Splunk Fundamentals 1

Outline

Module 1: Introducing Splunk
Module 2: Splunk Components
Module 3: Installing Splunk
Module 4: Getting Data In
Module 5: Basic Search
Module 6: Using Fields
Module 7: Best Practices
Module 8: Splunk's Search Language
Module 9: Transforming Commands
Module 10: Creating Reports and Dashboards
Module 11: Pivot and Datasets
Module 12: Creating and Using Lookups
Module 13: Creating Scheduled Reports and Alerts

Module 1
Introducing Splunk

Understanding Splunk

- What Is Splunk?
- What Data?
- How Does Splunk Work?
- How Is Splunk Deployed?
- What are Splunk Apps?
- What are Splunk Enhanced Solutions?

What Is Splunk?

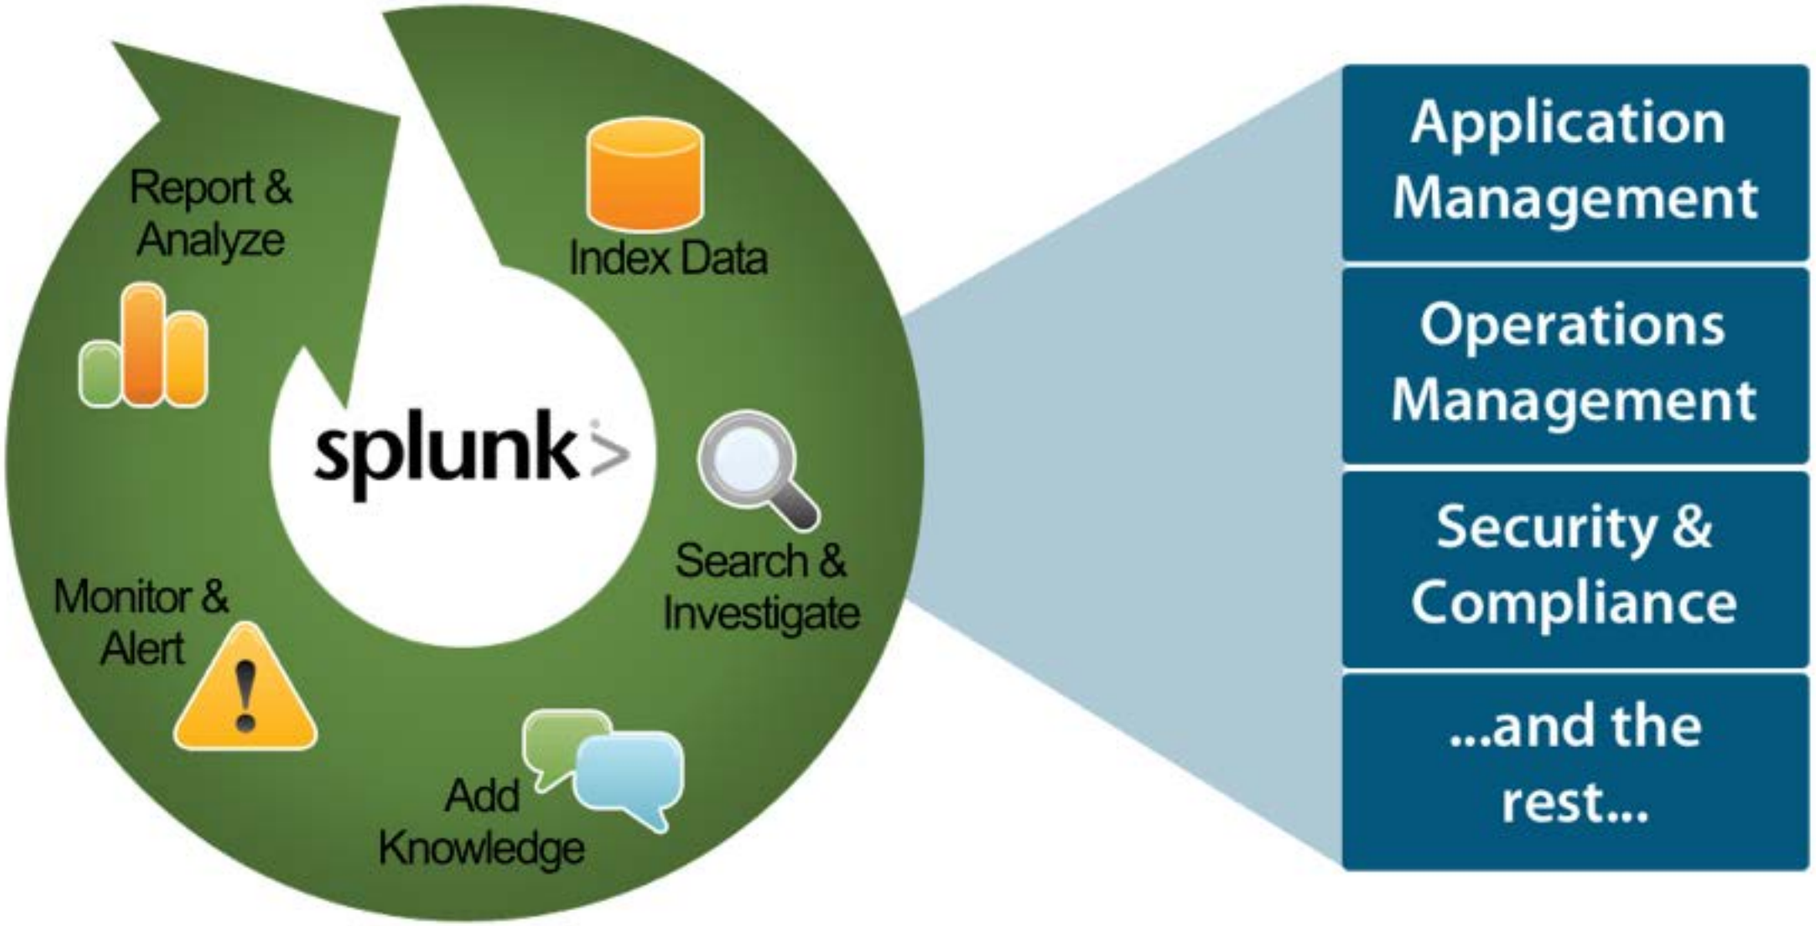

Aggregate, analyze, and get answers from your machine data

What Data?

Index **ANY** data from **ANY** source

- Computers
- Network devices
- Virtual machines
- Internet devices
- Communication devices
- Sensors
- Databases

- Logs
- Configurations
- Messages
- Call detail records
- Clickstream
- Alerts
- Metrics
- Scripts
- Changes
- Tickets

Note

For **lots** of ideas on data to collect in your environment, get the Splunk publication *The Essential Guide to Machine Data*.

How Does Splunk Work?

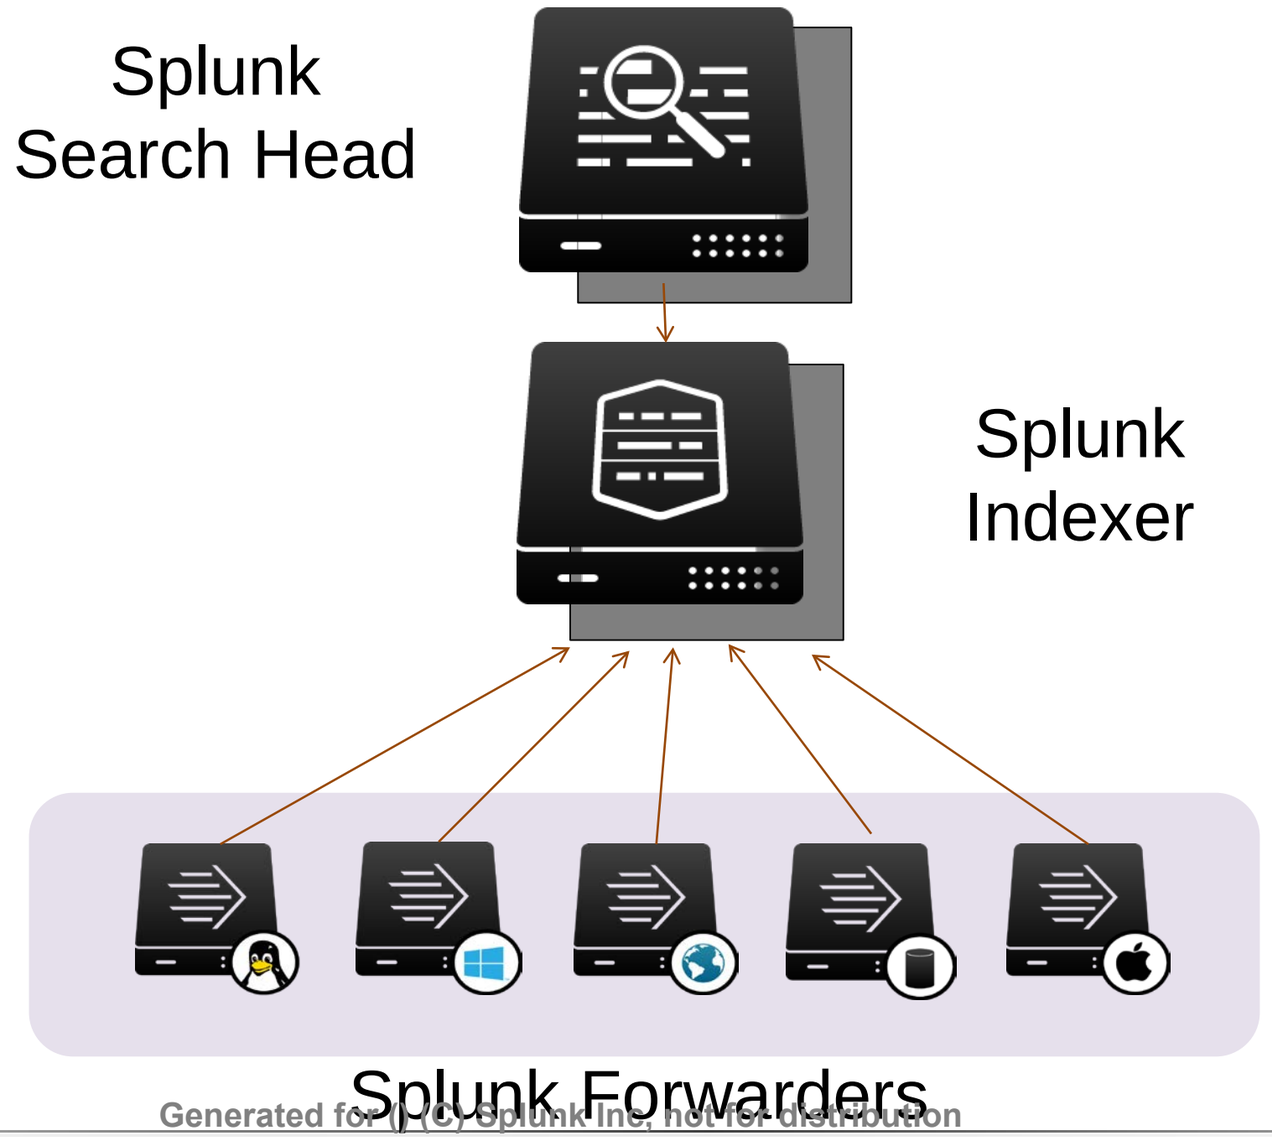

How Is Splunk Deployed?

- Splunk Enterprise

  Splunk components installed and administered on-premises

- Splunk Cloud
  - Splunk Enterprise as a scalable service
  - No infrastructure required

- Splunk Light

  Solution for small IT environments

What are Splunk Apps?

- Designed to address a wide variety of use cases and to extend the power of Splunk
- Collections of files containing data inputs, UI elements, and/or knowledge objects
- Allows multiple workspaces for different use cases/user roles to co-exist on a single Splunk instance
- 1000+ ready-made apps available on Splunkbase (splunkbase.com) or admins can build their own

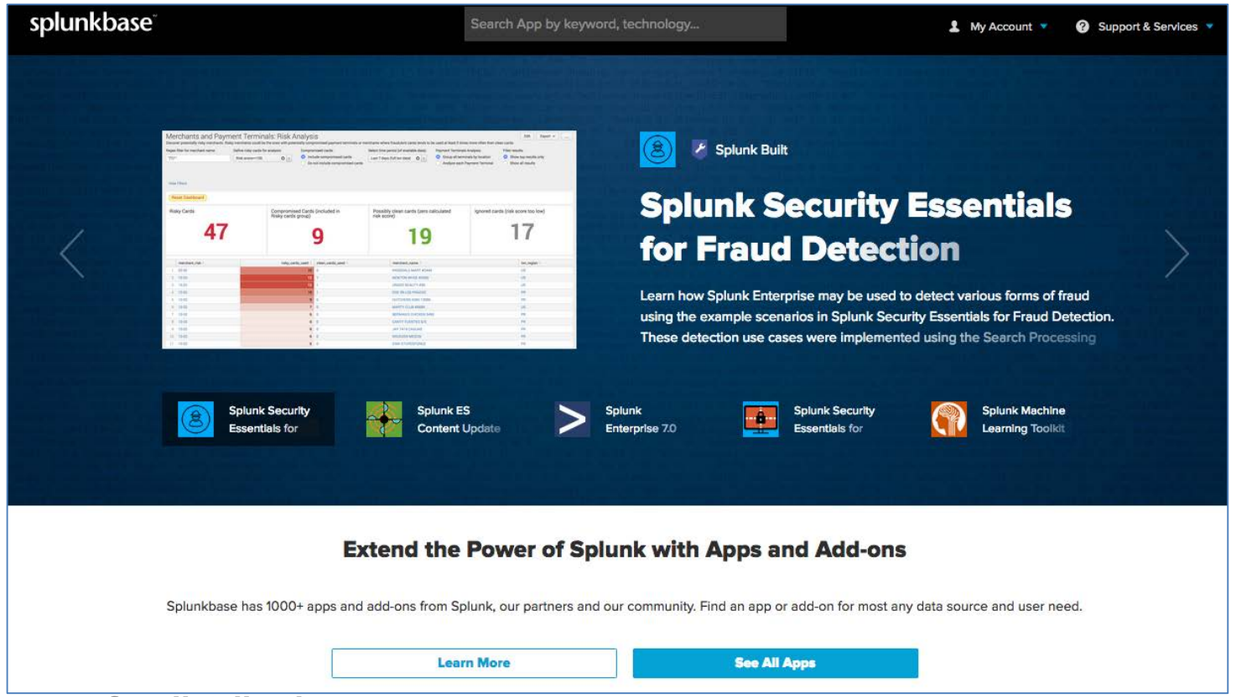

What are Splunk Enhanced Solutions?

- Splunk IT Service Intelligence (ITSI)
  - Next generation monitoring and analytics solution for IT Ops
  - Uses machine learning and event analytics to simplify operations and prioritize problem resolution

- Splunk Enterprise Security (ES)
  - Comprehensive Security Information and Event Management (SIEM) solution
  - Quickly detect and respond to internal and external attacks

- Splunk User Behavior Analytics (UBA)

  Finds known, unknown, and hidden threats by analyzing user behavior and flagging unusual activity

Note

For more info, see Appendix A: Splunk Premium Solutions and Apps.

Users and Roles

- Splunk users are assigned roles, which determine their capabilities and data access
- Out of the box, there are 3 main roles:
  - Admin
  - Power
  - User
- Splunk admins can create additional roles

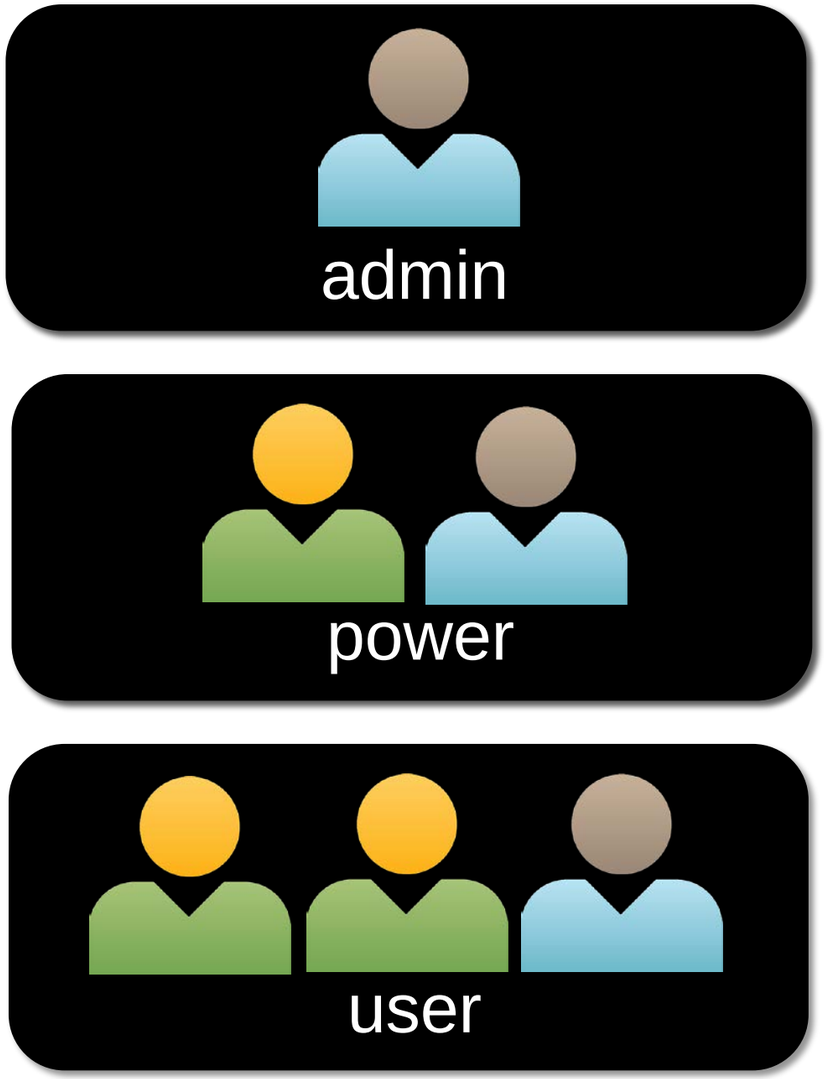

Note

In this class, the account you'll use for the lab exercises has been assigned the **Power** role.

Logging In

1. Log into Splunk with a web browser
2. The main view of your default app appears

   You or your organization may change your default app

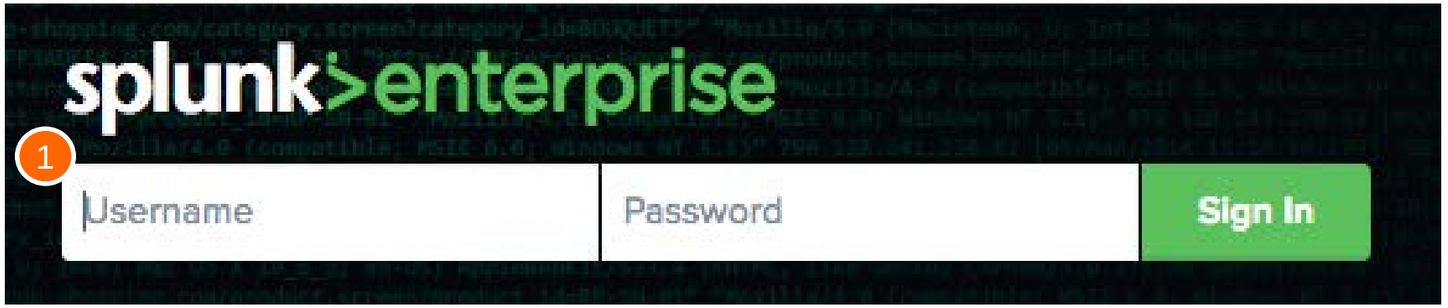

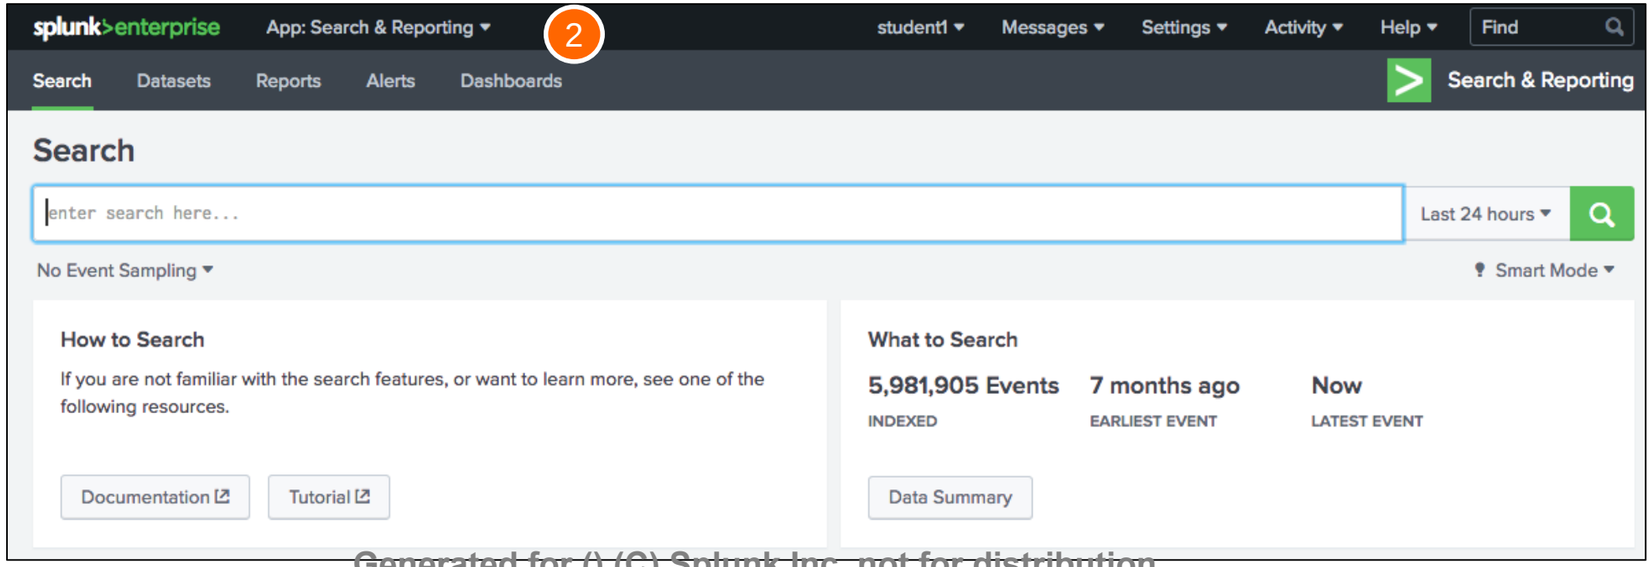

Choosing Your App

- Apps allow different workspaces for specific use cases or user roles to co-exist on a single Splunk instance
- In this class, you'll explore:
  - The Home app
  - The Search & Reporting app (also called the Search app)

Home app

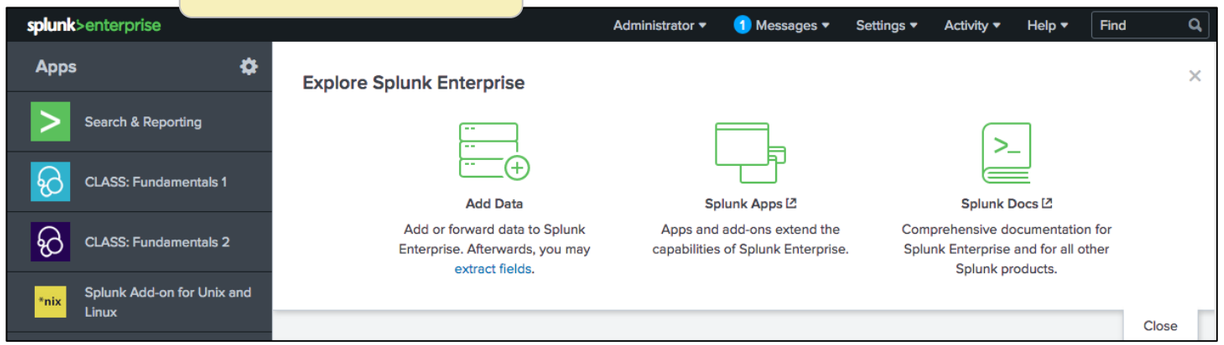

Within an app

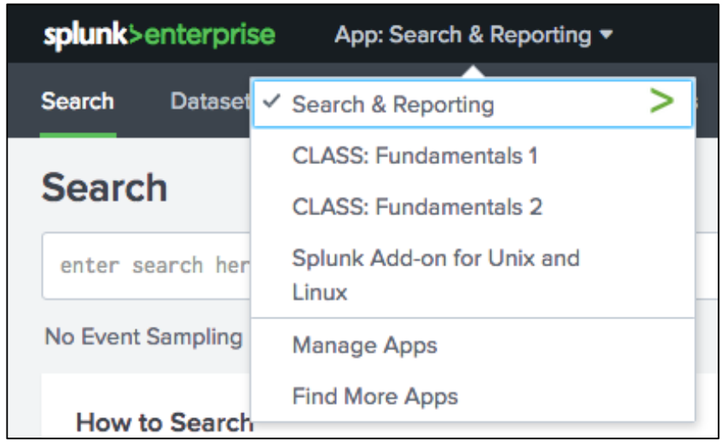

Note

For more info on apps, see.
http://docs.splunk.com/Documentation/Splunk/latest/Admin/Whatsanapp

Home App

You can always click the Splunk logo to return to whatever app is set as your default app.

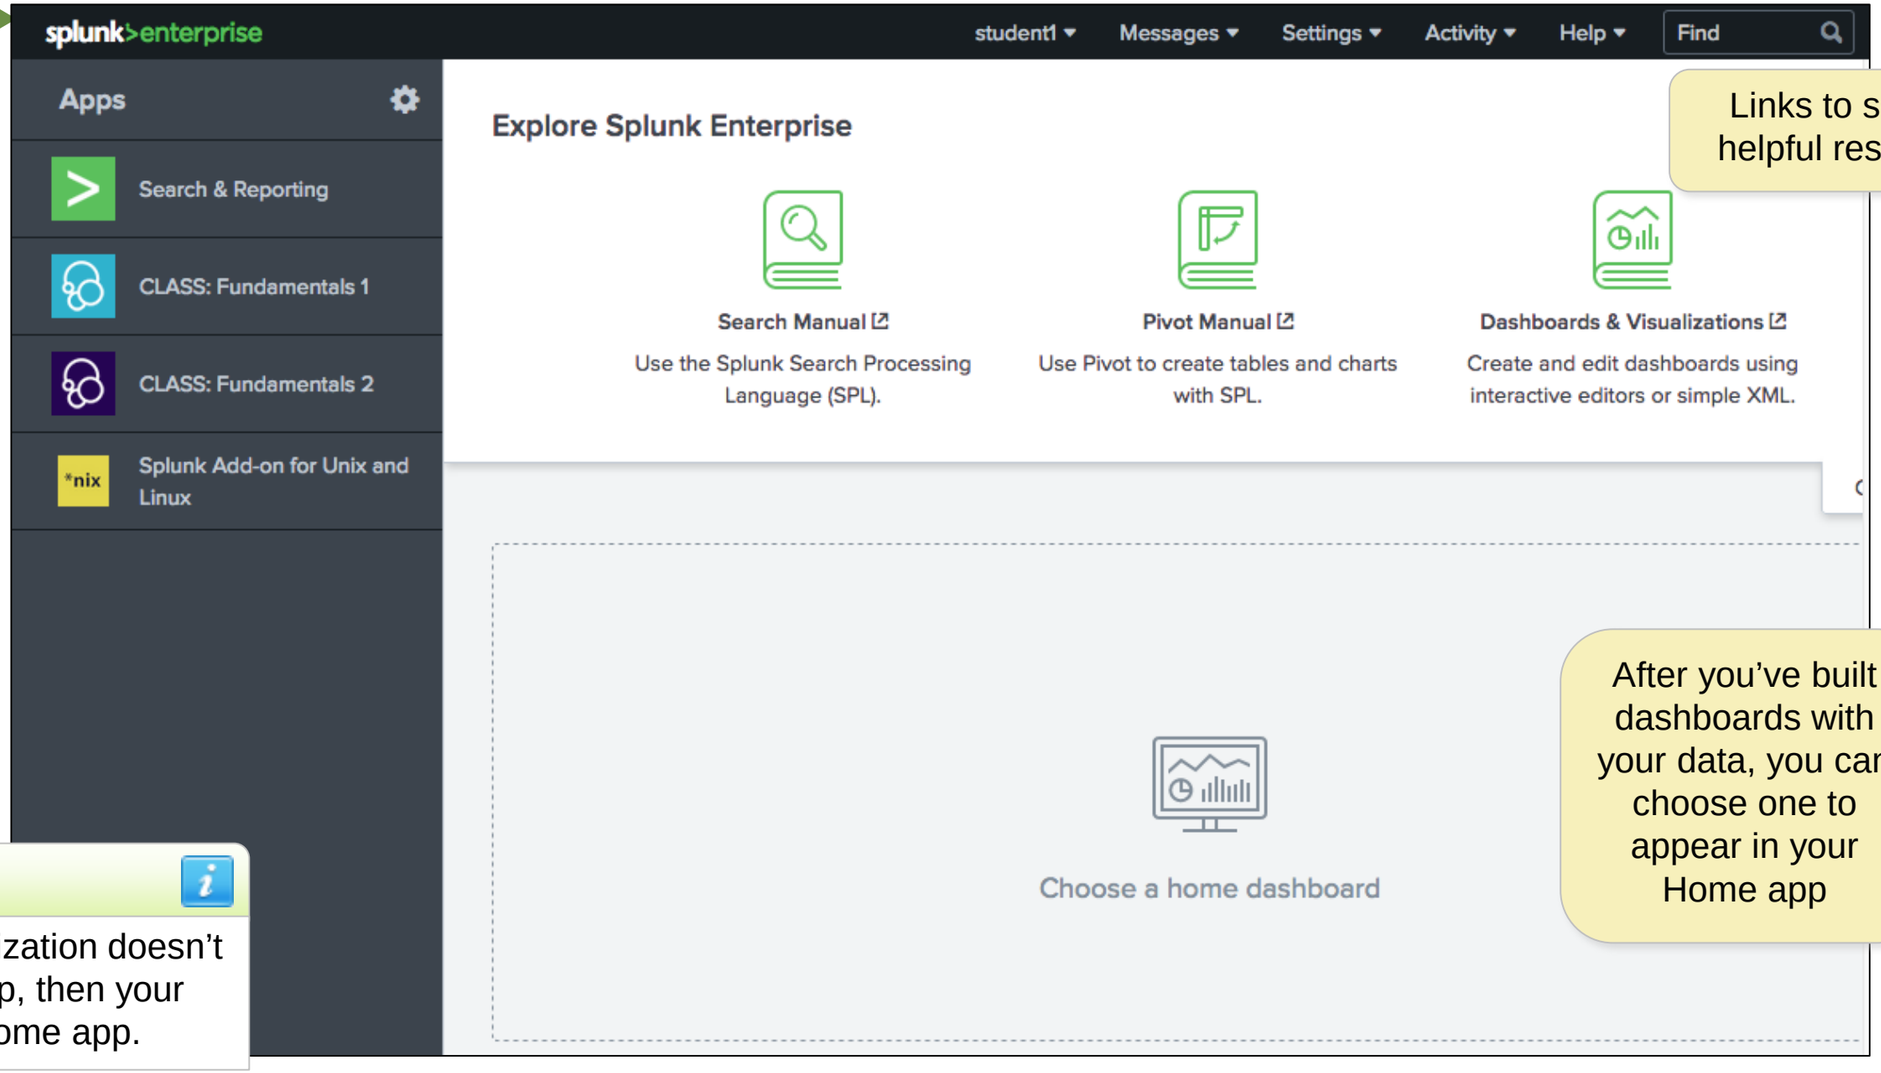

Links to several helpful resources

After you've built dashboards with your data, you can choose one to appear in your Home app

Note

If you or your organization doesn't choose a default app, then your default app is the Home app.

Search & Reporting App

- Provides a default interface for searching and analyzing data
- Enables you to create knowledge objects, reports, and dashboards
- Access by selecting the **Search & Reporting** button on the Home app or from an app view, select **Apps** > **Search & Reporting**

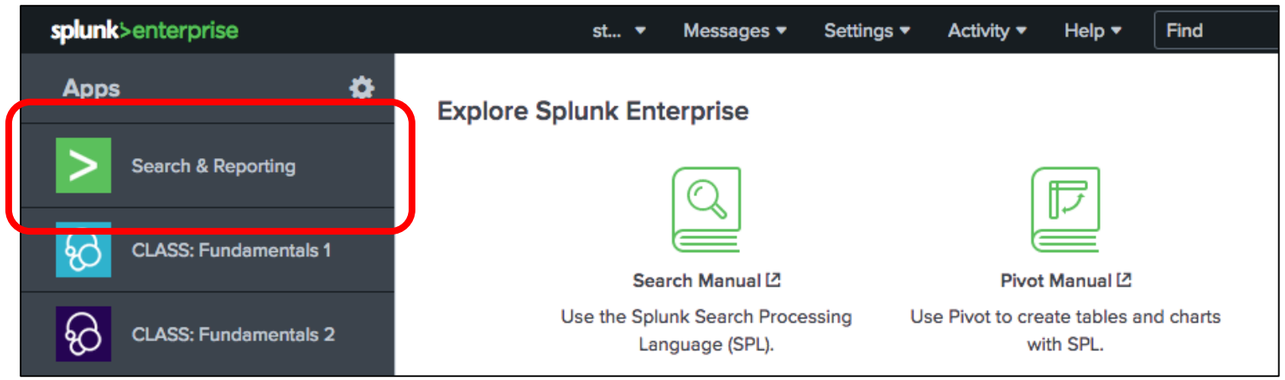

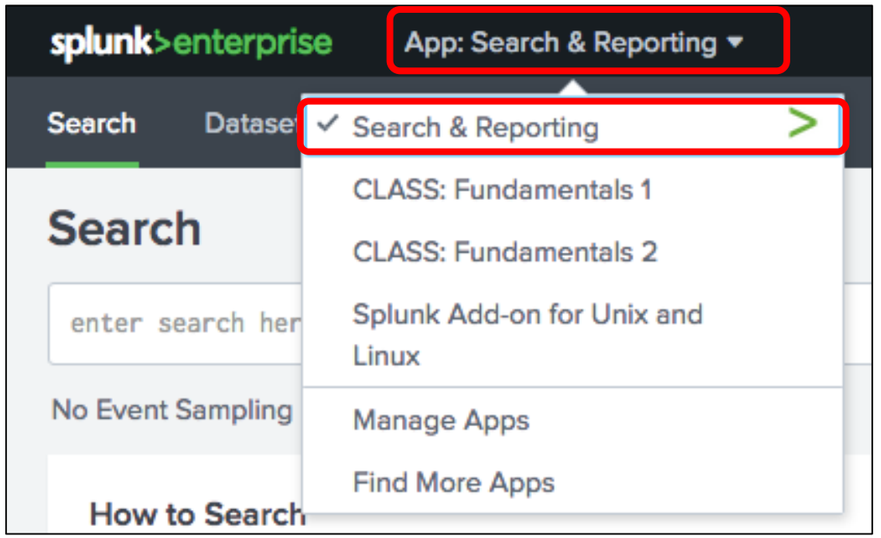

Search & Reporting App (cont.)

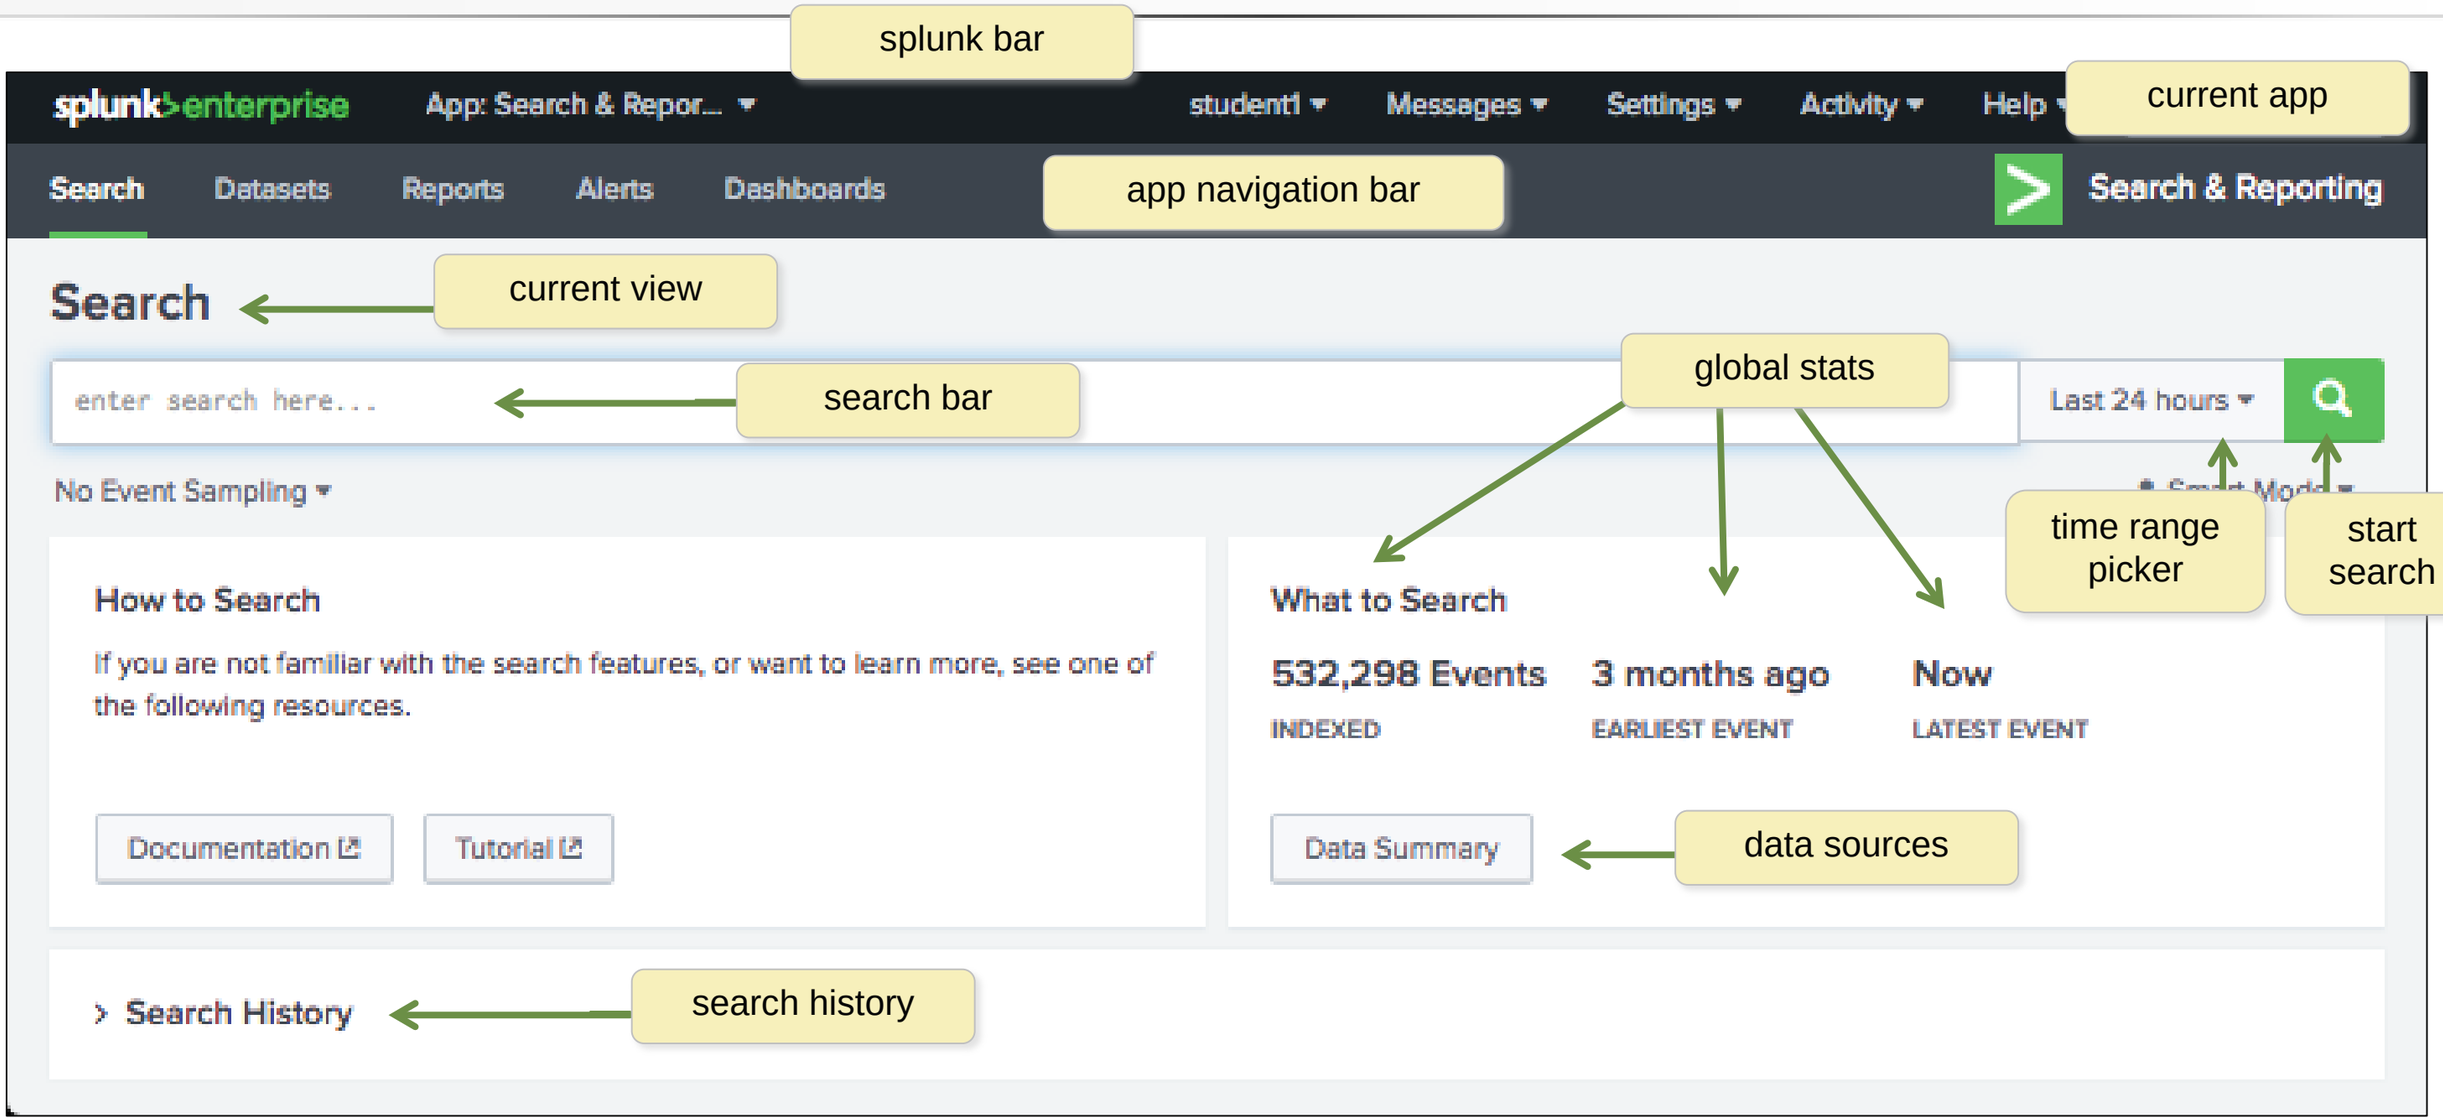

Data Summary Tabs

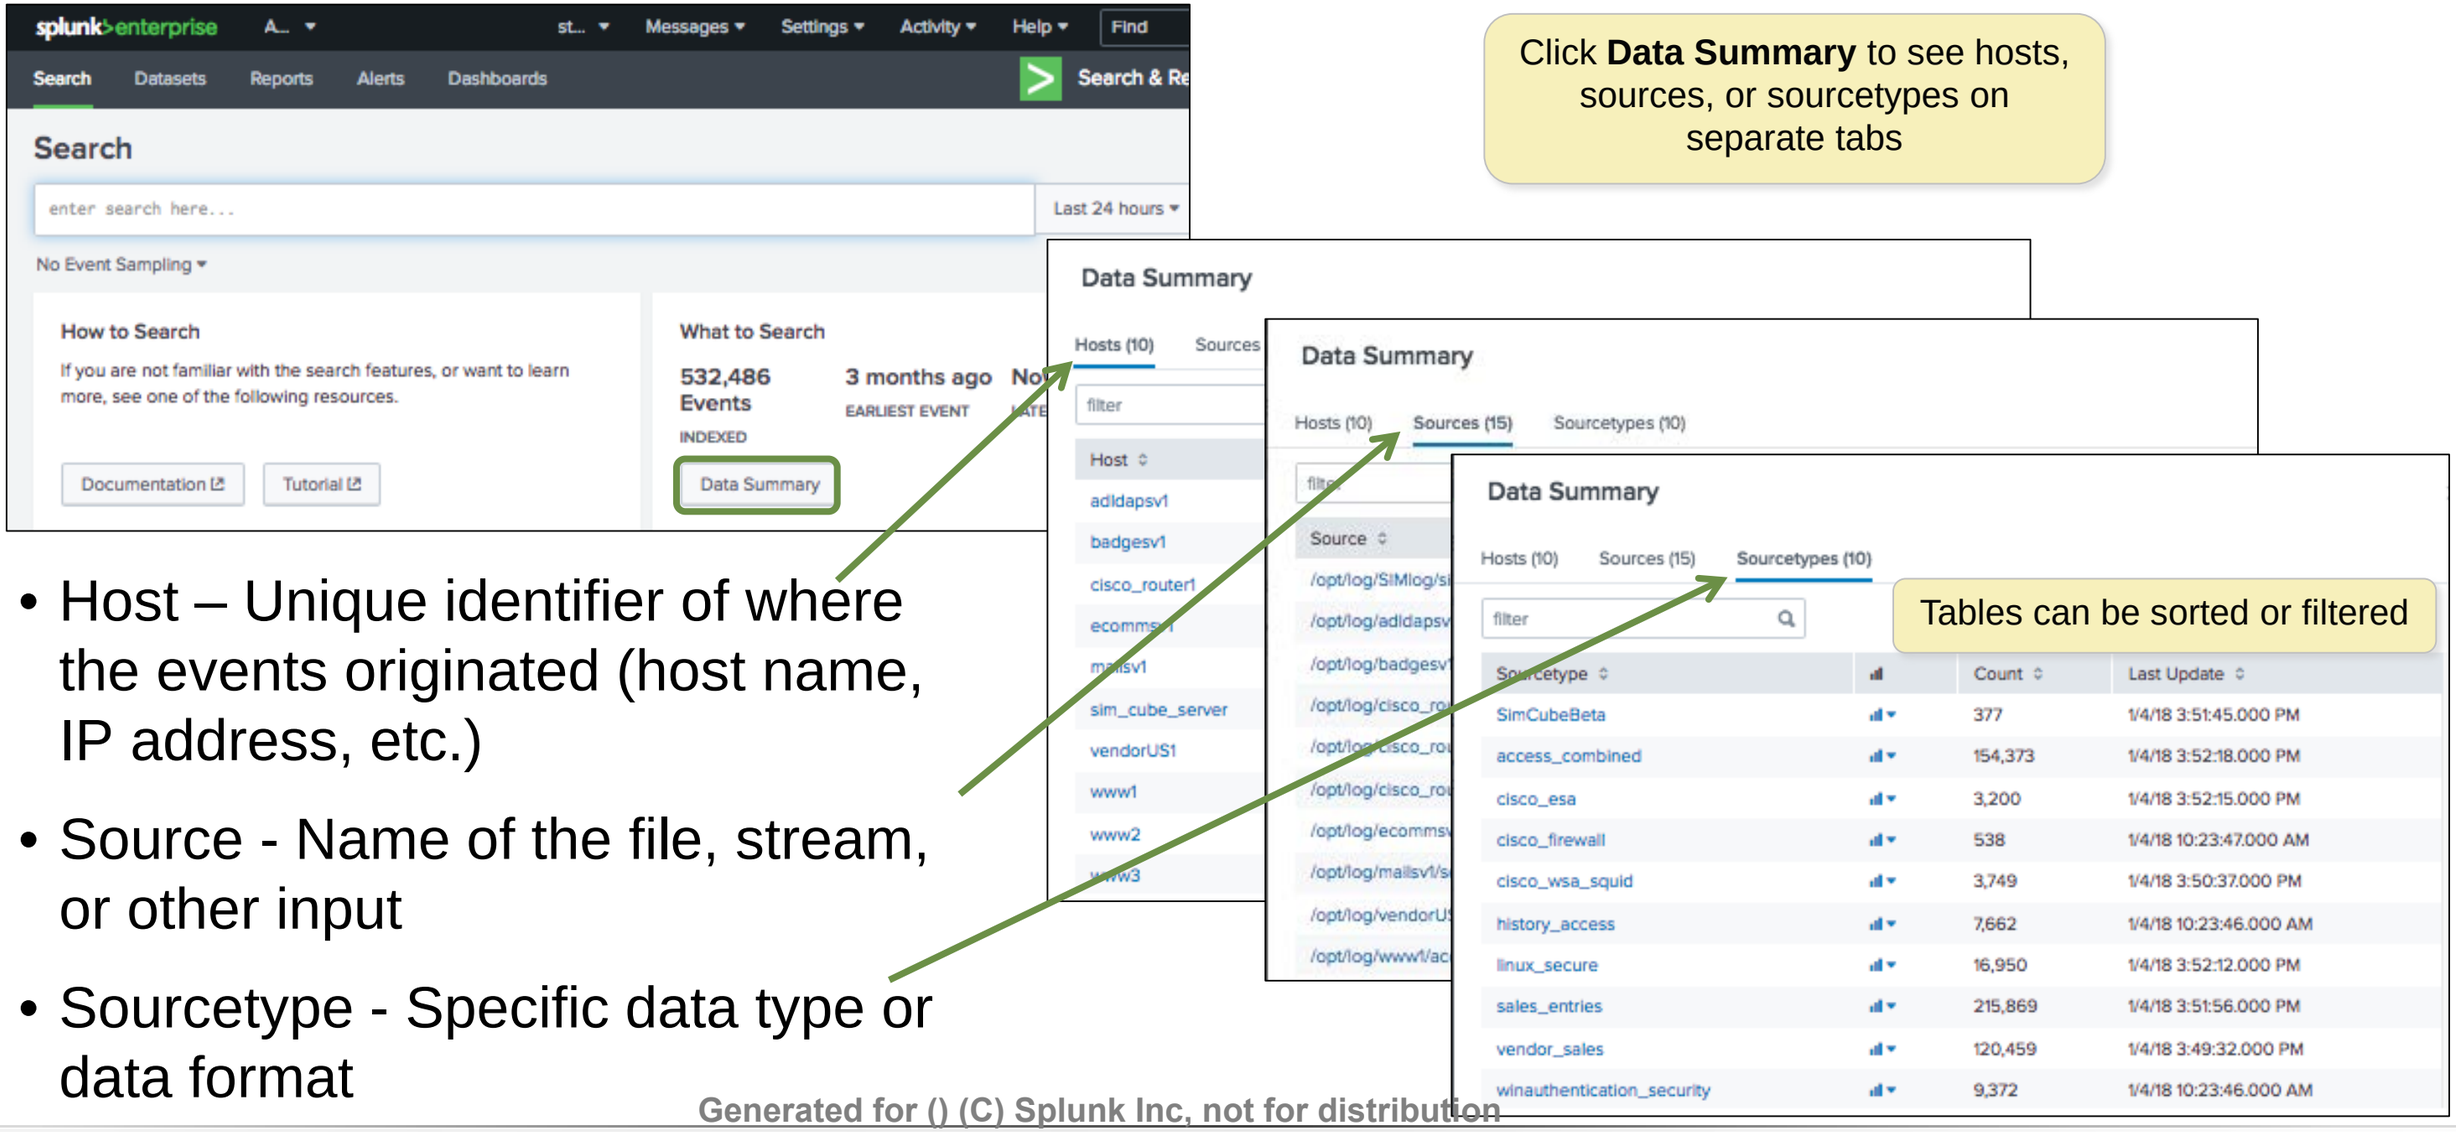

- Host – Unique identifier of where the events originated (host name, IP address, etc.)
- Source - Name of the file, stream, or other input
- Sourcetype - Specific data type or data format

Events Tab

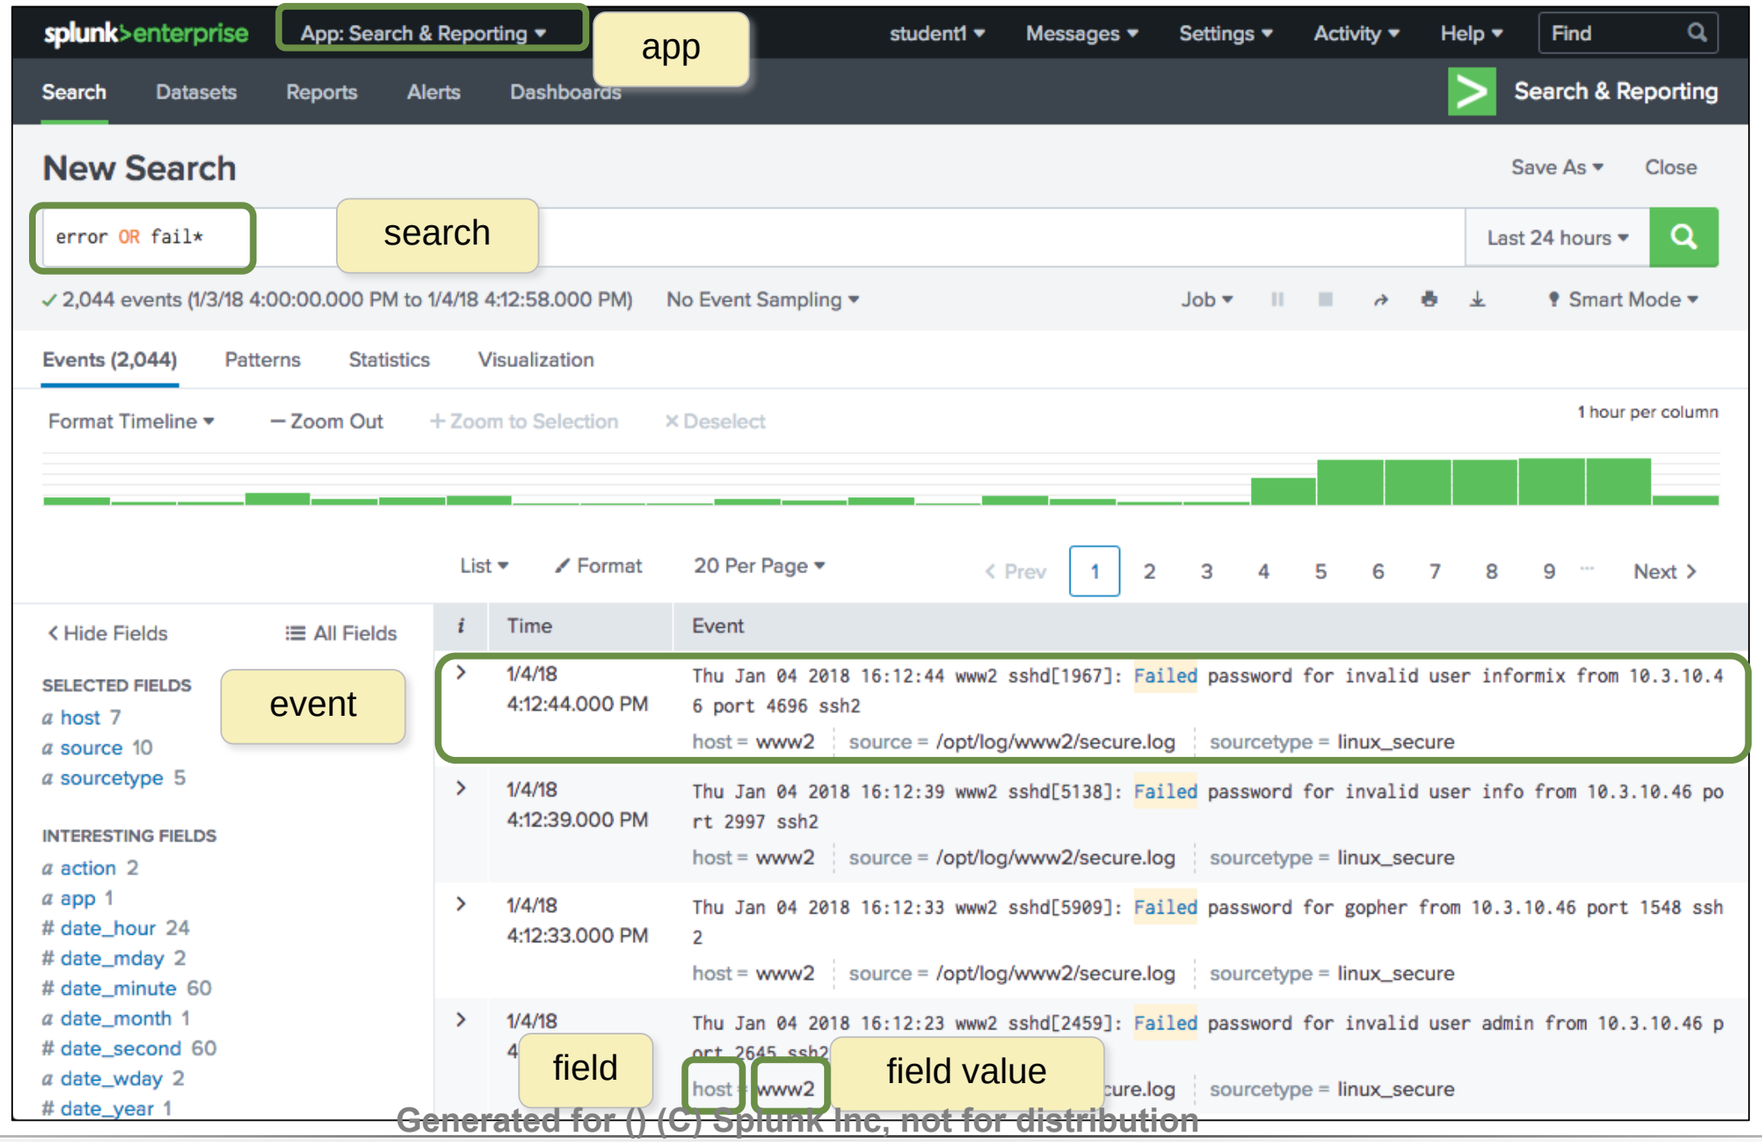

Course Scenario

- Use cases in this course are based on Buttercup Games, a fictitious gaming company
- Multinational company with its HQ in San Francisco and offices in Boston and London
- Sells products through its worldwide chain of 3rd party stores and through its online store

Your Role at Buttercup Games

- You're a Splunk power user
- You're responsible for providing info to users throughout the company
- You gather data/statistics and create reports on:
  - IT operations: information from mail and internal network data
  - Security operations: information from internal network and badge reader data
  - Business analytics: information from web access logs and vendor data

Callouts

Scenarios

- Many of the examples in this course relate to a specific scenario
- For each example, a question is posed from a colleague or manager at Buttercup Games

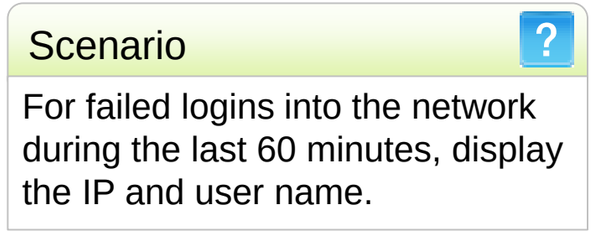

Notes & Tips

References for more information on a topic and tips for best practices

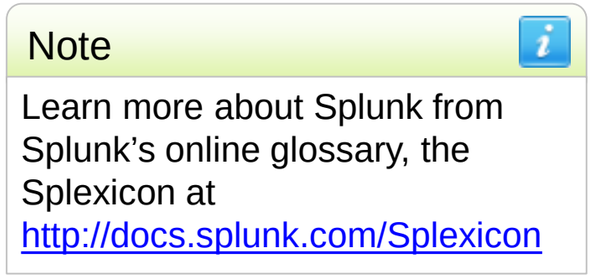

Module 2:
Splunk Components

Splunk Components

Splunk is comprised of three main processing components:

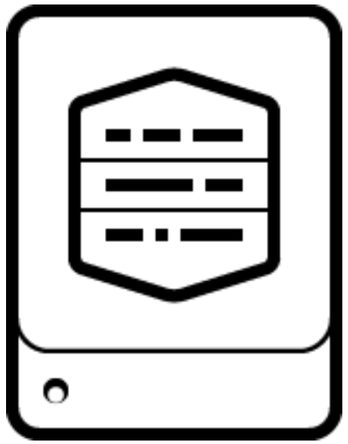

Indexer

Search Head

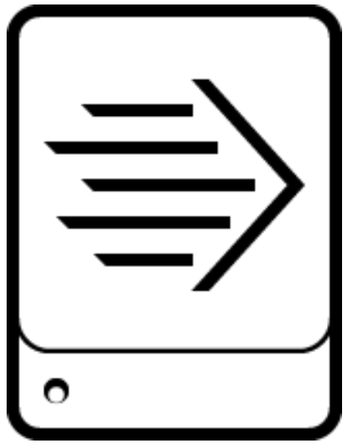

Forwarder

Splunk Components - Indexer

- Processes machine data, storing the results in indexes as events, enabling fast search and analysis

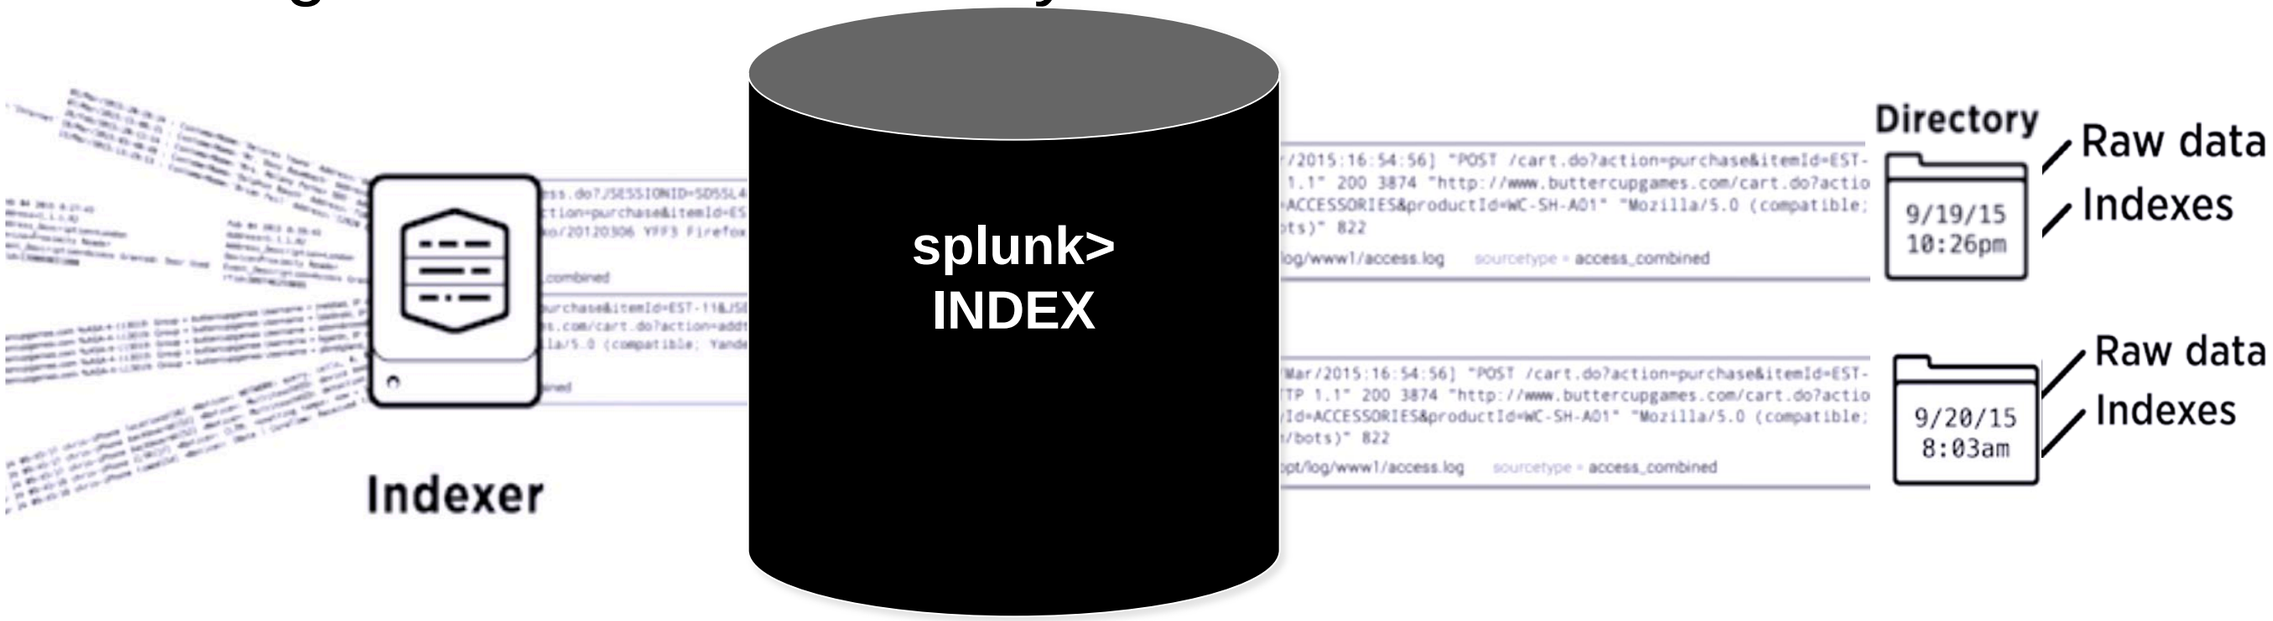

- As the Indexer indexes data, it creates a number of files organized in sets of directories by age
  - Contains raw data (compressed) and indexes (points to the raw data)

Splunk Components – Search Heads

- Allows users to use the Search language to search the indexed data
- Distributes user search requests to the Indexers
- Consolidates the results and extracts field value pairs from the events to the user
- Knowledge Objects on the Search Heads can be created to extract additional fields and transform the data without changing the underlying index data

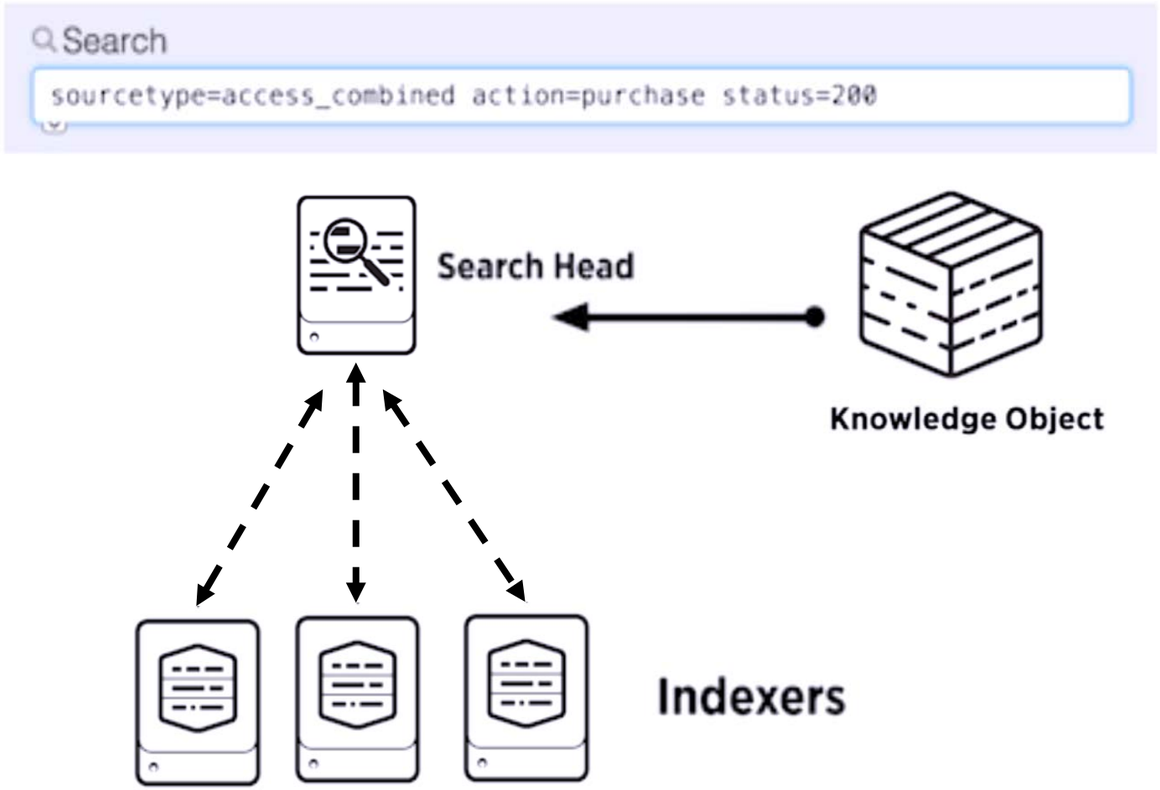

Splunk Components – Search Heads (cont.)

