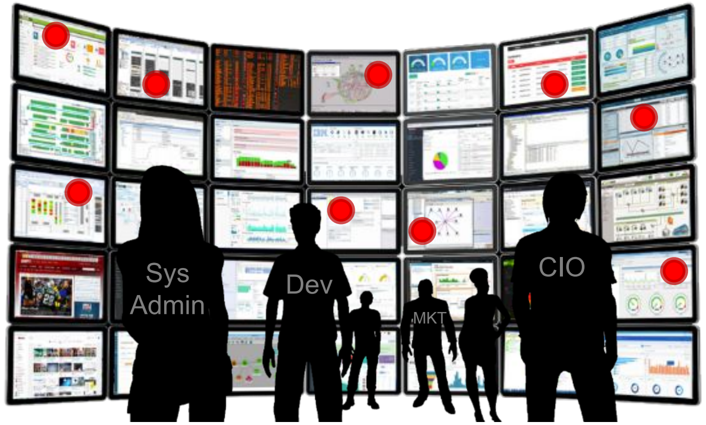

Too Many Tools

Siloed Views of Data

Too Much Event Noise

Long Times to Resolve

The impact can be significant

Putting revenue, customer experience, employee effectiveness & innovation at risk

Lost Revenue

Outages and incidents impact the services and apps driving revenues

Poor Customer Experiences

Customers click away and brand reputation is damaged

Decreased Employee Effectiveness

Teams thrash, finger-point, and key employees leave

Struggle to Innovate

IT spends too much time fixing problems instead of innovating and transforming

Why The Old Ways Disappoint

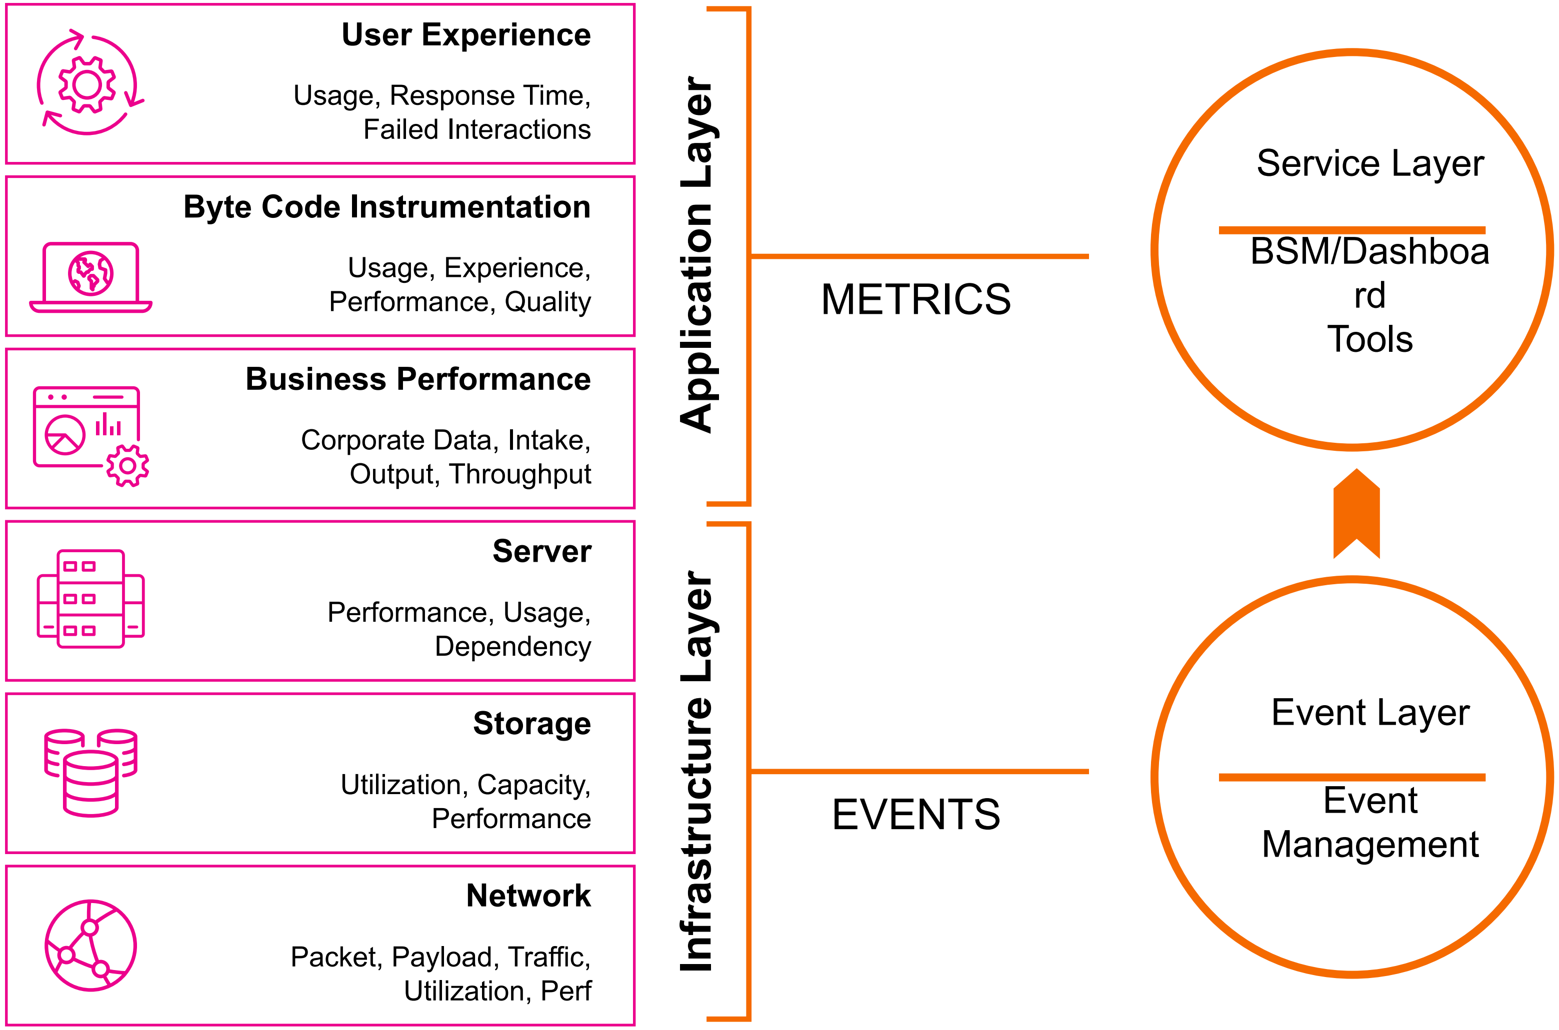

Challenges

- Many disparate components
- Brittle integrations
- Data is summarized and lost
- Longer root-cause identification
- End-to-end view challenging
- Labor-intensive to manage
- Not agile for digital business

IT Service Intelligence Platform Approach

Application Layer

User Experience
Usage, Response Time, Failed Interactions

Byte Code Instrumentation
Usage, Experience, Performance, Quality

Business Performance
Corporate Data, Intake, Output, Throughput

Infrastructure Layer

Server
Performance, Usage, Dependency

Storage
Utilization, Capacity, Performance

Network
Packet, Payload, Traffic, Utilization, Perf

MACHINE DATA

Splunk Approach:

- Single repository for ALL data
- Data in original raw format
- Machine learning
- Simplified architecture
- Fewer resources to manage
- Collaborative approach

Splunk IT Service Intelligence

Data-driven service monitoring and analytics

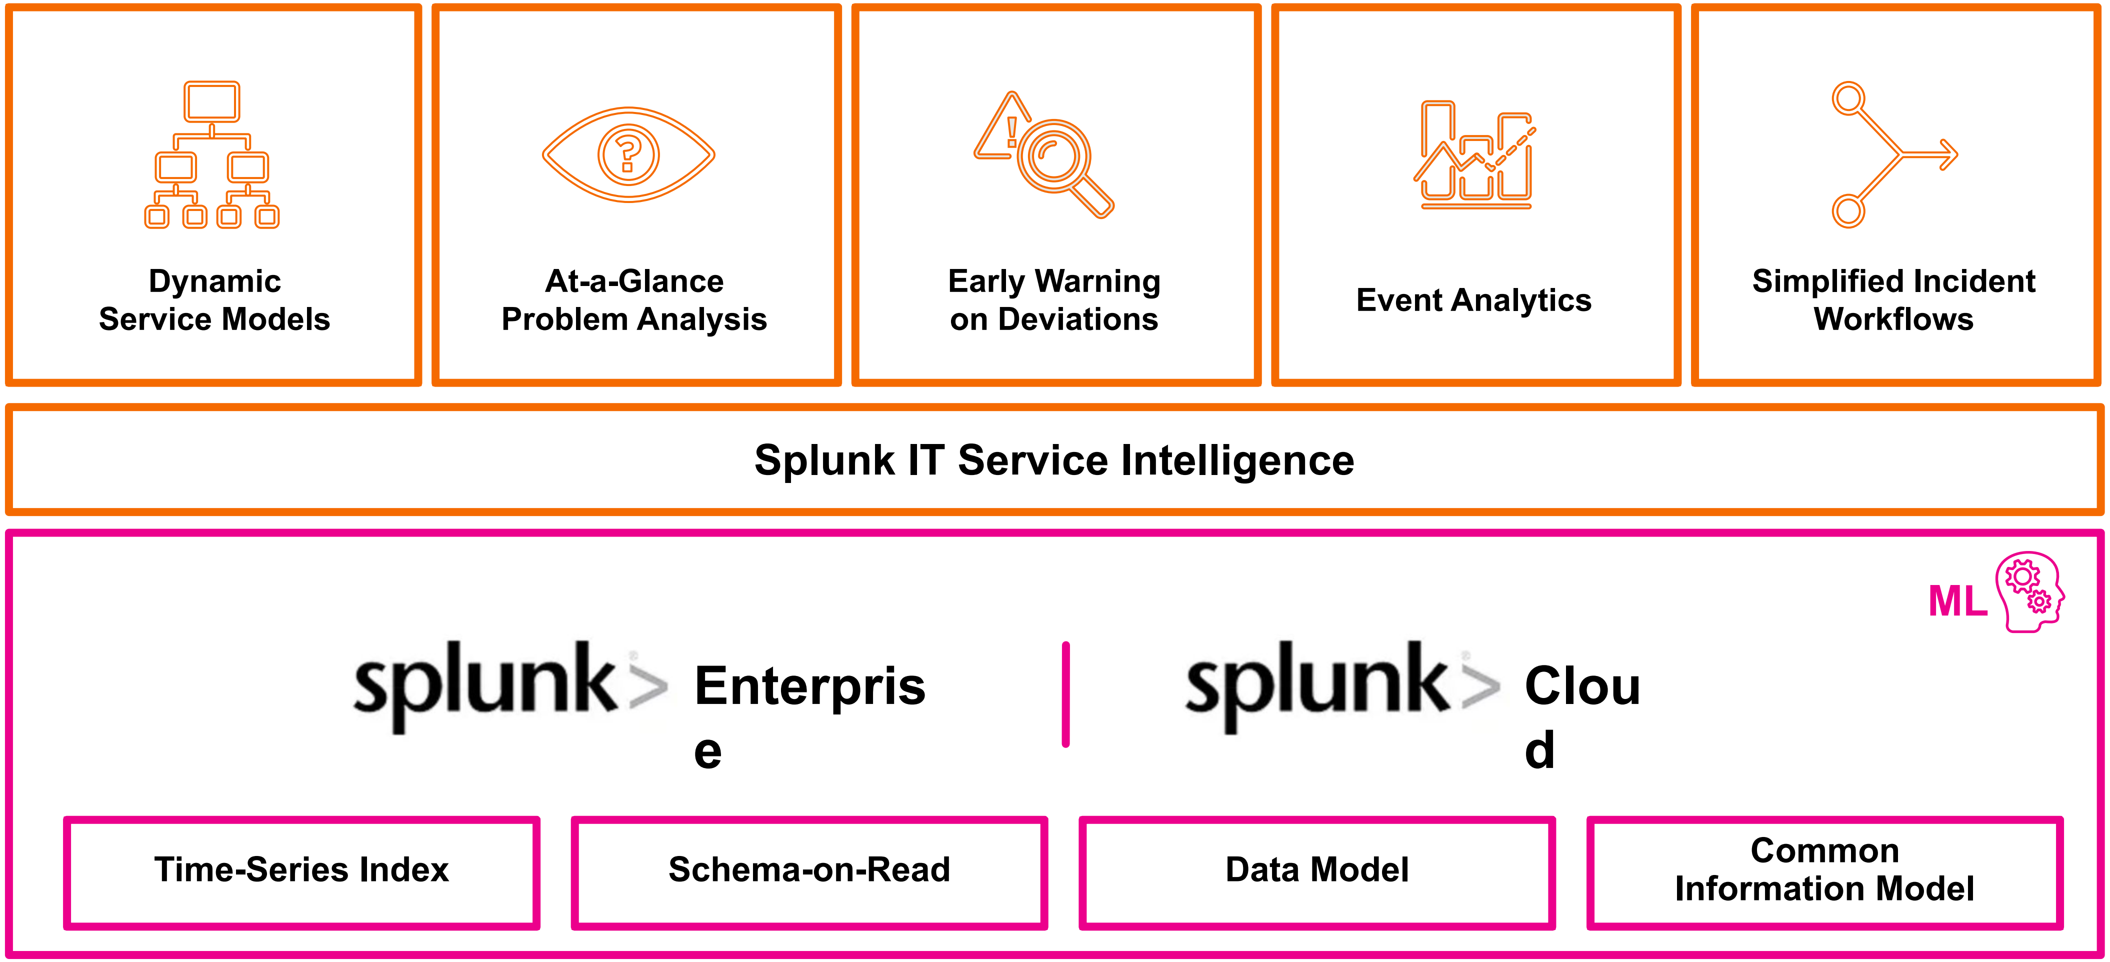

Key Terminology

What is a Service

A logical group of technology components that a user deems important to be monitored together. Services can encompass multiple tiers of the IT domain.

Online Store Service

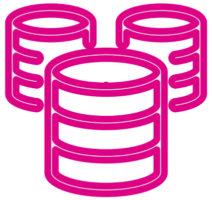

Database Service

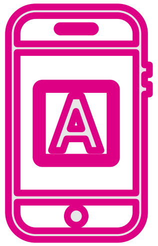

Mobile App Service

Call Centre Service

splunk> turn data into doing

Service Dependencies

Business Services

Application Services

Technical Services

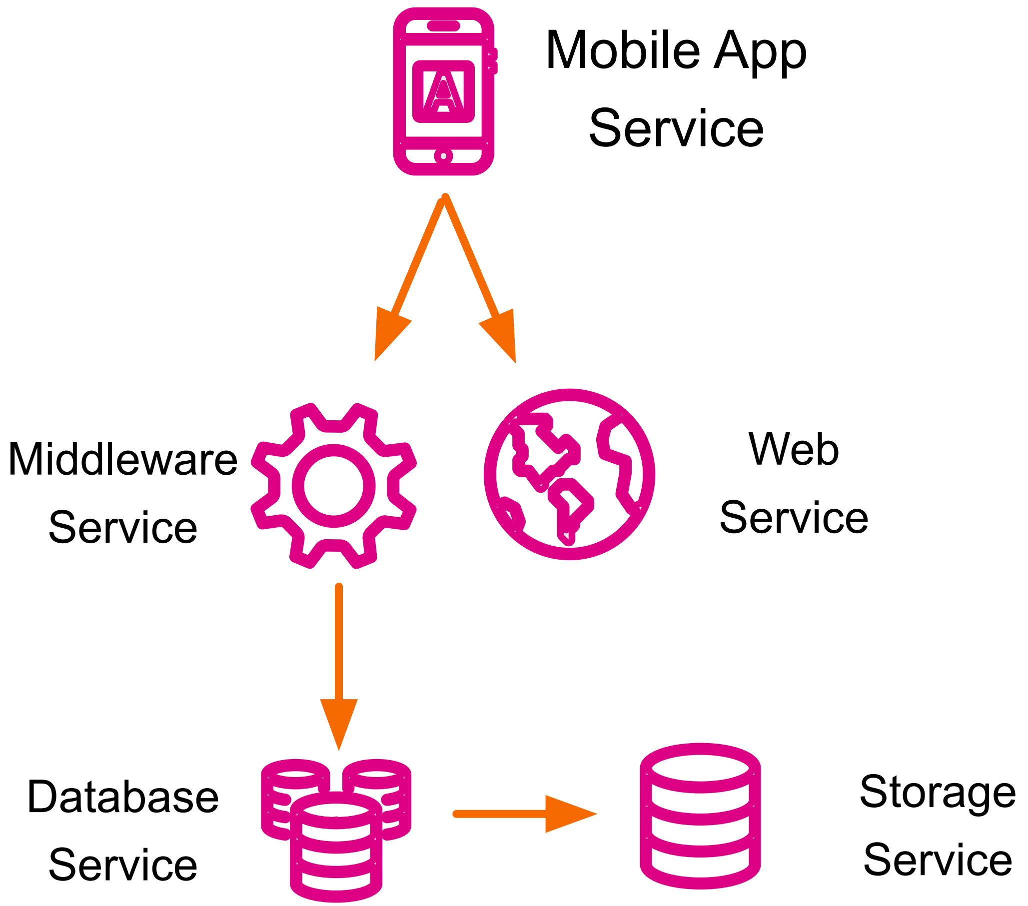

What is a KPI

A Key Performance Indicator and Health Scores constitute the means by which Services are monitored. A KPI is a Splunk saved search in ITSI that helps monitor a specific value.

Online Store Service

- Number of visitors
- Revenue
- Transactions per day

Database Service

- CPU Utilisation
- Query Response Time
- Storage Free Space

What is a Health Score

A Health score is a score from 0-100 (0 being critical and 100 being normal) that measures the health of a Service. It is calculated based on all KPIs importance and its status.

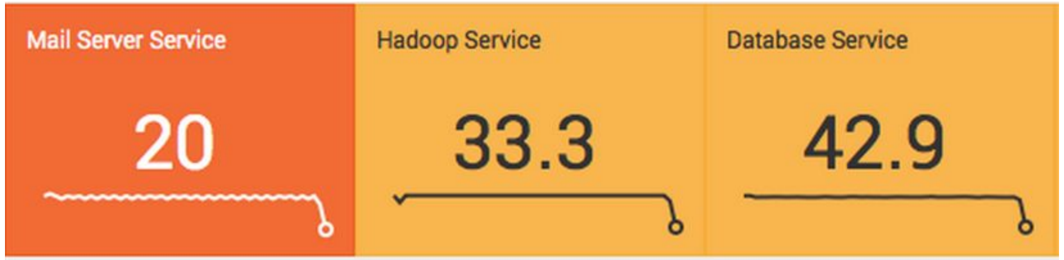

splunk> turn data into doing

What is an Entity

An **Entity** is an optional sub-element of a KPI.
A KPI can be filtered by entities.
ITSI can import entities from CMDBs & other sources

Severity	Entity Name	Value
Critical	mysql-02	0 %
Normal	mysql-01	33.26 %
Normal	mysql-03	30.7 %
Normal	mysql-04	28.05 %

ITSI Service Definition

To summarize a service is comprised of:

A Service

Health Score

0 or More

Entities

0 or More

KPIs

0 or More

Dependencies

Workshop Back Story

Workshop Back Story

- Buttercup Store has just deployed a new Web service, ***however*** the engineers forgot to include database monitoring.
- Splunk administrators have deployed collectors on the infrastructure and have requested that we build a service centric monitoring solution.
- The CxO has just also requested that we include some 'AI' as they read on a website Artificial Intelligence can solve everything!
- The lab exercises start on the back of a service decomposition workshop to identify all the components that make up this service.

Service Decomposition Outcome

Example diagram detailing the content of a service, typically this is a photo of a whiteboard

Business Layer

Visitors
Conversion Rate
Product Cost

Buttercup Store

Revenue
Revenue per Order
Successful Checkouts

Application Layer

Manufacturing

Availability
Cycle Time
Performance

Quality
Total Units
Units Produced

Web Front End

4xx Errors
5xx Errors
Bytes In
Bytes Out

Website Requests
CPU Utilisation%
Disk I/O
Response Time

Mobile App

4xx Error Count
5xx Error Count
Response Time

Process Layer

Product Catalog

5xx Errors Count
Active Session
Application Errors
Response Time

Order Management

Active Session
Application Health
Response Time

Cart Management

Active Session
Application Apdex Score
Application Health
Response Time

Authorization

Application Response
Transaction Time
Response Time
End User Response

Infrastructure Layer

On-Prem Database

Database response time
CPU Utilisation

Cloud Database

CPU Utilisation
Database Response
Memory Used
Database Queries

External Authentication

Authentication Errors
Response Time
Requests

Missing Service & KPIs

Let’s get stuck in to it!

Point your browser to : http://<aws_instance>:8000/en-US/

This will be the link you got when you registered.

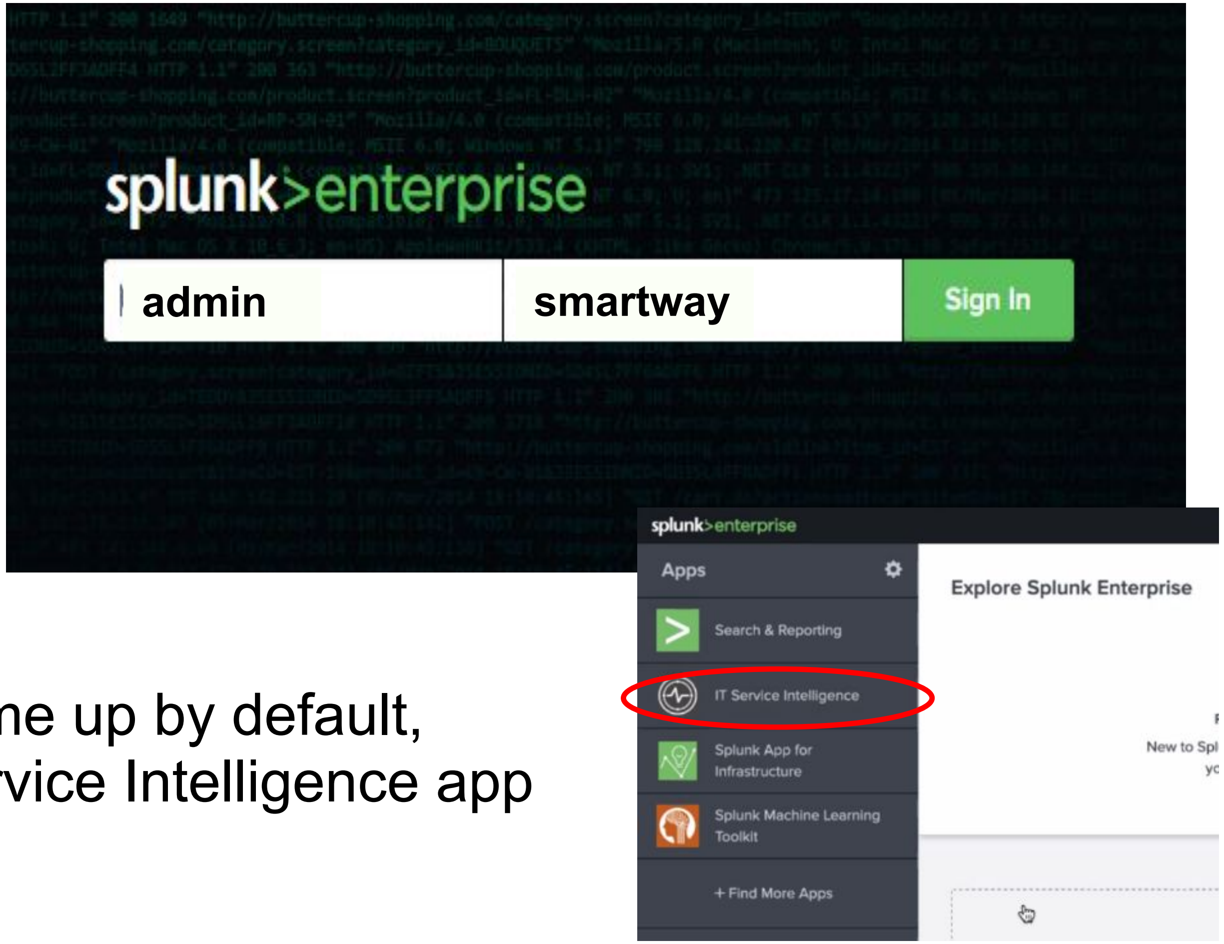

If it does not come up by default, select the IT Service Intelligence app

Services Exercise

Service Analyzer

The Service Analyzer view provides a visual representation of our services and the KPIs associated KPIs. The highest severity is brought to the top left of each group.

- Select the ‘Default Service Analyzer’ view.

- Click on the  icon to switch the visualization tree view

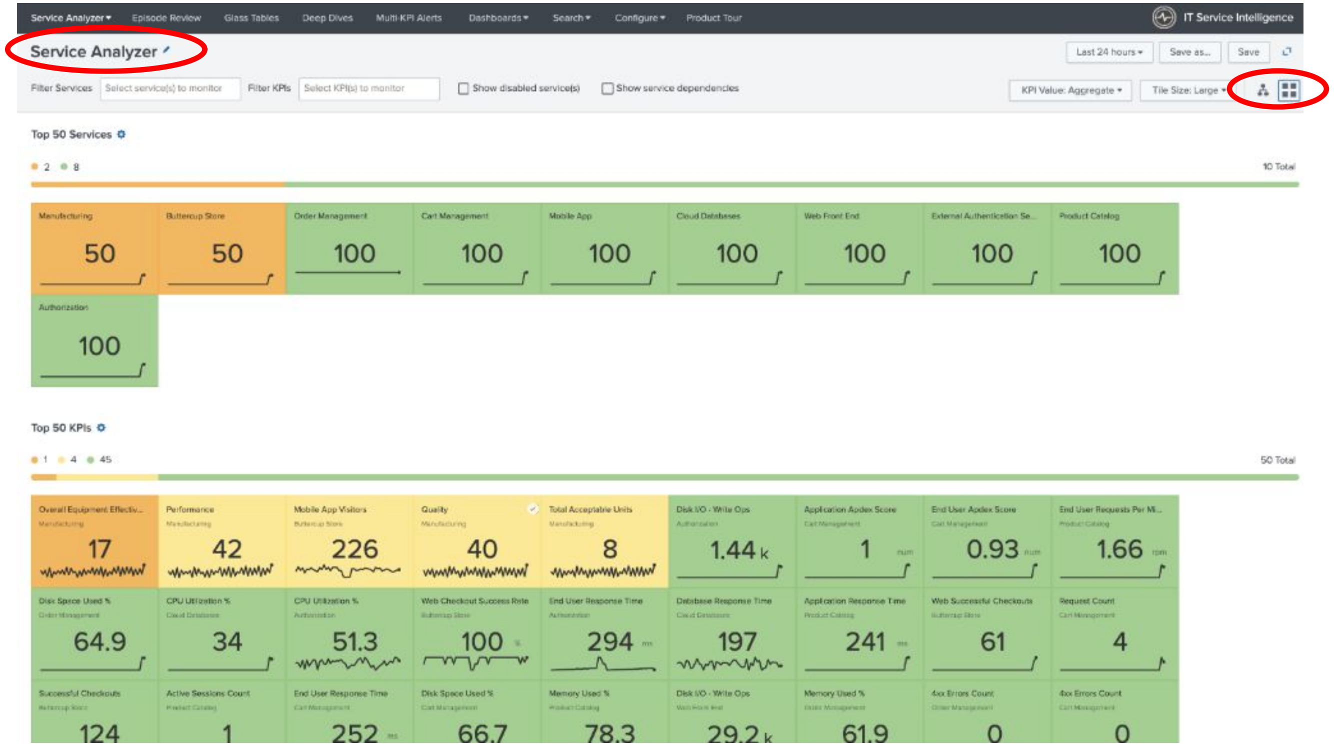

Service Tree

- The Service Analyzer tree view provides a visual representation of our services and the dependencies between them. Using this view you will also see the KPIs, entities, and most critical notable events associated with a service.

- The health of a service is affected by the health of a child service

- The tree can be built manually, however typically this is imported from a CMDB or a search

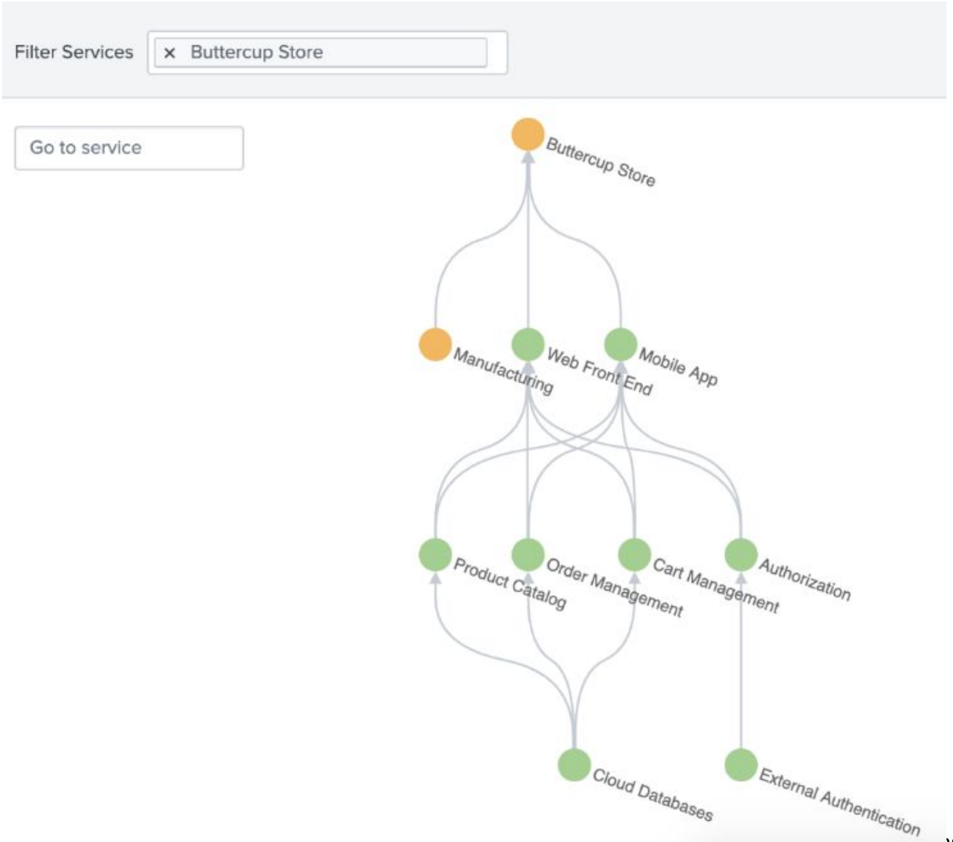

Service Tree

We can see that there is an issue with the Manufacturing service.

- Click on the effected service (***Manufacturing***) to investigate which KPIs have degraded.
- Review which entities are effected.

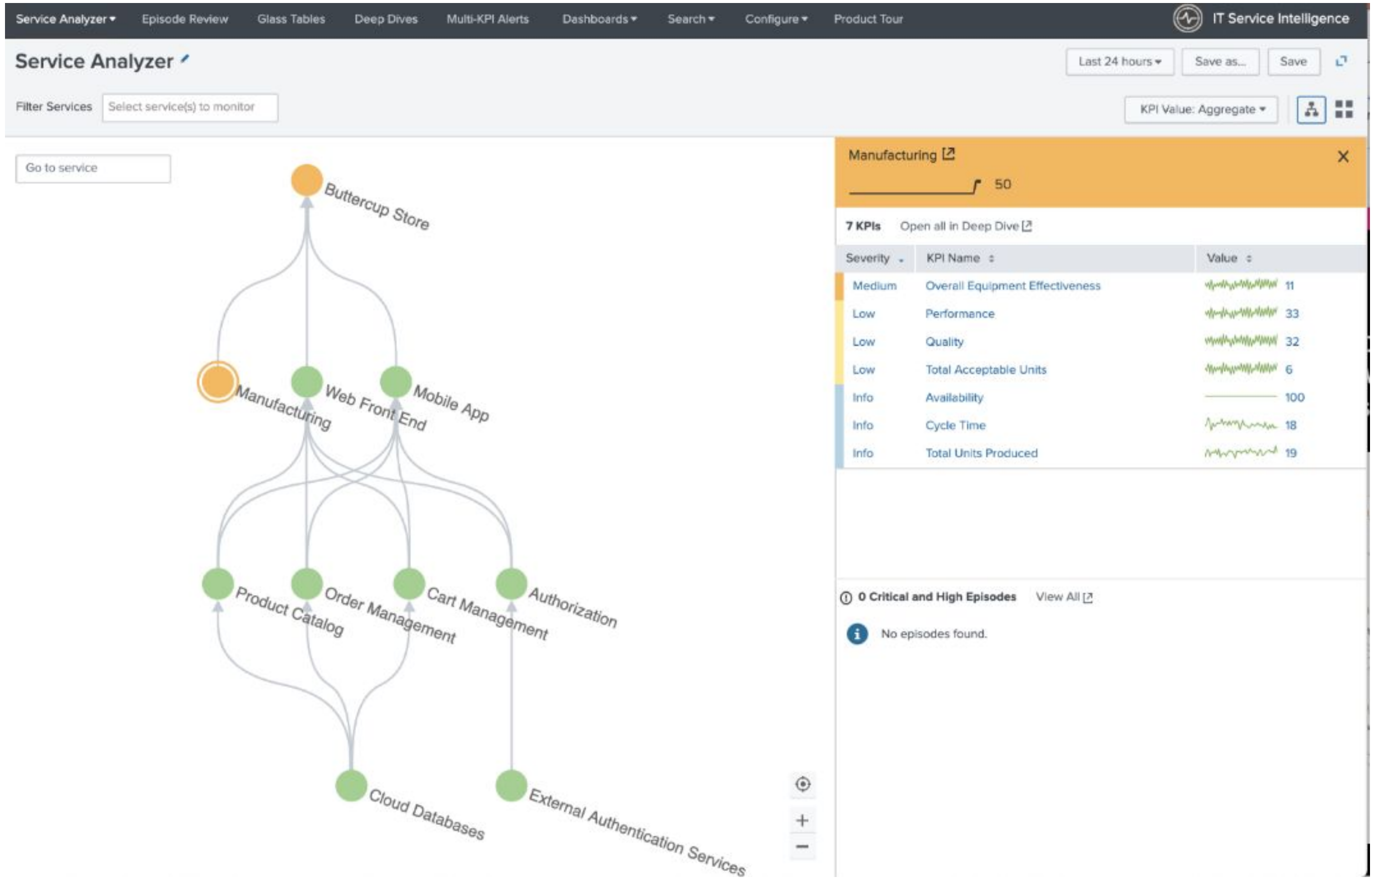

Services Lab

During the service decomposition workshop we identified a missing ‘On-Prem Database’ service.   In this lab we are going to create this new service, we will review the ITSI DB module and then select a service template.

- To create a service you can use (Technical) KPI’s that comes with modules:
  - OS: Linux, Unix & Windows
  - Web server: Apache & Microsoft IIS
  - Application server: Tomcat & Websphere
  - Database: Microsoft SQL & Oracle
  - Storage: Netapp ONTAP & EMC VNX
  - Load Balancer: F5 Big IP & Netscaler
  - Virtualisation: VMWare & Hyper=V

Services Lab

- Select the Configure menu + Services

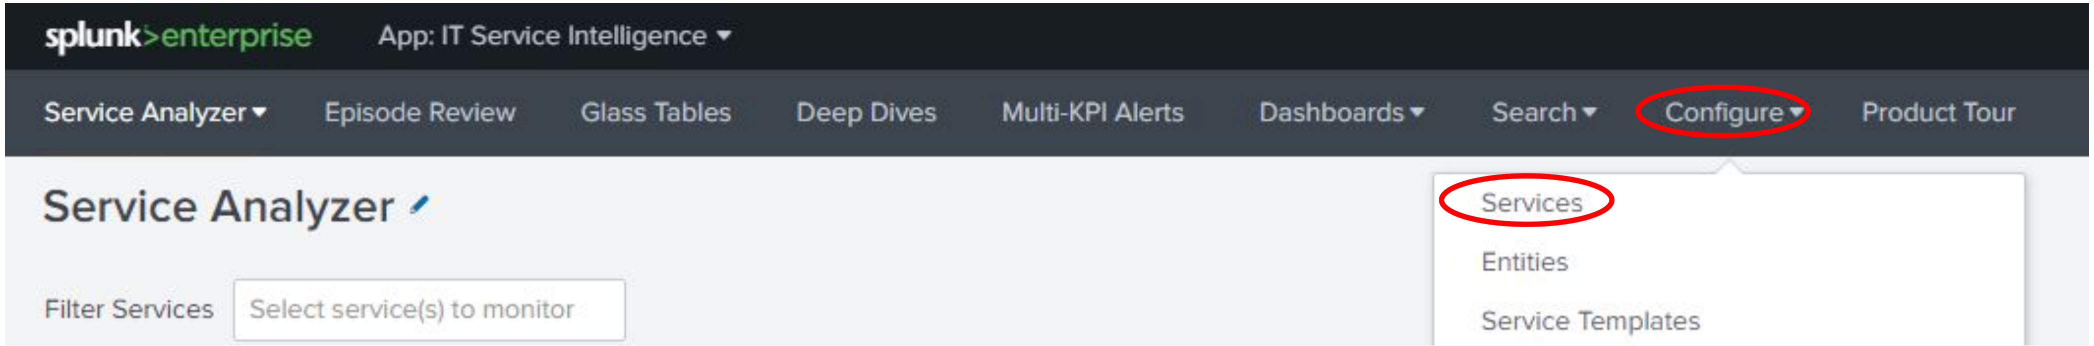

- Click Create Service > Create Service

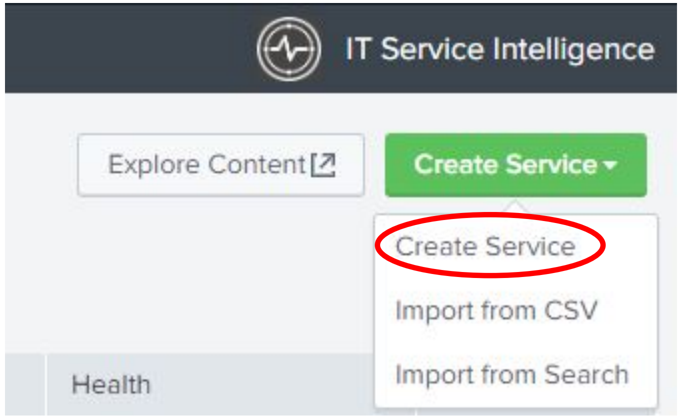

Services Lab

- Name the service as ‘On-Prem Database’
- Check ‘**Database**’ in modules list
  - Review the selected KPIs
  - *Do not hit the ‘Create’ button!*
- ITSI suggests best KPIs for database monitoring

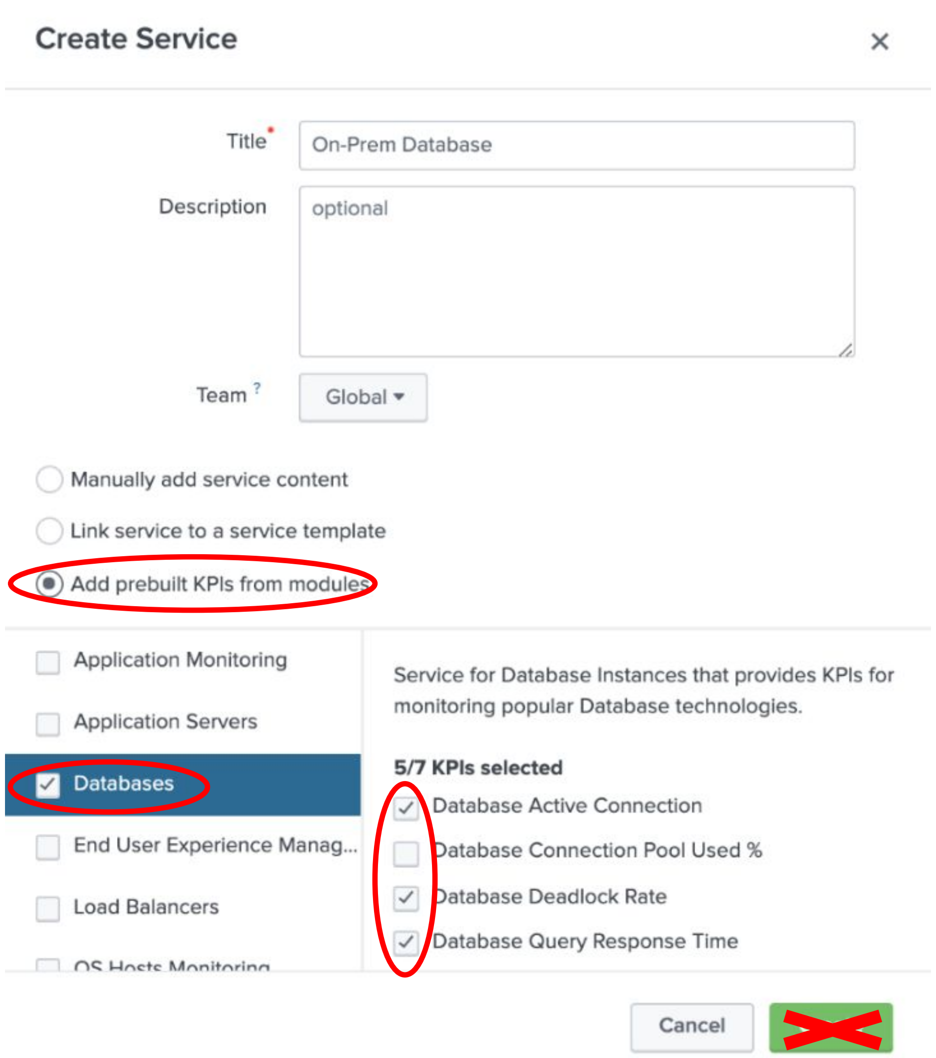

Services Lab

Instead we are going to utilize the ITSI service templates feature, this will build the service with predefined KPIs.

- Select ‘*Link service to a service template*’ button
- Choose ‘***On-Prem Database***’ template
- Click the ‘*Enable 7 days of backfill for all Service KPIs*’ option
  - *Please note the option is hidden below so you will need to scroll down!*
- Click ‘*Create*’ button

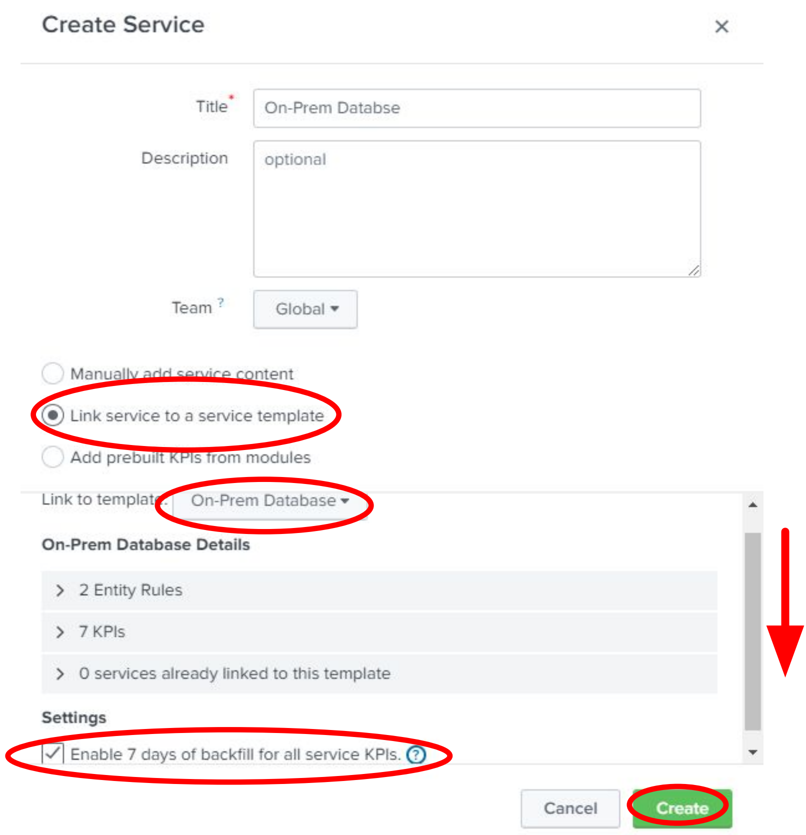

Services Lab

- The new ‘On-Prem Database’ service is based on a template, if you review the ‘Entities’ tab we can see that the entities are already filtered.

