.conf2014

YOUR DATA ADVENTURE

Speeding up Dashboards with Pivot

Rupak Pandya | OI Practice
Manager, Function1

FUNCTION1

www.function1.com

splunk>

Disclaimer

During the course of this presentation, we may make forward-looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only, and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.

Agenda

- Background
- Splunk Search
- Overview of Data Models
- Overview of the Pivot Search Command
- Demonstration

Meet Function1

- Founded in 2007, Function1 is an enterprise technology solution firm and has been a Preferred Splunk Partner since 2011
- We have 11 consultants in our Operational Intelligence group that specialize in delivering Splunk Professional Services

Meet Rupak Pandya

- Practice Manager of the Function1 OI Group
- Joined the Operational Intelligence team as an experienced consultant from a large, global consulting company
- Avid follower of all of the Washington, D.C. area sports teams

rupak@function1.com
301.452.2475 ext 24

Anatomy of a Splunk Search

```
index=_internal sourcetype=splunkd group=per_sourcetype_thruput |
eval host_series = host + "_" + series | stats sum(ev) by
host_series | rename sum(ev) AS "Total Events" | sort - "Total
Events"
```

> Search And Filter

> Enrich

> Report

> Format

Challenges

```
index=warum sourcetype="ri:pas:application" | lookup employee_details empusername as user OUTPUT empshiftstarttime empshiftendtime empshiftdays | eval hour = tonumber(strftime(_time,"%H")) | eval overnight = if(empshiftstarttime>empshiftendtime,1,0) | eval valid_time=if((overnight==1 AND (hour >= empshiftstarttime OR hour <= empshiftendtime)) OR (overnight==0 AND hour>=empshiftstarttime AND hour<empshiftendtime),1,0) | eval shiftday=if(overnight==1 AND hour < empshiftendtime,relative_time(_time,"-1d"),_time) | eval weekday = strftime(_time,"%a") | eval shiftday=strftime(shiftday,"%a") | eval valid_day=if(match(empshiftdays,shiftday),1,0) | search valid_time=0 OR valid_day=0 | bucket _time span=5m | stats count by _time user command object
```

- Complicated searches get very verbose
- Searchers need to understand data's structure
- Non-technical users might not have knowledge of underlying data
- Splunk admins do not always know what users will be searching on

There Has to be a Better Way…

Data Model Goals

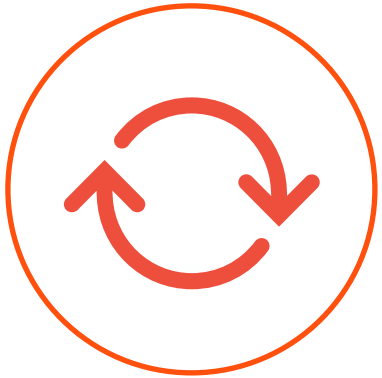

Make it easy to share/reuse domain knowledge

Admins/power users build data models

Non-technical users interact with data via pivot UI

Data Models

splunk>

What are Data Models?

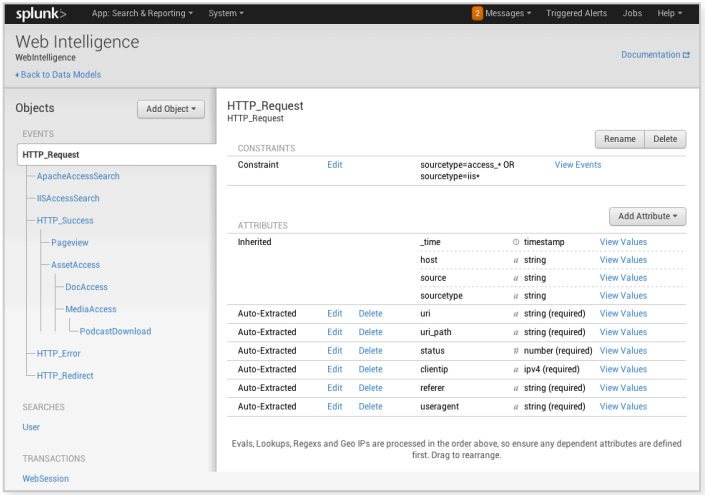

- Hierarchically structured search-time mapping of semantic knowledge about one or more datasets
  - Fields that data models use are called attributes
  - To create an effective data model, you must understand your data sources and your data semantics