Search Heads also provide tools to enhance the search experience such as reports, dashboards and visualizations

user	src_ip	app
nsharpe	68.144.35.139	sshd
zabbix	223.205.219.67	sshd
djohnson		sshd
whois	121.9.245.177	sshd
nsharpe		sshd
admin	200.6.134.23	
jean-luc	223.5.16.102	
inet	91.217.178.210	
mysql	217.132.169.69	

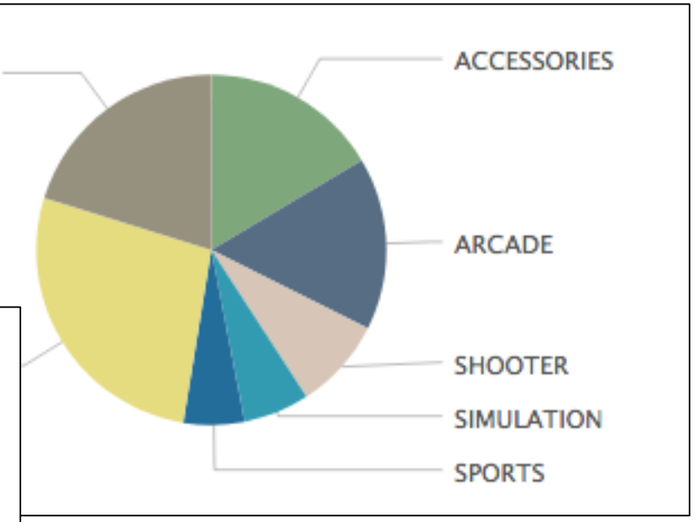

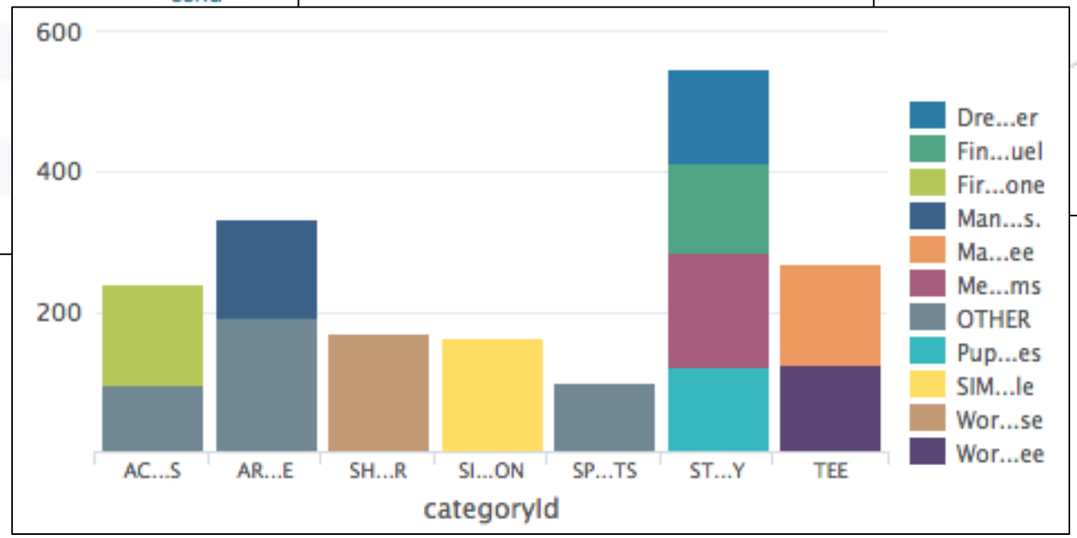

Splunk Components – Forwarders

- Splunk Enterprise instances that consume and send data to the index
- Require minimal resources and have little impact on performance
- Typically reside on the machines where the data originates
- Primary way data is supplied for indexing

IP = 10.3.10.6, Session disconnected. Session type = TPsecOve
IP = 10.1.10.216, Session connected. Session type = SSL, Dura
s, IP = 10.1.10.133, Session connected. Session type = IKE, Du
i, IP = 10.3.10.18, Session disconnected. Session type = IKE, D
= 10.1.10.211, Session connected. Session type = SSL, Duratio

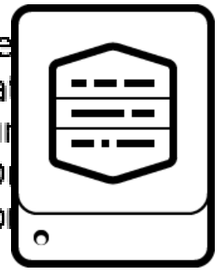

Web Server with Forwarder instance installed

Indexer

Additional Splunk Components

In addition to the three main Splunk processing components, there are some less-common components including :

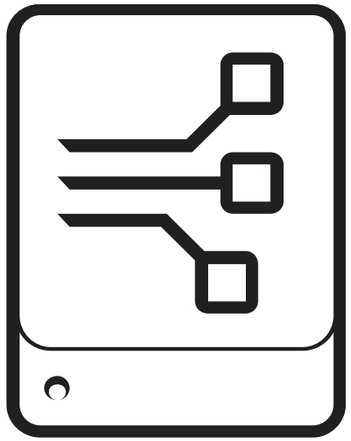

Deployment Server

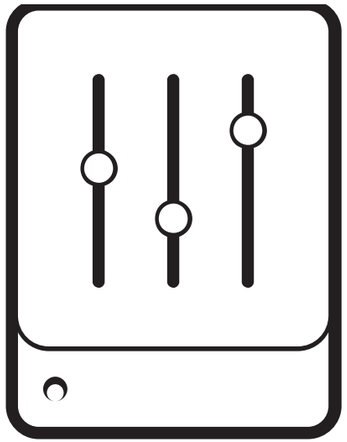

Cluster Master

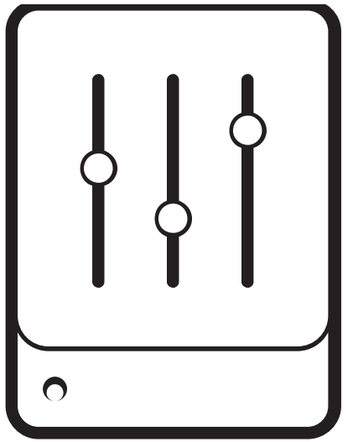

License Master

Splunk Deployment – Standalone

- **Single Server**
  - All functions in a single instance of Splunk
  - For testing, proof of concept, personal use, and learning
  - This is what you get when you download Splunk and install with default settings
- Recommendation
  - Have at least one test/development setup at your site

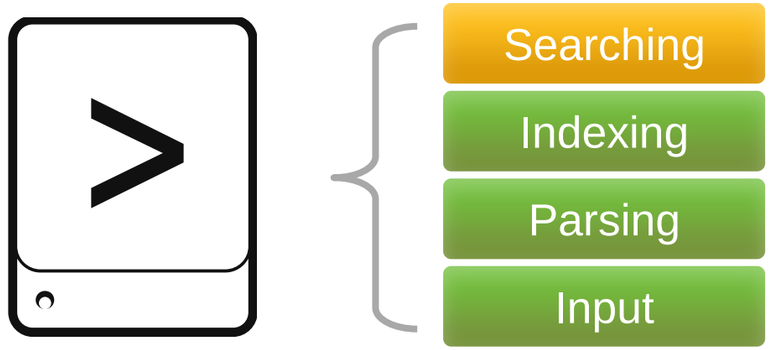

Splunk Deployment – Basic

- Splunk server
  - Similar to server in standalone configuration
  - Manage deployment of forwarder configurations
- Forwarders
  - Forwarders collect data and send it to Splunk servers
  - Install forwarders at data source (usually production servers)

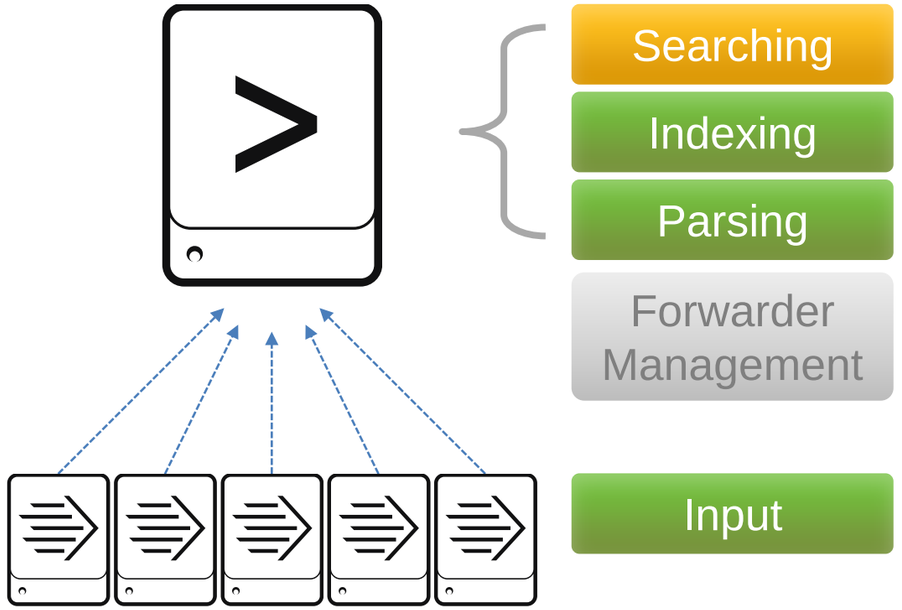

Basic Deployment for organizations:

- Indexing less than 20GB per day
- With under 20 users
- Small amount of forwarders

Splunk Deployment – Multi-Instance

- Increases indexing and searching capacity
- Search management and index functions are split across multiple machines

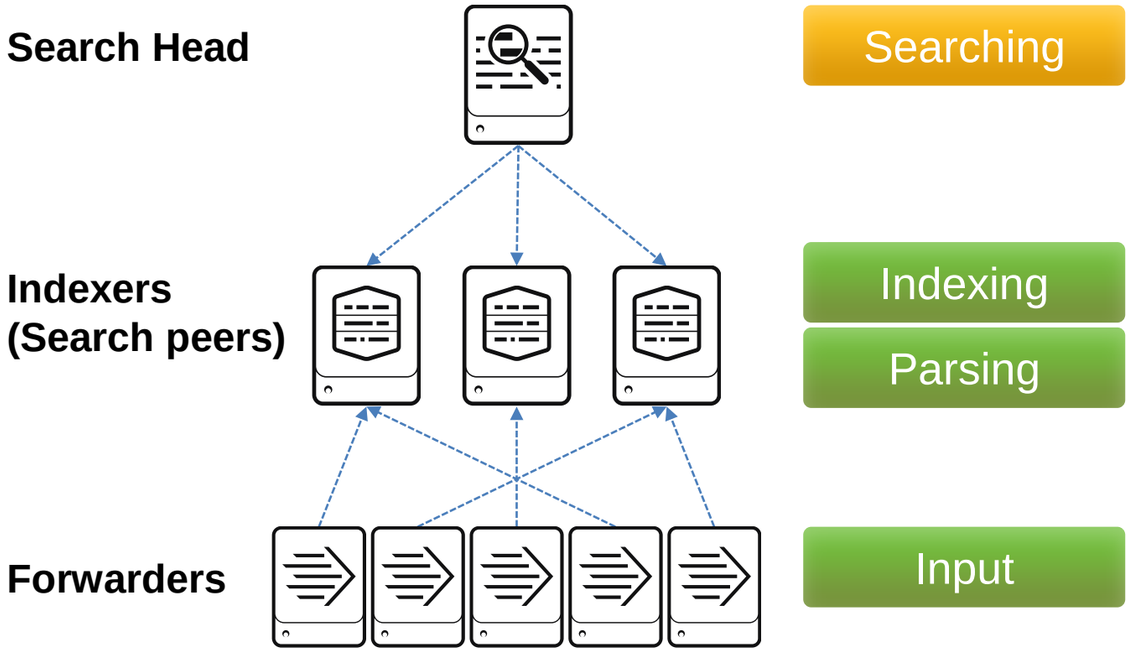

Deployment for organizations:

- Indexing up to 100 GB per day
- Supports 100 users
- Supports several hundred forwarders

Splunk Deployment – Increasing Capacity

- Adding a Search Head Cluster:
  - Services more users for increased search capacity
  - Allows users and searches to share resources
  - Coordinate activities to handle search requests and distribute the requests across the set of indexers
- Search Head Clusters require a minimum of three Search Heads
- A Deployer is used to manage and distribute apps to the members of the Search Head Cluster

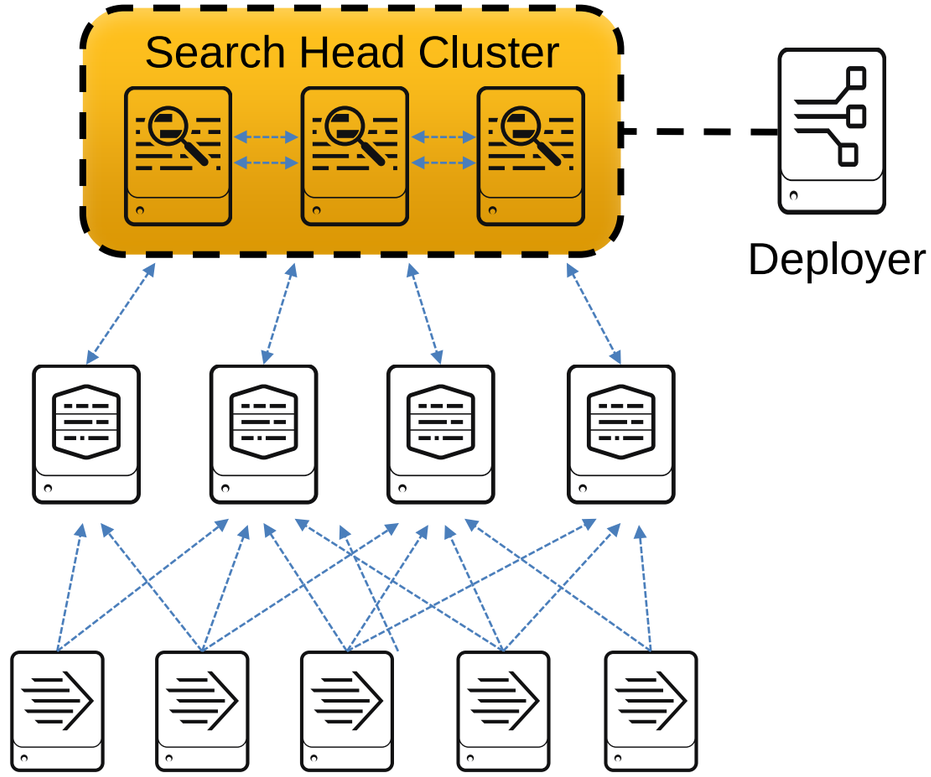

Splunk Deployment – Index Cluster

- Traditional Index Clusters:
  - Configured to replicate data
  - Prevent data loss
  - Promote availability
  - Manage multiple indexers
- Non-replicating Index Clusters
  - Offer simplified management
  - Do not provide availability or data recovery

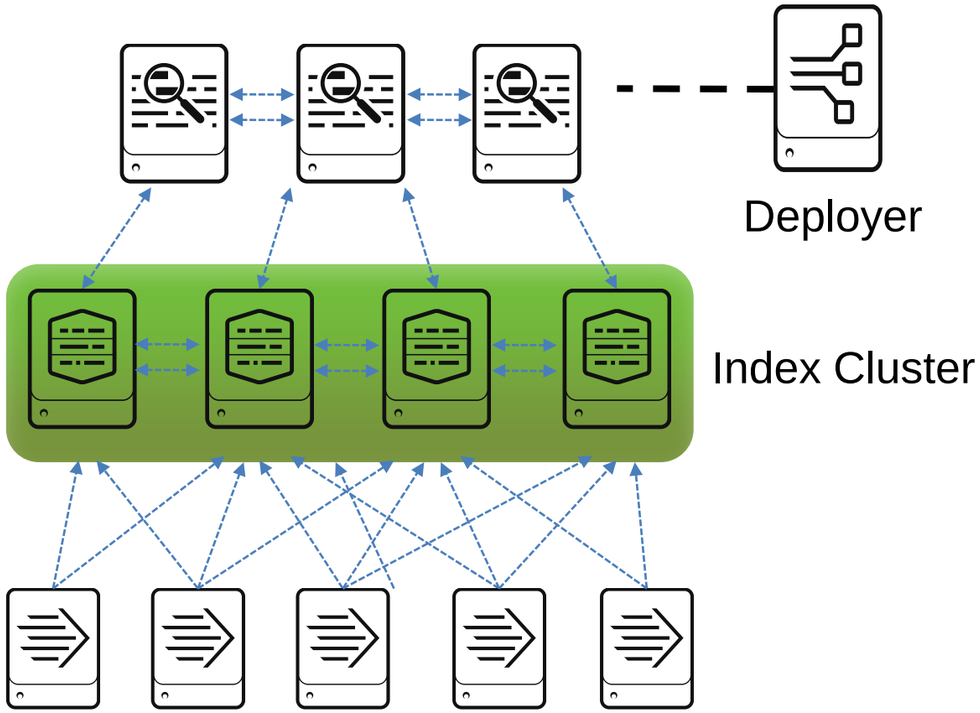

Module 3: Installing Splunk

Splunk Enterprise Install Package

There are multiple Splunk components installed from the Splunk Enterprise package

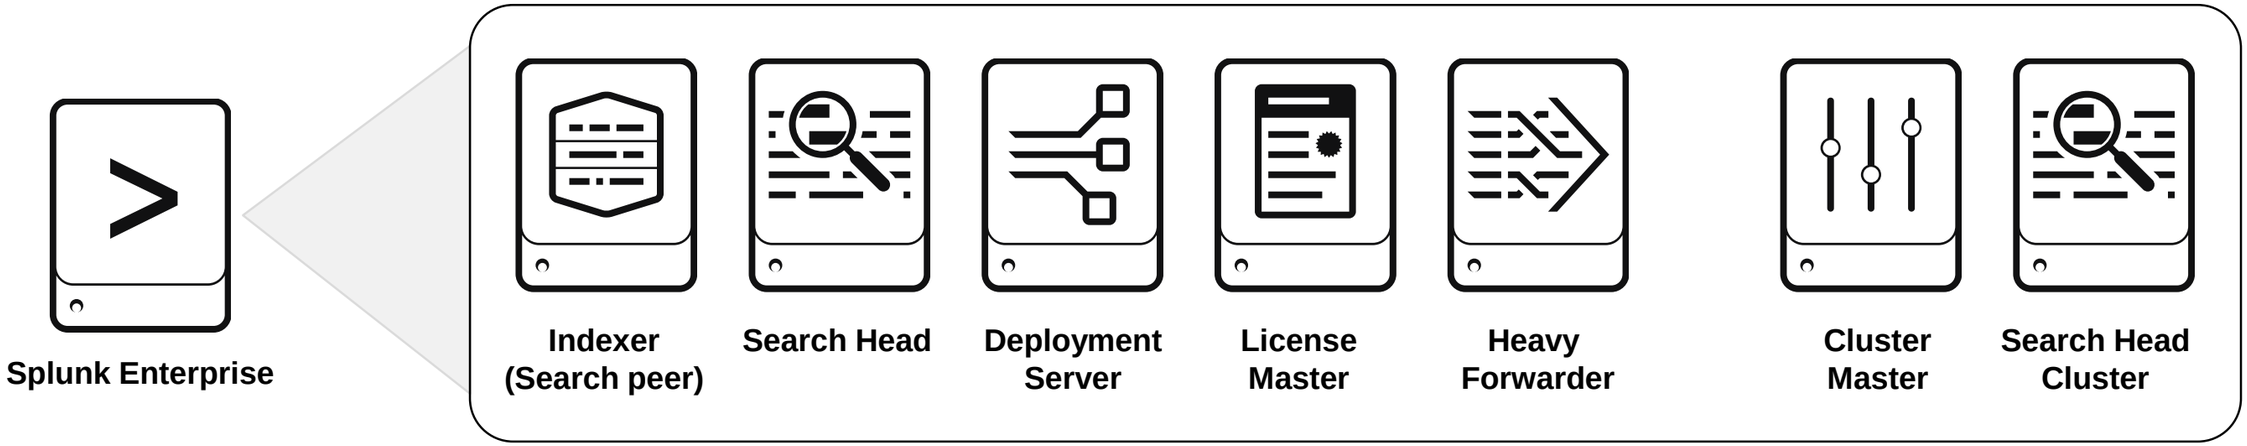

Splunk Enterprise Installation Overview

- Verify required ports are open (splunkweb, splunkd, forwarder) and start-up account
- Download Splunk Enterprise from www.splunk.com/download
- Installation: (as account running Splunk)
  - *NIX – un-compress the .tar.gz file in the path you want Splunk to run from
  - Windows – execute the .msi installer and follow the wizard steps
- Complete installation instructions at:

  docs.splunk.com/Documentation/Splunk/latest/Installation/Chooseyourplatform
- After installation:
  - Splunk starts automatically on Windows
  - Splunk must be manually started on *NIX until boot-start is enabled

Splunk Component Installation Overview

- Installing Splunk Enterprise as an Indexer or Search Head is identical to installing a single deployment instance
- The difference happens at a configuration level
  - Installation as configuration is an iterative and ongoing event as you build and scale your deployment
  - Administrators need to be in control of the environment to fulfill emerging needs
  - Before installing Indexers or Search Heads, be sure to keep in mind the different hardware requirements

Common Splunk Commands

splunk is the program in the **bin** directory to run the CLI

Command	Operation
splunk help	Display a usage summary
splunk [start | stop | restart]	Manage the Splunk processes
splunk start –-accept-license	Automatically accept the license without prompt
splunk status	Display the Splunk process status
splunk show splunkd-port	Show the port that the **splunkd** listens on
splunk show web-port	Show the port that Splunk Web listens on
splunk show servername	Show the servername of this instance
splunk show default-hostname	Show the default host name used for all data inputs
splunk enable boot-start -user	Initialize script to run Splunk Enterprise at system startup

Module 4
Getting Data In

Splunk Index Time Process

- Splunk index time process (data ingestion) can be broken down into three phases:
  1. **Input phase**: handled at the source (usually a forwarder)
     - The data sources are being opened and read
     - Data is handled as streams and any configuration settings are applied to the entire stream
  2. **Parsing phase**: handled by indexers (or heavy forwarders)
     - Data is broken up into events and advanced processing can be performed
  3. **Indexing phase**:
     - License meter runs as data and is initially written to disk, prior to compression
     - After data is written to disk, it **cannot** be changed

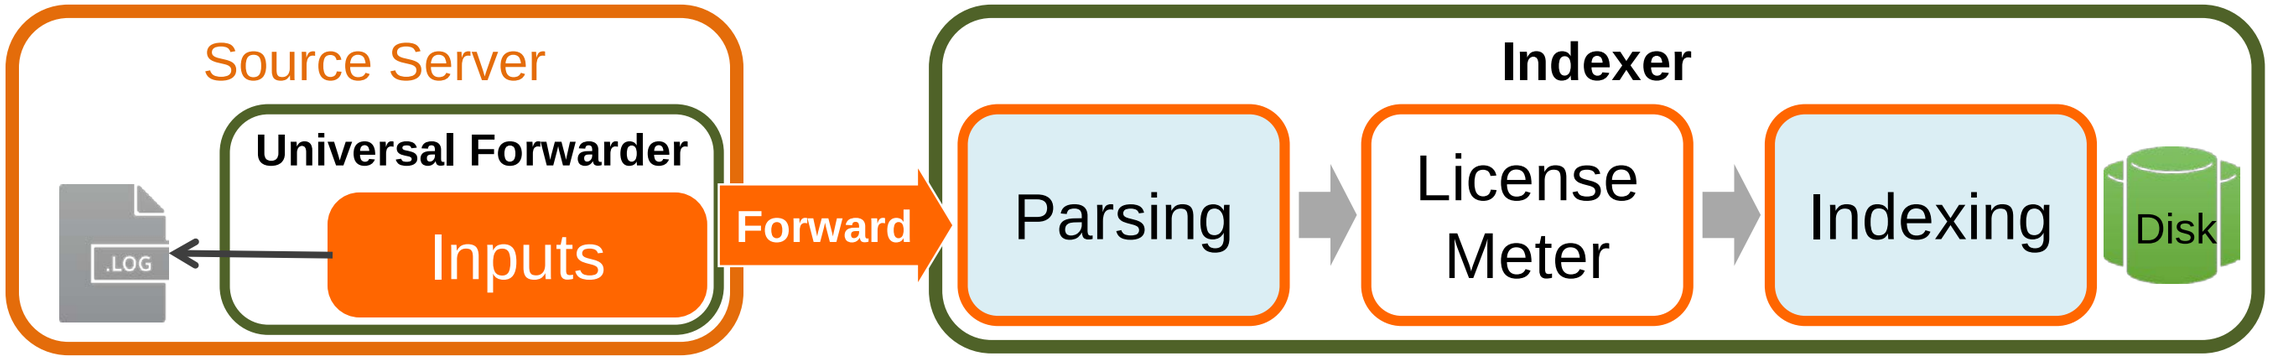

Data Input Types

- Splunk supports many types of data input
  - **Files and directories:** monitoring text files and/or directory structures containing text files
  - **Network data:** listening on a port for network data
  - **Script output:** executing a script and using the output from the script as the input
  - **Windows logs:** monitoring Windows event logs, Active Directory, etc.
  - **HTTP**: using the HTTP Event Collector
  - And more...
- You can add data inputs with:
  - Apps and add-ons from Splunkbase
  - Splunk Web
  - CLI
  - Directly editing **`inputs.conf`**

Default Metadata Settings

- When you index a data source, Splunk assigns metadata values
  - The metadata is applied to the entire source
  - Splunk applies defaults if not specified
  - You can also override them at input time or later

Metadata	Default
`source`	Path of input file, network hostname:port, or script name
`host`	Splunk hostname of the inputting instance (usually a forwarder)
`sourcetype`	Uses the source filename if Splunk cannot automatically determine
`index`	Defaults to **`main`**

Adding an Input with Splunk Web

- Splunk admins have a number of ways to start the **Add Data** page
  - Click the **Add Data** icon
    - On the admin's **Home** page
    - On the **Settings** panel
  - Select **Settings** > **Data inputs** > **Add new**

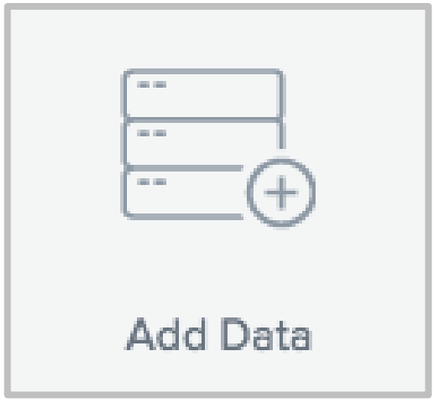

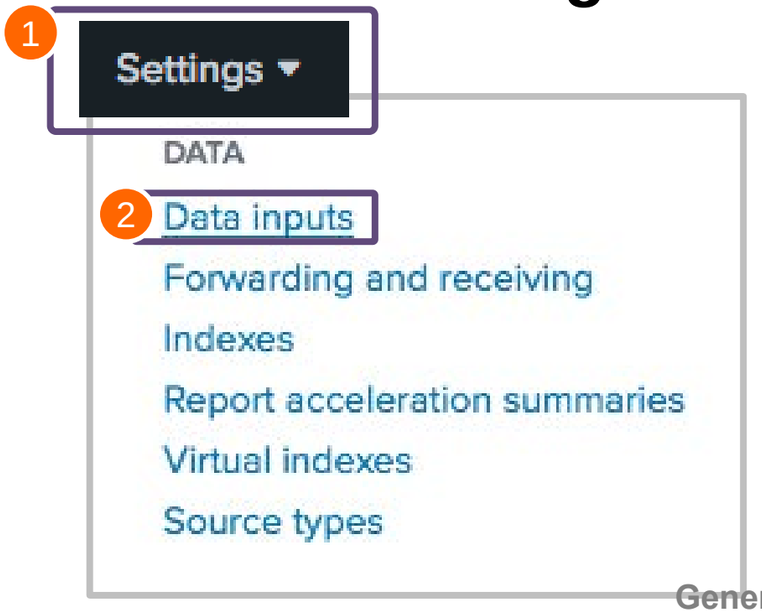

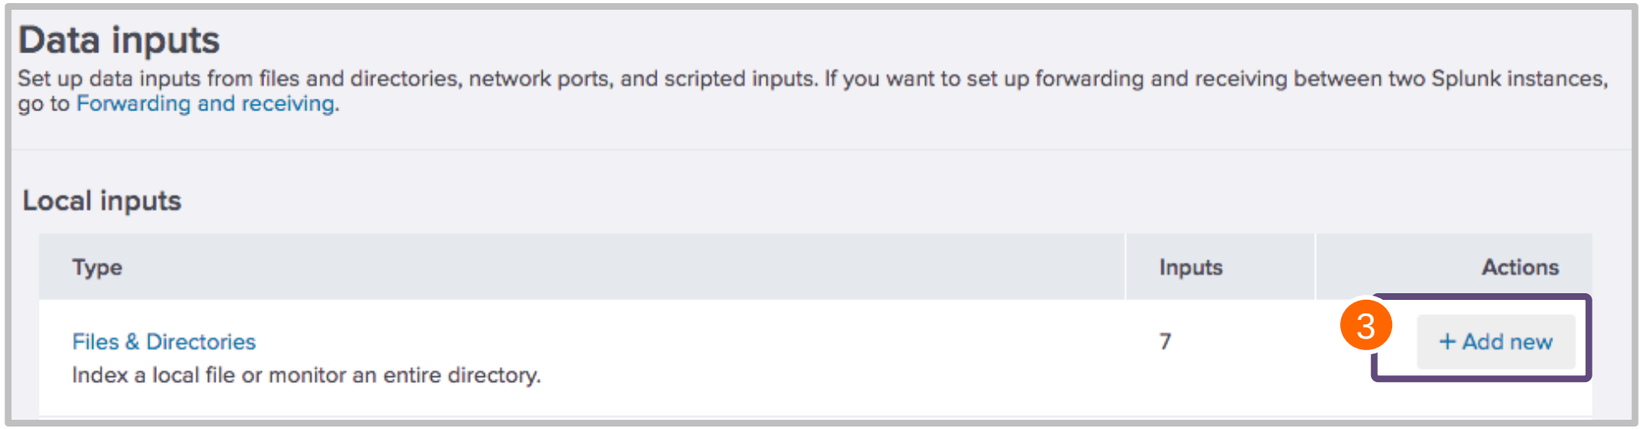

Add Data Menu

Add Data menu provides three options depending on the source to be used

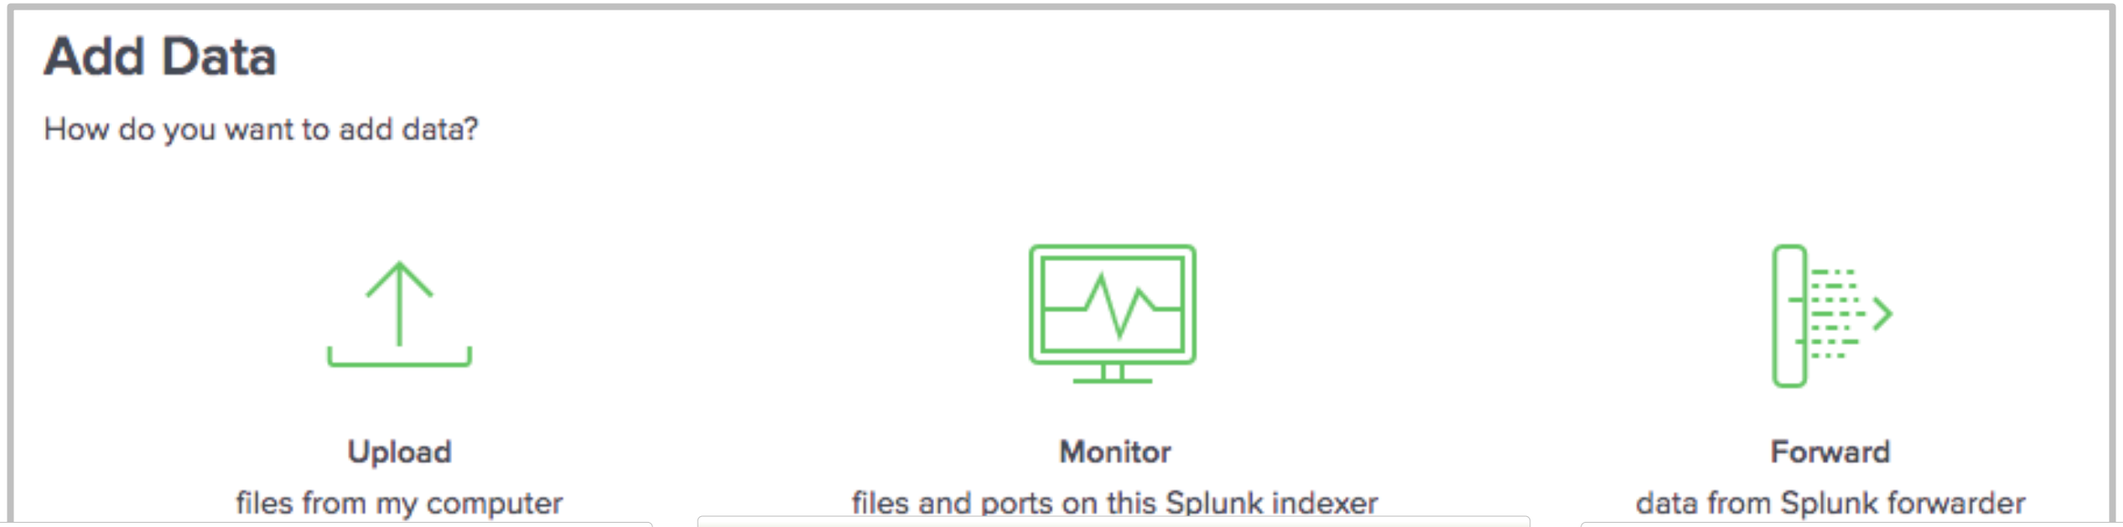

Upload Option

Upload allows uploading local files that only get indexed once. Useful for testing or data that is created once and never gets updated. Does not create **`inputs.conf`**.

Monitor Option

Provides one-time or continuous monitoring of files, directories, http events, network ports, or data gathering scripts located on Splunk Enterprise instances. Useful for testing inputs.

Forward Option

Main source of input in production environments. Remote machines gather and forward data to indexers over a receiving port.

Select Source

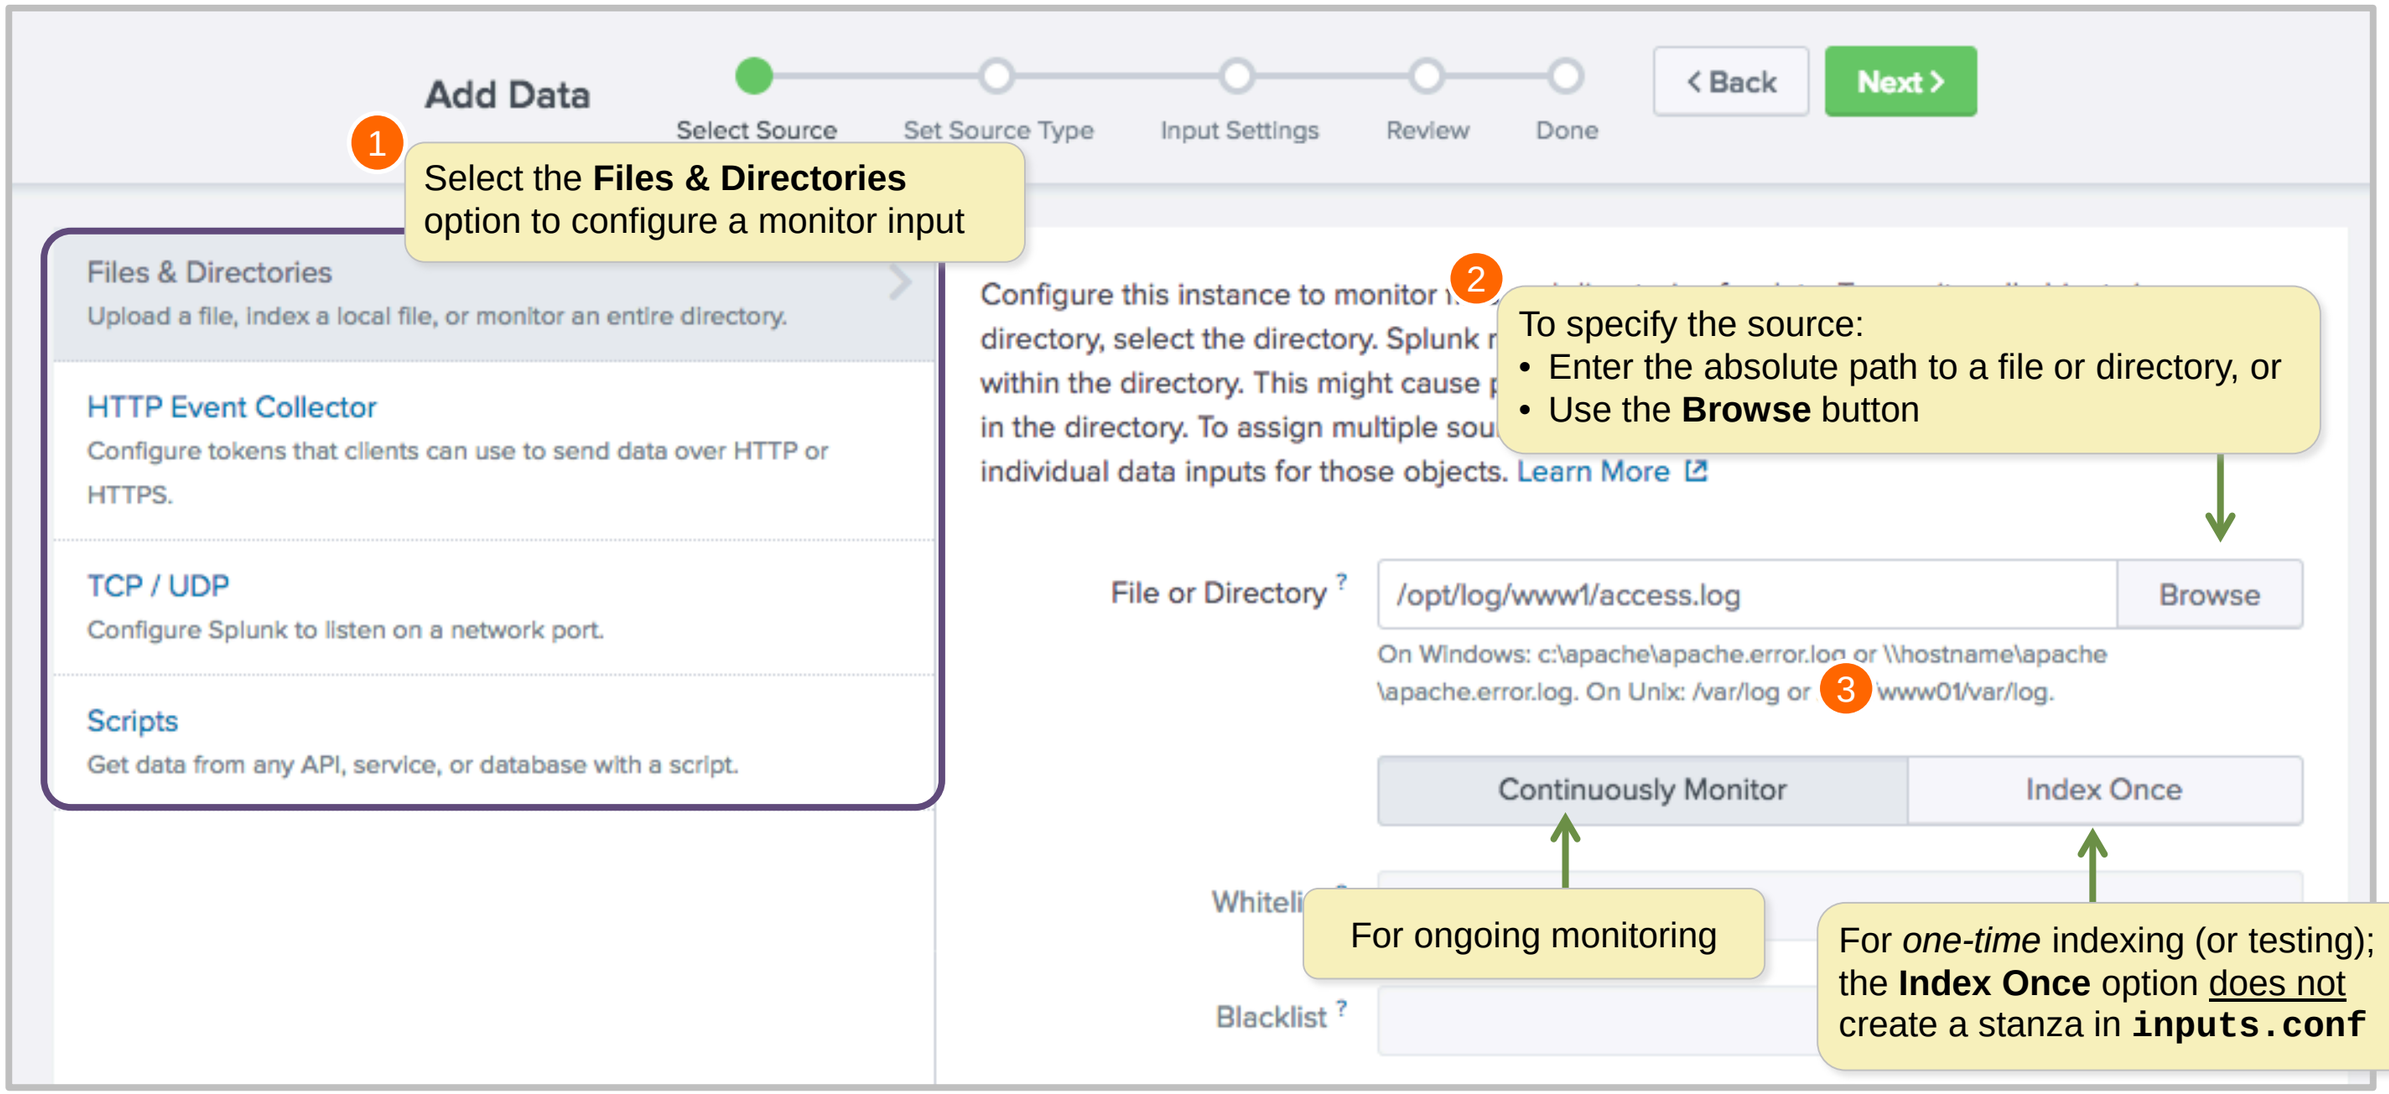

Set Source Type (Data Preview Interface)

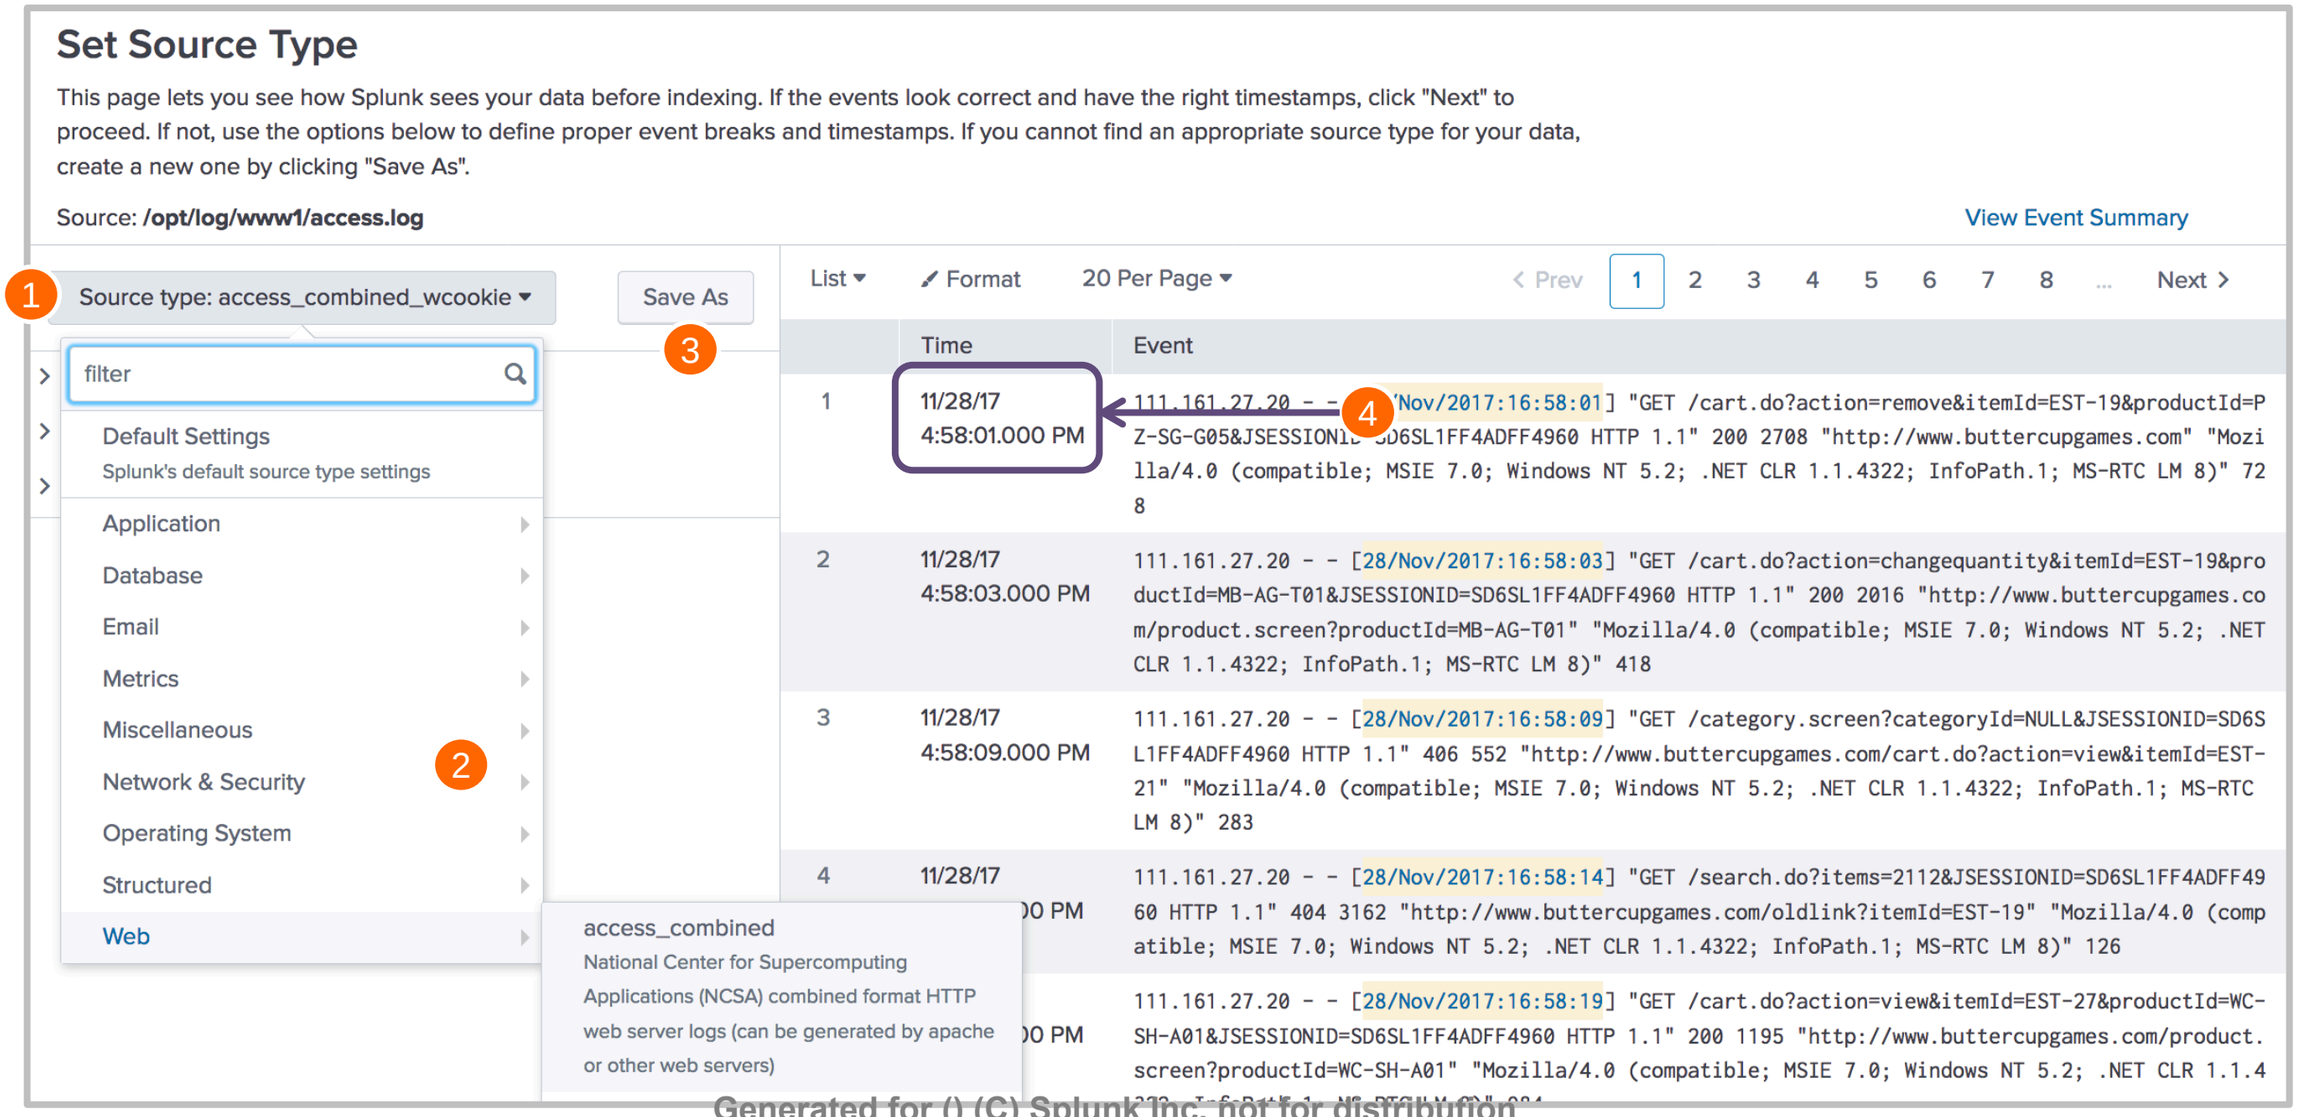

Set Source Type (cont.)