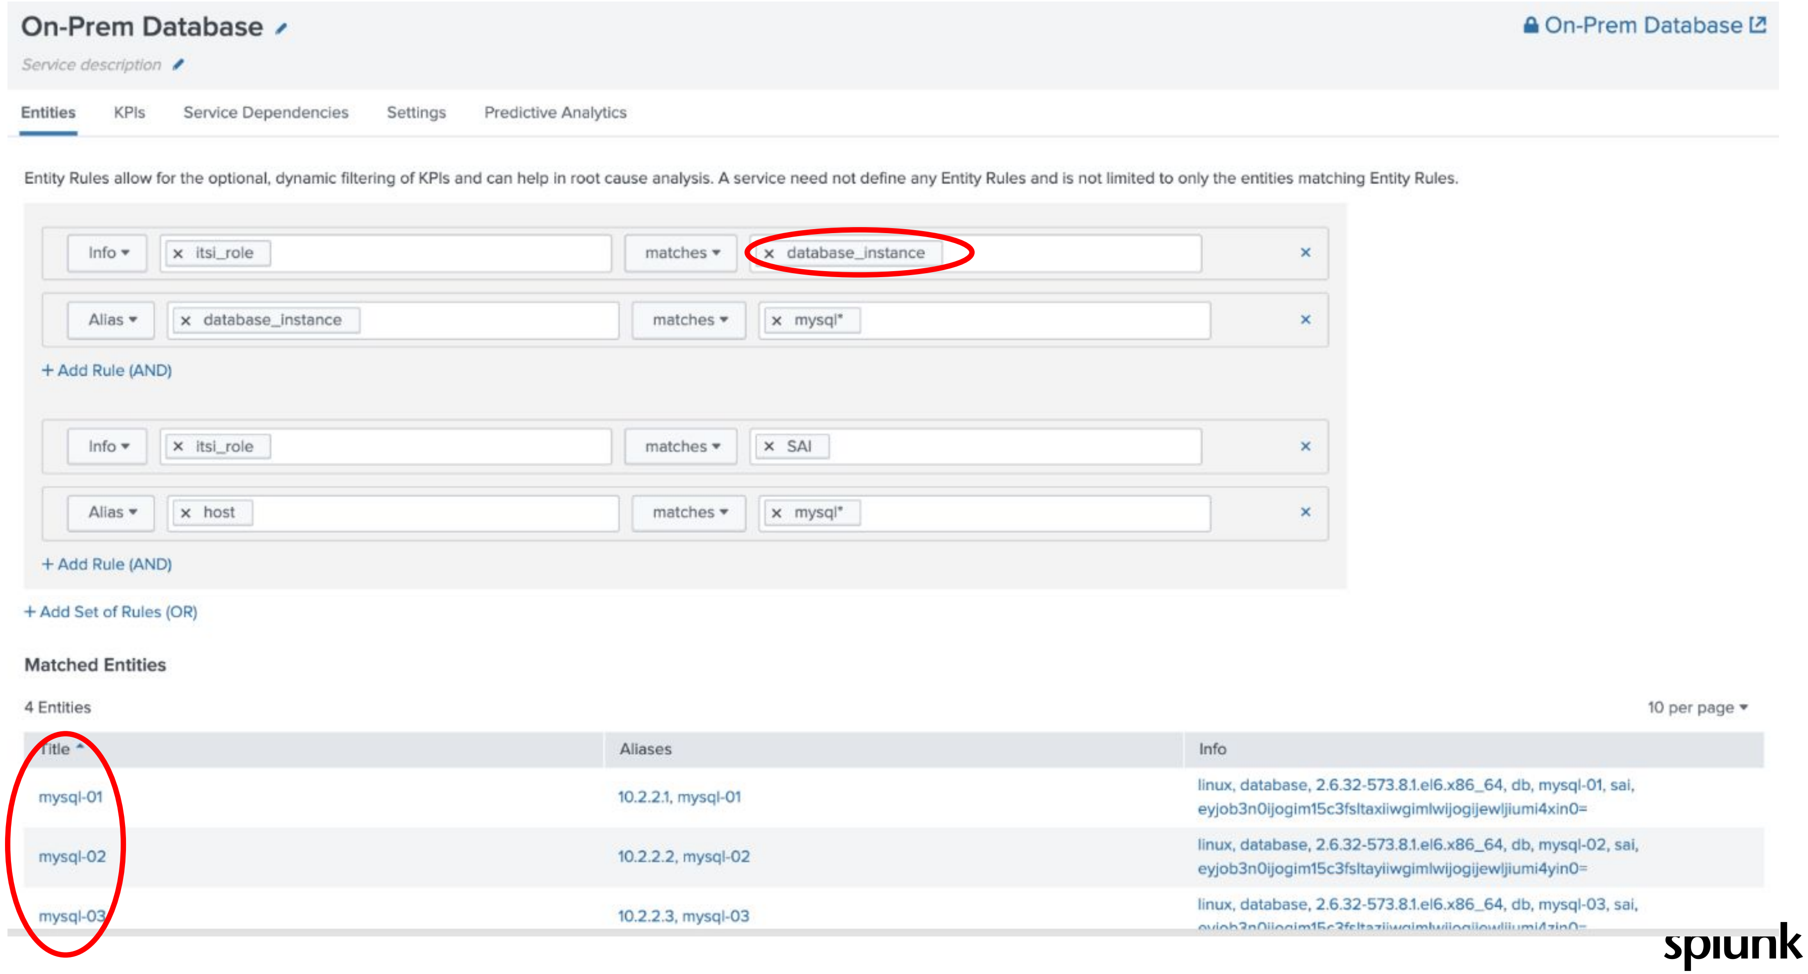

Services Lab

Under the KPIs tab we can see some KPIs that have been inherited from the service template, the padlocks indicate that changes to the template will be pushed to all linked services.

- Click 'Database Queries' KPIs to review.

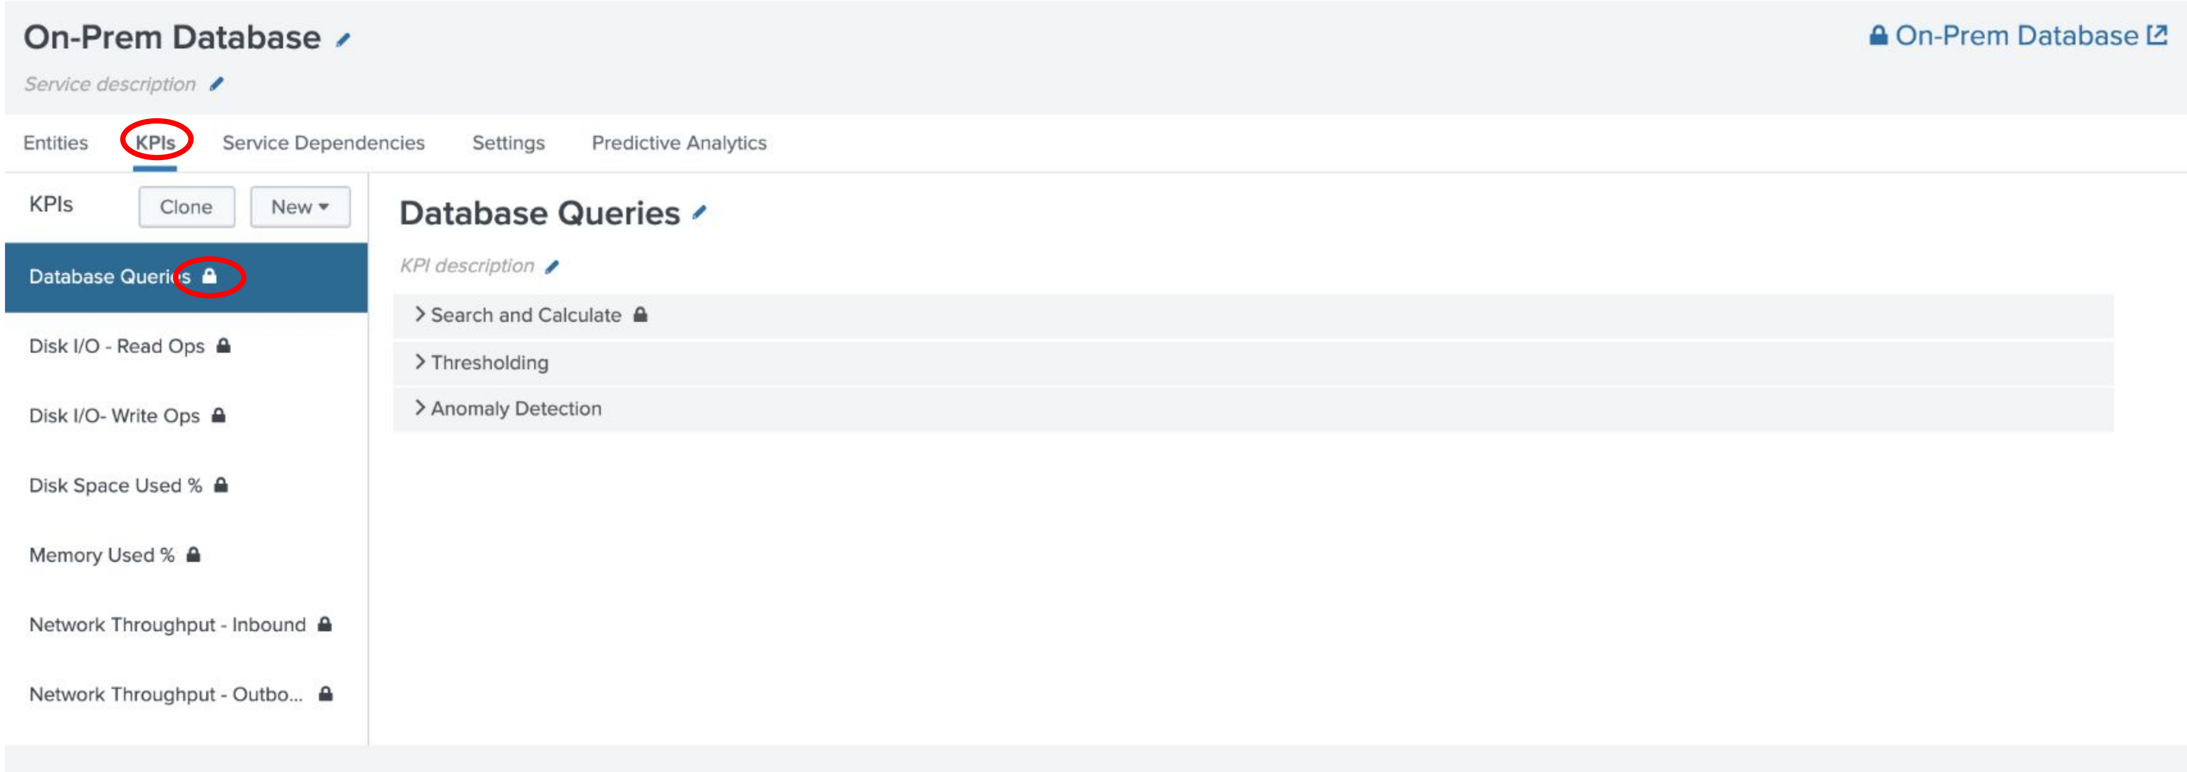

Services Lab

The ‘Settings’ tab enables configuration of the service attributes. The new (linked) database service is disabled by default.

- Switch to Setting tab
- Toggle status to ‘*Enable*’
- Investigate the effect changing the Importance and Simulated Severity has, on the Simulated Health Score
- ***Please do NOT enable Service Health Score backfill at this point***

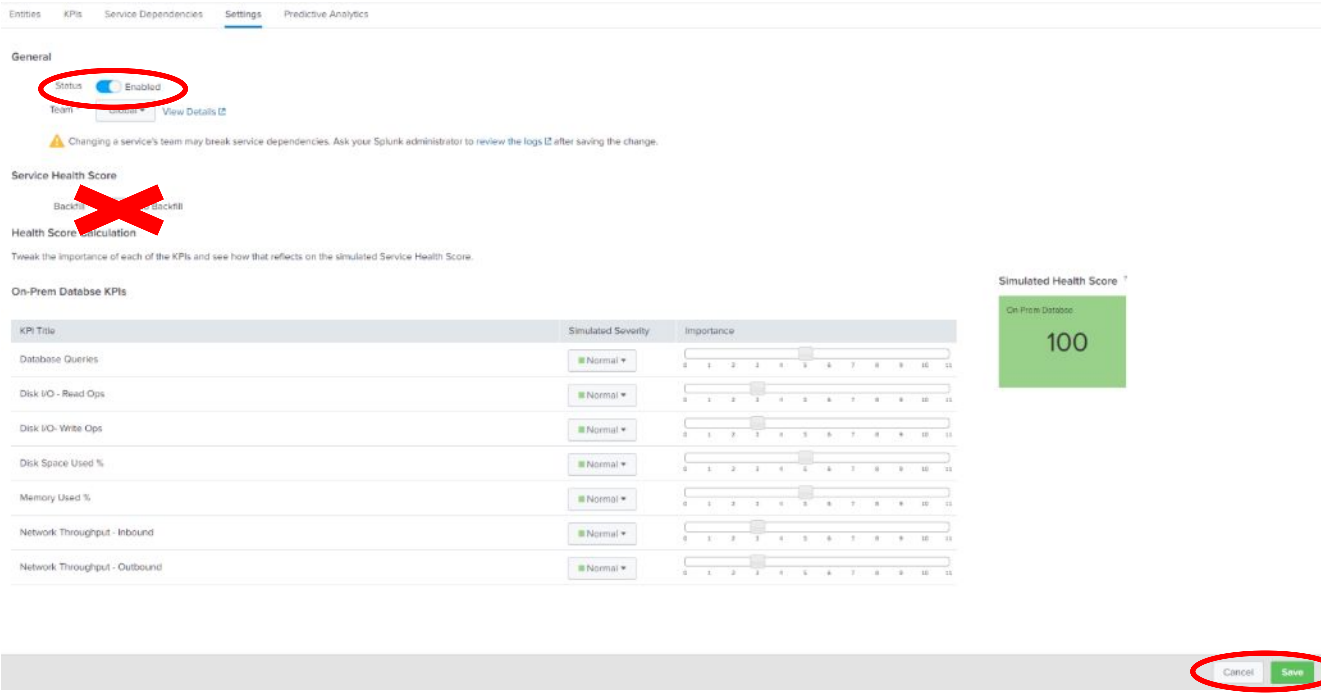

- Click the ‘Save’ button (Bottom-Right)

Services Lab

The new Database service will be a dependency of the **Authorization** service, any service health changes will be propagated to the parent service(s).

- Select *Configure > Services*
- Edit the *‘Authorization’ service*

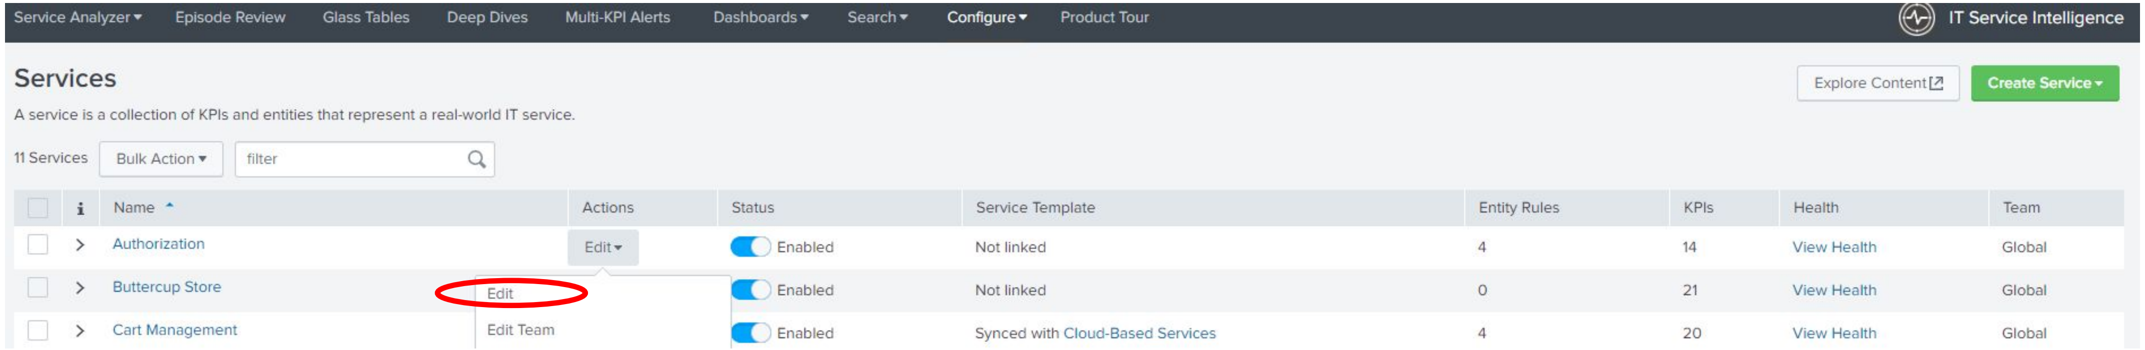

- *Click ‘Service Dependencies’ tab*

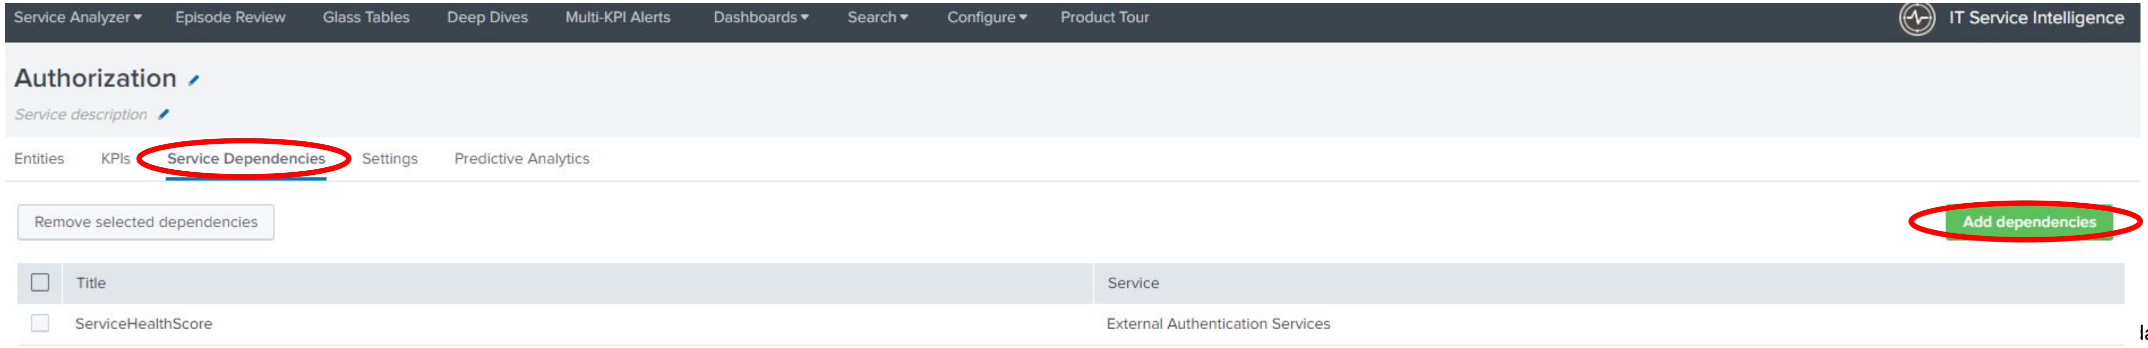

lata into doing

Services Lab

- Tick ‘On-Prem Database’ service
- Tick the ’ServiceHealthScore’
- Press ‘Done’ button
- Press ‘Save’ button

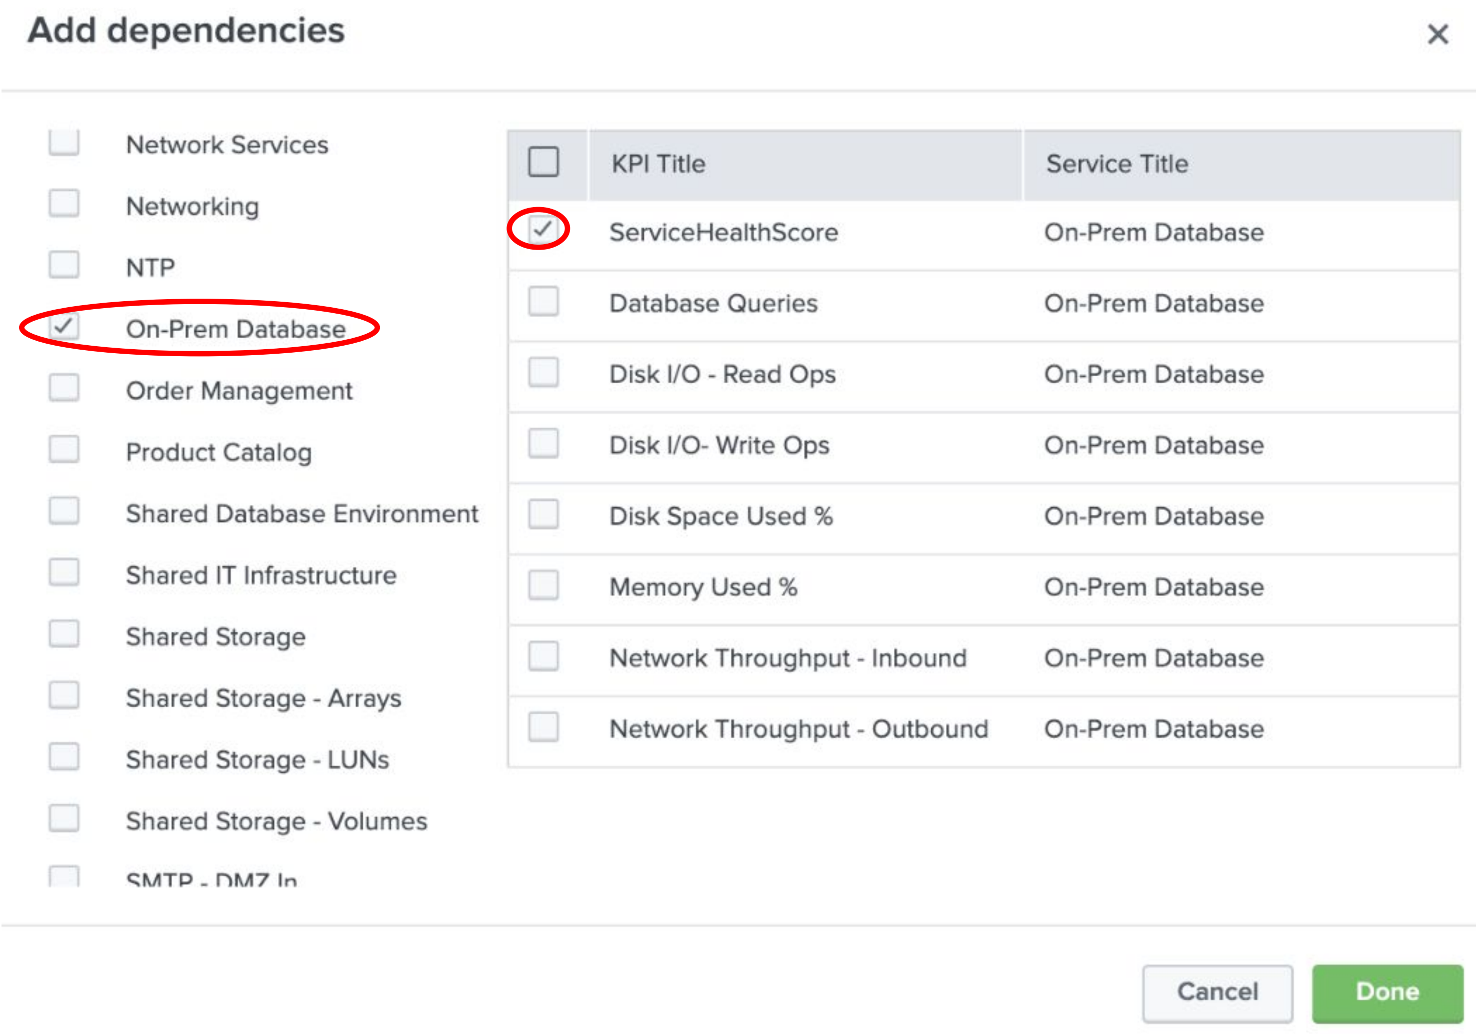

Check Service Tree

- Next go to Service Analyzer, and select the service tree
- Check if the On-Prem database is added

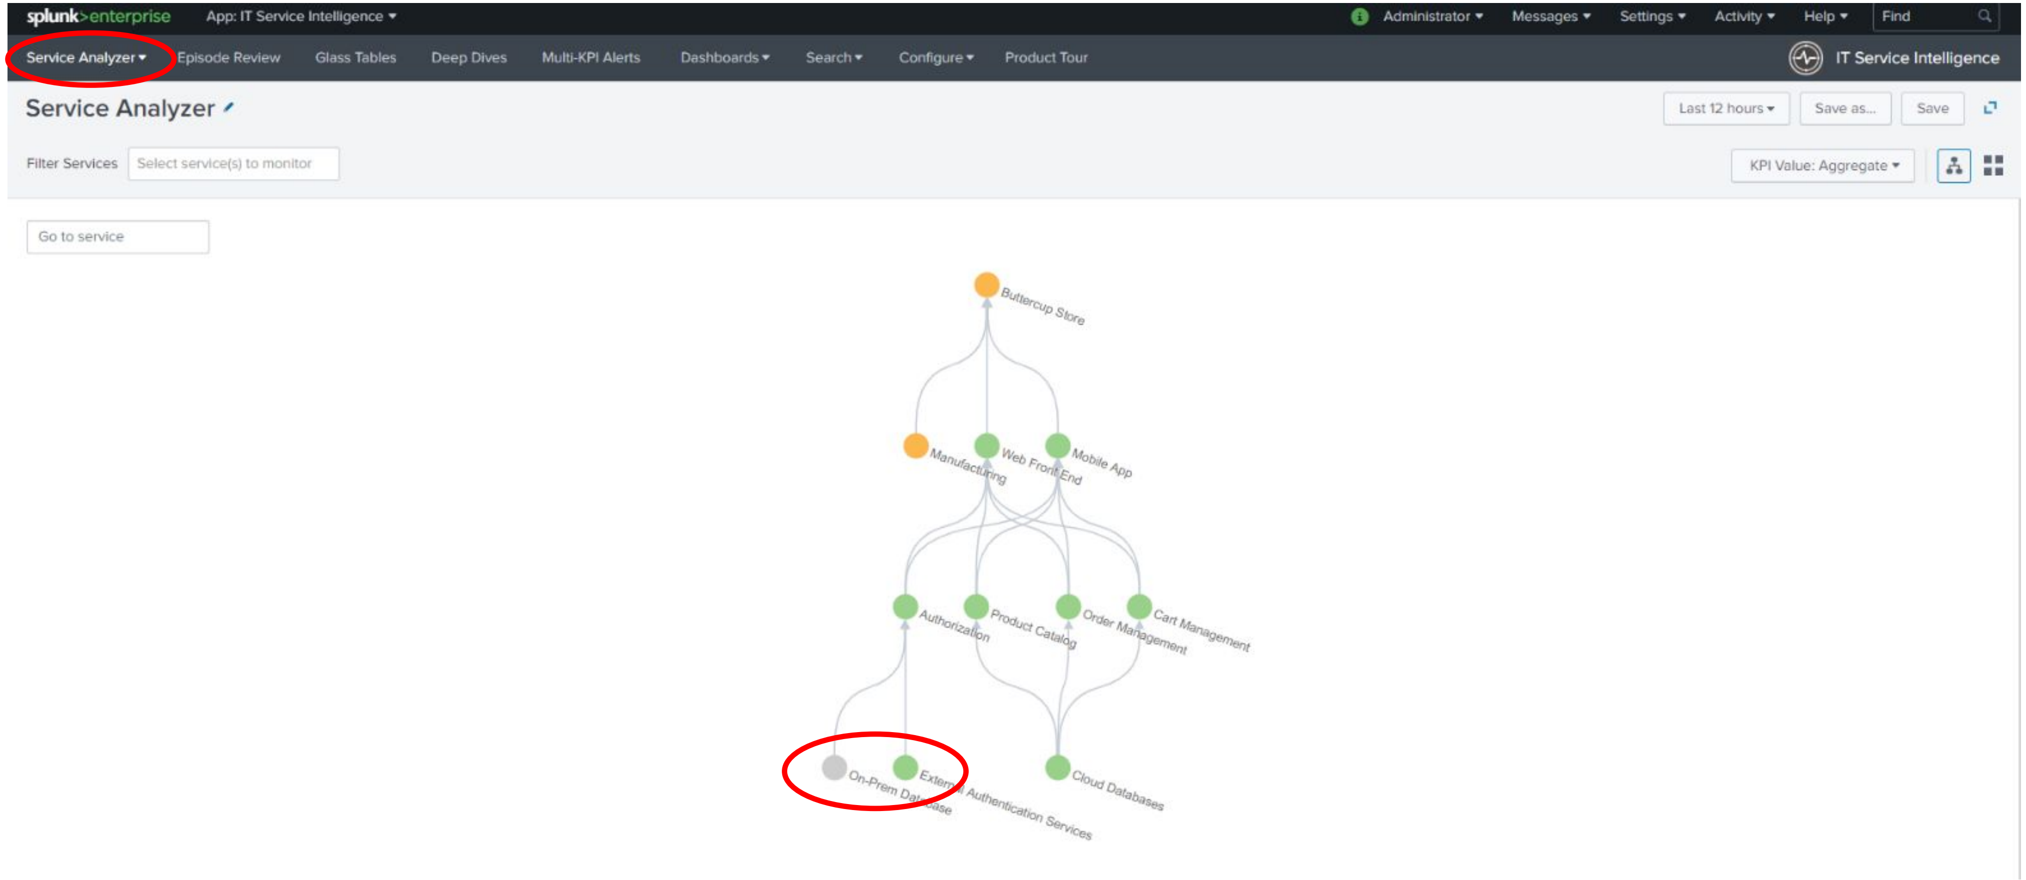

KPI Exercise

KPI Lab

The new ‘On-Prem Database’ service is based on a template however we need to add an extra ad-hoc KPI to monitor the CPU utilization.

- Select Configure > Services
- Select the ‘On-Prem Database’
- Select KPI tab
- Click New > Generic KPI
- Set Title to ‘CPU Utilization %’
- Click ‘Next’ button

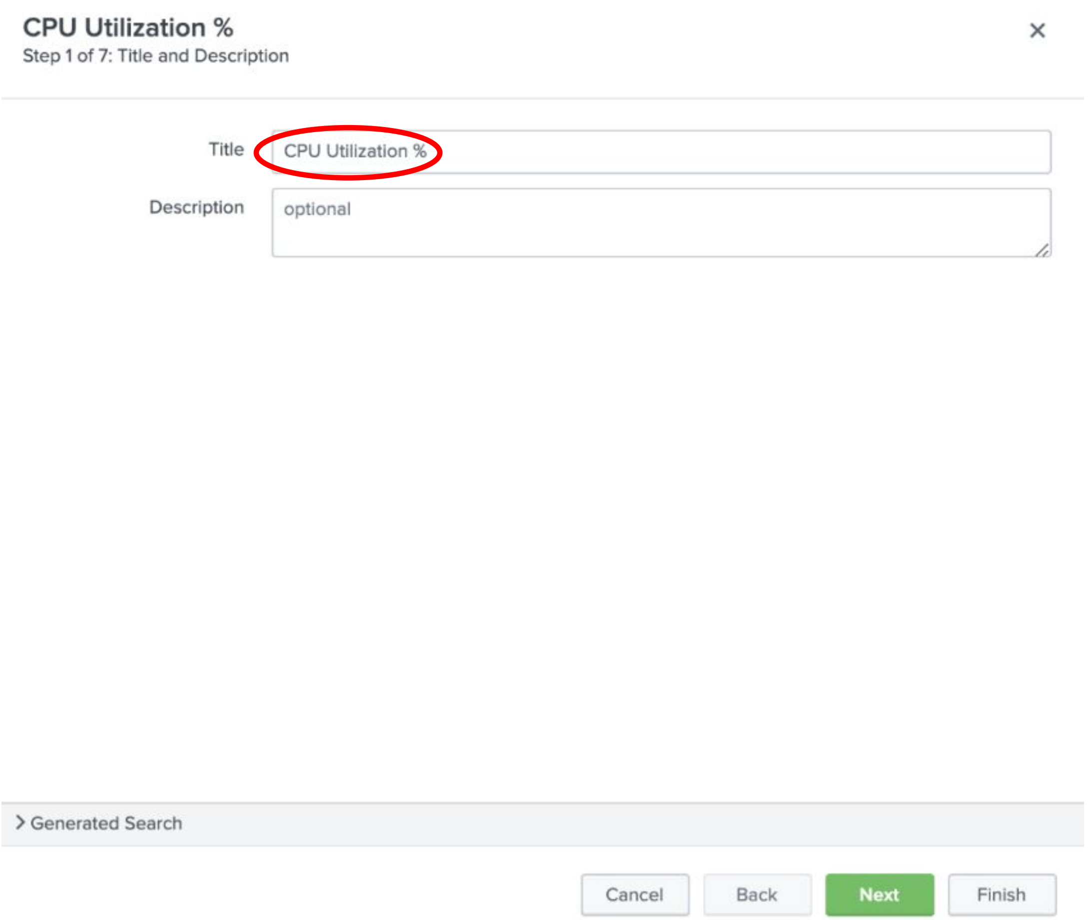

KPI Lab

The new KPI source could be driven by a data model, ad-hoc search or a base search. It is always best to utilize base searches as they can return multiple KPI metrics with a single search.

- Click 'Base Search'
- Select 'All Metrics'
- Select 'Average CPU (Linux)'

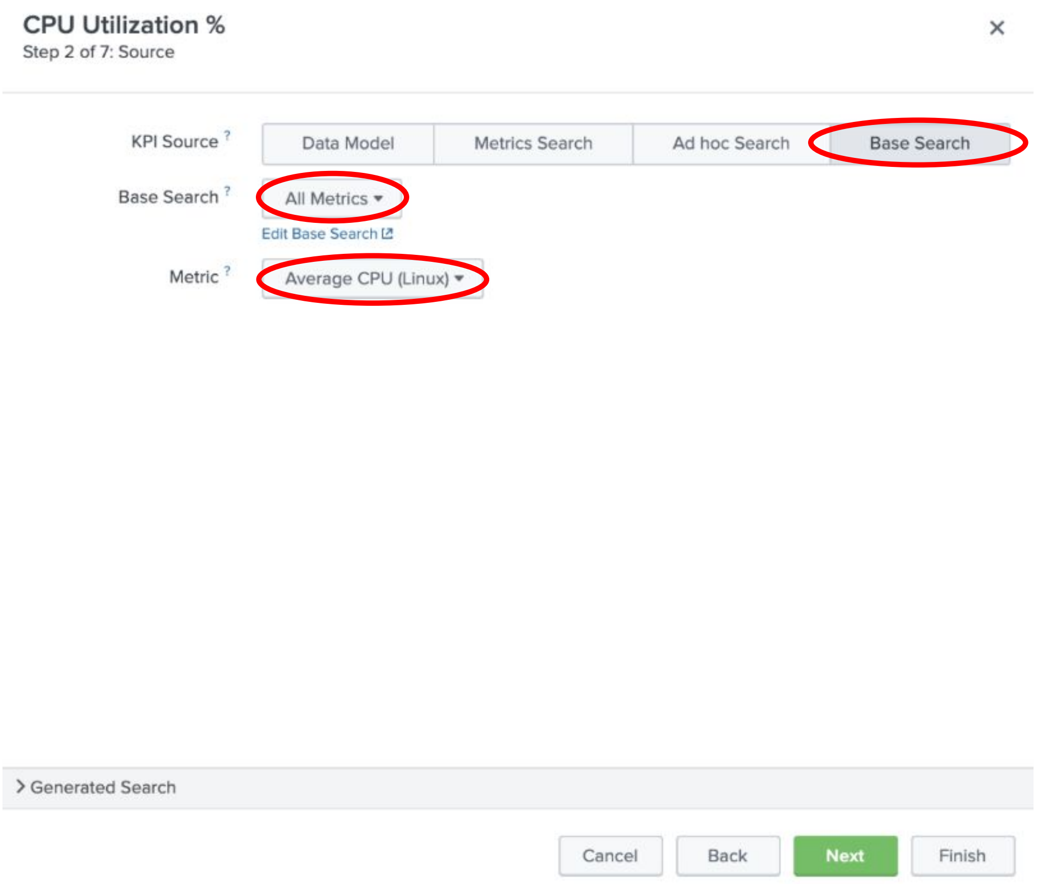

- Click 'Next' button

KPI Lab

- There is no option to split this base search.

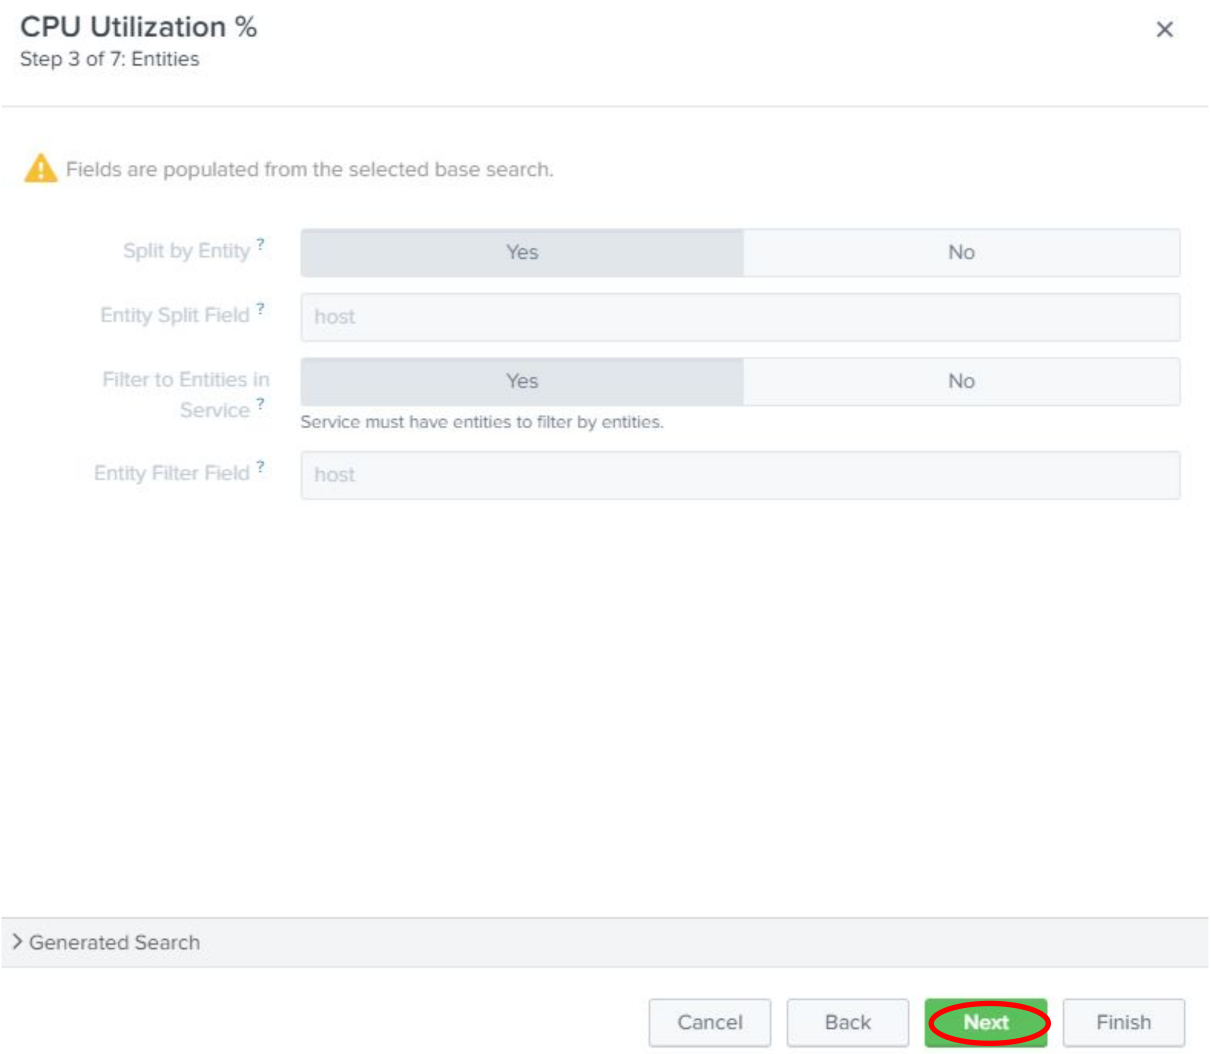

- Click ‘Next’ button

KPI Lab

We need to build the KPI calculation criteria, this includes how often data is collected, the entity calculation and the calculation window.

- We have decided that the metric will be collected every 5 minutes
- We want to perform an **Average** of the metric
- Over 5 min window
- Click 'Next' button

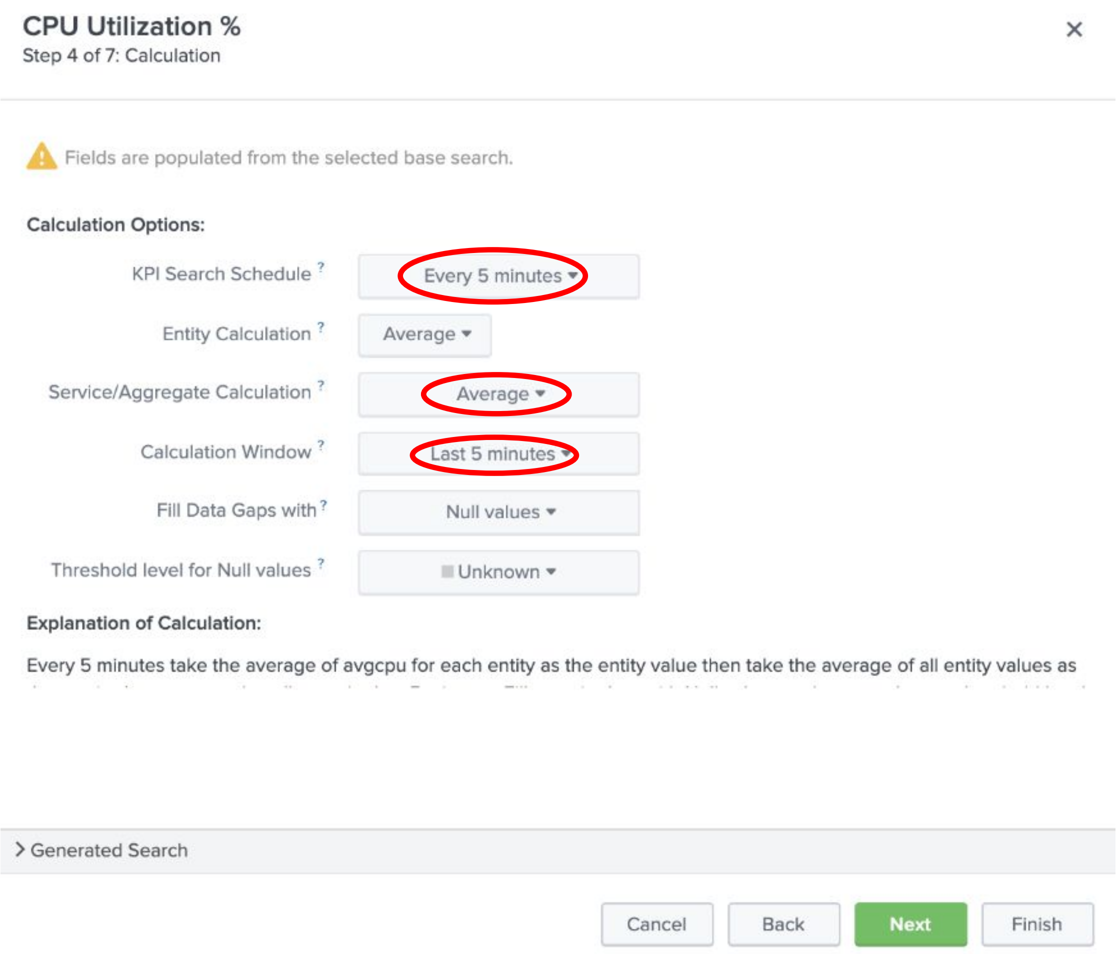

KPI Lab

The monitoring and unit fields will be populated from the base search.