A Data Model is a Collection of Objects

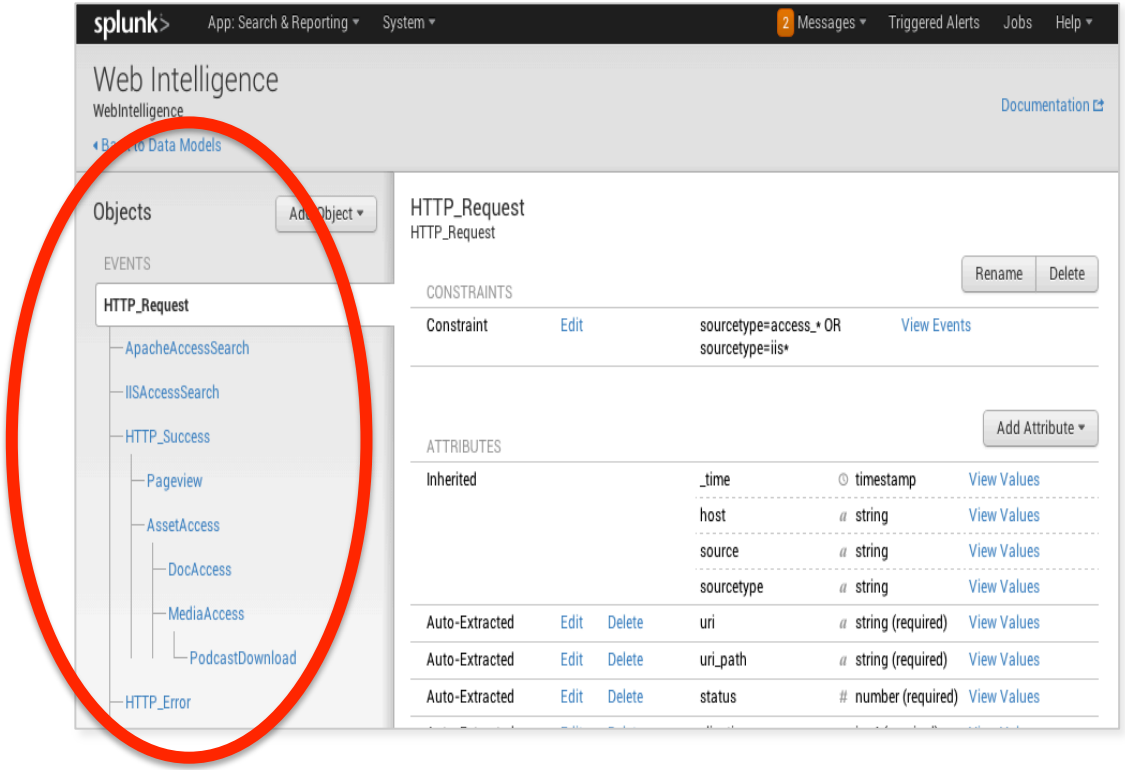

Objects Have Constraints and Attributes

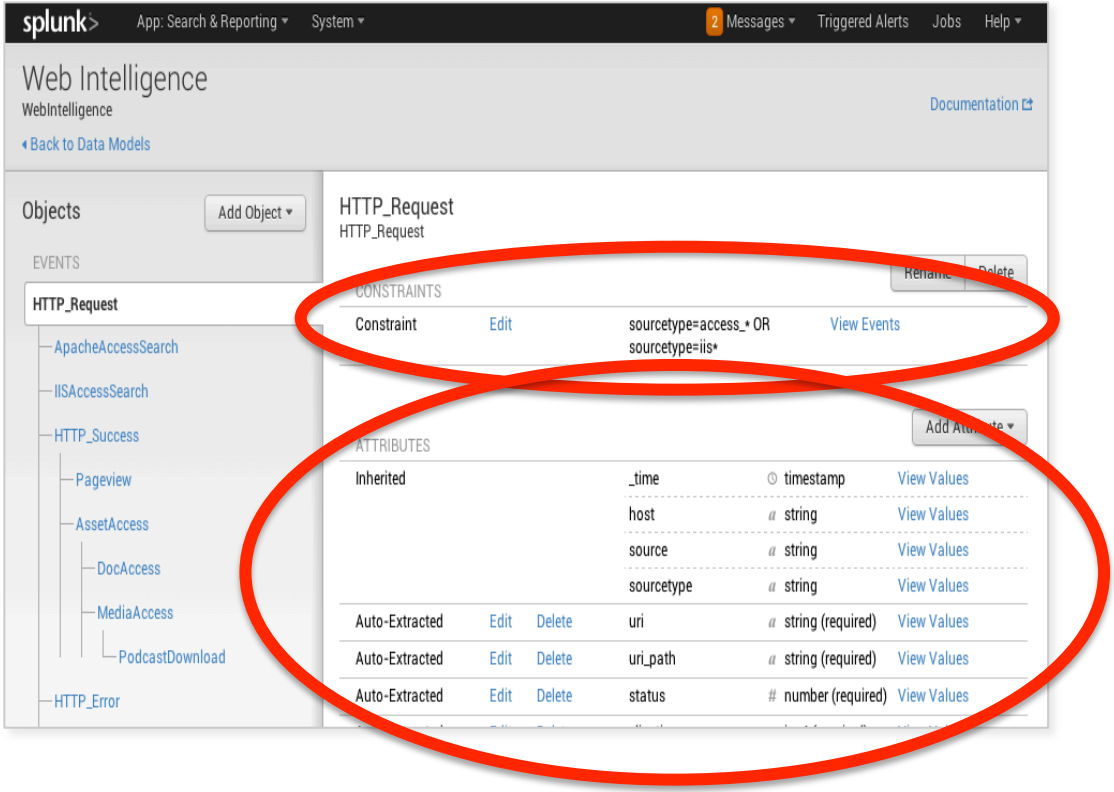

Child Objects Inherit Constraints and Attributes