1. Splunk automatically determines the source type for major data types when there is enough data
2. You can choose a different source type from the dropdown list
3. Or, you can create a new source type name for the specific source
4. **Data preview** displays how your processed events will be indexed
   - If the events are correctly separated and the right timestamps are highlighted, you can move ahead
     - If not, you can select a different source type from the list or customize the settings

Pretrained Source Types

- Splunk has default settings for many types of data
- The docs also contain a list of source types that Splunk automatically recognizes
- Splunk apps can be used to define additional source types

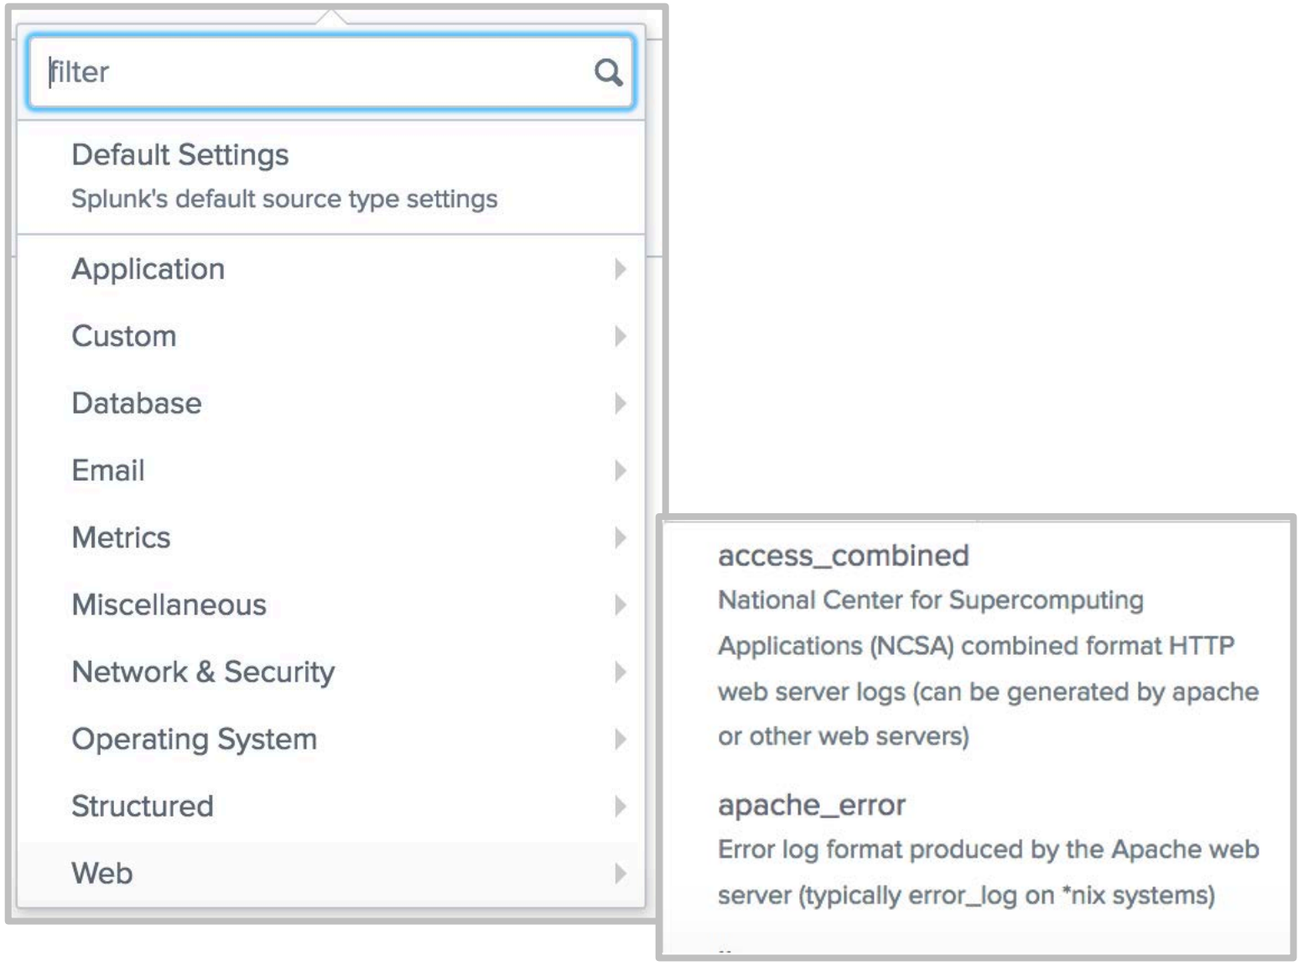

http://docs.splunk.com/Documentation/Splunk/latest/Data/Listofpretrainedsourcetypes

Input Settings

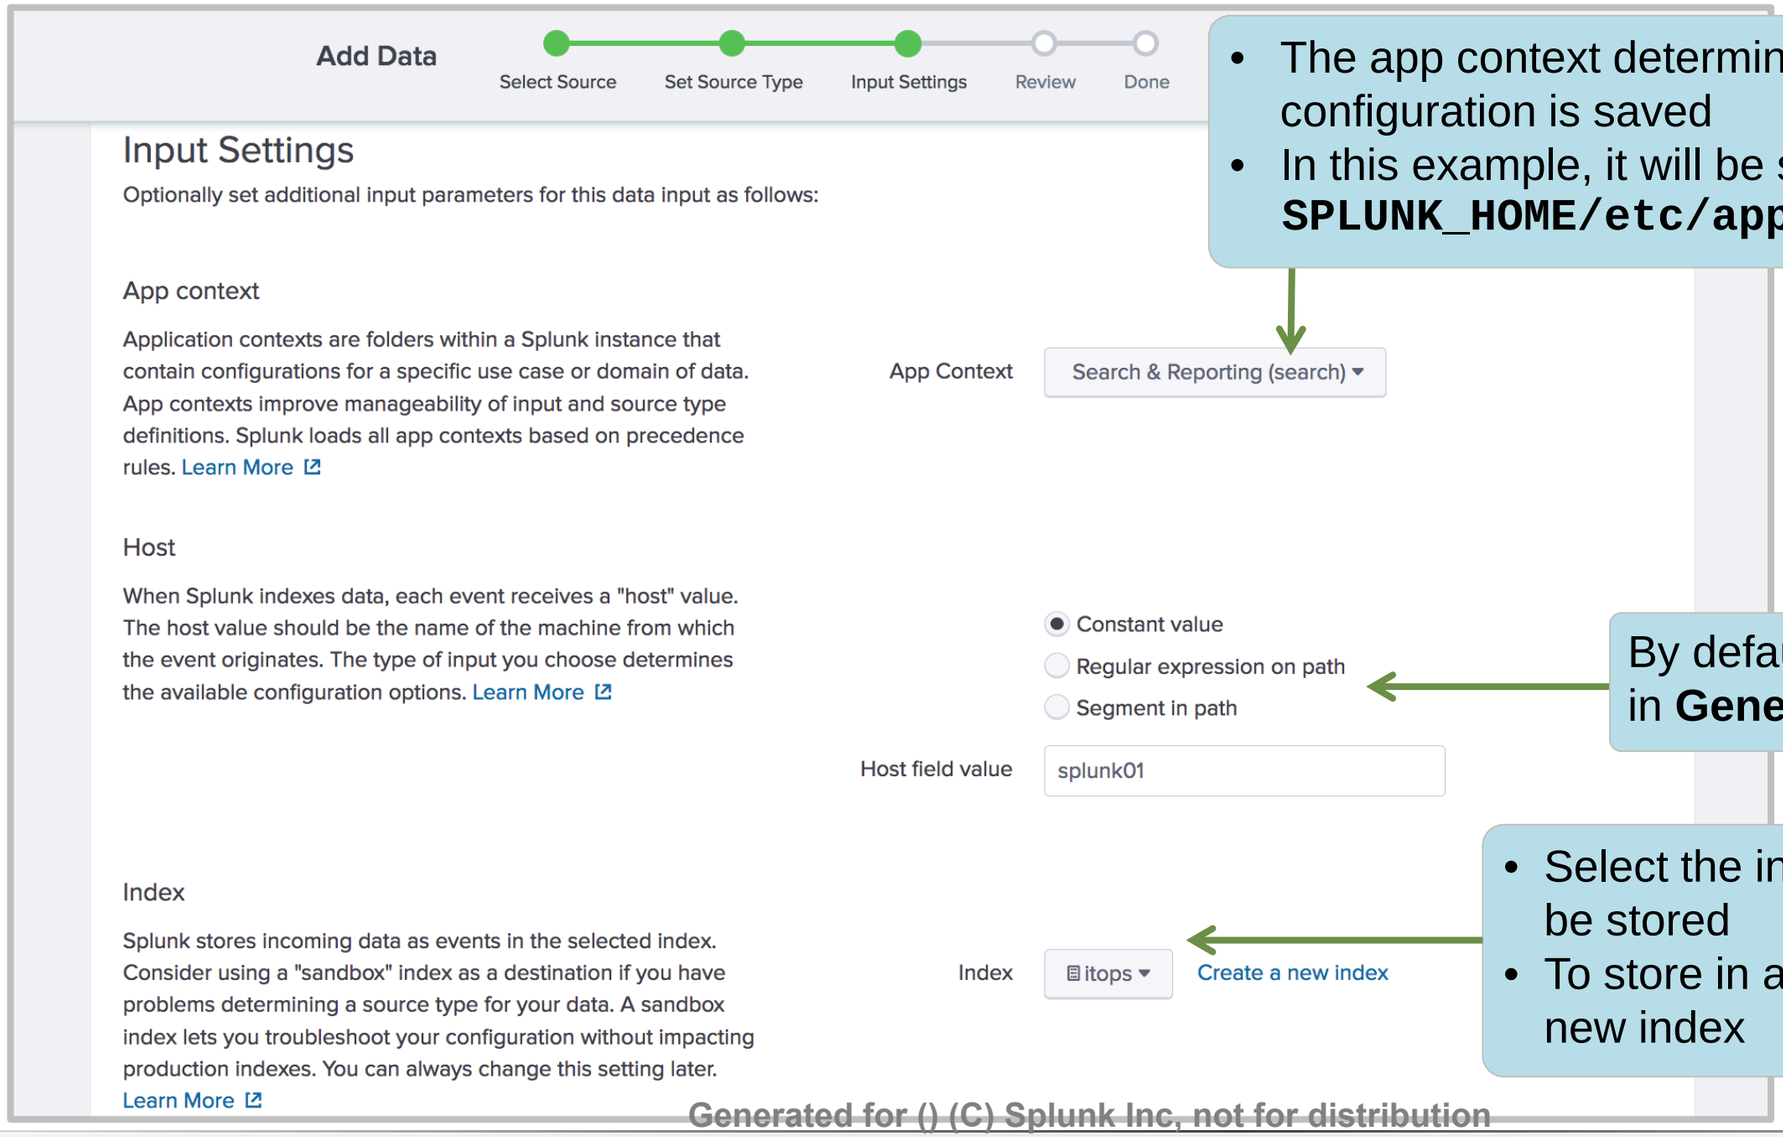

- The app context determines where your input configuration is saved
- In this example, it will be saved in: **`SPLUNK_HOME/etc/apps/search/local`**

By default, the default host name in **General settings** is used

- Select the index where this input should be stored
- To store in a new index, first create the new index

Review

- Review the input configuration summary and click **Submit** to finalize

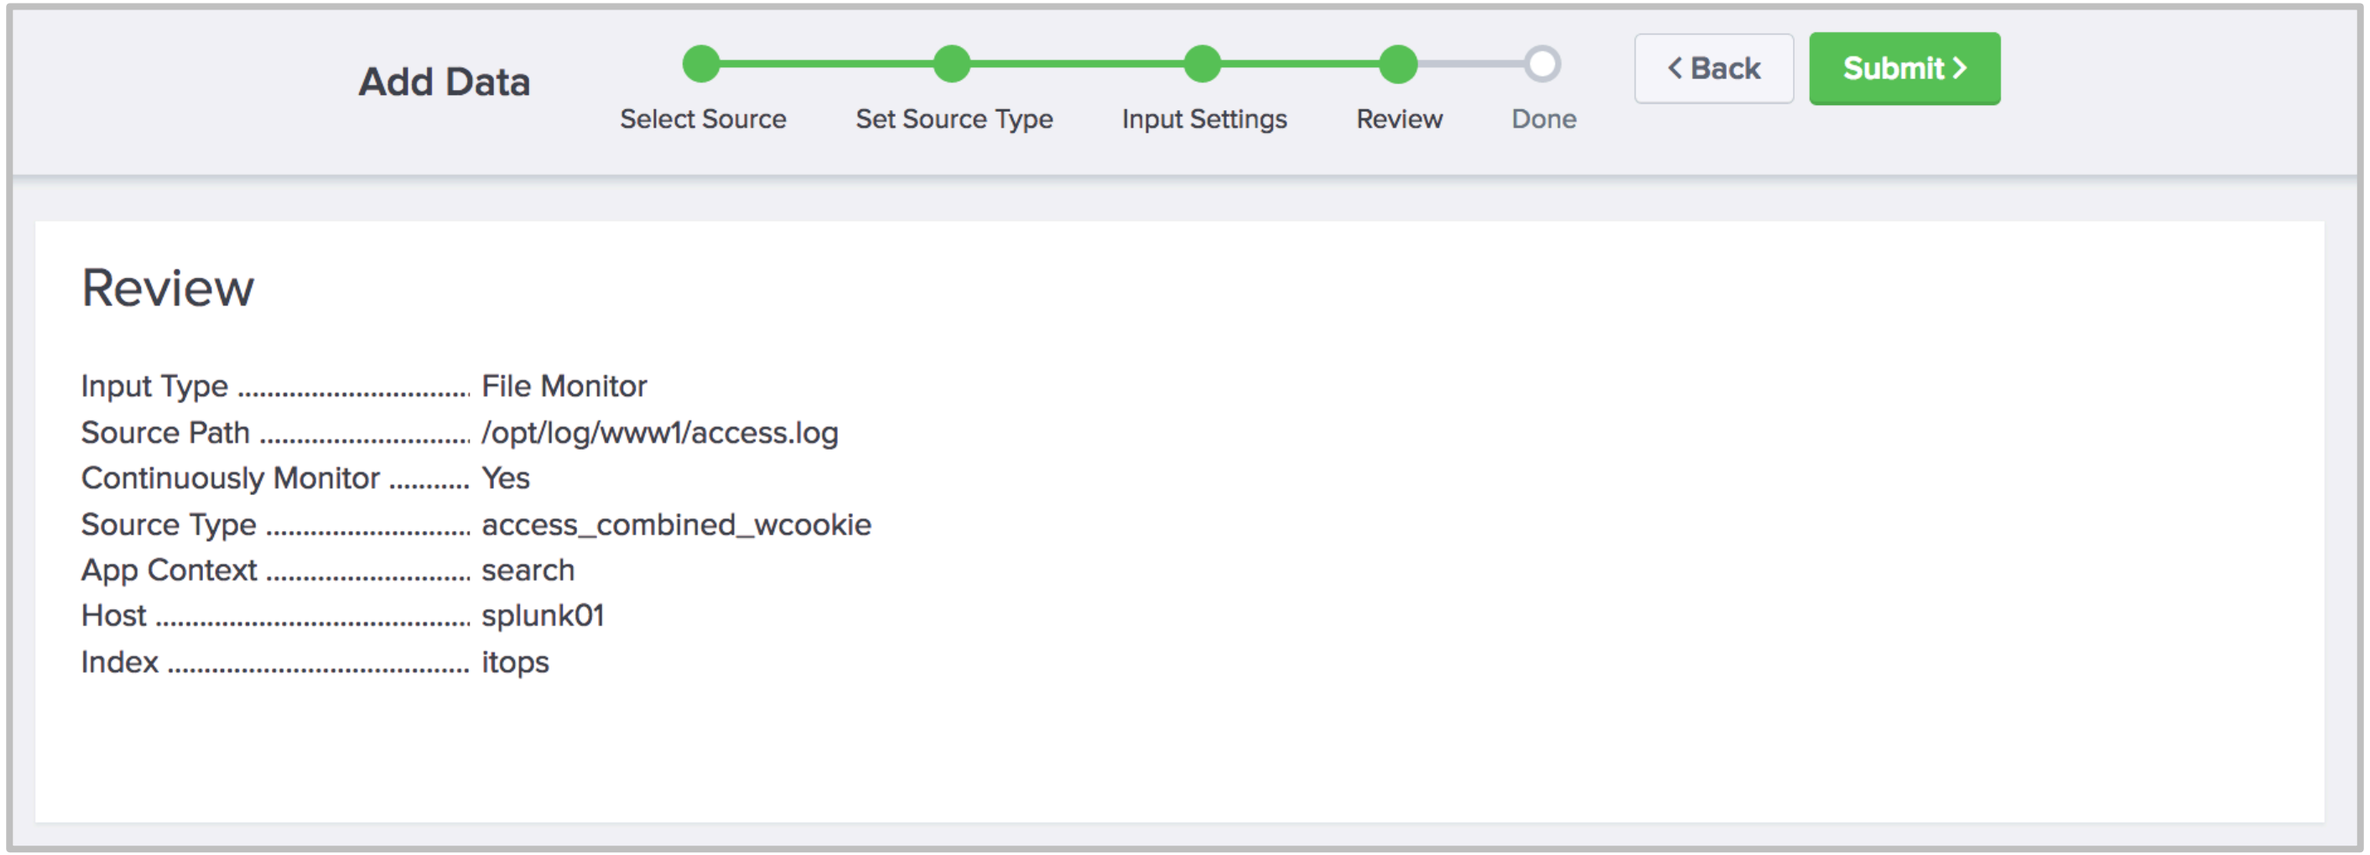

What Happens Next?

- Indexed events are available for immediate search
  - However, it may take a minute for Splunk to *start* indexing the data
- You are given other options to do more with your data

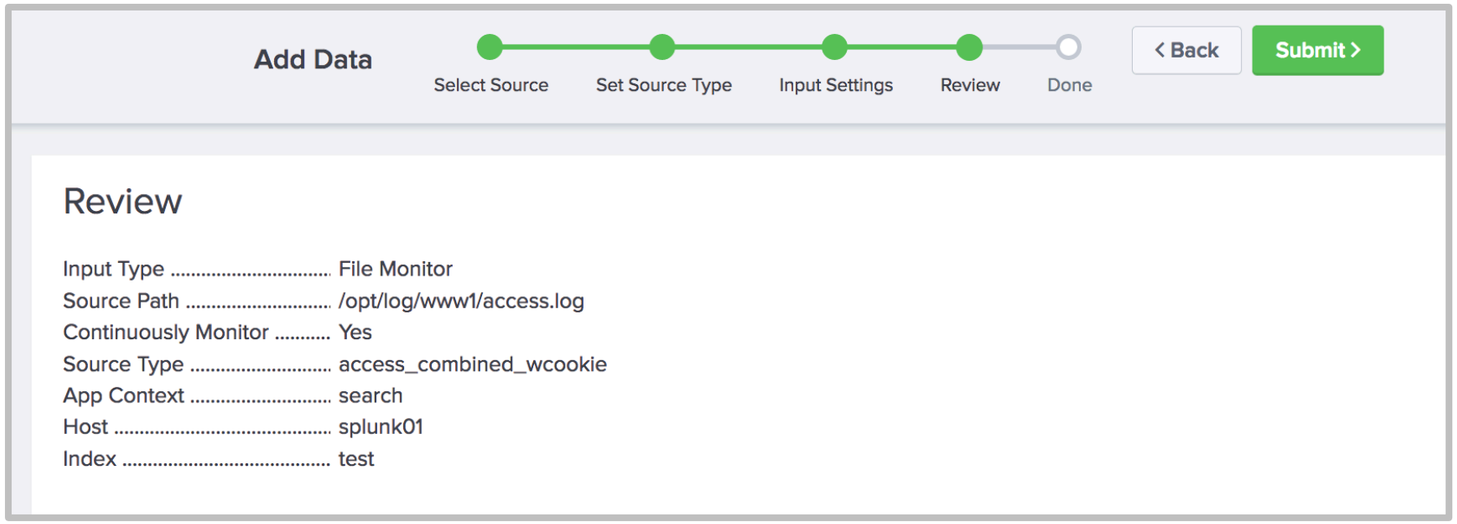

Module 5:
Basic Search

Search Assistant

- Search Assistant provides selections for how to complete the search string
- Before the first pipe (|), it looks for matching terms
- You can continue typing OR select a term from the list
  - If you select a term from the list, it is added to the search

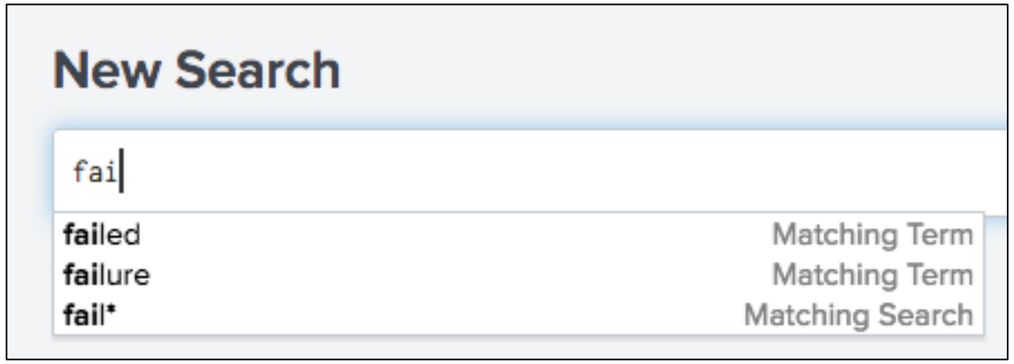

Search Assistant (cont.)

- After the first pipe, the Search Assistant shows a list of commands that can be entered into the search string
- You can continue typing OR scroll through and select a command to add
- If you mouse over a command, more information about the command is shown
- As you continue to type, Search Assistant makes more suggestions

A

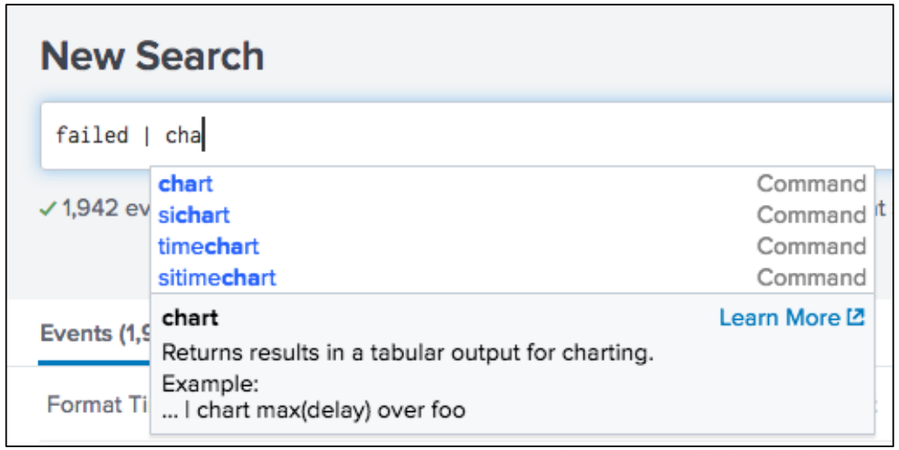

B

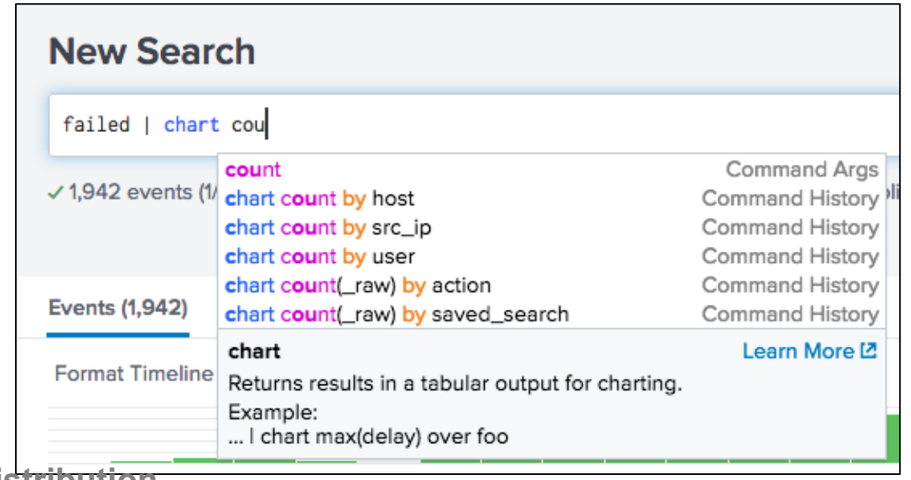

Search Assistant (cont.)

- Search Assistant is enabled by default in the **SPL Editor** user preferences
- By default, **Compact** is selected
- To show more information, choose **Full**

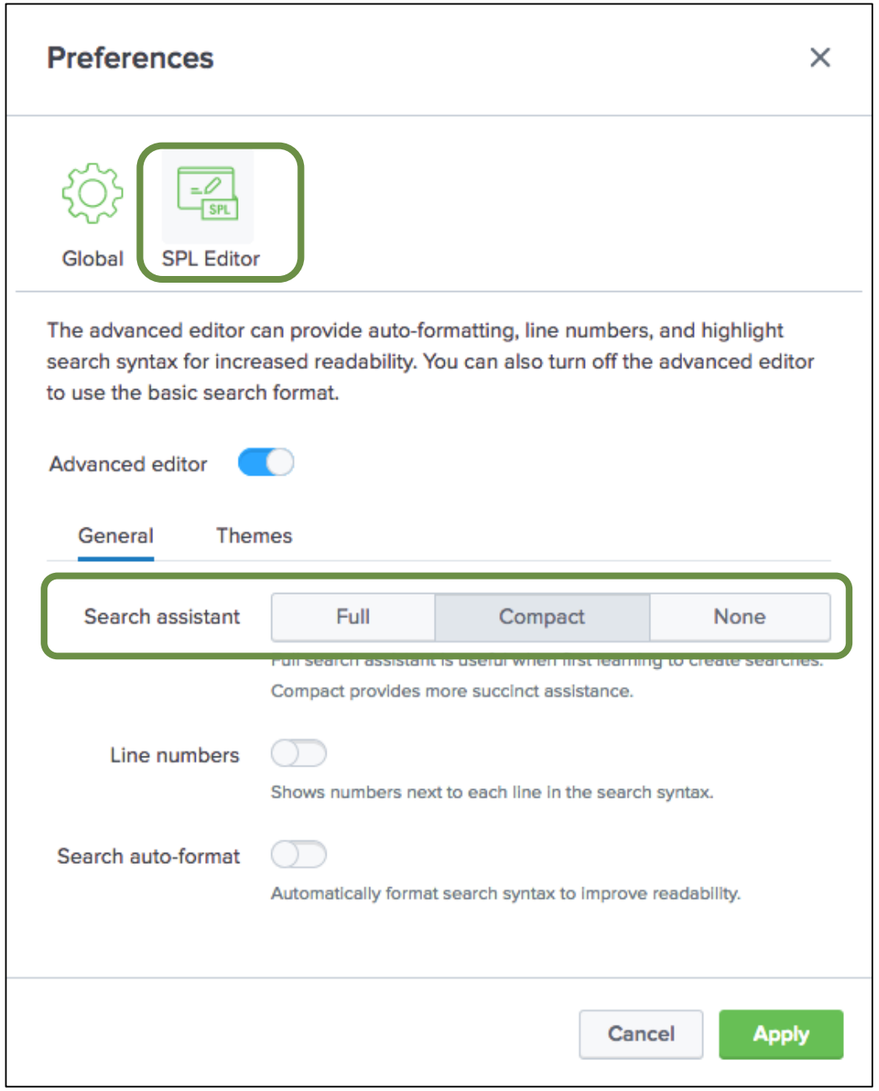

Compact Mode

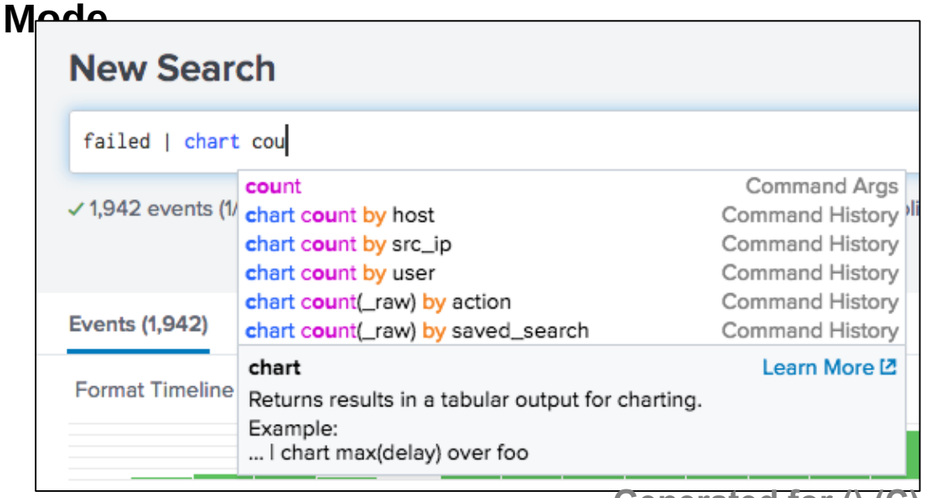

Search Assistant – Full Mode

- A To show more information, click **More »**
- B To show less information, click **« Less**
- C To toggle Full mode off, de-select **Auto Open**

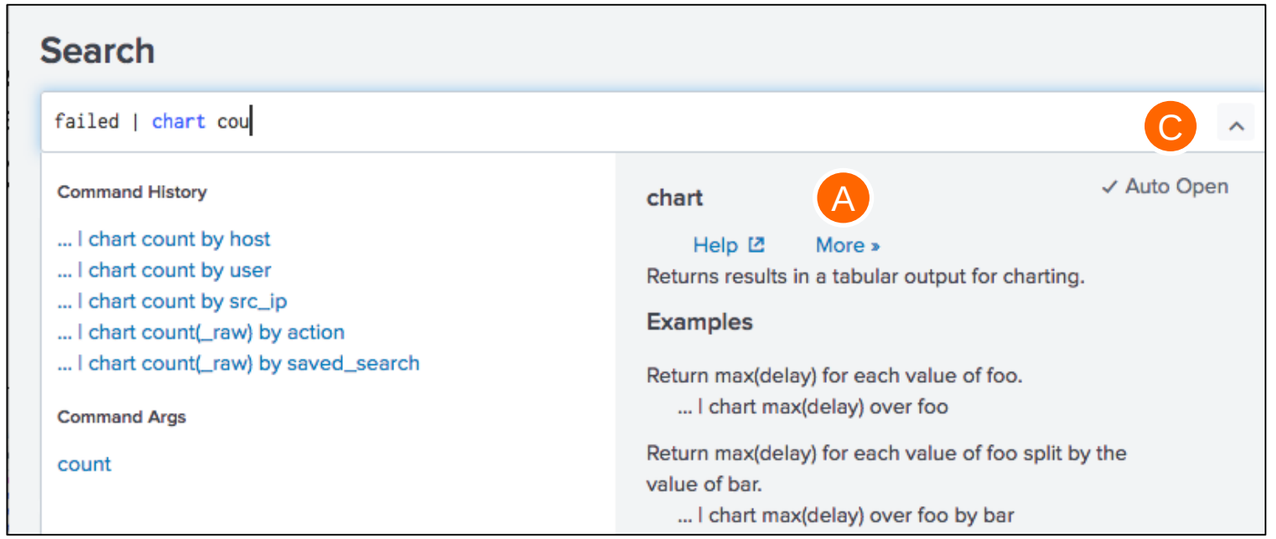

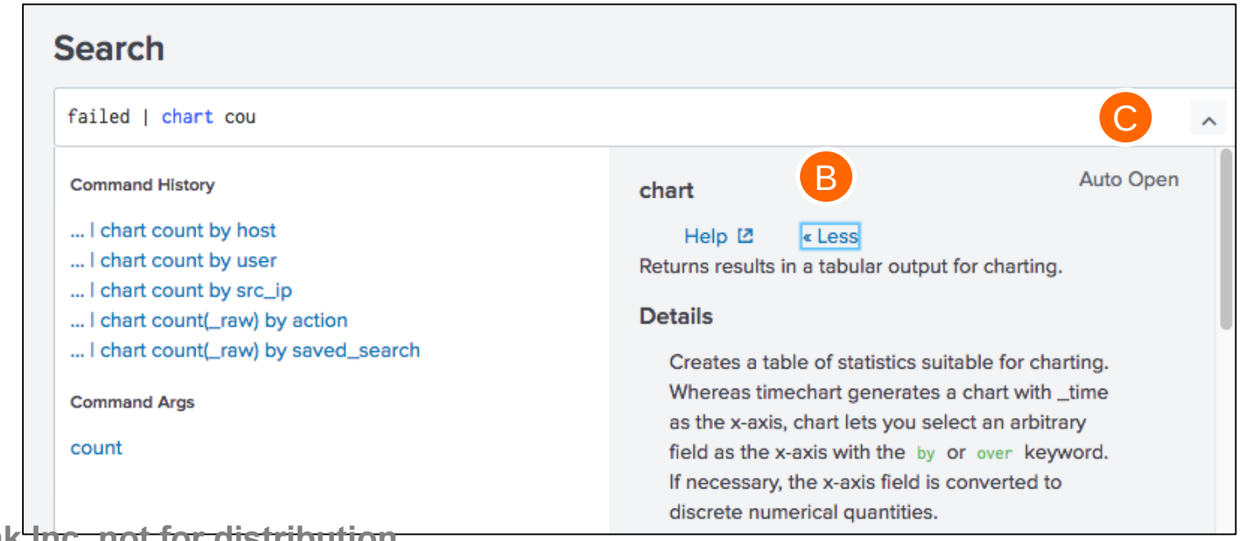

Search Assistant – Parentheses

- The Search Assistant provides help to match parentheses as you type
- When an end parenthesis is typed, the corresponding beginning parenthesis is automatically highlighted
  - If a beginning parenthesis cannot be found, *nothing* is highlighted

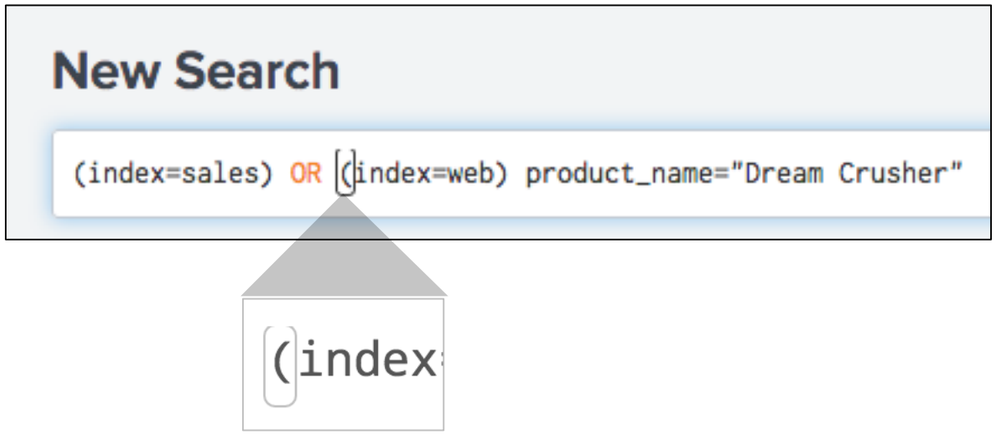

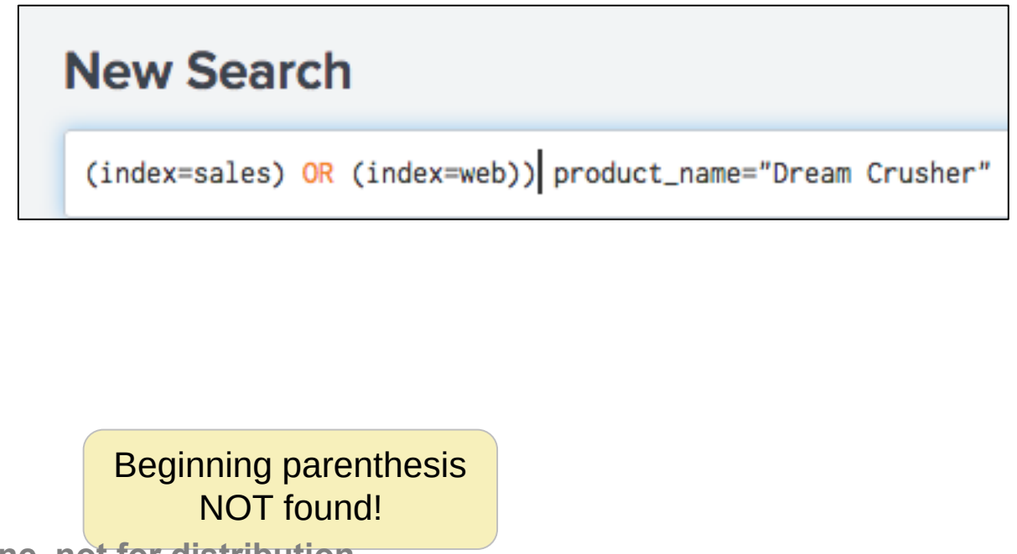

Viewing Search Results

- Matching results are returned immediately
- Displayed in reverse chronological order (newest first)
- Matching search terms are highlighted

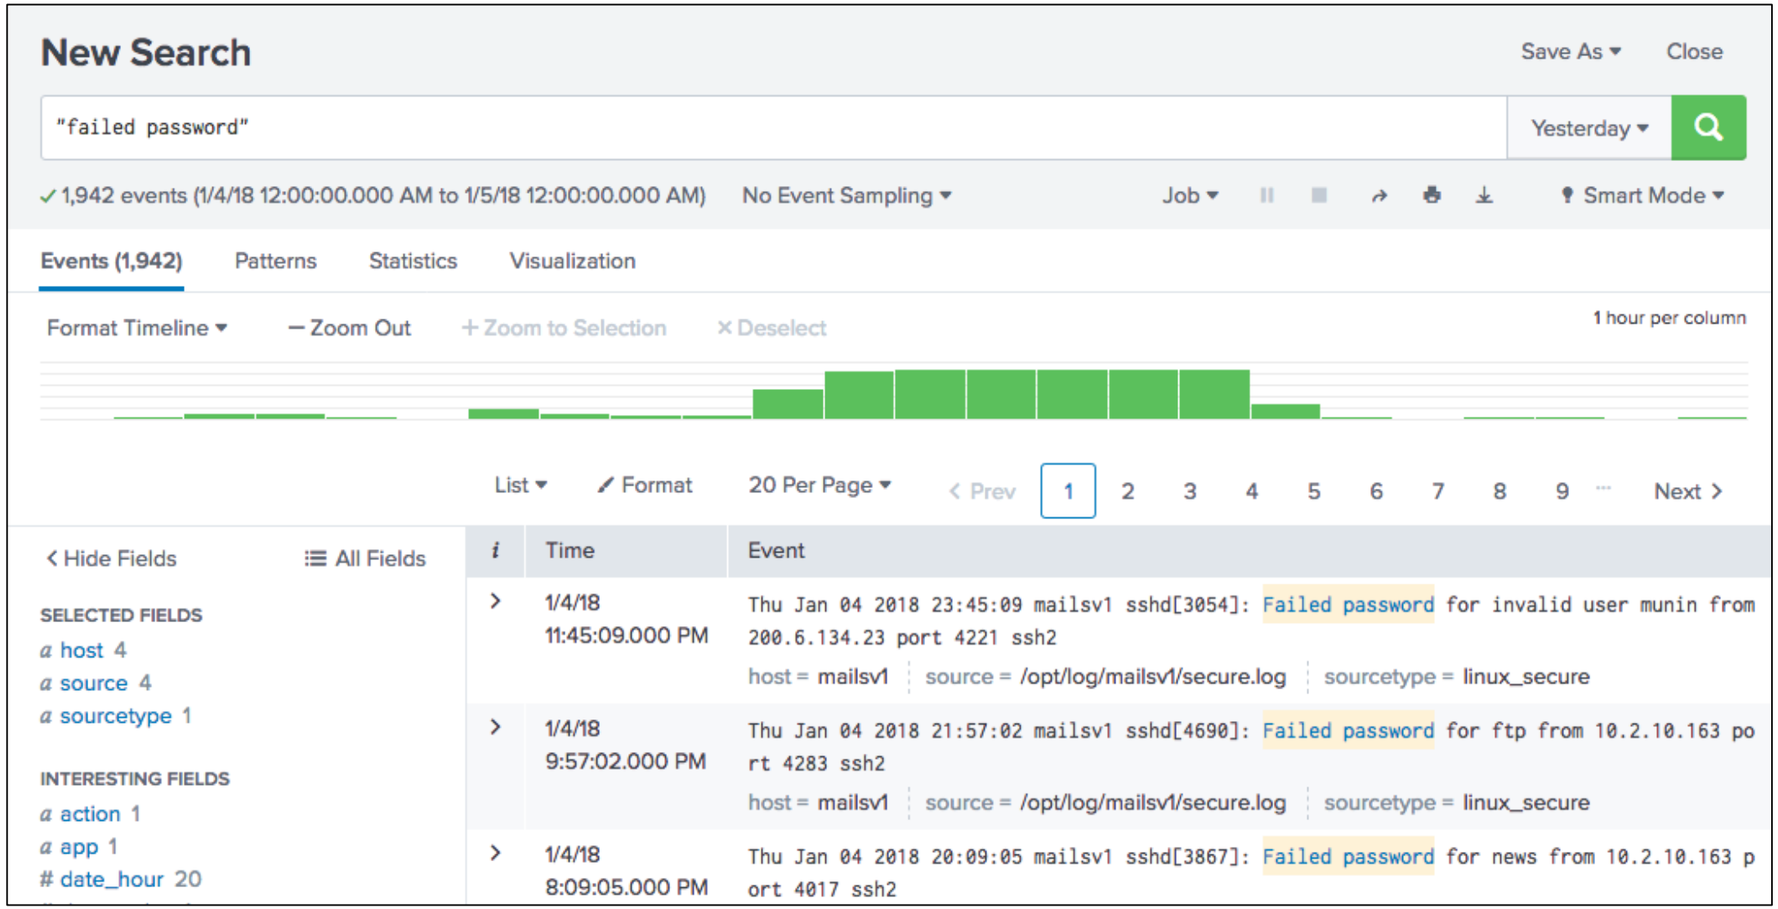

Viewing Search Results (cont.)

- Splunk parses data into individual events, extracts time, and assigns metadata
- Each event has:
  - timestamp
  - host
  - source
  - sourcetype
  - index

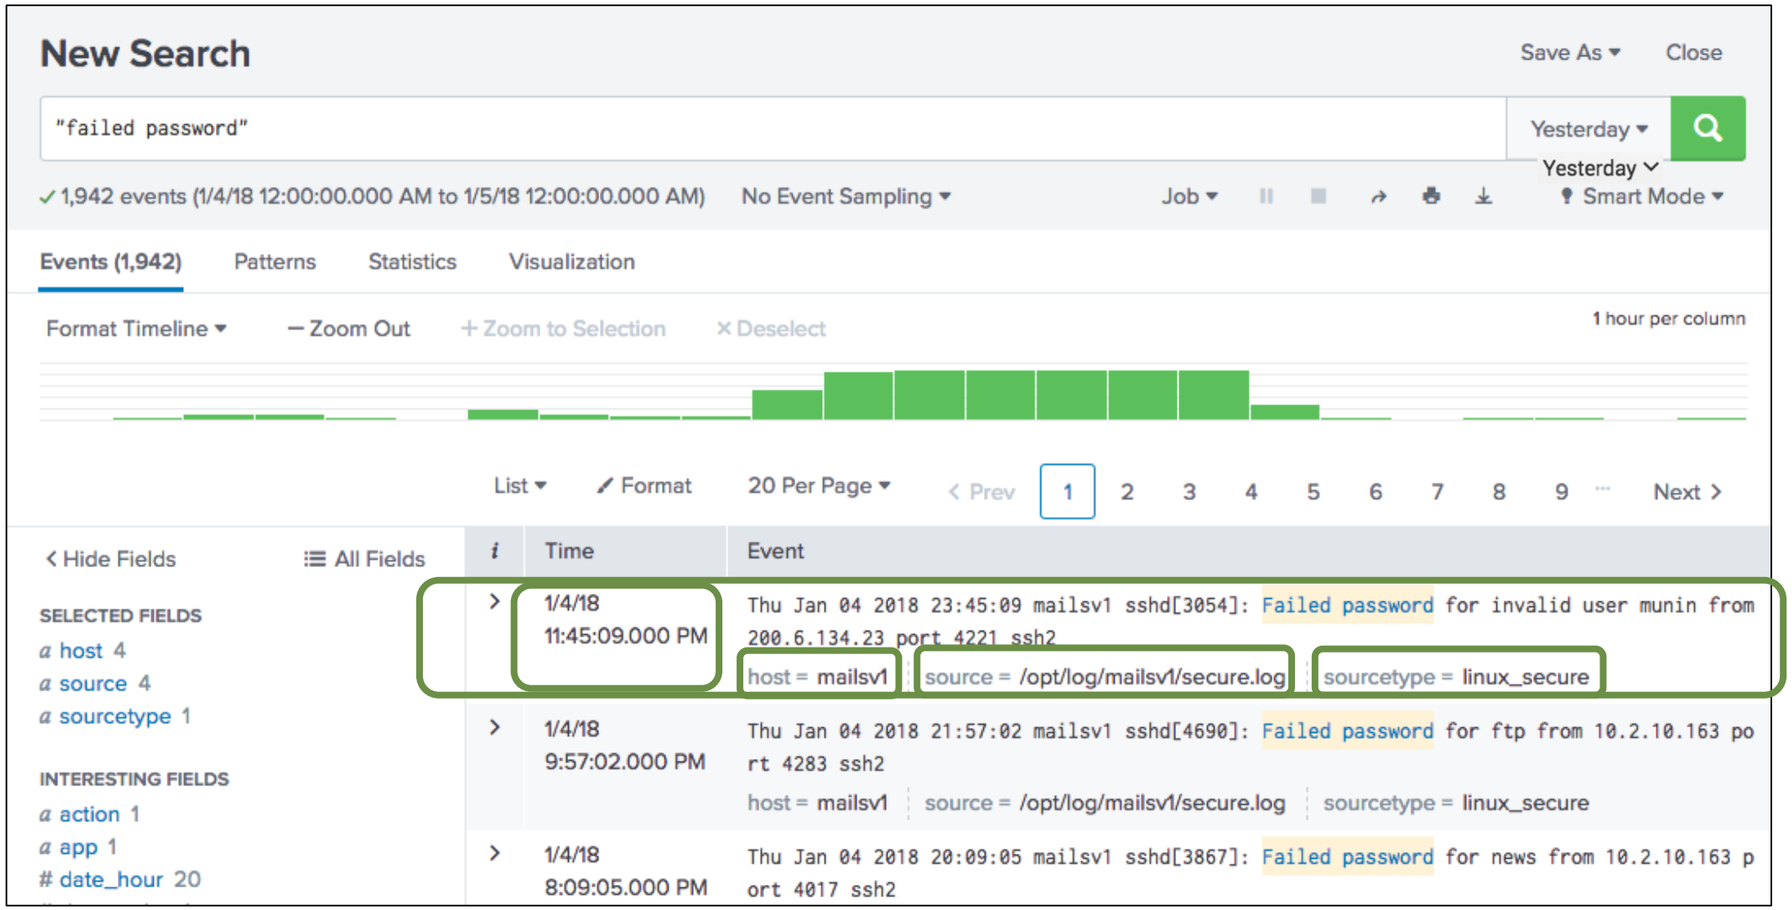

Viewing Search Results (cont.)

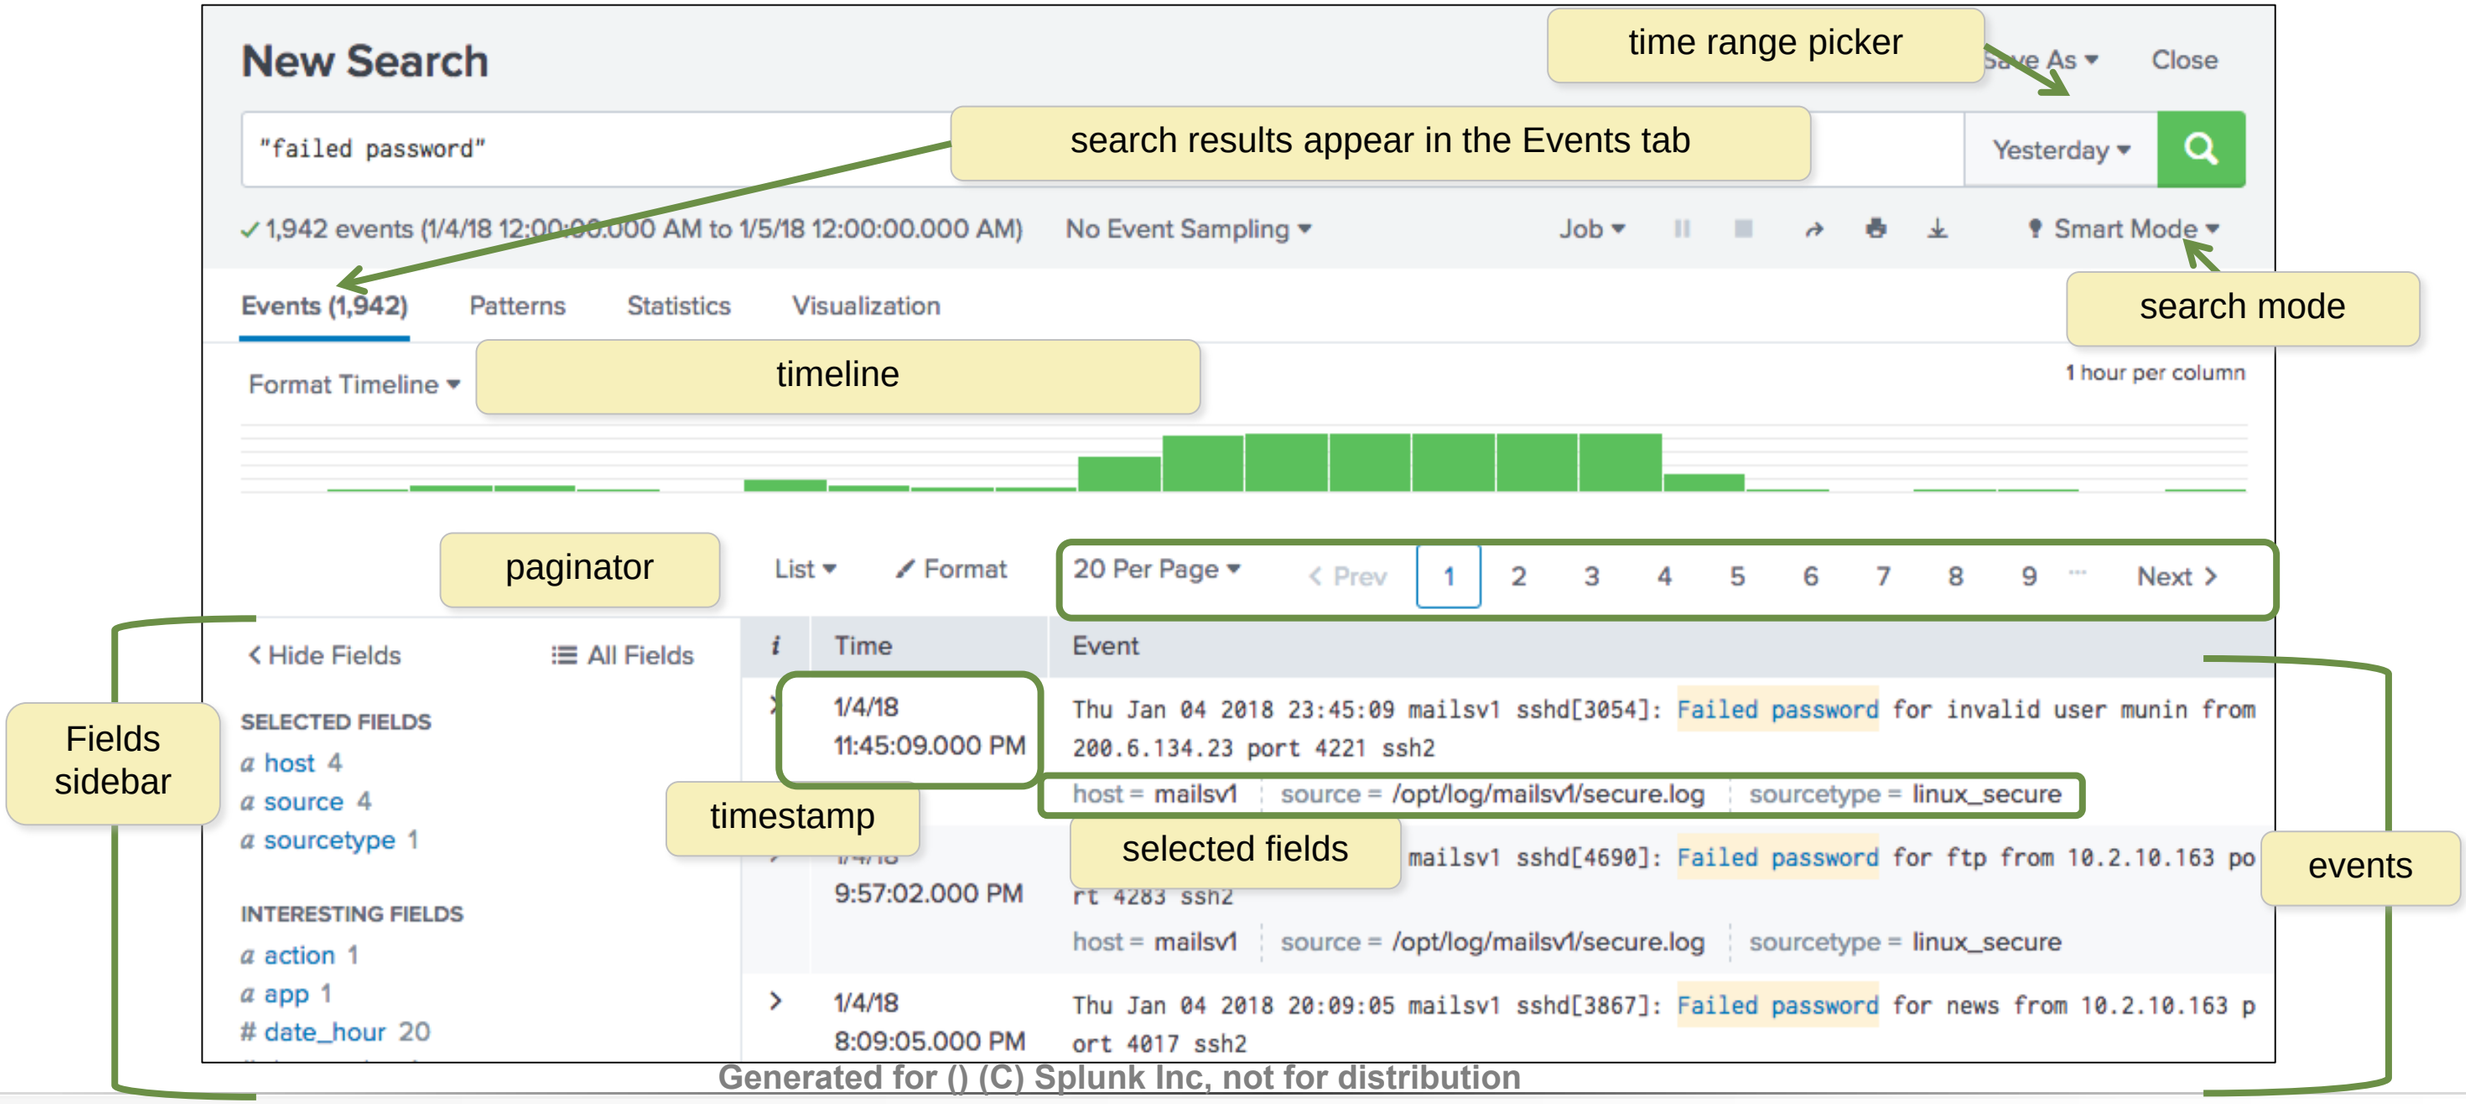

Using Search Results to Modify a Search

- When you mouse over search results, keywords are highlighted
- Click any item in your search results; a window appears allowing you to:
  - Add the item to the search
  - Exclude the item from the search
  - Open a new search including only that item

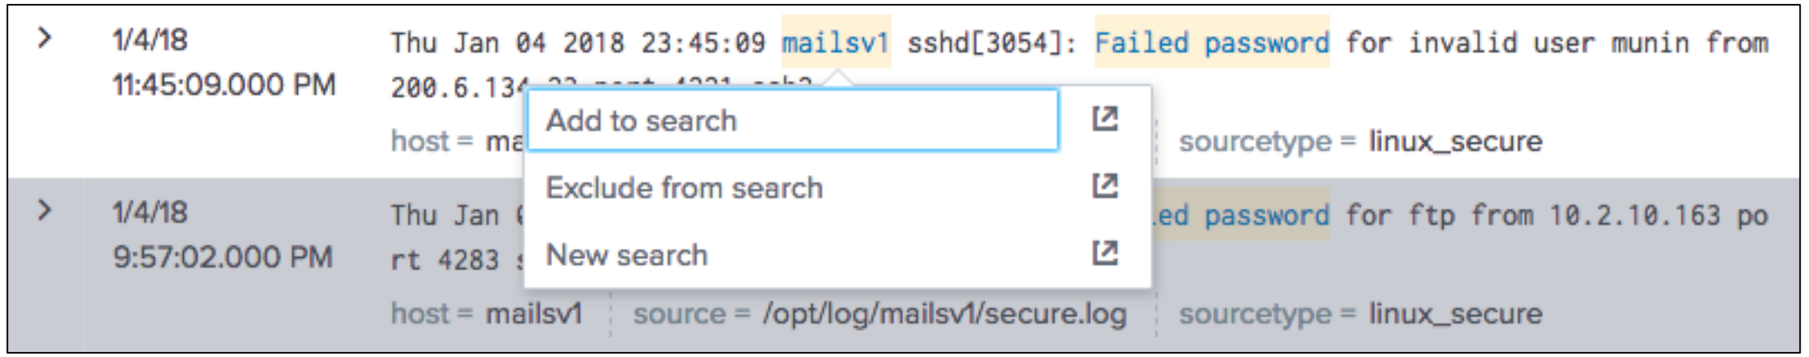

Changing Search Results View Options

You have several layout options for displaying your search results

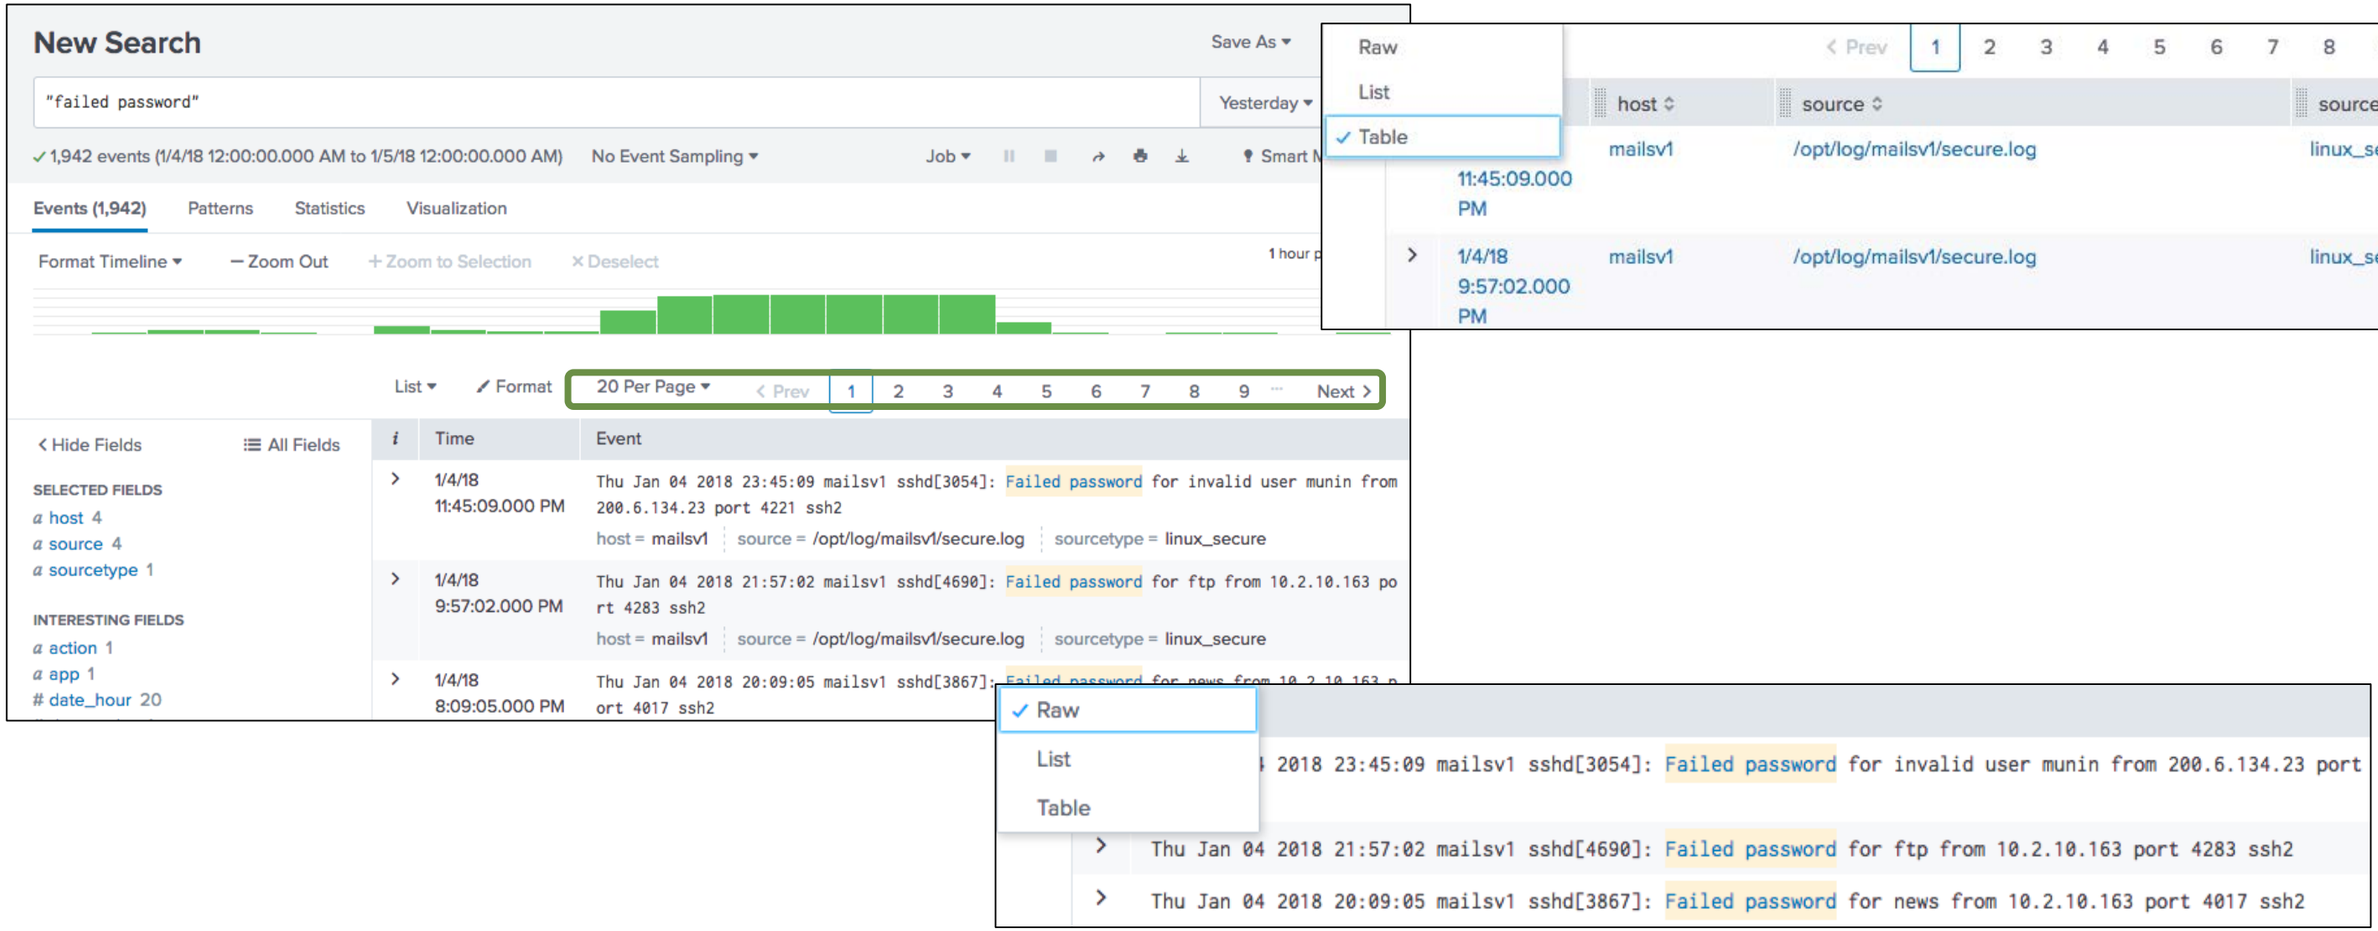

Selecting a Specific Time

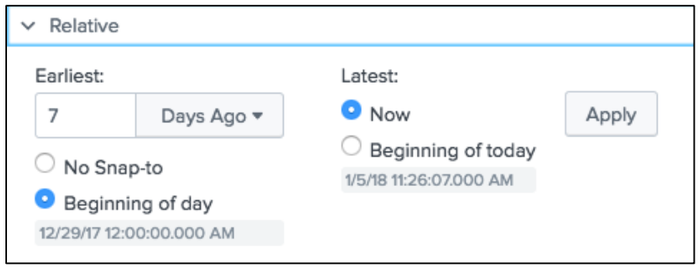

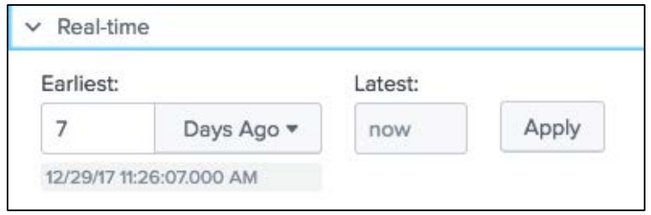

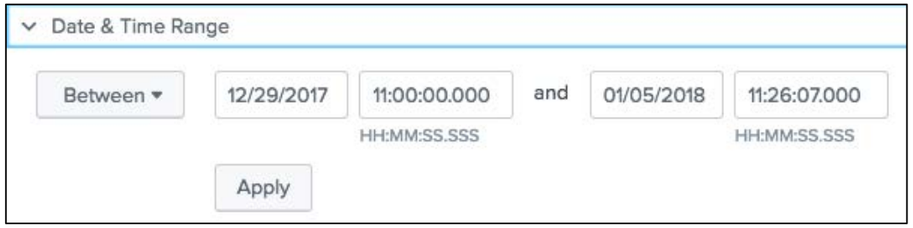

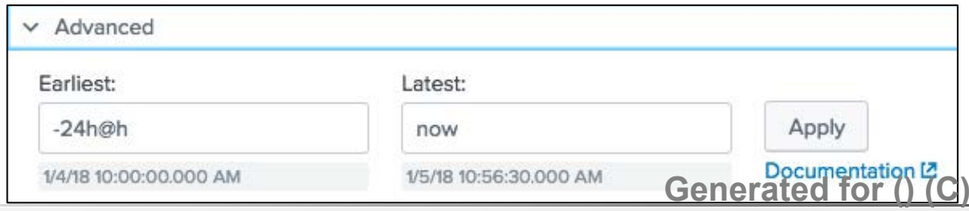

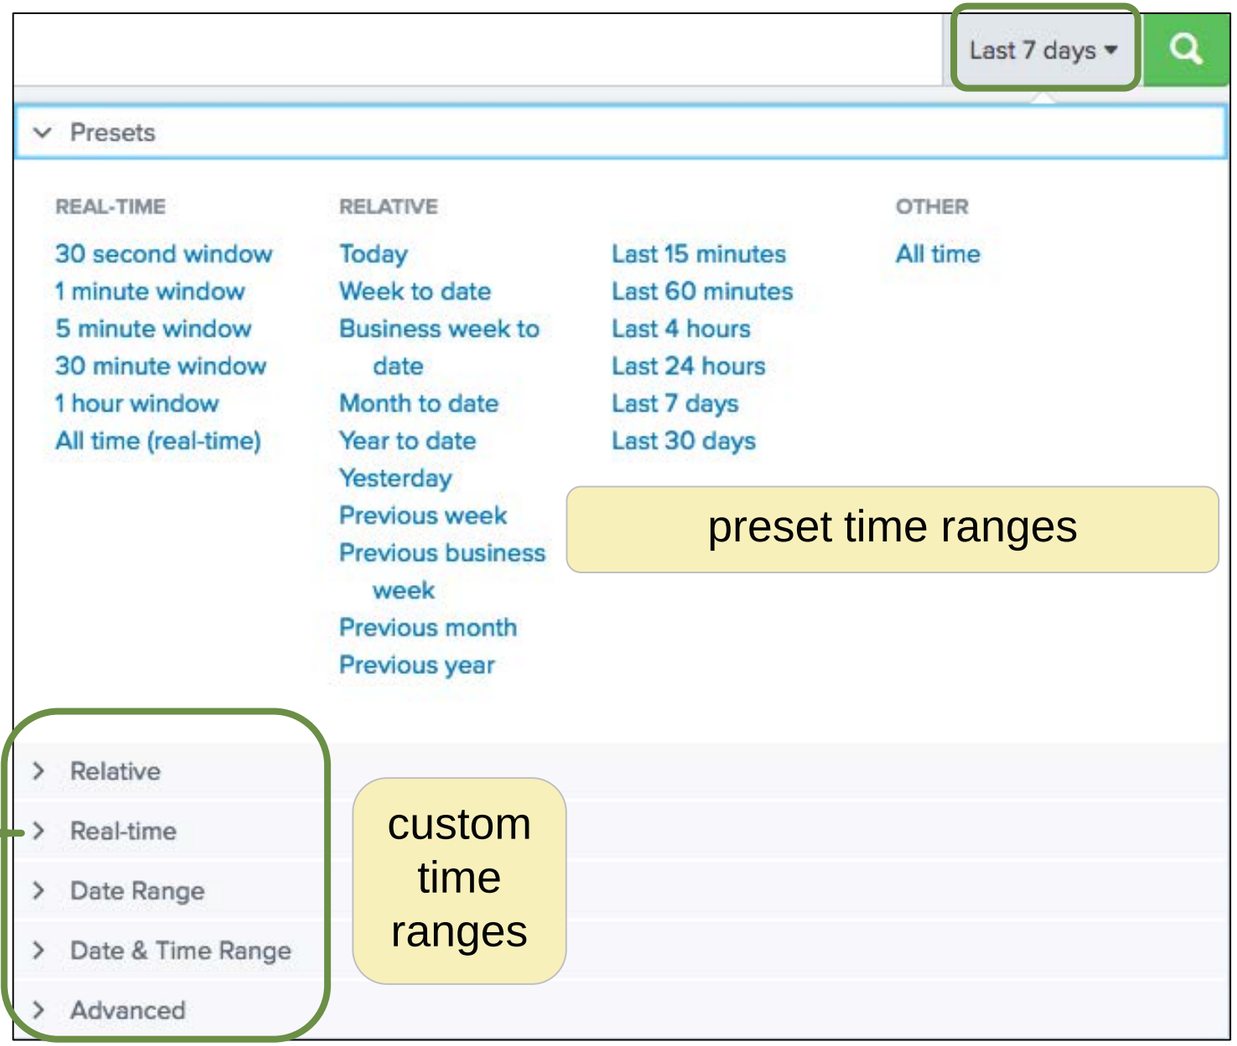

Time Range Abbreviations

- Time ranges are specified in the **Advanced** tab of the time range picker
- Time unit abbreviations include:

```
s = seconds   m = minutes   h = hours   d = days   w = week   mon = months   y = year
```

- **@** symbol "snaps" to the time unit you specify
  - Snapping rounds *down* to the nearest specified unit
  - Example: Current time when the search starts is 09:37:12

    -30m@h looks back to 09:00:00

Time Range: `earliest` and `latest`

- You can also specify a time range in the search bar
- To specify a beginning and an ending for a time range, use `earliest` and `latest`
- Examples:

`earliest=-h`	looks back one hour
`earliest=-2d@d latest=@d`	looks back from two days ago, up to the beginning of today
`earliest=6/15/2017:12:30:00`	looks back to specified time

Note

If time specified, it must be in MM/DD/YYYY:HH:MM:SS format.