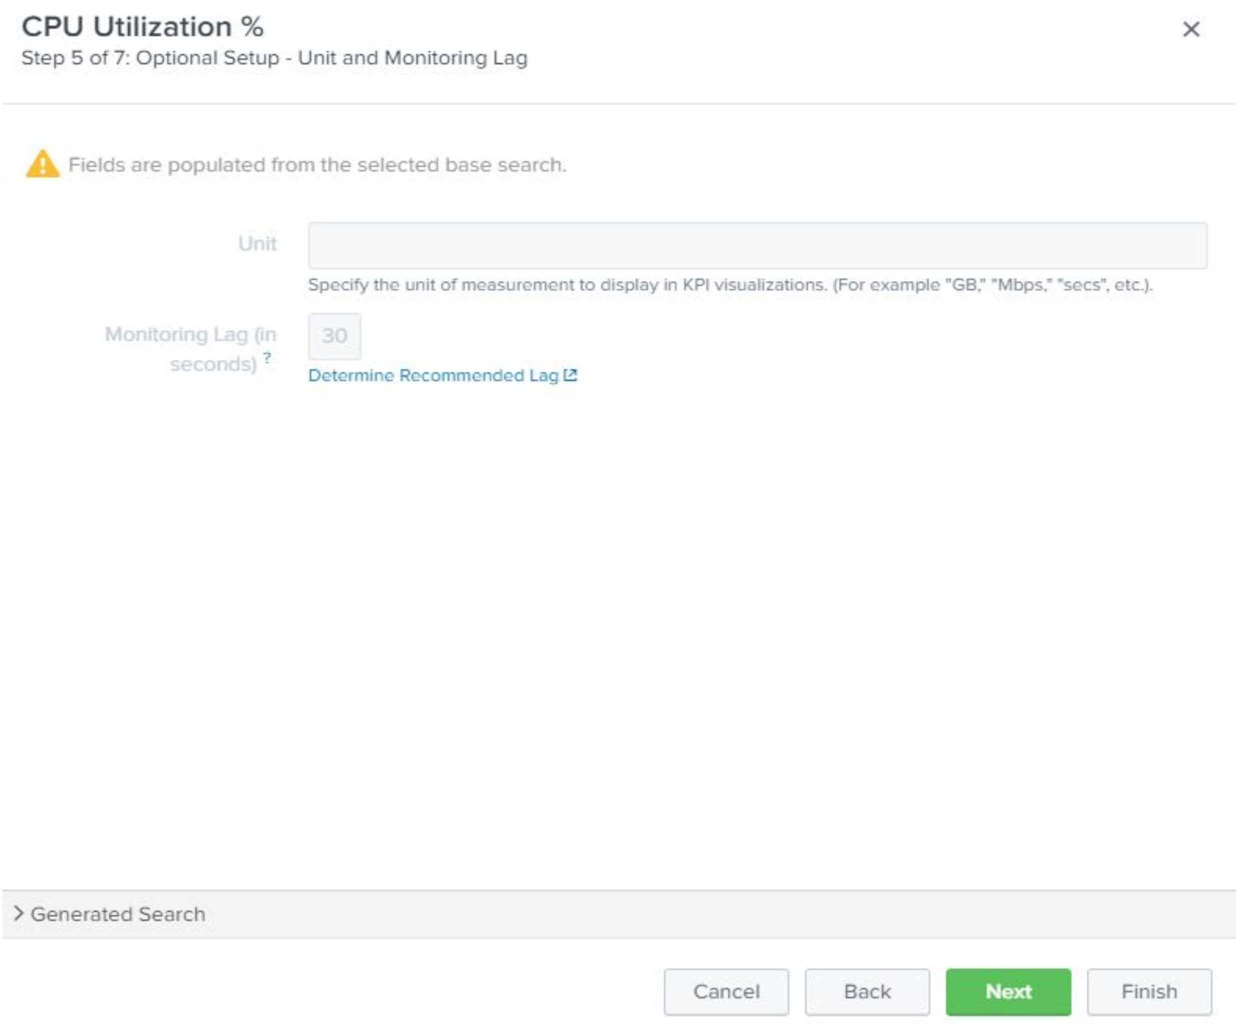

- Click ‘Next’ button

KPI Lab

We want this KPI to use data already ingested in Splunk over the last 7 days.

- Click ‘Enable Backfill’ button
- We will leave the backfill period as 7 days

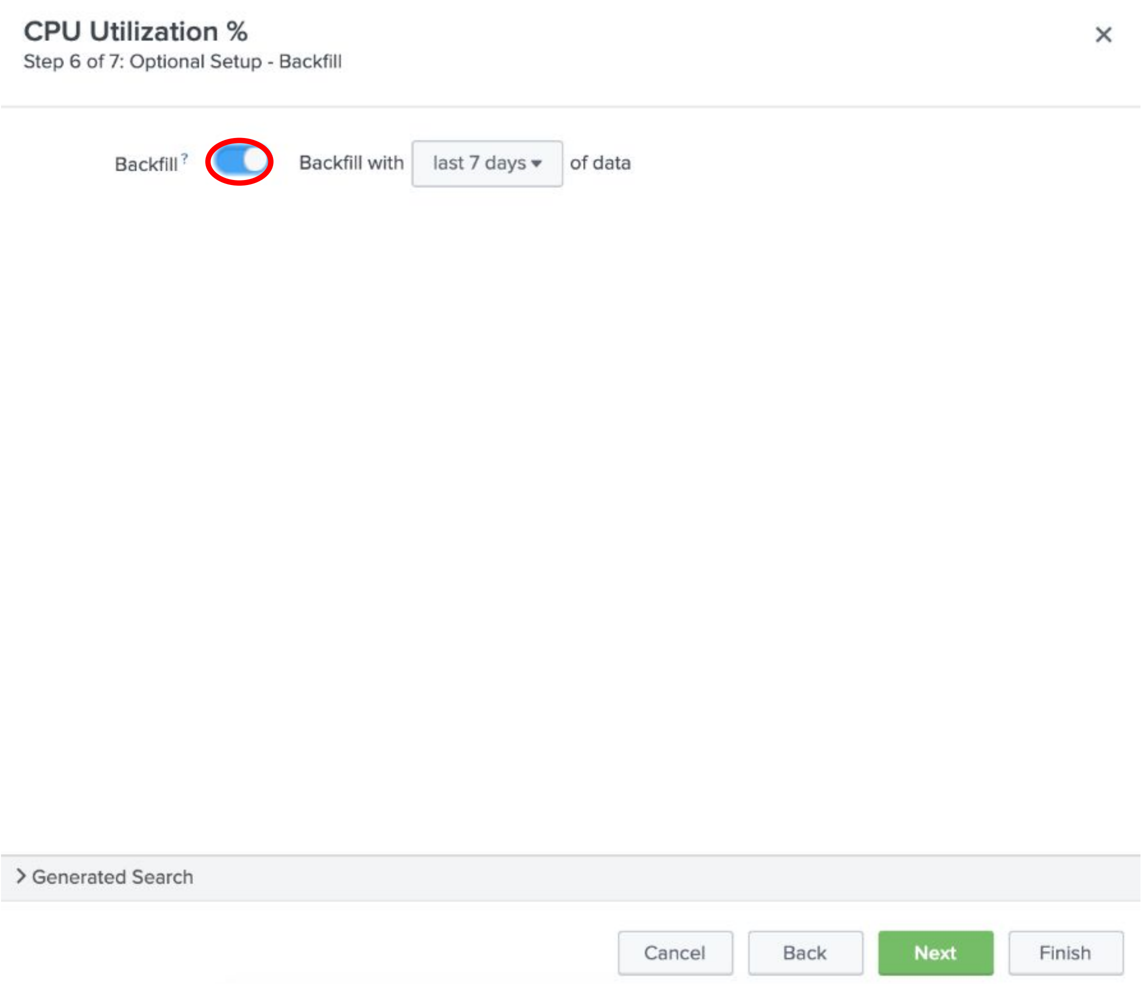

- Click ‘Next’ button

KPI Lab

We need to set some static thresholds for this new KPI

- Increase time to 4 hours
- Add & configure threshold:
  - Critical = 95
  - High = 90
  - Medium = 85
  - Low = 70
- *Click 'Finish' button*

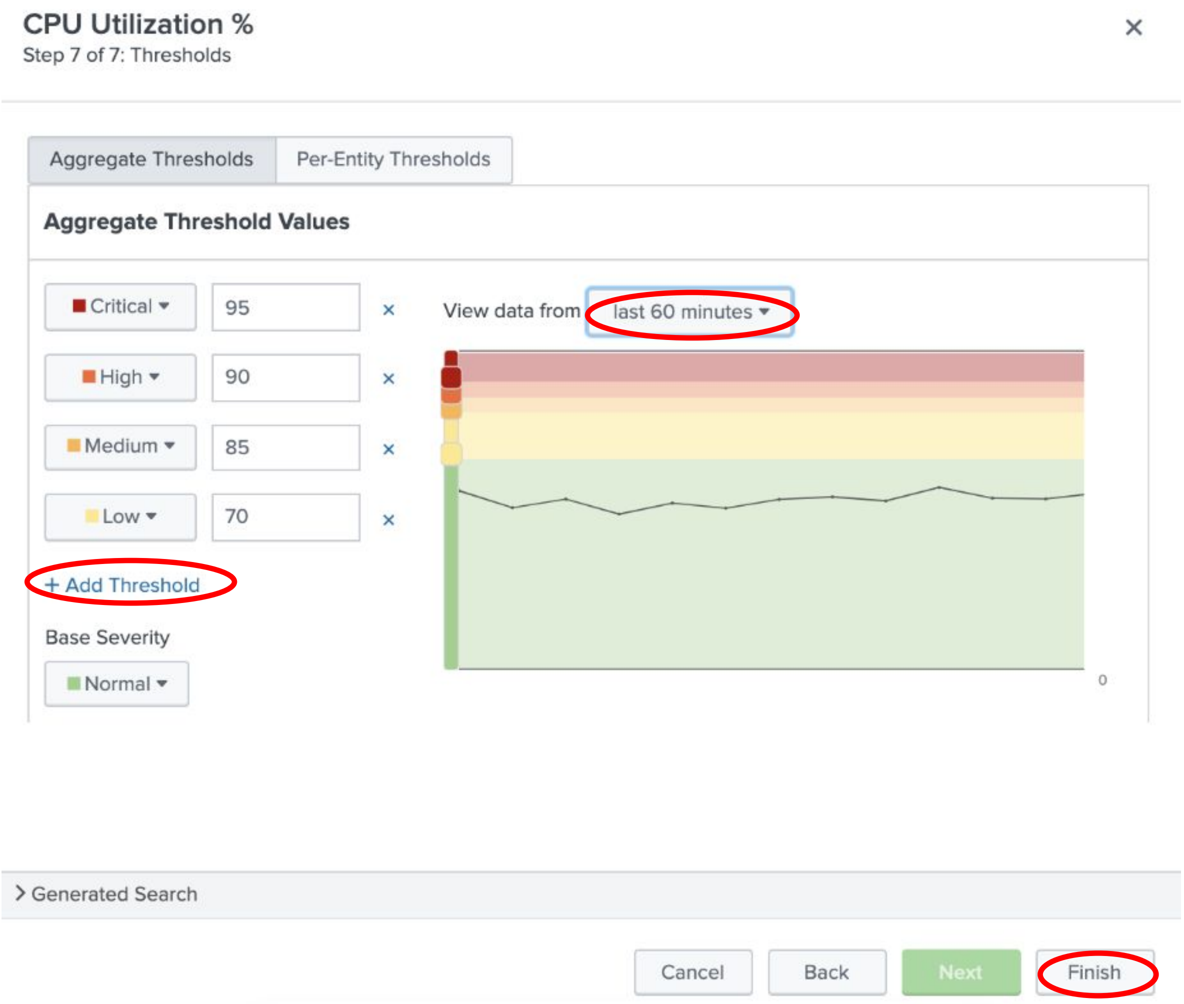

KPI Lab

- Note that the new KPI does not have a padlock icon. Inherited KPIs are locked to the service template so when changes are made these are pushed to the linked services, such as the one we are configuring.

- If you edit a locked KPI it will become an orphan and template changes no longer adopted

- *Click 'Save' button! Bottom Right*

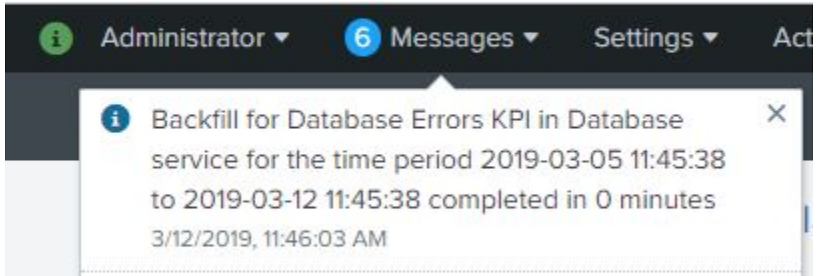

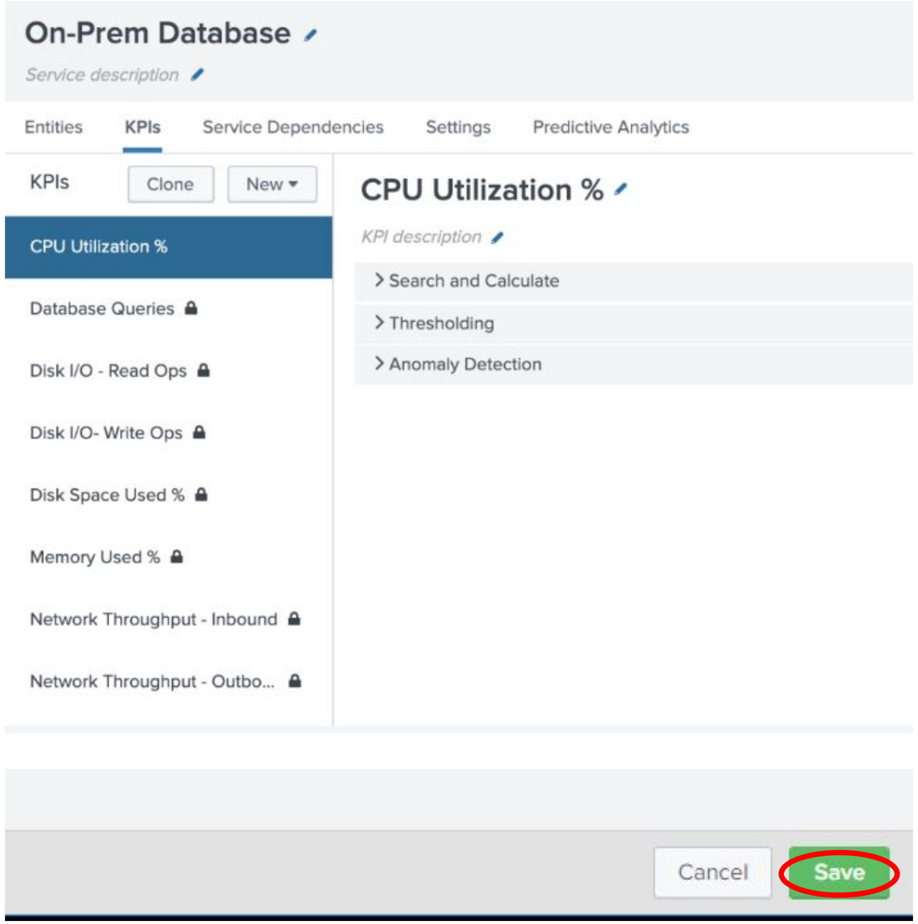

KPI Lab

An important dependency for the new ‘On-Prem Database’ service is response time, in this lab we will add an extra KPI: ‘Database Response Time’. For this KPI we use an ad-hoc search.

- Click Configure > Services
- Click ‘On-Prem Database’ & KPI tab
- Click New KPI > Generic KPI
- Title = ‘Database Response Time’

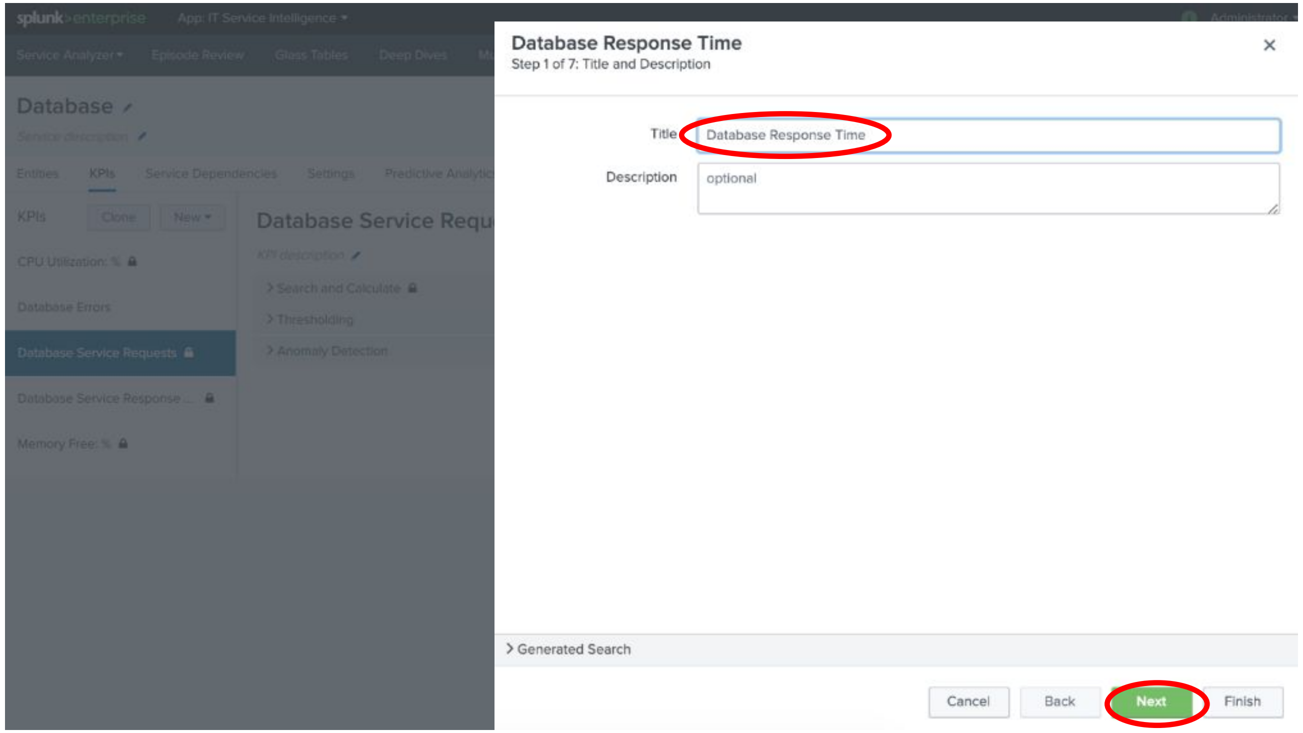

- Click ‘Next’ button

KPI Lab

The new KPI source could be driven by a data model, ad-hoc search or a base search. In this instance we will create the KPI using an ad-hoc search.

- Click ‘Ad-hoc Search’
- Enter the following search:
  - index=itsidemo sourcetype=stream:mysql query=*
- Enter ‘time_taken’ as the threshold field
- Click ‘Run Search’ button to test search

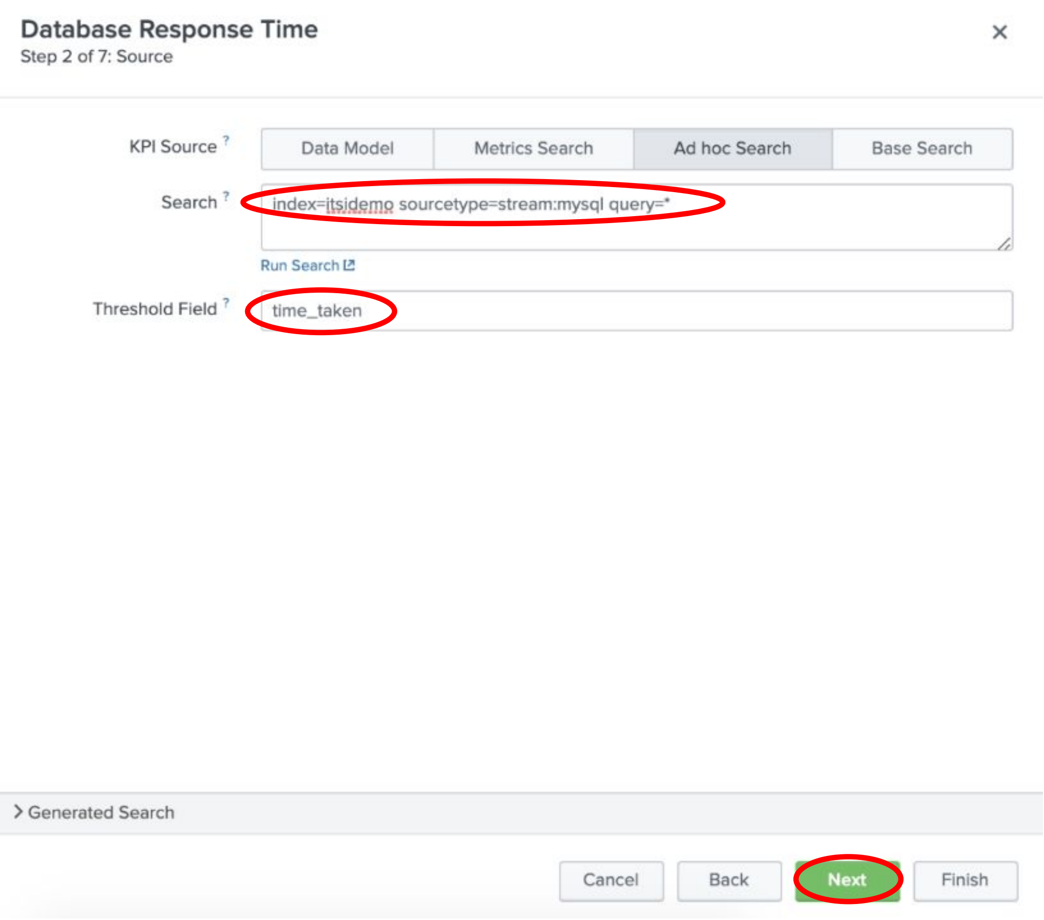

- Click ‘Next’ button

KPI Lab

We will split this KPI via the database name and filter via ip address.

- Select ‘Yes’ for split by entity
- Enter ‘dbname’ as the split by field
- Select ‘Yes’ for filter to entities in Service
- Type ‘ip’ for entity filter field

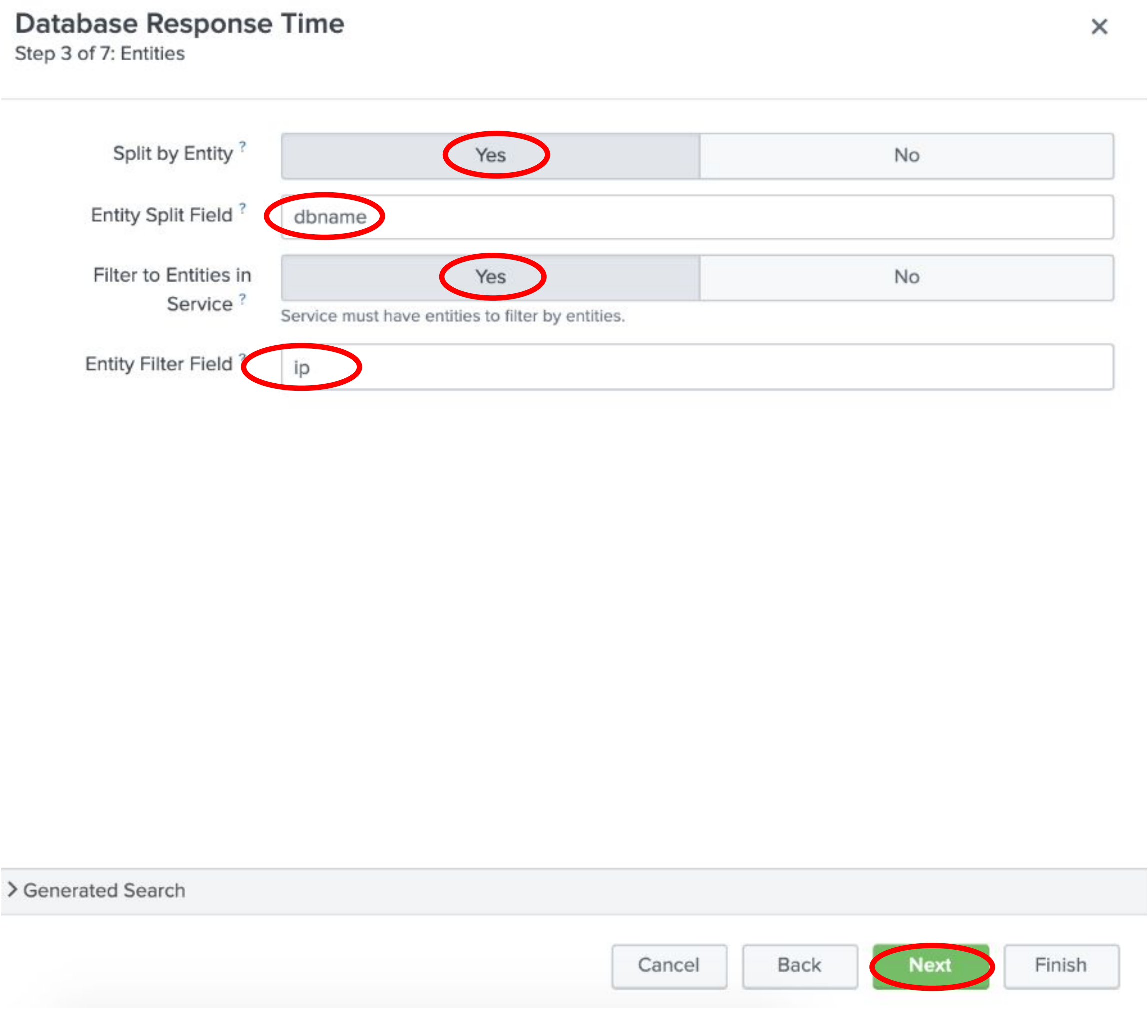

- Click Next

KPI Lab

The database response time KPI will have the following calculation options.

- Schedule = Every minute
- Entity Calculation = Average
- Aggregate Calculation = Average
- Calculation Window = 5 minutes

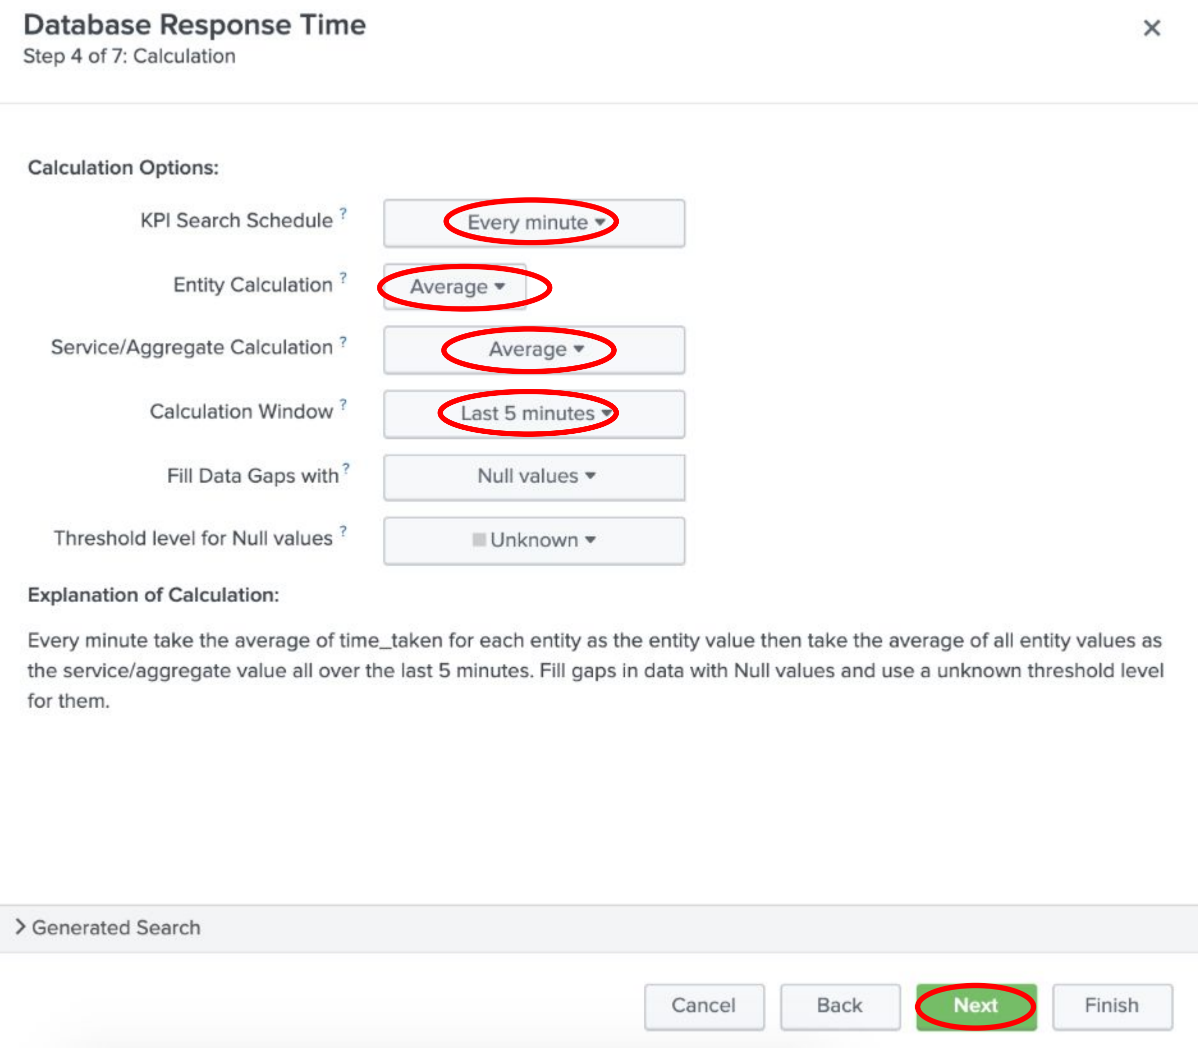

- Click Next

KPI Lab

We will leave the next screen with the default values.

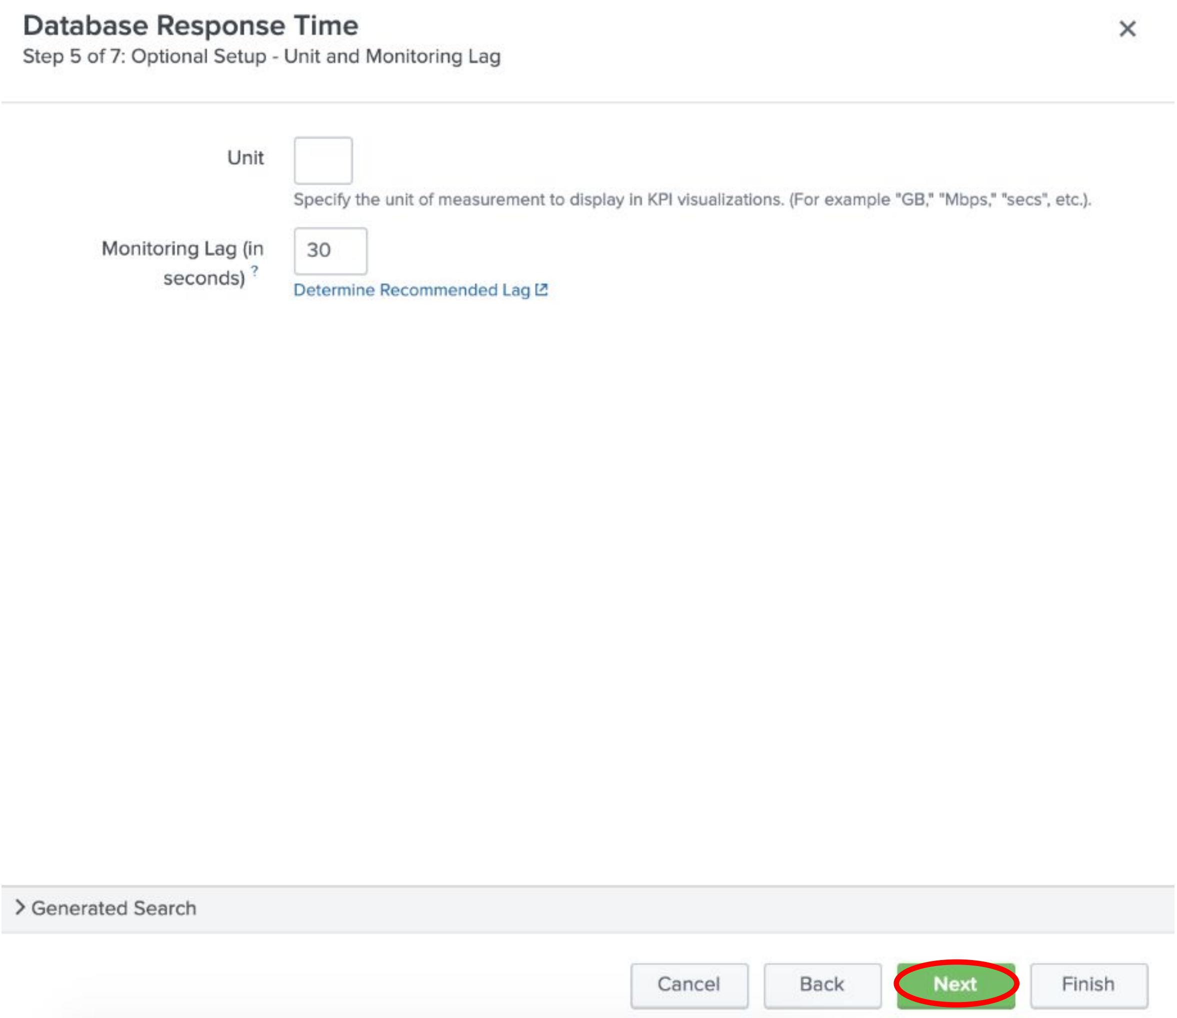

▶ Click Next

KPI Lab

We want this KPI to use data already ingested in Splunk over the last 7 days. This historical data will be used in the machine learning labs.

- Click 'Enable Backfill' button
- We will leave the backfill period as 7 days

- Click '*Next*' button

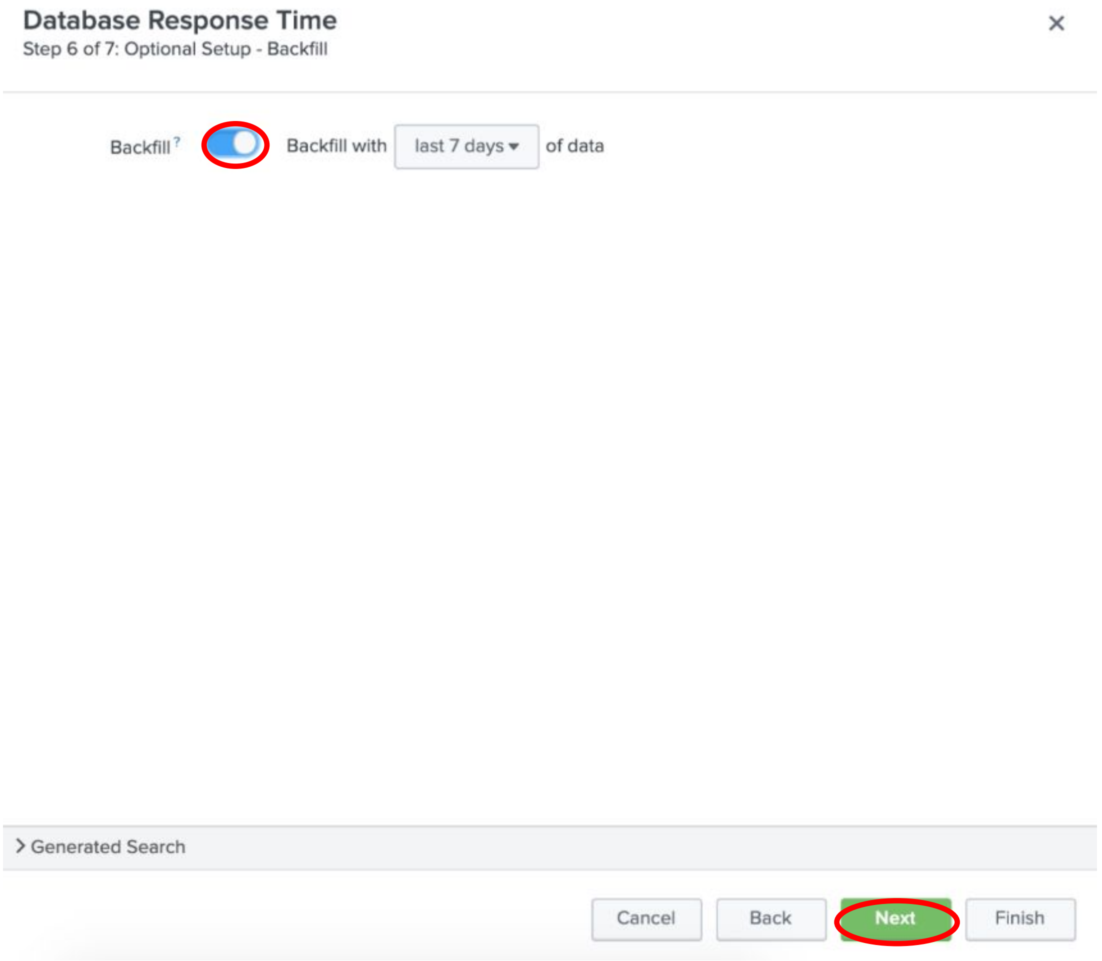

KPI Lab

We need to set some static thresholds for this new KPI

- Increase time to 4 hours
- Add & configure threshold:
  - High = 360
  - Medium = 310
  - Low = 200
- Click '*Finish*' button
- Do not forget to click '*Save*' button!!

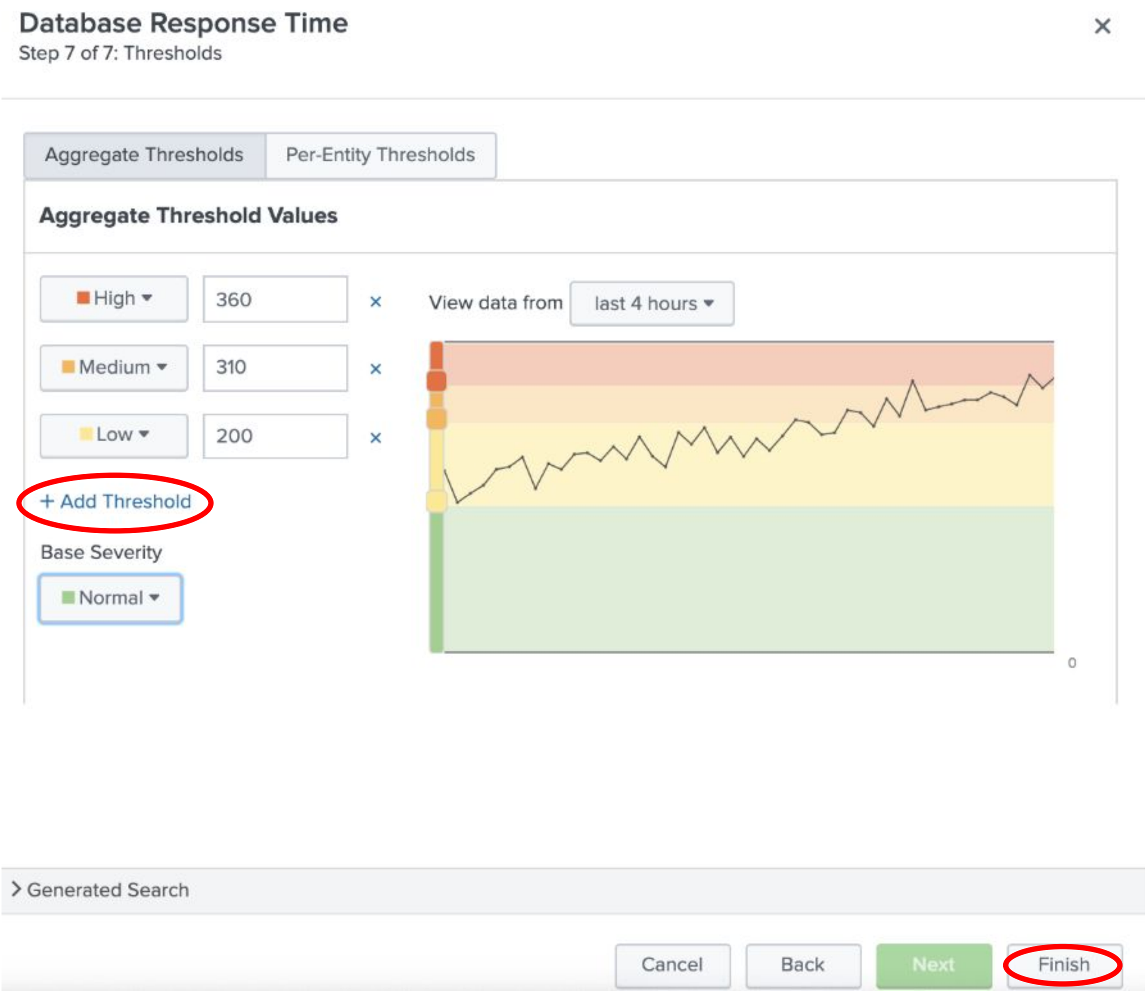

KPI Lab

We have two new KPIs for our On-Prem Database service, this new service can utilize past data via the Service Health Score backfill capability.

- Switch to the '*Settings*' tab
- Toggle backfill on (last 7 days)
- Click 'Save' button

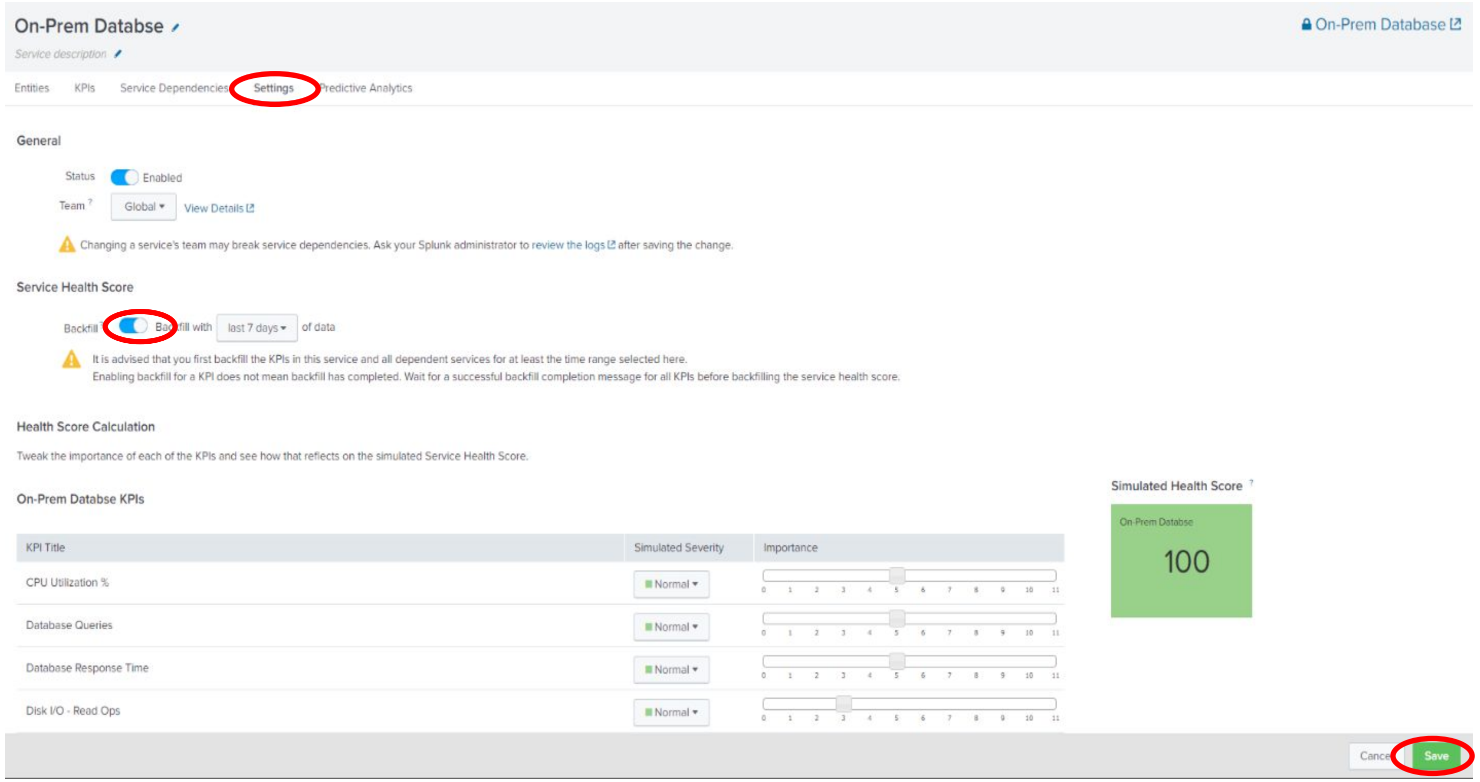

Infrastructure Monitoring

KPI Lab

There are now two new KPIs visible in the Service Analyzer for the On-Prem Database service, these KPIs contain historical data and we can see there is an issue with the CPU data.