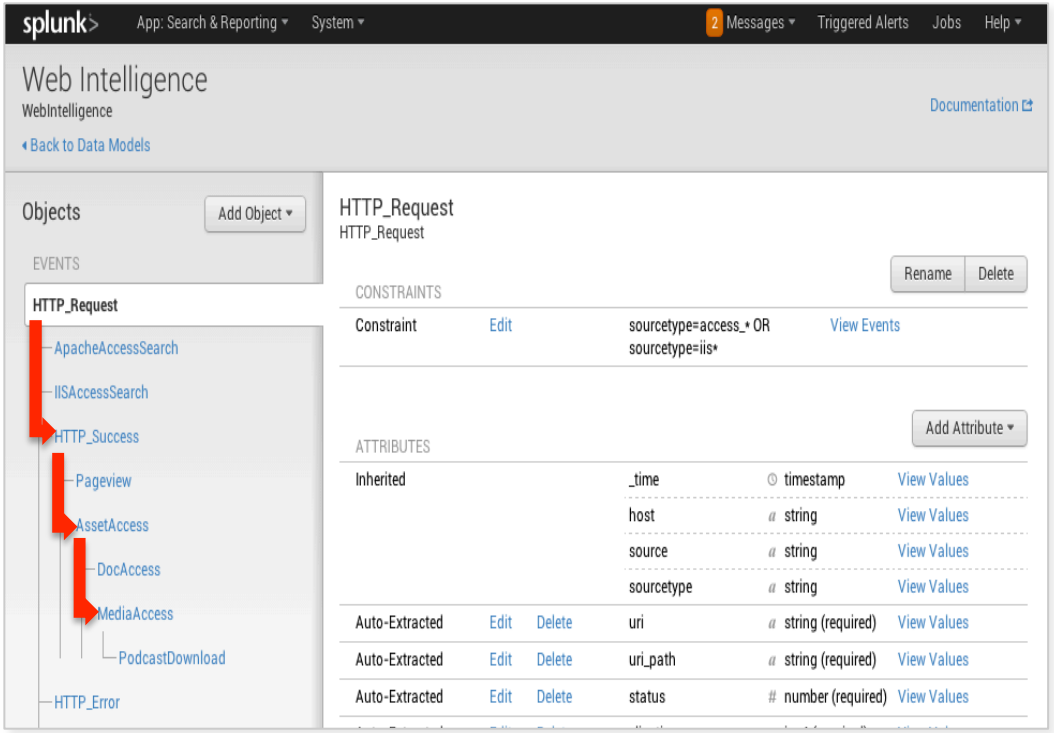

Child Objects Inherit Constraints and Attributes

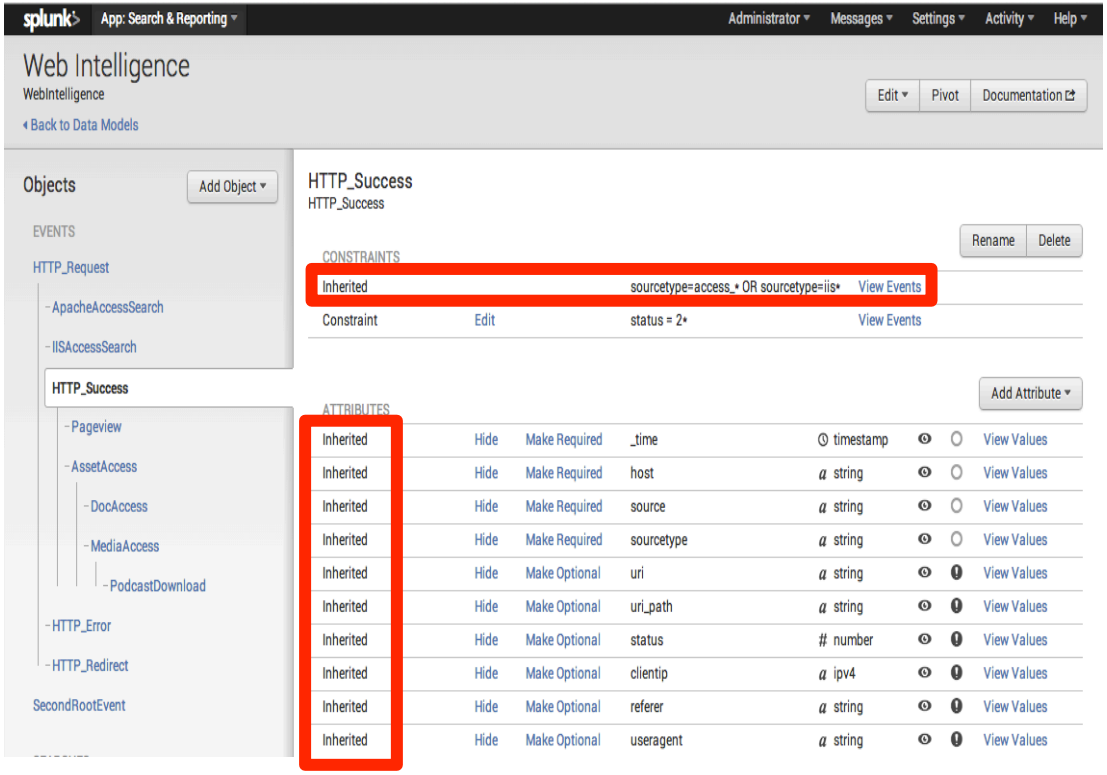

Building Data Models

splunk>

Three Root Object Types