Viewing the Timeline

- Timeline shows distribution of events specified in the time range
  - Mouse over for details, or single-click to filter results for that time period

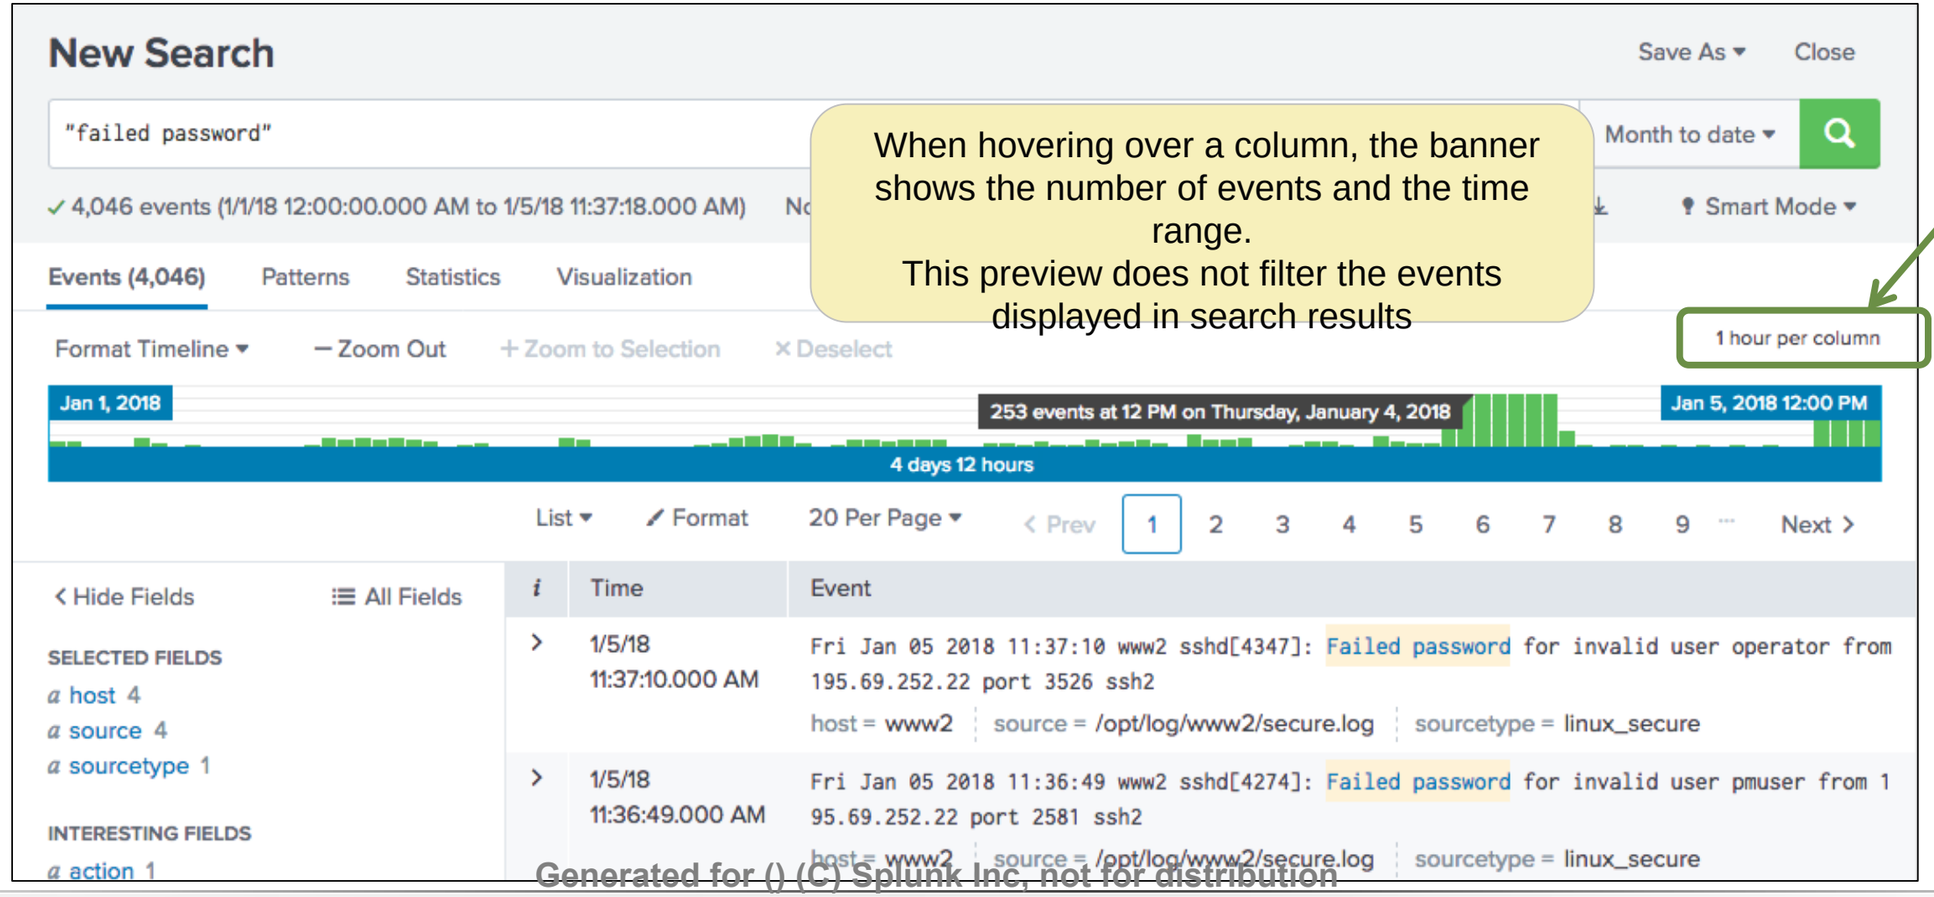

Viewing a Subset of the Results with Timeline

- To select a narrower time range, click and drag across a series of bars
  - This action filters the current search results
    - Does not re-execute the search
  - This filters the events and displays them in reverse chronological order (most recent first)

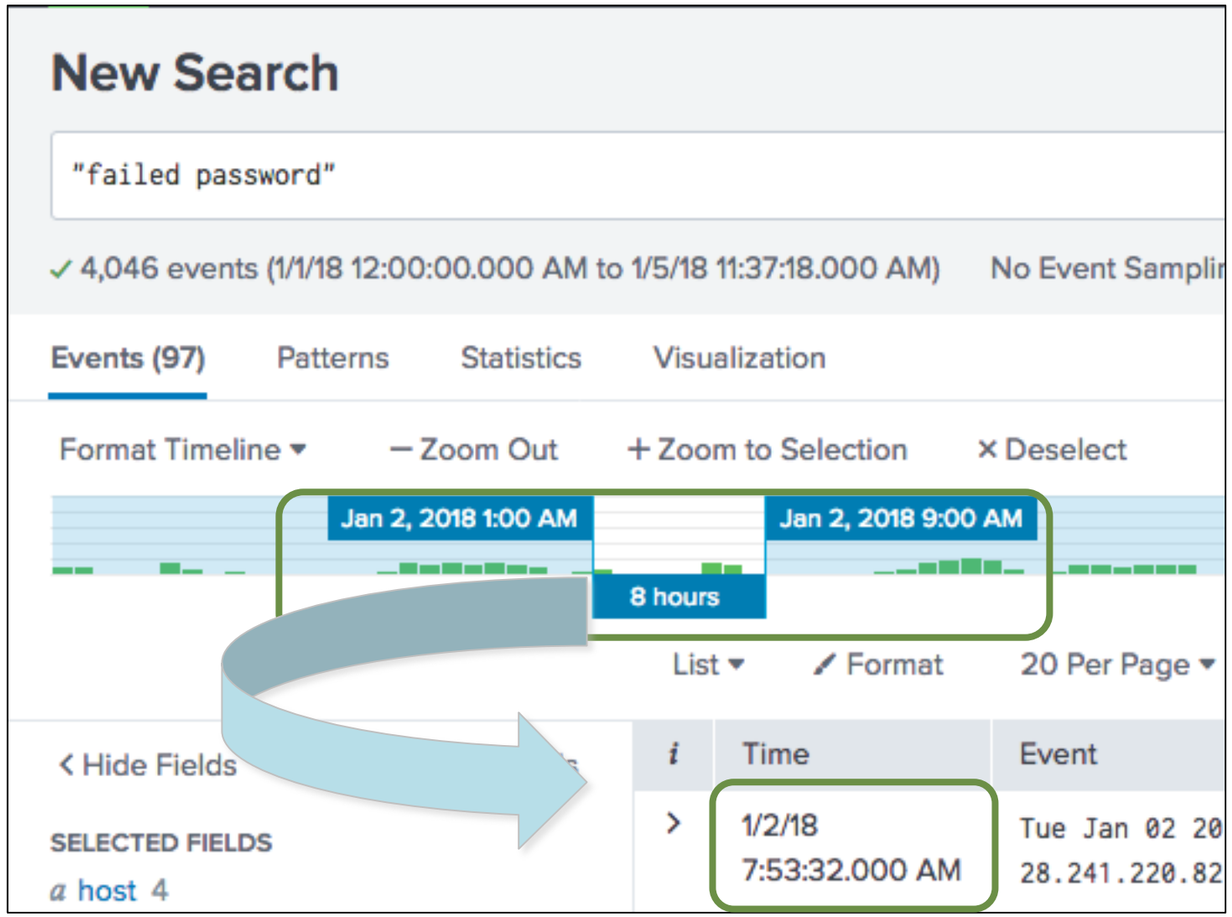

Using Other Timeline Controls

- **Format Timeline**
  - Hides or shows the timeline in different views
- **Zoom Out**
  - Expands the time focus and re-executes the search
- **Zoom to Selection**
  - Narrows the time range and re-executes the search
- **Deselect**
  - If in a drilldown, returns to the original results set
  - Otherwise, grayed out / unavailable

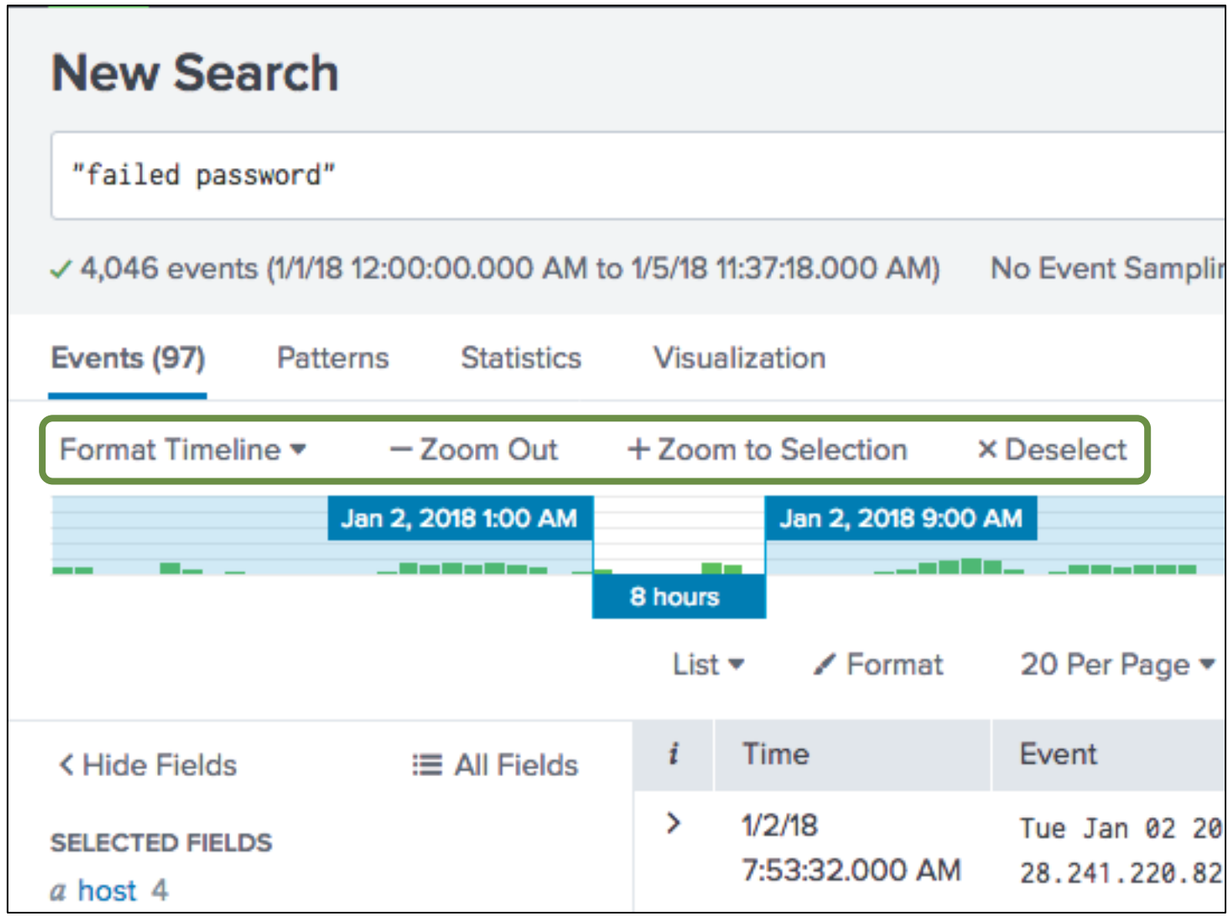

Controlling and Saving Search Jobs

- Every search is also a **job**
- Use the Job bar to control search execution
  - **Pause** – toggles to resume the search
  - **Stop** – finalizes the search in progress
    - Jobs are available for 10 minutes (default)
    - Get a link to results from the **Job** menu

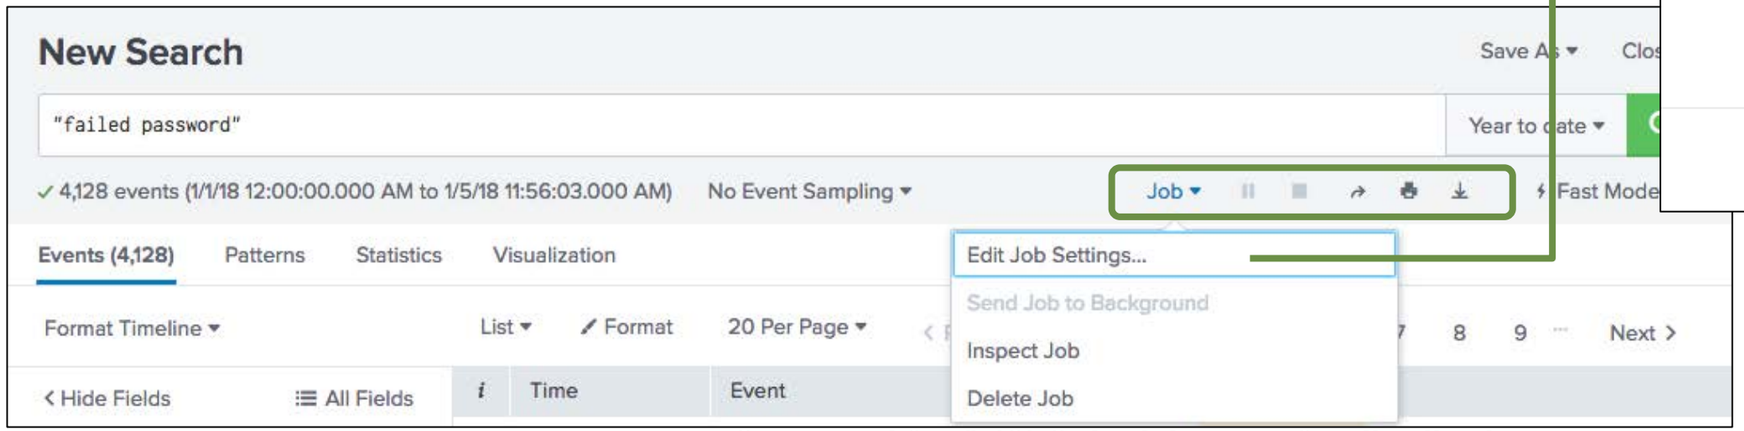

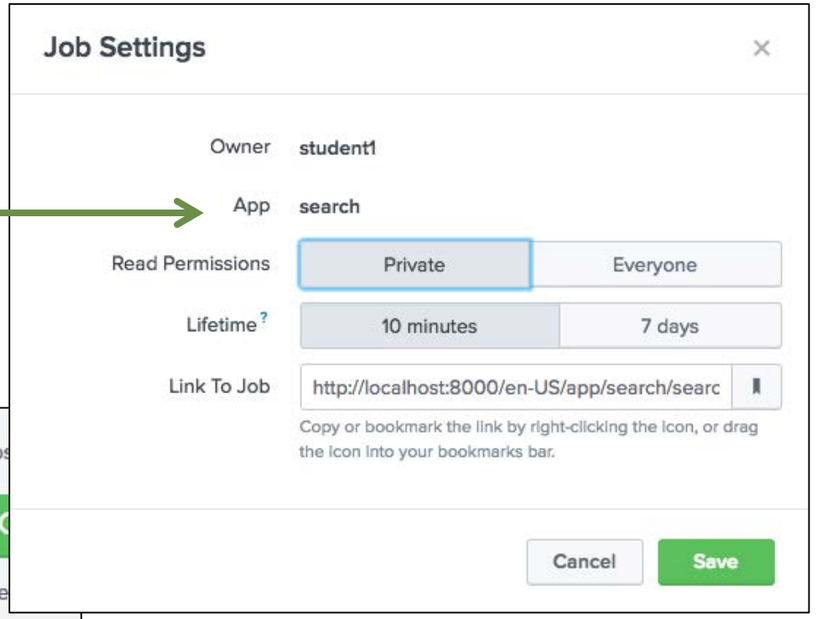

Setting Permissions

- **Private** [default]
  - Only the creator can access
- **Everyone**
  - All app users can access search results
- **Lifetime**
  - Default is 10 minutes
  - Can be extended to 7 days
  - To keep your search results longer, schedule a report

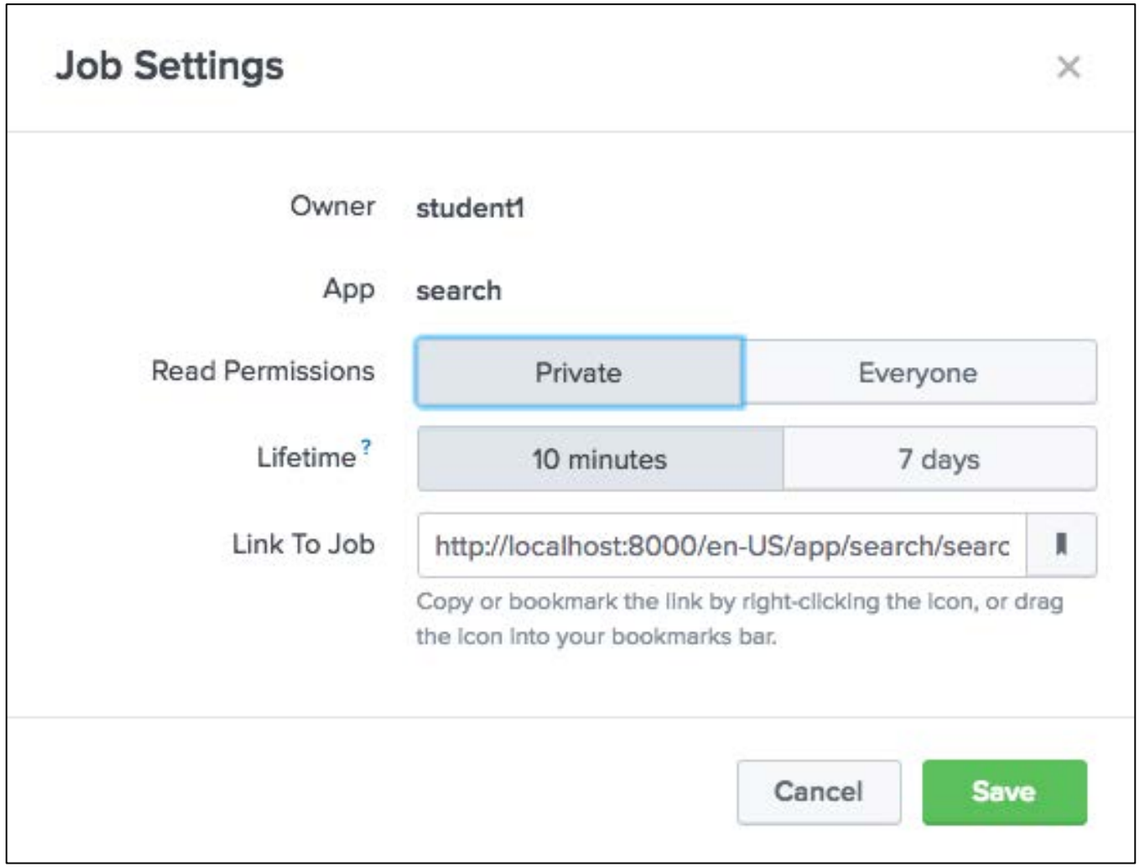

Sharing Search Jobs

- Use the Share button next to the Job bar to quickly:
  - Give everyone read permissions
  - Extend results retention to 7 days
  - Get a sharable link to the results
- Sharing search allows multiple users working on same issue to see same data
  - More efficient than each running search separately
  - Less load on server and disk space used

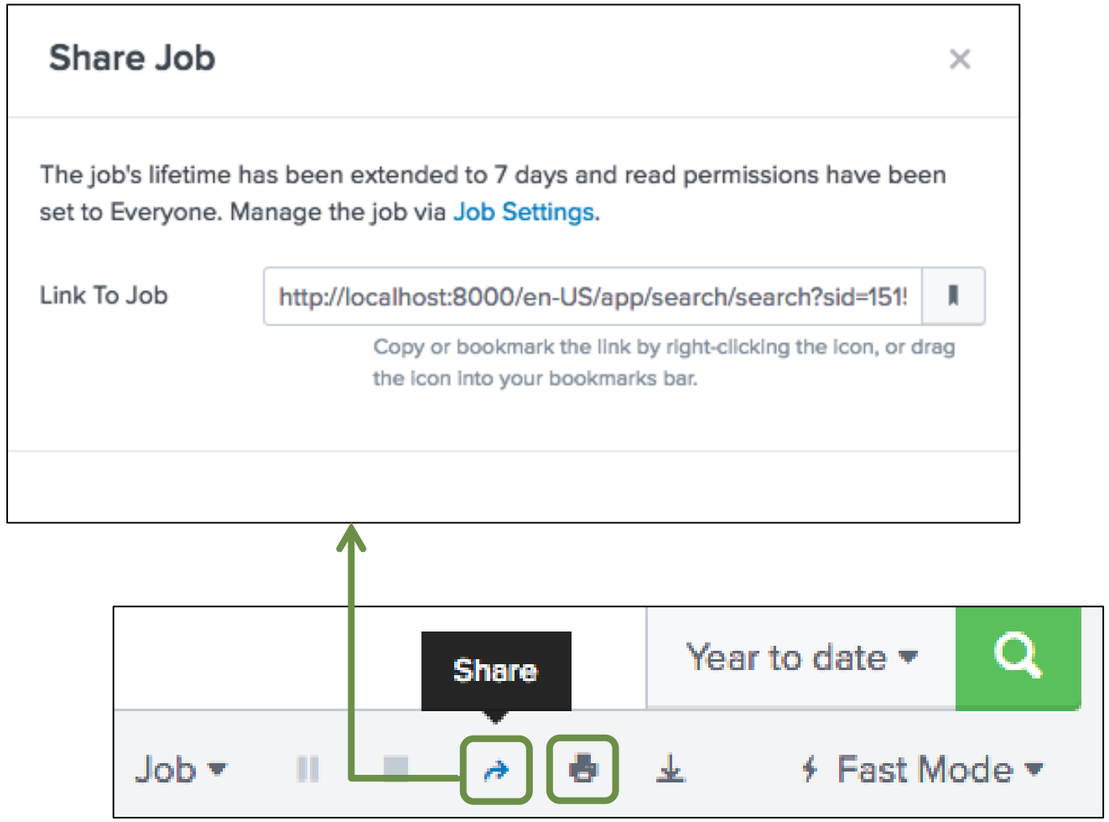

- Can also click printer icon to print results or save as PDF

Exporting Search Results

For an external copy of the results, **export** search results to Raw Events (text file), CSV, XML, or JSON format

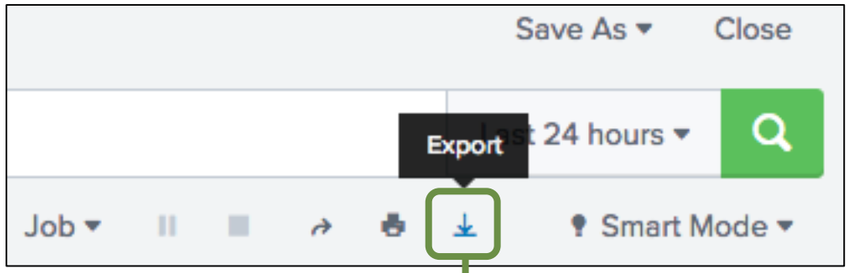

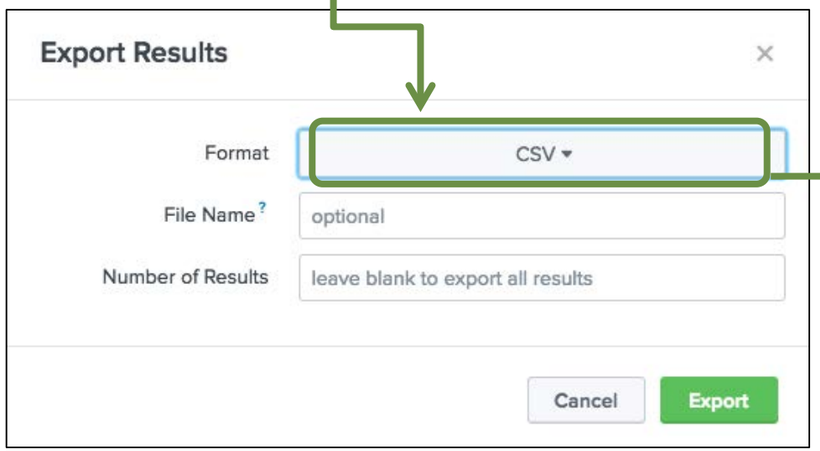

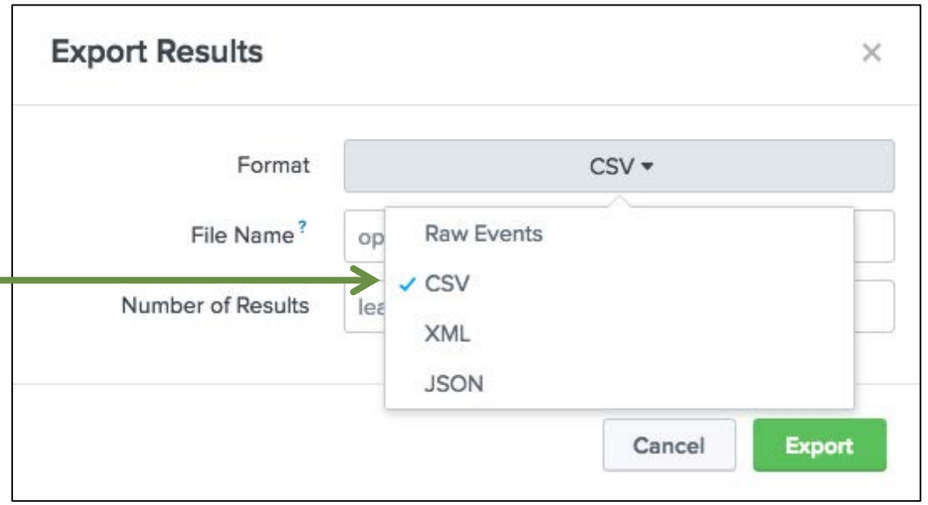

Note

Note that exporting the results of a large search is very memory-intensive!

Viewing Your Saved Jobs

- Access saved search jobs from the **Activity** menu
- The Search Jobs view displays jobs that:
  - You have run in the last 10 minutes
  - You have extended for 7 days
- Click on a job link to view the results in the designated app view

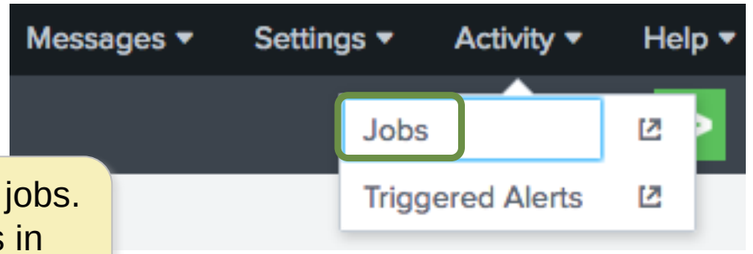

Click **Activity > Jobs** to view your saved jobs. Click the job's name to examine results in Search view. (The job name is the search string.)

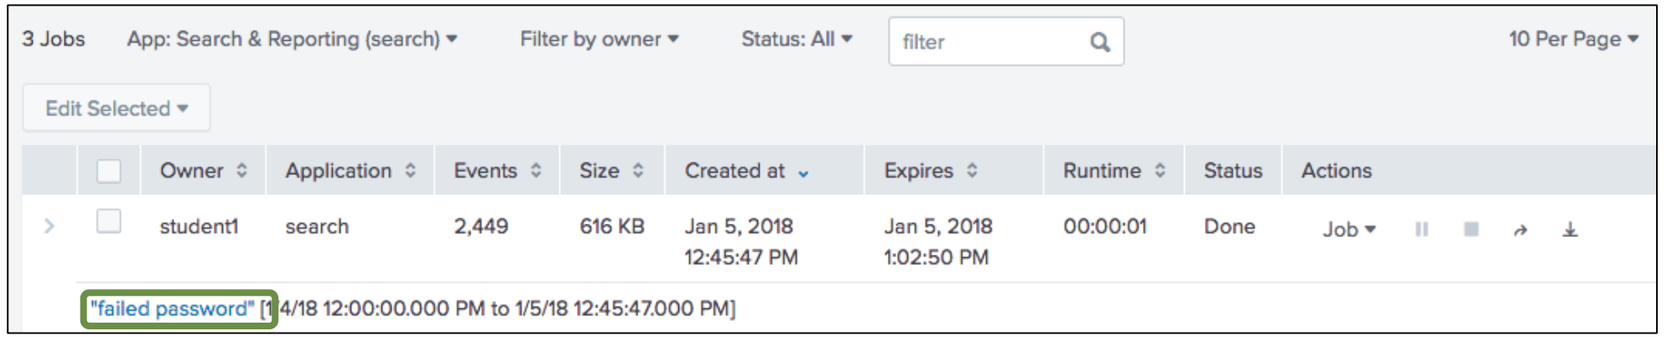

Viewing Your Search History

1. Search History displays your most recent ad-hoc searches – 5 per page

2. You can set a time filter to further narrow your results

3. Click the > icon in the leftmost column to expand long queries to display the full text

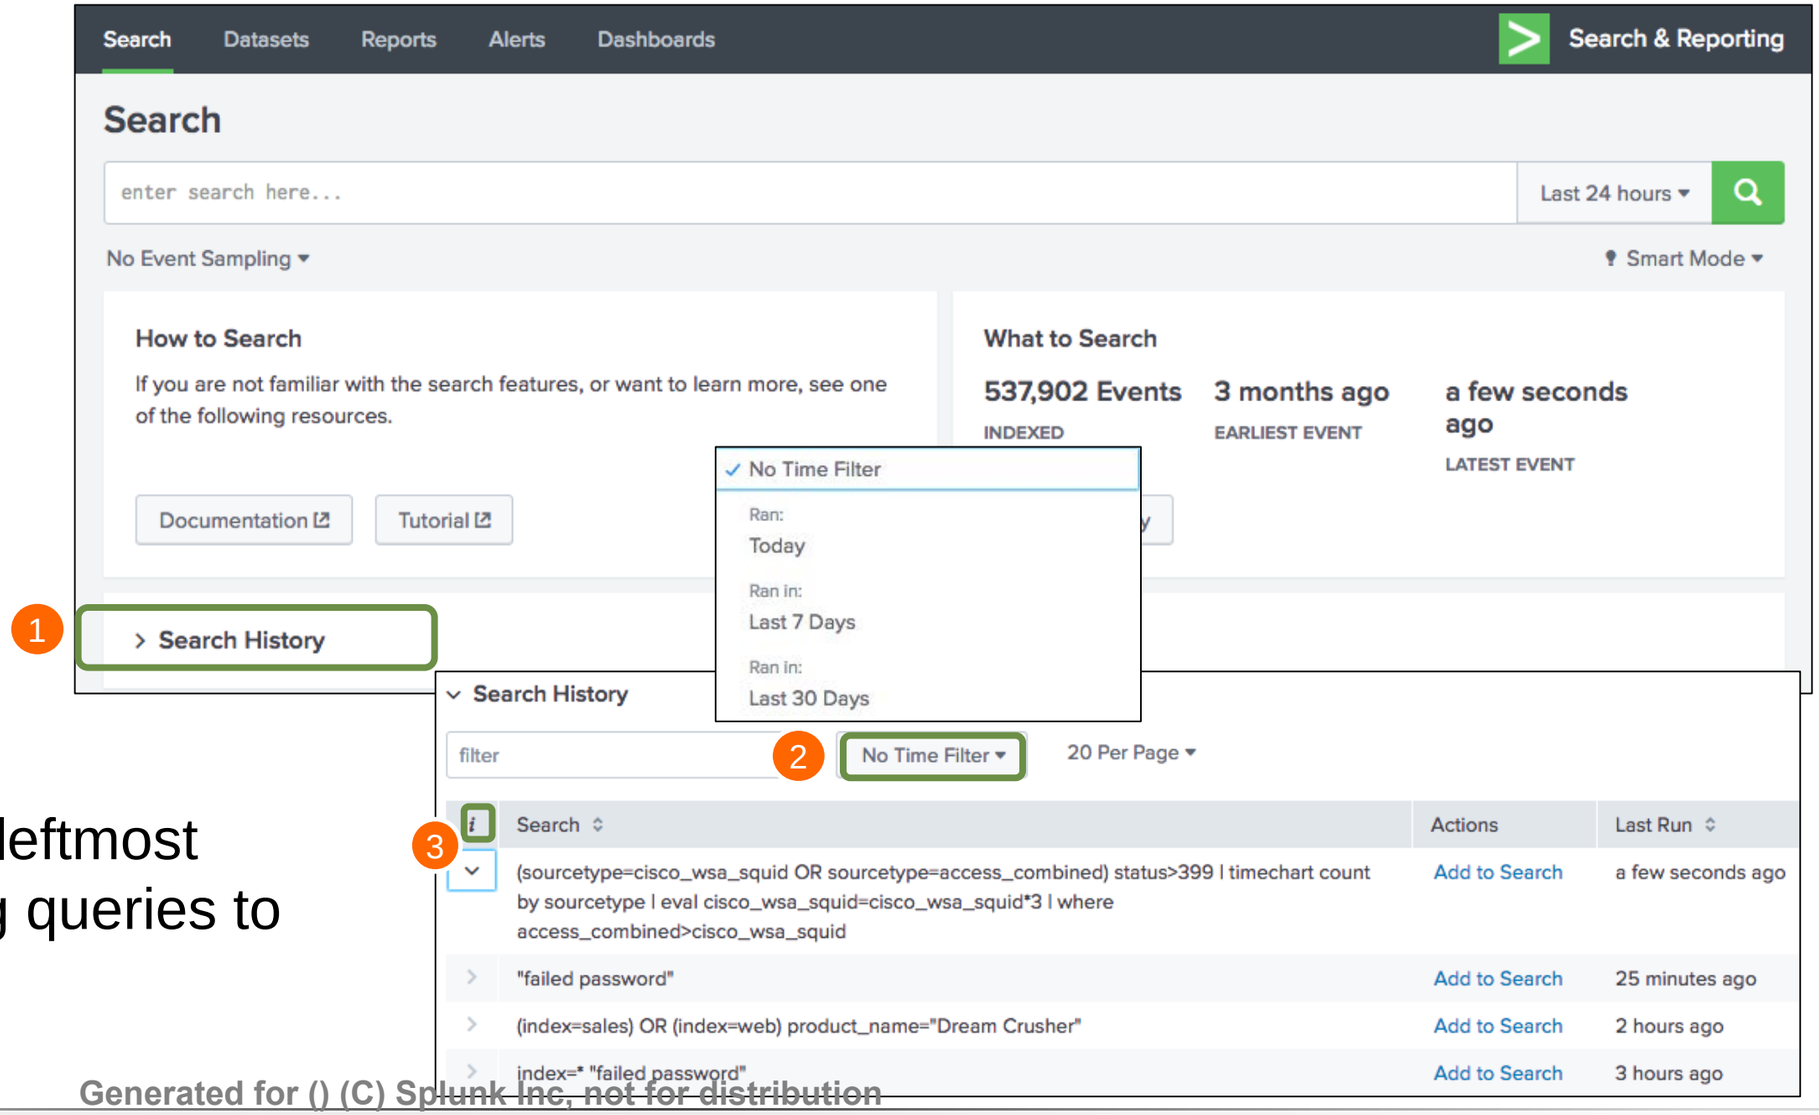

Module 6:
Using Fields in Searches

What Are Fields?

- Fields are searchable key/value pairs in your event data
  - Examples: `host=www1 status=503`
- Fields can be searched with their names, like separating an http status code of 404 from Atlanta's area code (`area_code=404`)
- Between search terms, `AND` is implied unless otherwise specified

Field Discovery

- Splunk automatically discovers many fields based on sourcetype and key/value pairs found in the data
- Prior to search time, some fields are already stored with the event in the index:
  - Meta fields, such as **`host`**, **`source`**, **`sourcetype`**, and **`index`**
  - Internal fields such as **`_time`** and **`_raw`**
- At search time, *field discovery* discovers fields directly related to the search's results
- Some fields in the overall data may not appear within the results of a particular search

Note

While Splunk auto-extracts many fields, you can learn how to create your own in the *Splunk Fundamentals 2* course.

Identify Data-Specific Fields

- Data-specific fields come from the specific characteristics of your data
  - Sometimes, this is indicated by obvious key = value pairs (**`action = purchase`**)
  - Sometimes, this comes from data within the event, defined by the sourcetype (**`status = 200`**)

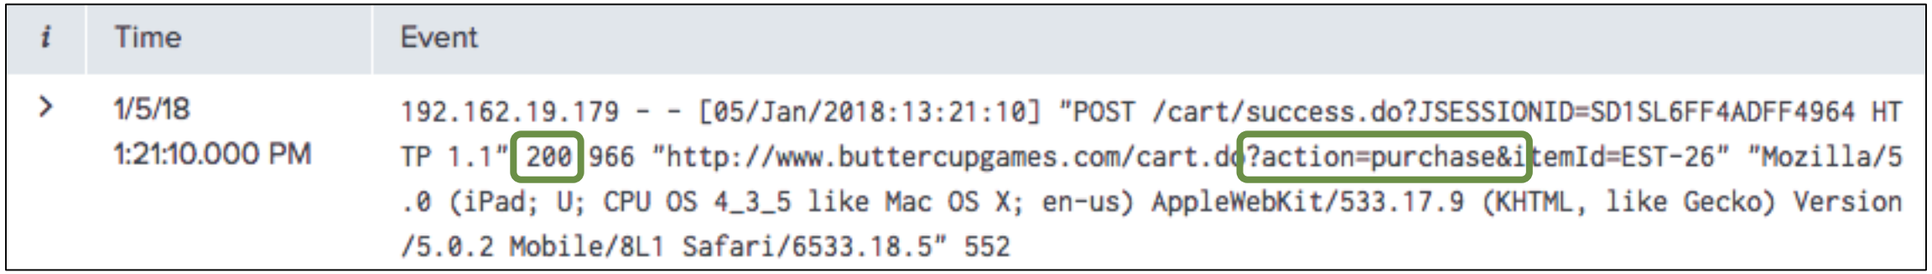

i	Time	Event
>	1/5/18 1:21:10.000 PM	192.162.19.179 - - [05/Jan/2018:13:21:10] "POST /cart/success.do?JSESSIONID=SD1SL6FF4ADFF4964 HTTP 1.1" 200 966 "http://www.buttercupgames.com/cart.do?action=purchase&itemId=EST-26" "Mozilla/5.0 (iPad; U; CPU OS 4_3_5 like Mac OS X; en-us) AppleWebKit/533.17.9 (KHTML, like Gecko) Version/5.0.2 Mobile/8L1 Safari/6533.18.5" 552

Note

For more information, please see:
http://docs.splunk.com/Documentation/Splunk/latest/Data/Listofpretrainedsourcetypes

Fields Sidebar

For the current search:

- **Selected Fields** – a set of configurable fields displayed for each event
- **Interesting Fields** – occur in at least 20% of resulting events
- **All Fields** link to view all fields (including non-interesting fields)

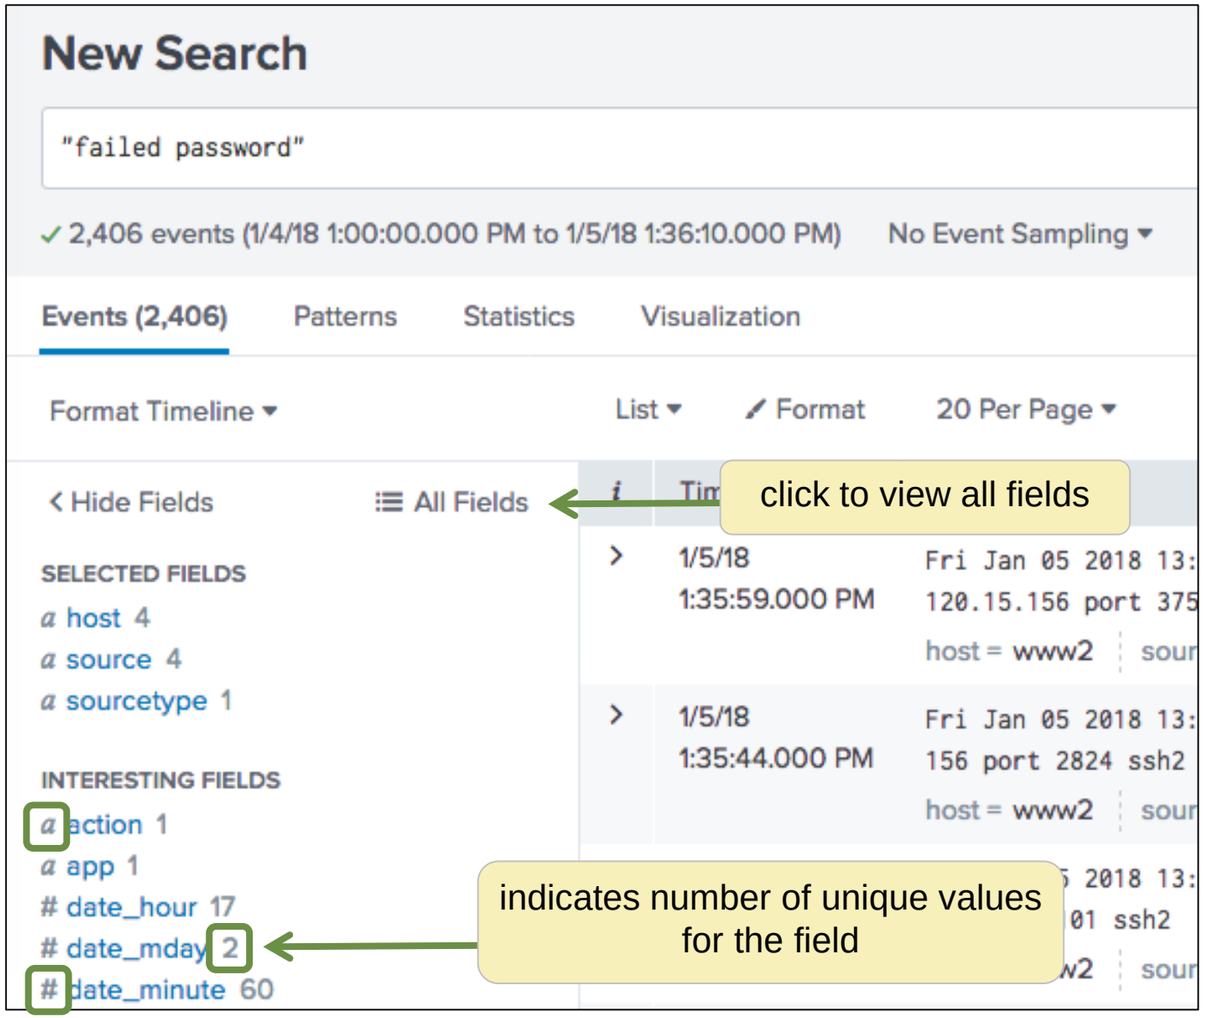

indicates the field's values are alpha-numeric

indicates that the majority of the field values are numeric

Describe Selected Fields

- Selected fields and their values are listed under every event that includes those fields
- By default, the selected fields are:
  - host
  - source
  - sourcetype
- You can choose any field and make it a selected field

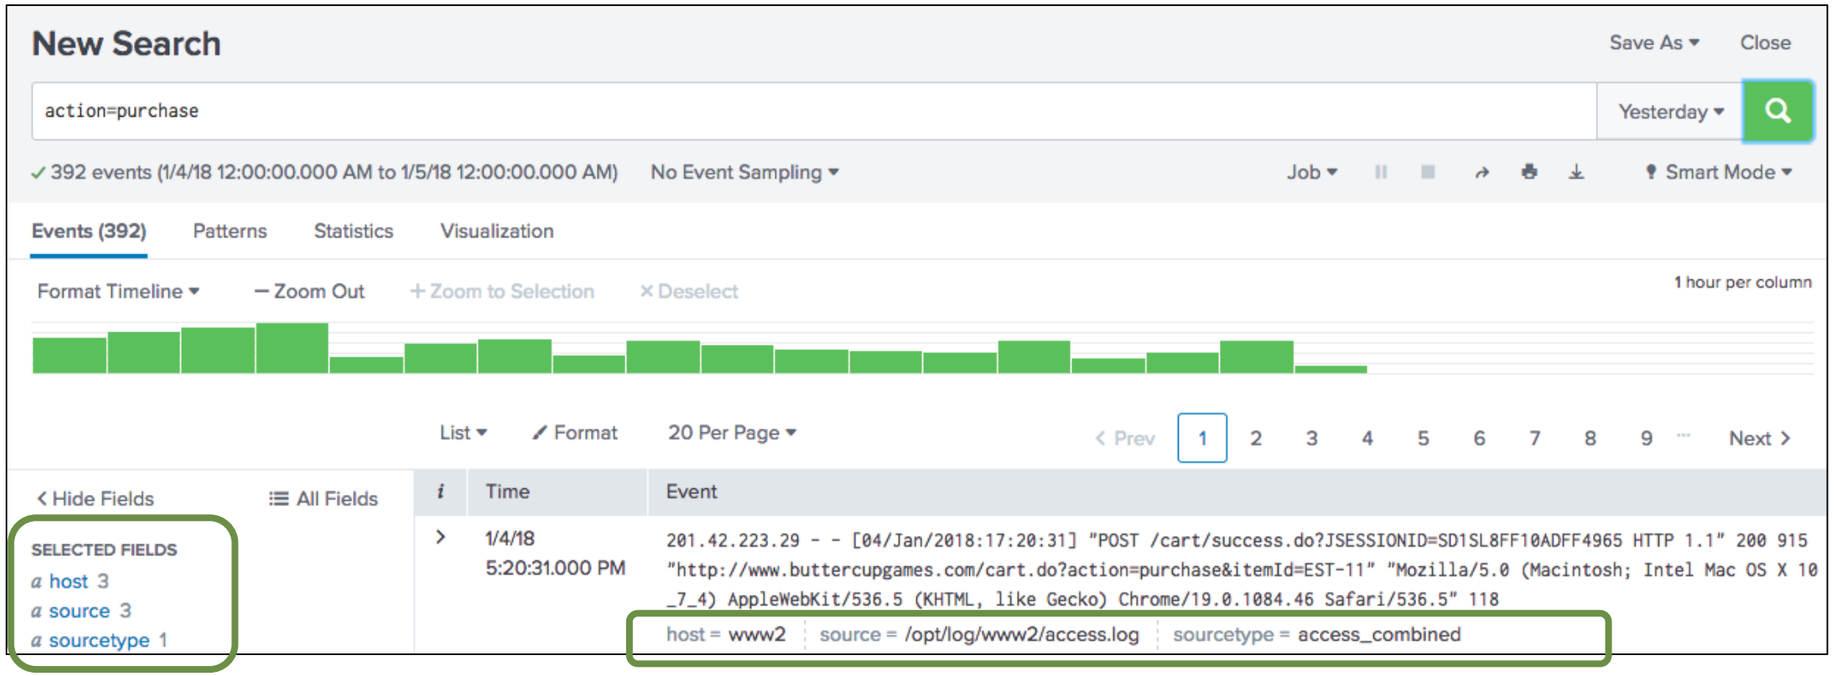

Make an Interesting Field a Selected Field

- You can modify selected fields
  1. Click a field in the Fields sidebar
  2. Click **Yes** in the upper right of the field dialog
- Note that a selected field appears:
  - In the Selected Fields section of the Fields sidebar
  - Below each event where a value exists for that field

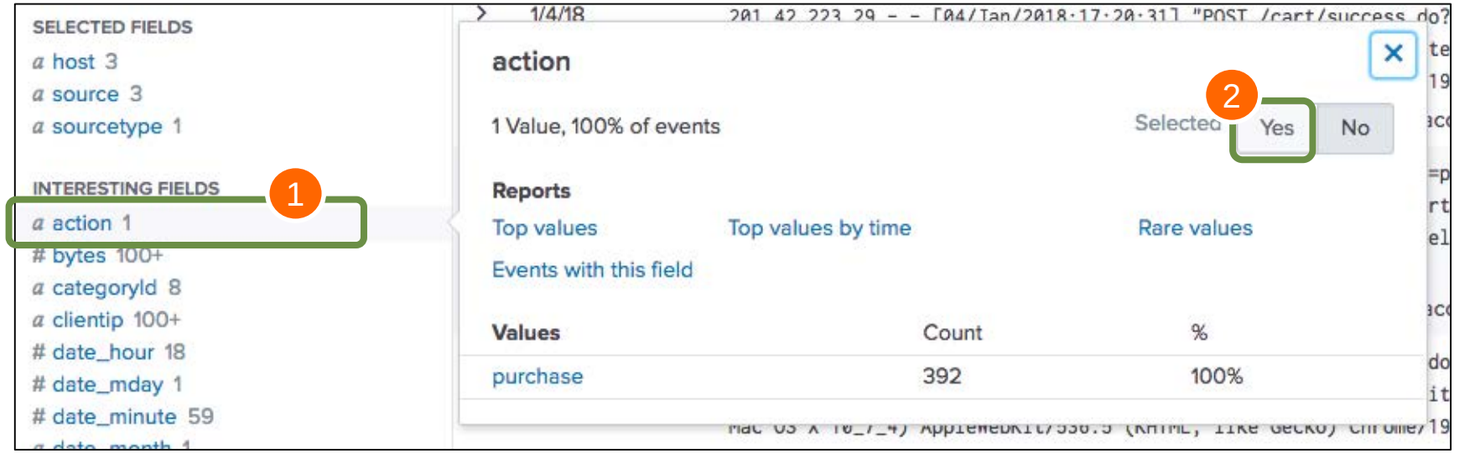

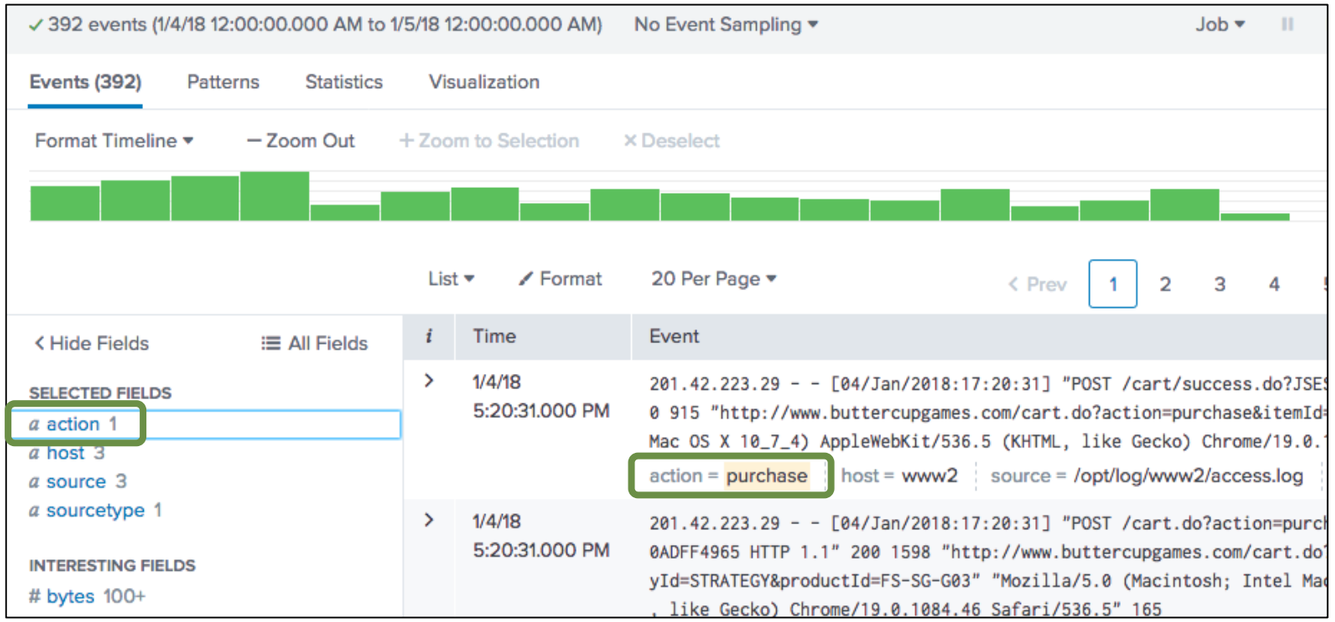

Make Any Field Selected

You can identify other fields as selected fields from All Fields (which shows all of the discovered fields)

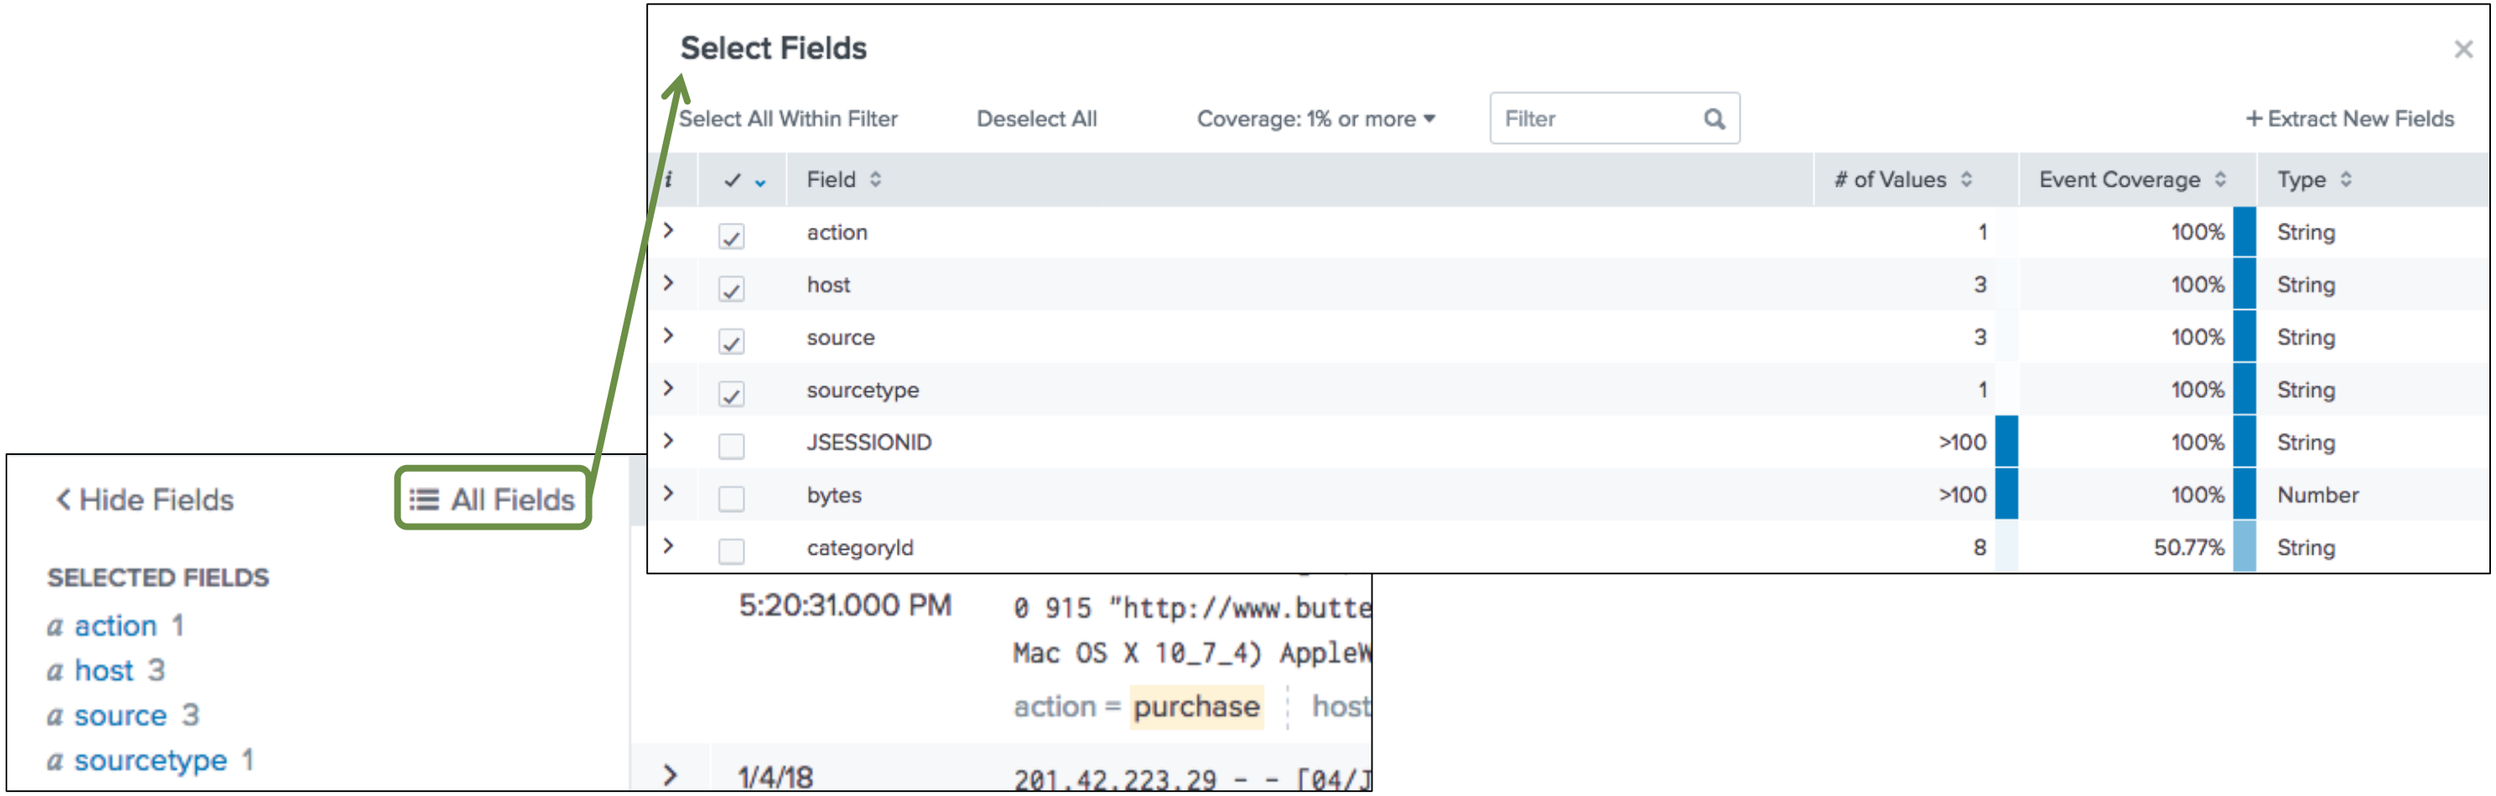

The Field Window

Select a field from the Fields sidebar, then:

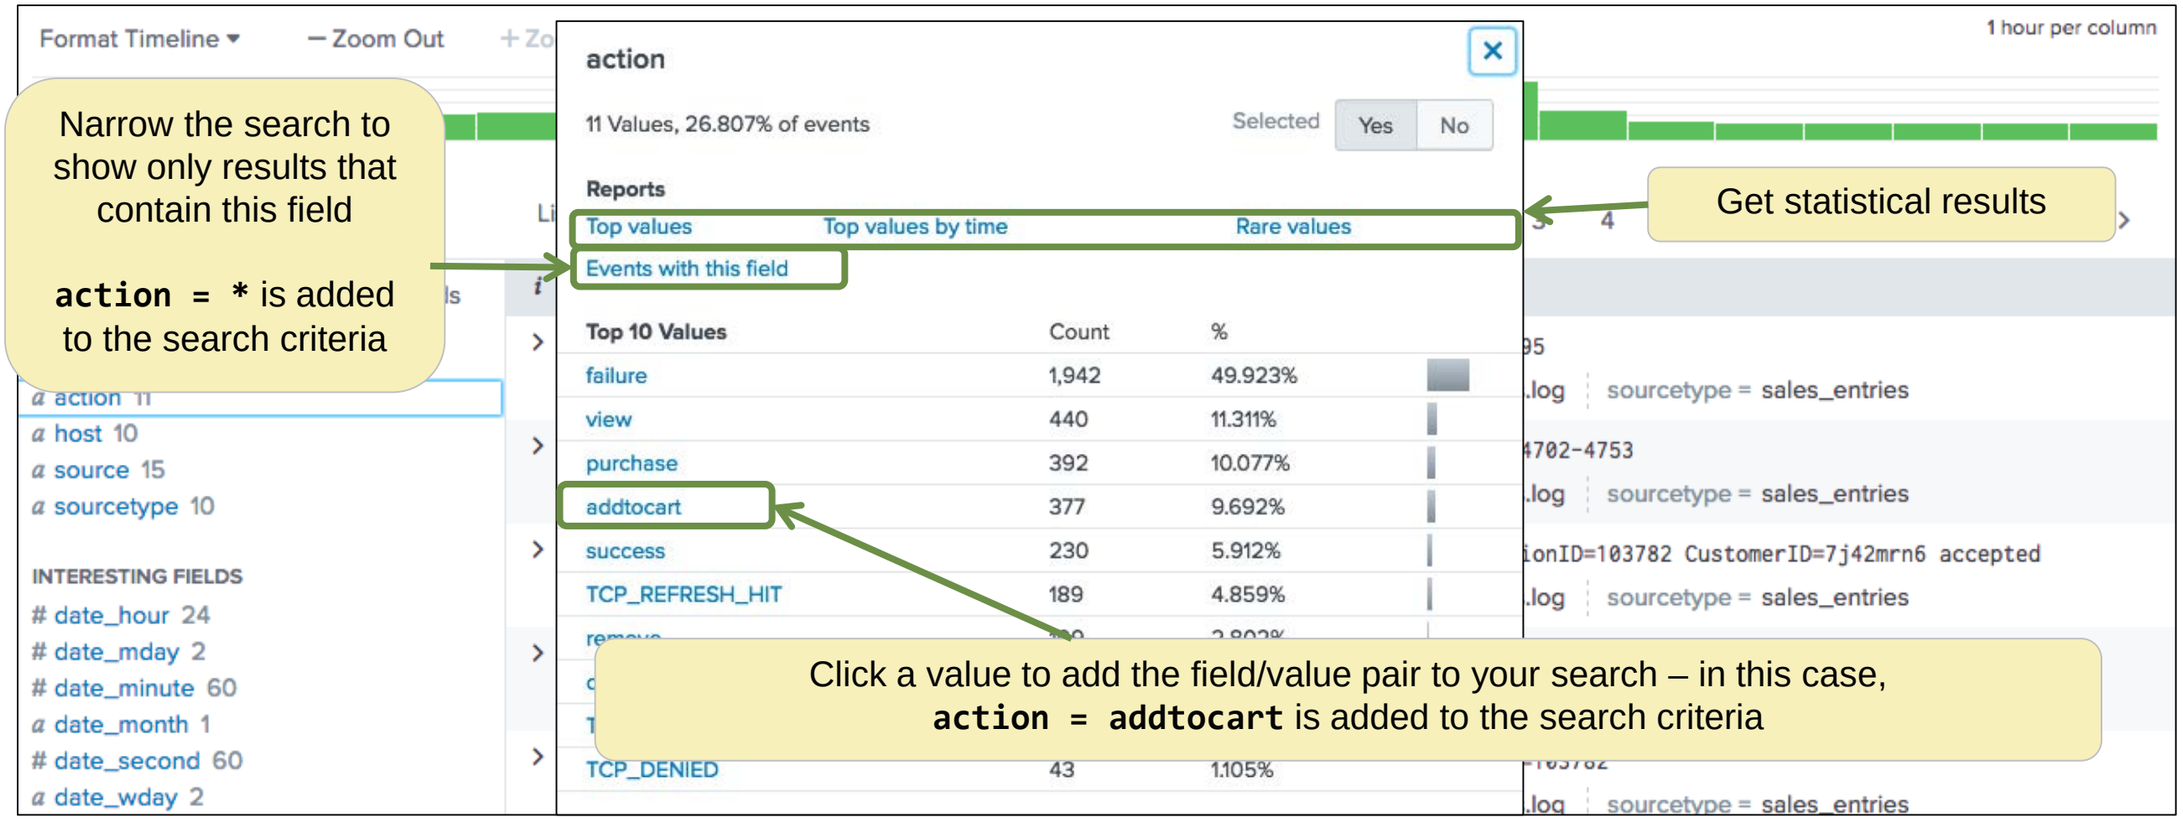

Using Fields in Searches

- Efficient way to pinpoint searches and refine results

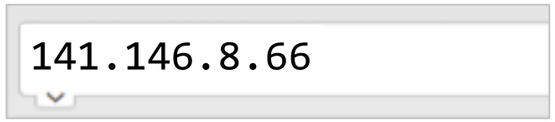

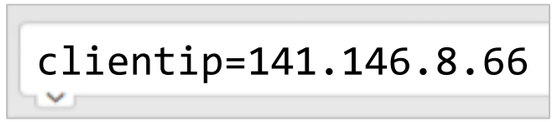

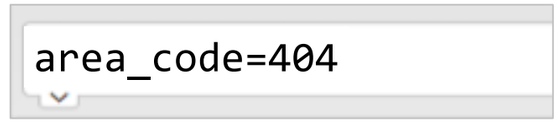

- Field names ARE case sensitive; field values are NOT
  - Example:

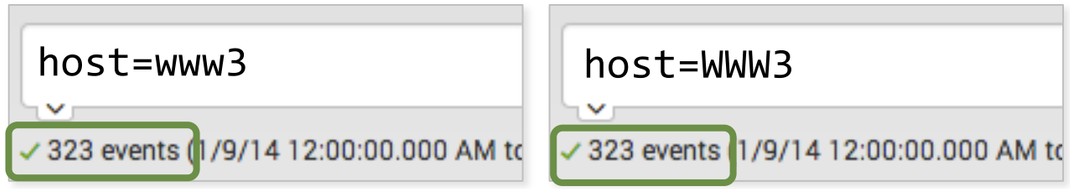

These two searches return results

This one does not return results

Using Fields in Searches (cont.)