- Navigate back to Service Analyzer
- Click ‘On-Prem Database’ service
- Select the new ‘CPU Utilization %’ KPI
- Click ‘mysql-02’ entity

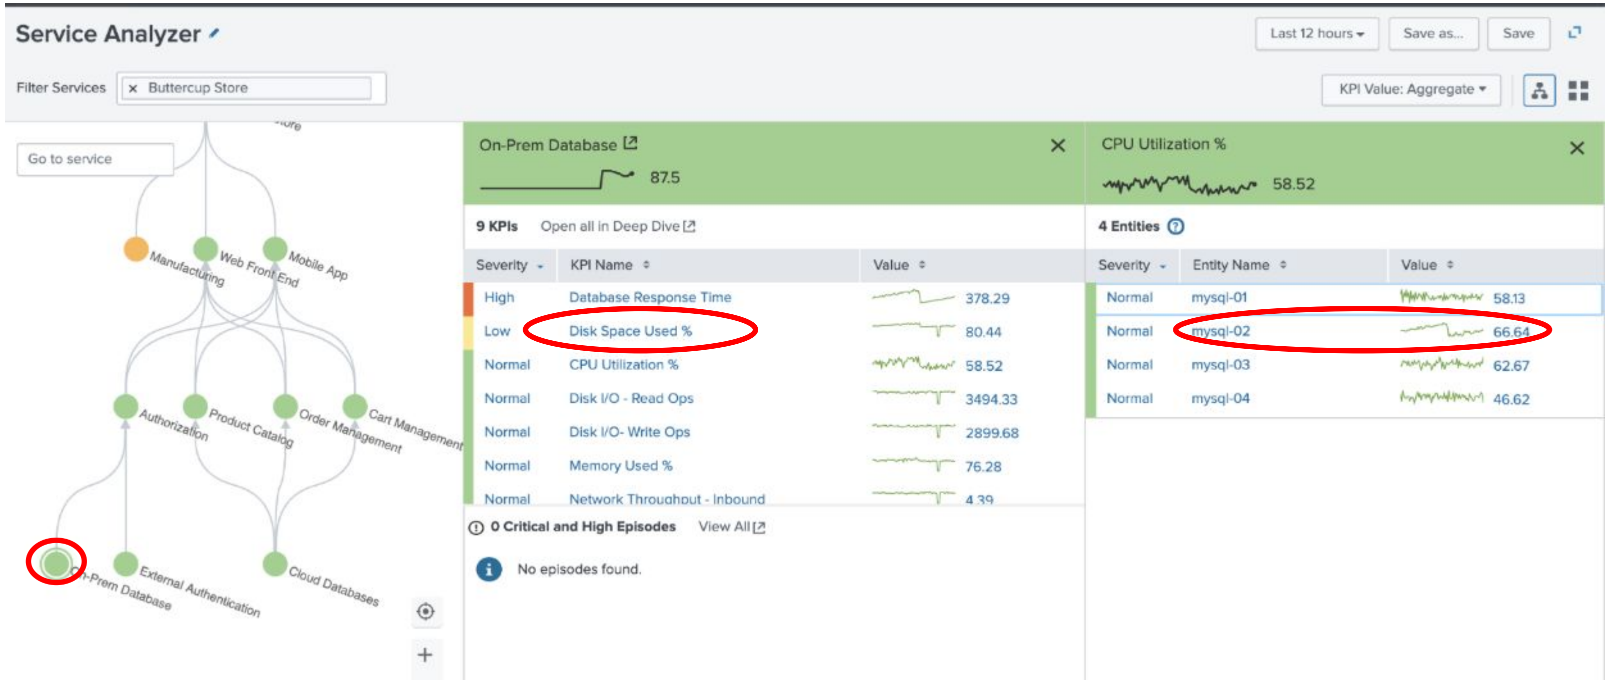

KPI Lab

The ‘Entities Details’ view is a high-level dashboard showing metric data for the selected entity.

- Click the Splunk App for Infrastructure for more metric information.

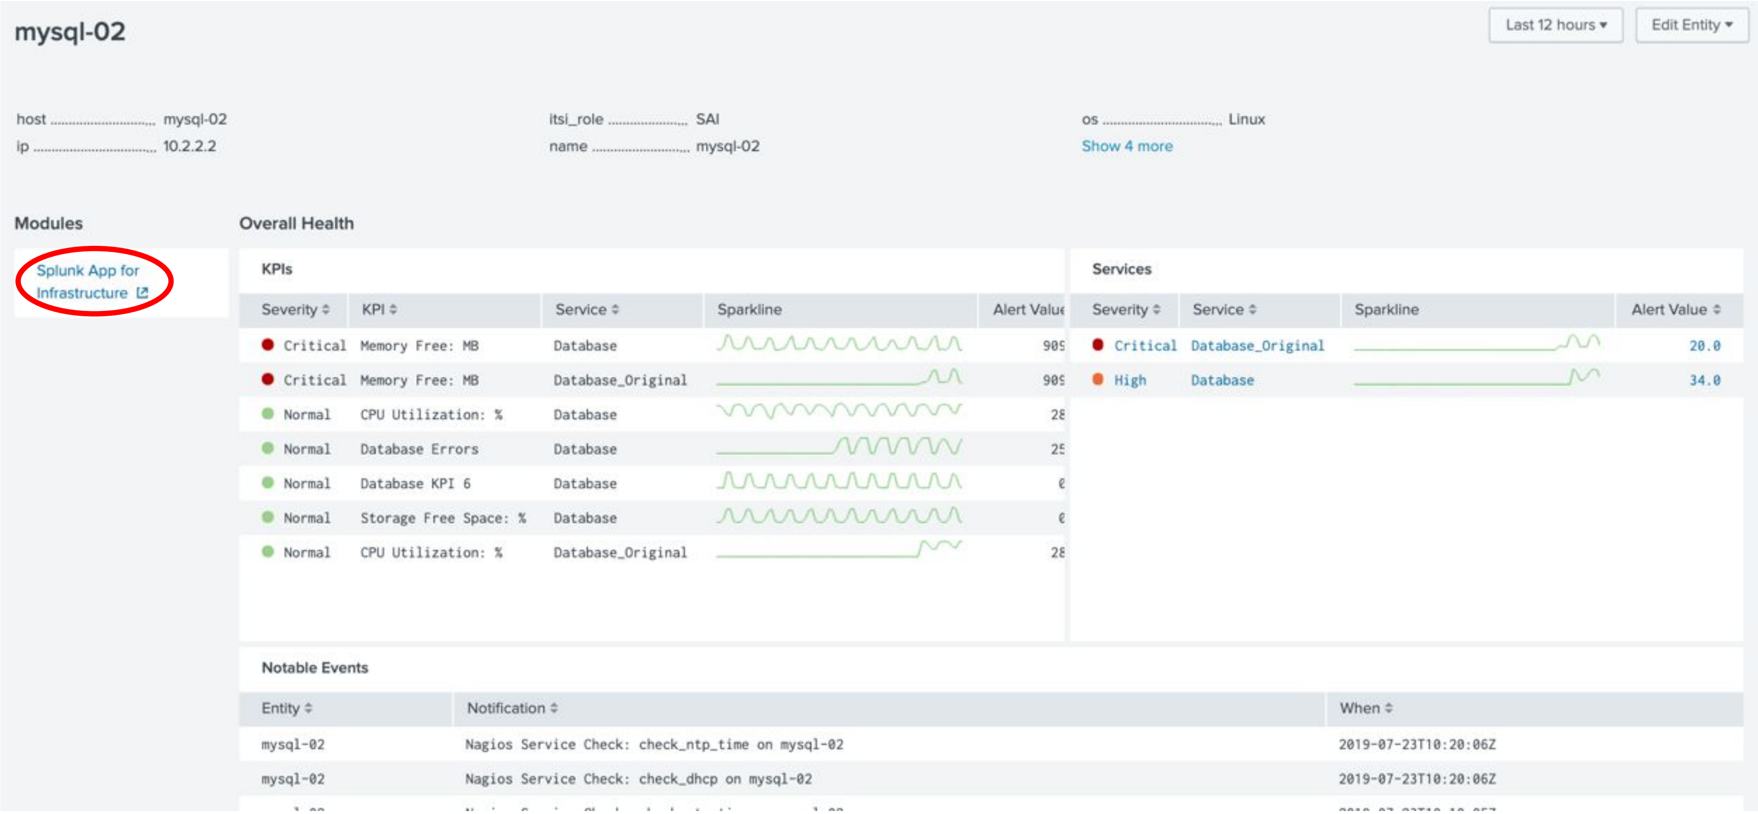

KPI Lab

The Splunk app for Infrastructure is a workspace to quickly drag and drop metric values.

- Click the 'analysis' tab to investigate the collected metrics for this entity.

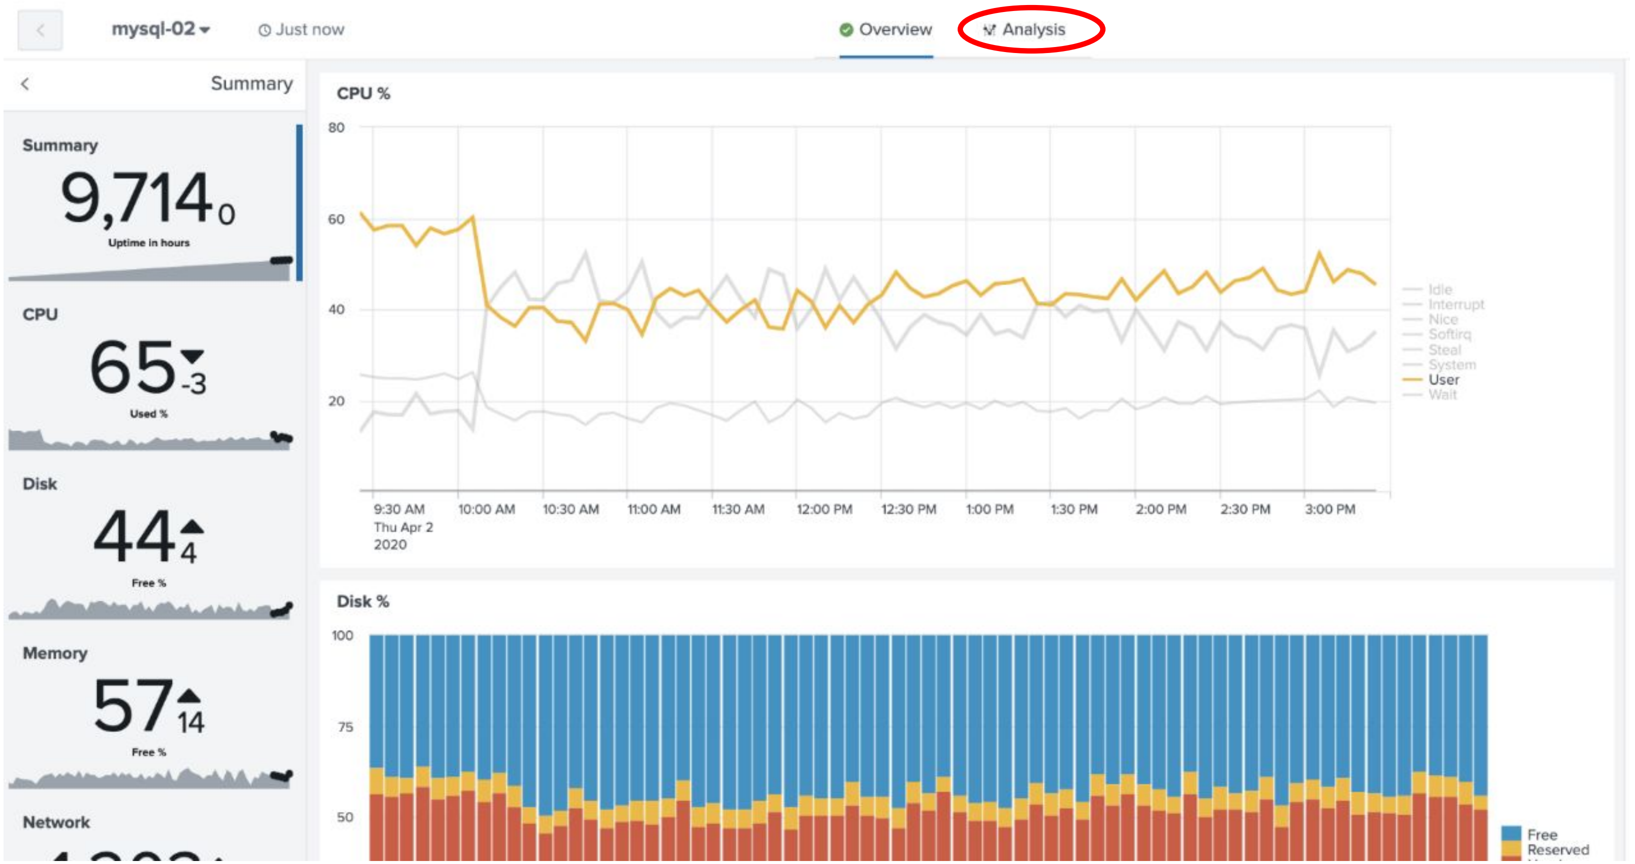

KPI Lab

- Challenge – If there is time attempt to recreating the following dashboard.

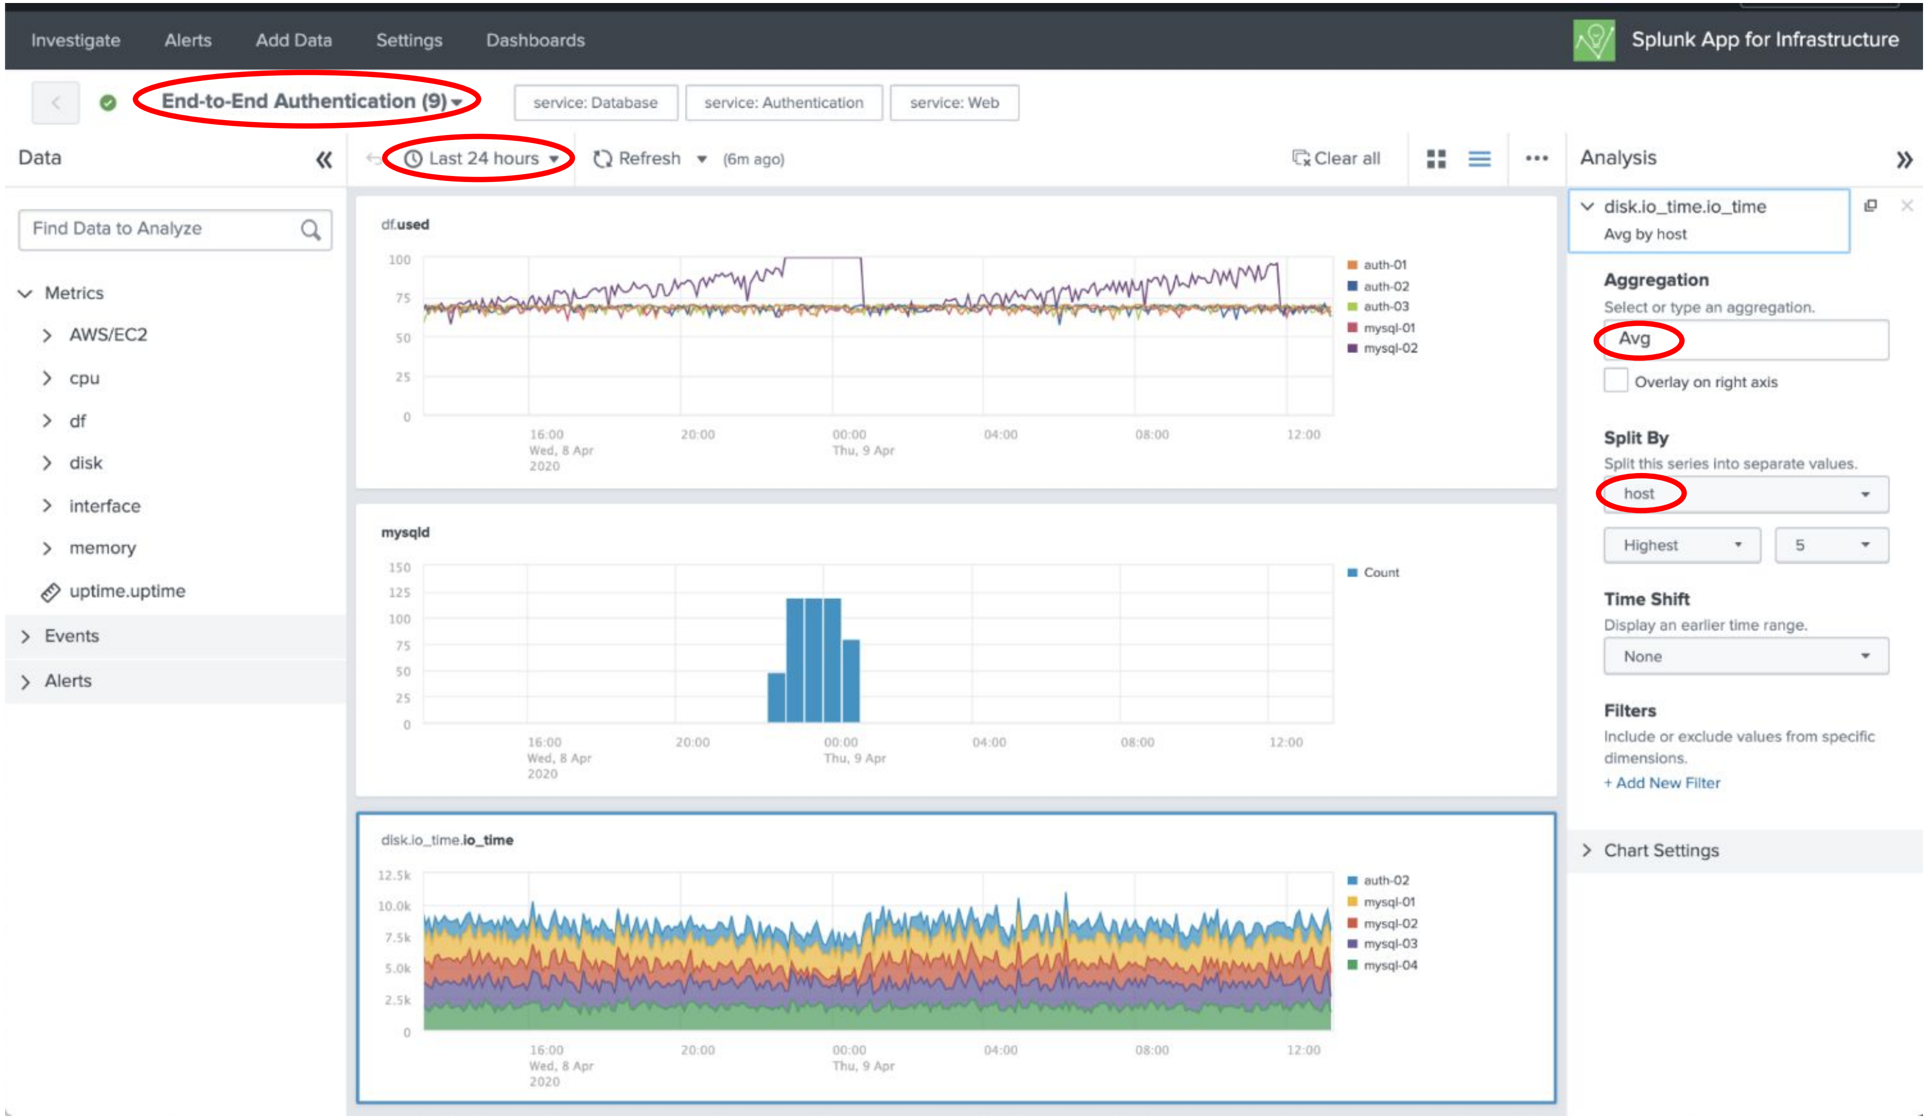

KPI Lab

- Navigate back to ITSI app for remaining labs.

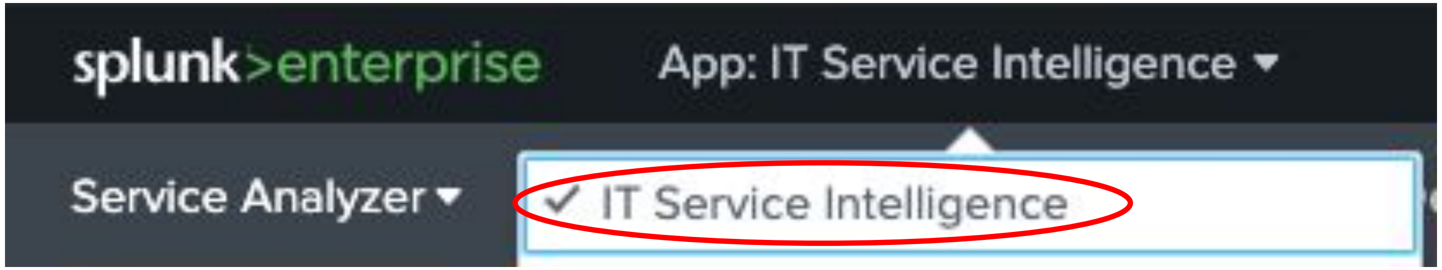

Machine Learning

Skill Areas for Machine Learning at Splunk

Splunk Expertise

Searching
Reporting
Alerting
Workflow

ITSI

Premium solutions provide out of the box ML capabilities.

MLTK

Splunk ML Toolkit facilitates and simplifies via examples & guidance

Data Science Expertise

Statistics/maths background
Algorithm selection
Model building

Domain Expertise

(IT, Security…)

Identify use cases
Drive decisions
Understanding of business impact

ITSI Machine Learning

Anomaly Detection

Deviation from past behavior
Deviation from peers
Unusual change in features
ITSI Anomaly Detection

Adaptive Thresholds

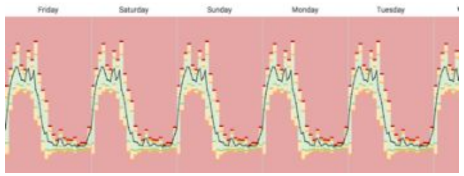

Adaptive Thresholds
What is normal behaviour and what is not normal
Ideal for cyclical and dynamic data

Predictions Analytics

Predict Service Health Score
Predicting events
Trend forecasting
Early warning of failure
Predictive maintenance

Event Clustering

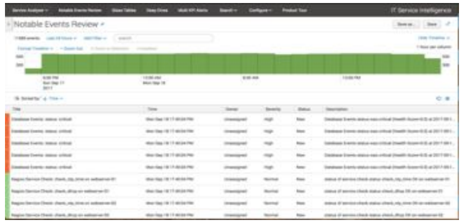

Identify peer groups
Event correlation
Reduce alert noise
ITSI Event Analytics

Machine Learning process

Ecosystem Splunk's App Ecosystem contains 1000's of free add-ons for getting data in, applying structure and visualizing your data giving you faster time to value.

MLTK The Machine Learning Toolkit delivers new SPL commands, custom visualizations, assistants, and examples to explore a variety of ml concepts.

Splunk Splunk Enterprise is the mission-critical platform for indexing, searching, analyzing, alerting and visualizing machine data.

Operationalized Data Science Pipeline

Collect Data	Clean & Munge	Search & Explore	Pre-processing Feature Selection	Choose Algorithm	Build, Test, Improve Models	Operationalize Monitor Alert	Visualize & Share
Ecosystem	Ecosystem	Splunk	MLTK	MLTK	MLTK	MLTK	Ecosystem
Splunk	Splunk		Splunk	Splunk	Splunk	Splunk	Splunk

splunk> turn data into doing

What Data Scientists Really Do

Data Preparation accounts for about 80% of the work of data scientists

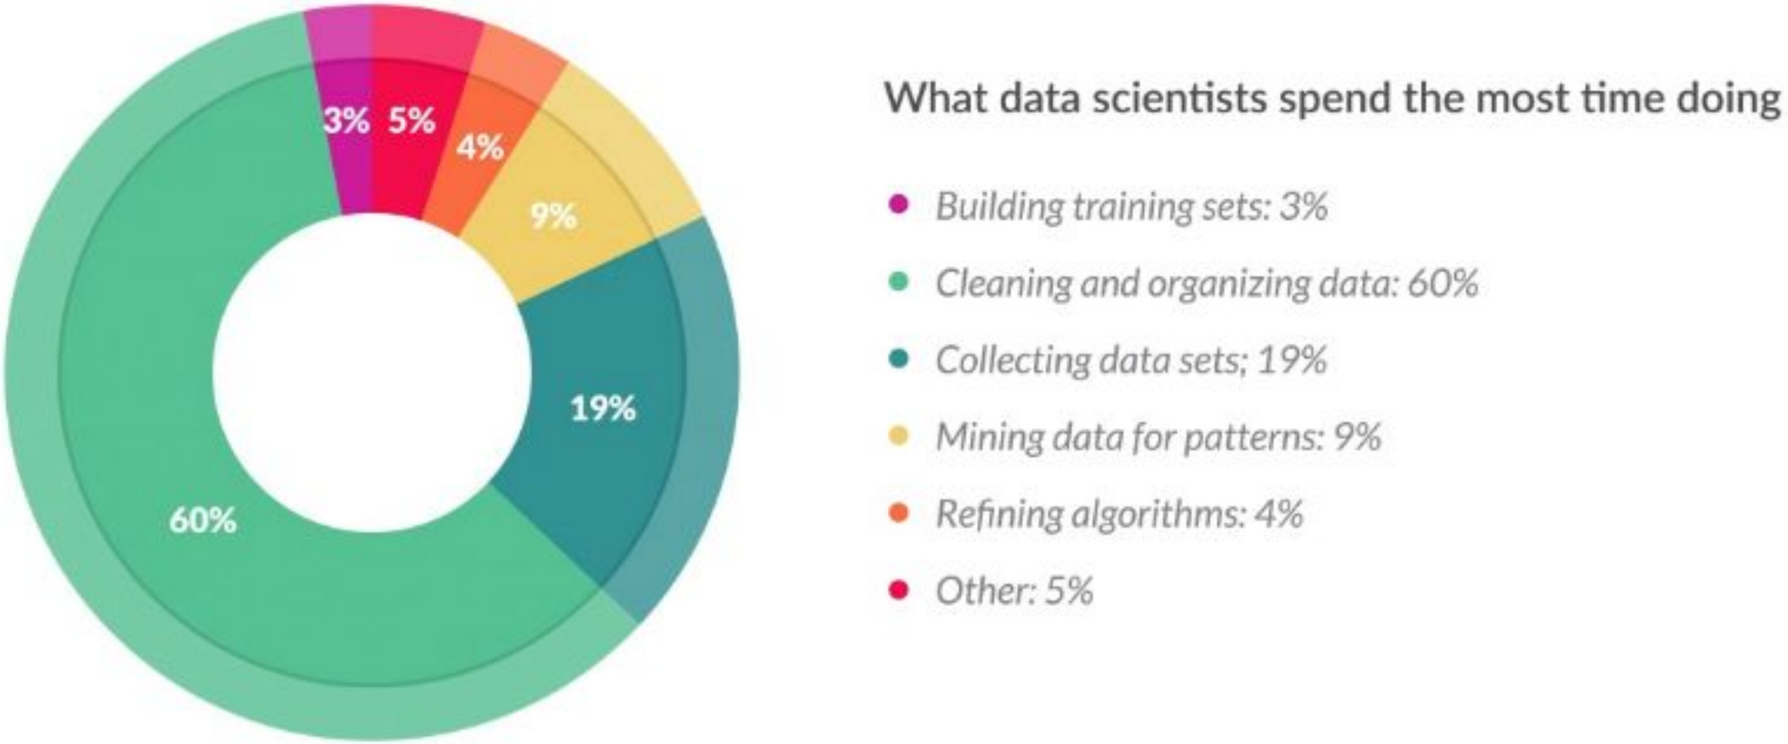

“Cleaning Big Data: Most Time-Consuming, Least Enjoyable Data Science Task, Survey Says”, Forbes Mar 23, 2016

What Data Scientists Really Do

Data Preparation accounts for about 80% of the work of data scientists

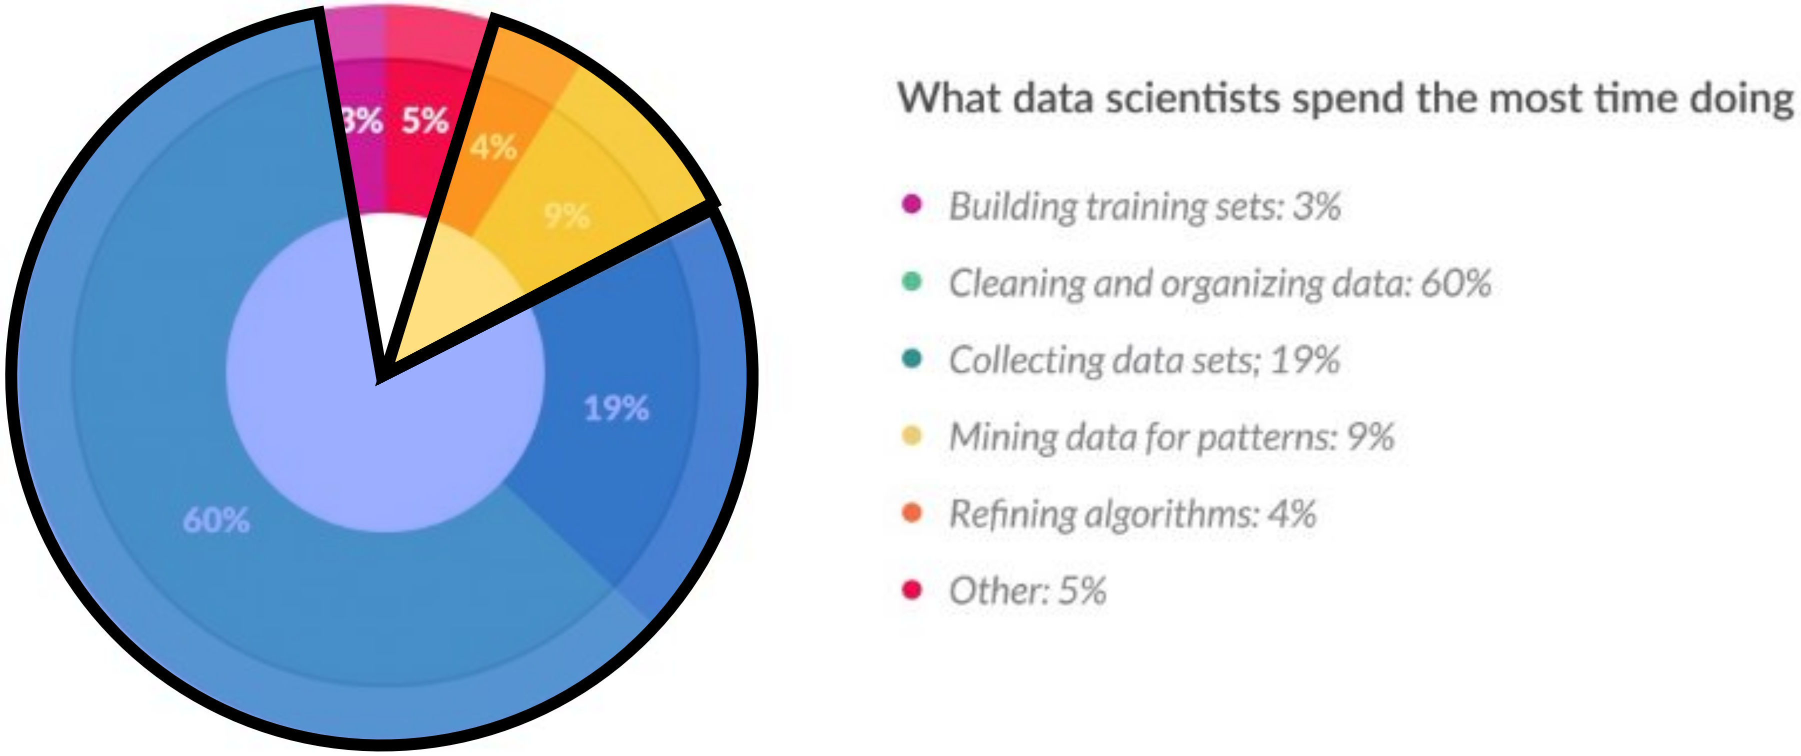

“Cleaning Big Data: Most Time-Consuming, Least Enjoyable Data Science Task, Survey Says”, Forbes Mar 23, 2016

Machine Learning Use Case

- There have been lots of social media comments regarding Buttercup' website availability, especially during the evenings.

- Luckily the website data is being ingested by Splunk, this sourcetype is called 'access_combined' and contains lots of information.

- The objective of this exercise to is utilise Splunk ITSI machine learning capabilities to identify normal behaviour for web response time.

- We will also look to see if the application servers supporting the website are functioning equally.

Machine Learning Lab

The Corporate Website Request KPI has no thresholding for alerting. In this lab we will configure machine learning to understand what normal looks like and alert us when the KPI falls outside this range.

- Go to the '*default*' service analyzer view and expand the tree view
- Click '*Web Front End*' service
- Select the '*Web Front End*' link

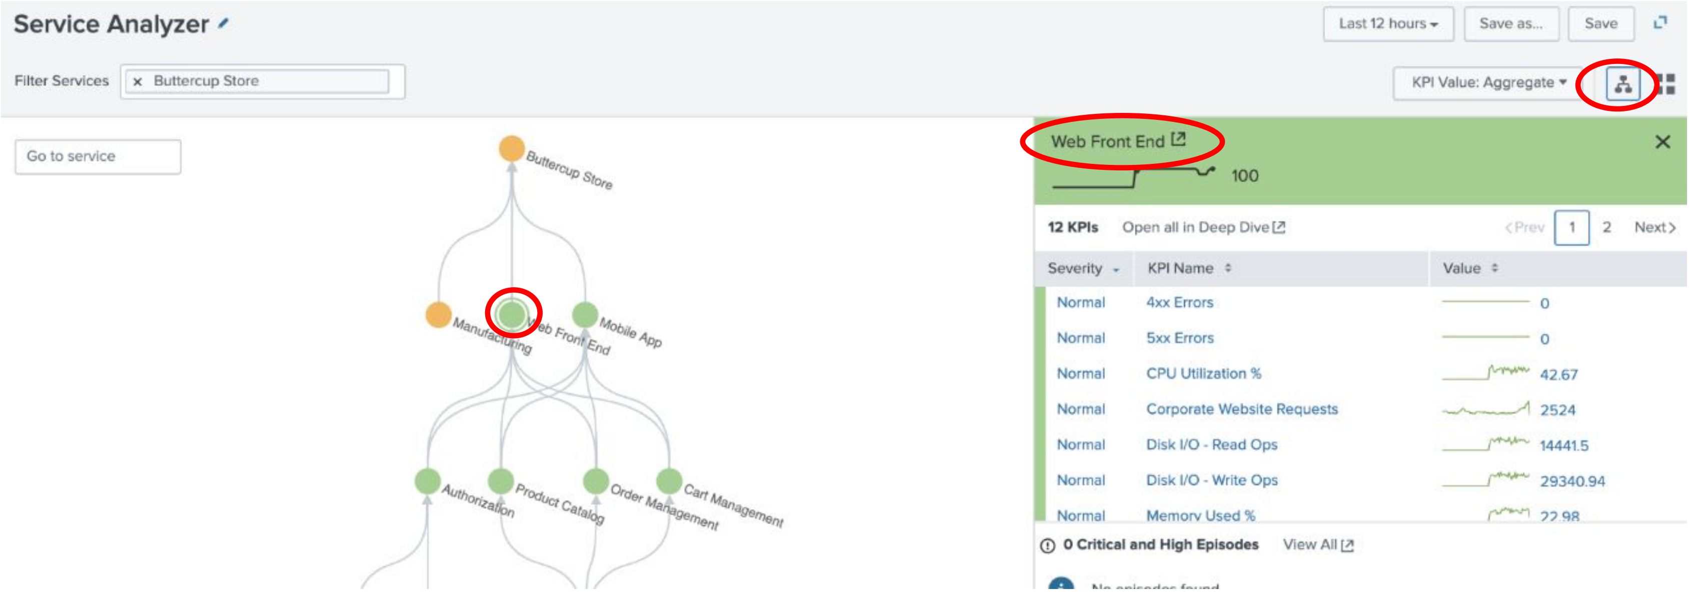

Machine Learning Lab

We need to instruct ITSI to use a template for thresholding, we will be using the adaptive standard deviation 3 hour working week template.

- Select the ‘*Corporate Website Requests*’ KPI

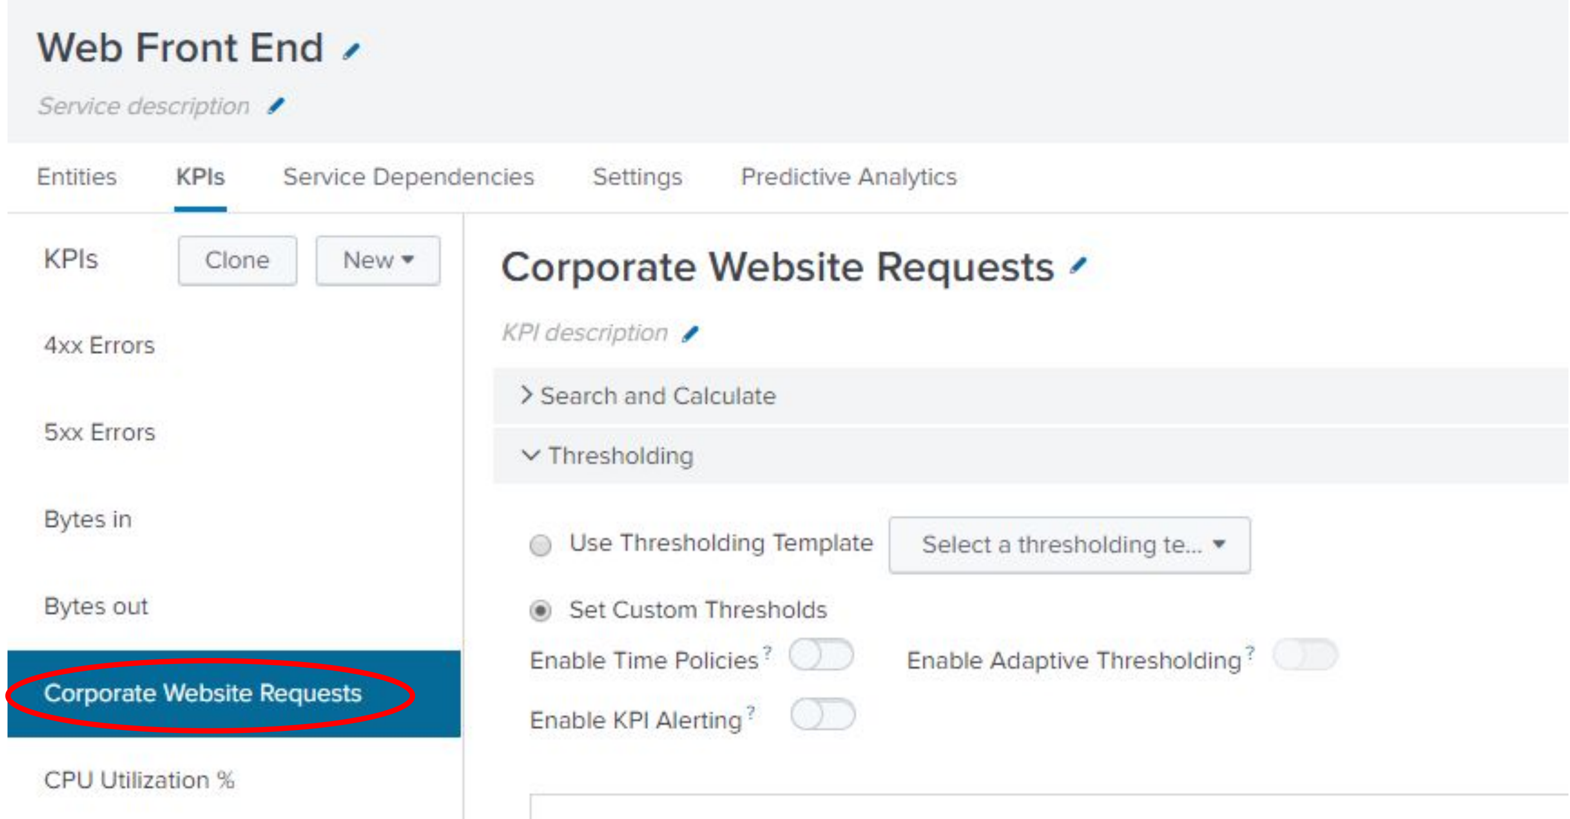

Machine Learning Lab

- Select ‘*use Thresholding Template*’ button
- Review the different options
- Select ‘*3-hour blocks work week*’
  - *(adaptive/stdev)*

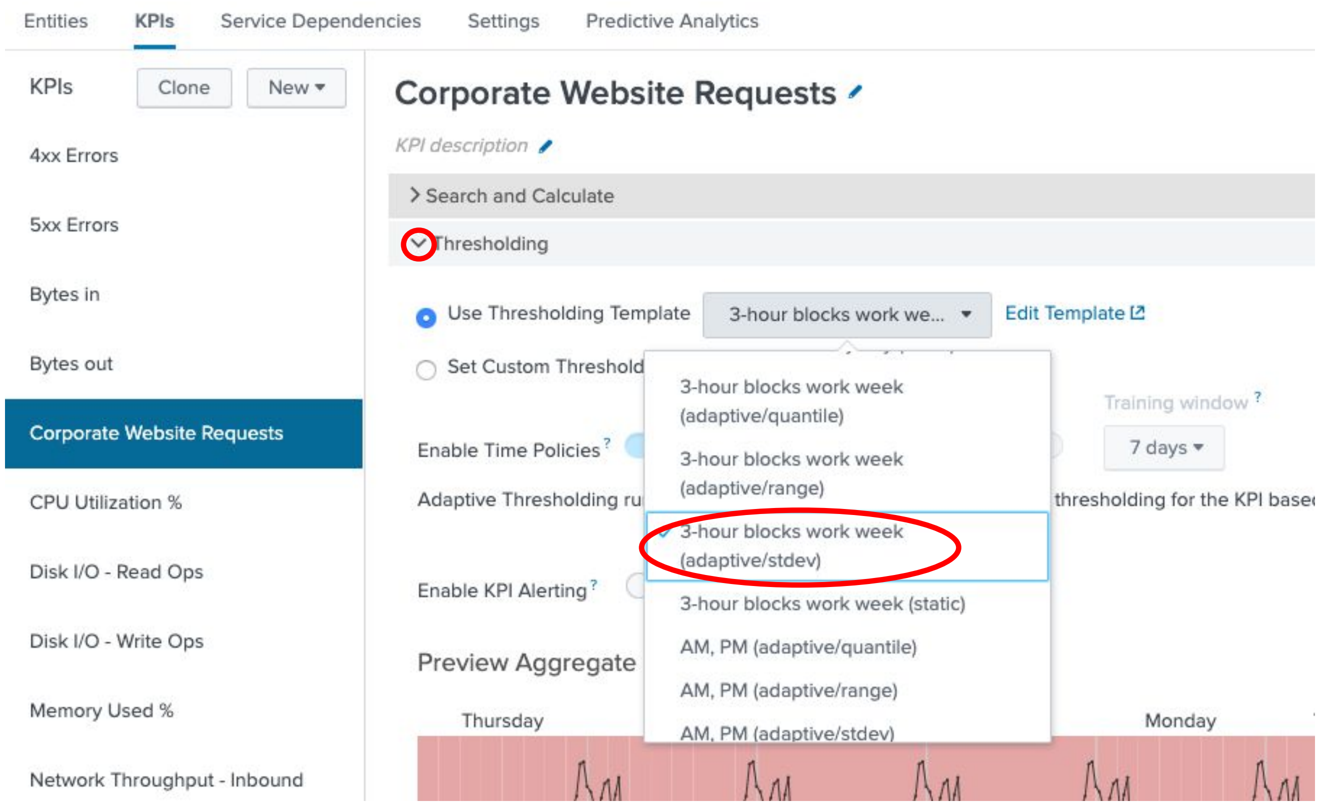

- Click ‘Apply’ button

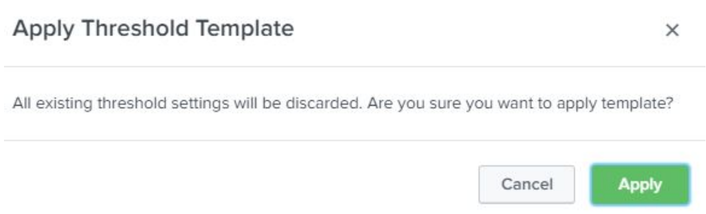

Machine Learning Lab

The built-in machine learning has configured thresholds, this is broken into 1-hour time ranges. However we want to use historical data to apply some adaptive thresholding.