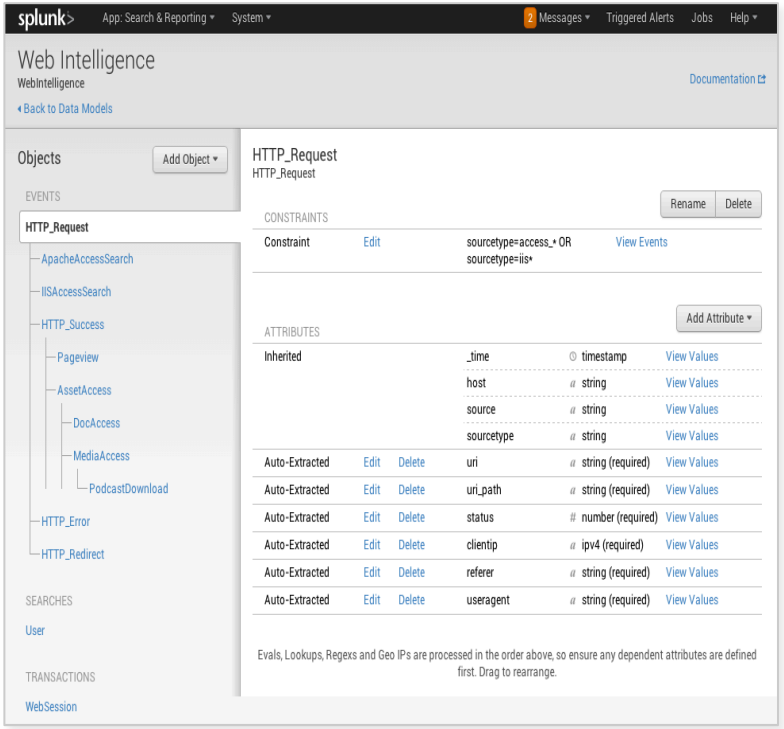

Event

- Maps to Splunk events
- Requires constraints and attributes

Three Root Object Types

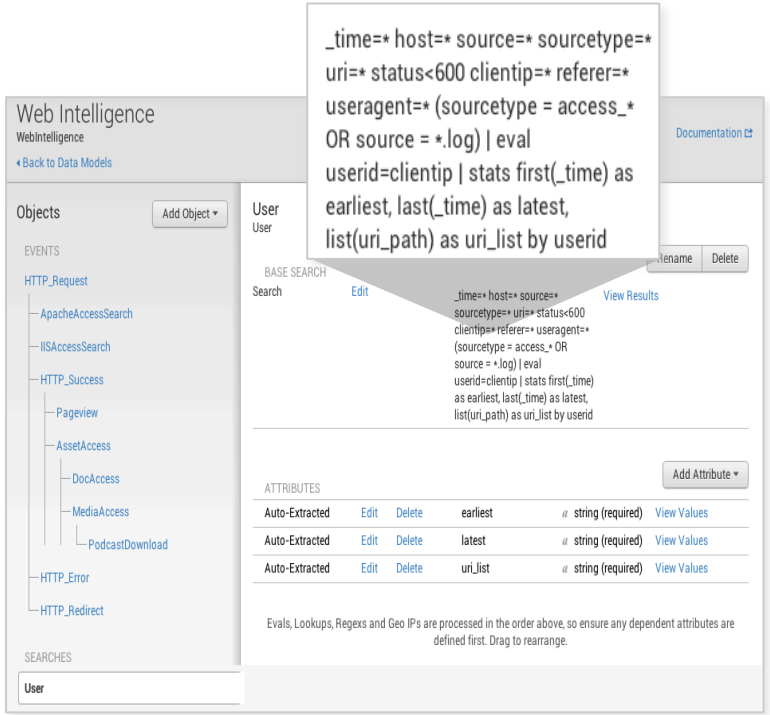

Event

- Maps to Splunk events
- Requires constraints and attributes

Search

- Maps to arbitrary Splunk search (may include generating, transforming and reporting search commands)
- Requires search string attributes