- For IP fields, Splunk is subnet/CIDR aware

```
clientip="202.201.1.0/24"
```

```
clientip="202.201.1.*"
```

- Use wildcards to match a range of field values
  - Example: **`user=*`** (to display all events that contain a value for user)

```
user=* sourcetype=access* (referer_domain=*.cn OR referer_domain=*.hk)
```

All time

- Use relational operators

With numeric fields

```
src_port>1000 src_port<4000
```

With alphanumeric fields

```
host!=www3
```


!= vs. NOT

- Both `!=` field expression and `NOT` operator exclude events from your search, but produce different results

- Example: `status != 200`
  - Returns events where `status` field exists and value in field doesn't equal 200

- Example: `NOT status = 200`
  - Returns events where `status` field exists and value in field doesn't equal 200 -- **and** all events where status field **doesn't** exist

!= vs. NOT (cont.)

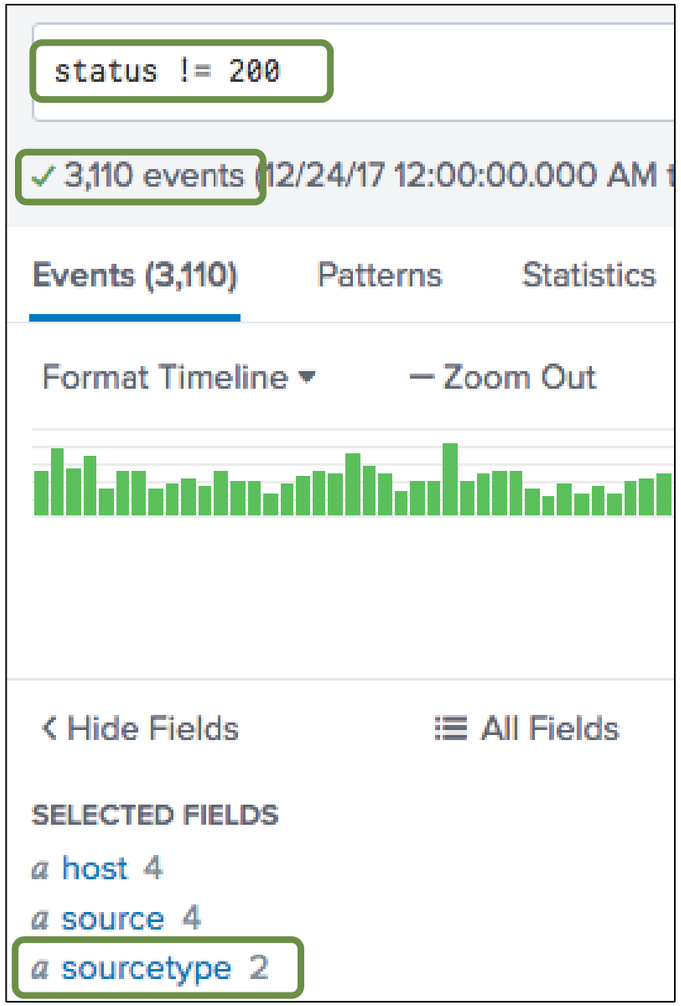

In this example:

- `status != 200` returns **3,110** events from **2** sourcetypes
- `NOT status=200` returns **66,855** events from **9** sourcetypes

Note

The results from a search using != are a **subset** of the results from a similar search using NOT.

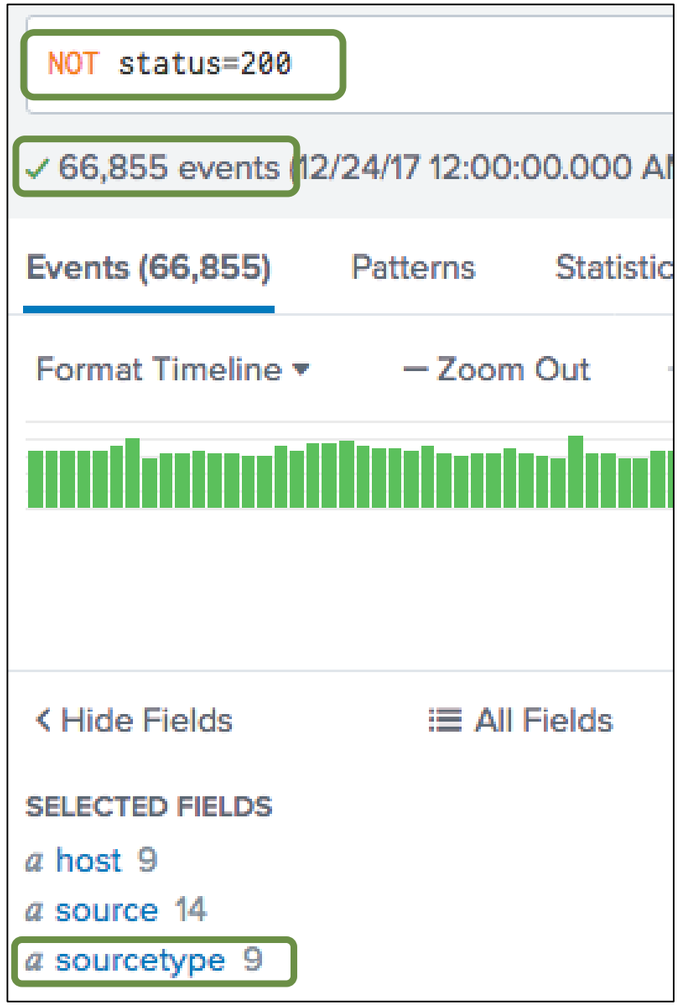

!= vs. NOT (cont.)

- Does `!=` and `NOT` ever yield the same results?
  - Yes, if you know the field you're evaluating always exists in the data you're searching
  - For example:
    - `index=web sourcetype=access_combined status!=200`
    - `index=web sourcetype=access_combined NOT status=200`

    yields same results because `status` field always exists in `access_combined` sourcetype

Search Modes: Fast, Smart, Verbose

- Fast: emphasizes speed over completeness
- Smart: balances speed and completeness (default)
- Verbose:
  - Emphasizes completeness over speed
  - Allows access to underlying events when using reporting or statistical commands (in addition to totals and stats)

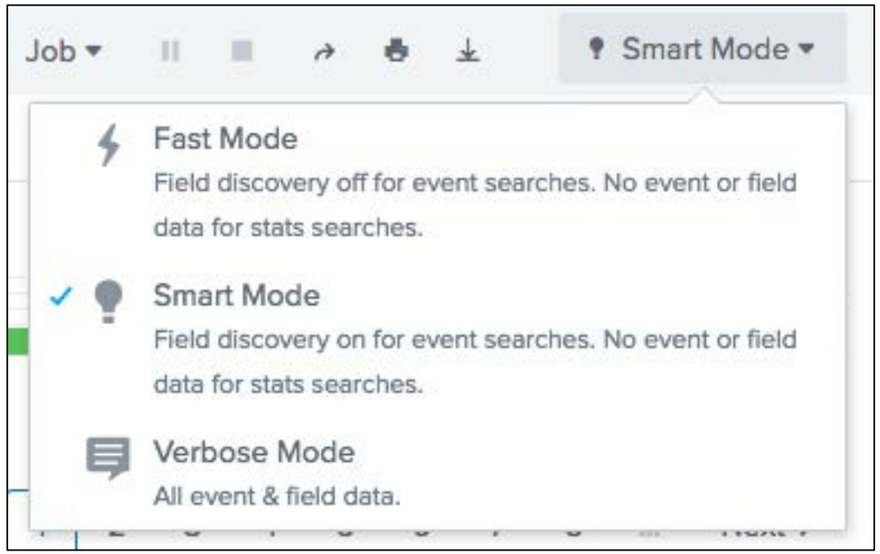

Note

You'll discuss statistical commands later in this course.

Module 7
Best Practices

Search Best Practices

- Time is the most efficient filter
- Specify one or more index values at the beginning of your search string
- Include as many search terms as possible

  If you want to find events with `"error" and "sshd"`, and 90% of the events include `"error"` but only 5% `"sshd"`, include both values in the search
- Make your search terms as specific as possible

  Searching for `"access denied"` is always better than searching for `"denied"`
- Inclusion is generally better than exclusion

  Searching for `"access denied"` is faster than searching for `NOT "access granted"`

Search Best Practices (cont.)

- Filter as early as possible

  For example, remove duplicate events, then sort

- Avoid using wildcards at the beginning or middle of a string
  - Wildcards at *beginning* of string scan all events within timeframe
  - Wildcards in *middle* of string may return inconsistent results
  - So use `fail*` (not `*fail` or `*fail*` or `f*il`)
- When possible, use OR instead of wildcards
  - For example, use `(user=admin OR user=administrator)` instead of `user=admin*`

Note

Remember, field names are case *sensitive* and field values are case *insensitive.*

Working with Indexes

- This search returns event data from the security index

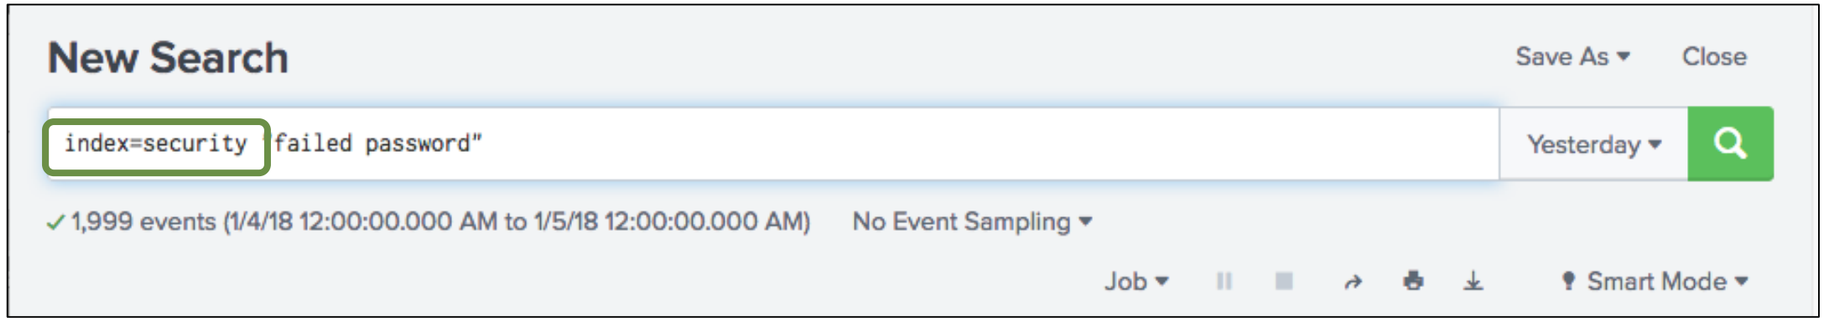

- It's possible to specify multiple index values in a search

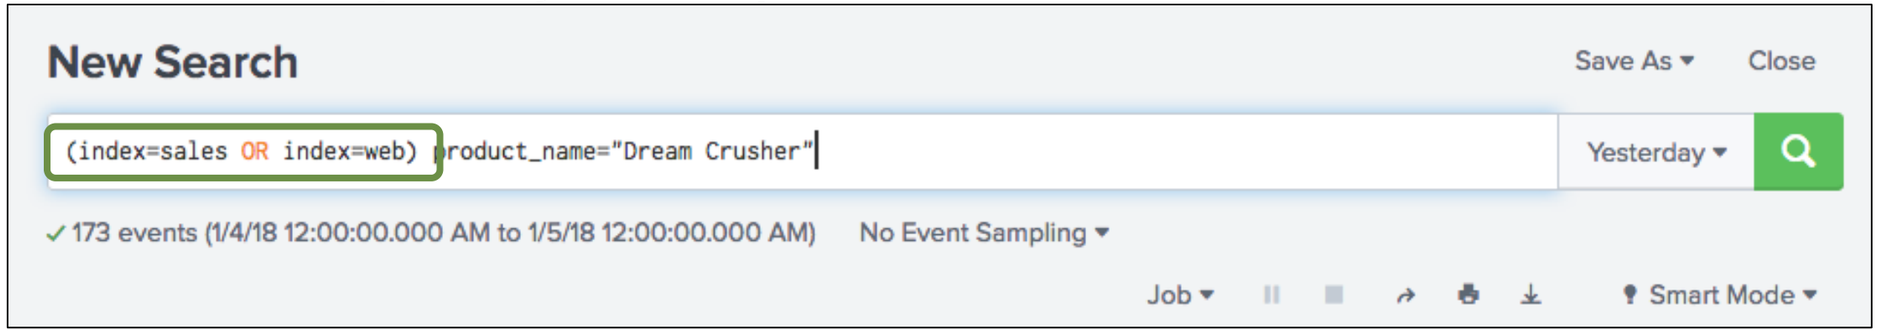

Working with Indexes (cont.)

- It's possible to use a wildcard (*) in index values

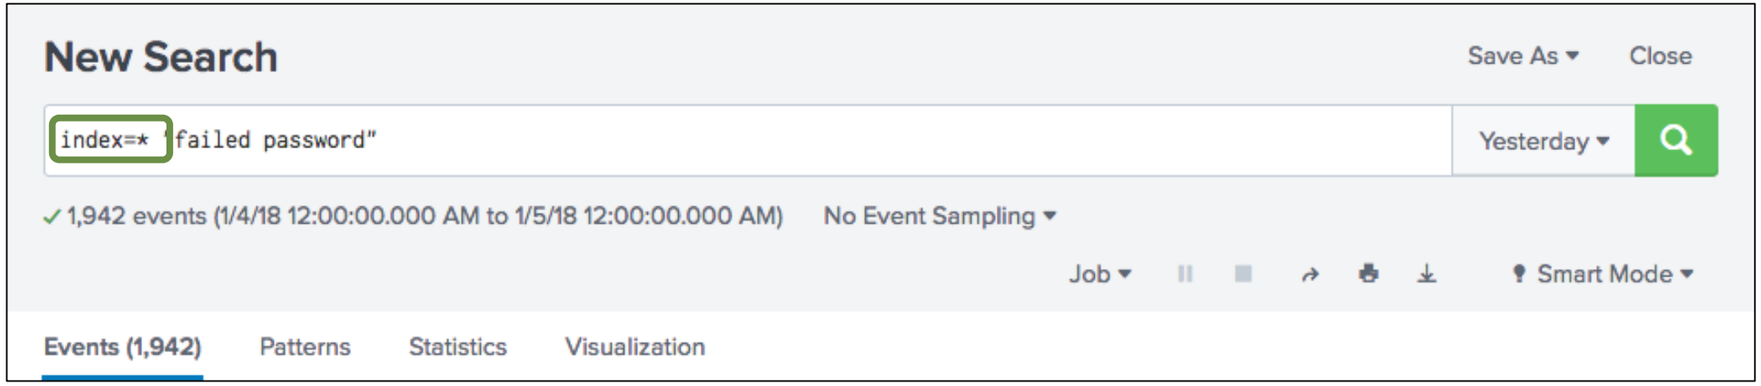

- It's also possible to search *without* an index—but that's inefficient and **not** recommended

Note 1

Although `index=*` is a valid search, better performance is always obtained by specifying one or more specific index values.

Note 2

For best performance, specify the index values at the **beginning** of the search string.

Viewing the Index Field

- The *index* always appears as a field in search results
- In the search shown here, no index was indicated in the search, so data is returned from two indexes: `web` and `sales`
- Remember, this practice is **not** recommended—it's always more efficient to specify one or more indexes in your search

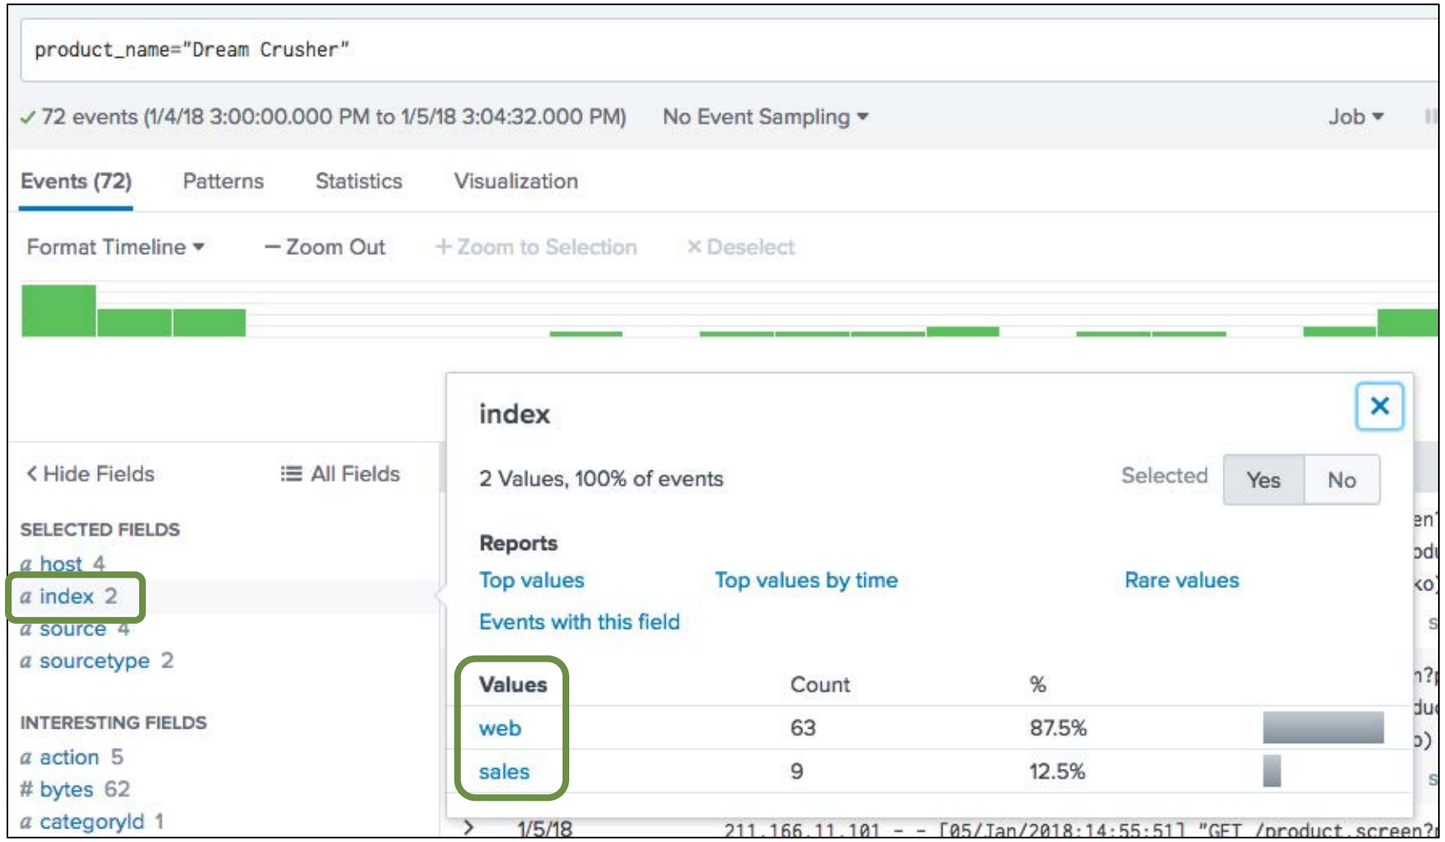

Module 8:
Splunk's Search Language

Search Language Syntax

This diagram represents a search, broken into its syntax components:

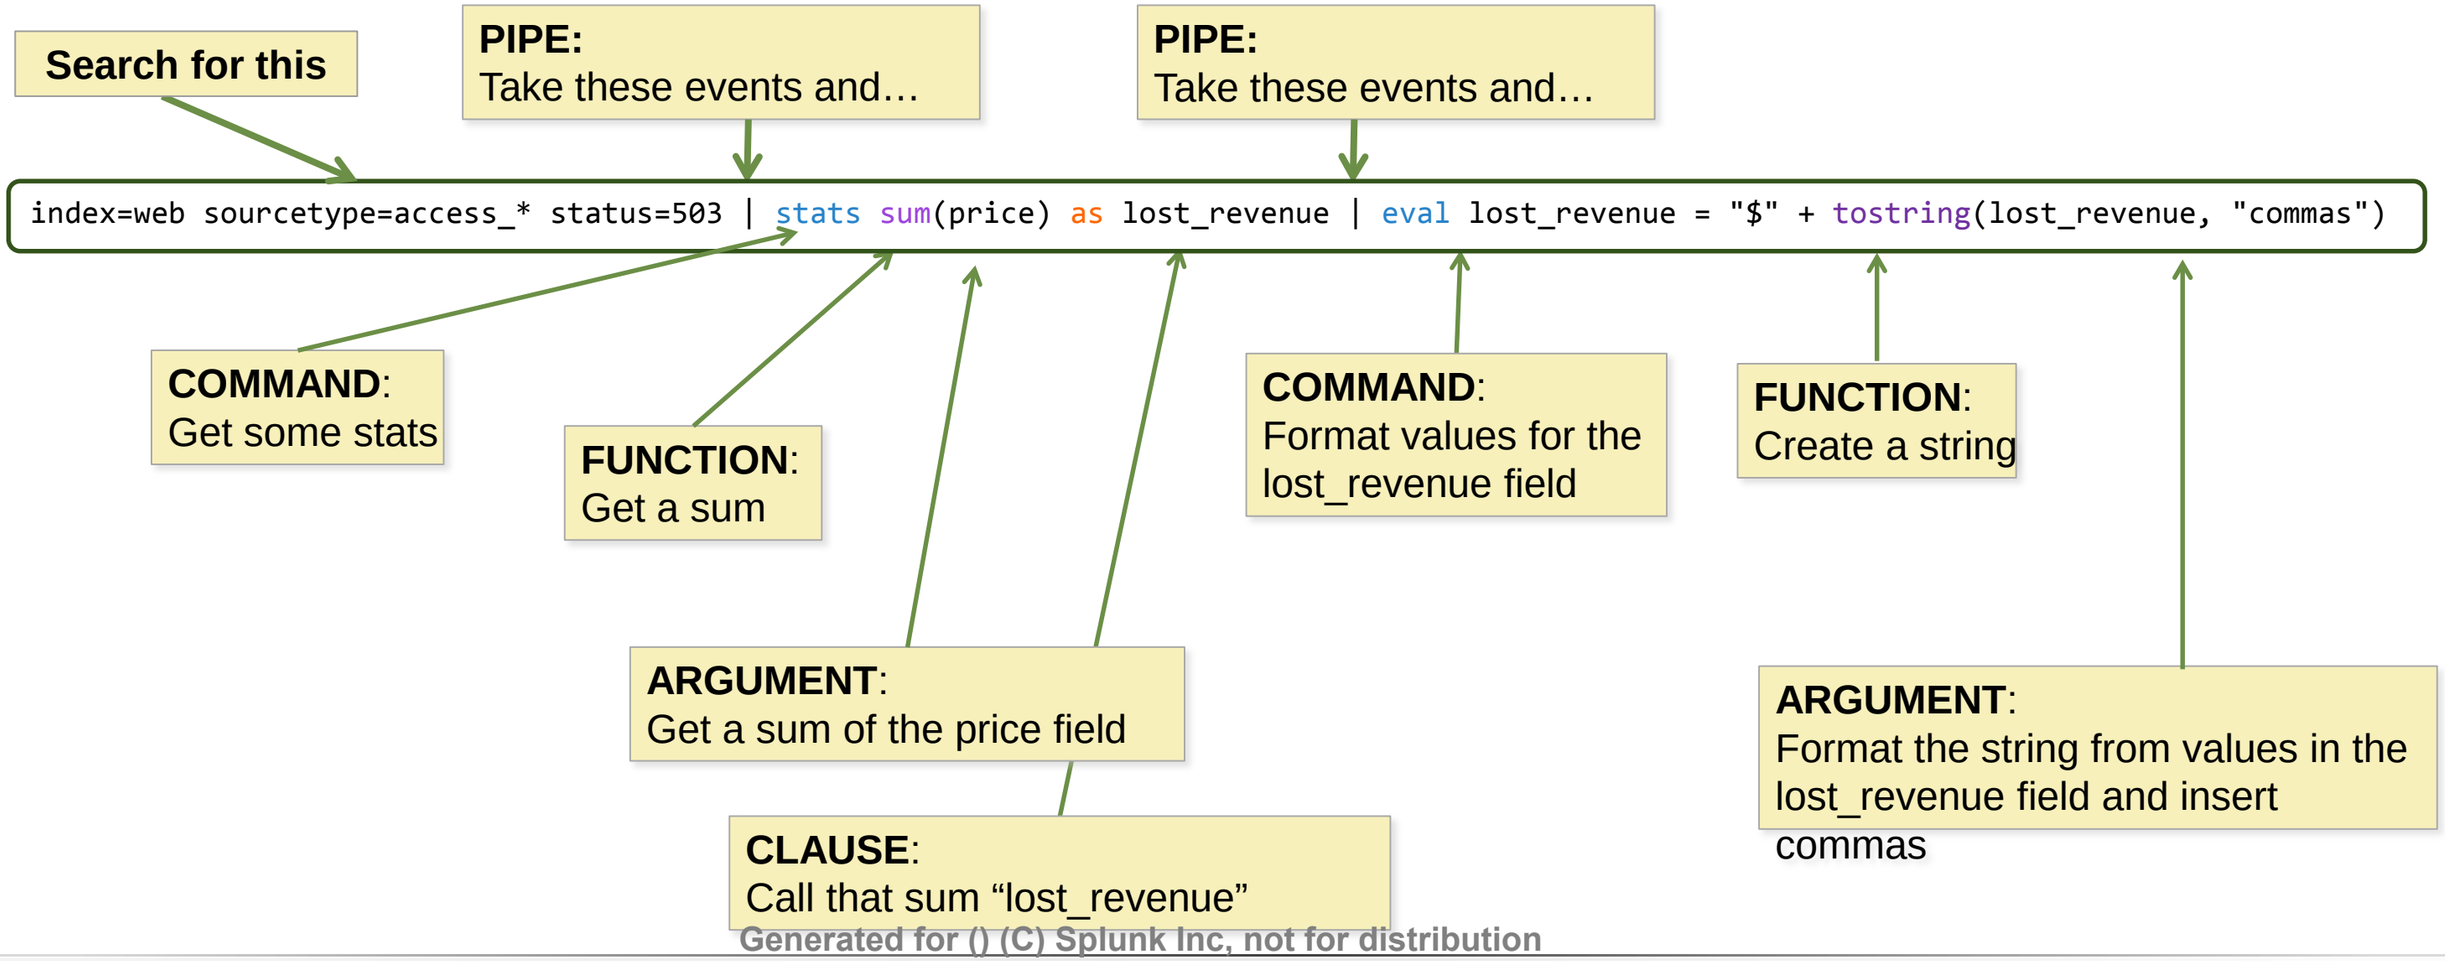

Search Language Syntax Components

- Searches are made up of 5 basic components
  1. **Search terms** – what are you looking for?
     - Keywords, phrases, Booleans, etc.
  2. **Commands** – what do you want to do with the results?
     - Create a chart, compute statistics, evaluate and format, etc.
  3. **Functions** – how do you want to chart, compute, or evaluate the results?
     - Get a sum, get an average, transform the values, etc.
  4. **Arguments** – are there variables you want to apply to this function?
     - Calculate average value for a specific field, convert milliseconds to seconds, etc.
  5. **Clauses –** how do you want to group or rename the fields in the results?
     - Give a field another name or group values by or over

The Search Pipeline

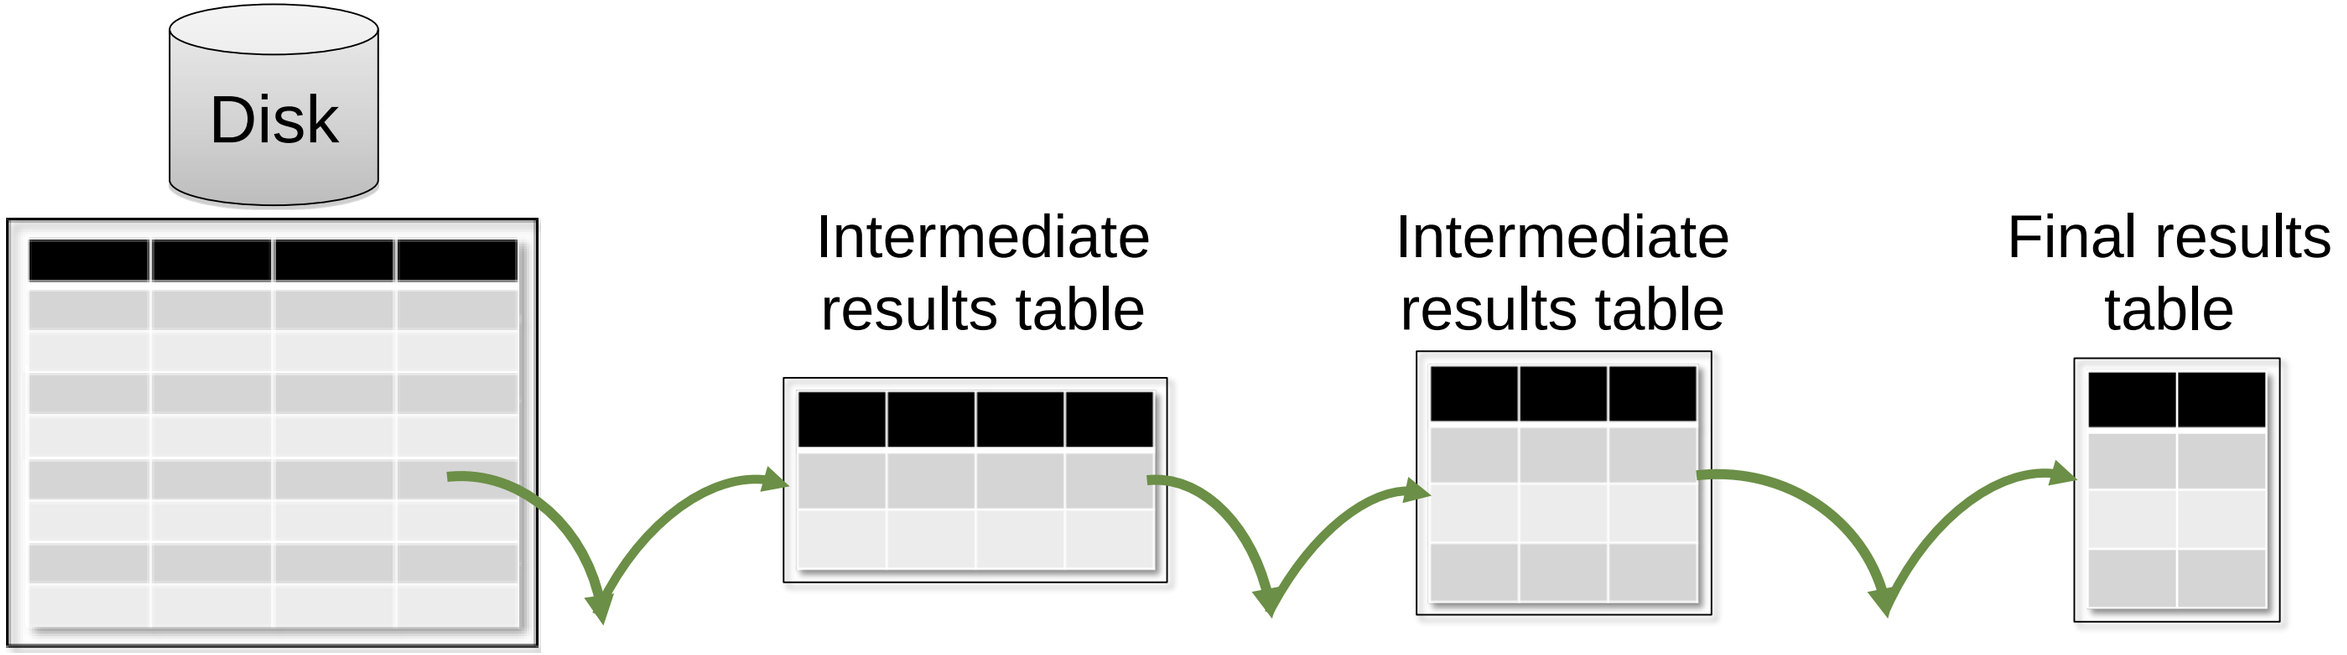


```
index=security sourcetype=linux_secure fail*   | top user   | fields – percent
```

Fetch events from disk that match

Summarize into table of top 10 users

Remove column showing percentage

Making the Pipeline More Readable

- Put each pipe in the pipeline on a separate line as you type by turning on auto-formatting
- Go to **Preferences > SPL Editor** and turn on Search auto-format

Instead of this:

New Search

```
index=web sourcetype=access_* status=503 | stats sum(price) as lost_revenue | eval lost_revenue = "$" + tostring(lost_revenue, "commas")
```

You'll get this:

New Search

```
index=web sourcetype=access_* status=503
| stats sum(price) as lost_revenue
| eval lost_revenue = "$" + tostring(lost_revenue, "commas")
```

Note

You can also use **Shift + Enter** to go to a new line.

Syntax Coloring

- By default, some parts of the search string are automatically colored as you type
- The color is based on the search syntax
- The rest of the search string remains black

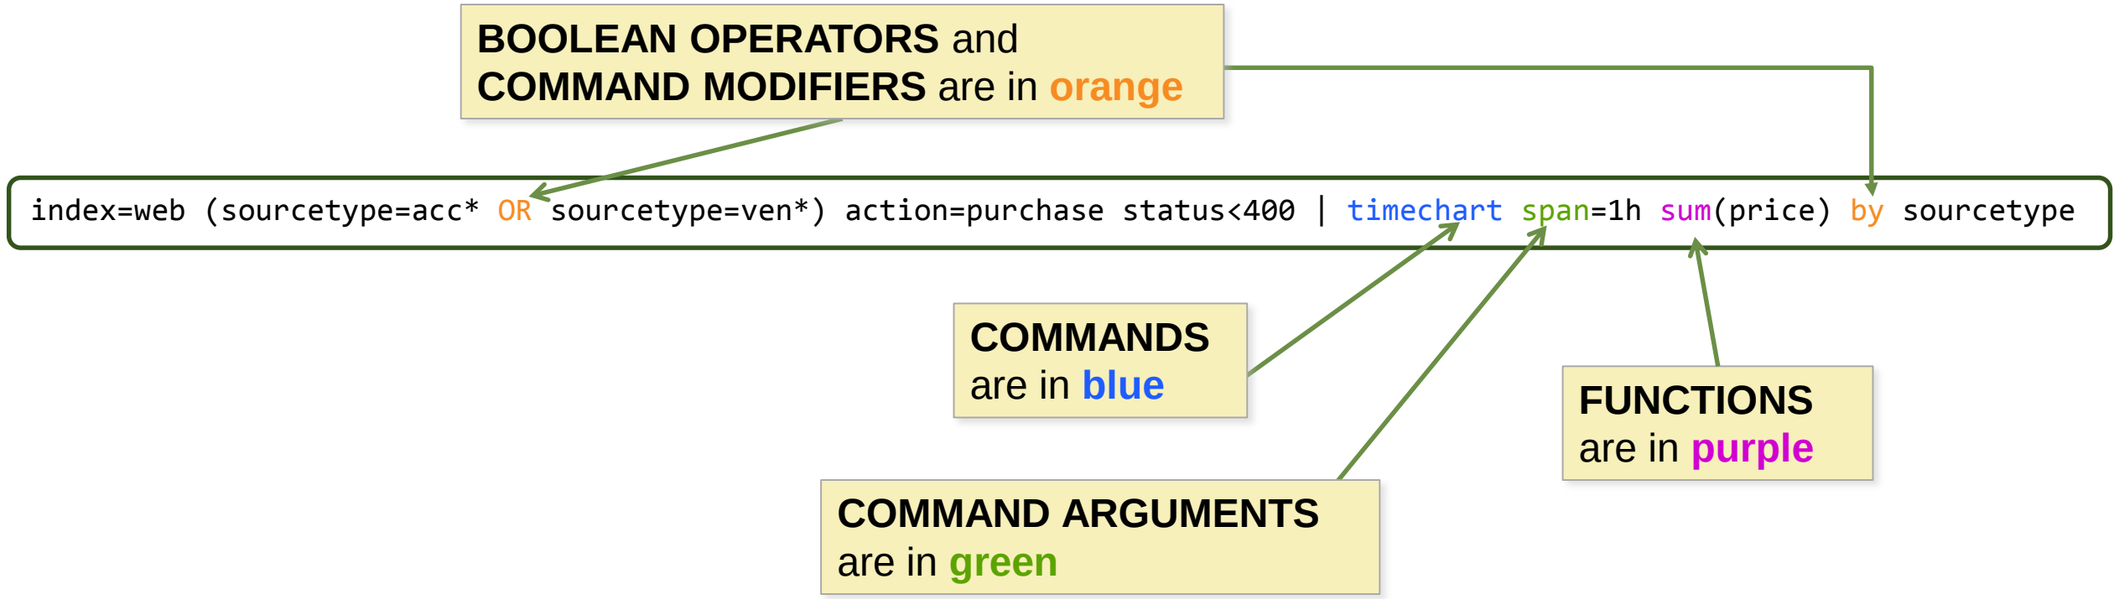

Syntax Coloring (cont.)

- You can turn off automatic syntax coloring
  1. Go to **Preferences** > **SPL Editor**
  2. Choose the Themes tab and select Black on White instead of the Light Theme default
  3. Click Apply
- You can also display colored text against a black background by selecting Dark Theme

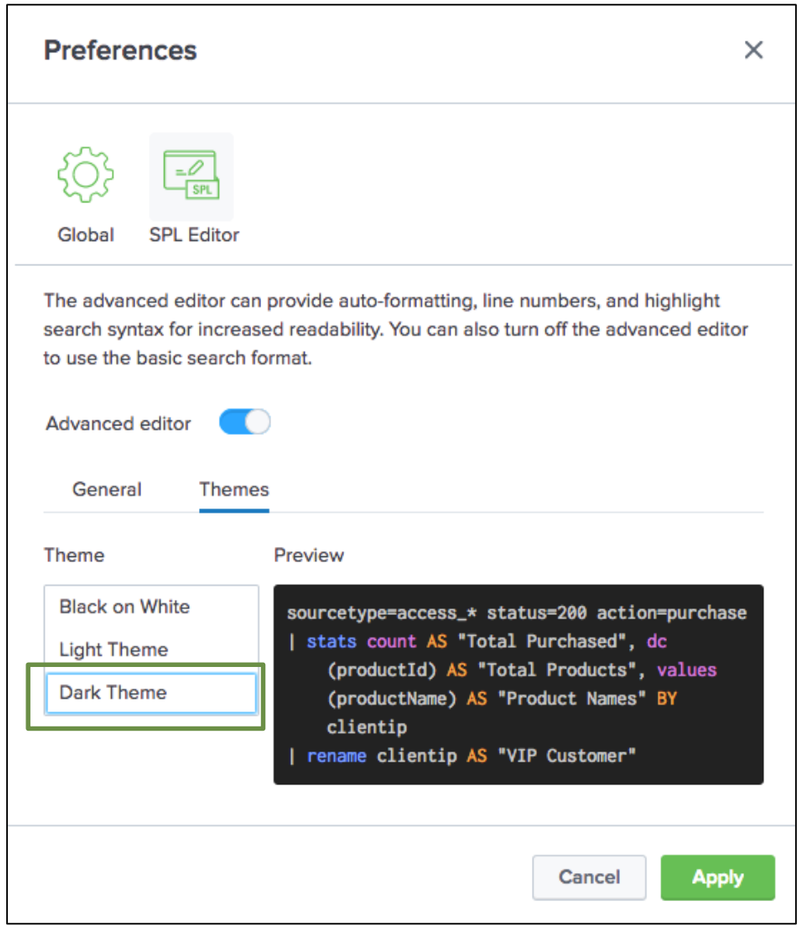


```
index=web (sourcetype=acc* OR sourcetype=ven*) action=purchase status<400 | timechart span=1h sum(price) by sourcetype
```

Creating a Table

- `table` command returns a table formed by only fields in the argument list
- Columns are displayed in the order given in the command
  - Column headers are field names
  - Each row represents an event
  - Each row contains field values for that event

Scenario

Display the `clientip`, `action`, `productId`, and `status` of customer interactions in the online store for the last 4 hours.

```
index=web sourcetype=access_combined
| table clientip, action, productId, status
```

clientip	action	productId	status
220.225.12.171	view	WC-SH-A02	200
220.225.12.171		WC-SH-G04	200
188.173.152.100		CU-PG-G06	200
188.173.152.100	purchase		200
188.173.152.100	purchase	WC-SH-T02	200
188.173.152.100	addtocart	WC-SH-T02	200
188.173.152.100		WC-SH-T02	200
188.173.152.100		WC-SH-A02	200

Renaming Fields in a Table

- To change the name of a field, use the `rename` command
- Useful for giving fields more meaningful names
- When including spaces or special characters in field names, use double straight quotes:

A `rename productId as ProductID`

B `rename action as "Customer Action"`

C `rename status as "HTTP Status"`

Scenario ?

Display the `clientip`, `action`, `productId`, and `status` of customer interactions in the online store for the last 4 hours.

```
index=web sourcetype=access_combined
| table clientip, action, productId, status
| rename productId as ProductID, A
  action as "Customer Action", B
  status as "HTTP Status" C
```

clientip	Customer Action (B)	ProductID (A)	HTTP Status (C)
12.130.60.4			200
27.102.11.11		SC-MG-G10	200
27.102.11.11		DC-SG-G02	200
27.102.11.11			200
27.102.11.11	view	FI-AG-G08	200
99.61.68.230	purchase		200
99.61.68.230	purchase	WC-SH-A01	200
99.61.68.230	addtocart	WC-SH-A01	200

Renaming Fields in a Table (cont.)

Once you rename a field, you can't access it with the original name

```
index=web sourcetype=access_combined
| table clientip, action, productId, status
| rename productId as ProductID,
  action as "Customer Action",
  status as "HTTP Status"
| table action, status
```

No results found.

```
index=web sourcetype=access_combined
| table clientip, action, productId, status
| rename productId as ProductID,
  action as "Customer Action",
  status as "HTTP Status"
| table "Customer Action", "HTTP Status"
```

Customer Action	HTTP Status
purchase	200
purchase	200
addtocart	200
	200
	200
view	200
view	200
	200

Using the `fields` Command

- Field extraction is one of the most costly parts of a search
- `fields` command allows you to include or exclude specified fields in your search or report
- To include, use `fields +` (default)
  - Occurs before field extraction
  - Improves performance
- To exclude, use `fields -`
  - Occurs after field extraction
  - No performance benefit
  - Exclude fields used in search to make the table/display easier to read

`fields` Command – Examples

Using command improves performance—only specified fields extracted

Scenario

Display network failures during the previous week.

```
index=security
sourcetype=linux_secure
(fail* OR invalid)
```

Returned **6,567** results by scanning **6,567** events in **1.425** seconds:

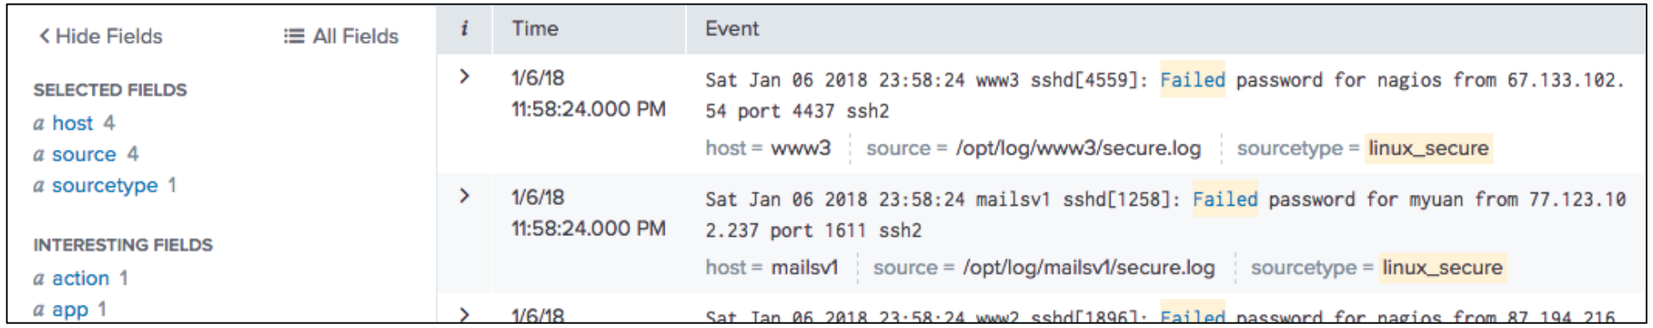

Scenario

Display network failures during the previous week. Retrieve only `user`, `app`, and `src_ip`.

```
index=security
sourcetype=linux_secure
(fail* OR invalid)
| fields user, app, src_ip
```

Returned **6,567** results by scanning **6,567** events in **0.753** seconds:

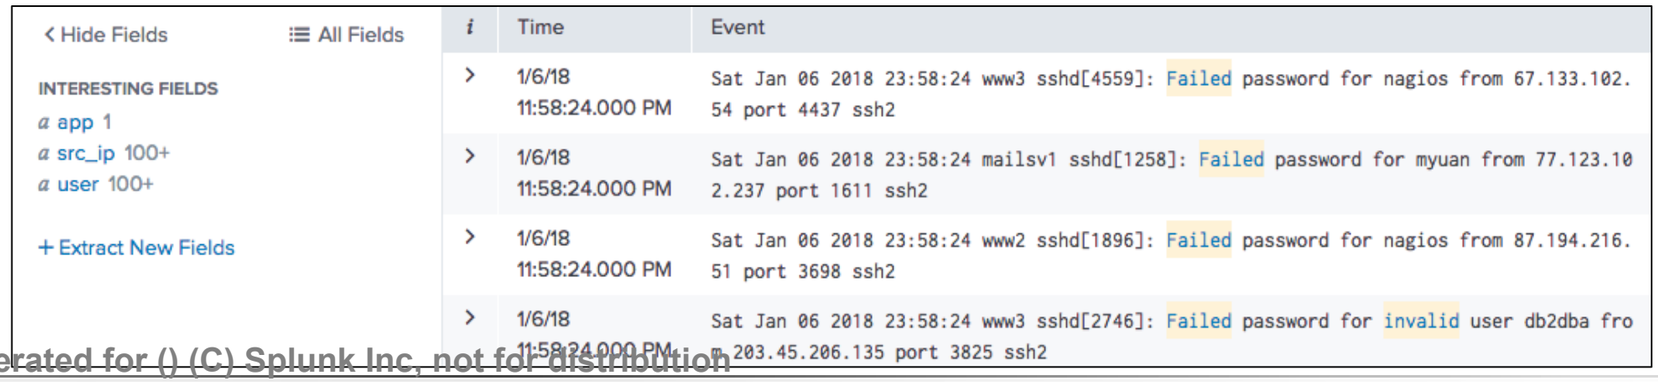

Using the dedup Command

Use dedup to remove duplicates from your results

```
index=sales sourcetype=vendor_sales Vendor=Bea* | table Vendor, VendorCity, VendorStateProvince, VendorCountry
```

Vendor	VendorCity	VendorStateProvince	VendorCountry
Beach Games	Miami	Florida	United States
Beads & Games	Ft. Riley	Kansas	United States
Beach Games	Miami	Florida	United States
Beantown Games	Boston	Massachusetts	United States
Beantown Games	Boston	Massachusetts	United States
Beach Games	Fort Lauderdale	Florida	United States

…| `dedup Vendor | table` …

Vendor	VendorCity	VendorStateProvince	VendorCountry
Beach Games	Miami	Florida	United States
Beads & Games	Ft. Riley	Kansas	United States
Beantown Games	Boston	Massachusetts	United States
Beauty Games	Butte	Montana	United States

…| `dedup Vendor, VendorCity | table` …

Vendor	VendorCity	VendorStateProvince	VendorCountry
Beach Games	Miami	Florida	United States
Beads & Games	Ft. Riley	Kansas	United States
Beach Games	Fort Lauderdale	Florida	United States
Beantown Games	Boston	Massachusetts	United States
Beauty Games	Butte	Montana	United States

Using the `sort` Command

- Use `sort` to order your results in `+ ascending` (default) or `– descending`
- To limit the returned results, use the `limit` option

```
... | sort limit=20 –categoryId, product_name
... | sort 20 count
```

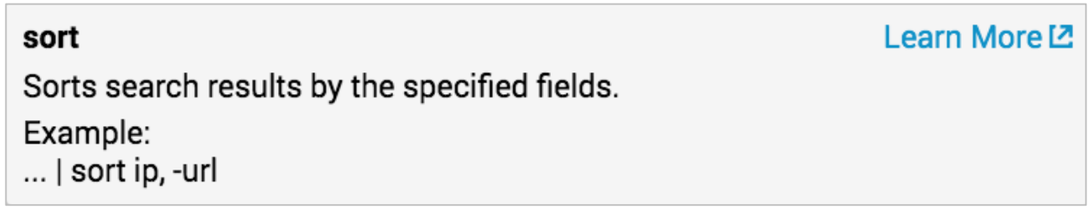

Using the `sort` Command (cont.)

`sort -/+<fieldname>` sign followed by fieldname sorts results in the sign's order

`sort -/+ <fieldname>` sign followed by **space** and then fieldname applies sort order to **all** following fields without a different explicit sort order

```
index=sales sourcetype=vendor_sales
Vendor=Bea*
| dedup Vendor, VendorCity
| table Vendor, VendorCity,
VendorStateProvince, VendorCountry
| sort -Vendor, VendorCity
```

Vendor	VendorCity	VendorStateProvince	VendorCountry
Beauty Games	Butte	Montana	United States
Beantown Games	Boston	Massachusetts	United States
Beads & Games	Ft. Riley	Kansas	United States
Beach Games	Fort Lauderdale	Florida	United States
Beach Games	Miami	Florida	United States

```
index=sales sourcetype=vendor_sales
Vendor=Bea*
| dedup Vendor, VendorCity
| table Vendor, VendorCity,
VendorStateProvince, VendorCountry
| sort - Vendor, VendorCity
```

Vendor	VendorCity	VendorStateProvince	VendorCountry
Beauty Games	Butte	Montana	United States
Beantown Games	Boston	Massachusetts	United States
Beads & Games	Ft. Riley	Kansas	United States
Beach Games	Miami	Florida	United States
Beach Games	Fort Lauderdale	Florida	United States

Useful References

- Search Reference: http://docs.splunk.com/Documentation/Splunk/latest/SearchReference
- Search Quick Reference: http://docs.splunk.com/Documentation/Splunk/latest/SearchReference/SplunkEnterpriseQuickReferenceGuide

Module 9
Transforming Commands

Getting Top Values

- The `top` command finds the most common values of a given field in the result set
- By default, output displays in table format

src_ip	count	percent
132.55.227.221	105	27.131783
214.156.206.45	9	2.325581
118.6.85.68	9	2.325581
87.194.216.51	8	2.067183
108.65.113.83	6	1.550388
89.11.192.18	5	1.291990
211.166.11.101	5	1.291990
195.69.160.22	5	1.291990
81.18.148.190	4	1.033592
223.205.219.67	4	1.033592

Scenario

Determine which IP addresses generated the most attacks in the last 60 minutes,

```
index=security sourcetype=linux_secure
(fail* OR invalid)
| top src_ip
```


top Command

- By default, returns top 10 results
- Automatically returns count and percent columns
- Common constraints: `limit` `countfield` `showperc`
- `top` command with `limit=20` is automatically added to your search string when you click Top values in a field window

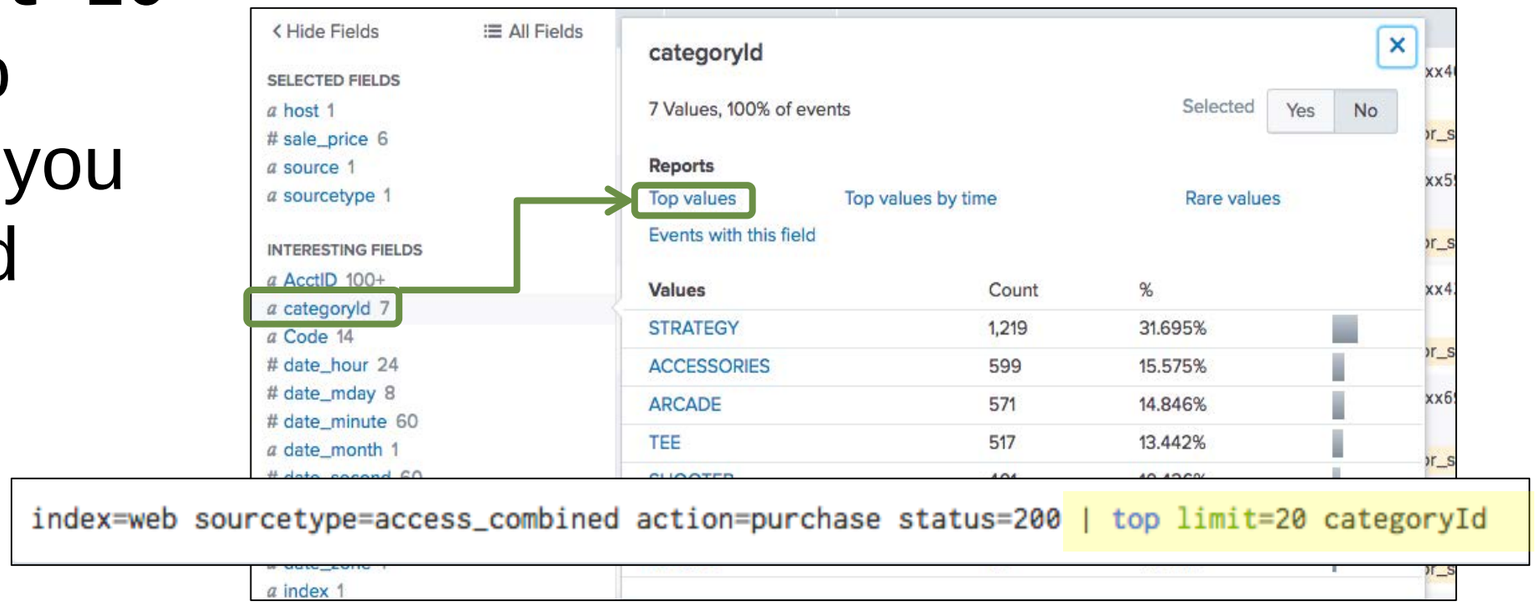

Note

Creating top values reports from field windows was discussed in Module 4.

top Command – Single Field

- Control # of results displayed using `limit`
- `limit=#` returns this number of results
- `limit=0` returns unlimited results

```
index=security sourcetype=linux_secure
(fail* OR invalid)
| top limit=5 src_ip
```

Scenario

During the last hour, display the top 5 IPs that generated the most attacks.

src_ip	count	percent
132.55.227.221	106	27.390181
118.6.85.68	9	2.325581
87.194.216.51	8	2.067183
214.156.206.45	8	2.067183
108.65.113.83	7	1.808786

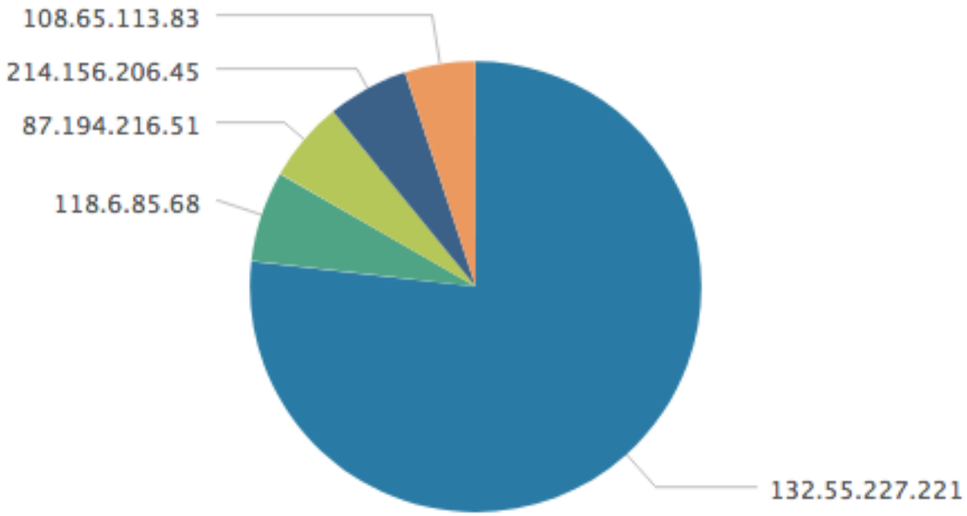

top Command – Multiple Fields

- If the showperc is not included – or it is included and set to t – a percent column is displayed
- If showperc=f, then a percent column is NOT displayed

Scenario

Display the top 3 common values for users and web categories browsed during the last 24 hours.

```
index=network sourcetype=cisco_wsa_squid
| top cs_username x_webcat_code_full limit=3
```

A

B

cs_username	x_webcat_code_full	count	percent
kjoslin@buttercupgames.com	Sports and Recreation	81	6.617647
kperna@buttercupgames.com	Shopping	61	4.983660
apreusig@buttercupgames.com	Arts and Entertainment	58	4.738562

top Command – Single Field with by Clause

Scenario 1

Display the top 3 web categories browsed by each user during the last 24 hours.

```
index=network sourcetype=cisco_wsa_squid
| top x_webcat_code_full by cs_username limit=3
```

B A

A B

cs_username	x_webcat_code_full	count	percent
acurry@buttercupgames.com	Uncategorized URLs	5	35.714286
acurry@buttercupgames.com	Shopping	4	28.571429
acurry@buttercupgames.com	Health and Nutrition	3	21.428571
adombrowski@buttercupgames.com	Uncategorized URLs	1	50.000000
adombrowski@buttercupgames.com	Computers and Internet	1	50.000000
apreusig@buttercupgames.com	Arts and Entertainment	58	95.081967

Scenario 2

Display the top 3 users for each web category during the last 24 hours.

```
index=network sourcetype=cisco_wsa_squid
| top cs_username by x_webcat_code_full limit=3
```

C D

D C

x_webcat_code_full	cs_username	count	percent
Advertisements	gnooteboom@buttercupgames.com	5	41.666667
Advertisements	iking@buttercupgames.com	3	25.000000
Advertisements	tzielinski@buttercupgames.com	2	16.666667
Arts and Entertainment	fbryan@buttercupgames.com	61	46.923077
Arts and Entertainment	moh@buttercupgames.com	6	4.615385
Arts and Entertainment	pbunch@buttercupgames.com	4	3.076923

top Command – Renaming countfield Display

- By default, the display name of the countfield is count
- countfield=*string* renames the field for display purposes

Scenario

Display the top 3 user/web categories combinations during the last 24 hours. Rename the count field and show count, but not the percentage.

```
index=network sourcetype=cisco_wsa_squid
| top cs_username x_webcat_code_full limit=3 A
  countfield="Total Viewed" B showperc=f
```

cs_username	x_webcat_code_full	Total Viewed
kjoslin@buttercupgames.com	Sports and Recreation	81
kperna@buttercupgames.com	Shopping	61
apreusig@buttercupgames.com	Arts and Entertainment	58

Note

A Boolean can be t/f, true/false, as well as 1/0.

rare Command

- The rare command returns the least common field values of a given field in the results
- Options are identical to the top command

Scenario

Identify which product is the least sold by Buttercup Games vendors over the last 60 minutes.

```
index=sales sourcetype=vendor_sales
  | rare product_name showperc=f limit=1
```

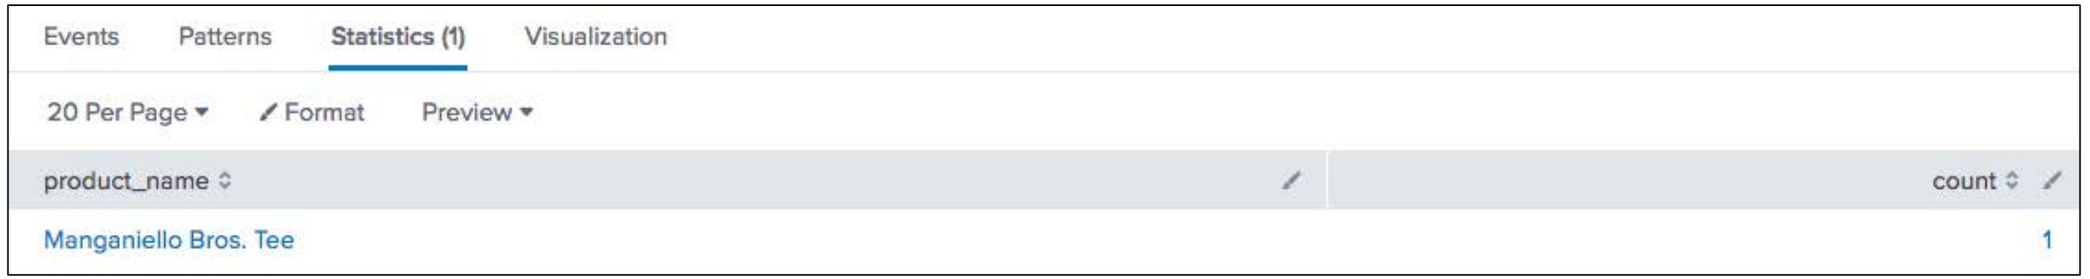

`stats` Command

- `stats` enables you to calculate statistics on data that matches your search criteria
- Common functions include:
  - `count` – returns the number of events that match the search criteria
  - `distinct_count, dc` – returns a count of unique values for a given field
  - `sum` – returns a sum of numeric values
  - `avg` – returns an average of numeric values
  - `list` – lists all values of a given field
  - `values` – lists unique values of a given field

Note

To view all of the functions for `stats`, please see:
http://docs.splunk.com/Documentation/Splunk/latest/SearchReference/CommonStatsFunctions

stats Command – count

- count returns the number of matching events based on the current search criteria
- Use the as clause to rename the count field

Scenario

Count the invalid or failed login attempts during the last 60 minutes.

```
index=security sourcetype=linux_secure
(invalid OR failed)
| stats count
```

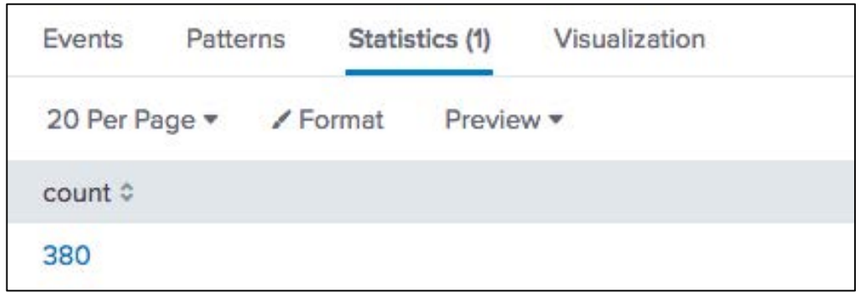


```
index=security sourcetype=linux_secure
(invalid OR failed)
| stats count as "Potential Issues"
```

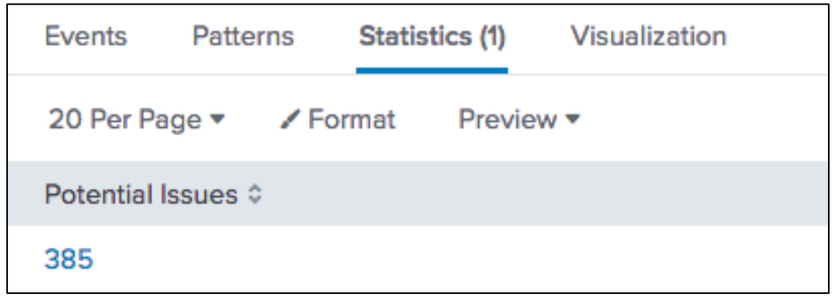

`stats` Command – `count(field)`

Adding a *`field`* as an argument to the `count` function returns the number of events where a value is present for the specified field

Scenario

Count the number of events during the last 15 minutes that contain a vendor action field. Also count the total events.

```
index=security sourcetype=linux_secure
| stats count(vendor_action) as ActionEvents, A
  count as TotalEvents B
```

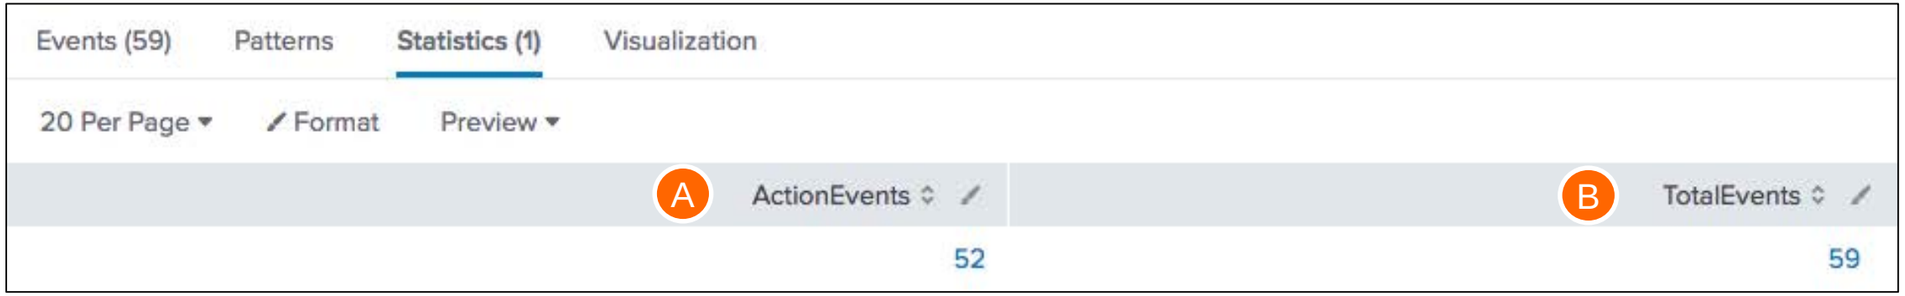

stats Command – by *fields*

Scenario

Count the number of events by user, app, and vendor action during the last 15 minutes.

```
index=security sourcetype=linux_secure
| stats count by user, app, vendor_action
```

- by clause returns a count for each value of a named field or set of fields
- Can use any number of fields in the by *field* list

user	app	vendor_action	count
administrator	sshd	Failed	1
agushto	sshd	Failed	1
backup	sshd	Failed	1
bin	sshd	Failed	2
brian	sshd	Failed	1
db	sshd	Failed	2
db2fenc1	sshd	Failed	2
db4	sshd	Failed	1
dba	sshd	Failed	1
demon	sshd	Failed	1

`stats` Command – `distinct_count(field)`

- `distinct_count()` or `dc()` provides a count of how many unique values there are for a given field in the result set
- This example counts how many unique values for `s_hostname`

Scenario

How many unique websites have employees visited in the last 4 hours?

```
index=network sourcetype=cisco_wsa_squid
 | stats dc(s_hostname) as "Websites visited:"
```

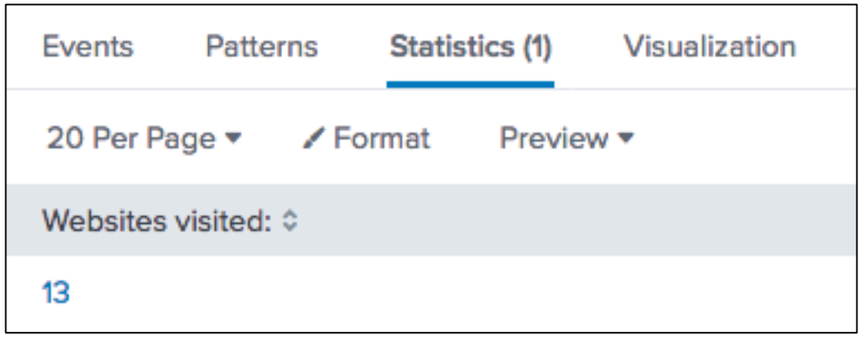

stats Command – sum(*field*)

Scenario

How much bandwidth did employees consume at each website during the past week?

```
index=network sourcetype=cisco_wsa_squid
| stats sum(sc_bytes) as Bandwidth by s_hostname
| sort -Bandwidth
```

A B C

For fields with a numeric value, you can sum the actual values of that field

Events | Patterns | **Statistics (520)** | Visualization

20 Per Page ▾ | Format | Preview ▾

s_hostname	Bandwidth
www.archerytalk.com	6330733
www.infoblox.com	2818279
www.bjc-aces.com	1849884
www.bradblog.com	1489356
www.ncsl.org	1425778

stats Command – sum(*field*) – (cont.)