- Open the ‘Configuration Thresholds for Time Policies’ box
- Review different times
  - Choose ‘*Weekdays, 9am-12am*’
- Click ‘*Apply Adaptive Thresholding*’ button
  - Wait 30 Seconds
- Notice threshold
- Click ‘Save’ button

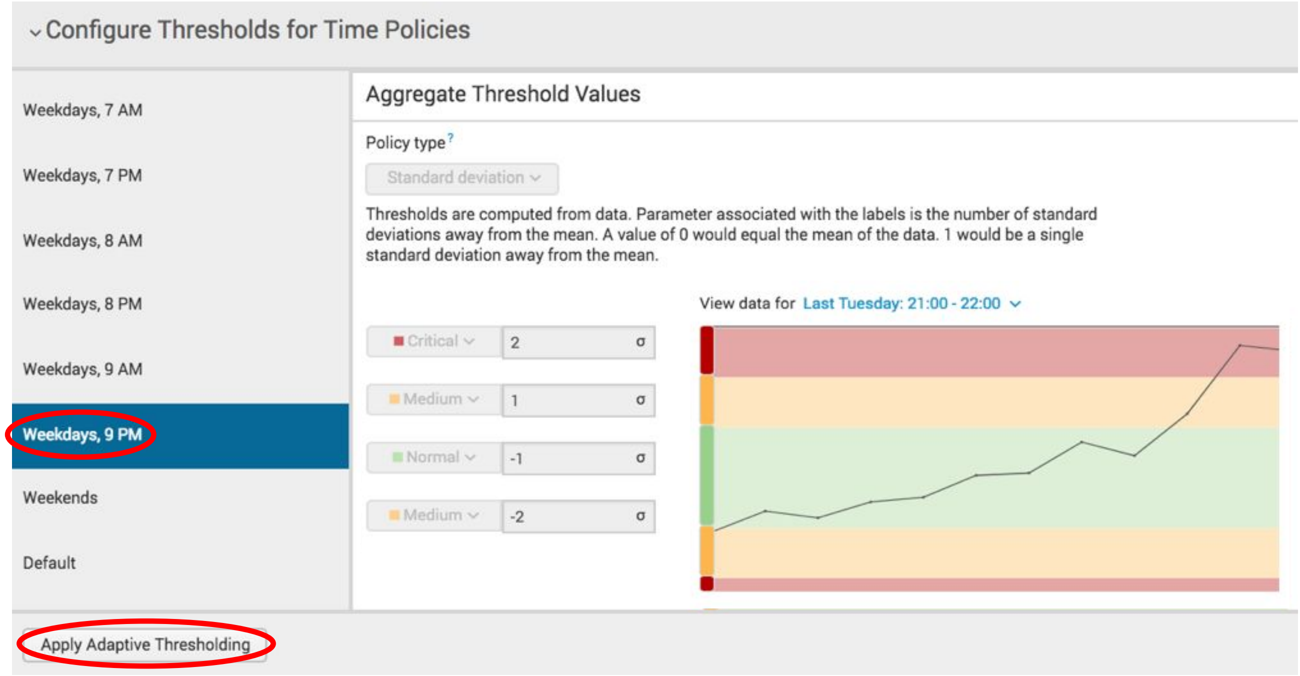

Machine Learning Lab

Internal SLAs do not apply on weekends, tracking is still necessary, but management wants deliberately higher values.

We will now modify the new adaptive thresholds to increase the weekend ranges, this will result in a custom template.

- Click on '*Thresholding*' arrow
- Review the preview window
  - Note weekend range
- Click '*Set Custom Threshold*'

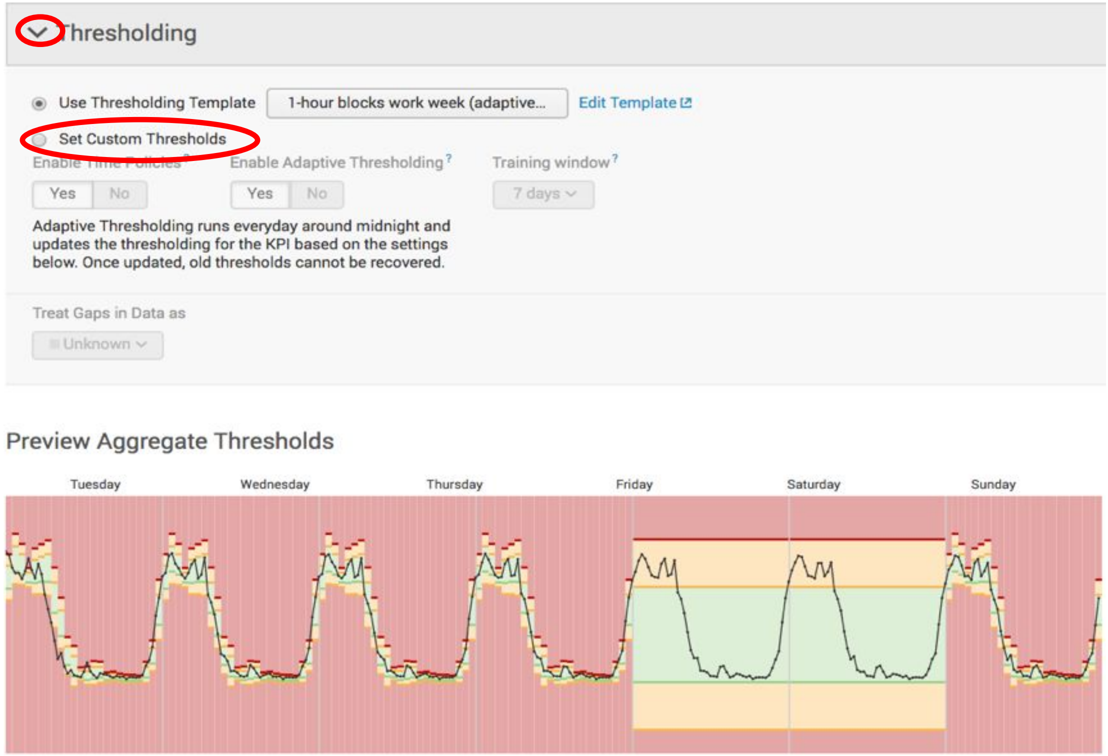

Machine Learning Lab

We will now modify the new adaptive thresholds to increase the weekend ranges, this will result in a custom template.

- Expand '*Configure Thresholds for Time Policies*'
- Click '*Weekends*'
- Move all sliders as show
- Click '*Apply Adaptive Thresholding*'

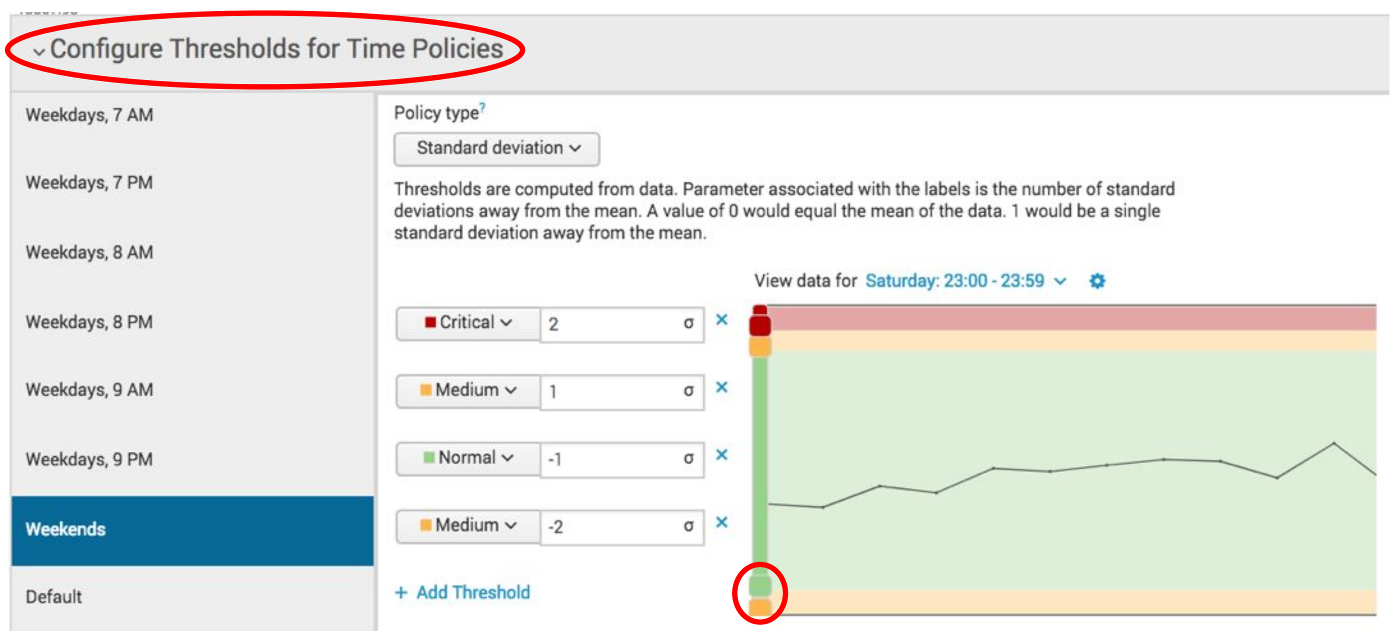

Machine Learning Lab

The preview aggregate threshold window now shows that during the weekend we will not receive an warning alert for the website response time KPI.

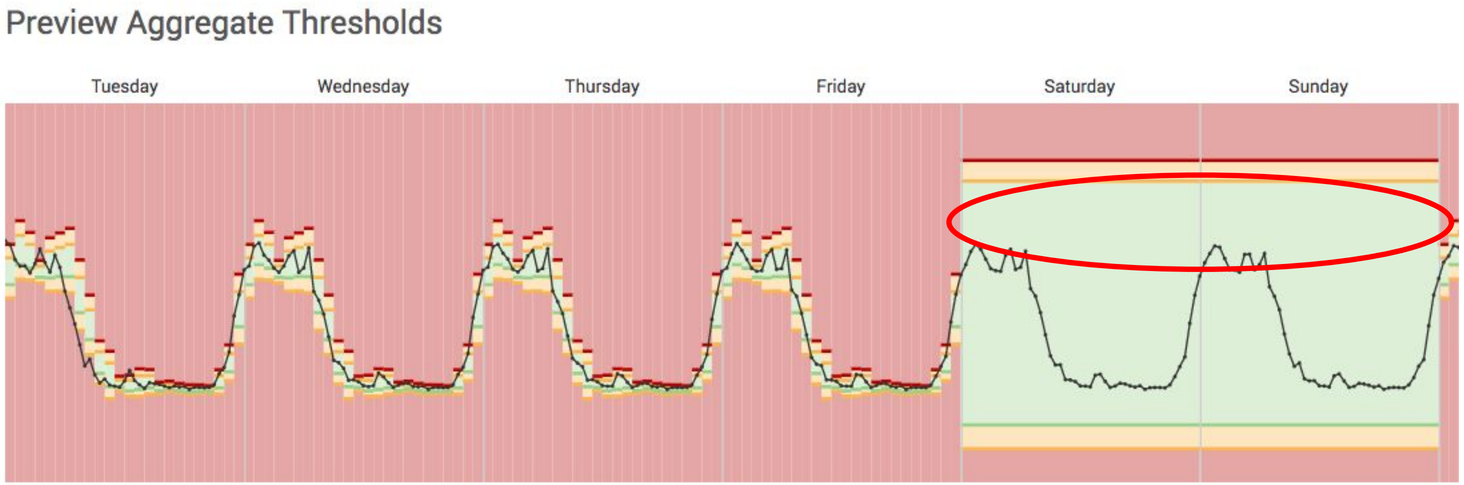

- Click ‘Save’ button

Deep Dive

Deep Dive Use Case

- Typically when organisations have outages they create a war room to identify the root cause as quickly as possible, this involves bringing together many business & technical stake holders at great expense.

- The deep dive capabilities within ITSI brings together multiple data sources into a single visualization. The correlation of data streams enable quick identification of root cause and effect on the business.

- In this lab we will build a deep dive visualization for the new On-Prem Database service, this will bring business and technical KPIs together with raw event data.

- *Extra – Once this lab is completed review the comparisons options.*

Deep Dive Lab

- Navigate to the '*Default Service Analyzer*' view
- Toggle to '*Tree View*' mode
- Click '*On-Prem Database*' service
- Click '*Open all in deep dive*'

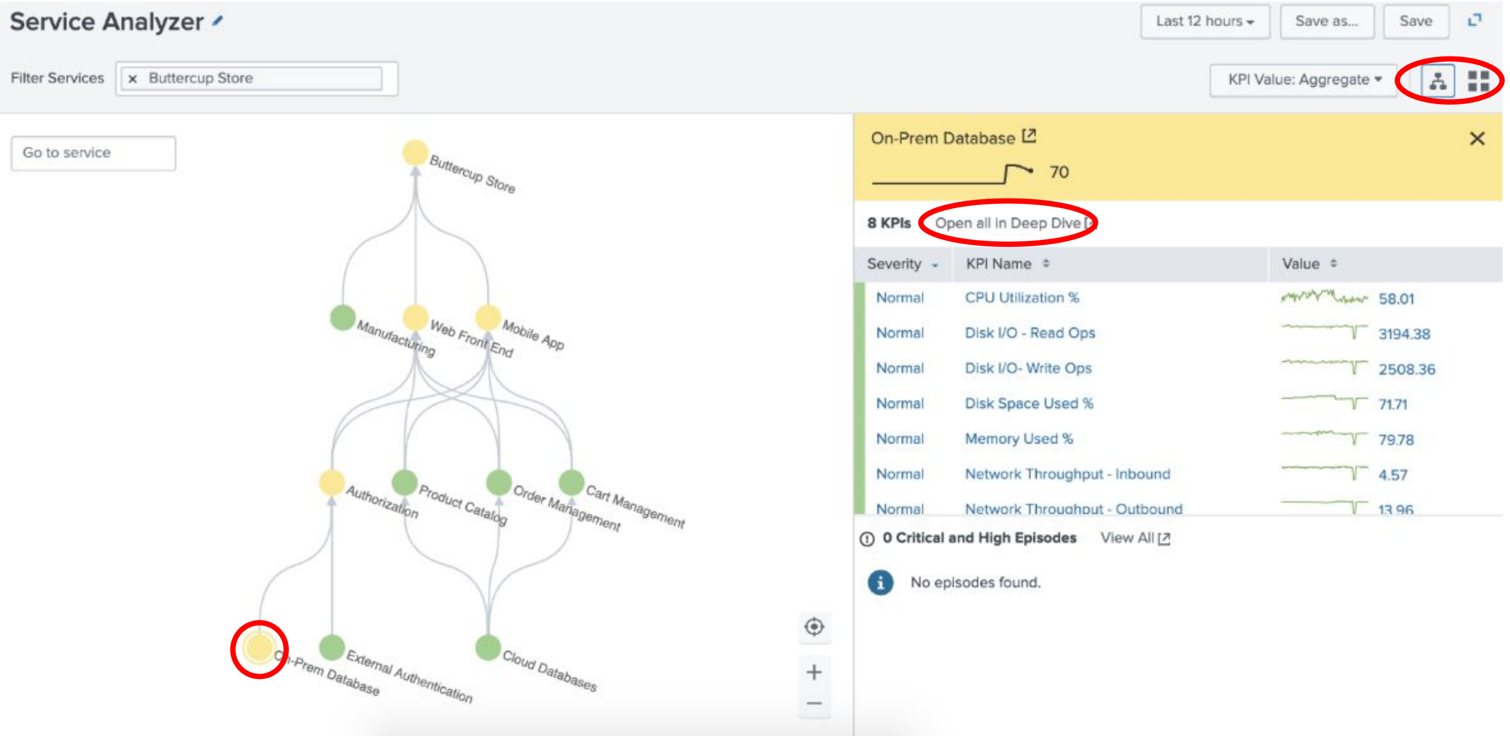

Deep Dive Lab

This deep dive view is used to bring all the relevant data to run an efficient war room, we can add/remove swim lanes to make the visualization even more useful.

- Select the three swim lanes
- Bulk Actions > Delete

Deep Dive Lab

To understand the impact we need to add some business KPI to this deep dive, this will speed up investigations and diagnosis.

- Either navigate up the service tree or change focus to '*Buttercup Store*'
- Click + on the following KPIs;
  - Revenue
  - Revenue per Order
  - Successful Checkouts
- Move lanes as per image

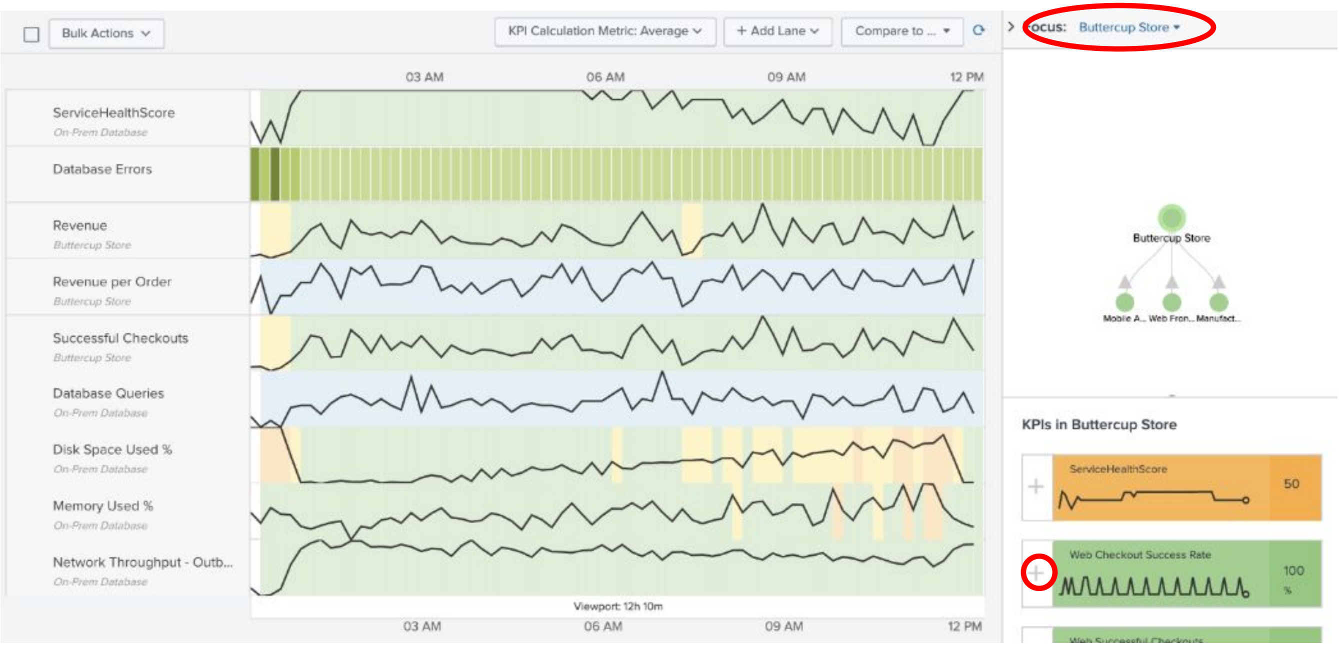

Deep Dive Lab

To enable investigation into anomalous activity in your KPIs we can drill down on KPIs to gain deeper insights.

- Select ‘Disk Space used %’
- Select the COG icon next to Disk Space used %
- Select Lane Overlay options
- Select Enable Overlays ‘Yes’

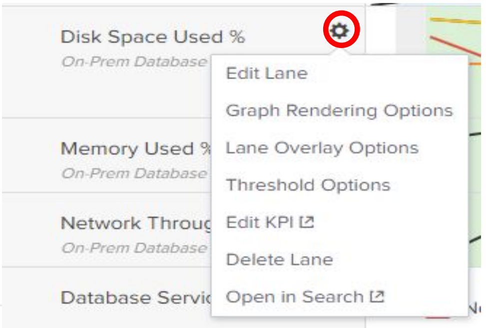

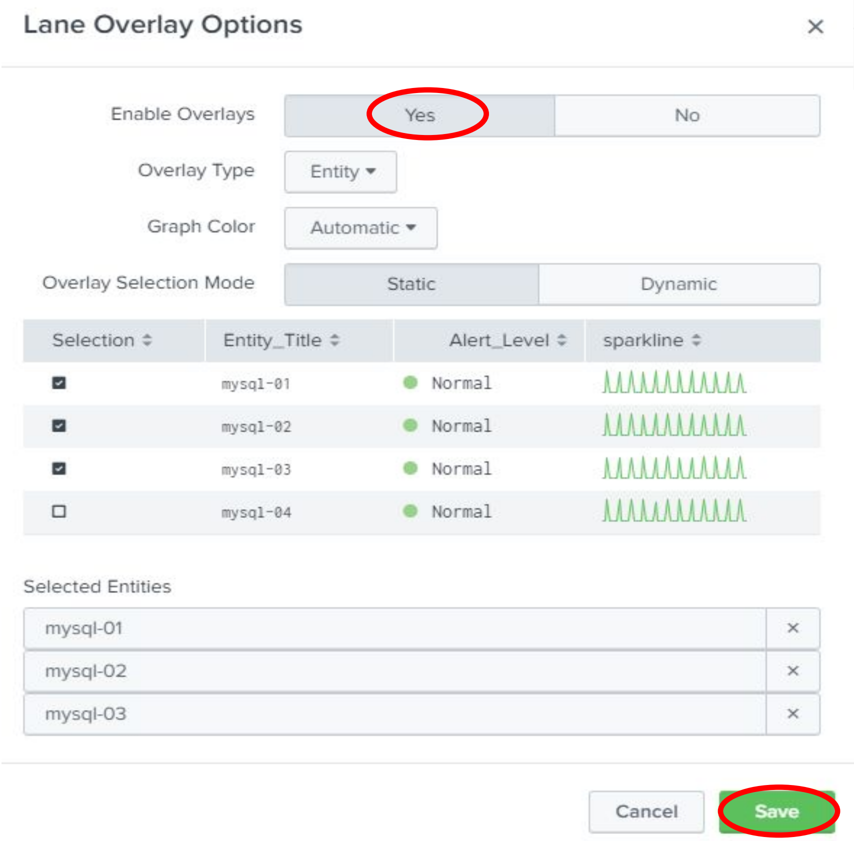

Deep Dive Lab

- Hover over one of the lines from ‘Disk Space Use %’. Notice the drill down possibility to the Splunk App for Infrastructure

Deep Dive Lab

- Select 'Add Lane'
- Click 'Add Event Lane'
- Input 'Database Service Errors' Title
- Event Search:
  - *'sourcetype="stream:mysql" status>200'*
- Click 'Create Lane'

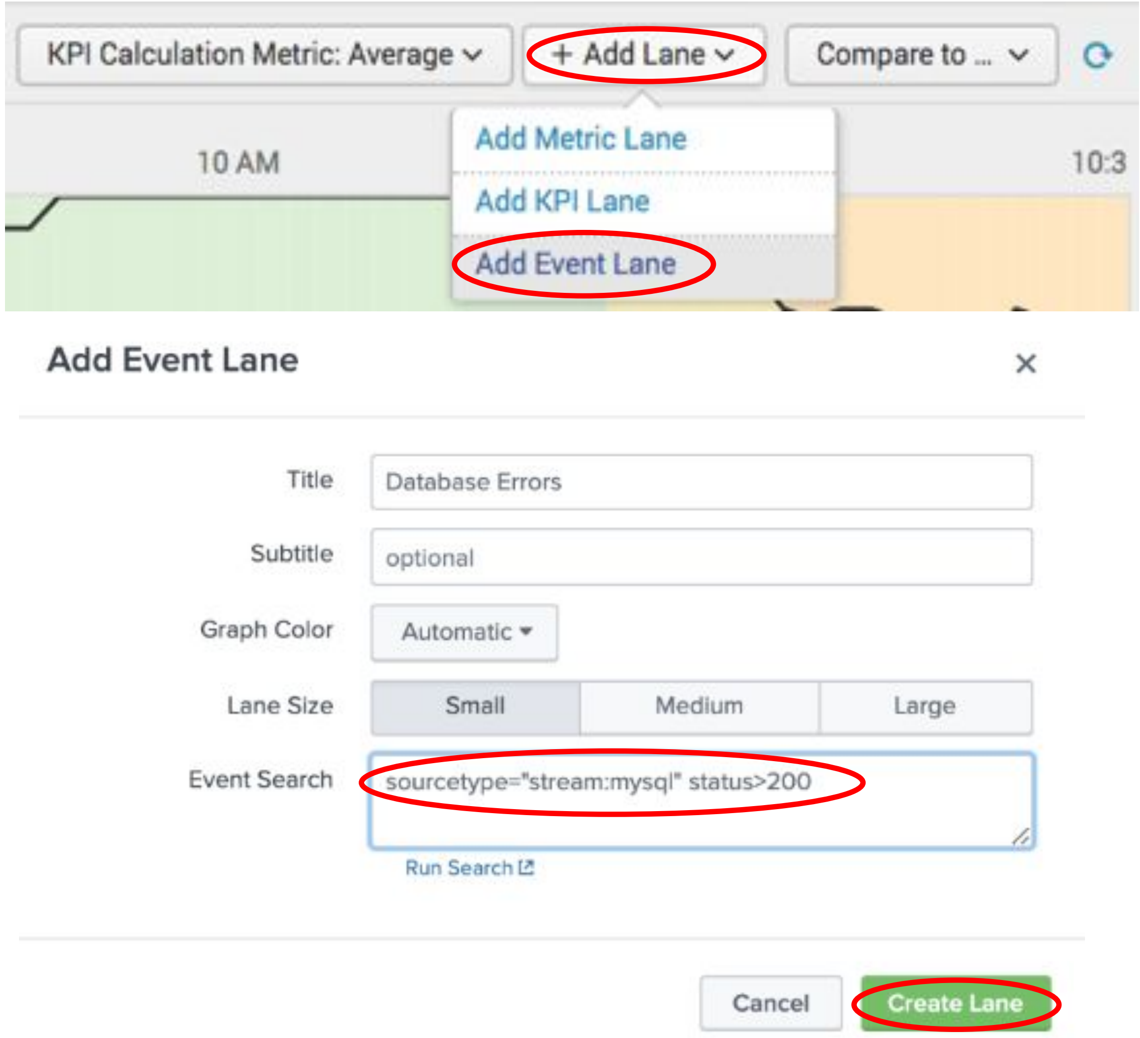

Deep Dive Lab

- Extra
  - Investigate blue event lane
  - Compare to yesterday

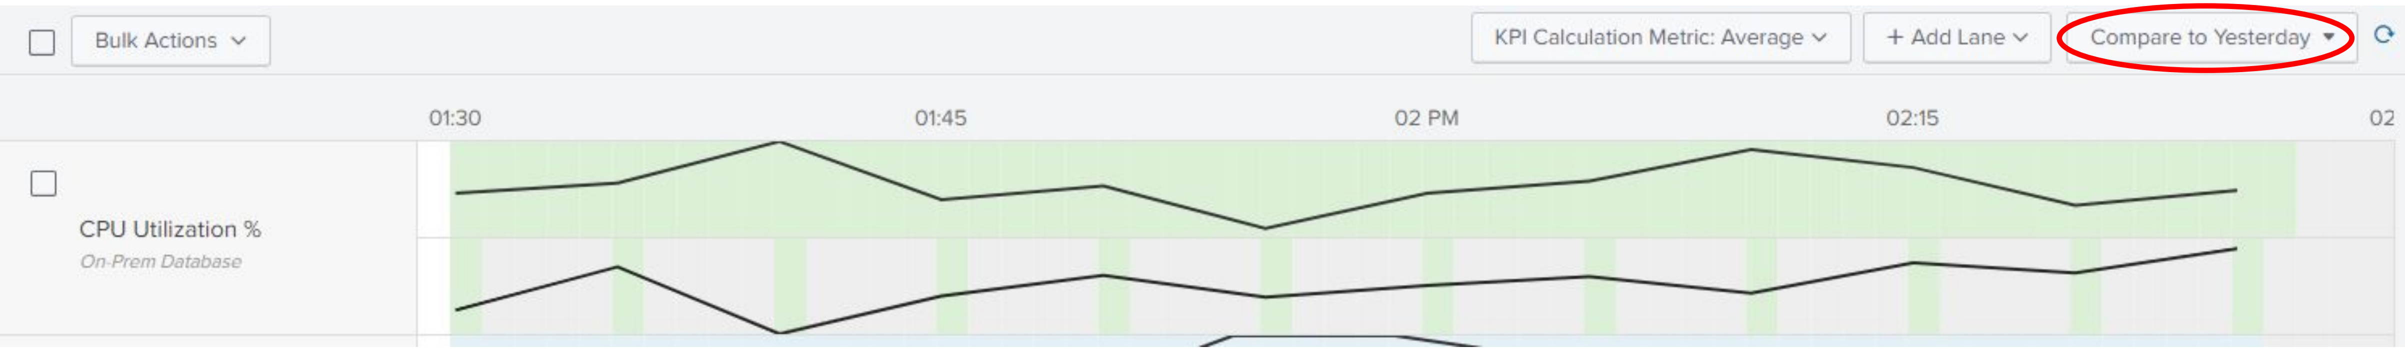

- *Note: adjust the Time-Picker to only show the last 60 minutes.*

Deep Dive Lab

Once you have finished investigating the Deep Dive dashboard

- make sure you click ‘*Save As*’
- And save ‘*Database Deep Dive – <Your Initials>*’

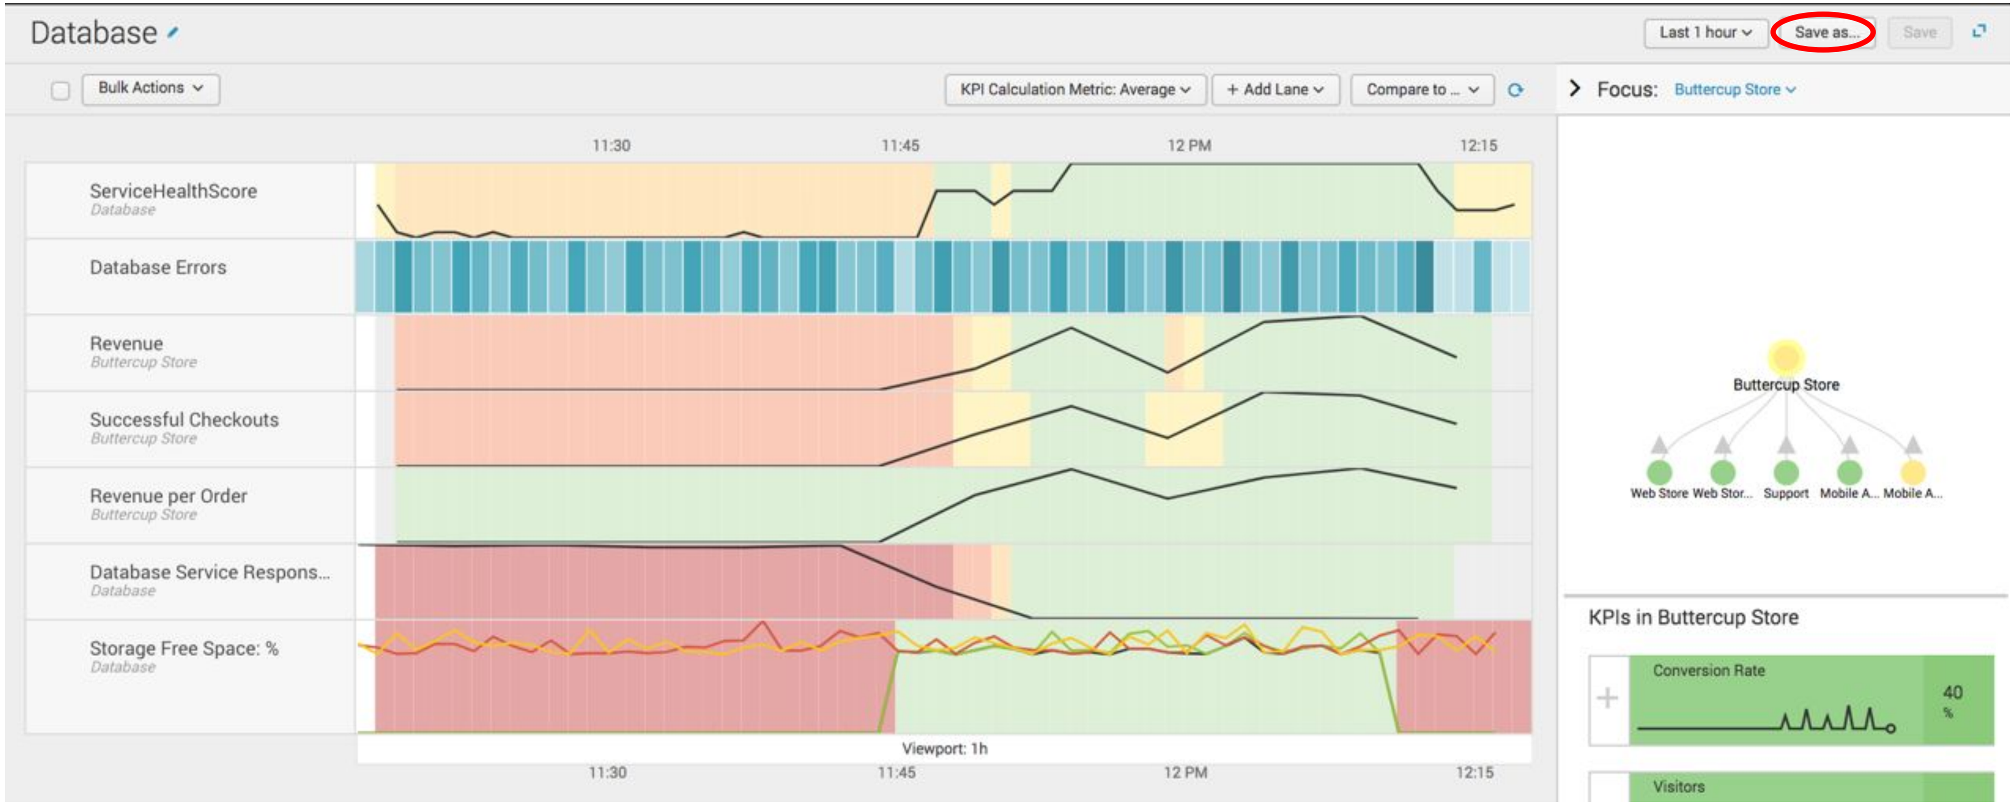

Glass Tables

When to use Glass Tables

- Executive overview and business metrics
- Highly complex and valuable services
- Services that fail often or result in War Rooms
- Services that have recurring outages
- Main use case to stream-line root cause investigation
- When you want to link to Splunk Enterprise or other tools

- *Note : It is best practice to import a background with most of your graphical design, and then drag KPIs onto the canvas.*


Current Operational Health
99
LPAR Availability
1
LPAR Near Capacity
24
Available LPARS
63.5
CICS Txn Per Sec (Thousands)
99 0.1%
Database Avail
42 1.2%
Help Desk Tickets
1,112
Revenue Impact
Operational Performance
IBM Z
98
890
82
27
43
117
63
8
382
77
4921
1324
182
67
78
27
SERVICES
STATUS
Personal Banking
Credit
Home Loans
Insurance
Financial Planning
Mobile
MQ Bus
Batch Processing
DB2
CICS
QA 01
DEV 01
splunk >

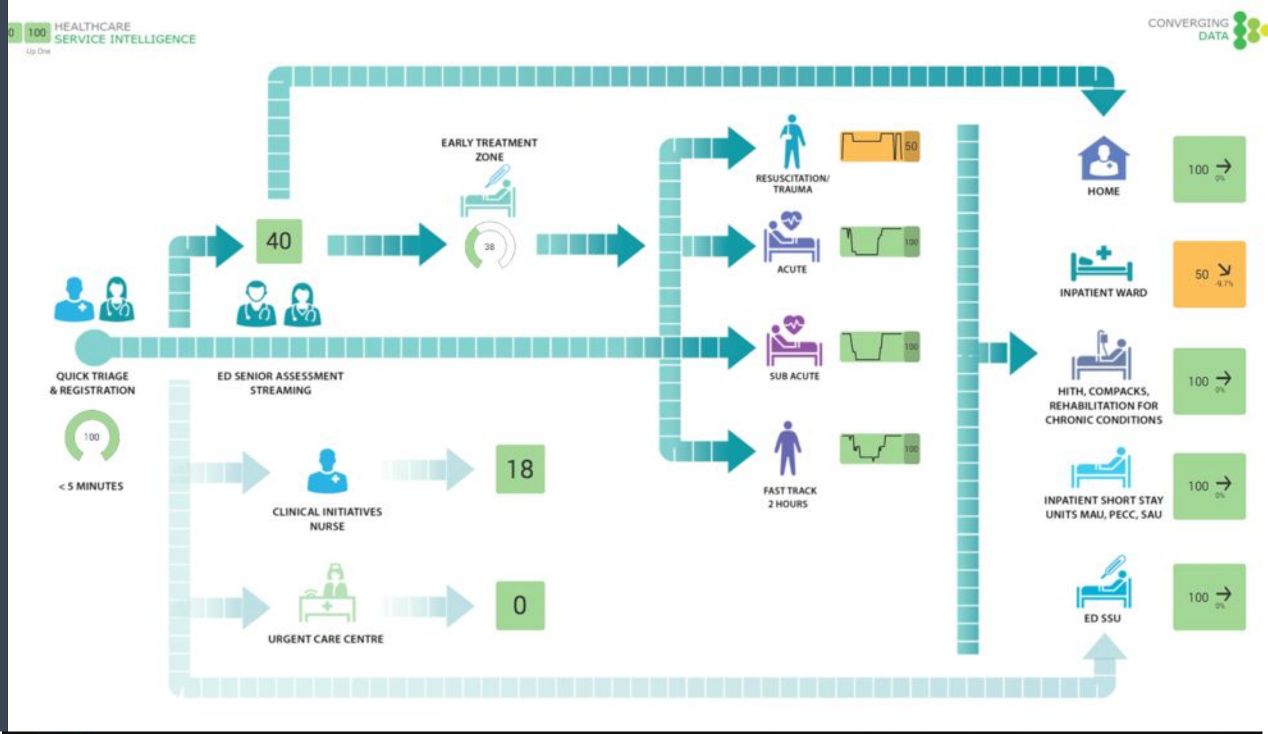
HEALTHCARE SERVICE INTELLIGENCE
CONVERGING DATA
QUICK TRIAGE & REGISTRATION
< 5 MINUTES
ED SENIOR ASSESSMENT STREAMING
40
EARLY TREATMENT ZONE
CLINICAL INITIATIVES NURSE
18
URGENT CARE CENTRE
0
RESUSCITATION/ TRAUMA
ACUTE
SUB ACUTE
FAST TRACK 2 HOURS
HOME
INPATIENT WARD
HITH, COMPACKS, REHABILITATION FOR CHRONIC CONDITIONS
INPATIENT SHORT STAY UNITS MAU, PECC, SAU
ED SSU
100
50

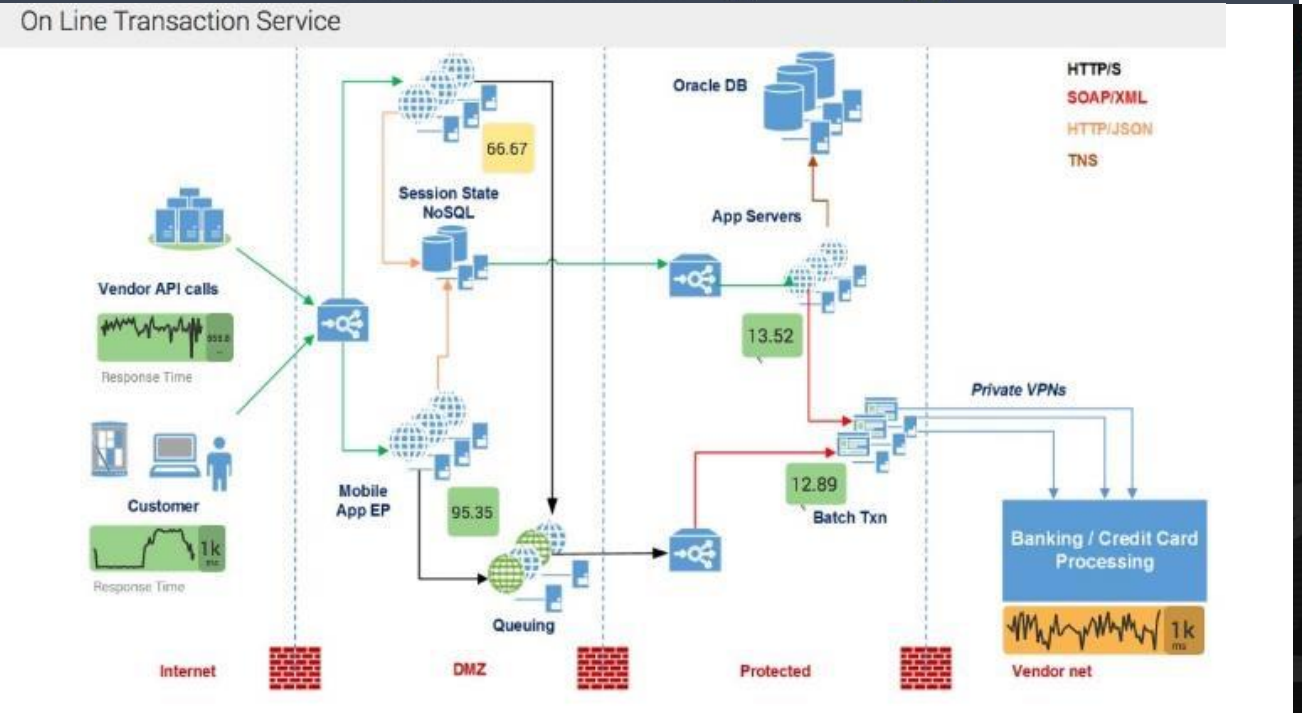
On Line Transaction Service
HTTP/S
SOAP/XML
HTTP/JSON
TNS
Oracle DB
66.67
Session State NoSQL
App Servers
Vendor API calls
Response Time
13.52
Customer
Response Time
1k
Mobile App EP
95.35
12.89
Batch Txn
Private VPNs
Queuing
Banking / Credit Card Processing
1k
Internet
DMZ
Protected
Vendor net


NATIONAL PARK SERVICE
National Park Visitor Services
Park Services Goals
Parks Status
91
Current Visitor Level (Thousands)
97.1
Customer Sat %
22.5
Current Employees (Thousands)
231
Current Volunteers (Thousands)
12 1.2%
Call Wait Time (s)
Park Services
2.7 1.2%
Consumer Complaints
510
Total Accidents
1,112
Total Cost USD$
82
Budget Trend
Customer Experience
68
Social Media NPS
98
Web Site Avail
94
Ticket Purchasing
Infrastructure
88
Roads / Paths
87
Lodging
47
Food / Water
Resources
Recyling Program
YTD Goal
Previous Year
YTD Trend
H2O Usage
YTD
YTD Forecast
Prev Year
Power Usage
425
1,365
Air Quality
123.7
Avg Park Pollution Content
splunk >

Glass Table Use Case

- The business leaders would like a high-level dashboard showing the key functions and services of the organisation.

- The objective of this exercise is to complete the existing IT operations dashboard (Digital Transaction Flow) with the new On-Prem Database service healthscore, including a drill down to a deep dive.