Three Root Object Types

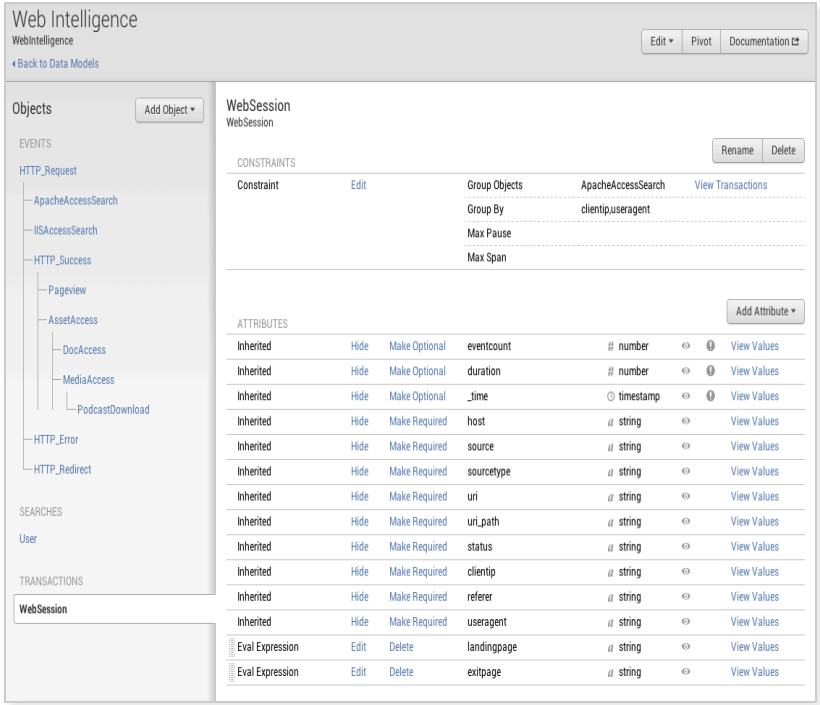

Event

- Maps to Splunk events
- Requires constraints and attributes

Search

- Maps to arbitrary Splunk search (may include generating, transforming and reporting search commands)
- Requires search string attributes

Transaction

- Maps to groups of Splunk events or groups of Splunk search results
- Requires objects to group, fields/ conditions to group by, and attributes

Anatomy of a Data Model

Data models are comprised of one or more objects

Hierarchical Parent/ Child Relationship

The child objects inherit the constraints (searches, filter on the data) and attributes.nt object

Where are they Stored?

- Each data model is a separate JSON file
- Lives in [app_name]/[local|default]/data/models
- Editing this file is not supported!