Scenario

Report the number of retail units sold and sales revenue for each product during the previous week.

```
index=sales sourcetype=vendor_sales
| stats A count(price) as "Units Sold"
B sum(price) as "Total Sales"   by product_name C
| sort -"Total Sales" D
```

- A A single `stats` command
- B can have multiple functions
- C The by clause is applied to both functions
- D `sort` Total Sales in descending order

product_name	Units Sold (A)	Total Sales (B)
Dream Crusher	78	3119.22
World of Cheese	78	1949.22
Manganiello Bros.	45	1799.55
SIM Cubicle	72	1439.28
Final Sequel	55	1374.45
Mediocre Kingdoms	50	1249.50
Orvil the Wolverine	30	1199.70
Benign Space Debris	31	774.69
Curling 2014	28	559.72
World of Cheese Tee	47	469.53

C D

stats Command – avg(*field*)

- The avg function provides the average numeric value for the given numeric field
- An event is not considered in the calculation if it:
  - Does not have the field
  - Has an invalid value for the field

Scenario

What is the average bandwidth used for each website usage type?

```
index=network sourcetype=cisco_wsa_squid
| stats avg(sc_bytes) as "Average Bytes" A
  by usage B
```

usage	Average Bytes
Borderline	18870.117341640707
Business	12957.355321020228
Personal	12915.358326596604
Unknown	12747.842960288808
Violation	8831.08888888889

B A

`stats` Command – `list(field)`

- `list` function lists all field values for a given field
- This example lists the websites visited by each employee
  - Security logs generate an event for each network request
    - The same hostname appears multiple times
  - To return a list of “unique” field values, use the `values` function

Scenario

Which websites has each employee accessed during the last 60 minutes?

```
index=network sourcetype=cisco_wsa_squid
| stats list(s_hostname) as "Websites visited:"
  by cs_username
```

cs_username	Websites visited:
	www.filmschoolrejects.com
apucci@buttercupgames.com	www.zimbio.com
arangel@buttercupgames.com	www.cnet.com
basselin@buttercupgames.com	www.americangangster.net www.americangangster.net www.americangangster.net www.blossomfloristla.com 84654321.cn
bgenin@buttercupgames.com	www.ambrosiasw.com www.ambrosiasw.com www.ambrosiasw.com

stats Command – values(*field*)

Scenario

Display by IP address the names of users who have failed access attempts in the last 60 minutes.

```
index=security sourcetype=linux_secure fail*
| stats values(user) as "User Names",
  count(user) as Attempts by src_ip
```

values function lists unique values for the specified field

src_ip	User Names	Attempts
107.3.146.207	administrator pat uni	3
109.169.32.135	harrypotter irc	2
110.138.30.229	daemon	1
111.161.27.20	administrator daemon games	3
112.111.162.4	noone	1

Formatting `stats` Tables

- Tables created with `stats` commands can be formatted
- Color code data in each column, based on rules you define
- Add number formatting (e.g. currency symbols, thousands separators)
- Can also format data on a per-column basis by clicking the icon above that column

product_name	Units Sold	Total Sales
Benign Space Debris	16	399.84
Curling 2014	35	699.65
Dream Crusher	23	919.77
Final Sequel	27	674.73

Formatting `stats` Tables – Examples

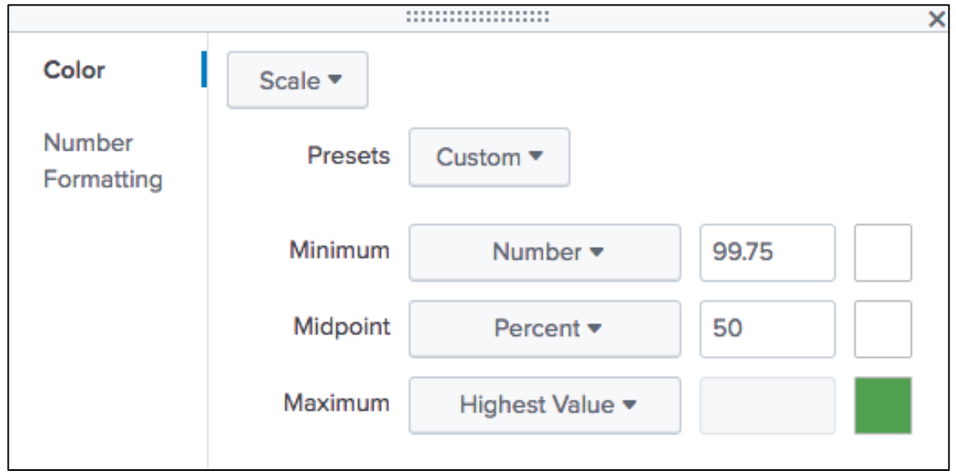

product_name	Units Sold	Total Sales
Benign Space Debris	16	399.84
Curling 2014	35	699.65
Dream Crusher	23	919.77
Final Sequel	27	674.73
Fire Resistance Suit of Provolone	17	67.83
Holy Blade of Gouda	44	263.56

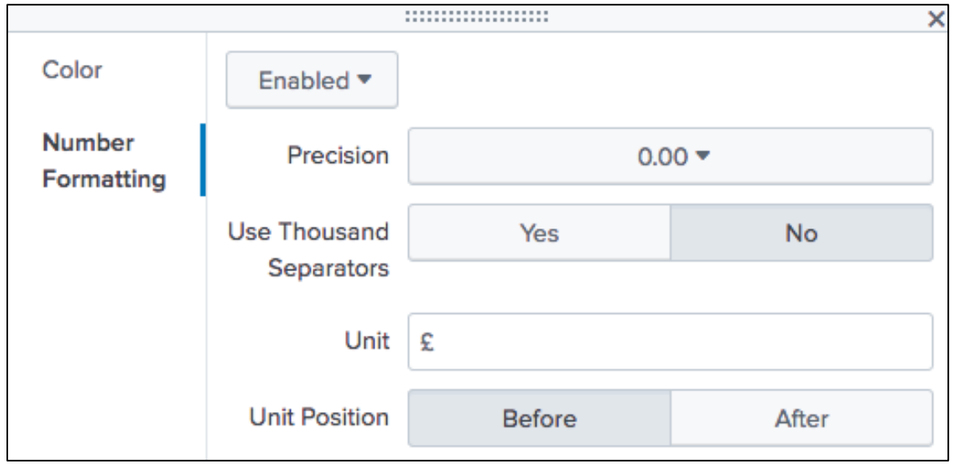

Total Sales
£ 399.84
£ 699.65
£ 919.77
£ 674.73
£ 67.83
£ 263.56
£ 999.75
£ 359.64

Module 10:
Creating Reports and Dashboards

What Are Reports?

- Reports are saved searches
- Reports can show events, statistics (tables), or visualizations (charts)
- Running a report returns fresh results each time you run it
- Statistics and visualizations allow you to drill down by default to see the underlying events
- Reports can be shared and added to dashboards

Smart Naming

- Before you begin using Splunk on the job, define a naming convention so you can always find your reports and tell them apart
- For example, you can create something simple like this:
  - <group>_<object>_<description>
    - **group**: the name of the group or department using the knowledge object such as sales, IT, finance, etc.
    - **object**: report, dashboard, macro, etc.
    - **description**: WeeklySales, FailedLogins, etc.
  - Using this example, a quarterly sales report can be identified as:
    - Sales_Report_QuarterlySalesRevenue

Note

If you set up naming conventions early in your implementation, you can avoid some of the more challenging object naming issues. The example is a suggestion. The details are found in the Splunk product documentation: http://docs.splunk.com/Documentation/Splunk/latest/Knowledge/Developnamingconventionsforknowledgeobjecttitles

Creating a Report from a Search

1. Run a search
2. Select **Save As**
3. Select **Report**

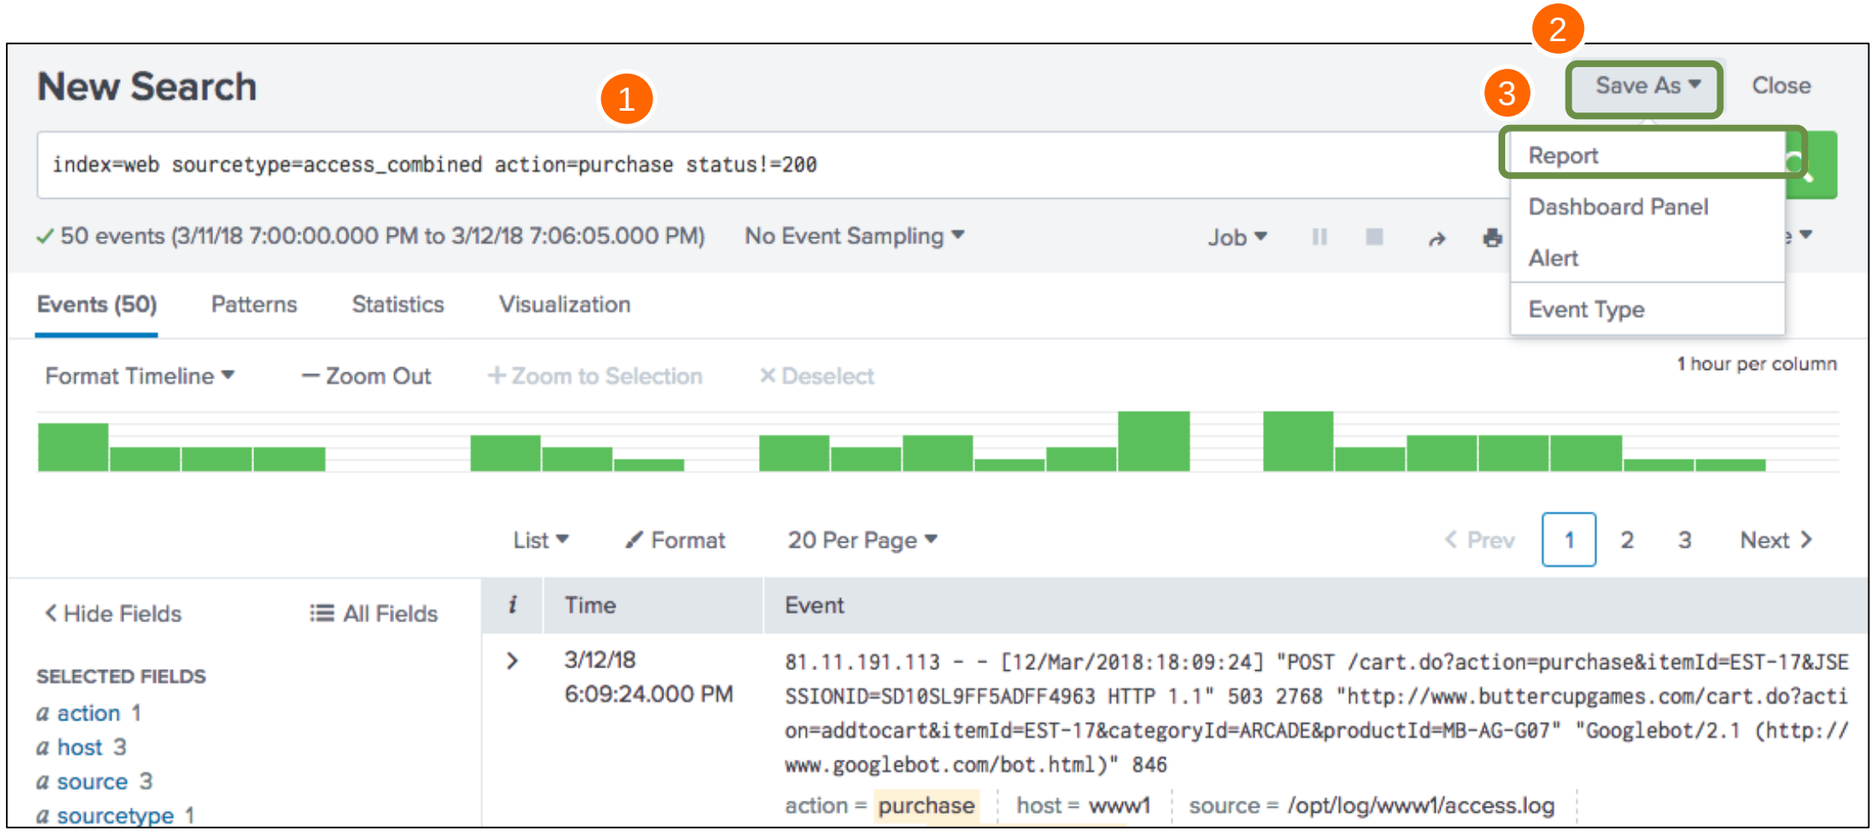

Creating a Report from a Search (cont.)

- A Give the report a meaningful title (required)
- B Specify a description (optional)
- C Select whether to include or not to include a time range picker
  - The report is saved with the time range that was selected when it was created
  - Adding a time range picker allows you to adjust the time range of the report when you run it

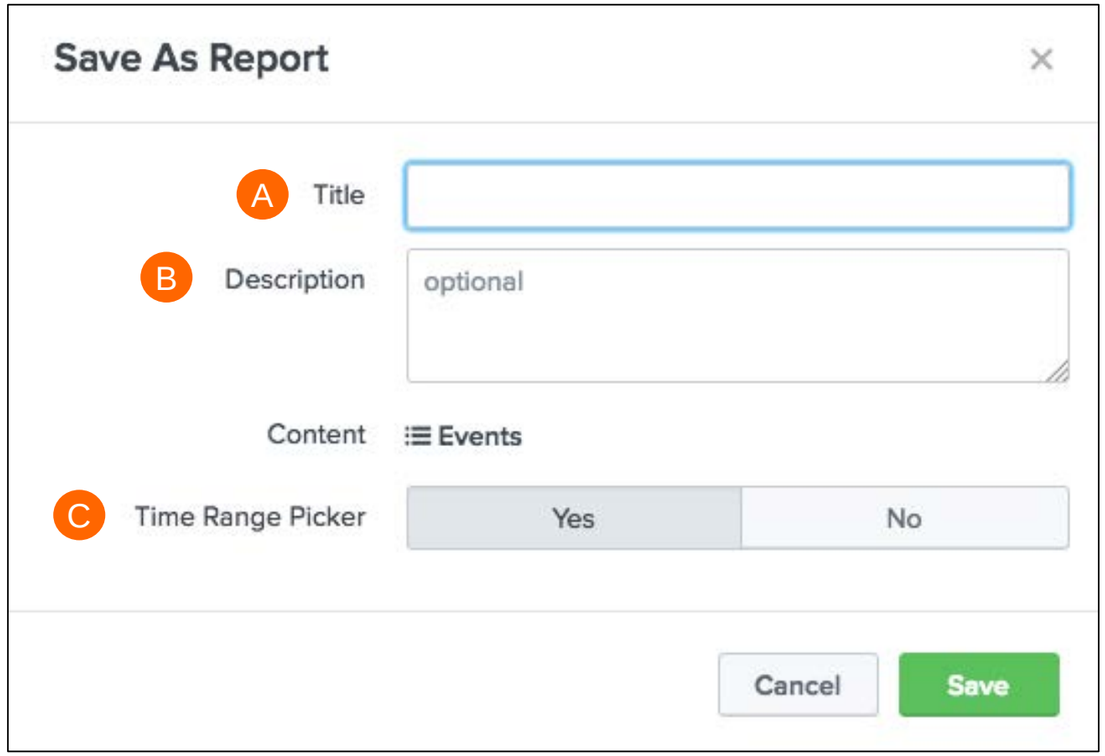

Creating a Report from a Search (cont.)

- You can change Additional Settings, as well as use the dialog buttons:
  - Click **Continue Editing** to make changes to your report
  - Click **Add to Dashboard** to add your report to a dashboard
  - Click **View** to display your report or run it again

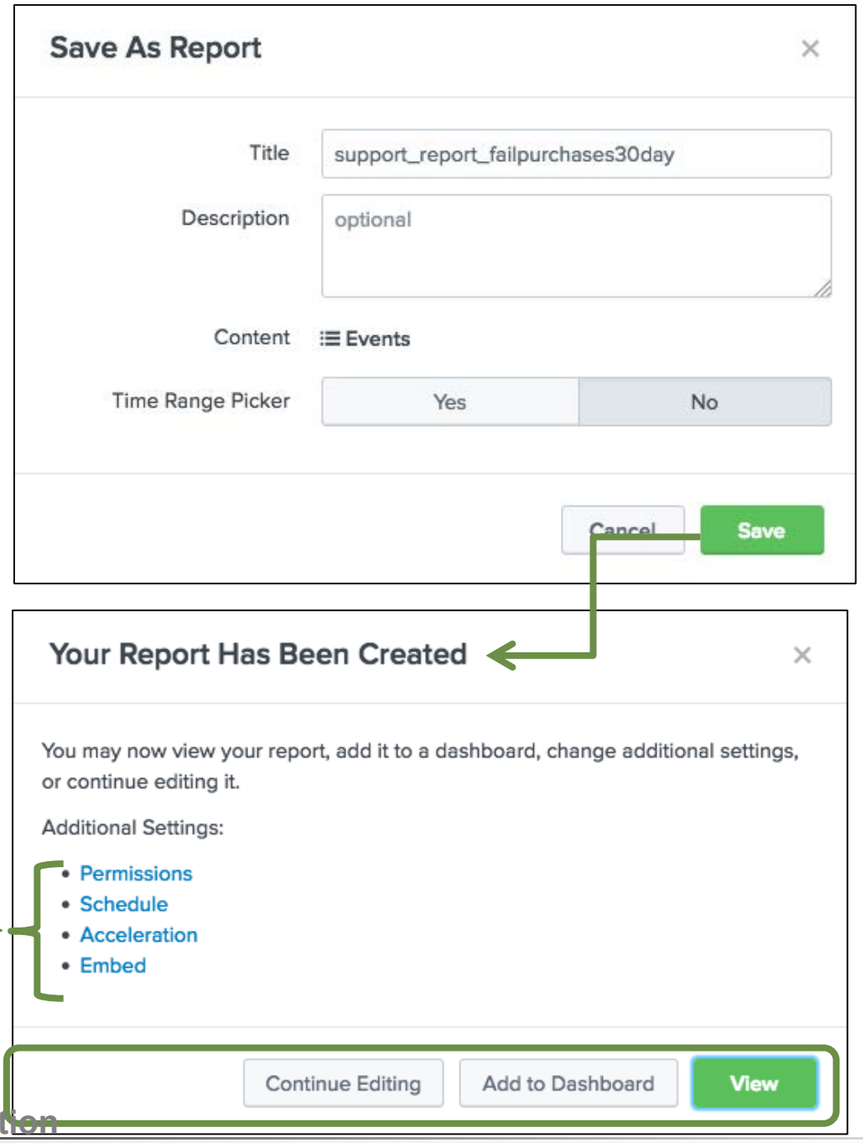

Running Reports

- Click **Reports**, then click the report title to run it
  - The report runs using the time range that was specified when it was saved
- Use the time range picker to change the time range of the report (if available)

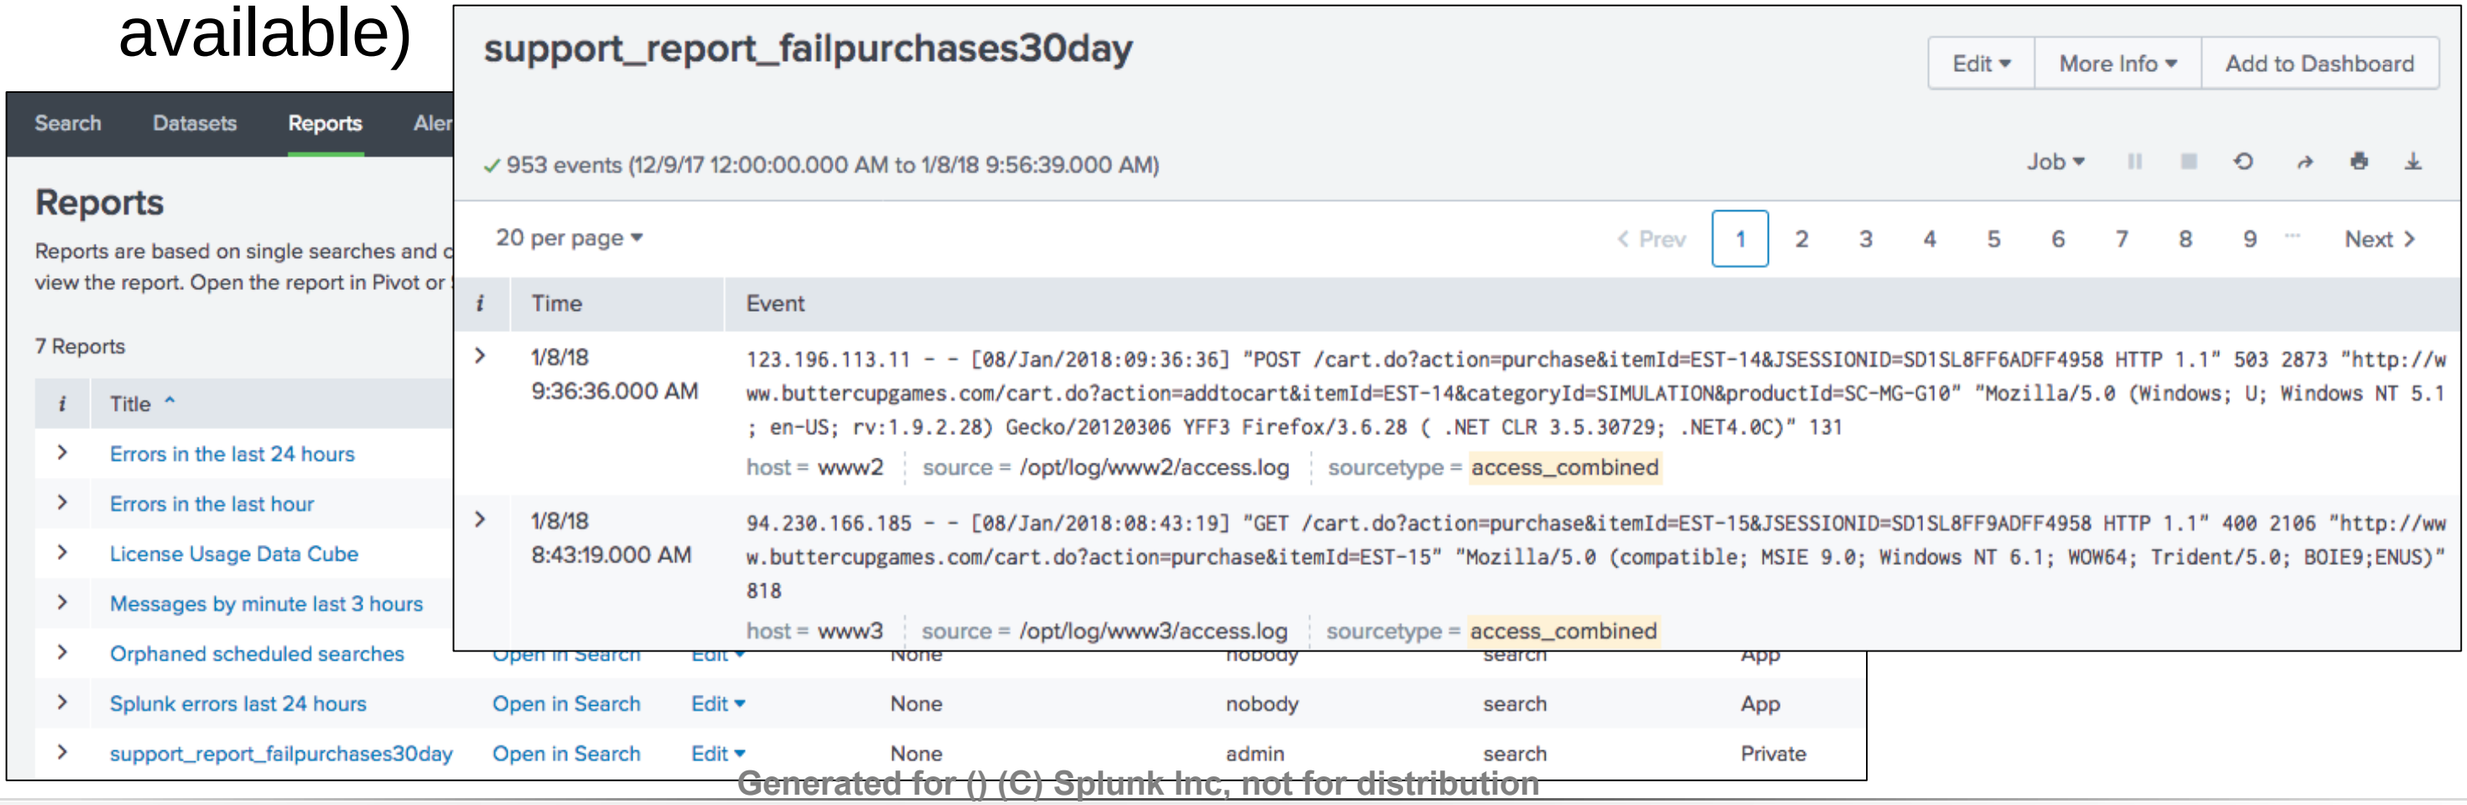

Editing Reports

- To edit a report's underlying search, select **Edit > Open in Search**
  - You can then edit and re-save, not save, or save-as a new report
- You can also edit the description, permissions, schedule, and acceleration
- Additionally, you can clone or delete the report

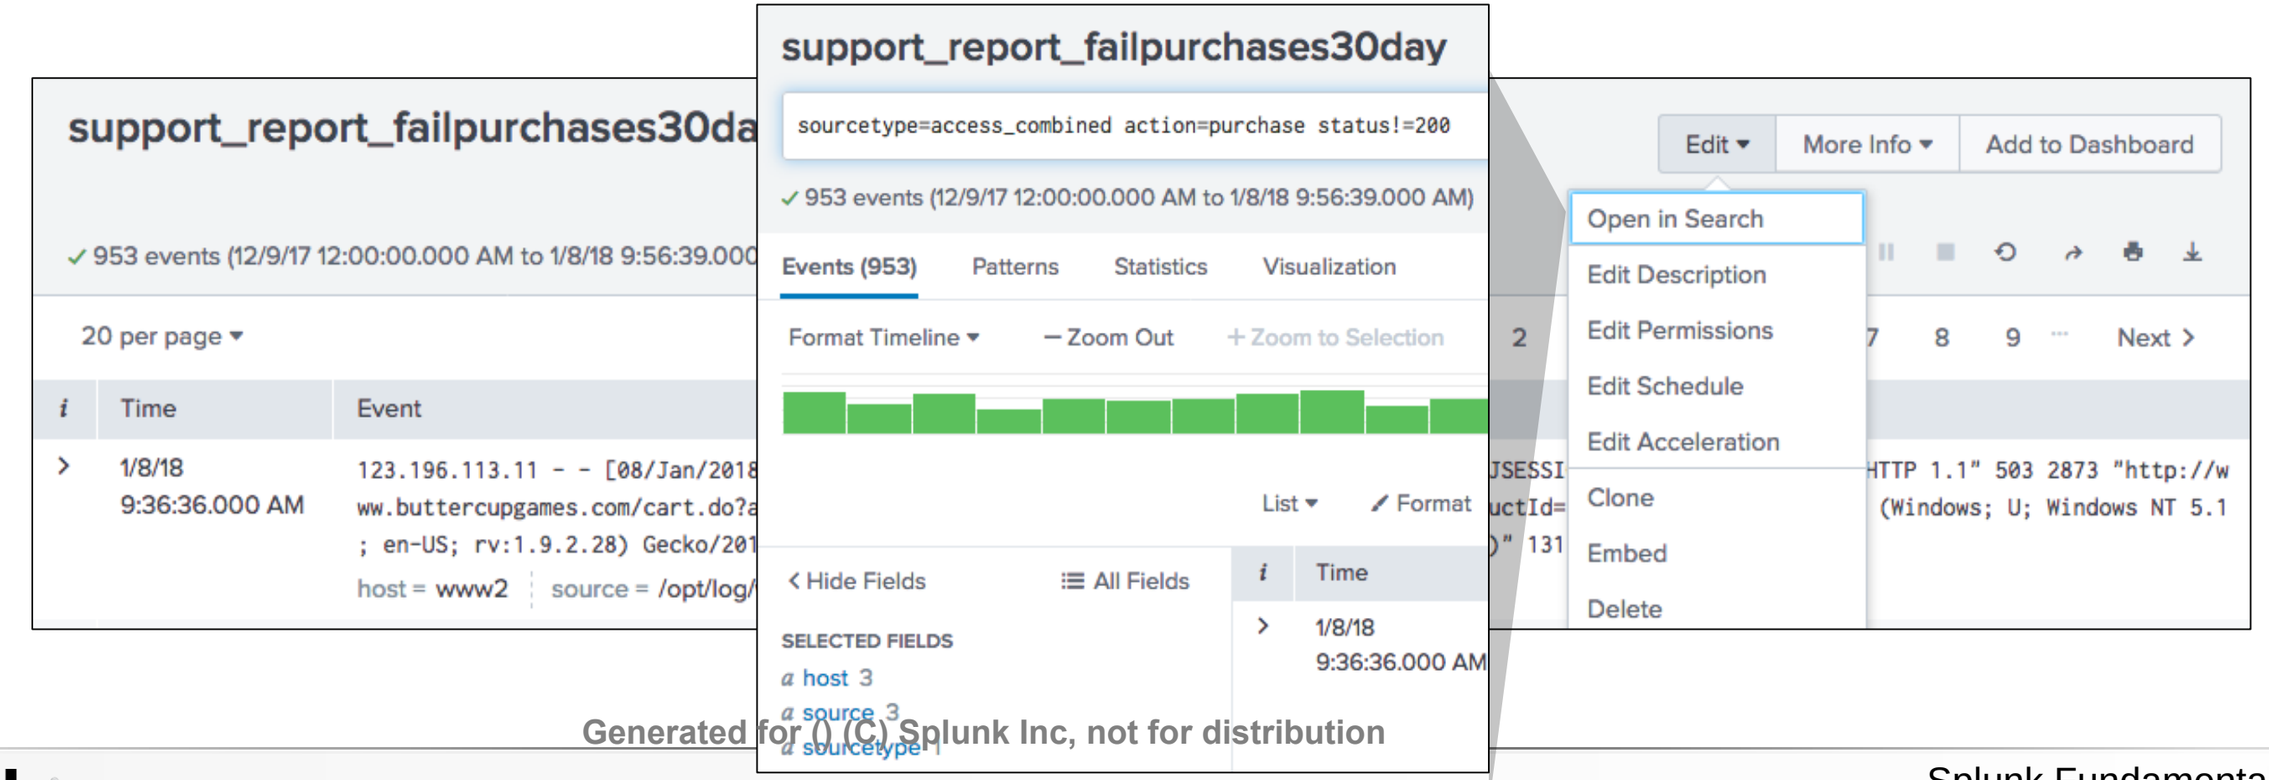

Creating Tables and Visualizations

Three main methods to create tables and visualizations in Splunk are:

1. Select a field from the fields sidebar and choose a report to run
2. Use the Pivot interface
   - Start with a dataset

     or

   - Start with Instant Pivot
   - See Module 11 in this presentation for more information about Pivot
3. Use the Splunk search language transforming commands in the Search bar

Viewing Tables and Visualizations

- Statistical reports leverage Splunk's built-in visualizations or table format
- These views give you insights into your organization’s data

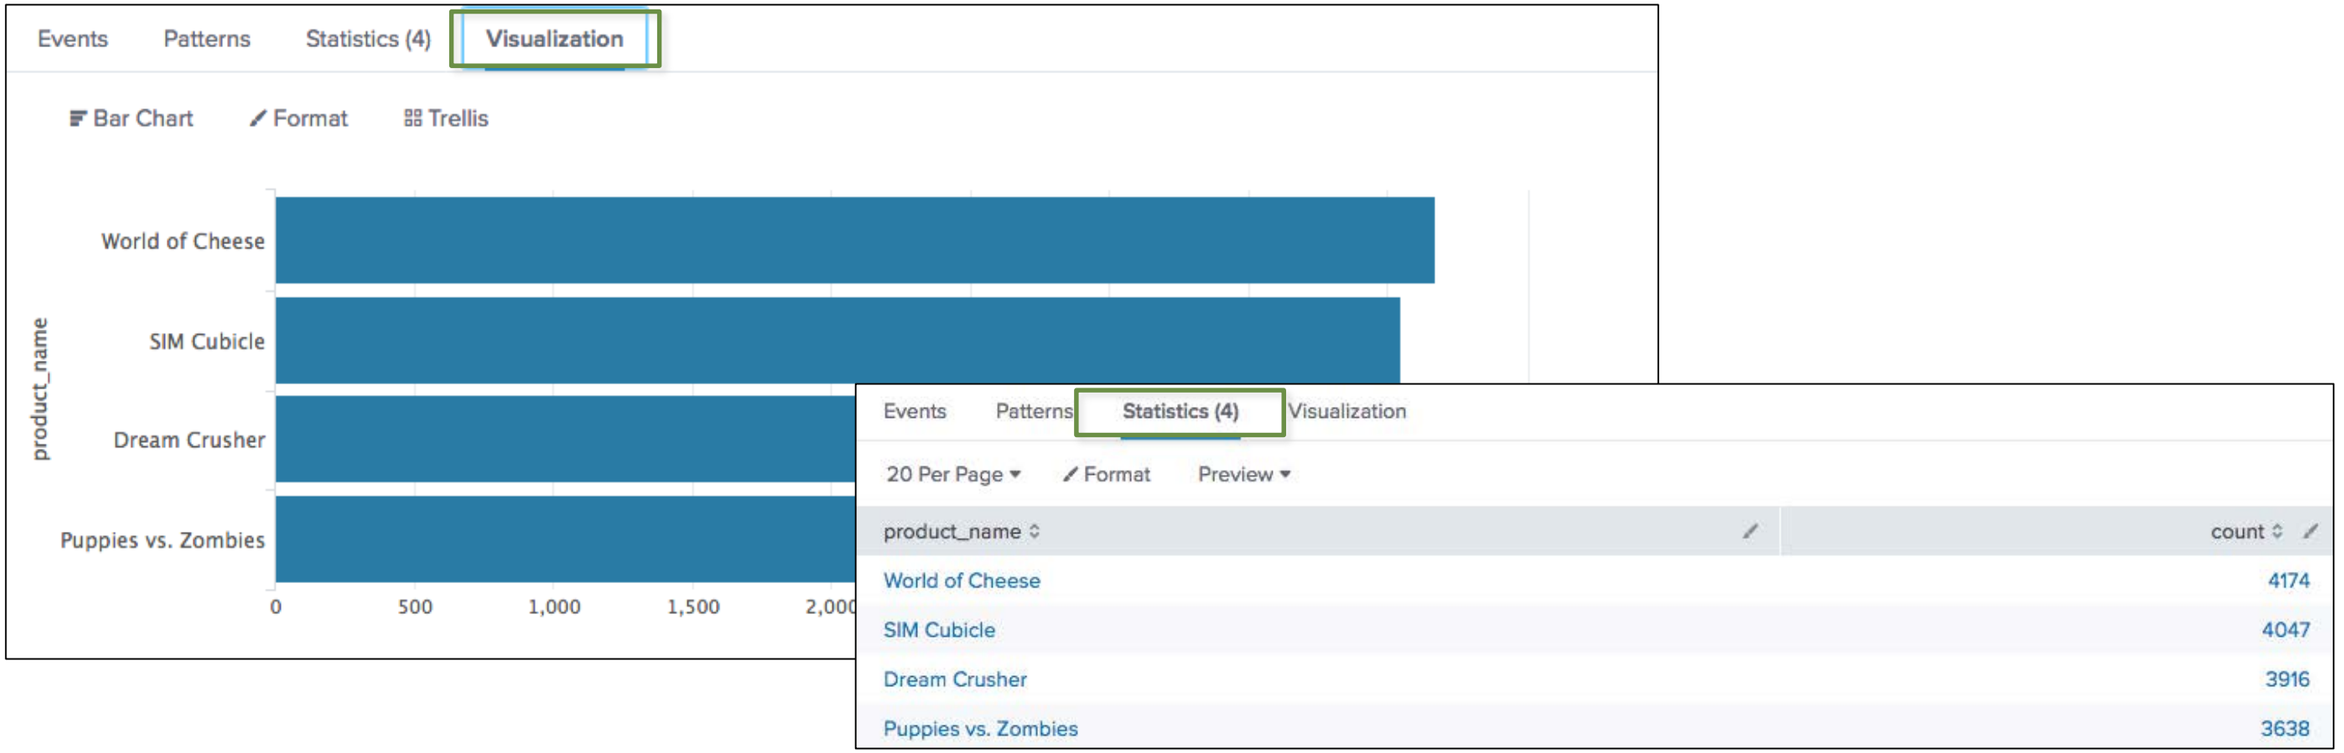

Creating Reports From the Field Window

- Numeric fields: choose from six report types with mathematical functions, such as average, maximum value, and minimum value
- This example generates a report that shows the average over time
  - This is known as a **timechart**

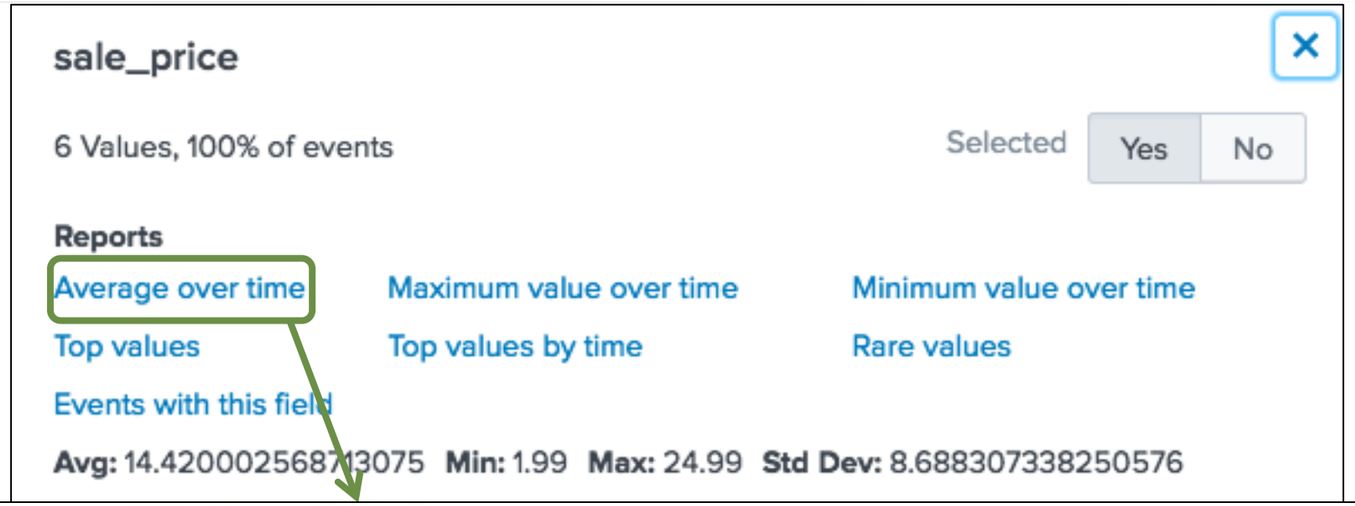

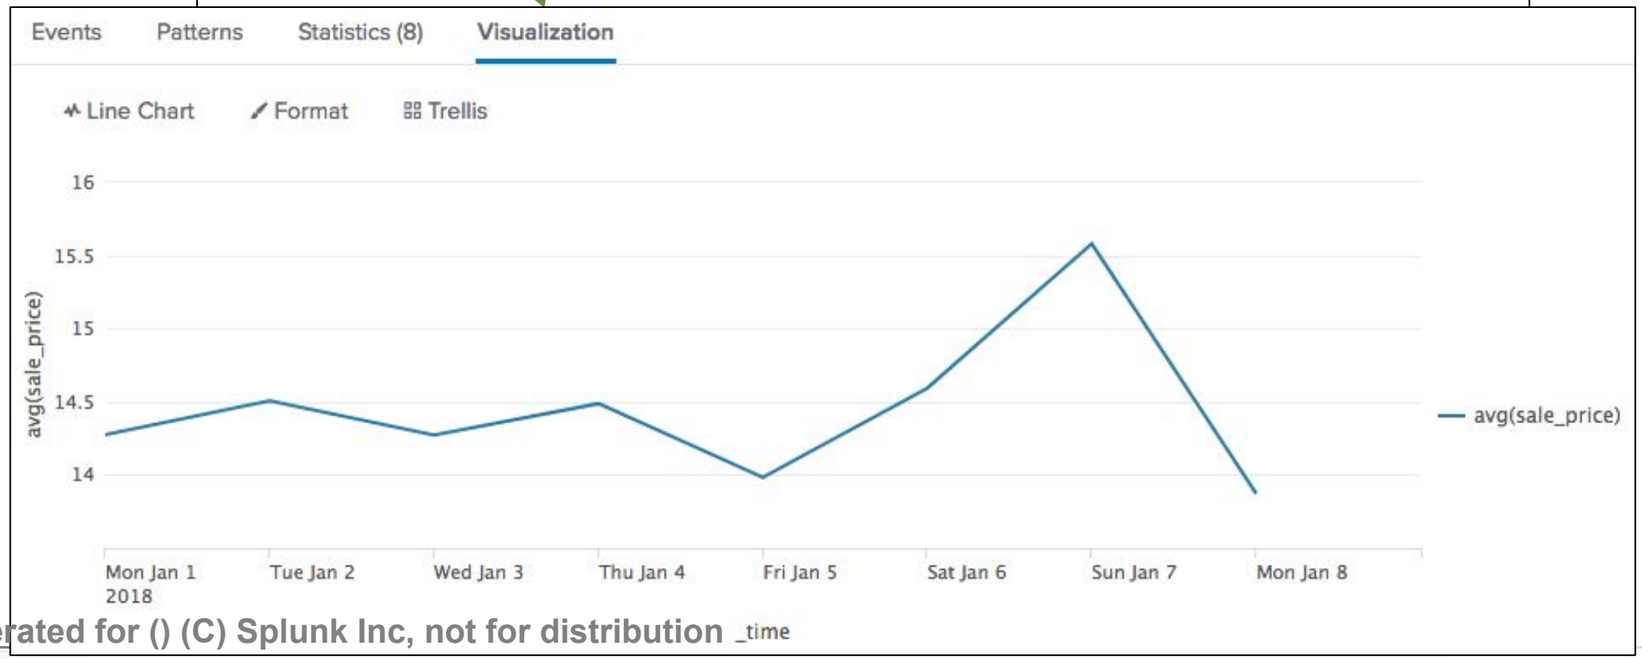

Creating a Top Values Report

- For alphanumeric character fields, there are only 3 available reports
- In this example, you want a report that shows the top **categories** purchased
  1. Run basic search: `sourcetype=access_combined status=200 action=purchase`
  2. Click the **categoryId** field
  3. Click **Top values**

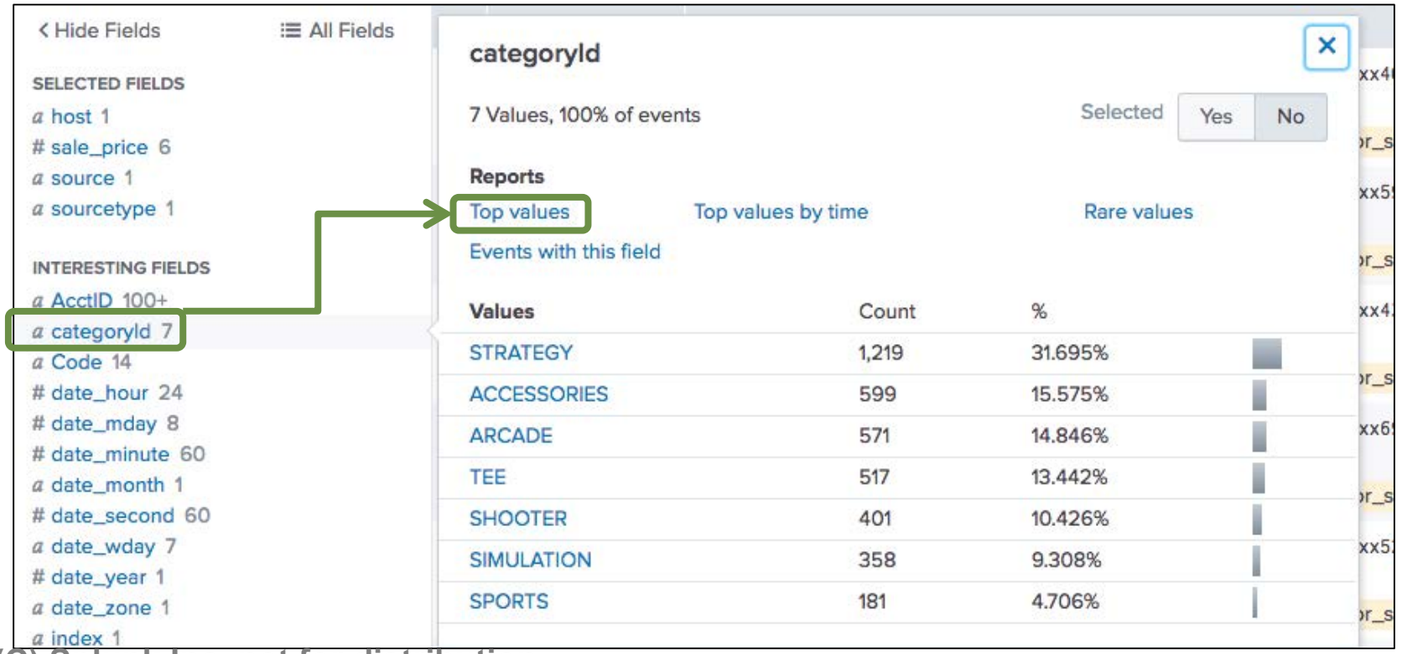

Creating a Top Values Report (cont.)

4. The `top` command with `limit=20` is added to the search string
5. A bar chart is returned on the Visualizations tab, displaying the top categories purchased

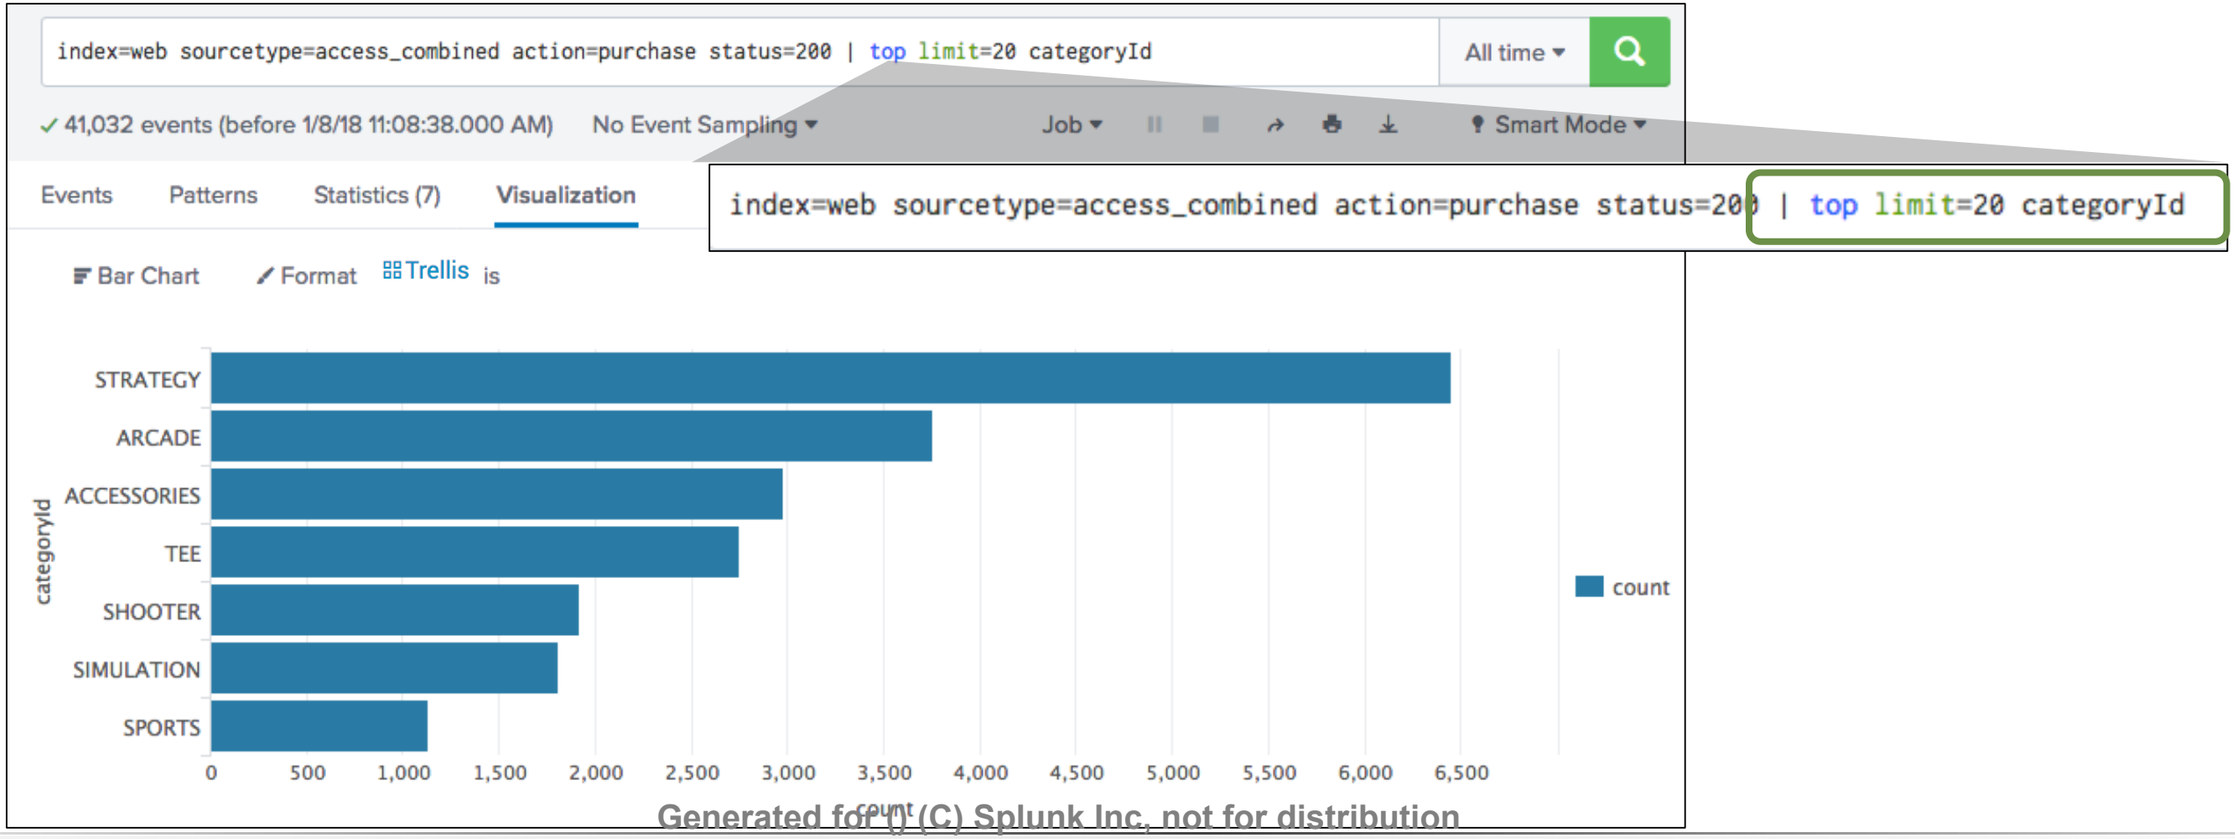

Changing the Visualization

- Select a visualization from the visualization type dropdown menu
- In this example, the column chart is changed to a pie chart

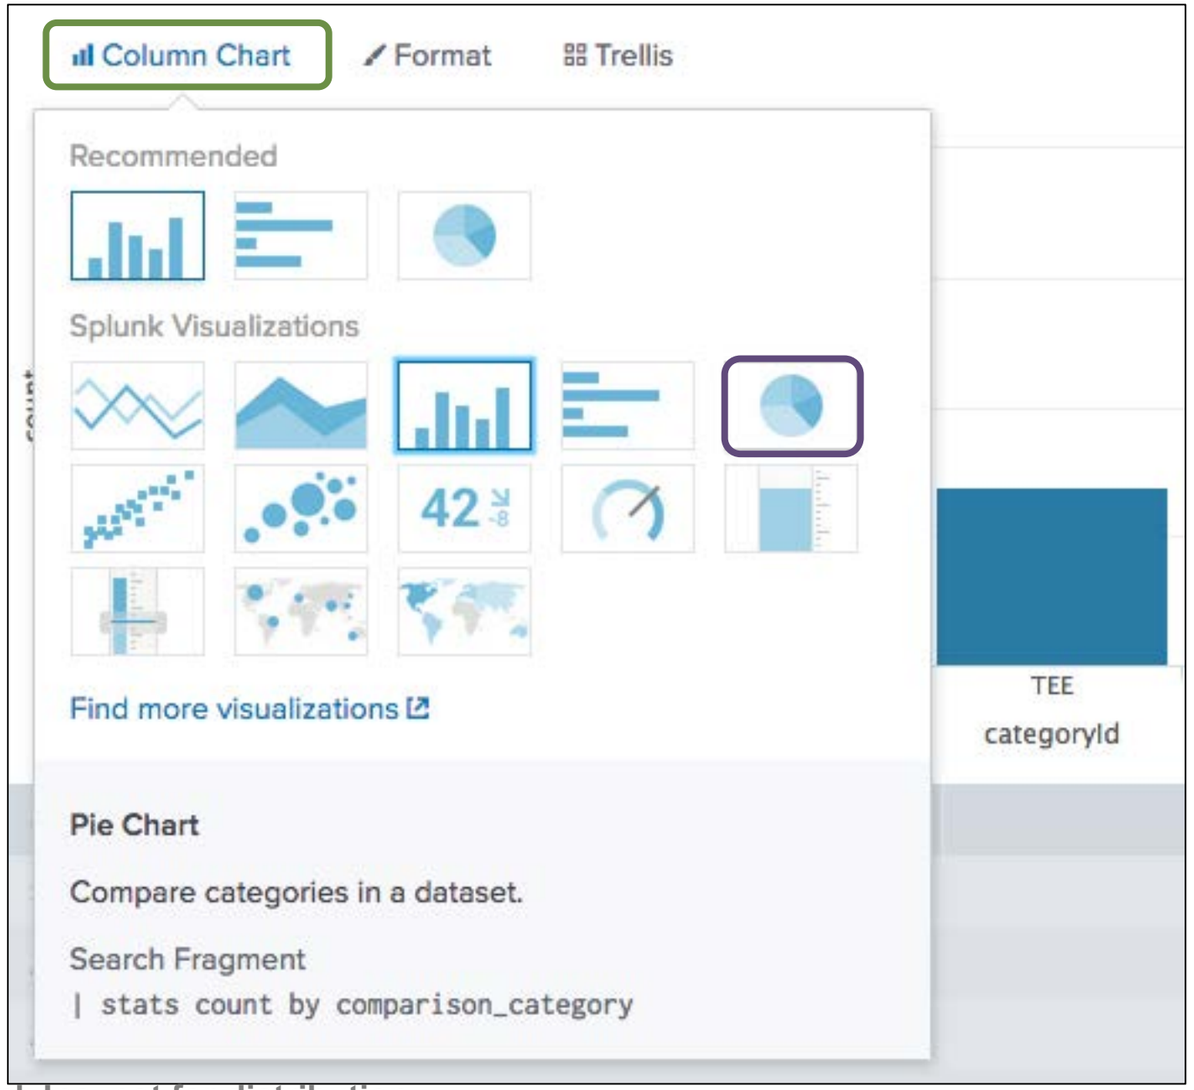

Changing the Visualization Format

- The **Format** menu allows you to change formatting options
- For example, for bar and column charts:
  - The **General** tab allows you to change Stack and Multi-series modes
  - The **X-Axis** and **Y-Axis** tabs allow you to change the axis labels and orientation
  - **Chart Overlay** allows you to add context to the chart by overlaying other field values

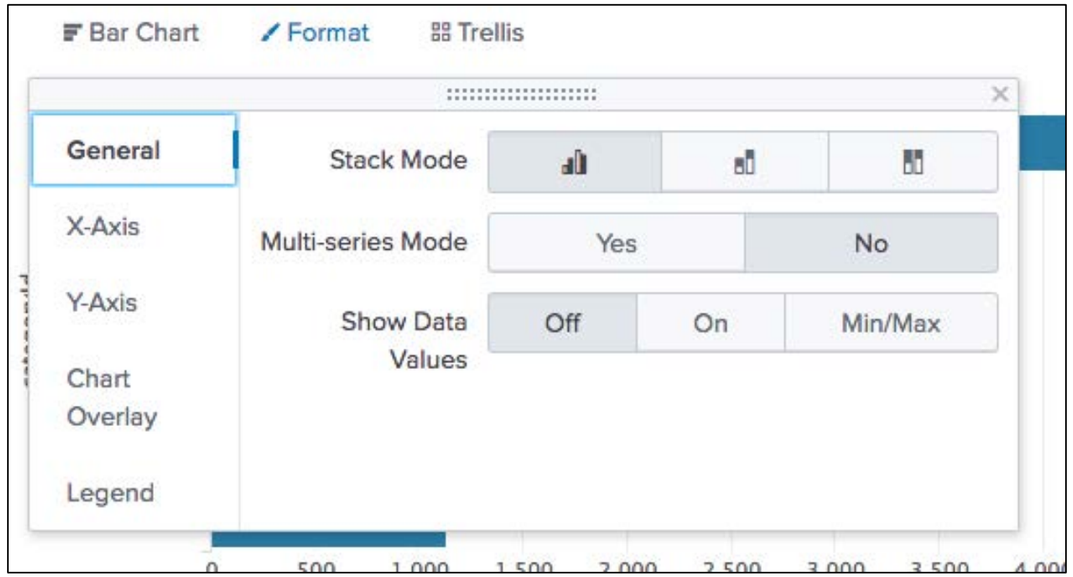

  - The **Legend** tab allows you to position the visualization legend as desired
  - **Stack Mode** allows you to stack colors to improve column chart readability when several colors are involved

Changing the Visualization Format (cont.)