- Second part of this lab is to create a business focus dashboard;
  - Import your group logo
  - Include several KPIs that only relate to the business
  - Keep technical KPIs to a minimum
  - Link existing dashboard to one of the services

Glass Table Lab

- Select *Glass Table menu*
- *Edit '4. Digital Transaction Flow'*

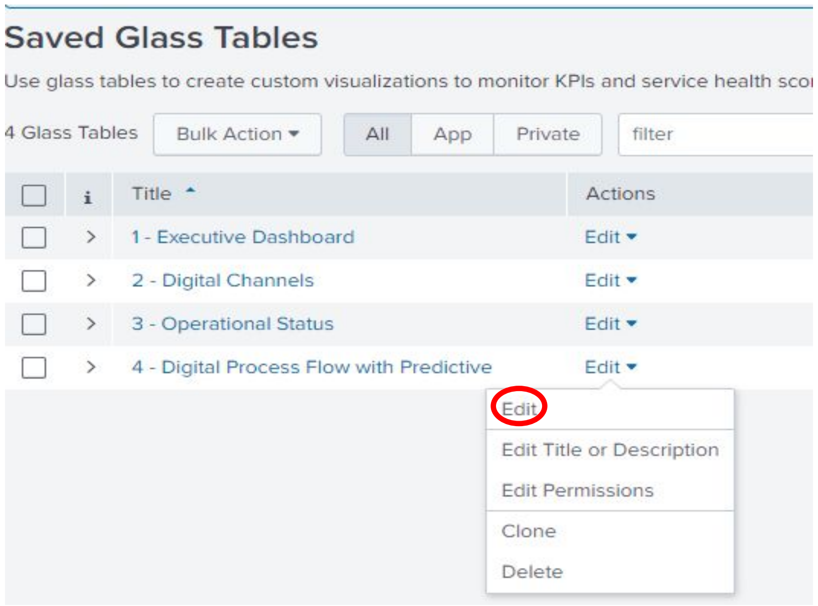

- *Click*  *icon to fit to page*
- Note the tool pallet icons
- *KPIs can be dragged onto canvas*
- *Configuration panel on right*

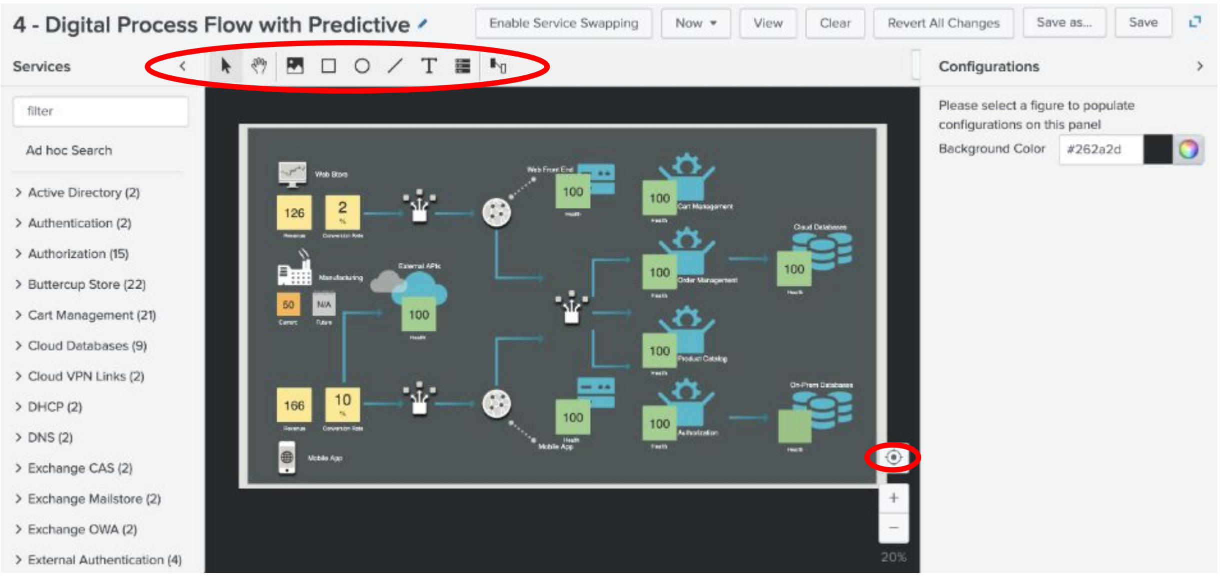

Glass Table Lab

- Expand ‘On-Prem Database’
- Drag ‘ServiceHealthScore onto the canvas
  - Configure attributes;
    - Label Box = Off
      - *Scroll down the dialog box*
    - Drilldown = On
    - Select ‘Saved Deep Dive’
      - Select ‘On-Prem Database DD’
    - Change visualisation type
      - Single Value
  - Click Update

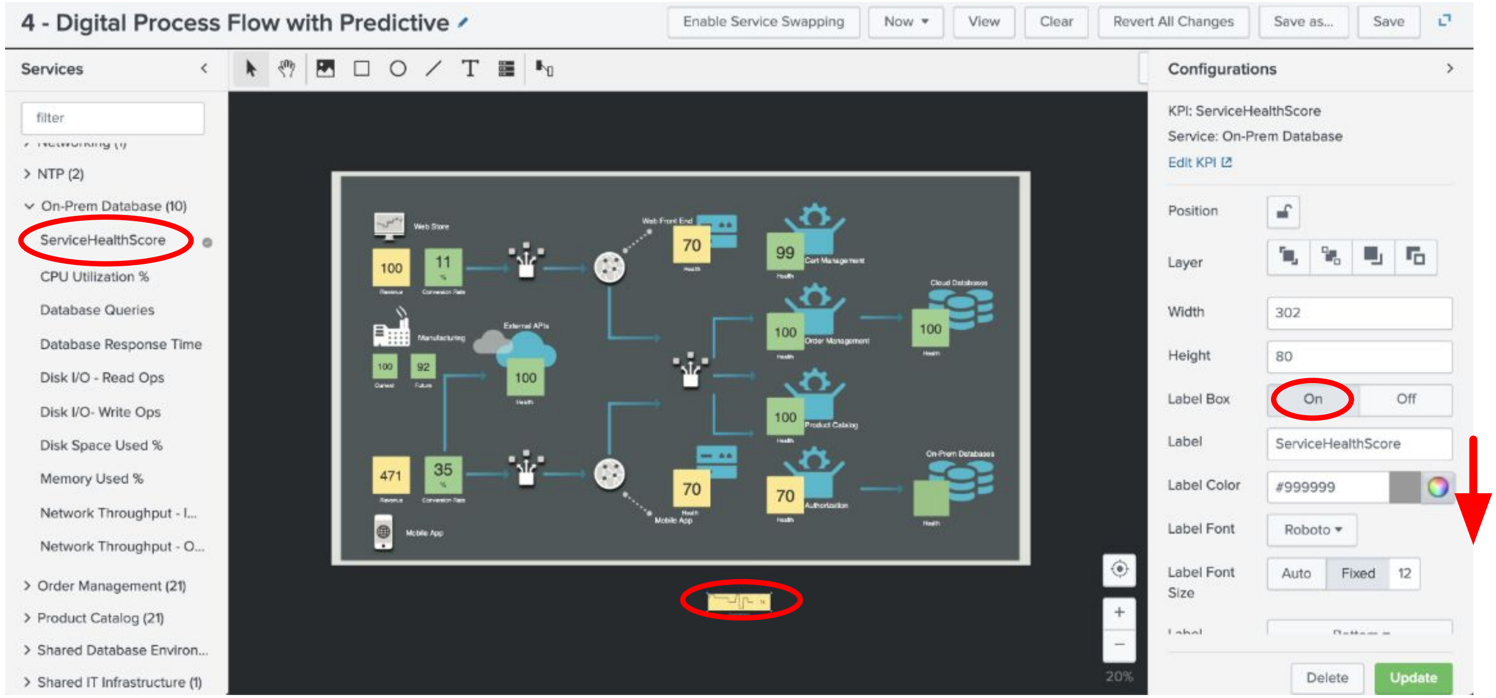

Glass Table Lab

- Move Database healthscore next to the Database icon (green box)

- On the right-hand panel
- Configure attributes;
  - Width = 180
  - Height = 180

- Click Update
- Click Save

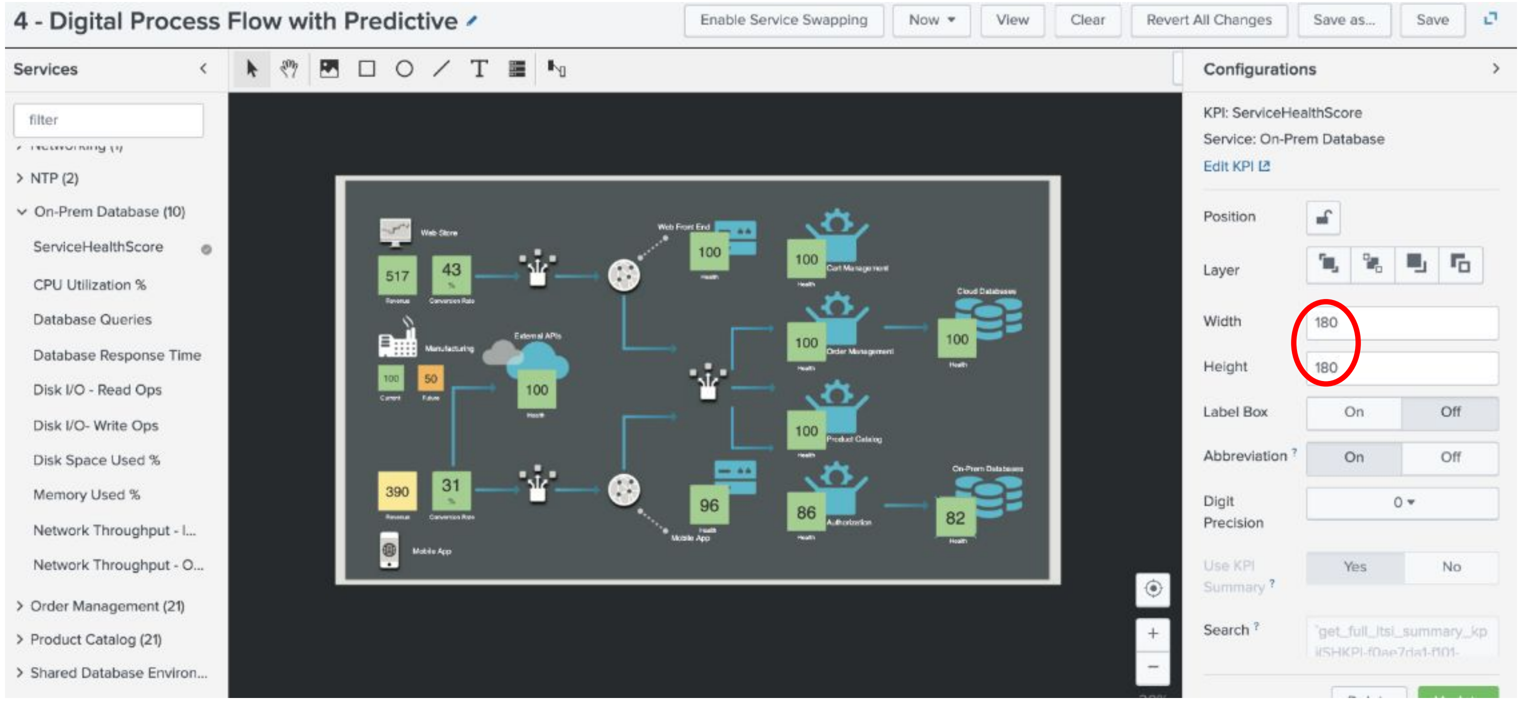

Glass Table Lab (Extra)

- Now it's your turn…you have 10 minutes

- Google and import your company logo
- Drop onto canvas;
  - Couple of icons
  - Link them
  - Business health scores
- Rename to reflect your business, be creative!

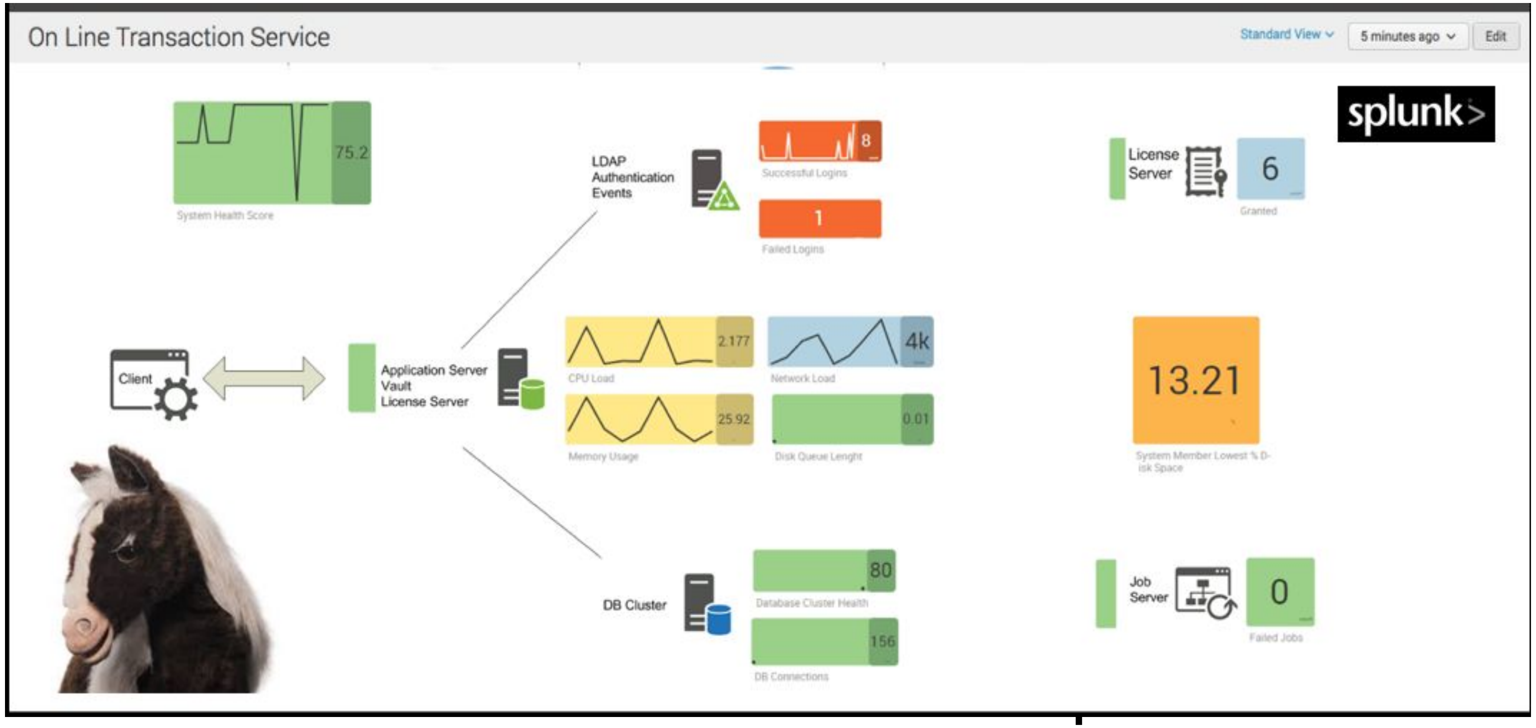

Machine Learning: Predictive Analytics

(or Imminent Outage Prediction)

What’s Imminent Outage Prediction?

- Using historical data from KPIs for a service and some clever ML algorithms, you can sometimes **predict an outage** 20-30 minutes before it happens!
- Works best when a service has 5+ good KPIs and 1+ week of historical data
- The algorithm looks for recognizable/predictable KPI behavior, which comes before the service's aggregate health score changes.
  - For example: before the last outage, CPU usage went up AND garbage collection times increased AND session counts dropped...

Workshop Use Case

- The IT operations team are struggling to resolve issues with the company manufacturing service, typically outages are reported via customers contacting the service desk to complain.
- They would like to use machine learning to predict health score degradation 30 minutes before it causes a service outage.
- The objective of this lab will be to use the ‘**Manufacturing**’ health score to build a predictive algorithm model to predict future issues.

Predictive Analytics Lab

We need to use machine learning to build a model for the Manufacturing service. This model will be used in the second part of this lab.

- Select ‘*configure*’ menu
- Pick ‘*Services*’ item
- Click edit ‘*Manufacturing*’ service
- Select ‘*Edit*’ menu

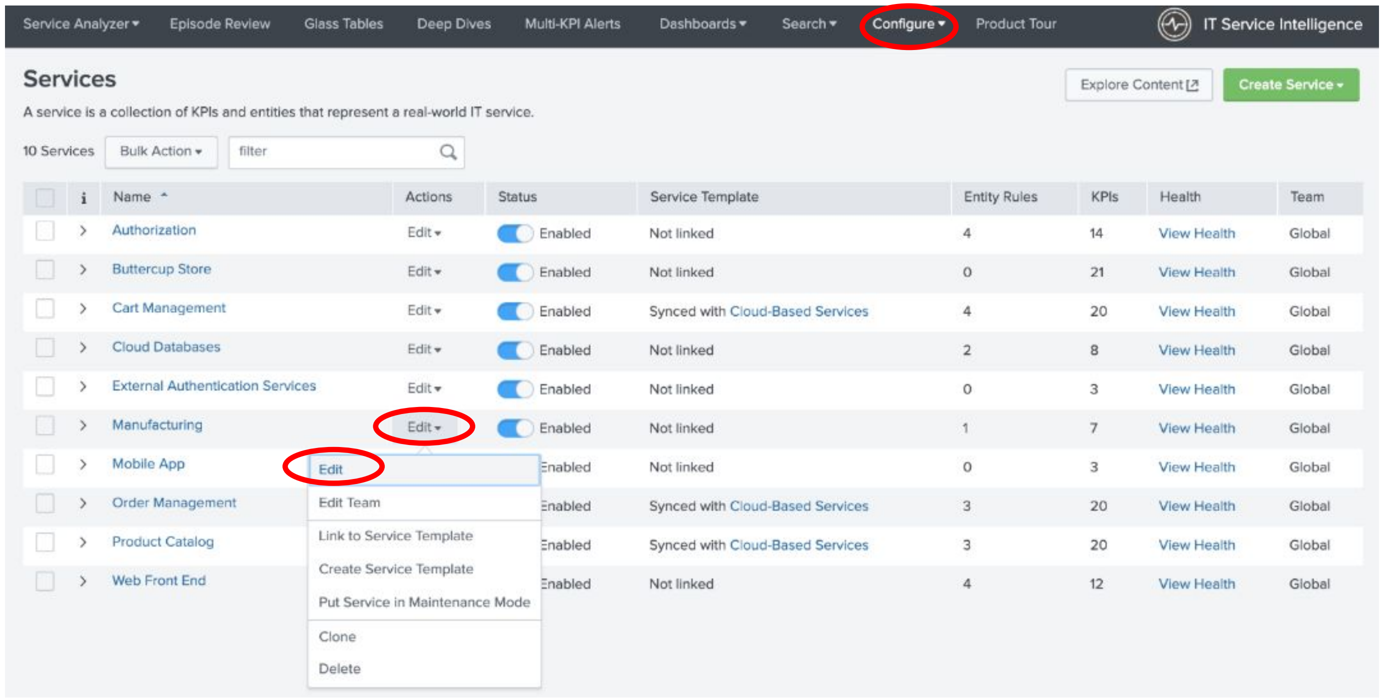

Predictive Analytics Lab

ITSI Predictive Analytics uses machine learning algorithms to predict the future health score of your service. This screen we will train and test different machine learning algorithms to determine which one will give the most accurate prediction.

- Select '*Predictive Analytics*'
- Time = 14 Days
- Algorithm Type = Regression
- Algorithm = Linear Regression
- Click 'Train' button

- Once the model has run, investigate below.
- Click '*Save*'

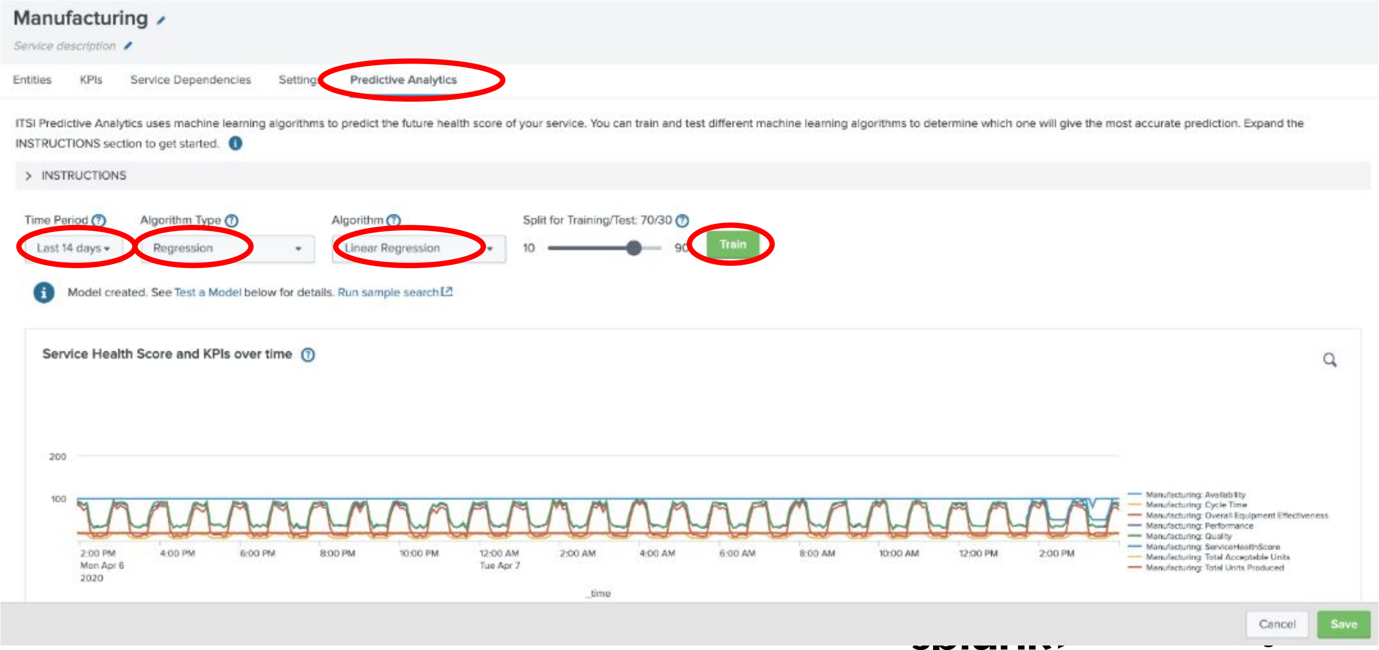

Predictive Analytics Lab

- Repeat the previous steps for the other two algorithms over 14-day period.
- Random Forest Regressor
- Click *’Train*’ button
- Remember ‘*Save*’ button !

- Gradient Boosting Regressor
- Click *’Train*’ button
- Remember ‘*Save*’ button !

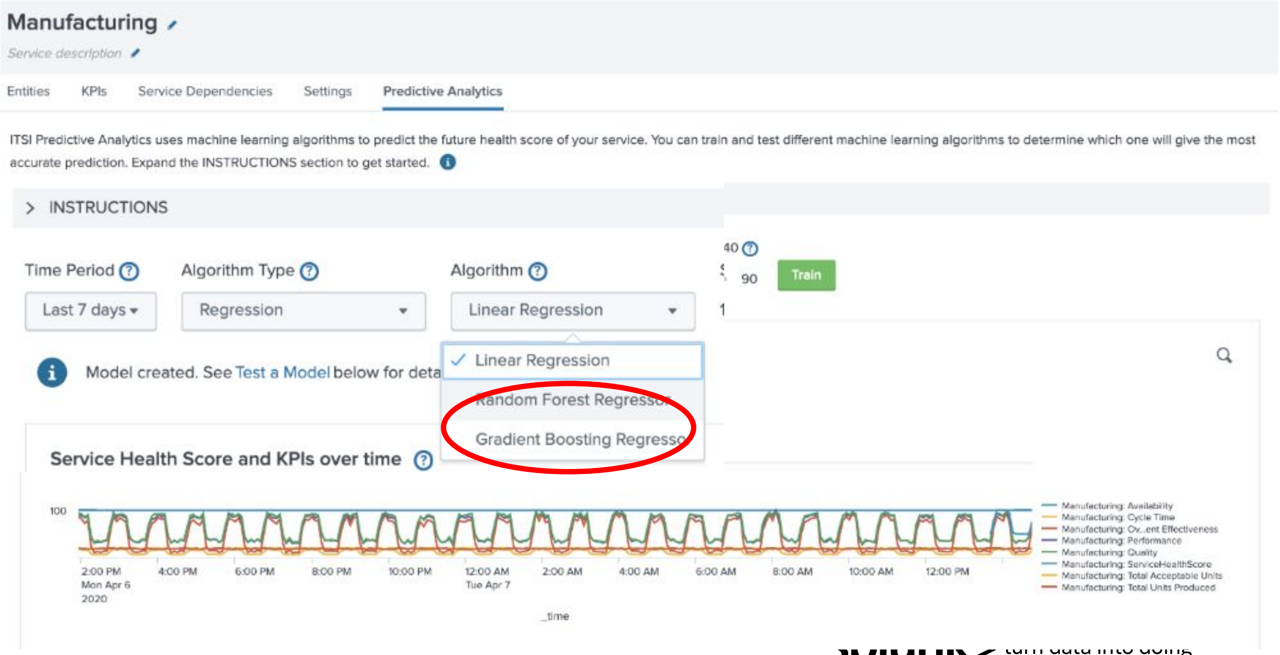

Predictive Analytics Lab

In this lab we are going to review the predictive analytics value for the **Manufacturing** service using the recommended model.

- Select *Dashboards > Predictive Analytics*
- *Select the 'Manufacturing' service*
- *Select the recommended algorithm model*

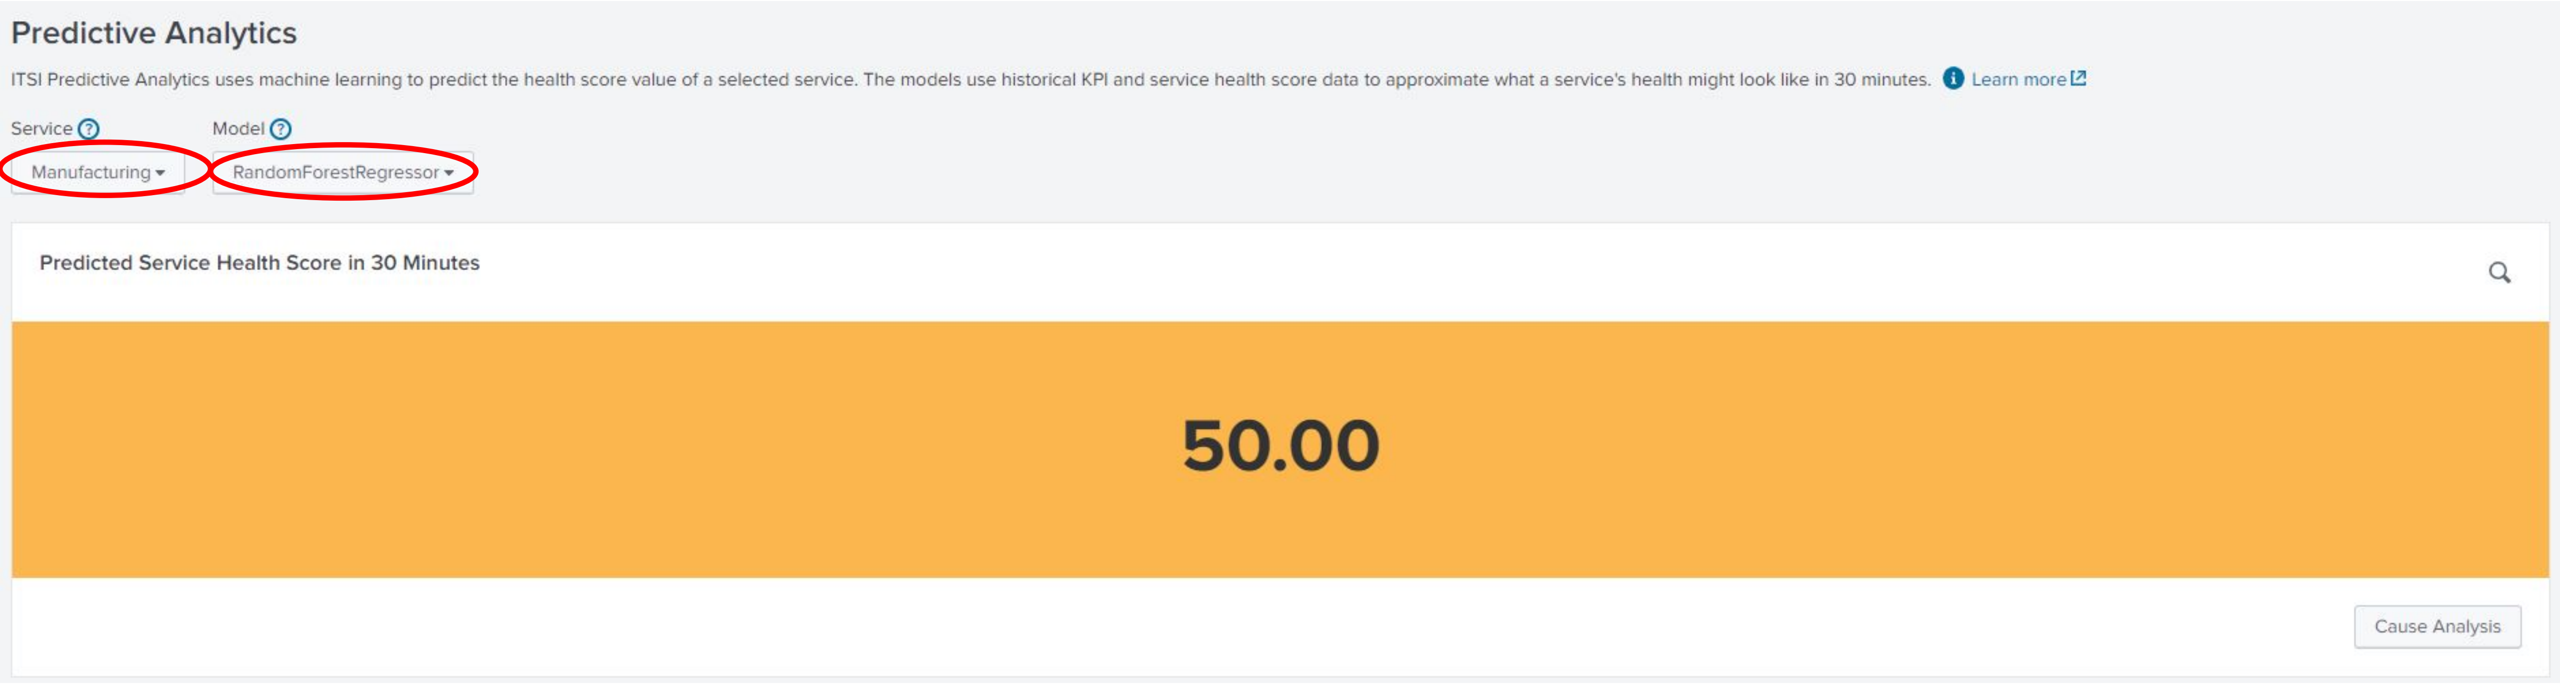

Predictive Analytics Lab

Once we have selected a model ITSI will calculate the future healthscore

- *Click 'Cause Analysis' button to review the suggested KPIs*

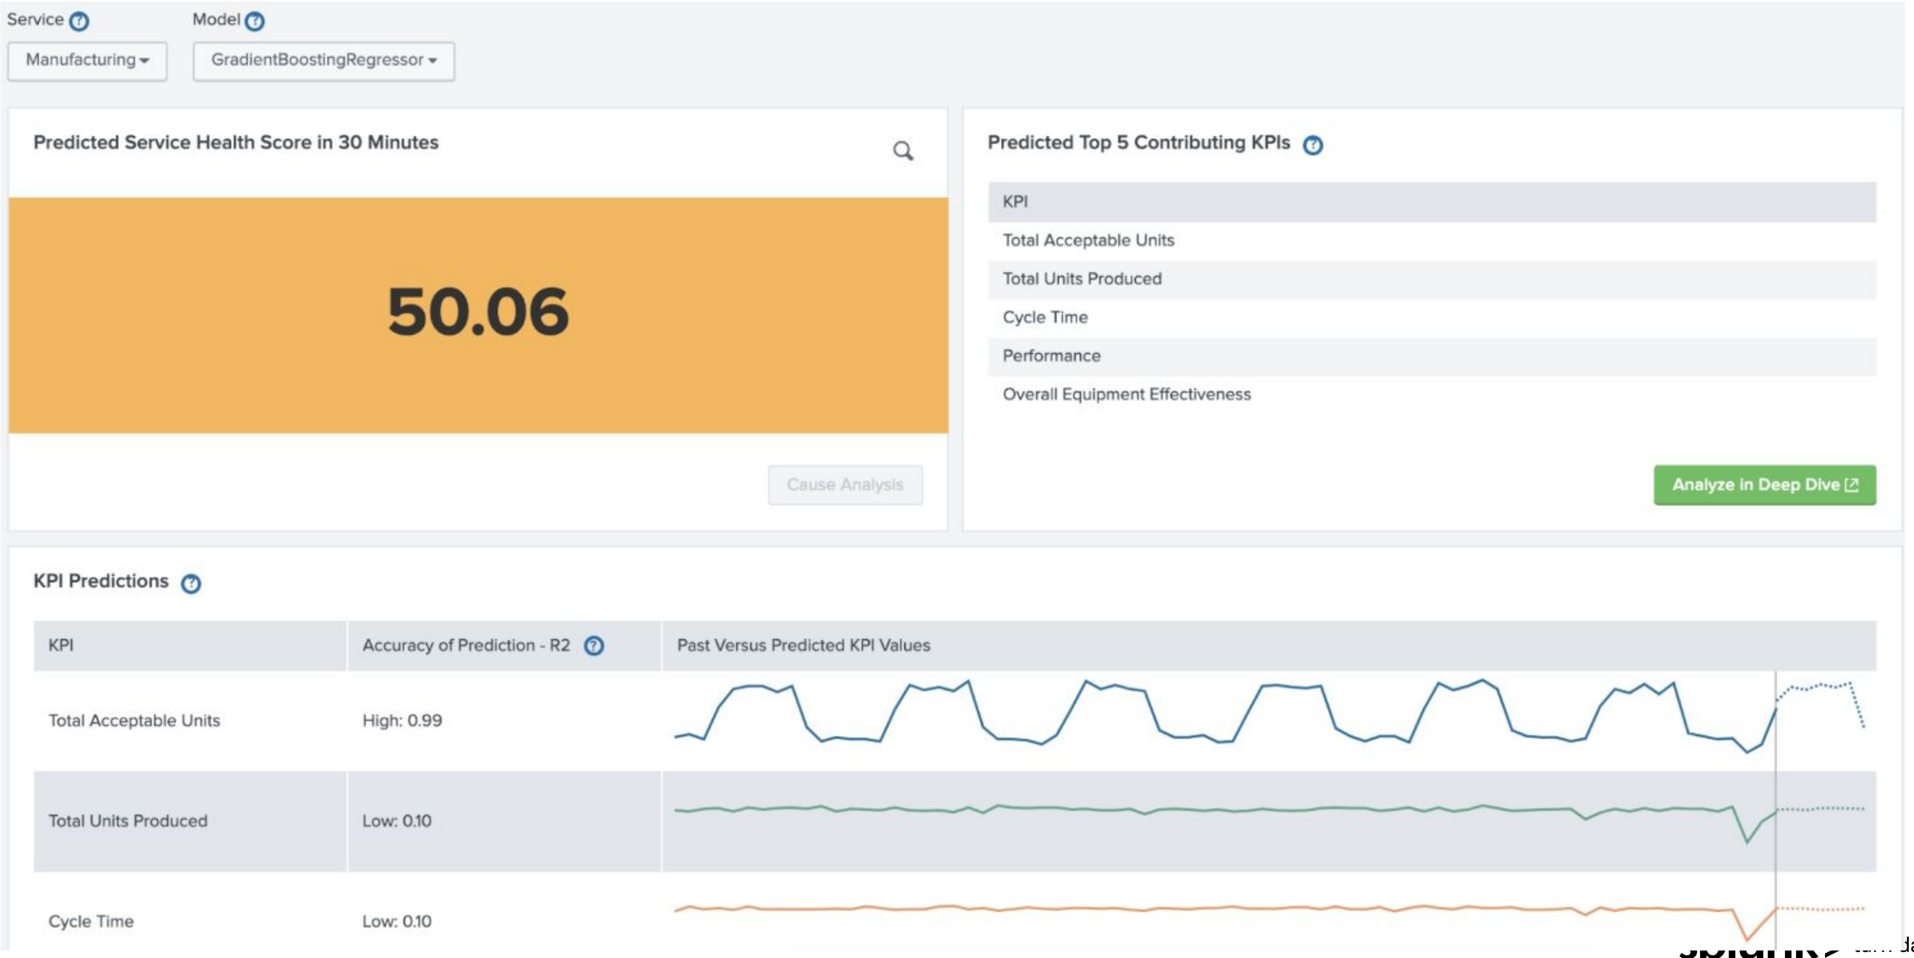

Predictive Analytics Lab

- Click the spyglass to review the SPL
- This SPL is already being used in our glass table

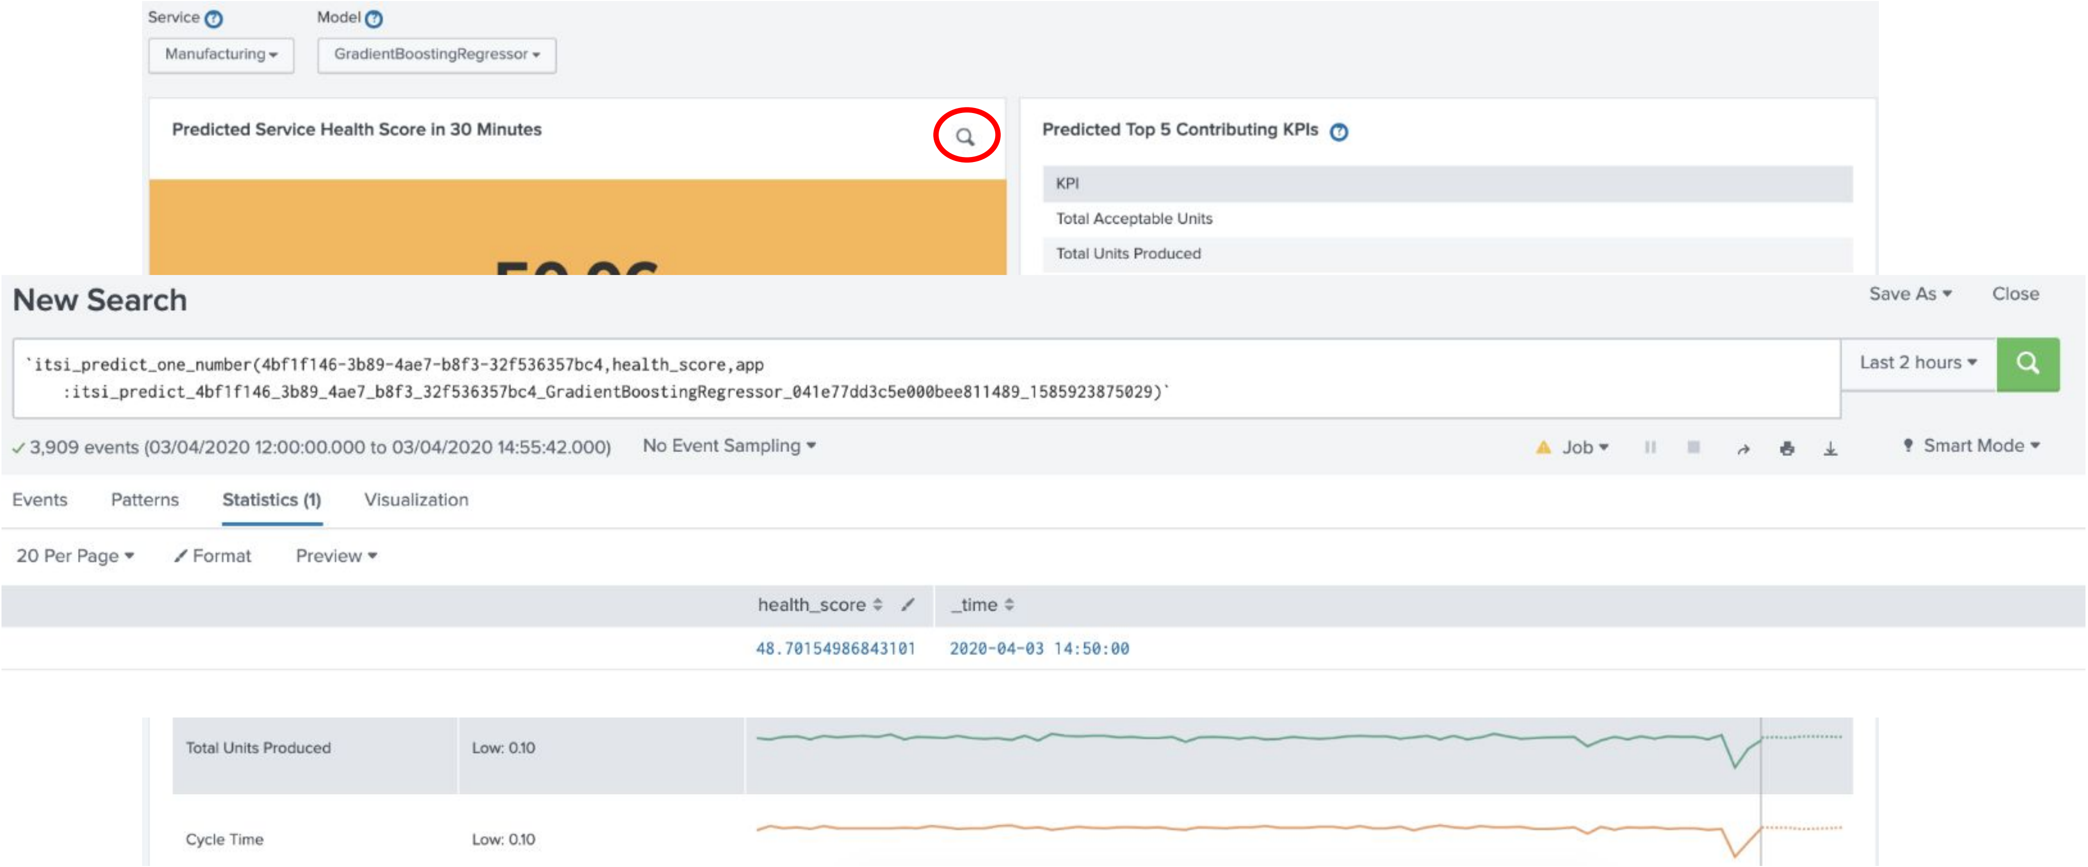

Predictive Analytics Lab

- Select ‘Glass Tables > Digital Transaction Flow with Predictive’
- Click ‘Edit’ button
- Review the ‘Database Future Health-Score attributes’

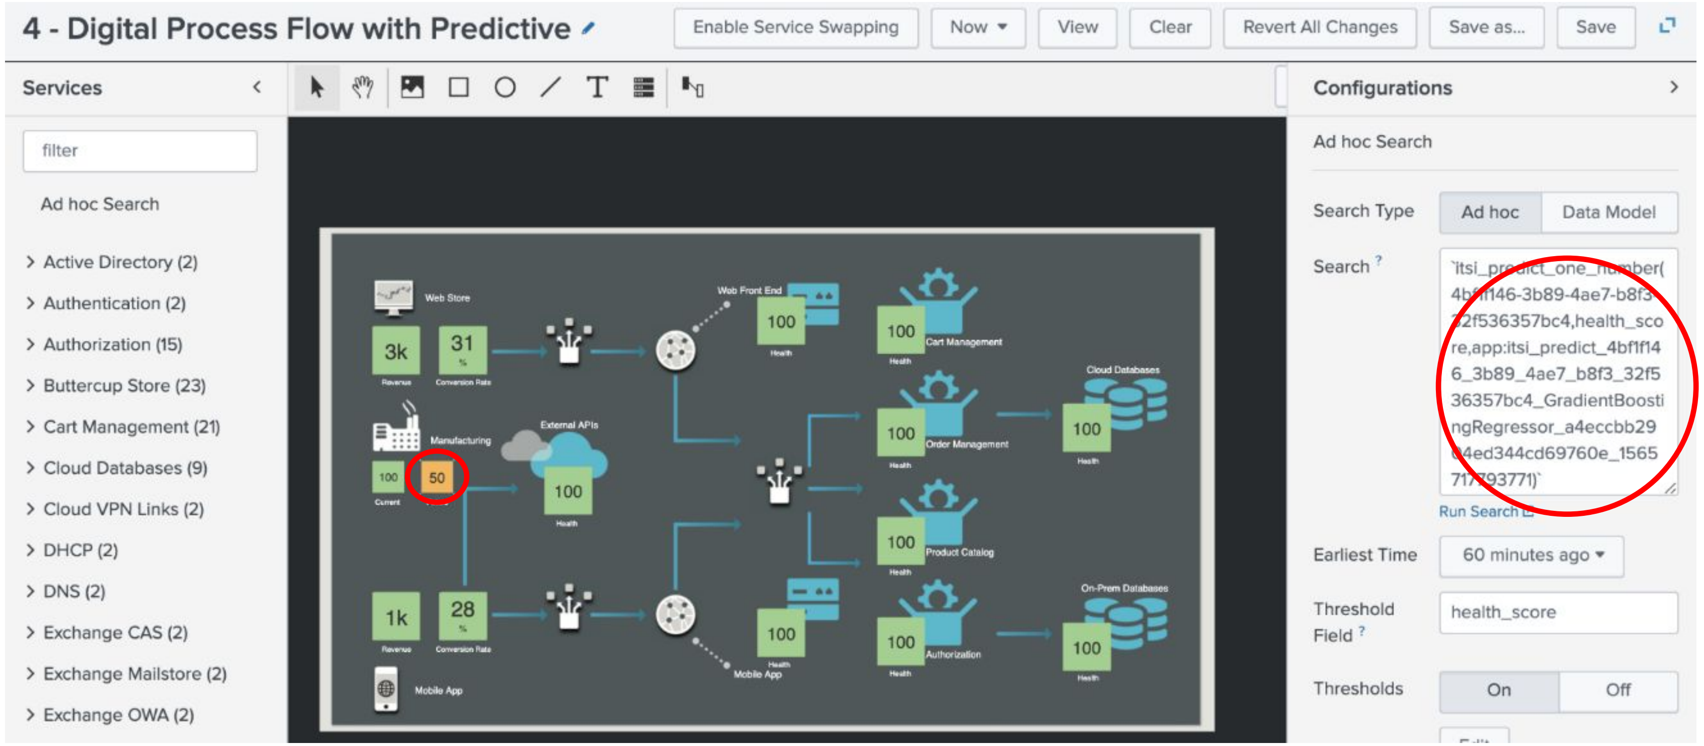

Machine Learning: Event Analytics

Event Analytics Lab

- The database team has expressed frustration with alerts originating from Nagios. The high volume of alerts is leading to alert fatigue and they lack the contextual information necessary to make them actionable.
- While the plan is to consolidate monitoring tools, they have asked if we can provide immediate relief using ITSI to group events together and reduce noise.
- Also, with hundreds of database instances to manage, tracking which alerts are associated to critical systems vs non-critical systems is tribal knowledge for the database team. They love the service tree view in ITSI and asked if it's possible to see the services affected by an alert.
- Interestingly enough, you heard the same challenges expressed by the web team who uses New Relic to monitor the application code.