Data Model Acceleration

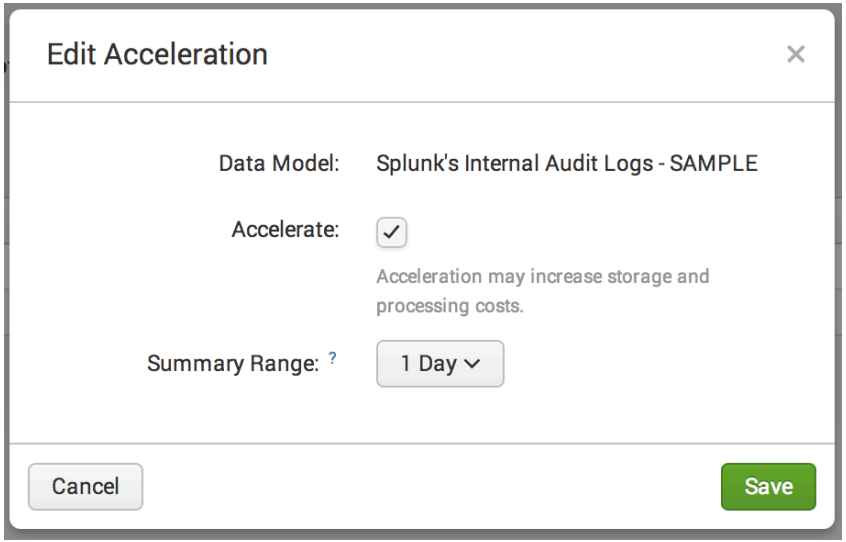

- Data model acceleration is a tool that you can use to speed up data models that represent extremely large datasets
- Data model acceleration summaries take the form of time-series index files (.tsidx) stored on Indexers
- Cannot edit accelerated models

Data Model Acceleration

Admin/Power User

Turn on acceleration via UI

Setting written to conf file

splunk> Server

Poll: are there new accelerated models?

Kick off collection

Non-technical User

Run a pivot report

ACCELERATION

Run search using on-disk acceleration

NO ACCELERATION

Kick off ad-hoc acceleration and run search

Best Practices

- Use event objects as much as possible: benefit from data model acceleration
- Minimize object hierarchy depth where possible: Constraint based filtering is less efficient as you move down the tree
- Data model acceleration is most efficient if the root event object being accelerated includes in its initial constraint search the index(es) that should be searched over

Things to Watch For…

- Data model acceleration only affects the first event object hierarchy in a data model
- Object constraints and attributes cannot contain pipes or subsearches
- Lookups used in attributes must be globally visible (or at least visible to app using the data model)
- No versioning on data models or objects

Pivot

splunk>

How Does Pivot Work?

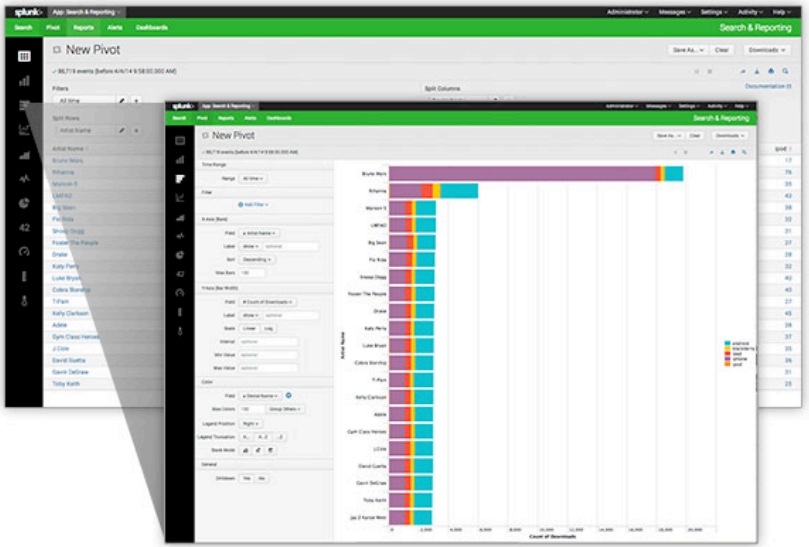

- Data models used to define the broad category of event data
- Hierarchically arranged collections of data model objects to subdivide the original datasets
- Define the attributes that you want Pivot to return results on