- Show Data Values determines whether to show data values in the visualization
- If Min/Max is selected, data is only shown on the bars containing the minimum and maximum values

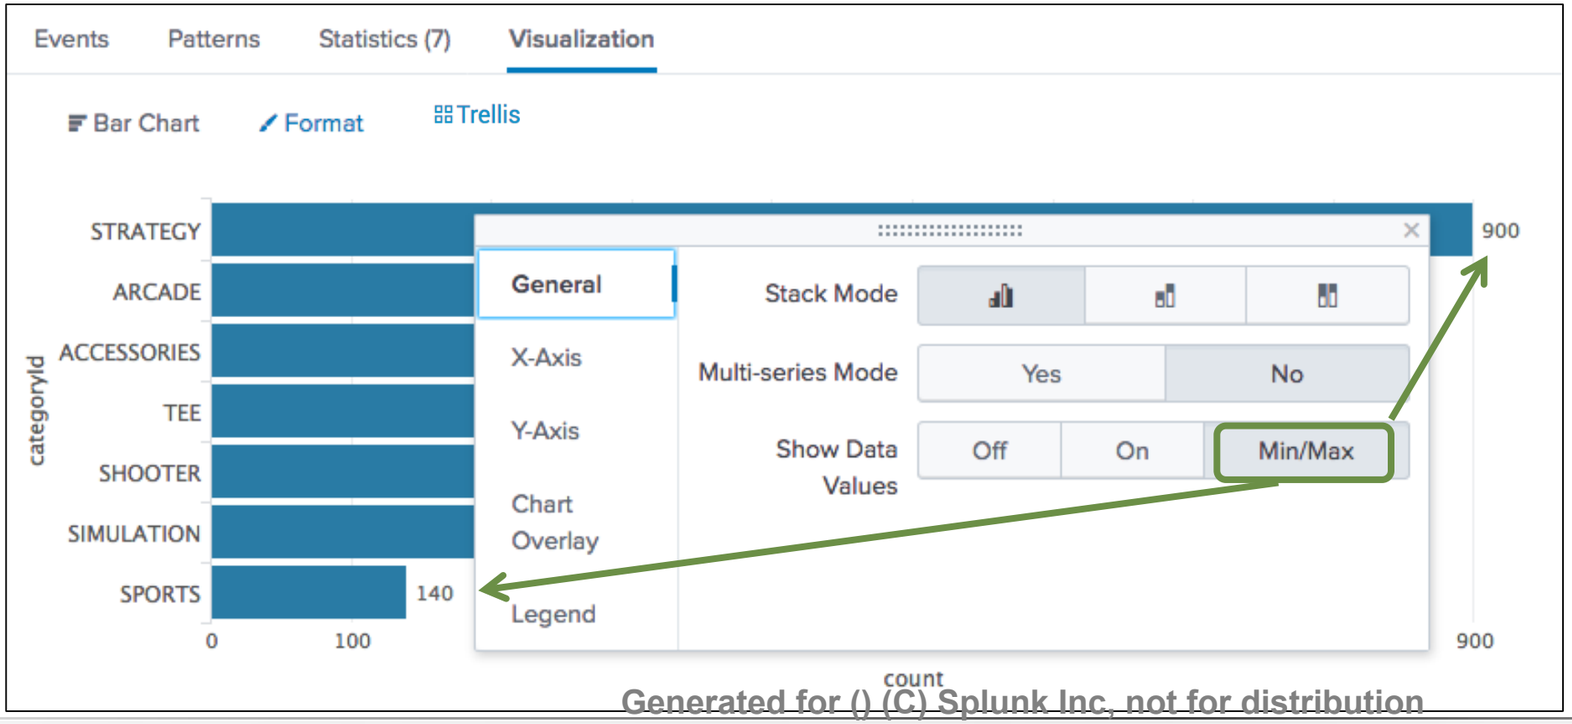

Note

When you make a change to the visualization settings – such as Min/Max – the visualization updates immediately.

Note

Learn more about modes and axes in the *Splunk Fundamentals 2* course. These modes require more sophisticated searches.

Viewing Results as a Table

Switch to the **Statistics** tab to view the results as a table

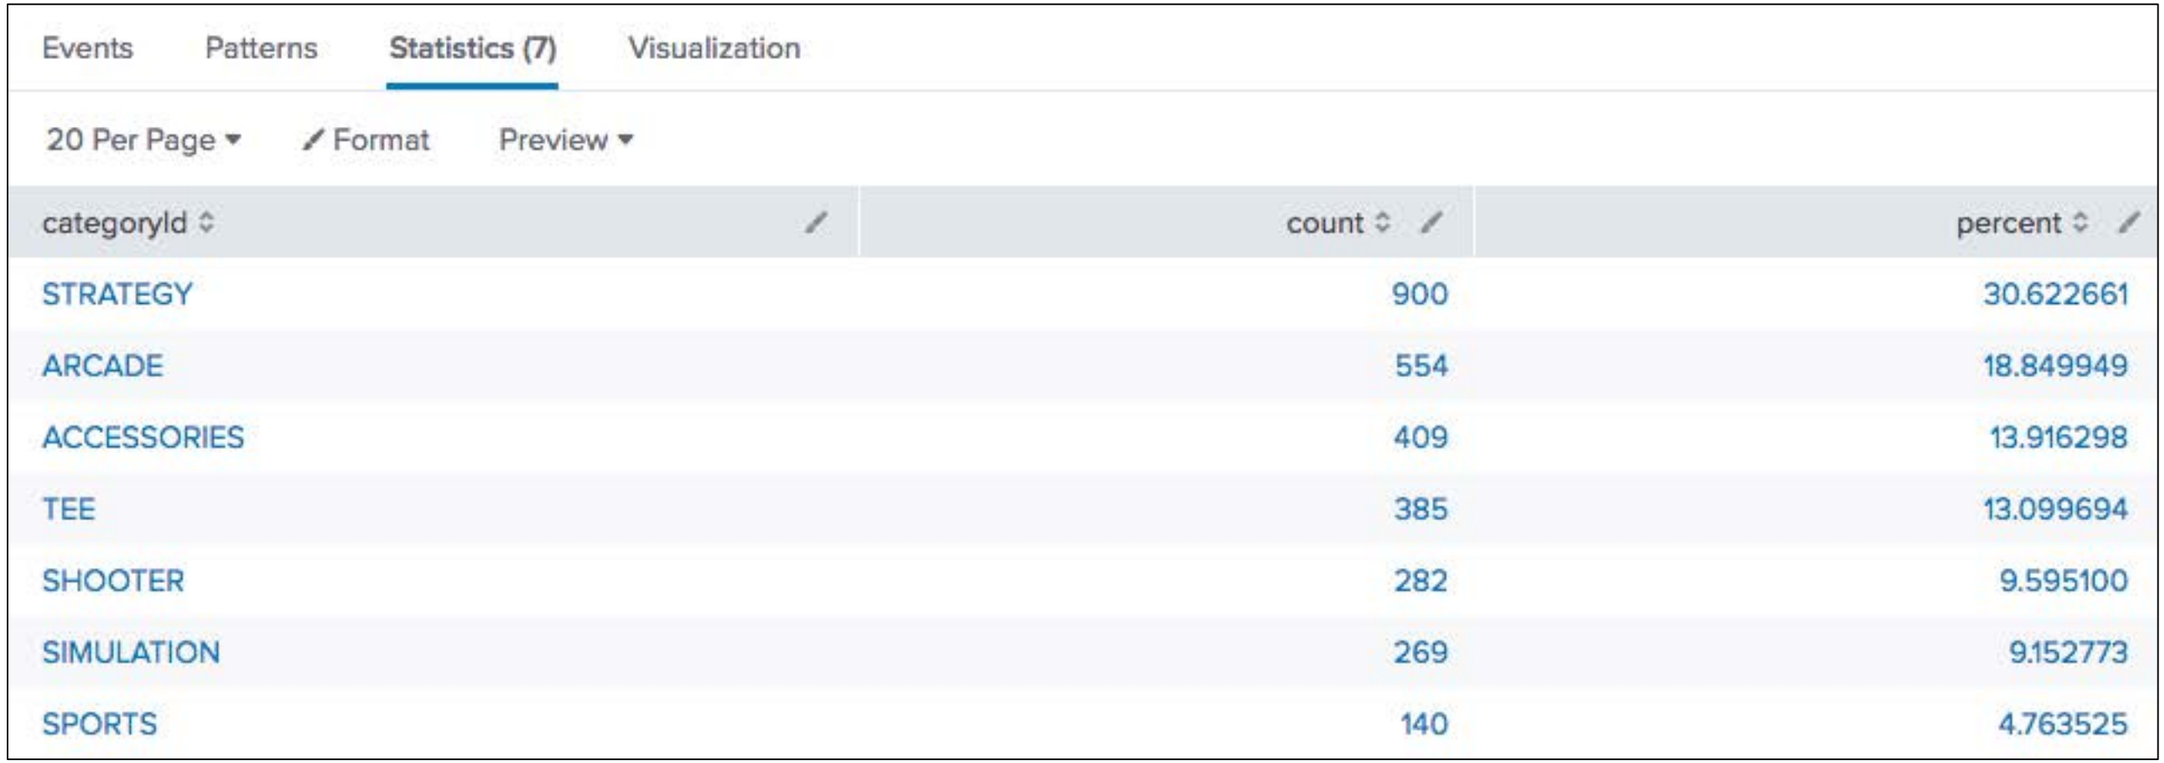

categoryId	count	percent
STRATEGY	900	30.622661
ARCADE	554	18.849949
ACCESSORIES	409	13.916298
TEE	385	13.099694
SHOOTER	282	9.595100
SIMULATION	269	9.152773
SPORTS	140	4.763525

Using the Data Overlay Format

- **Heat map** highlights outstanding values

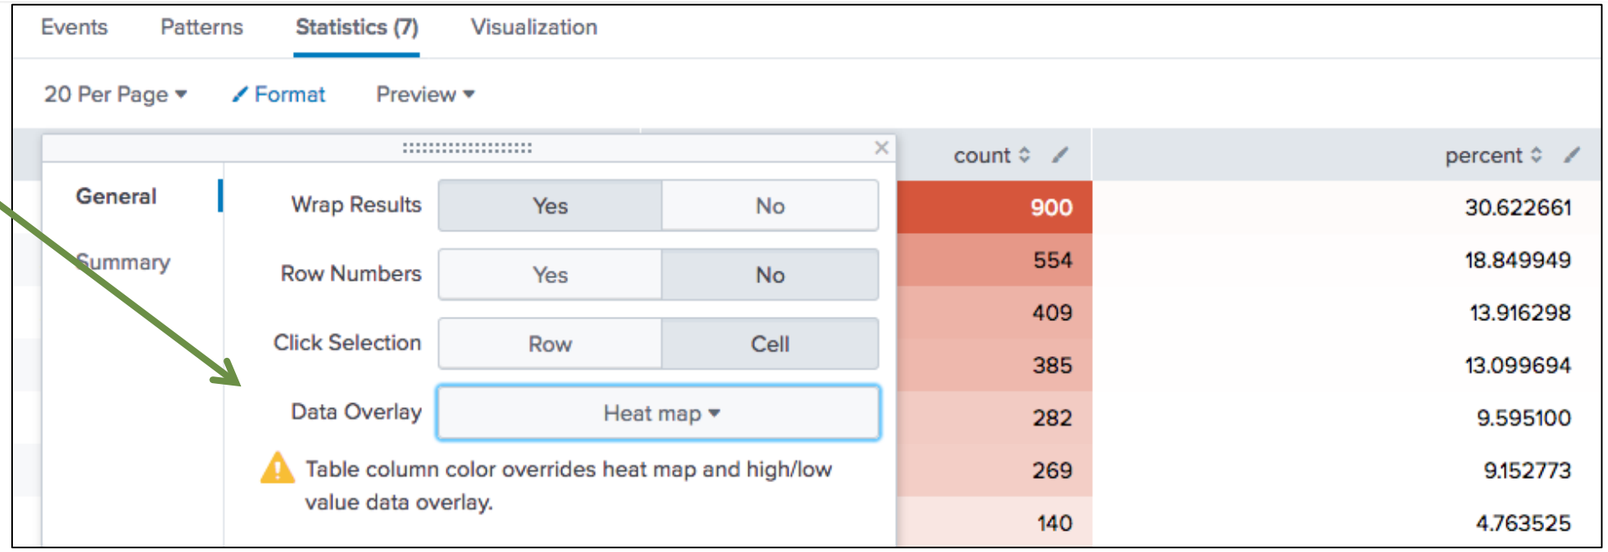

- **High and low values** highlights max and min of non zero values

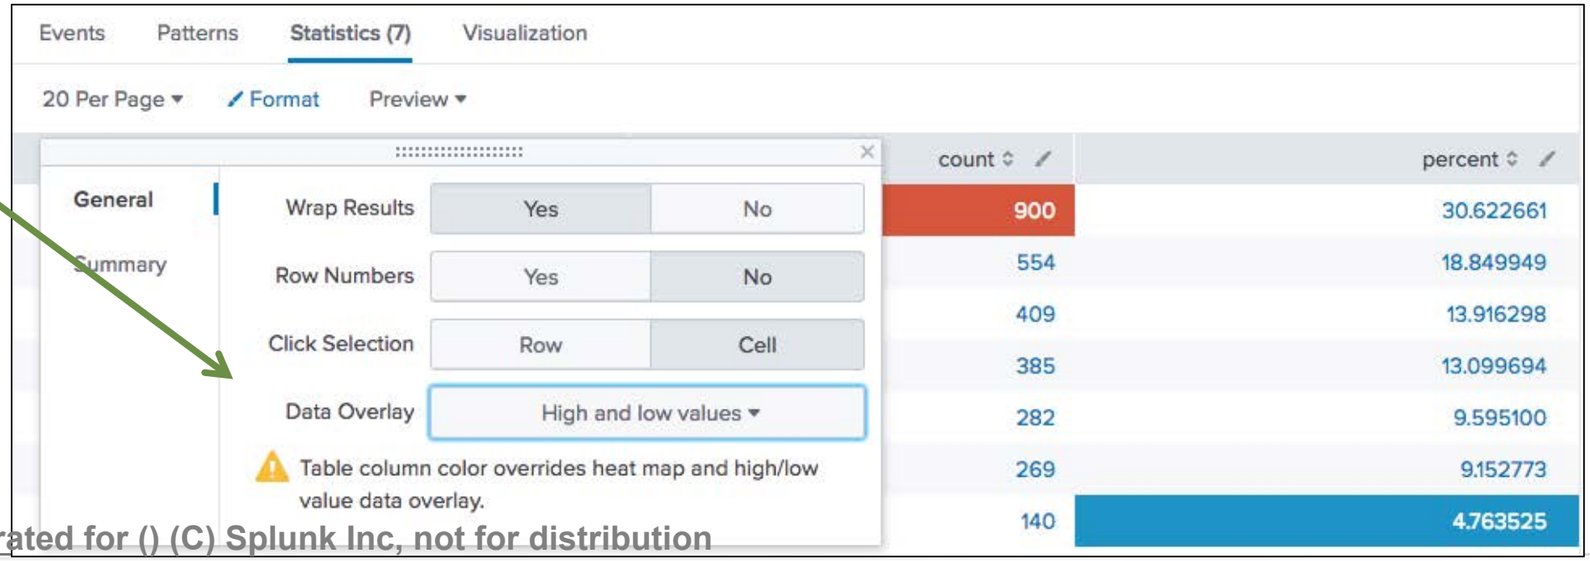

What Is a Dashboard?

- A dashboard consists of one or more panels displaying data visually in a useful way – such as events, tables, or charts
- A report can be used to create a panel on a dashboard

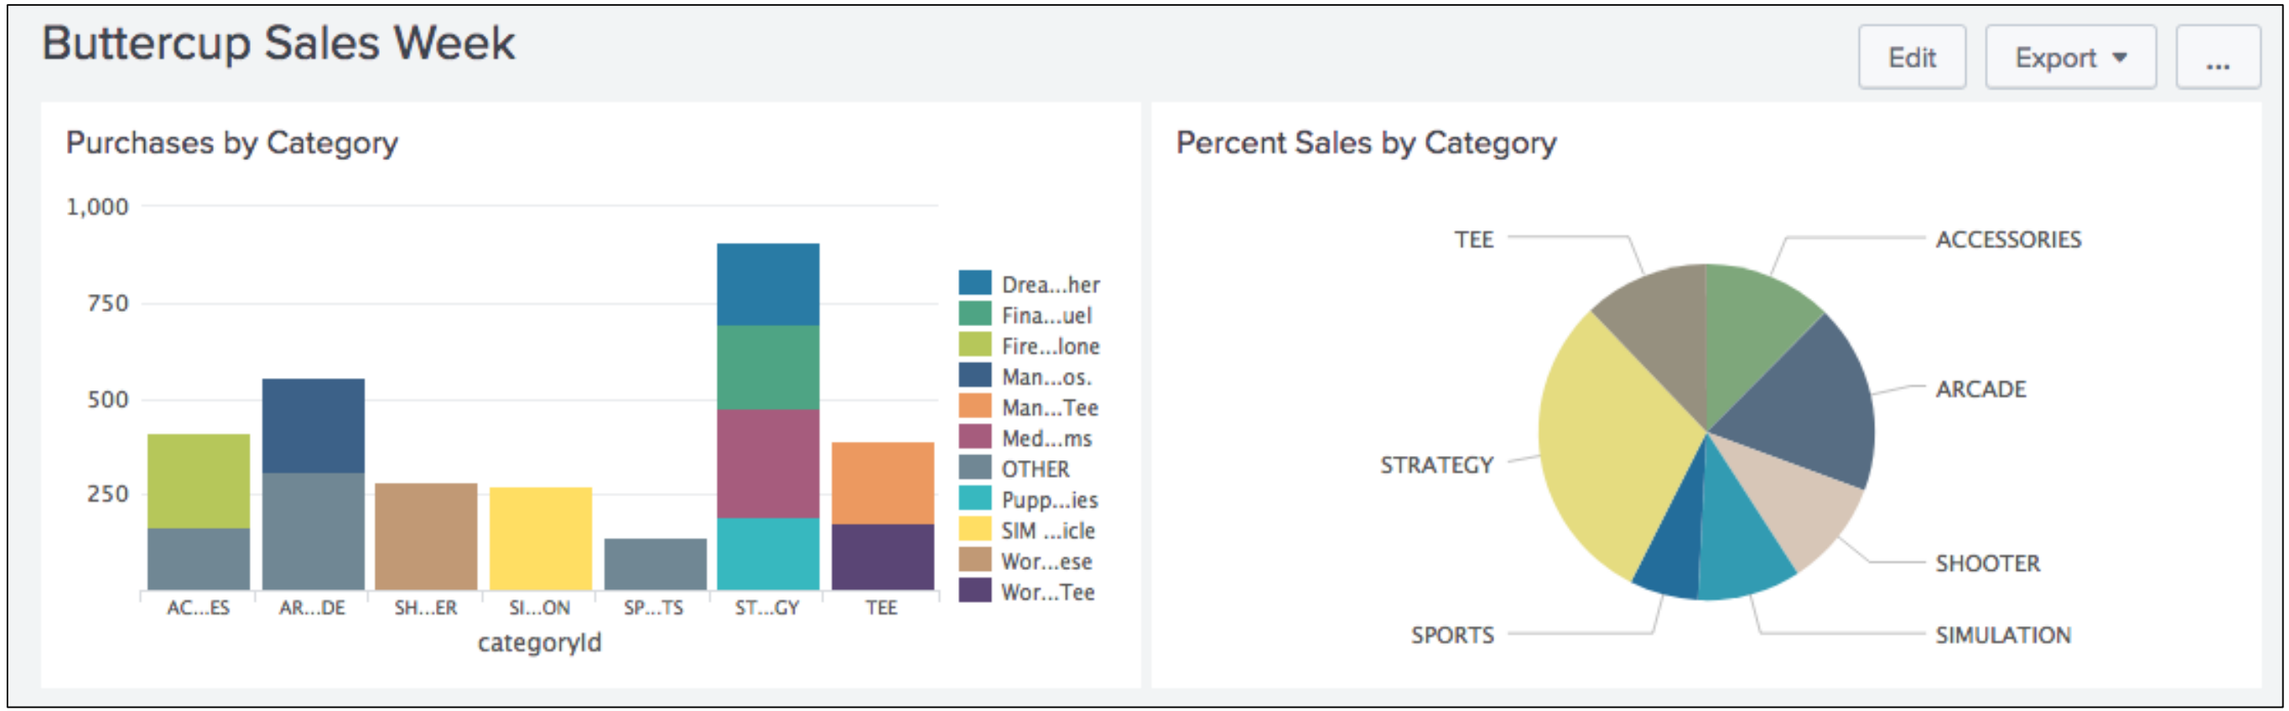

Adding a Report to a Dashboard

In the report, click Add to Dashboard to begin

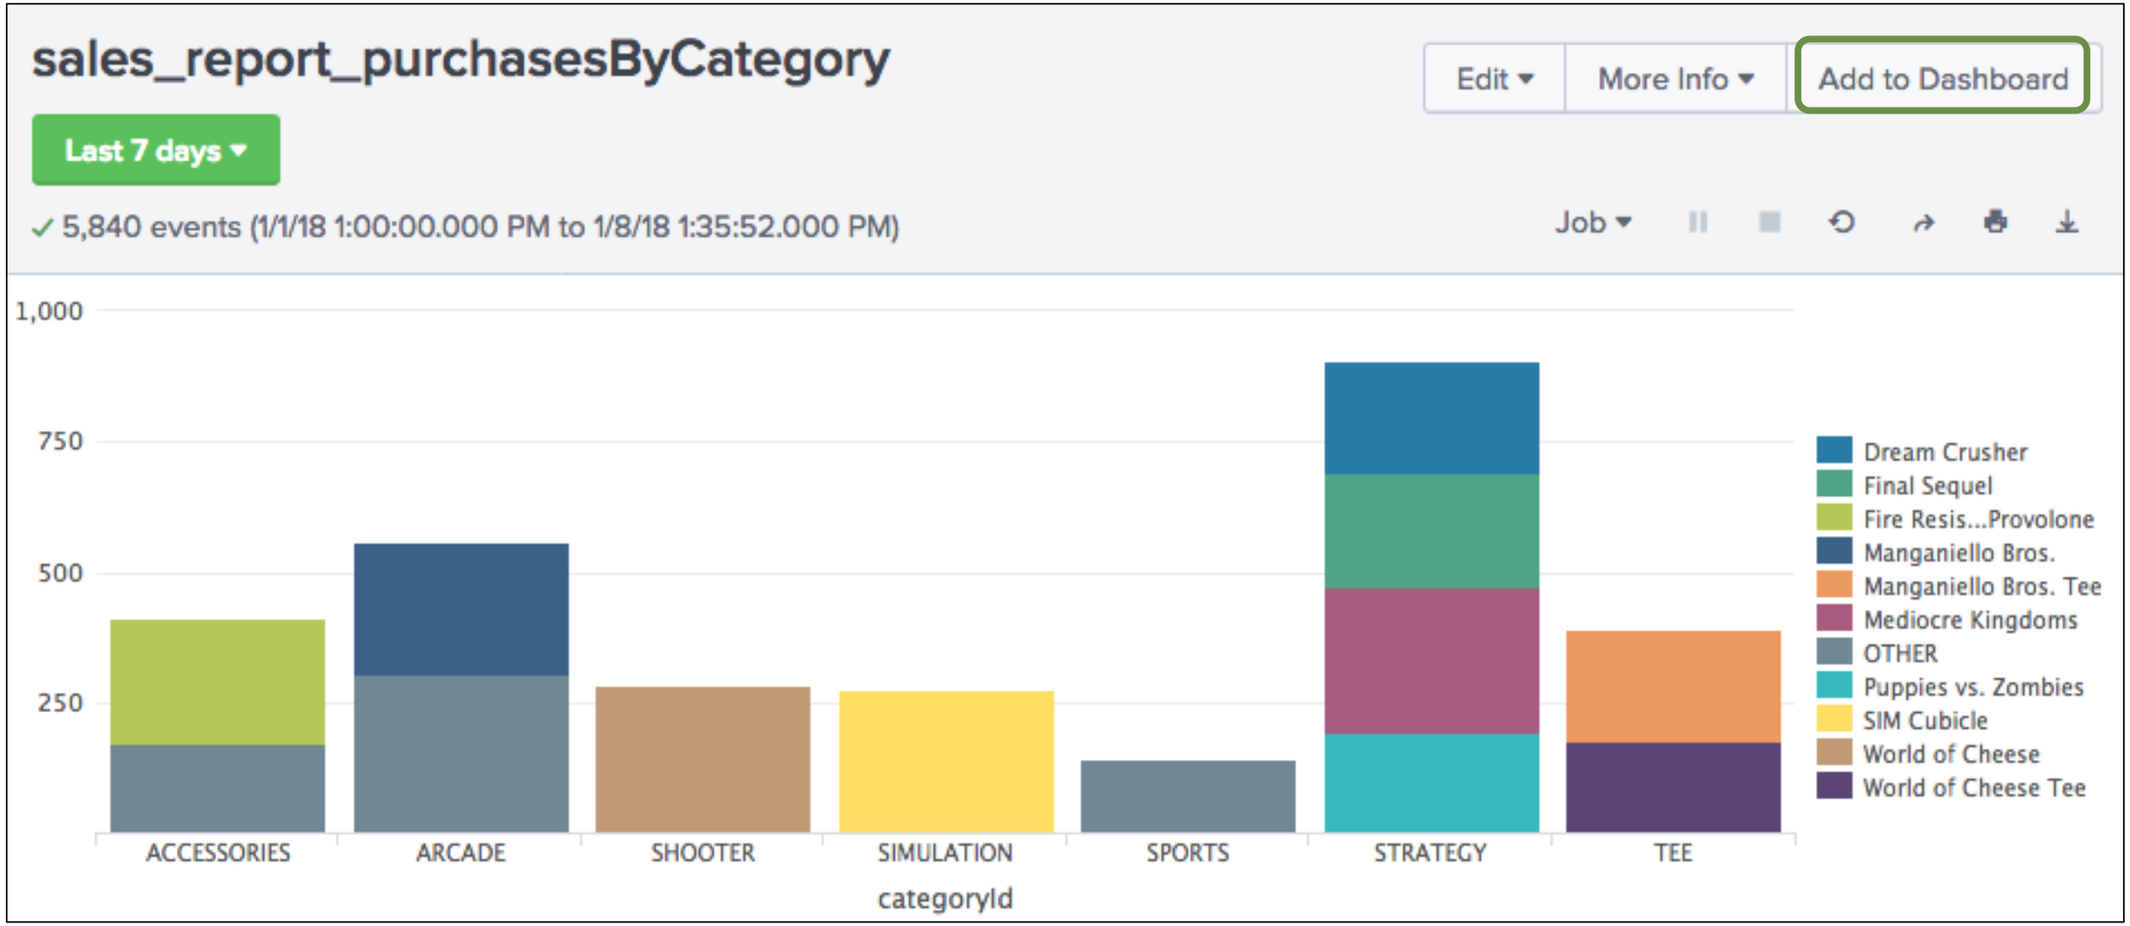

Adding a Report to a Dashboard (cont.)

- A Name the dashboard and optionally provide a description
- B Change the permissions (use Private until tested)
- C Enter a meaningful title for the panel
- D For **Panel Powered By,** click **Report**
- E For Panel Content, click **Column Chart** to display the visualization in the dashboard

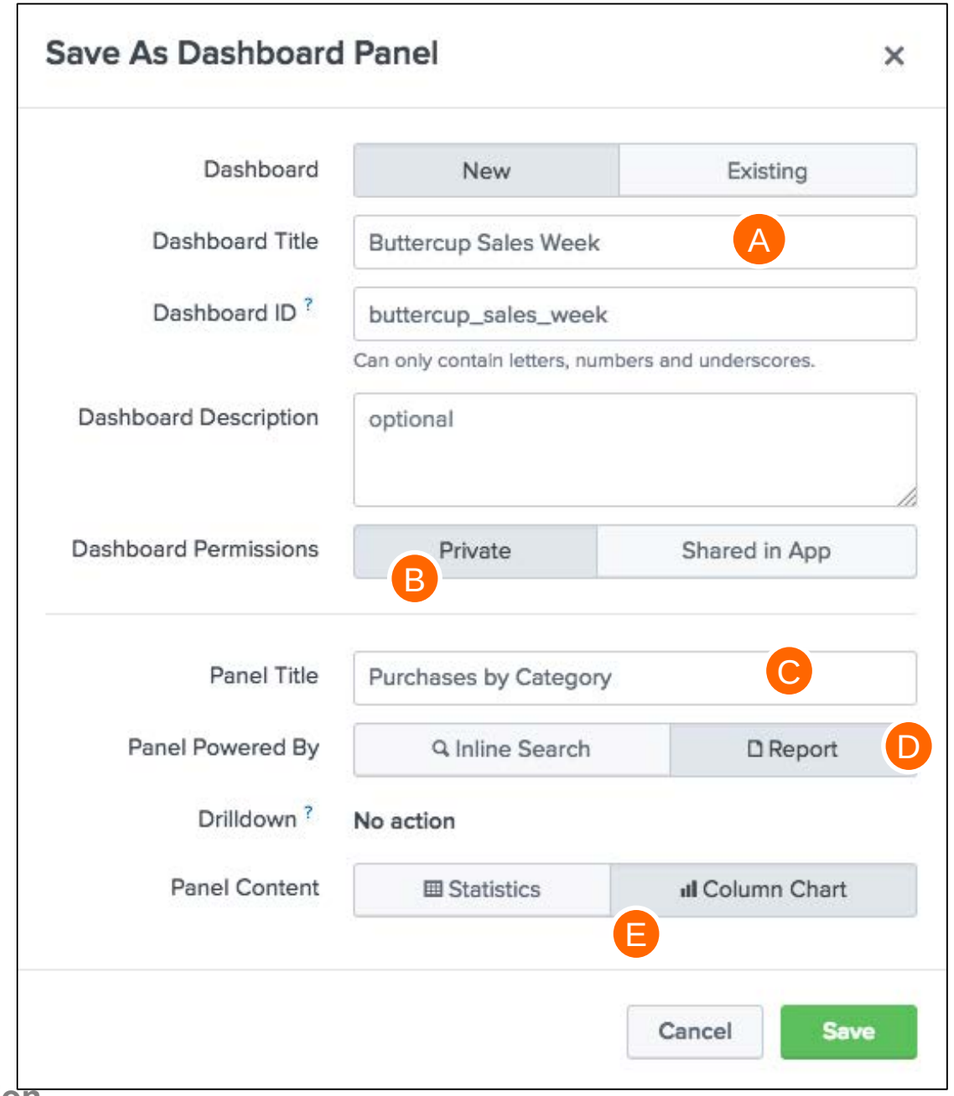

Note

The Dashboard ID is automatically populated with a unique value used by Splunk and should not be changed.

Adding a Report to a Dashboard (cont.)

After it is saved, you can view the dashboard immediately, or select the dashboard from the **Dashboards** view

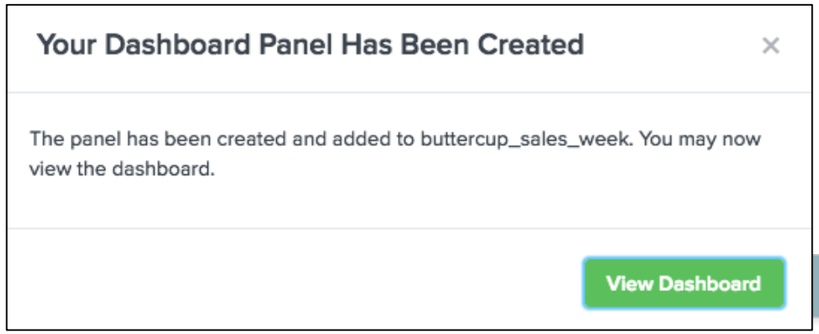

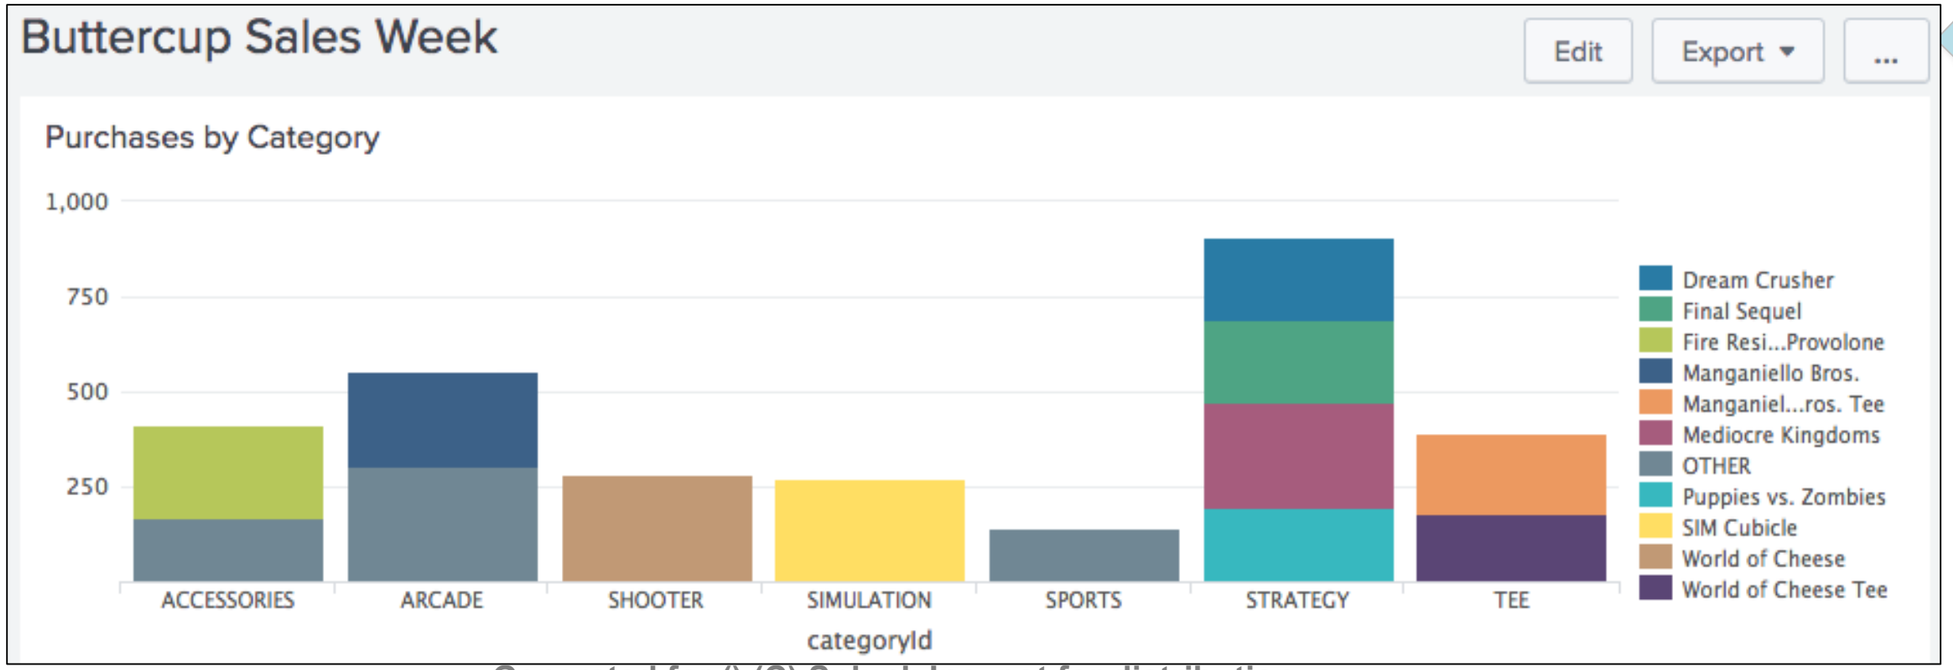

Why Create Panels from Reports?

- It is efficient to create most dashboard panels based on reports because
  - A single report can be used across different dashboards
  - This links the report definition to the dashboard
- Any change to the underlying report affects every dashboard panel that utilizes that report

Editing Panels

- After saving the panel, a window appears from which you can view the updated dashboard
- Click Edit to customize the dashboard

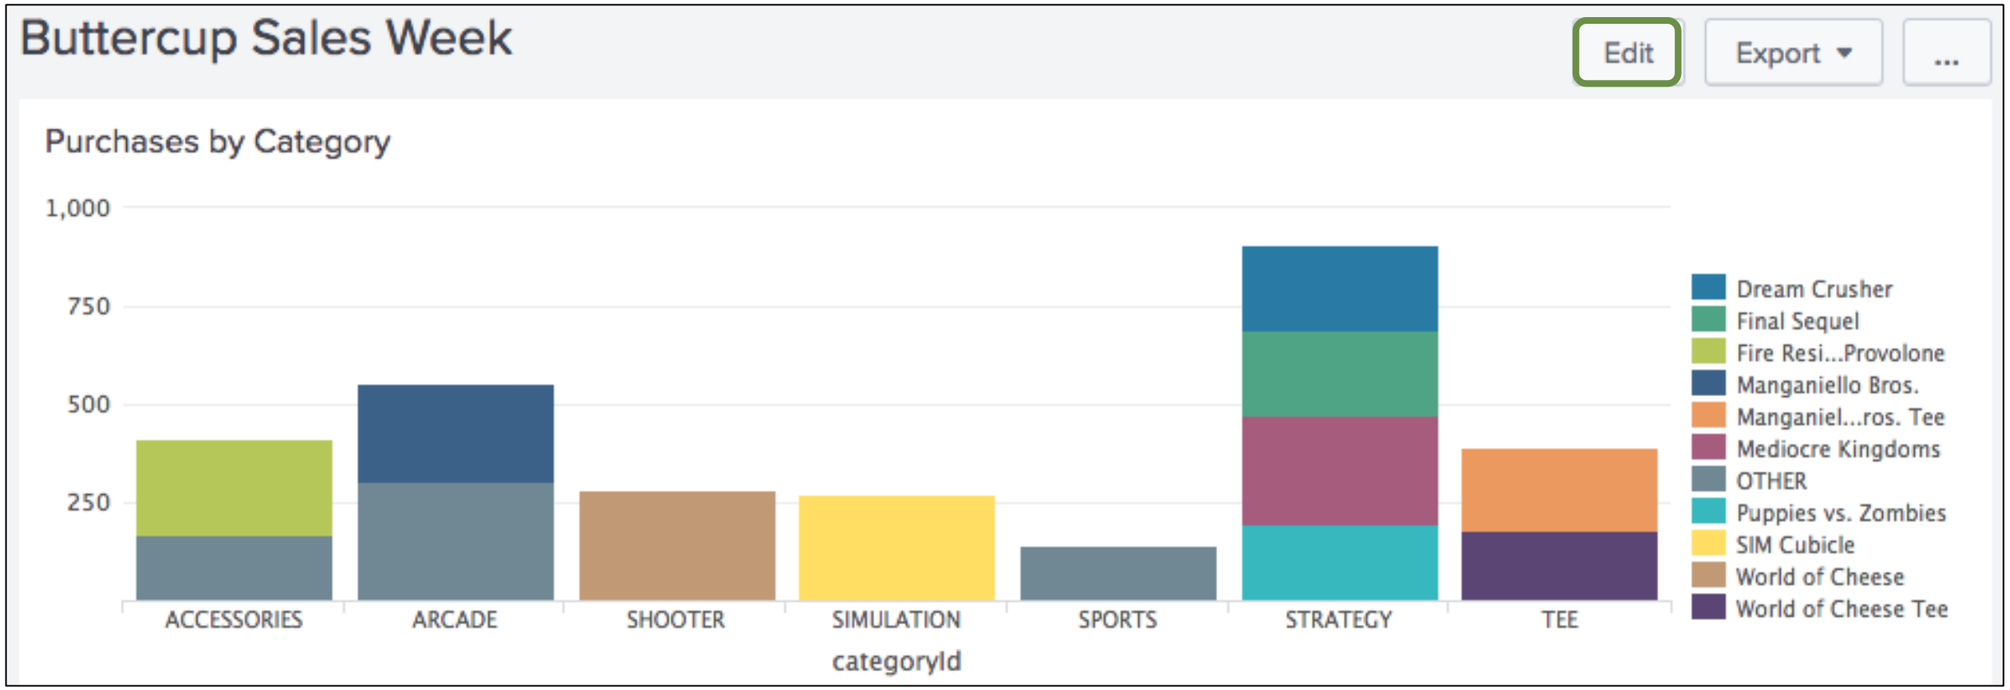

Editing Panel Layout

Click on the dotted bar on a panel to drag the panel to a new location

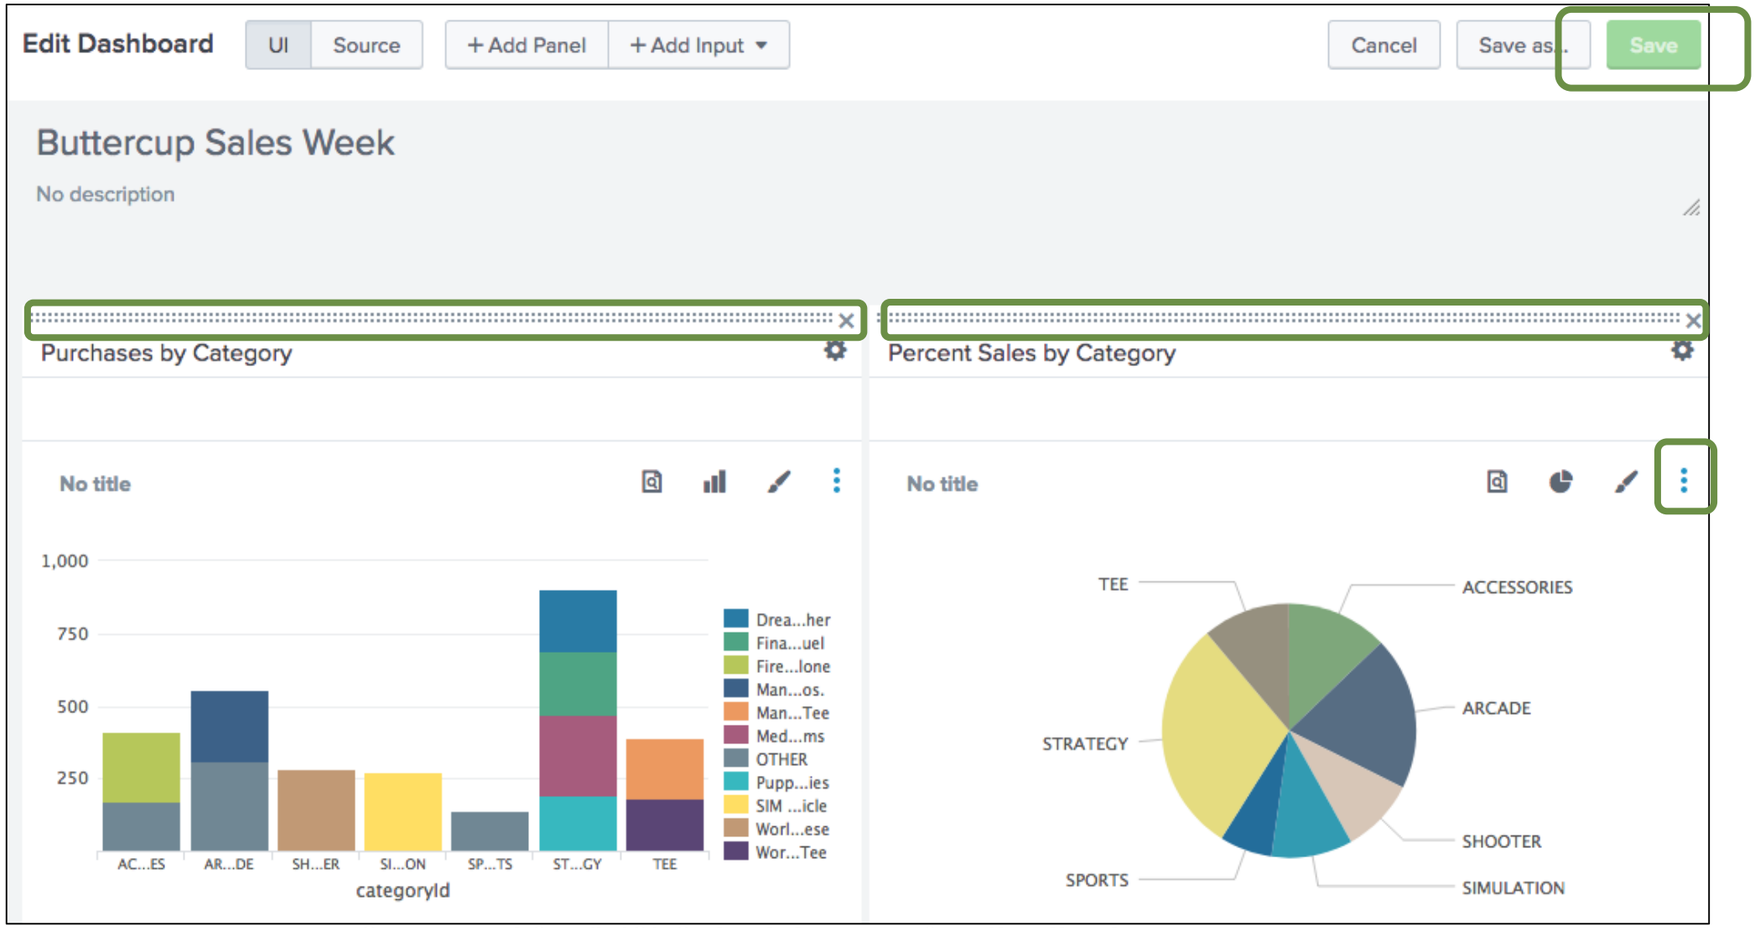

More Options icon (discussed on next slide)

Edit Panel Drilldown Options

1. In Edit Dashboard mode, click the More Options icon on any panel and select Edit Drilldown
2. In Drilldown Editor, select Link to search to access search directly from visualization

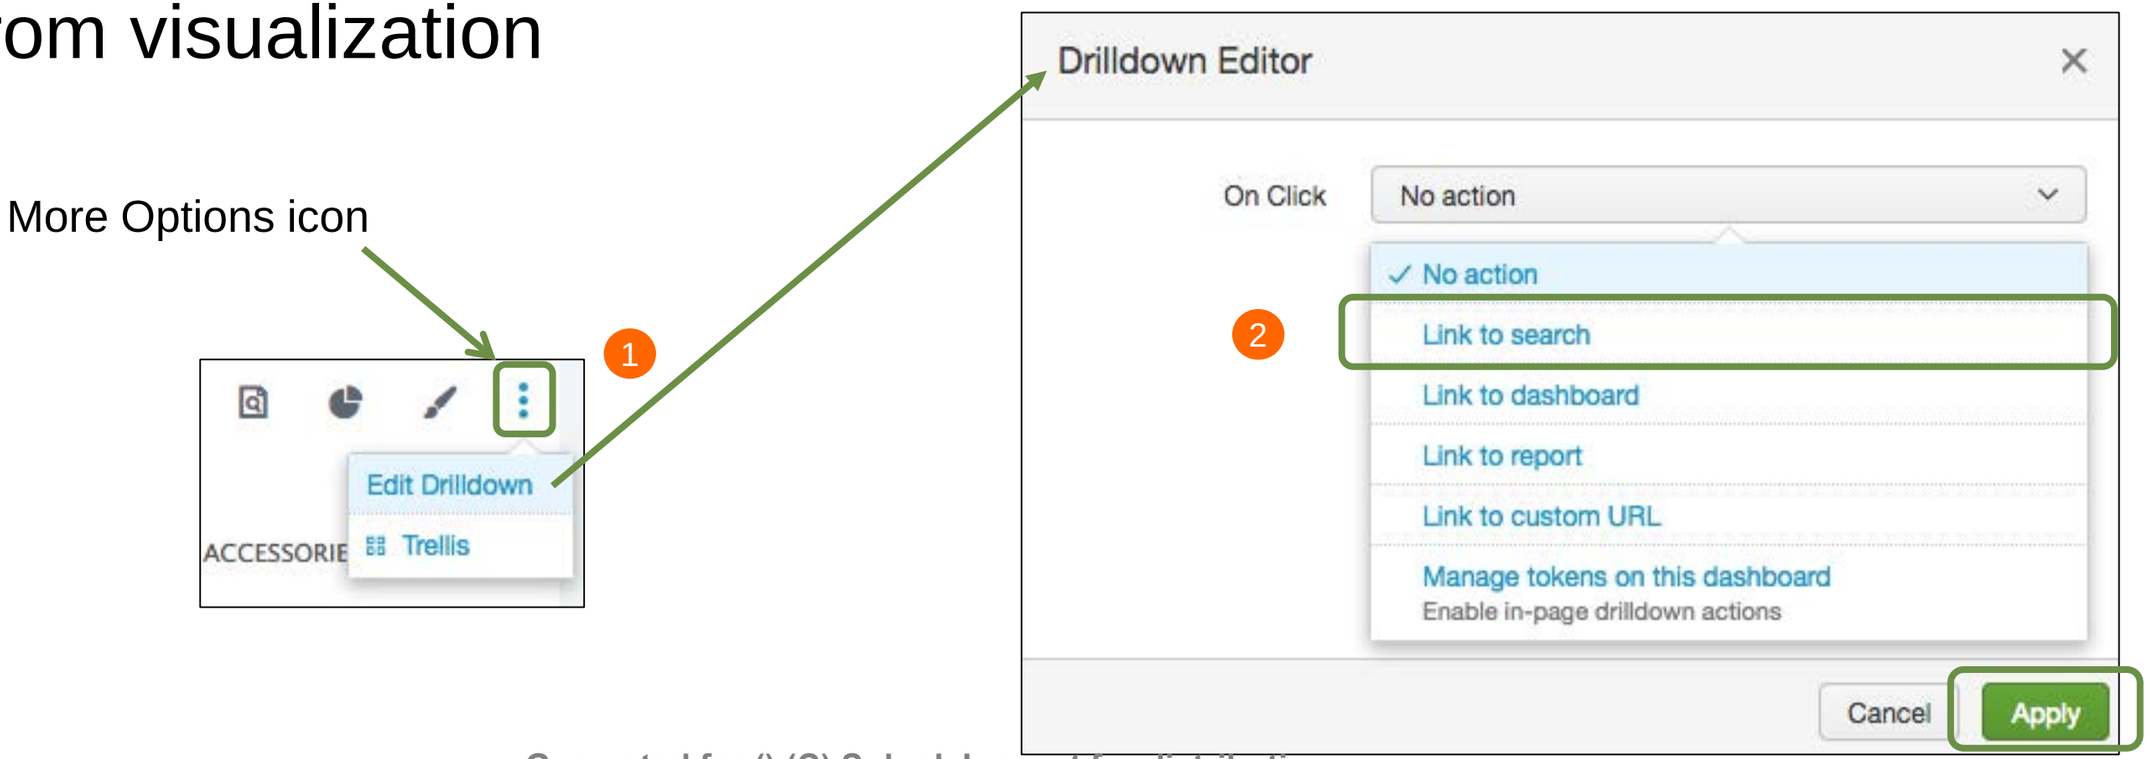

Drilldown from Visualization to Search

Once drilldown option is set, click an object in a chart or table to see its underlying events in Search view

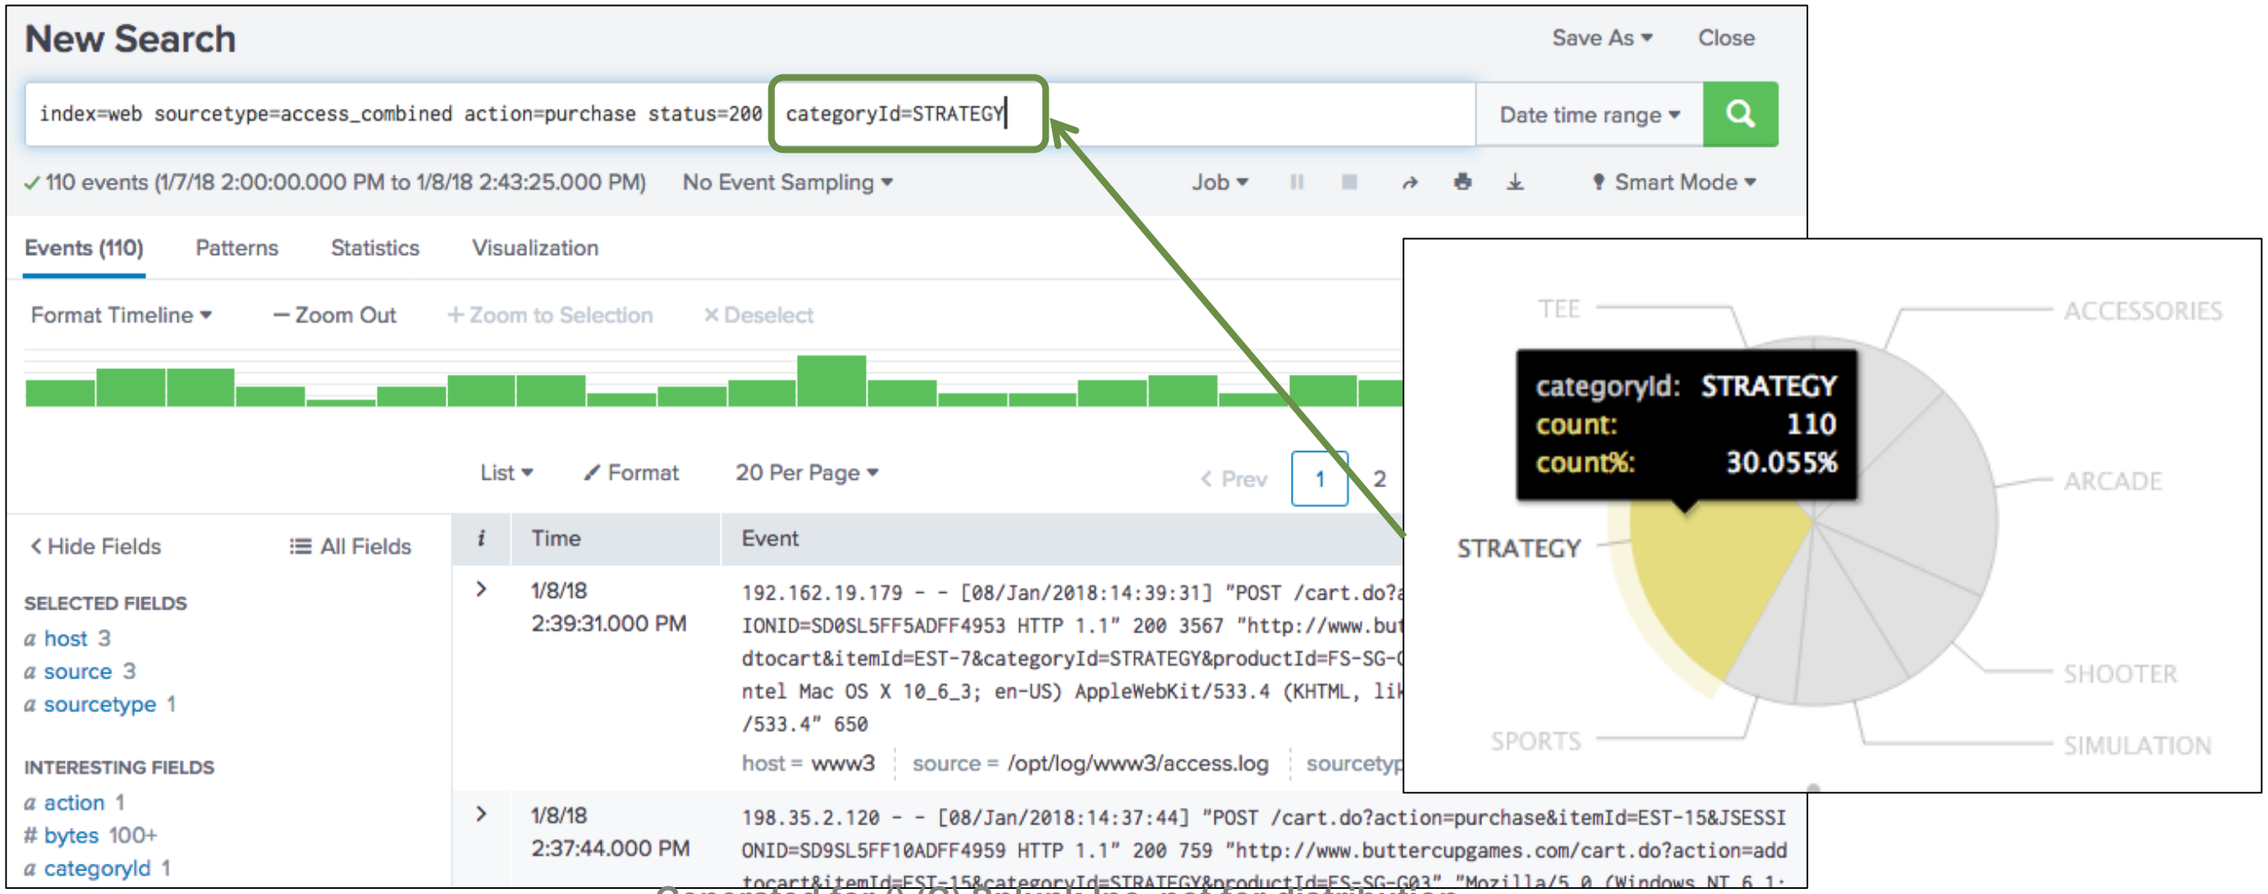

Clone a Dashboard

1. Click the ellipsis menu (**...**) and select **Clone**
2. Change the **Title** as desired and click **Clone Dashboard**