Event Analytics Lab

- Often, each team or layer of the stack has their own monitoring tools producing their own alerts. Each tool creating a silo of information that other teams don't get insight.

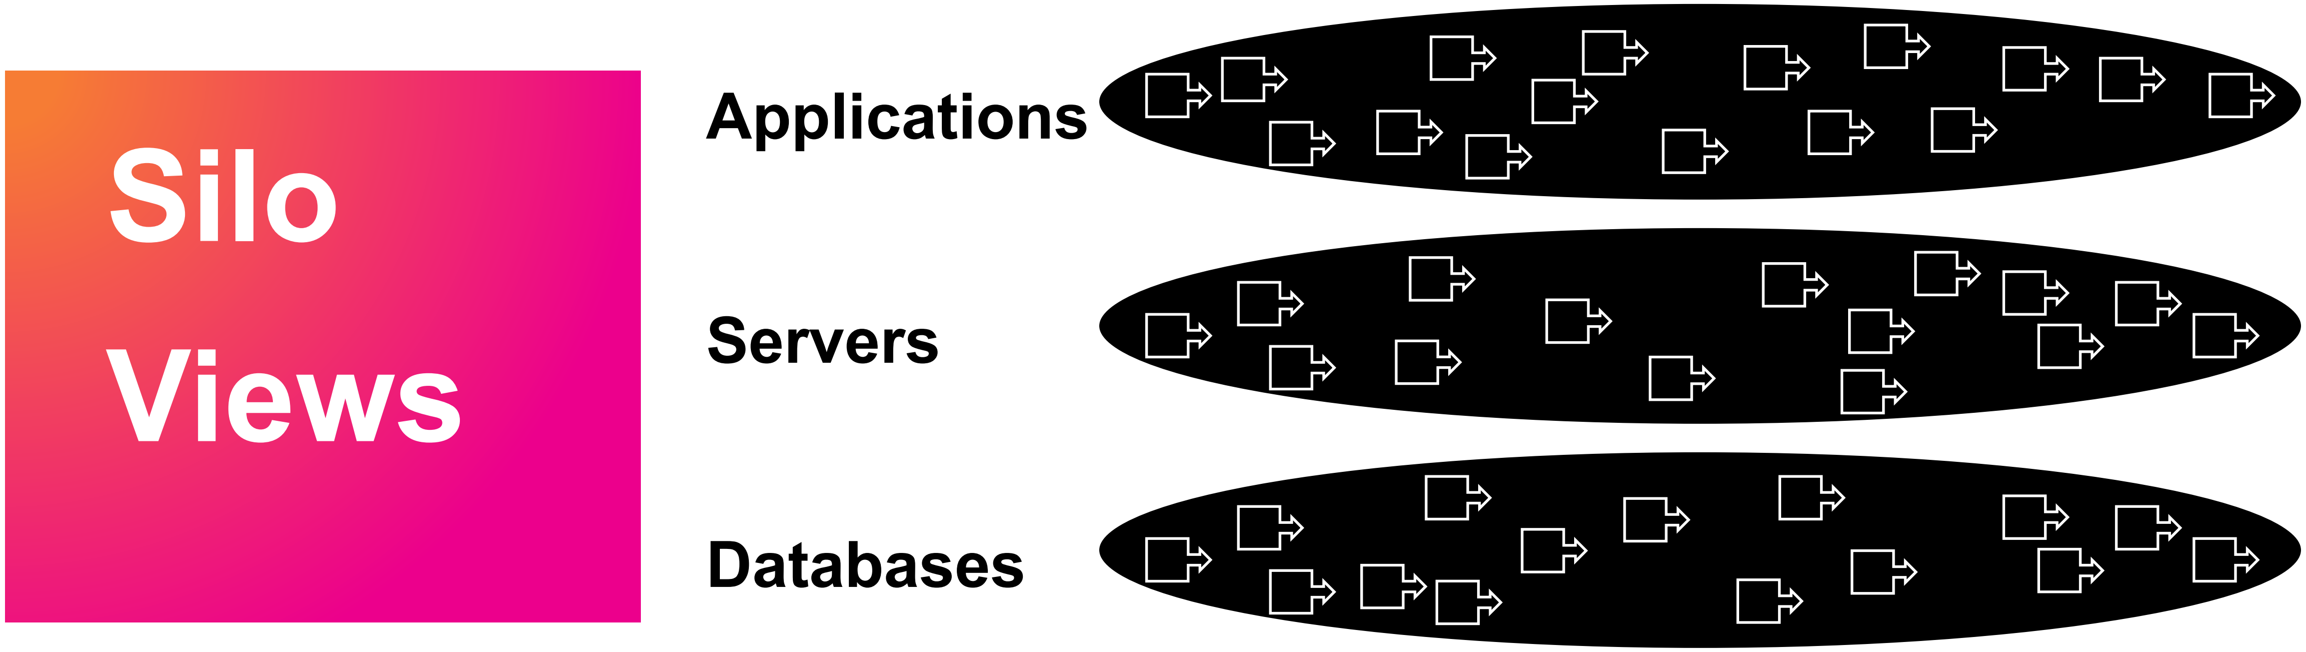

Event Analytics Lab

- If we can cluster events by time, we can immediately reduce noise. The two database events in orange happened around the same time and therefore, may very well be related to the same incident

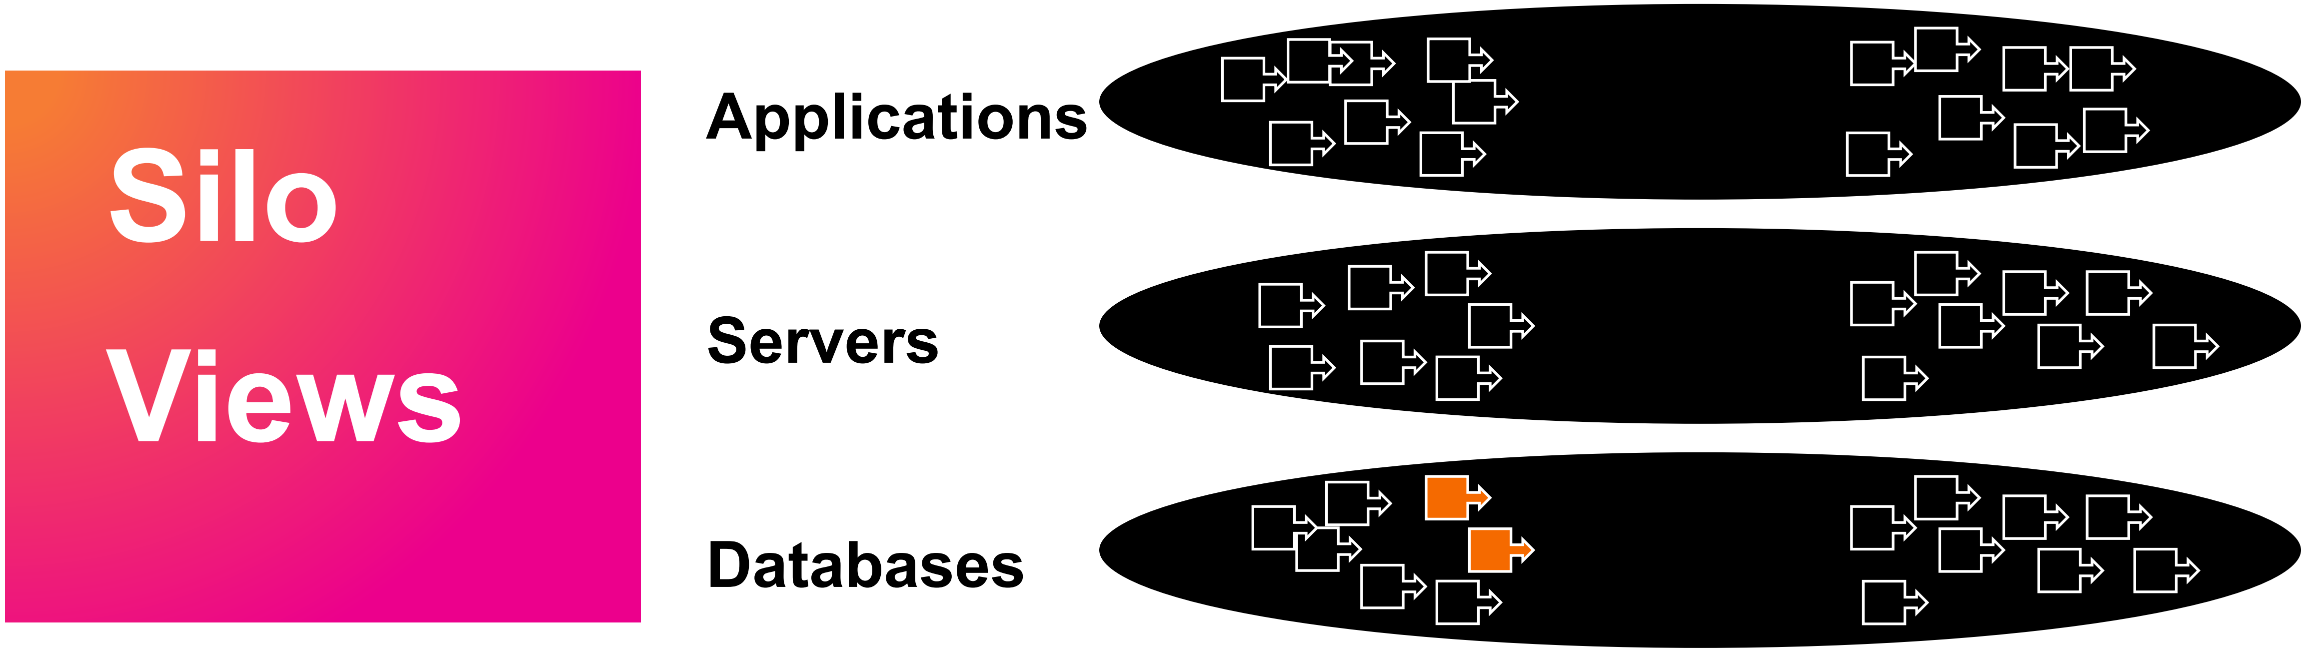

Event Analytics Lab

- The three events in green occurred later after a pause in the flow of events. Probably a different issue all together. So grouping events together by time is a very powerful noise reduction technique.

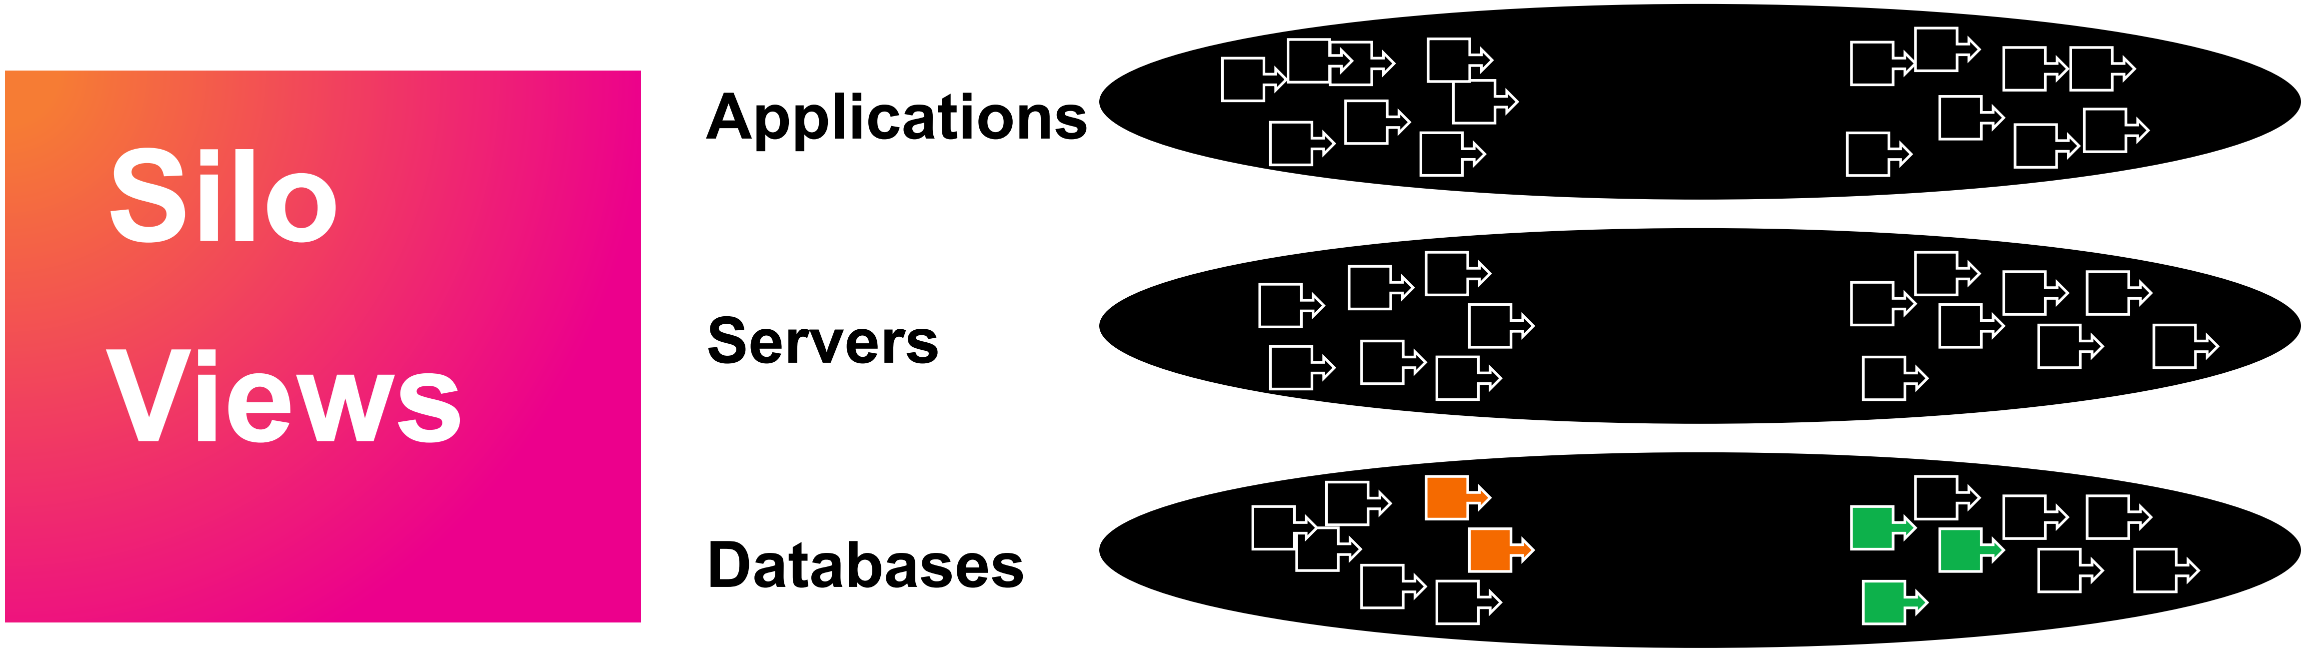

Event Analytics Lab

- We can extend the power of grouping beyond just time to further reduce noise. What if all of the events in orange we associated to the same database instance. Same timeframe… Same instance… that's probably all a related incident.

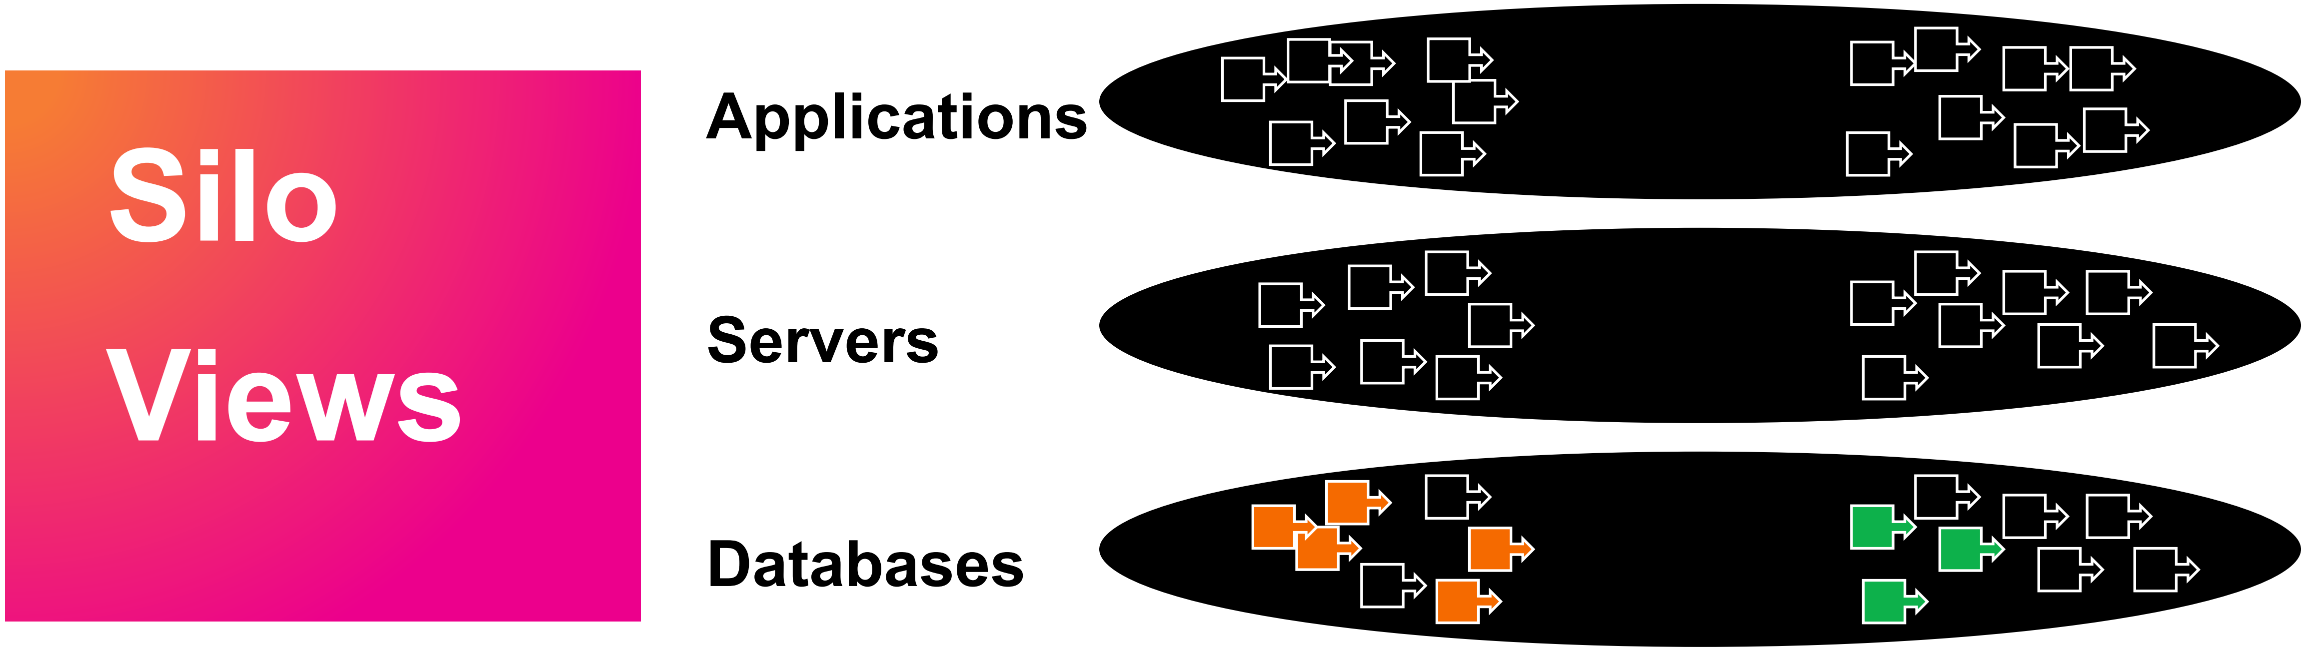

Event Analytics Lab

- We can extend the grouping across siloed monitoring tools, and across layers of the stack. What if I told you that all the events in orange were associated with machines that run the Buttercup Store.

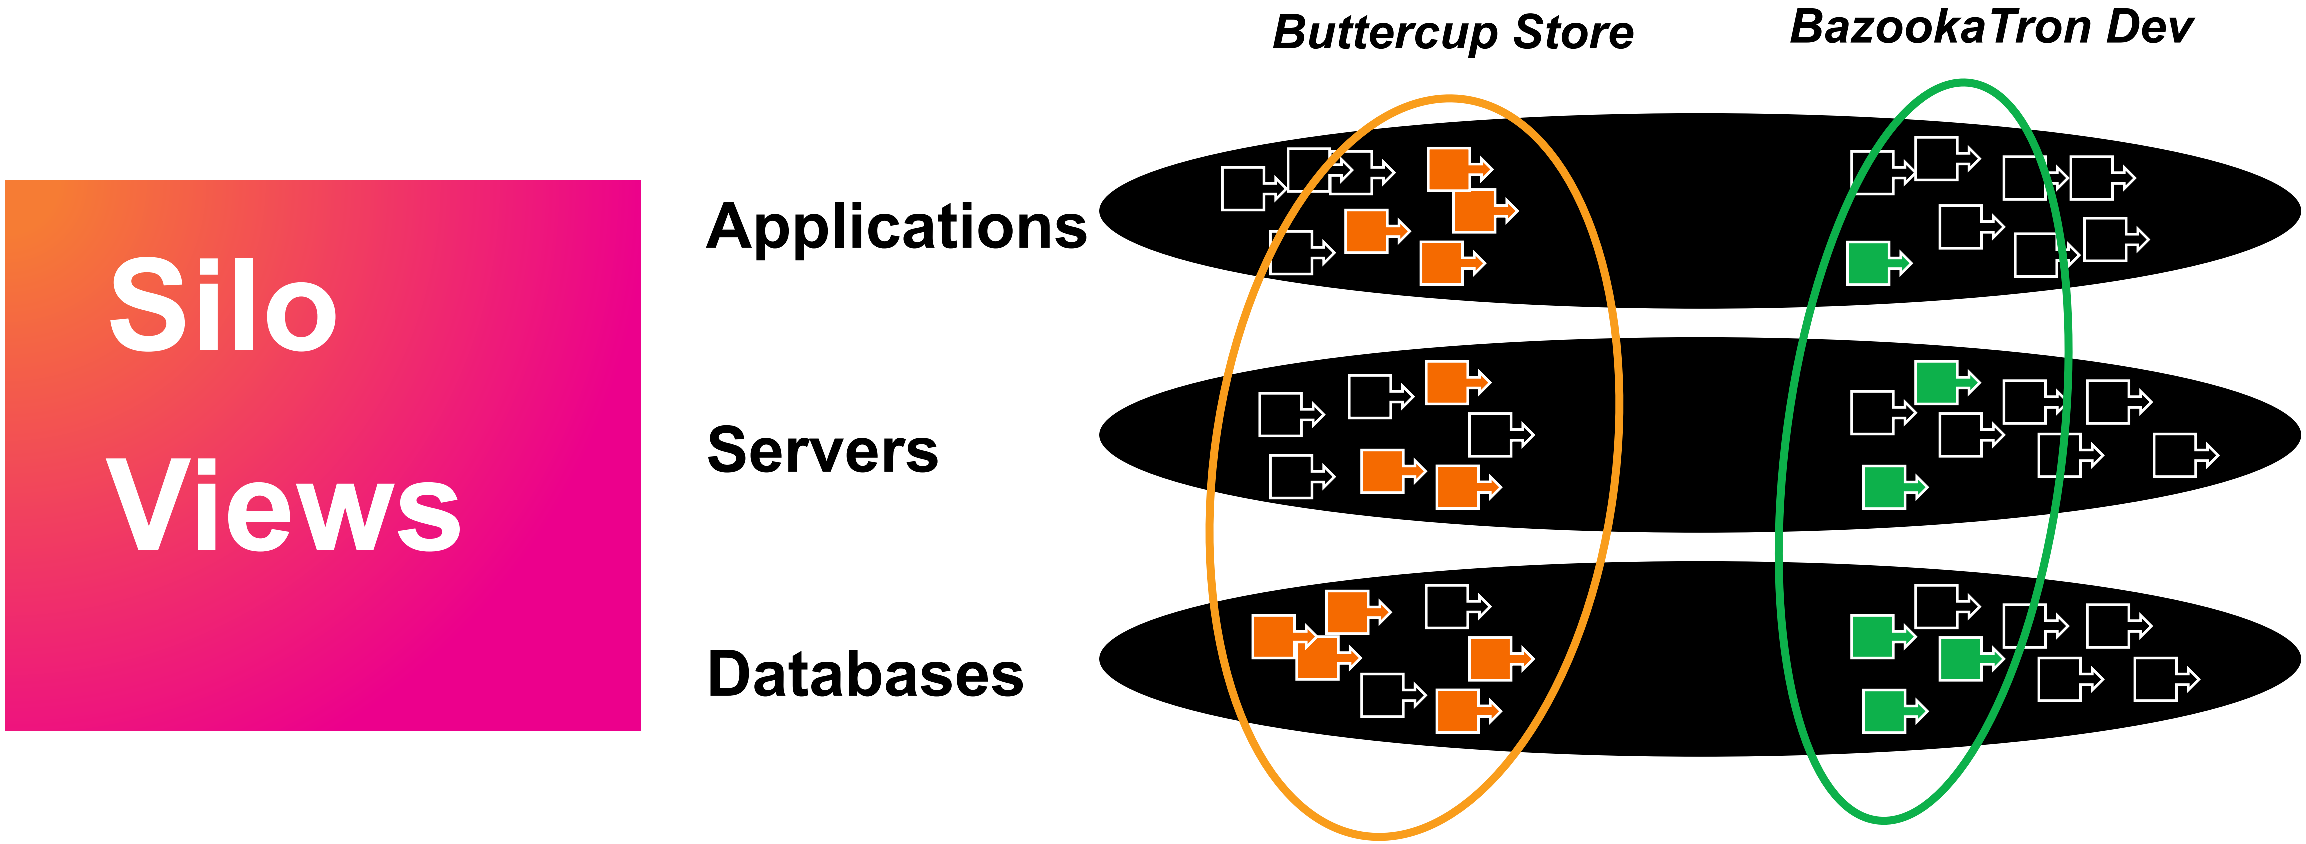

A quick terminology check

- **Alert** - Describes a state change for a target entity
  - Examples of different alerts:: Server42 is down, Filesystem is full
- **Alarm** - Specific alert for a target entity, can change severity/state over time
  - Example of single alarm: Server42 is “down”, then later “up”
- **Time-series data or Events**- The stuff that Splunk indexes
  - Includes traditional alerts & alarms, as well as logged data, metrics, wire data and more
- **Notable Event** - An actionable message (Splunk ITSI & Splunk ES)
  - Intended specifically for humans in Operations
- **Episode** - a group of Notable Events
- **Incident** - Unplanned interruption of an IT or business service
  - ServiceNow, Remedy

Event Analytics Lab

In this exercise we will group the database teams Nagios events together based on time and associate them to a service

Clean and prepare “raw” alert events
- SPL

Create Notable Events from alerts
- *Correlation Search*

Apply Service Context & Configure Event Grouping
- Notable Event Aggregation Policies

Review episodes

EA Lab Step 1 : Clean and prepare "raw" alert events

Check raw alerts → Normalizing → Filter duplicates → Add Service Context

| index=itsidemo sourcetype=nagios perfdata=SERVICEPERFDATA

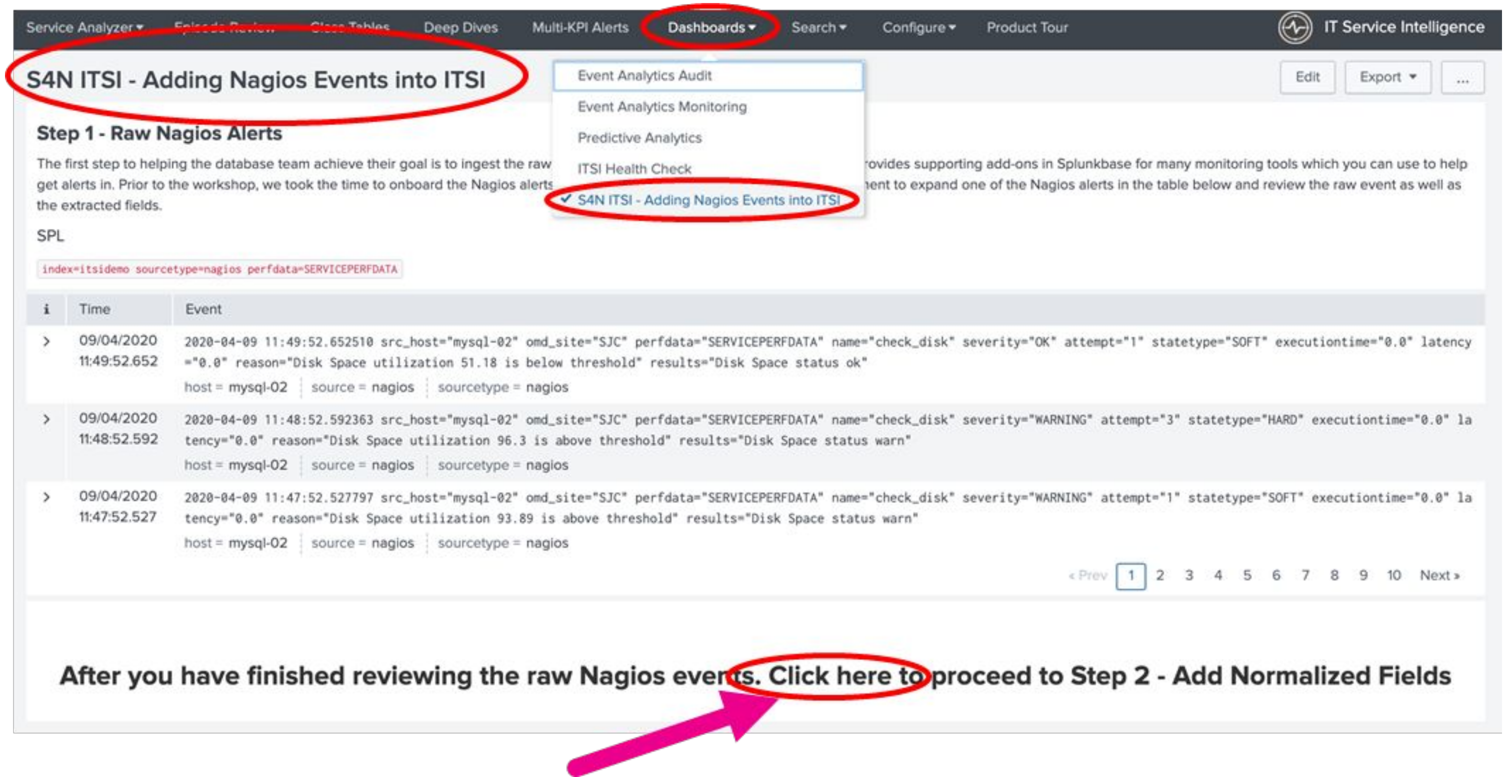

EA Lab Step 2 : Add Normalized Fields

Check raw alerts → **Normalizing** → Filter duplicates → Add Service Context

Goal: Normalize the data to make it possible to correlate data from different sources

```
| eval norm_severity=case(severity=="CRITICAL",6,severity=="WARNING",4, severity="OK", 2)
| eval norm_instance= src_host
| eval norm_test=name
```

Step 2 - Add Normalized Fields

While each monitoring tool will express events differently, they all communicate the same fundamental information. Such as, how severe is the event? To which machine is the event associated? What type of check or test was performed? To facilitate the grouping of multiple events from multiple monitoring tools, we must normalize this key information so that a common set of field names and values is used. The SPL below creates these normalized severity, instance, and test fields.

SPL

```
index=itsidemo sourcetype=nagios perfdata=SERVICEPERFDATA
| eval norm_severity=case(severity=="CRITICAL",6,severity=="WARNING",4, severity="OK", 2)
| eval norm_instance=src_host
| eval norm_test=name
```

_time	name	severity	norm_severity	norm_instance	norm_test	total
2020-04-15 10:02:27.773	check_disk	OK	2	mysql-02	check_disk	1
2020-04-15 10:01:27.710	check_disk	WARNING	4	mysql-02	check_disk	1
2020-04-15 09:56:27.396	check_disk	OK	2	mysql-02	check_disk	1
2020-04-15 09:55:27.334	check_disk	WARNING	4	mysql-02	check_disk	1
2020-04-15 09:51:27.088	check_disk	OK	2	mysql-02	check_disk	1
			594			157

« Prev 1 2 3 4 5 6 7 8 9 10 Next »

After you have finished reviewing the new normalized fields, click here to proceed to Step 3 - Deduplicate events

EA Lab Step 3 : Deduplicate events

Check raw alerts → Normalizing → **Filter duplicates** → Add Service Context

Goal: Removing duplicated events

| dedup consecutive=true src_host severity name

_time	name	norm_instance	severity	norm_severity	total
2020-04-15 10:02:27.773	check_disk	mysql-02	OK	2	1
2020-04-15 10:01:27.710	check_disk	mysql-02	WARNING	4	1
2020-04-15 09:56:27.396	check_disk	mysql-02	OK	2	1
2020-04-15 09:55:27.334	check_disk	mysql-02	WARNING	4	1
2020-04-15 09:51:27.088	check_disk	mysql-02	OK	2	1
				102	34

« Prev 1 2 3 4 5 6 7 Next »

Finally, after you have finished reviewing the deduplicated events, click here to proceed to Step 4 - Add Service Context

EA Lab Step 4 - Add Service Context

Check raw alerts → Normalizing → Filter duplicates → **Add Service Context**

Goal: Add Service Context makes correlation possible for different alerts in a service

| `apply_entity_lookup(host)`
| `get_service_name(serviceid,service_name)`

▶ Click on the magnifying glass

_time	name	norm_instance	severity	norm_severity	service_name
2020-04-15 10:02:27.773	check_disk	mysql-02	OK	2	On-Prem Database
2020-04-15 10:01:27.710	check_disk	mysql-02	WARNING	4	On-Prem Database
2020-04-15 09:56:27.396	check_disk	mysql-02	OK	2	On-Prem Database
2020-04-15 09:55:27.334	check_disk	mysql-02	WARNING	4	On-Prem Database
2020-04-15 09:51:27.088	check_disk	mysql-02	OK	2	On-Prem Database

« Prev 1 Open in Search 6 7 Next »

a minute ago

Event Analytics Lab

Check raw alerts → Normalizing → Filter duplicates → **Add Service Context**

- Select the SPL Query and copy to clipboard, we will need it in the next chapter

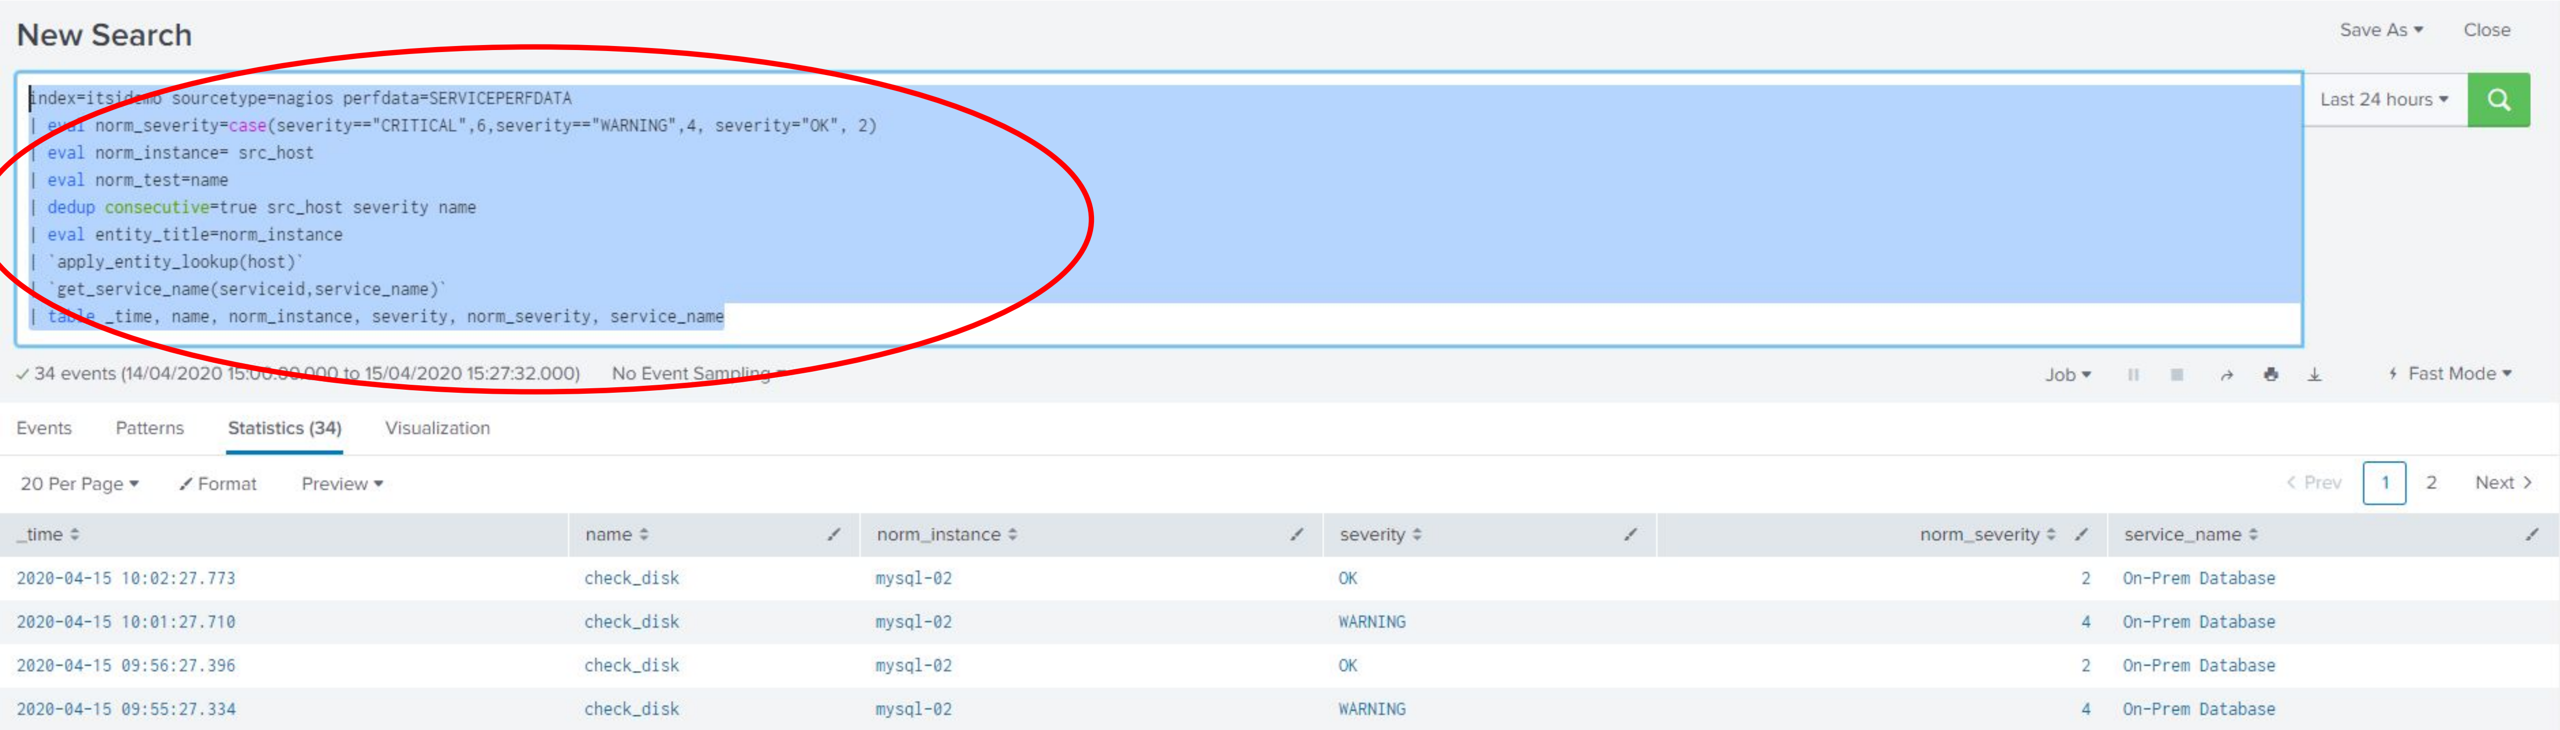

_time	name	norm_instance	severity	norm_severity	service_name
2020-04-15 10:02:27.773	check_disk	mysql-02	OK	2	On-Prem Database
2020-04-15 10:01:27.710	check_disk	mysql-02	WARNING	4	On-Prem Database
2020-04-15 09:56:27.396	check_disk	mysql-02	OK	2	On-Prem Database
2020-04-15 09:55:27.334	check_disk	mysql-02	WARNING	4	On-Prem Database

Event Analytics Lab

In this exercise we will group the database teams Nagios events together based on time and associate them to a service

Clean and prepare “raw” alert events
- SPL

Create Notable Events from alerts
- *Correlation Search*

Apply Service Context & Configure Event Grouping
- Notable Event Aggregation Policies

Review episodes

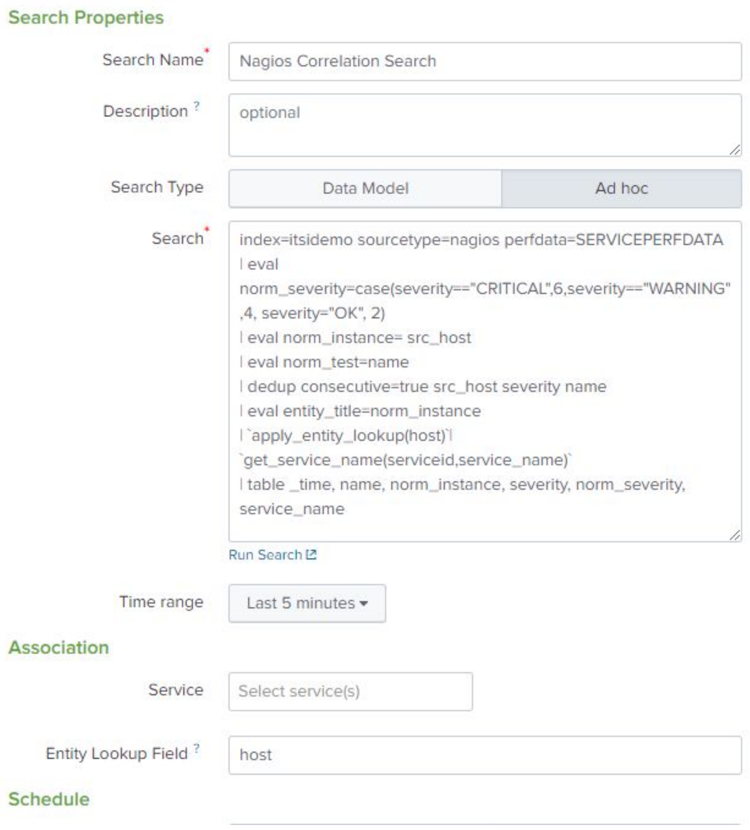

Event Analytics Lab

- Navigate to Configure -> Correlation Searches
- Create New Search -> Create Correlation Search