Pivot Editor

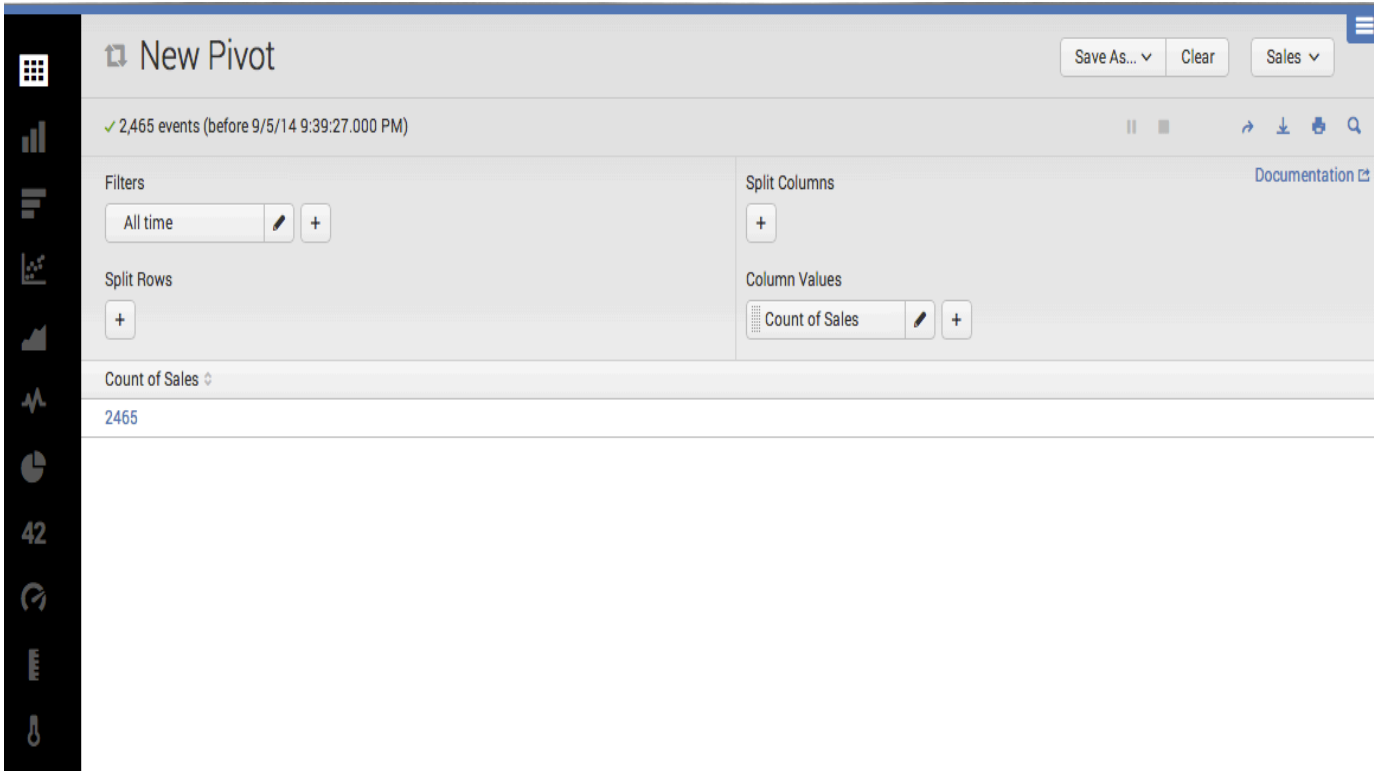

Search Commands

- Search commands that allow you to utilize data models
  - datamodel
  - tstats
  - pivot

Pivot

- Required arguments
  - datamodel-name
  - objectname
  - Pivot search
    - Has its own syntax that is different than the Splunk Search Processing Language
    - There are various elements that can be used here such as cell values, rows, columns, filters, limits, row and column formatting, and row sort options
- Command that fuels the Pivot Editor

Using the Pivot Command on a Custom Dashboard

- Create a search sing the Pivot Editor
- Click the Open in Search magnifying glass
- Create a dashboard with multiple user inputs
- Use the |pivot search you created with the Pivot Editor as the search for your dashboard panel
- Update the |pivot search with the filters you wish to provide your users

Demo

This demo will show you how to use the power of accelerated data models on a custom dashboard

splunk>

THANK YOU

splunk>