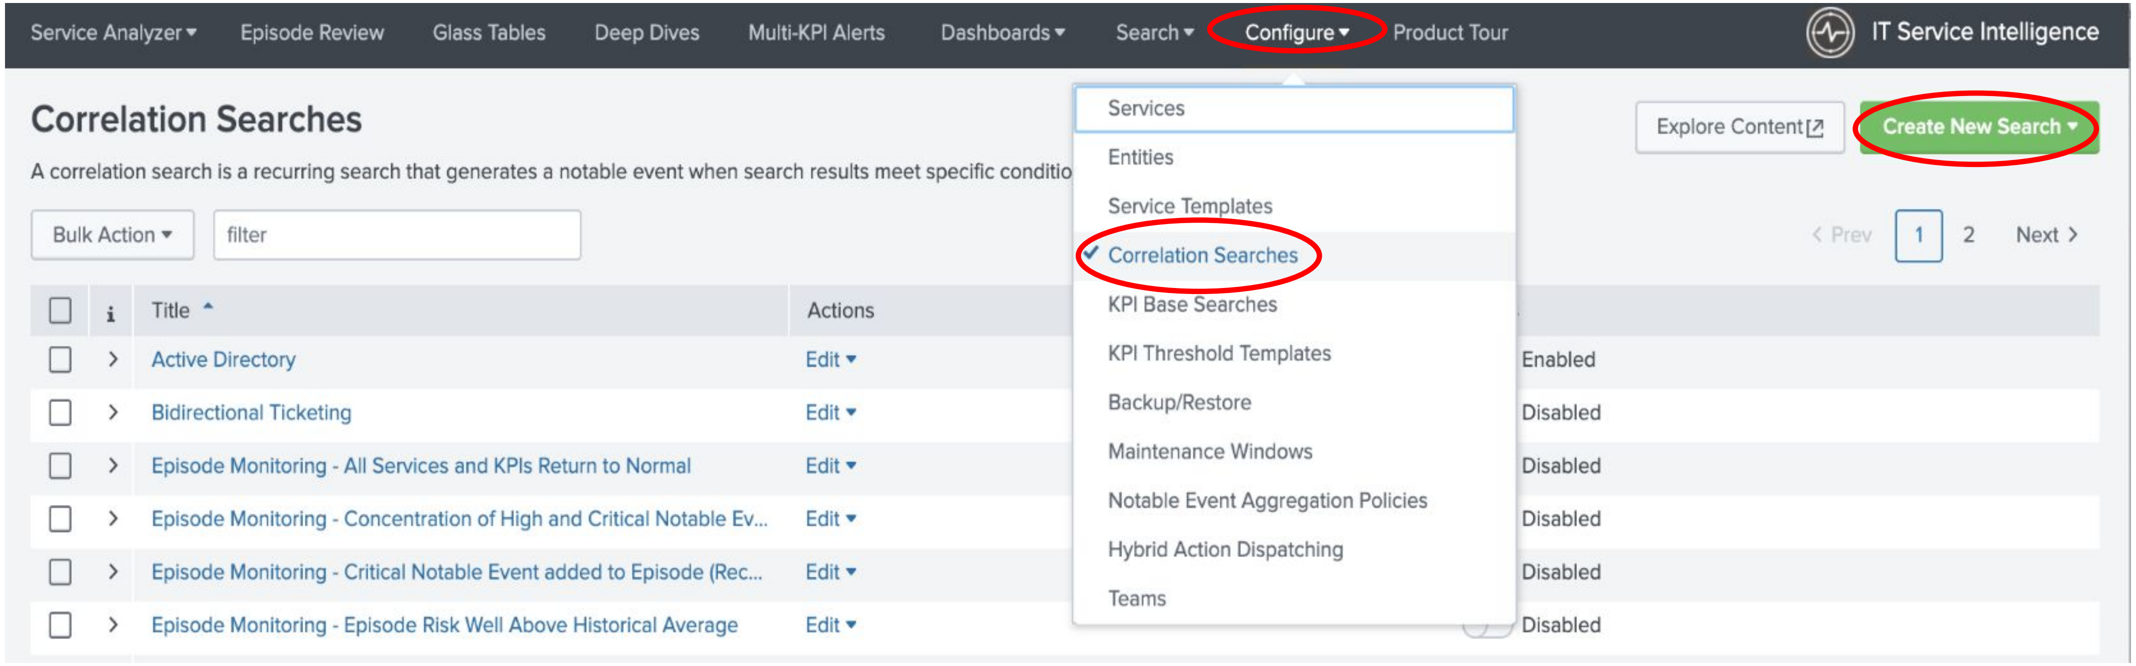

Event Analytics Lab: correlation search

- Name your search
- *Paste the SPL in search:*

index=itsidemo sourcetype=nagios perfdata=SERVICEPERFDATA

| eval norm_severity=case(severity=="CRITICAL",6,severity=="WARNING",4, severity="OK", 2)

| eval norm_instance= src_host

| eval norm_test=name

| dedup consecutive=true src_host severity name

| eval entity_title=norm_instance

| `apply_entity_lookup(host)`

| `get_service_name(serviceid,service_name)`

| table _time, name, norm_instance, severity, norm_severity, service_name

- *Time range: Last 5 minutes (select 'relative' in time picker)*
- *Run Every: 5 minutes*
- *Entity Lookup Field: host □to link with th Service Context*
- *Scroll down*

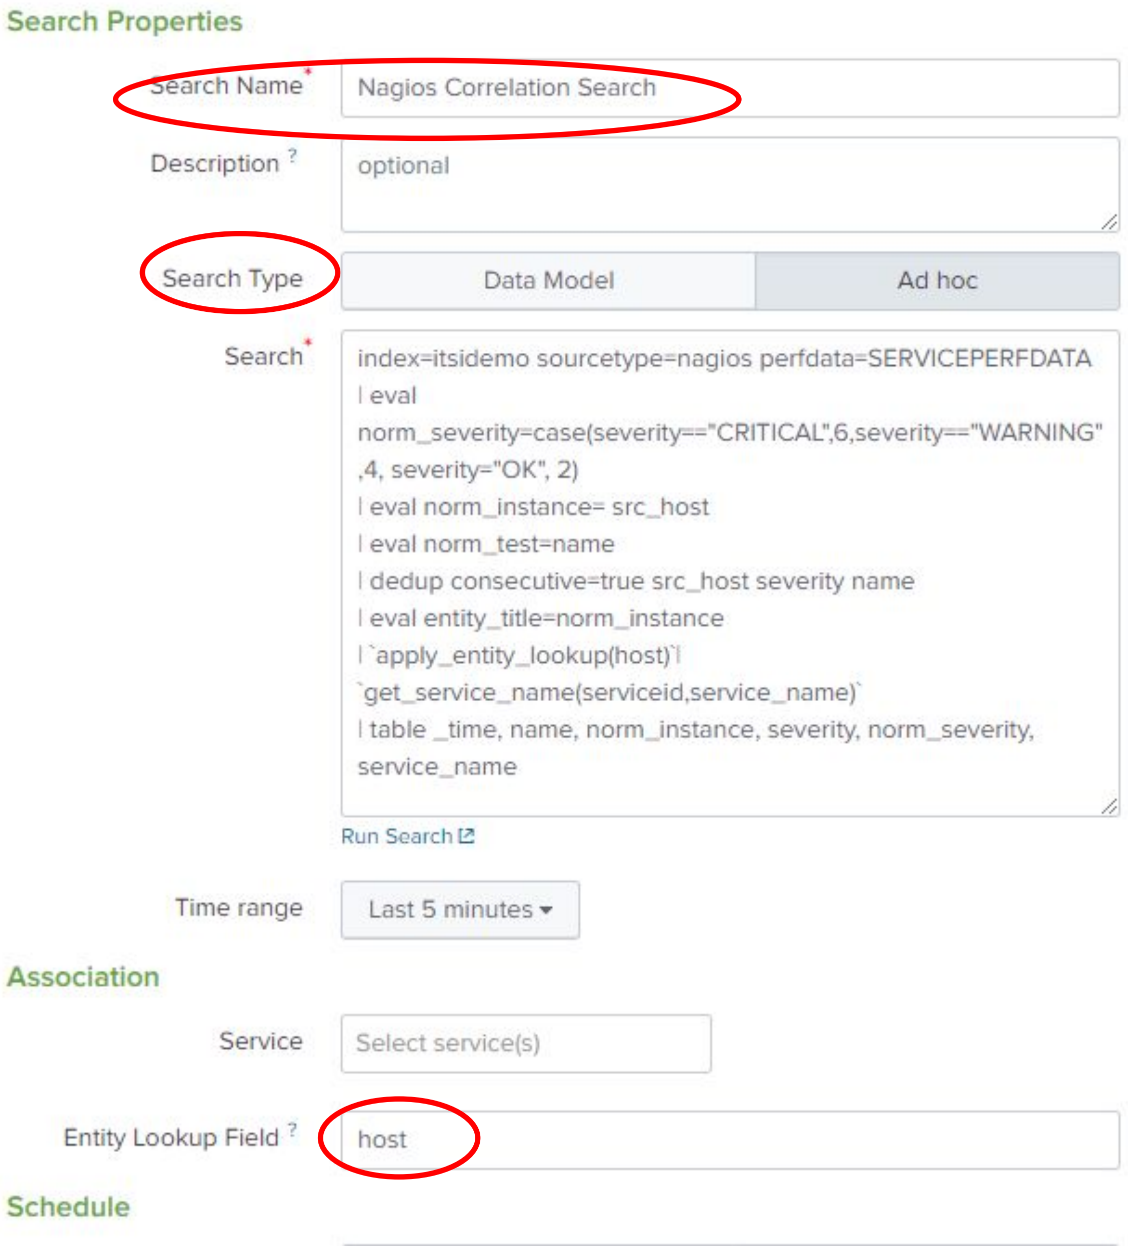

Event Analytics Lab: correlation search

- *Populate Notable Event Title*

  Nagios alert from %norm_instance%

- *Populate Notable Event Desc.*

  Nagios alert from %norm_instance%. %norm_test (%severity%)

- *Severity: Advanced Mode*
- *Severity: %norm_severity%*

- *Save*

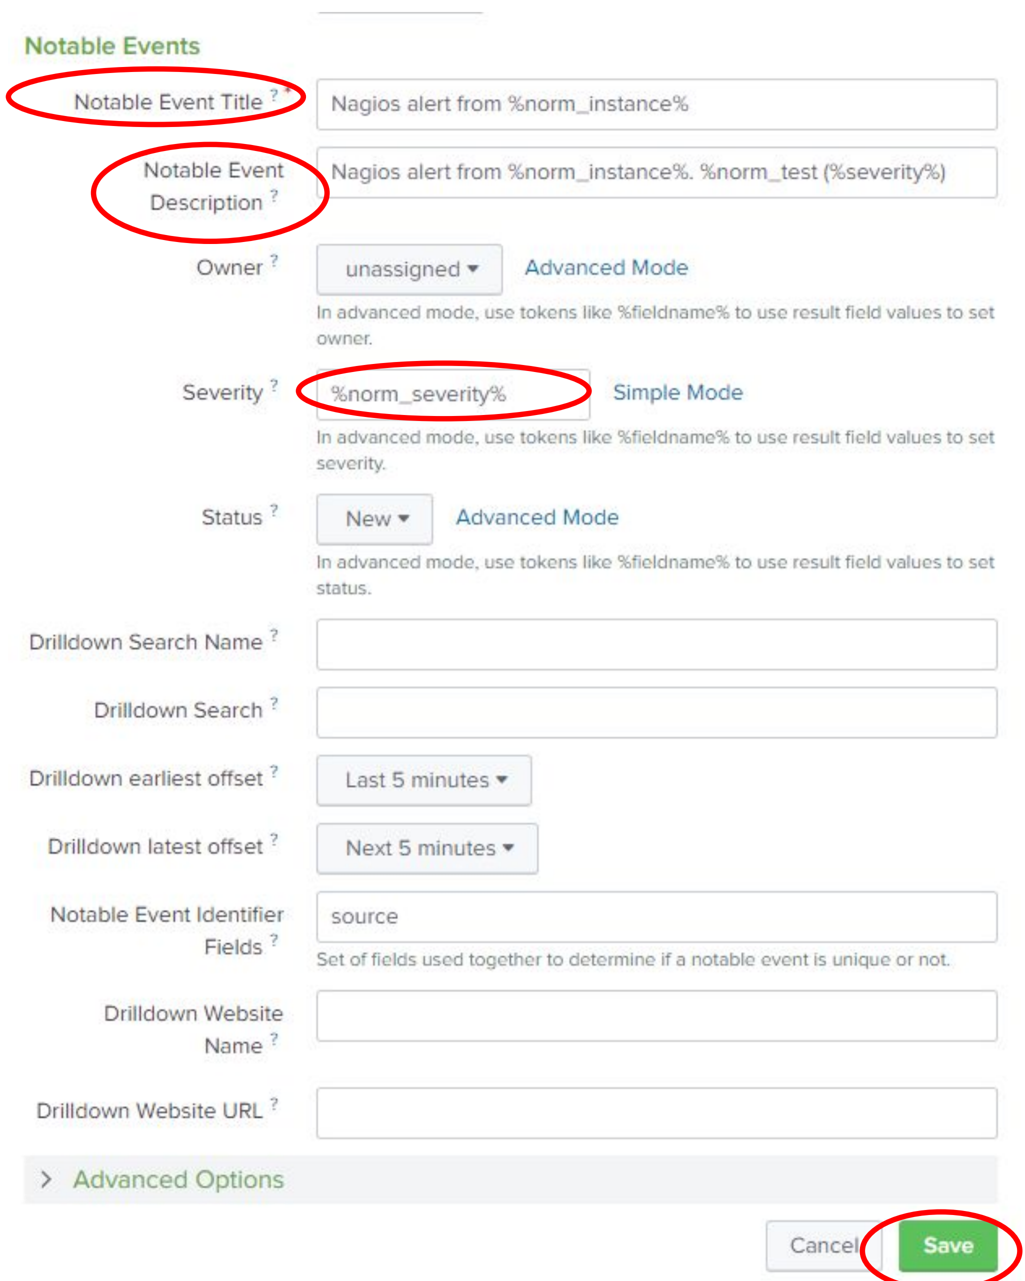

Event Analytics Lab: Notable Event Aggregation Policies

In this exercise we will group the database teams Nagios events together based on time and associate them to a service

Clean and prepare “raw” alert events
- SPL

Create Notable Events from alerts
- *Correlation Search*

Apply Service Context & Configure Event Grouping
- Notable Event Aggregation Policies

Review episodes

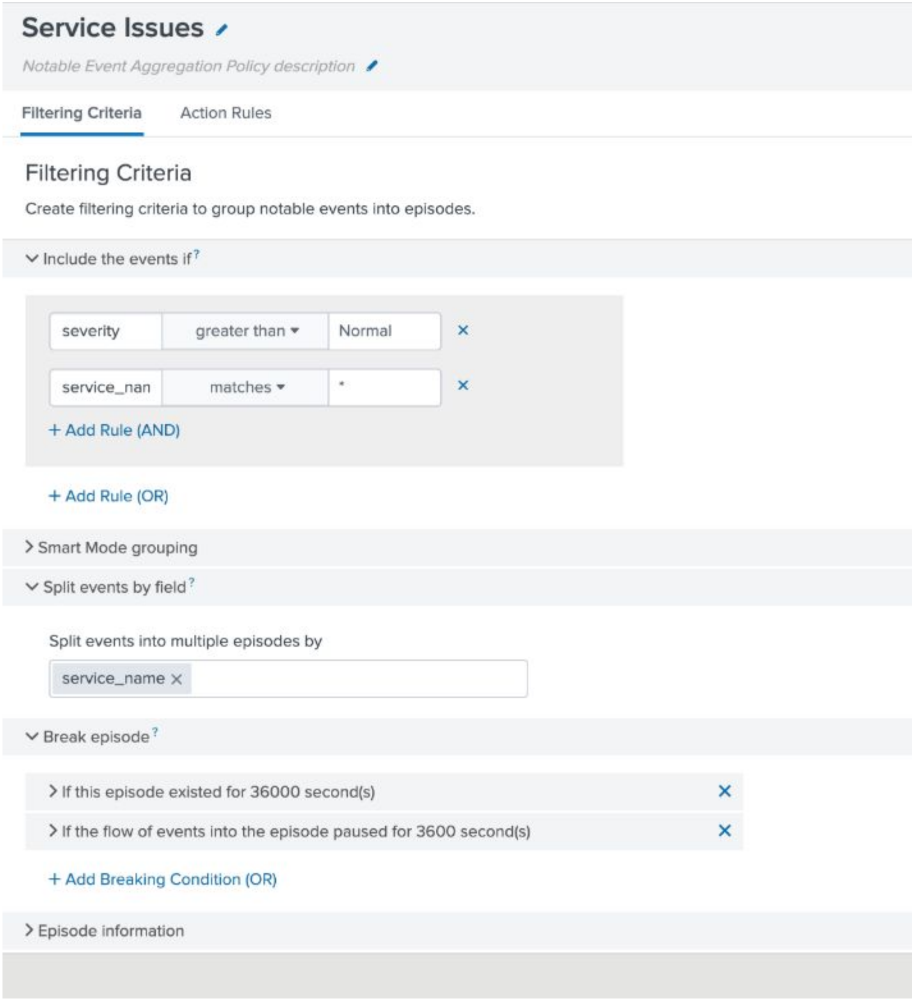

Event Analytics Lab: Notable Event Aggregation Policies

- Navigate to Configure -> Notable Event Aggregation Policies
- Edit *"Service Issues"* policy

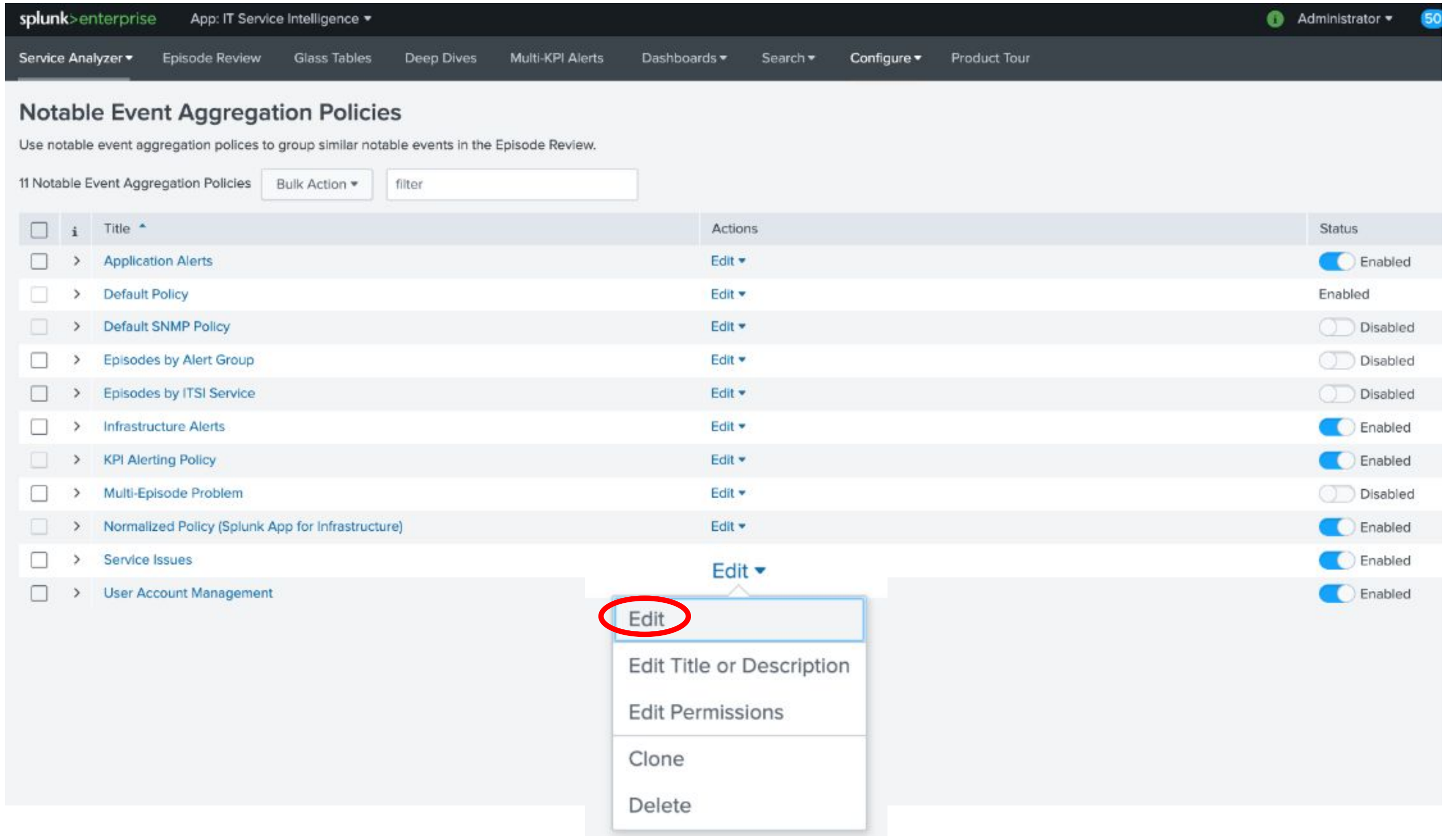

Event Analytics Lab: Notable Event Aggregation Policies

- *Review "Include the events if" configuration*
- *Review "Split events by field" configuration*
- *Review "Break episode" configuration*
- *Modify the "Break episode" configuration*
  - *If the flow of events paused for 3600 seconds*

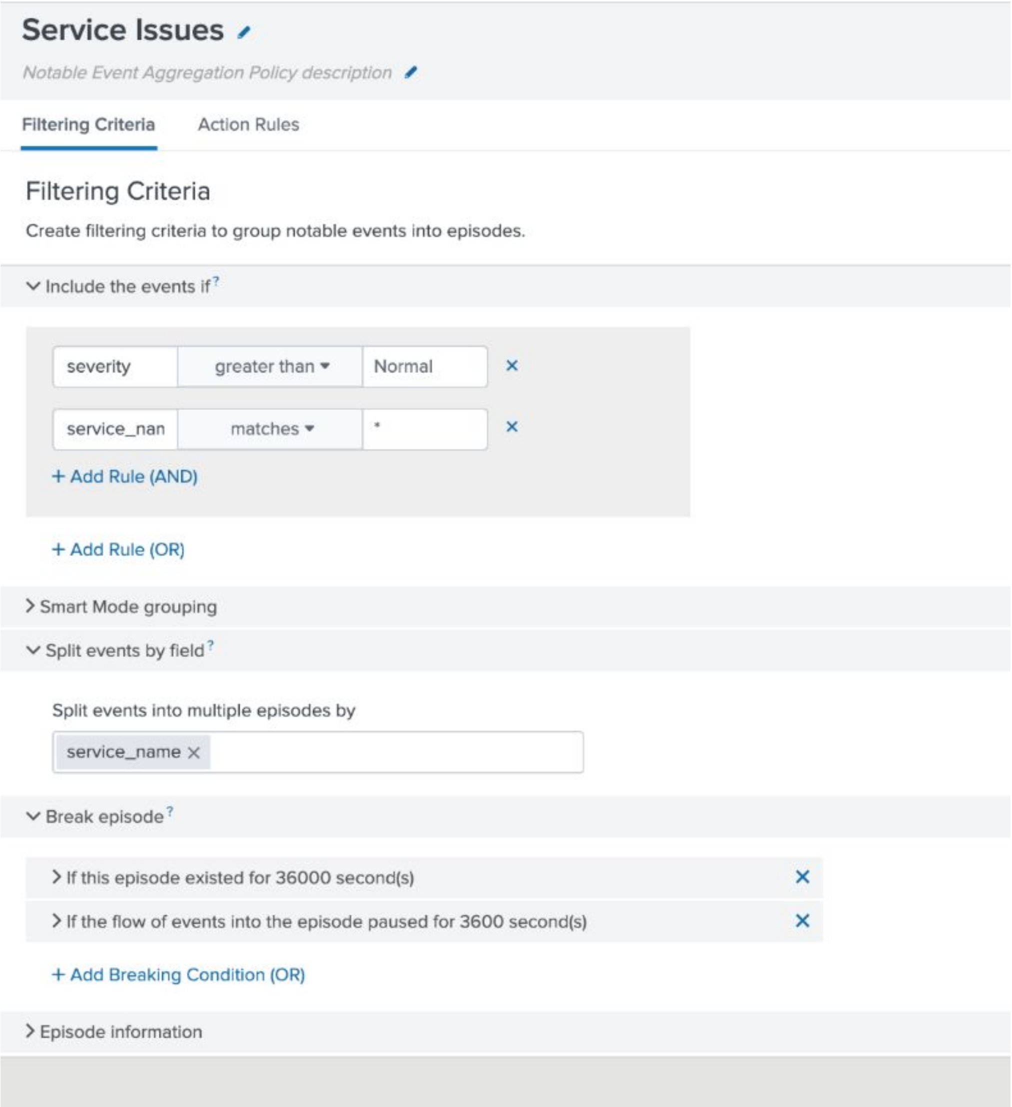

Event Analytics Lab: Notable Event Aggregation Policies

- You can see in the preview results how ITSI is now grouping events together based on your configurations.

- Preview Results

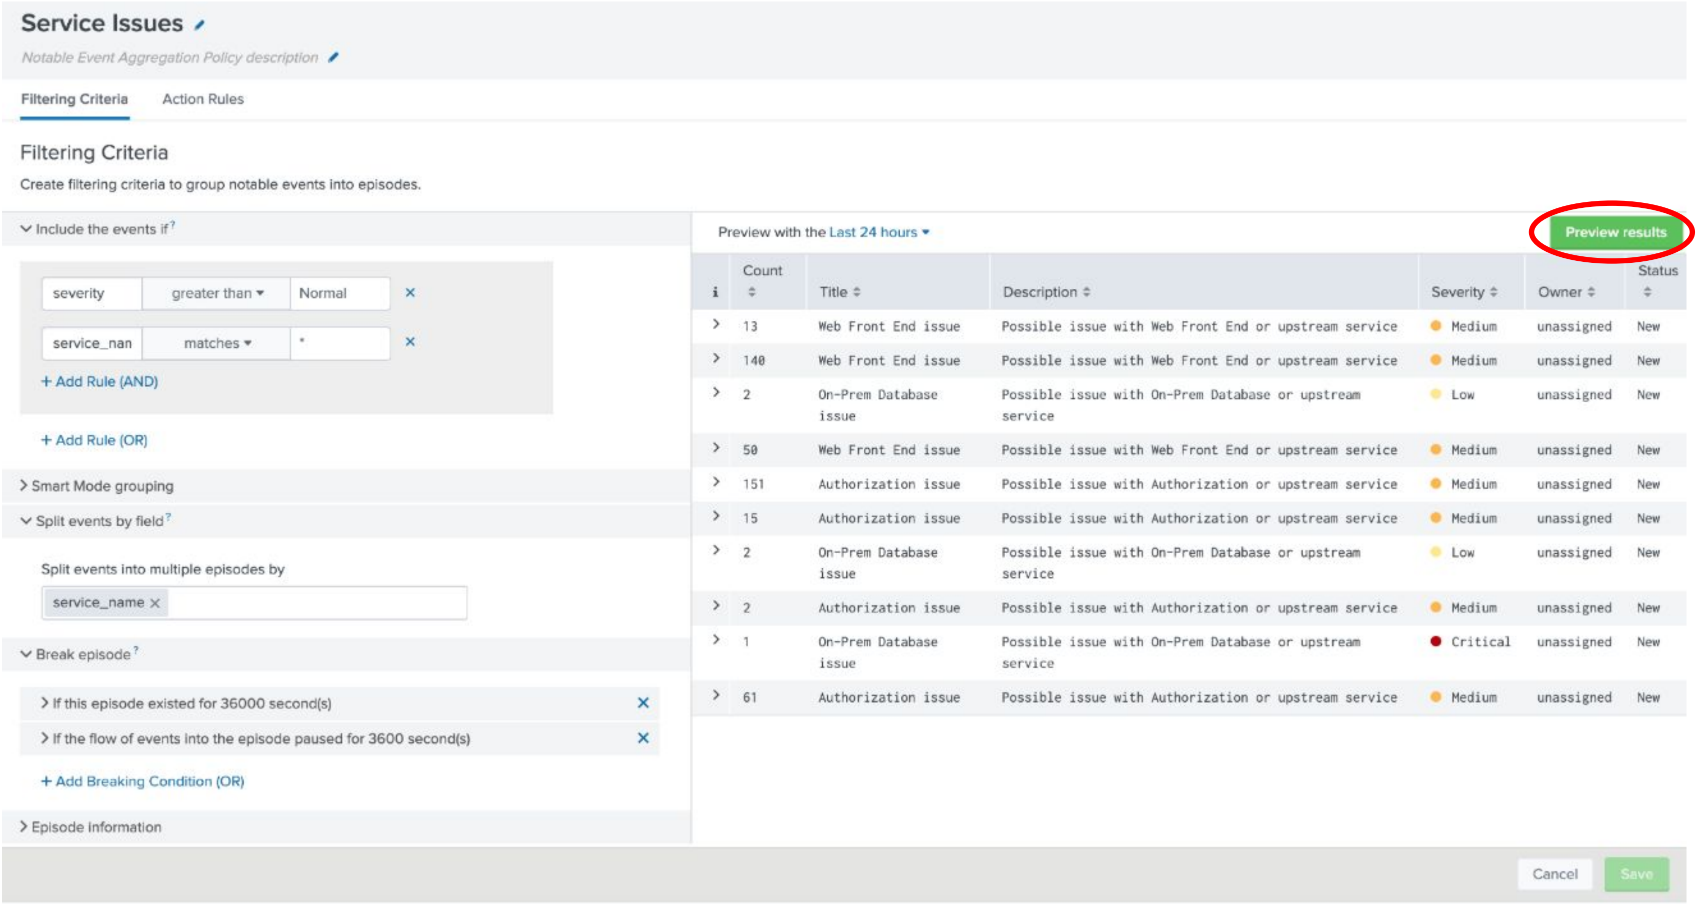

- Click Cancel

Event Analytics Lab

In this exercise we will group the database teams Nagios events together based on time and associate them to a service

Clean and prepare “raw” alert events
- SPL

Create Notable Events from alerts
- *Correlation Search*

Apply Service Context & Configure Event Grouping
- Notable Event Aggregation Policies

Review episodes

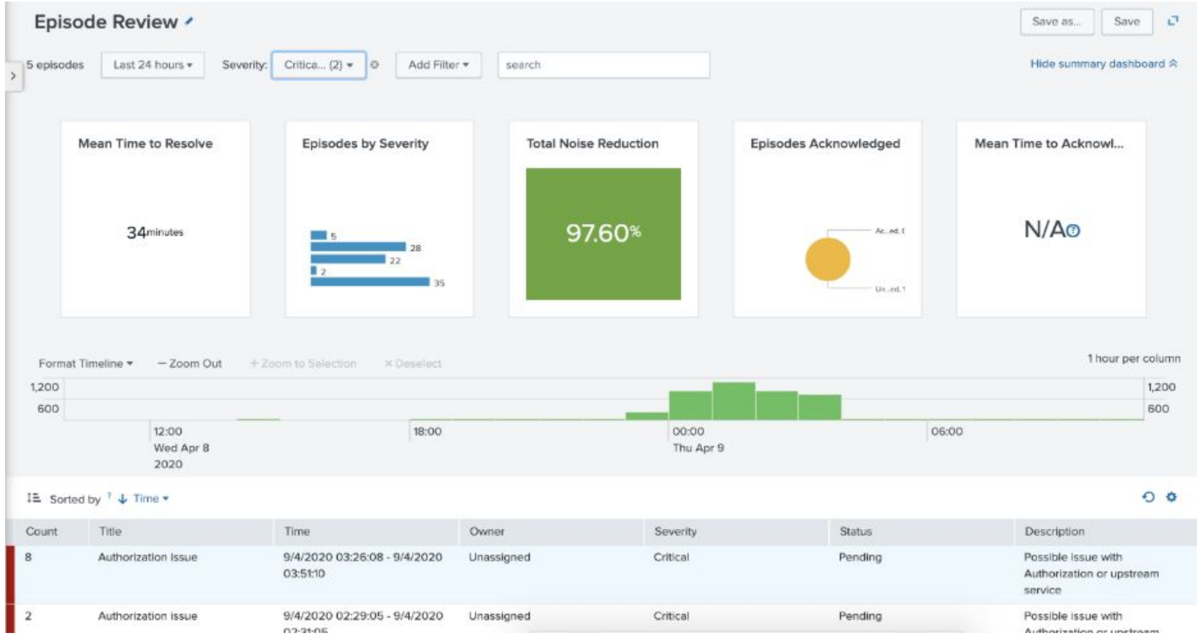

Event Analytics Lab: Episode View

When notable events are grouped by aggregation policies, the resulting groups are called episodes, you can think of an episode as an incident. The episode review page provides a great deal of information in a heads-up display and is like the cockpit view for Operations teams.

- Navigate to Episode Review
- Note the Noise Reduction

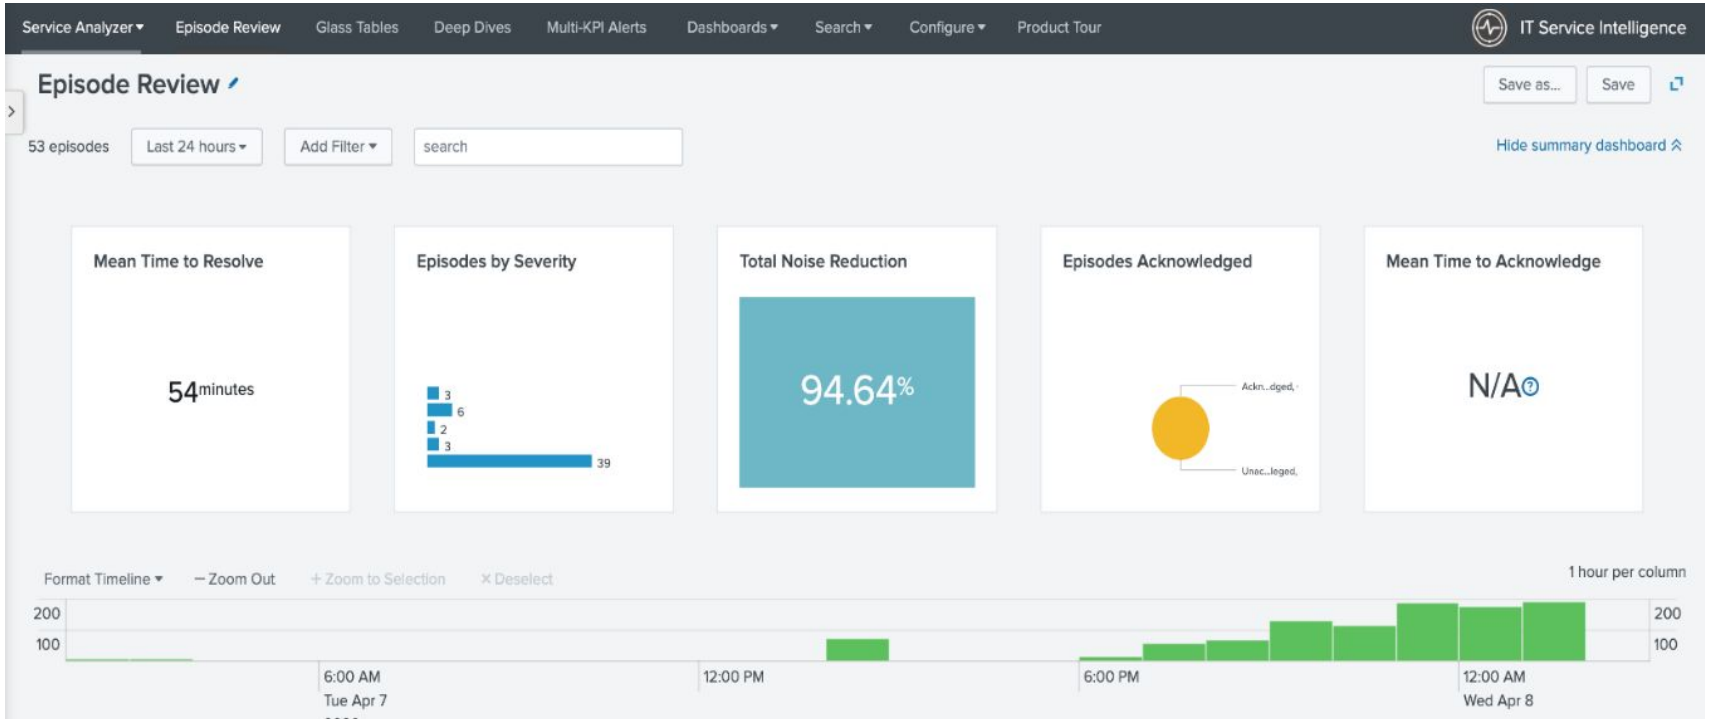

Event Analytics Lab: Episode View

- *Scroll down and review list*
- *Modify filter as shown*
  - *Critical & High only*

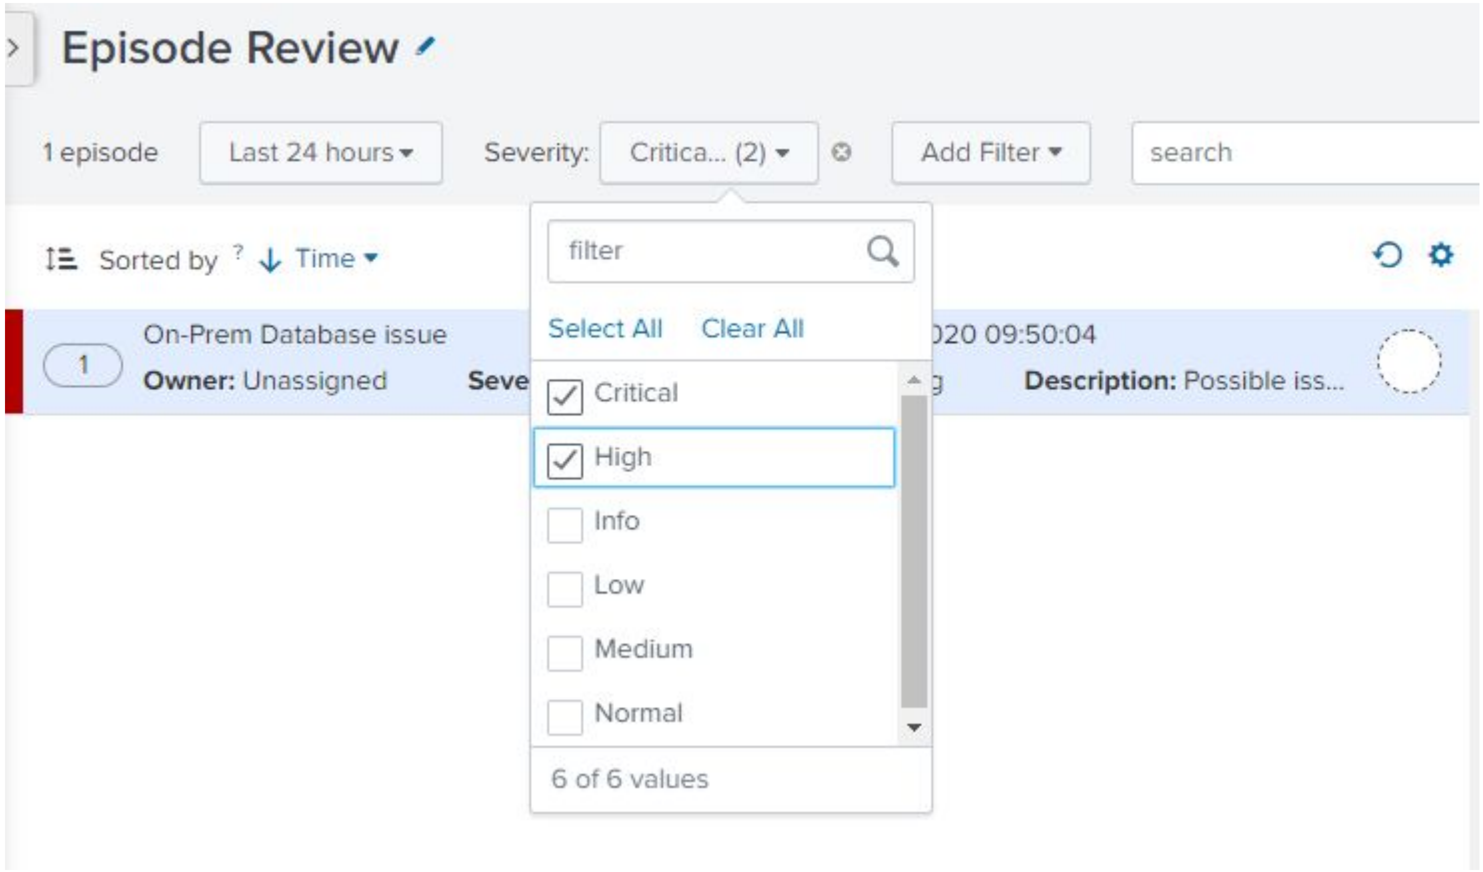

Event Analytics Lab: Episode View

- We will now review an episode to better understand the flow of events, and we will then ensure someone has ownership.

- Click on '*Authorization Issue*' episode
- Review the details for each tab
- Add your name to the comments

- Change to '*In Progress*'
- Review possible Actions

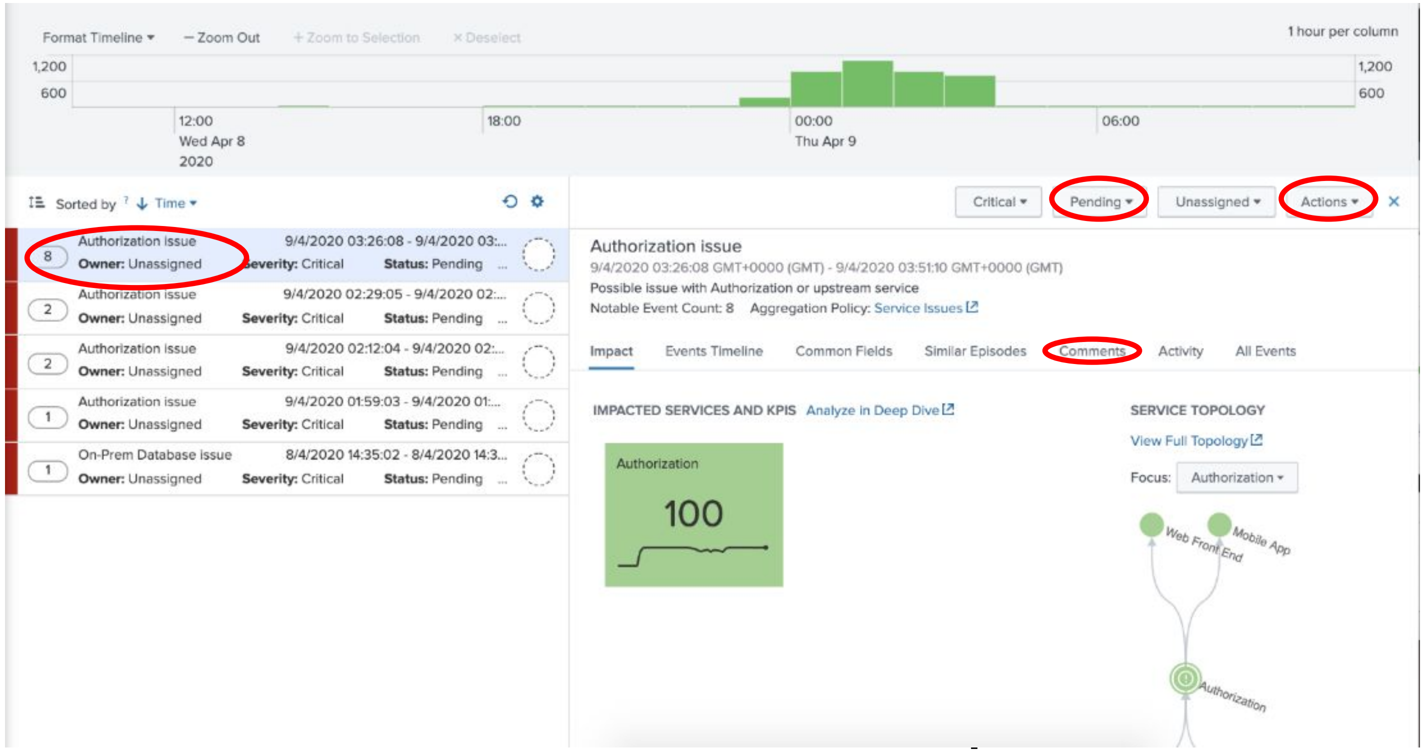

Next Steps

Somerford's Approach to ITSI

Somerford’s Glass Table Methodology

Project Phases

Initiation

- Kick-Off Call
- Pre-Workshop Planning
- ITSI Product Walkthrough

Analysis

- Service Identification Workshop
- Entity Strategy Workshop
- Service Decomposition Workshops

Data Onboarding & ITSI Installation

- ITSI Installation
- Data Onboarding
- Implement Entity Management Strategy

Delivery

- Delivery Planning
- ITSI Services Development
- Progress Review
- ITSI Service Review
- ITSI UI Development

Handover

- Project Closeout
- Enablement
- BAU Transition
- Onboarding Users
- Follow-up Reviews

Service Identification

- Which business and technology services are candidates?
- Which services shall we do first?

Service Decomposition

- What are the business services to be monitored?
- What technology services make up those business services?
- What are the components of the technology services?
- What do you care about? What are the metrics?
- What are the KPIs?
- What data drives each KPI?

Thank You!

Please get in touch with us
if you have any questions.

paul.winchester@somerfordassociates.com

@Somerford_Ltd

Somerford Associates Limited