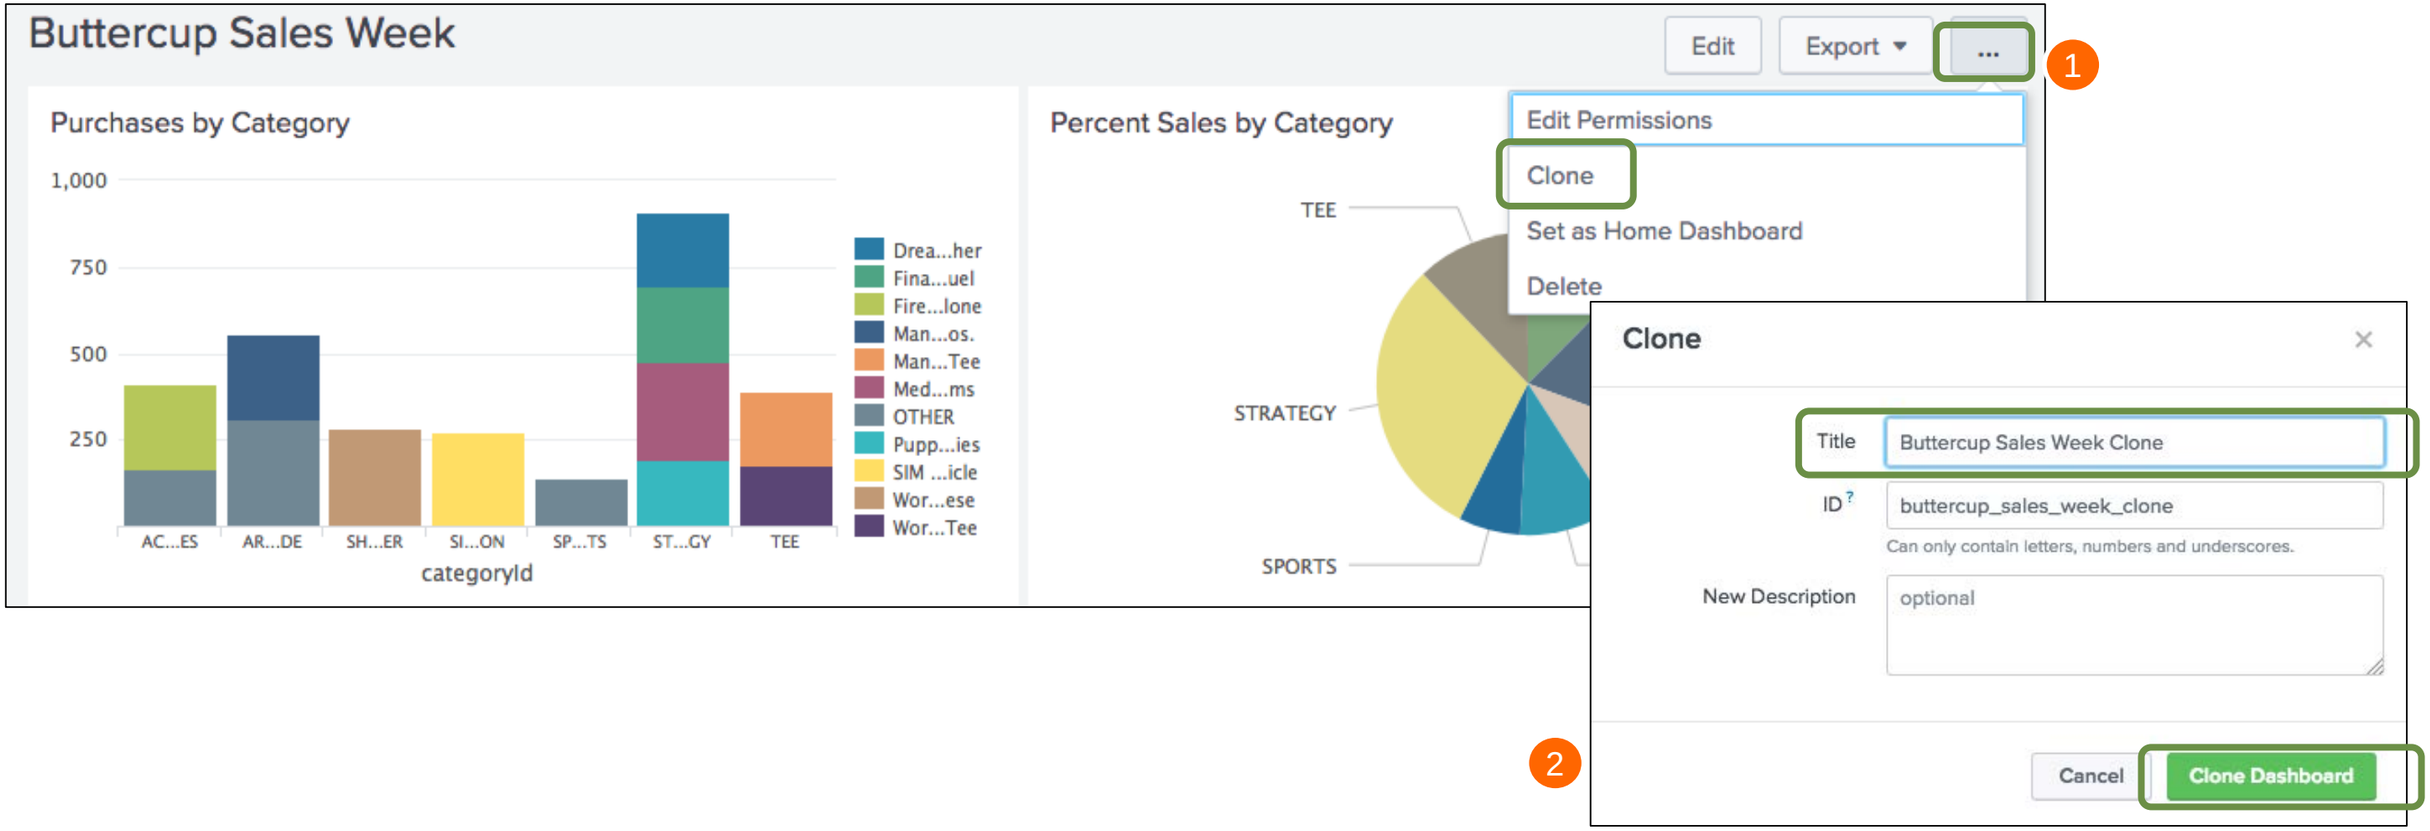

Export a Dashboard

Dashboards can be exported as PDF or printed

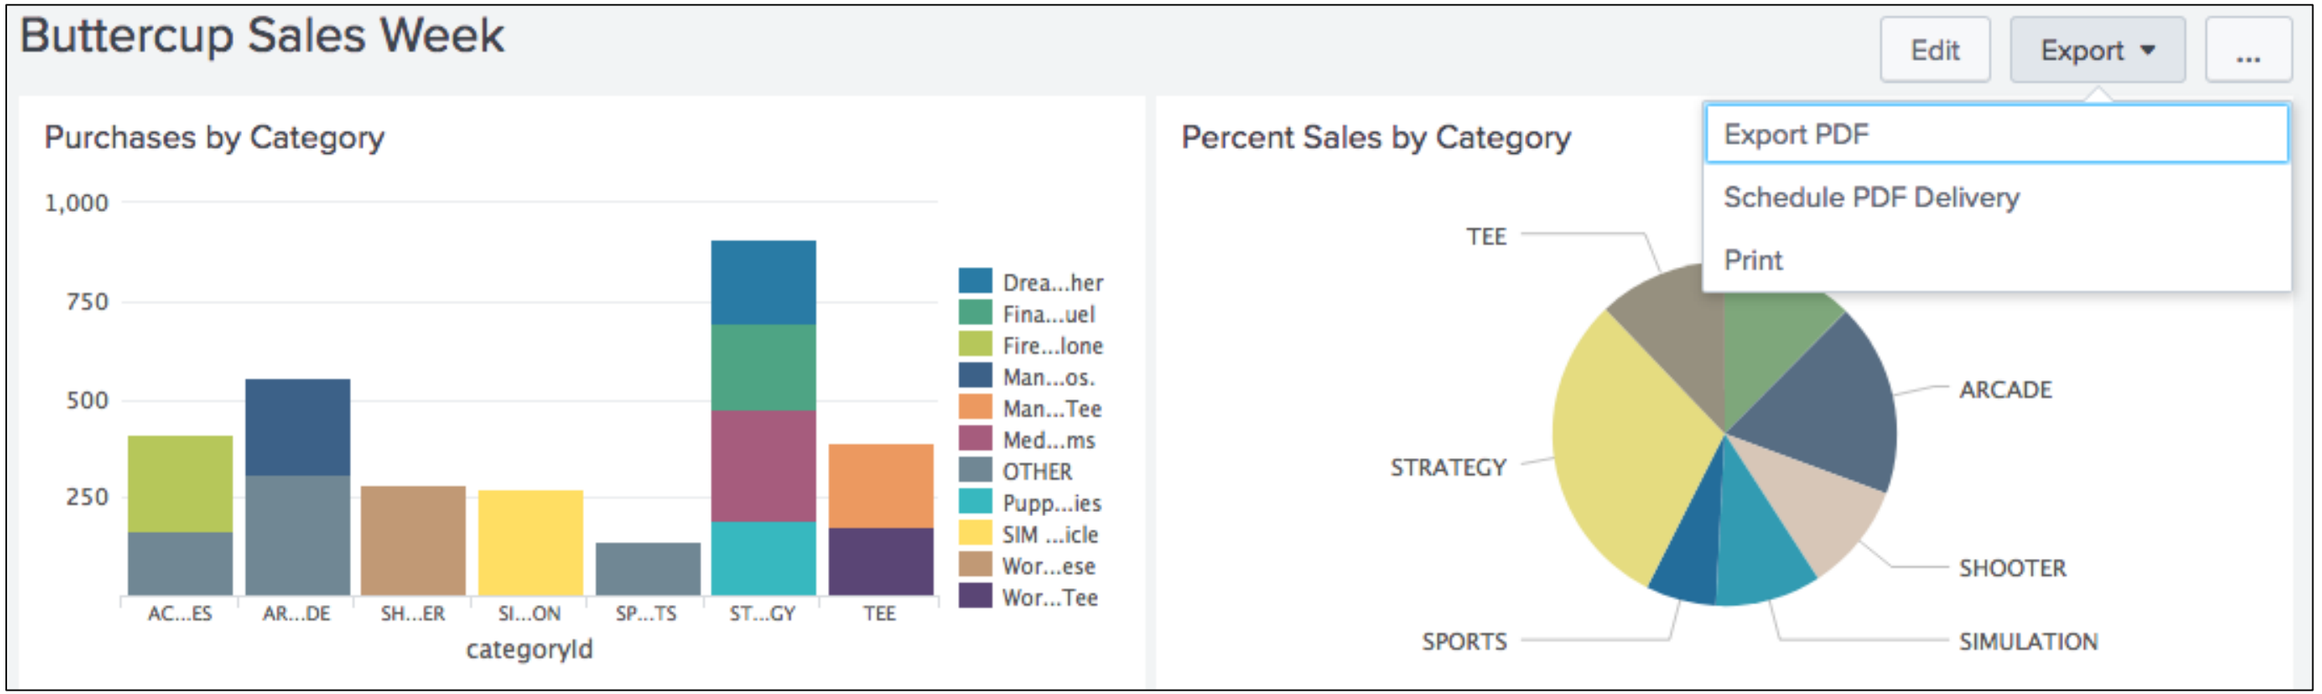

Set the Default Dashboard

- Set a dashboard to appear by default in the bottom panel of your home view
- From the Home app, click **Choose a home dashboard**

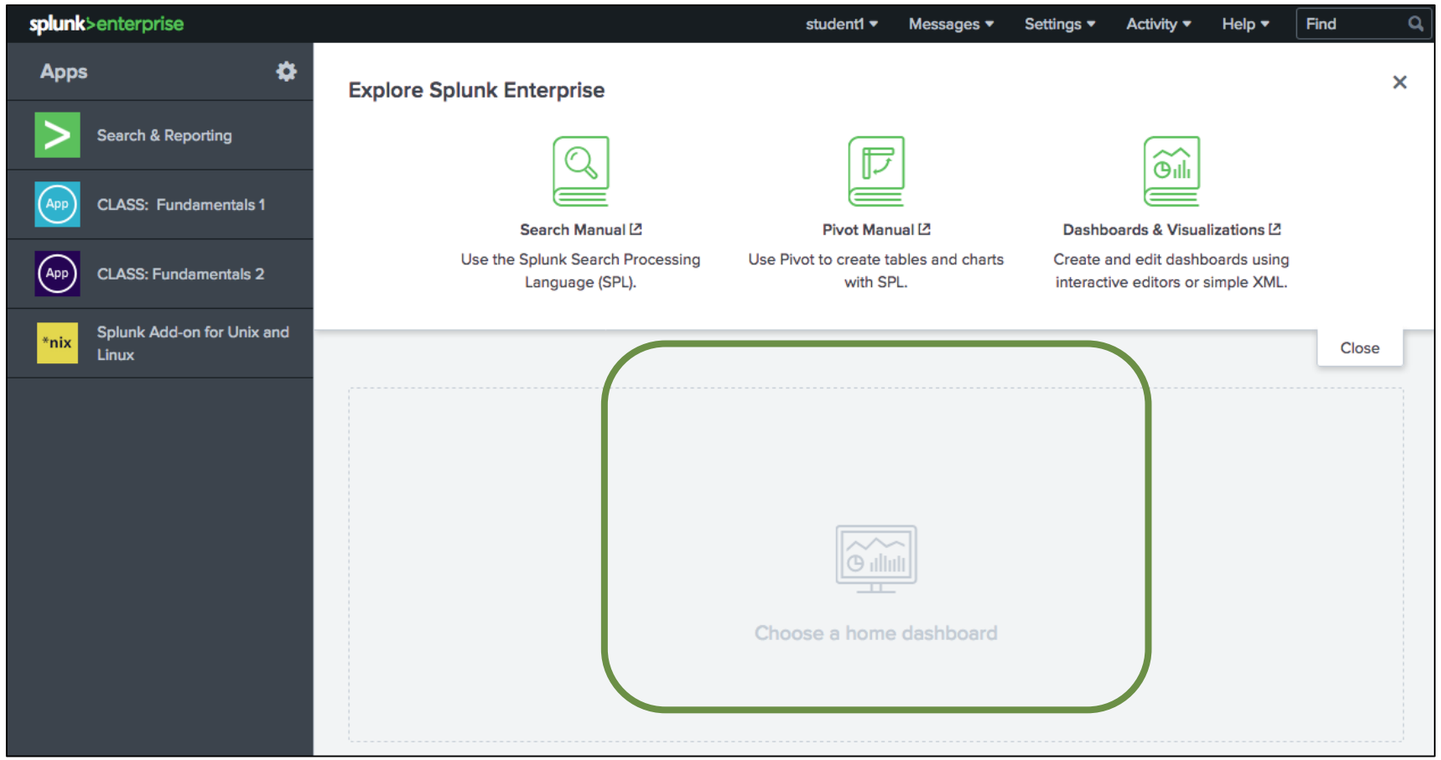

View Your Default Dashboard

After you've set a dashboard as default, your home view may look like this:

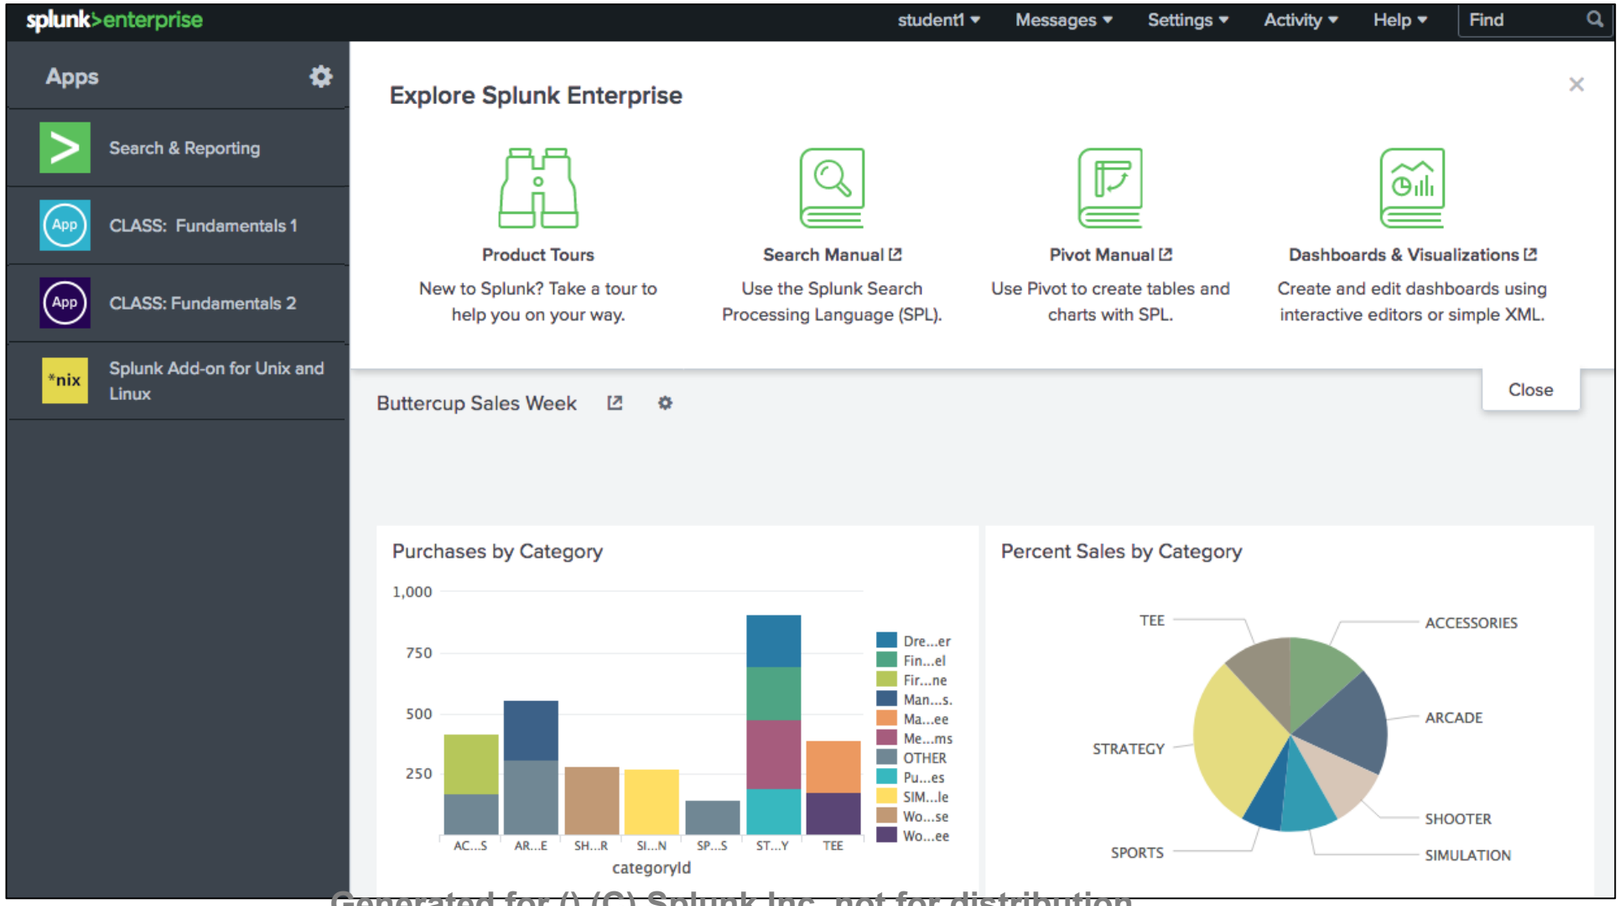

Module 11
Pivot & Datasets

Selecting a Dataset

1. From the Search & Reporting app, select the **Datasets** tab
   - Displays a list of available lookup table files ("lookups") and data models
   - Each lookup and data model represent a specific category of data
   - Prebuilt lookups and data models make it easier to interact with your data

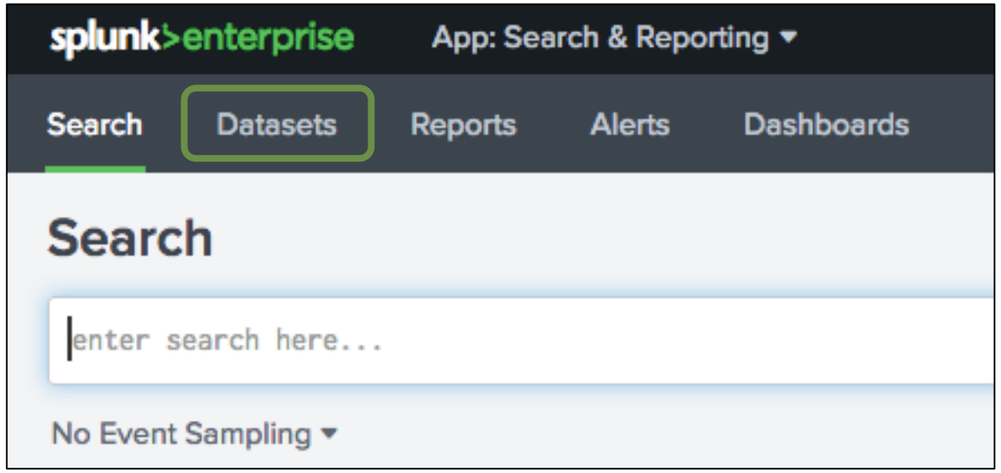

2. Click **Explore > Visualize with Pivot**

Opening in Pivot

- The Pivot automatically populates with a count of events for the selected object
- In this example, it shows all successful purchase requests for all time

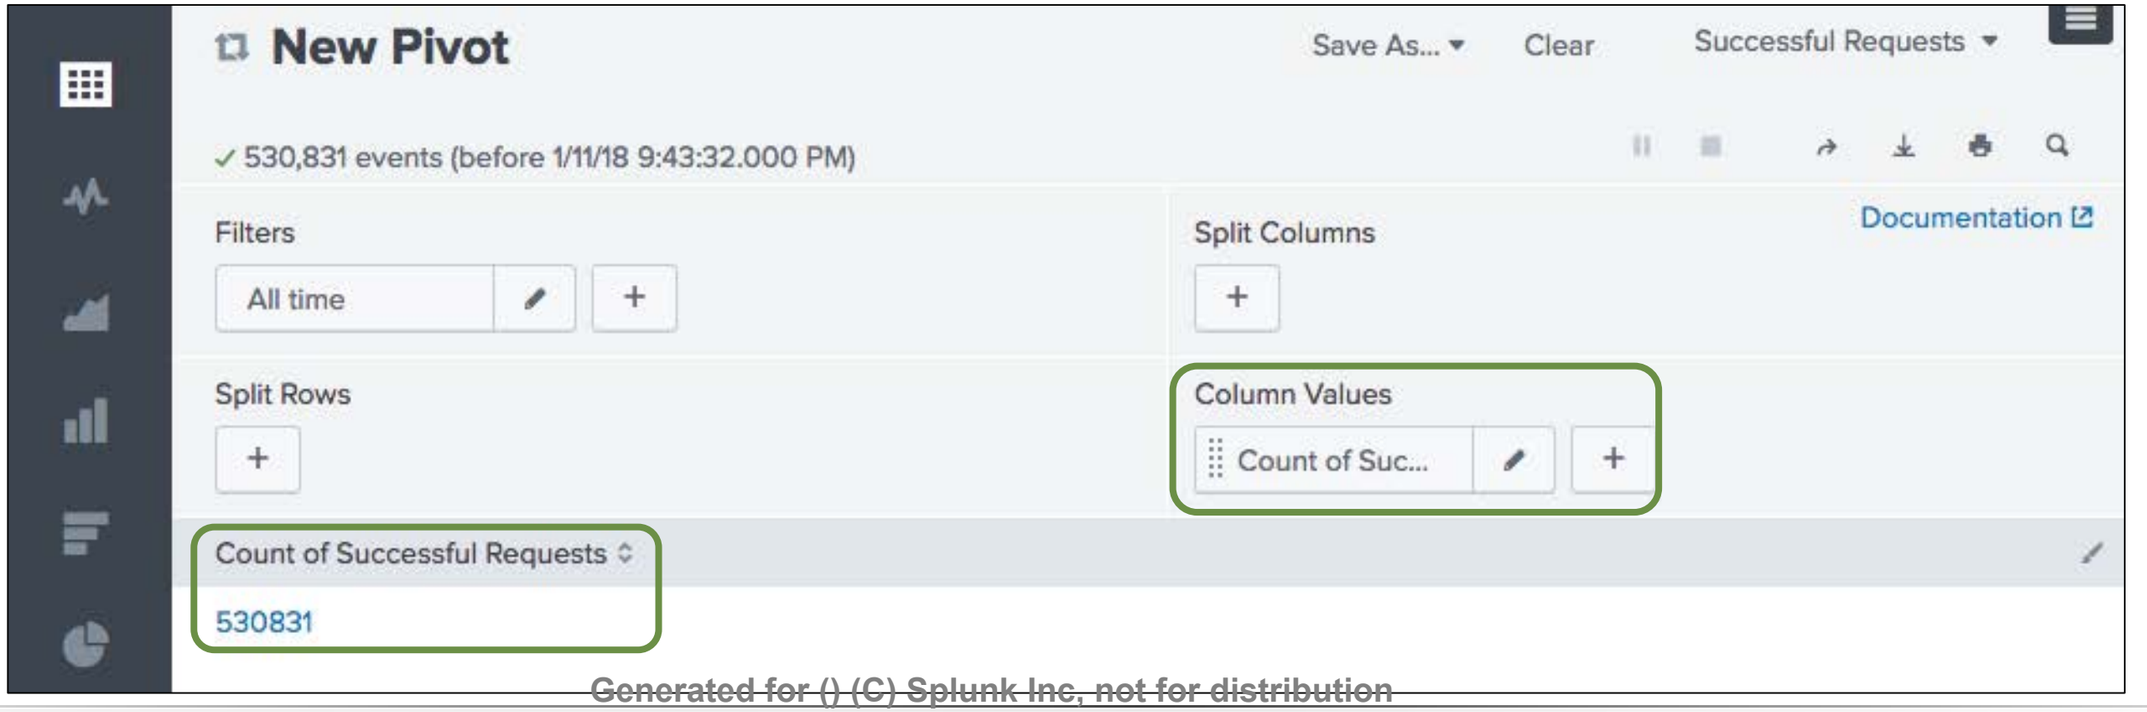

Selecting a Time Range

- The default is **All time**
- Click the pencil icon to select the desired time range
- The pivot runs immediately upon selecting the new time range

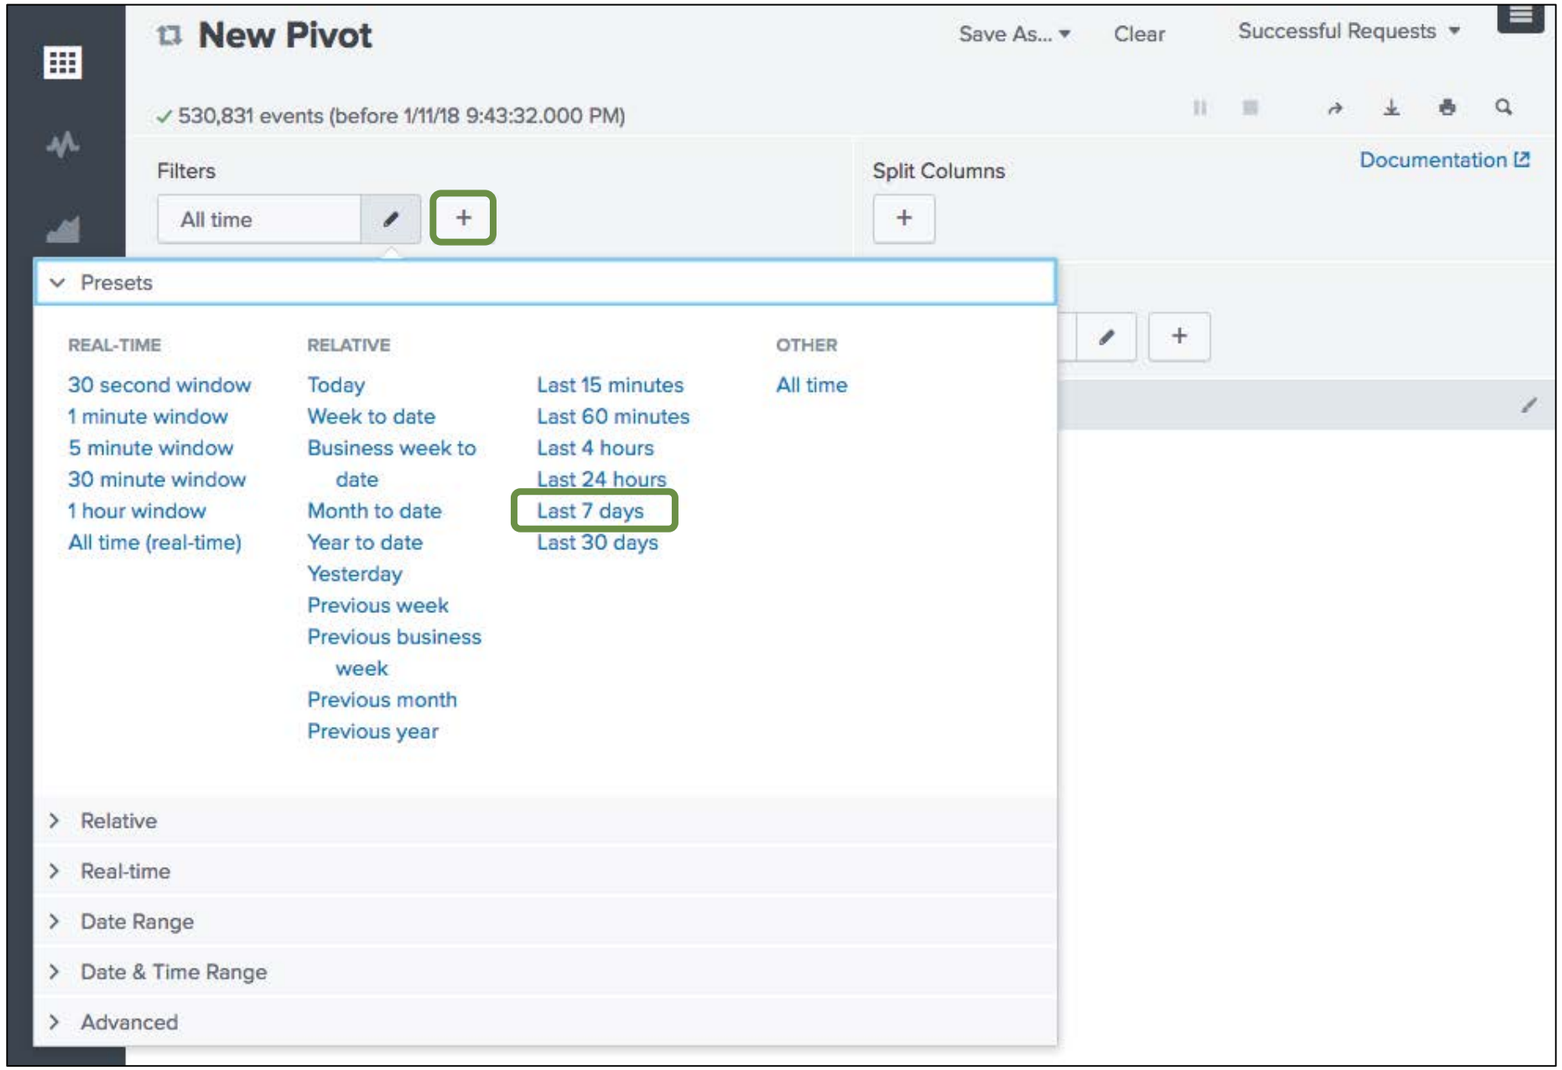

Split Rows

- Click + under **Split Rows** for a list of available attributes to populate the rows
- In this example, the rows are split by the **category** attribute, which lists:
  - Each game category on a separate row
  - A count of successful requests for each game category

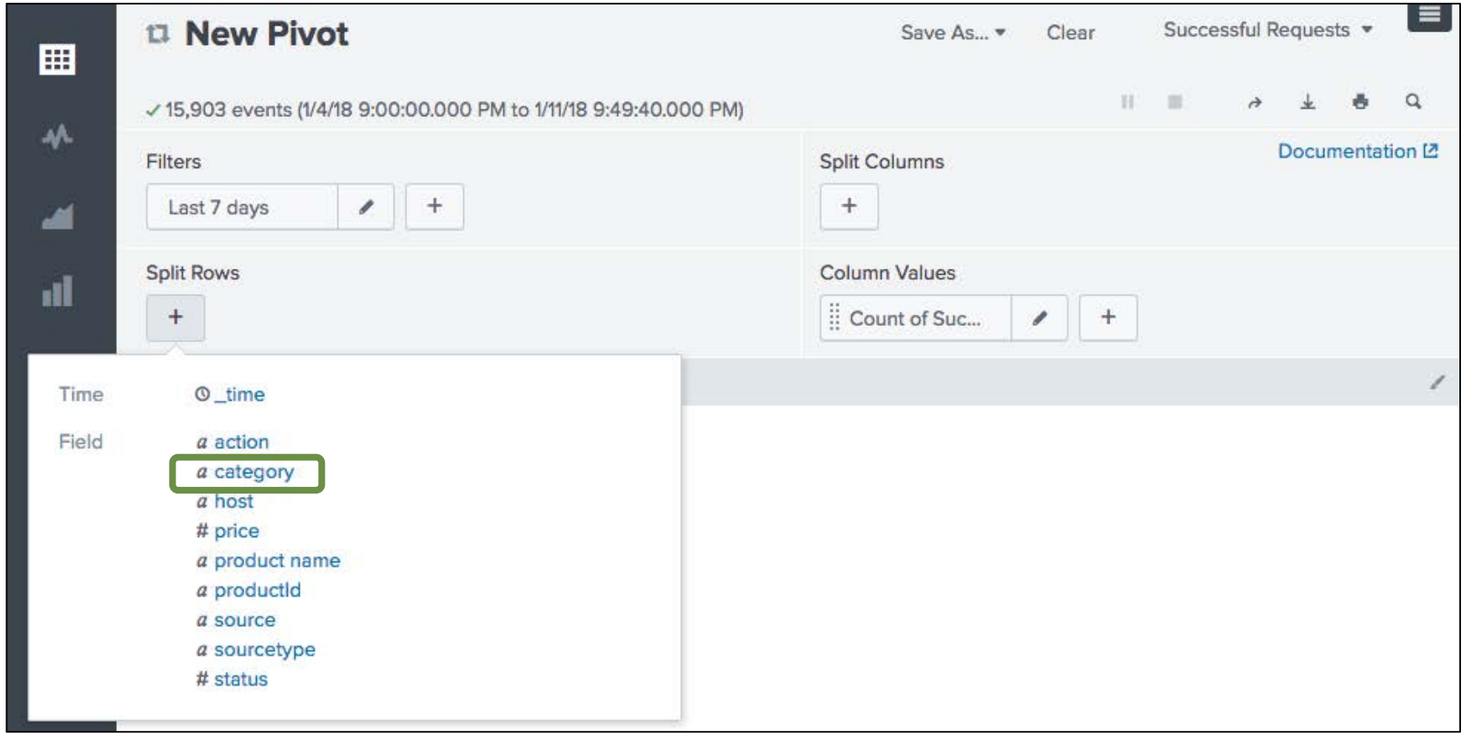

Split Rows (cont.)

- Once selected, you can:
  - Modify the label
  - Change the sort order
    - **Default** – sorts by the field value in ascending order
    - **Ascending** - sorts by the count in ascending order
    - **Descending** – sorts by the count in descending order
  - Define maximum # of rows to display
- Click **Add to Table** to view the results

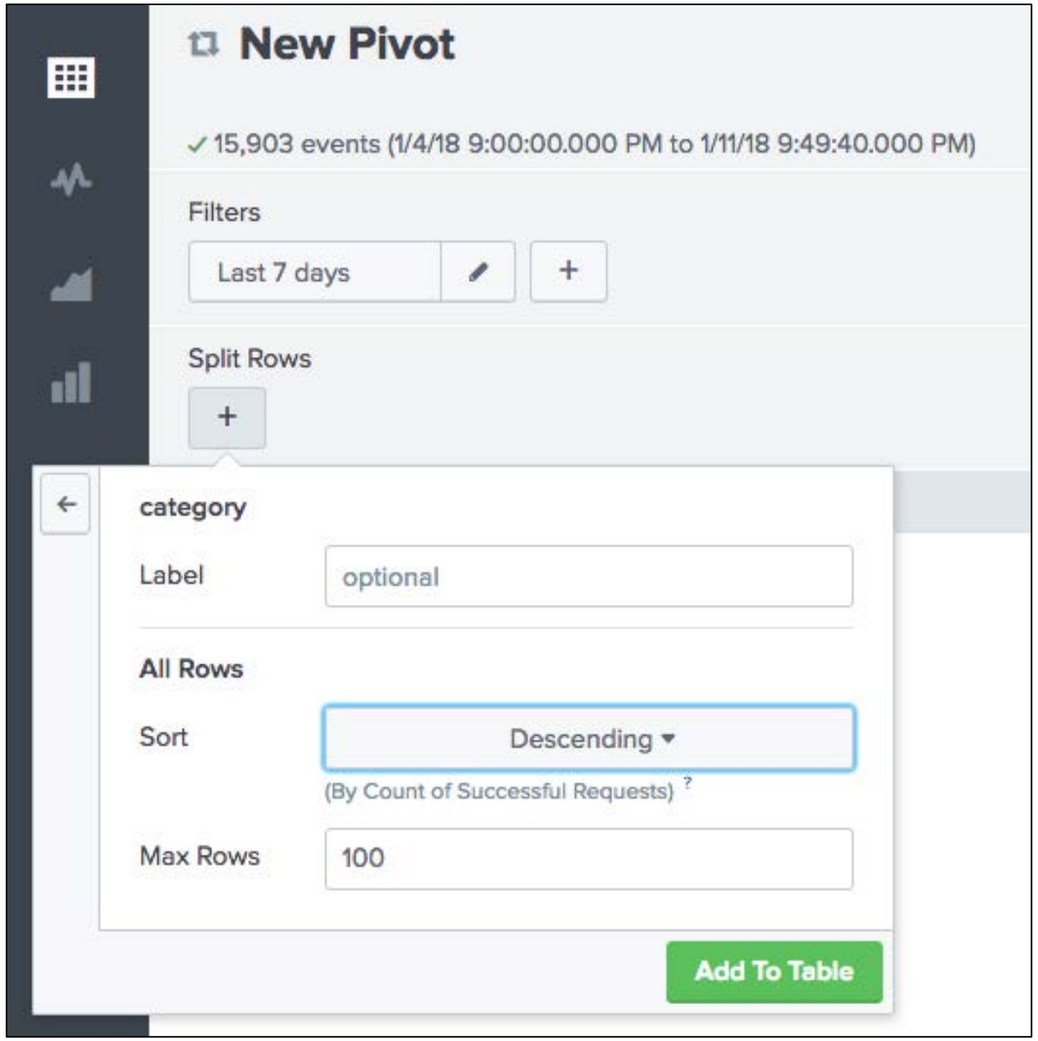

Results

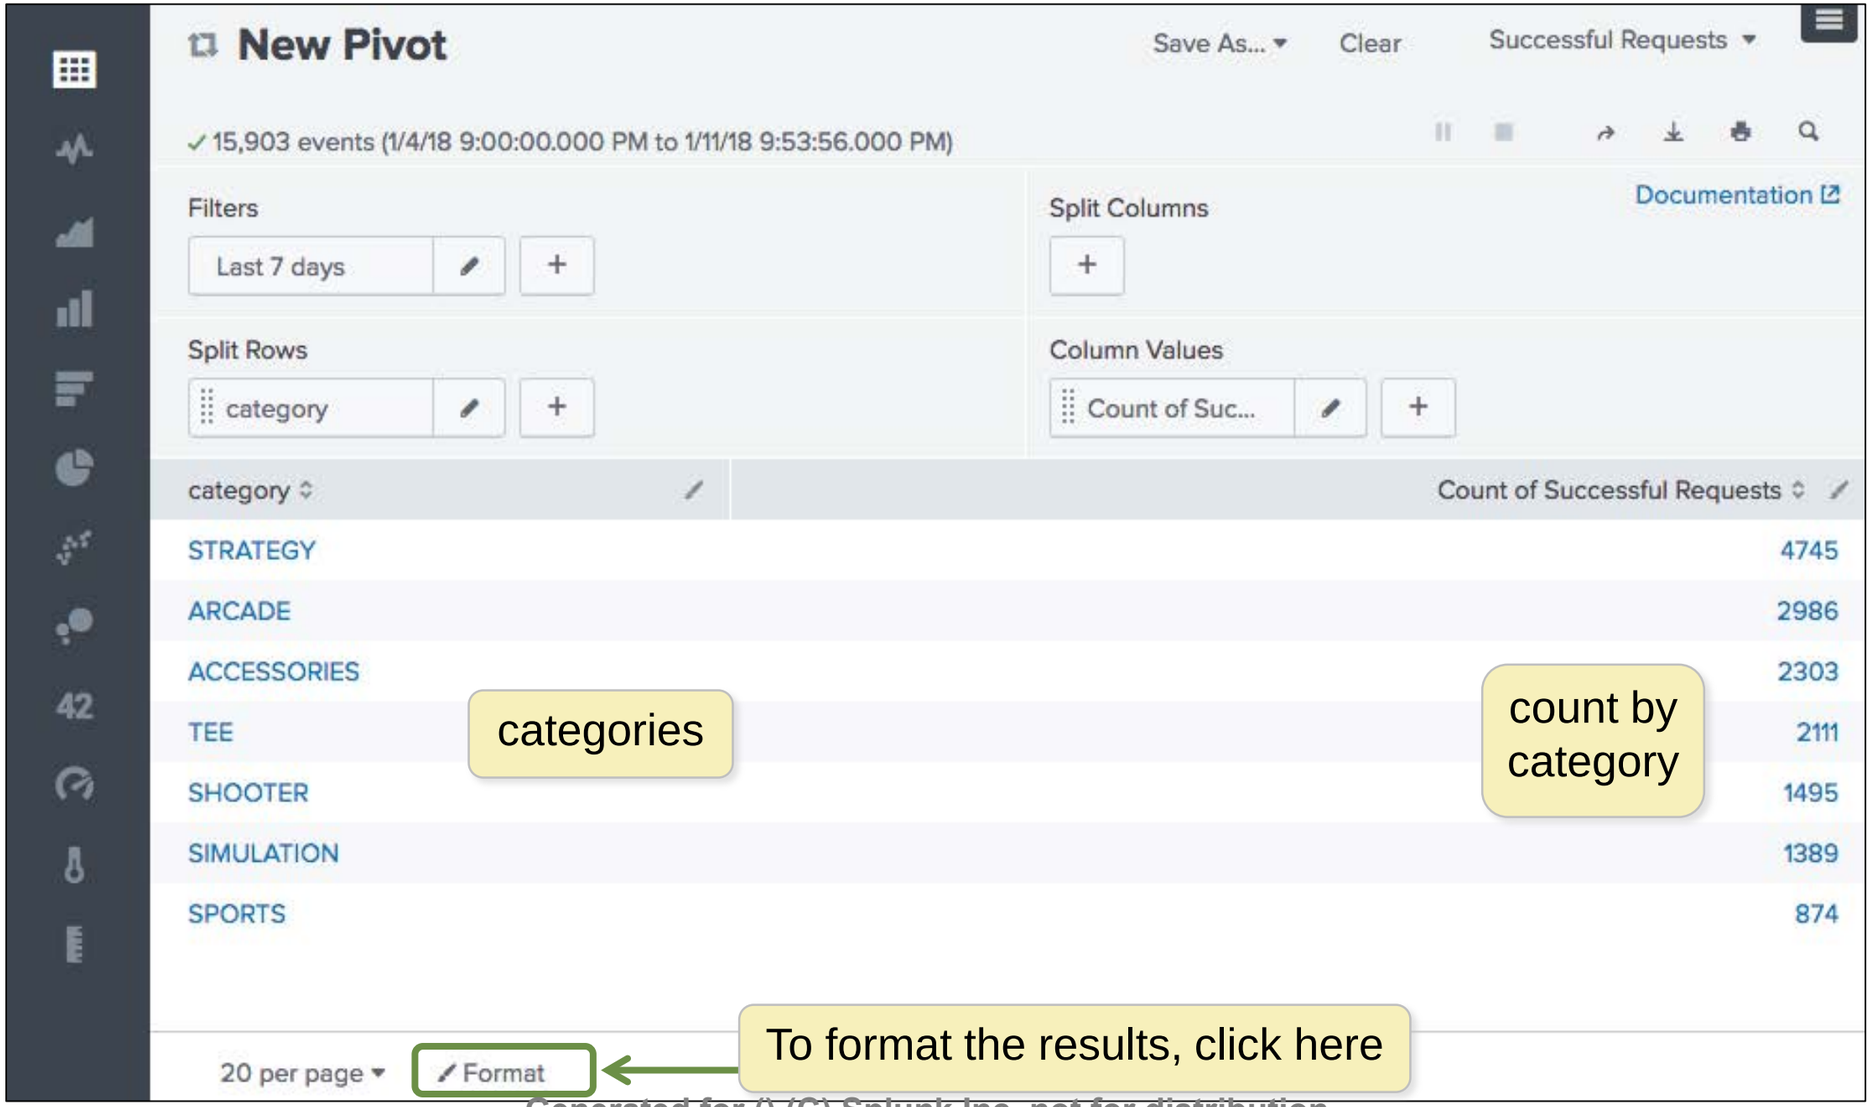

Formatting the Results

For example, to add totals on the Summary tab, click Yes next to Totals

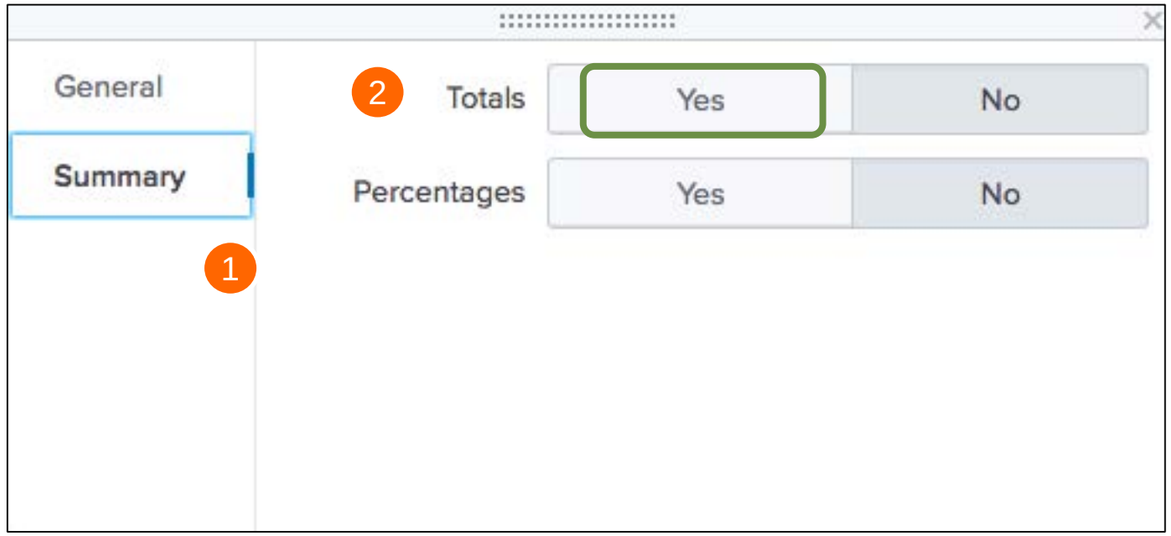

Updated Results (with Total)

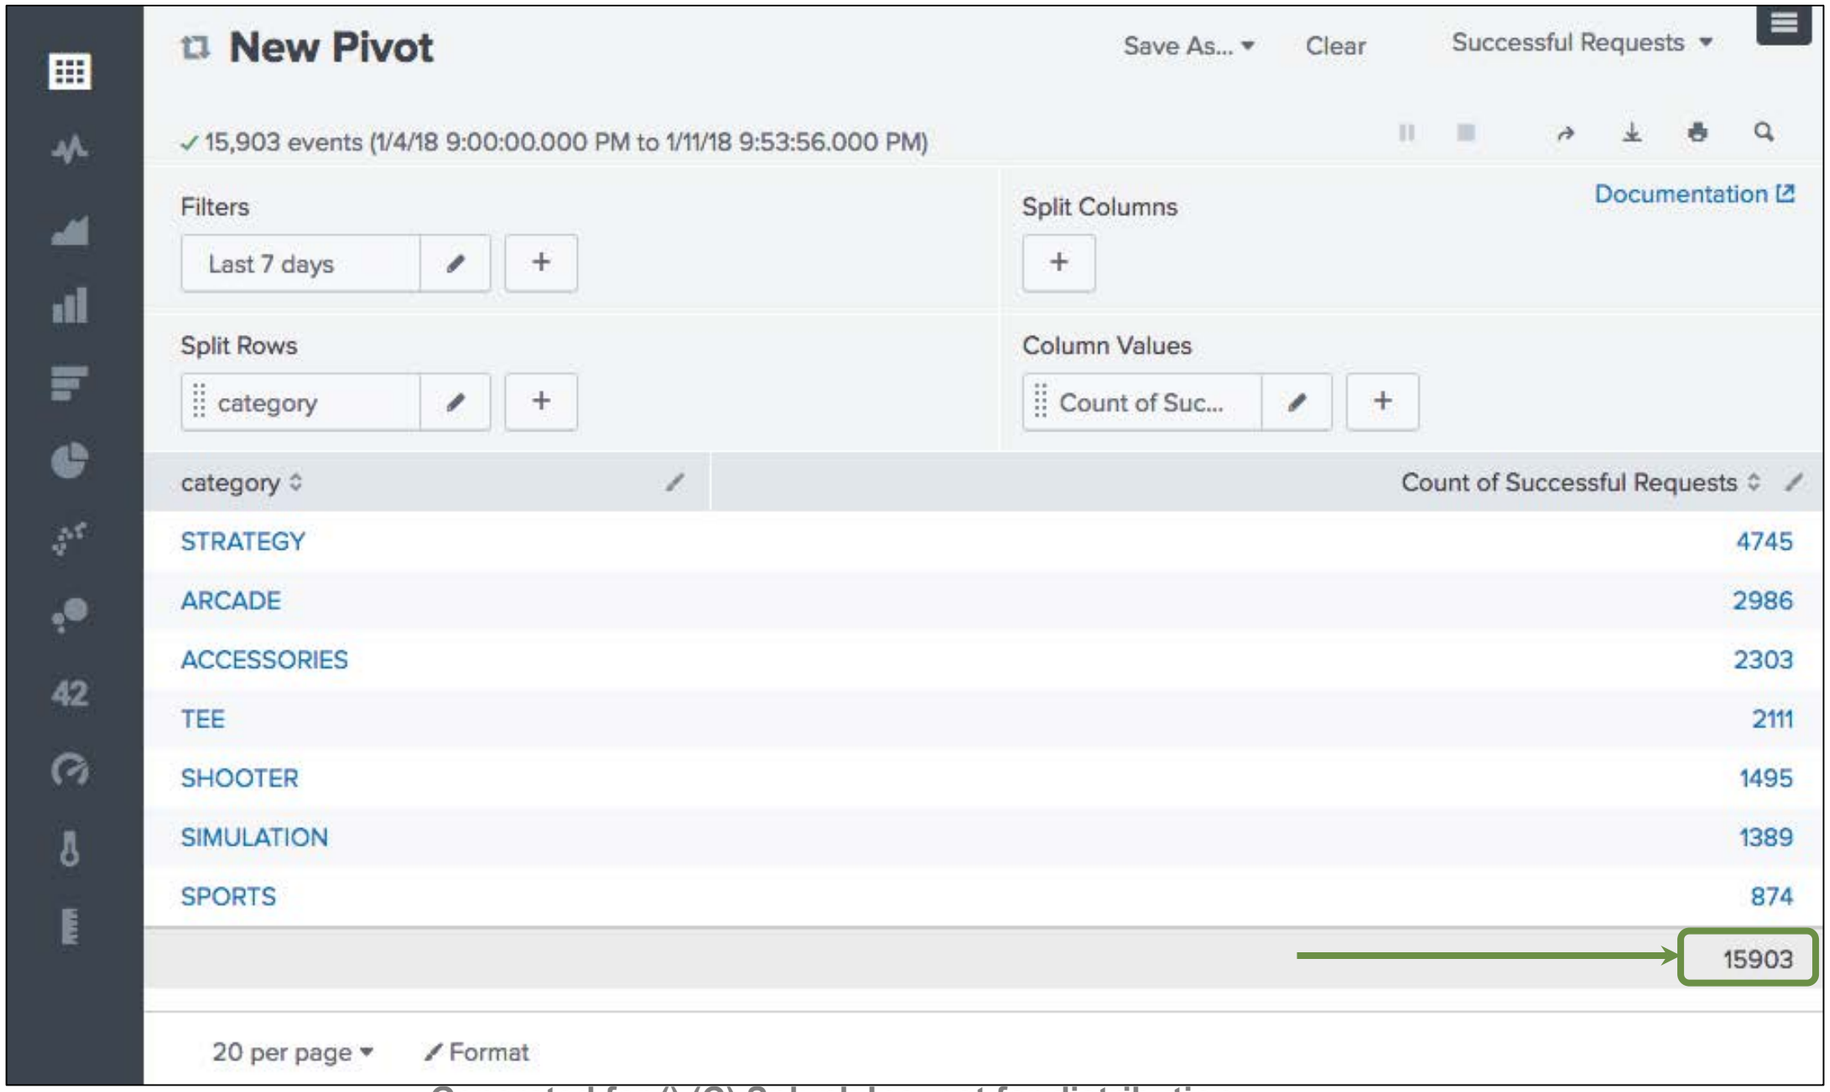

Split Columns

- Click + under **Split Columns** and select the desired split
- Specify the maximum number of columns and whether you want Totals

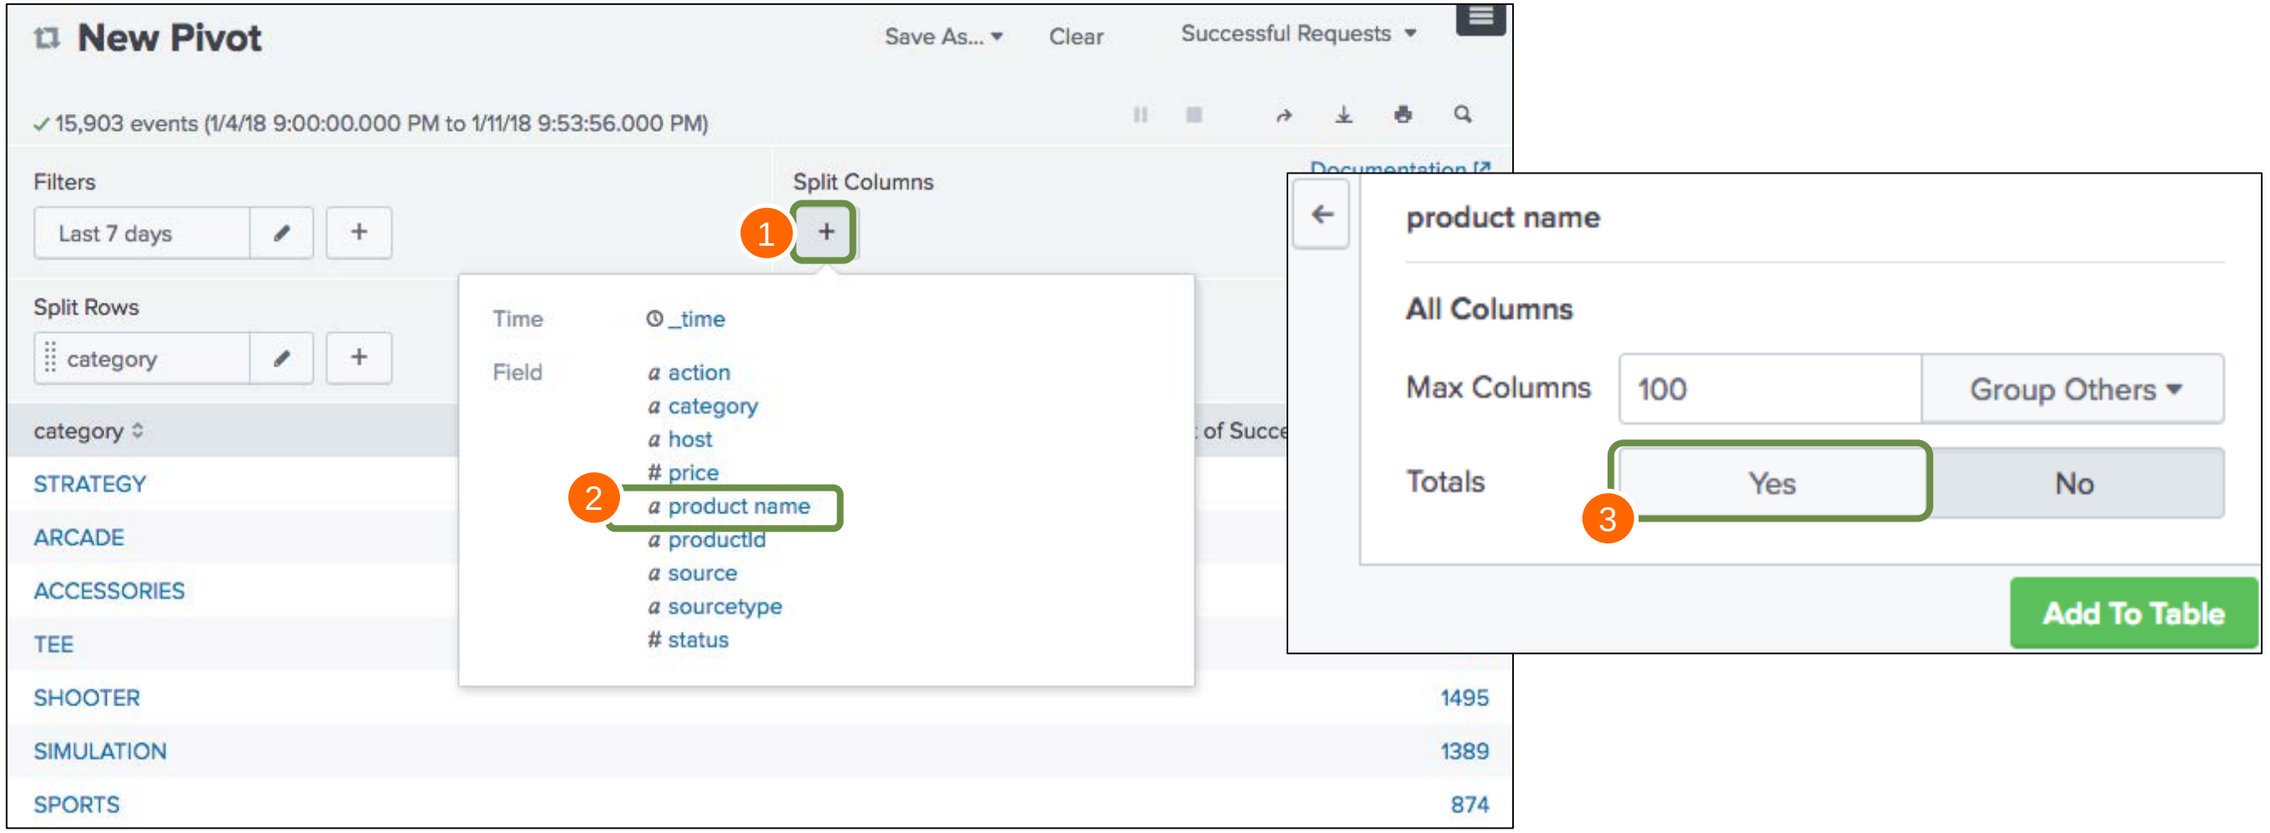

Results

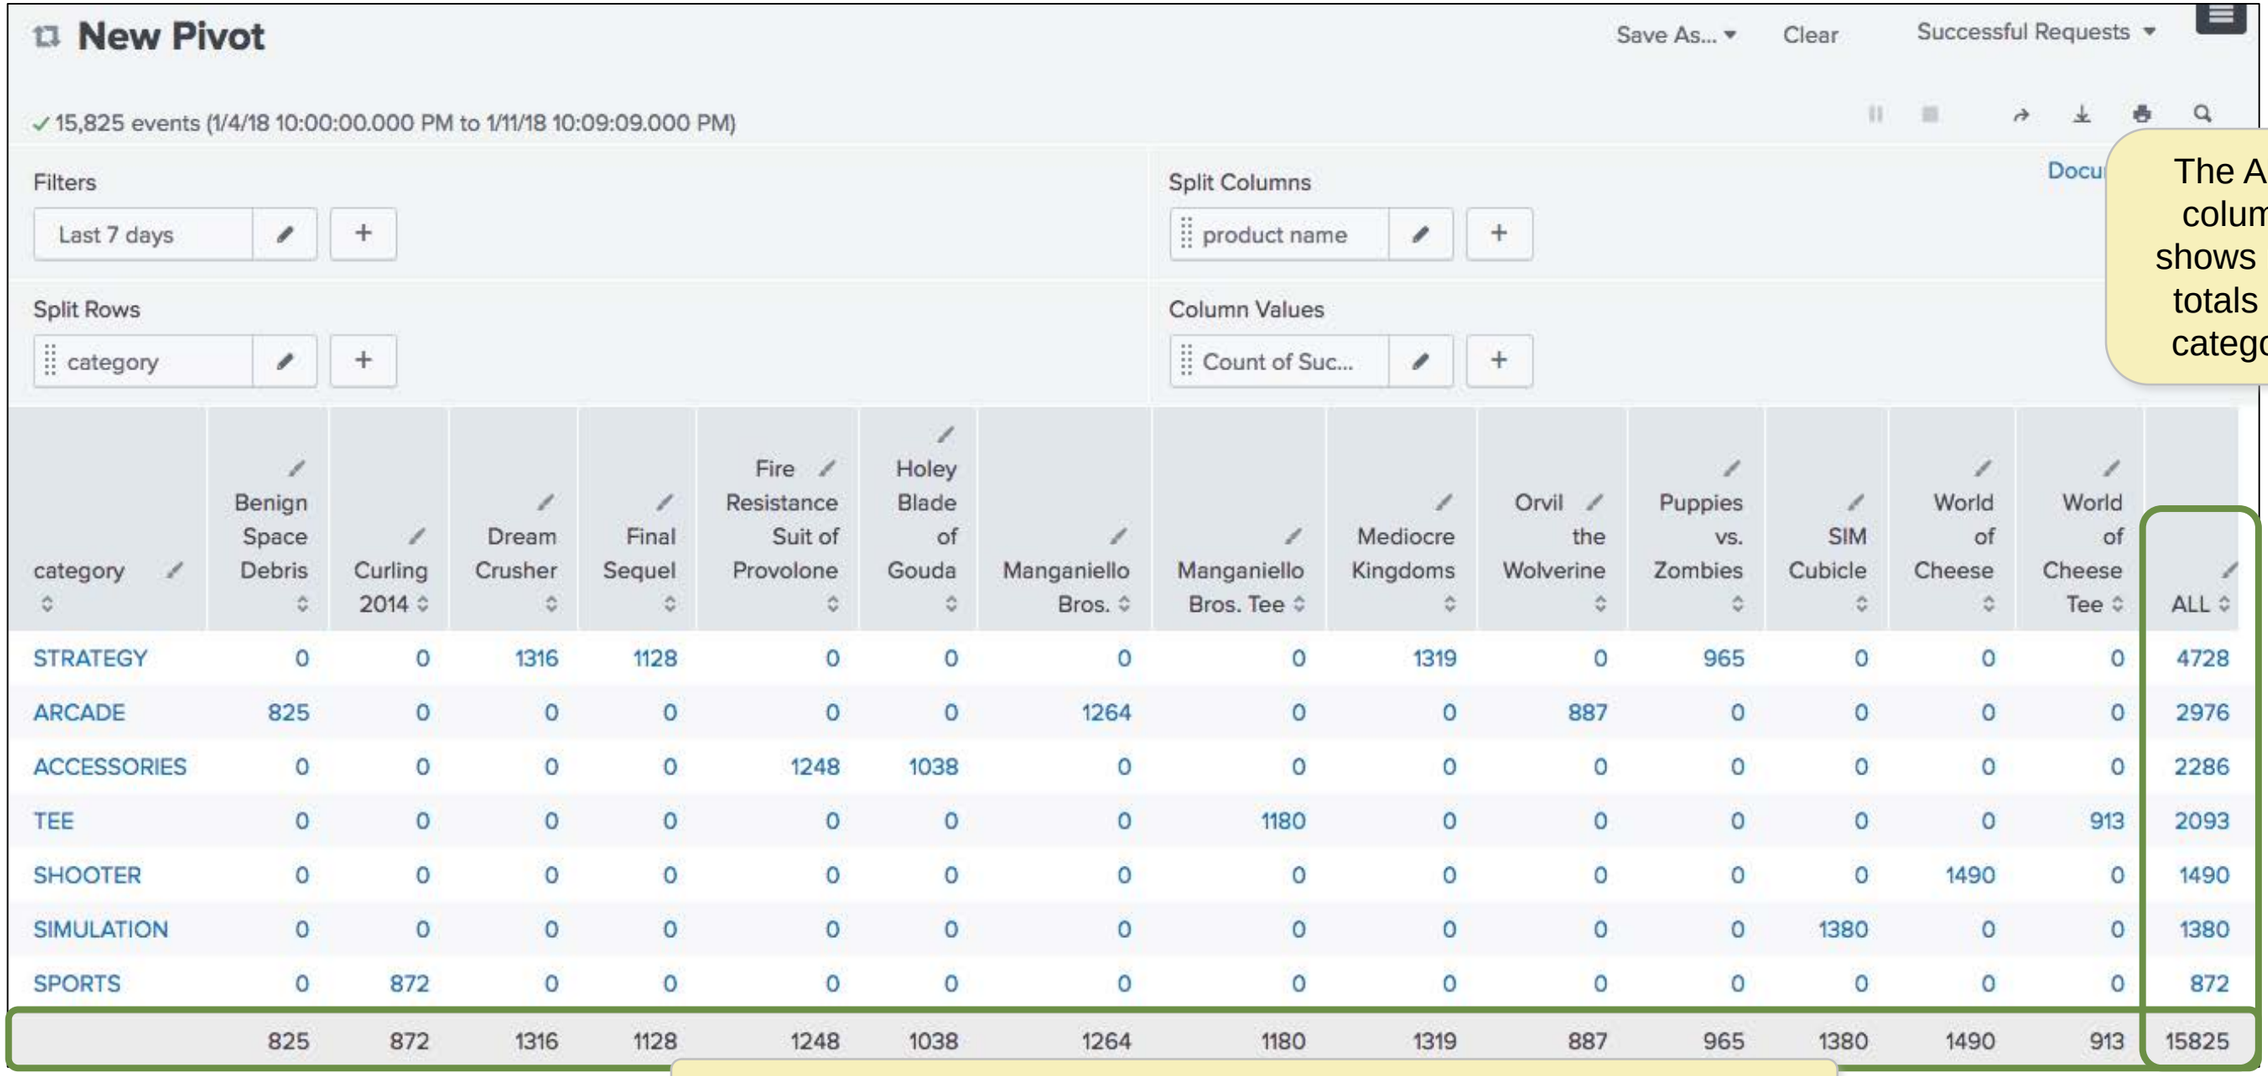

category	Benign Space Debris	Curling 2014	Dream Crusher	Final Sequel	Fire Resistance Suit of Provolone	Holey Blade of Gouda	Manganiello Bros.	Manganiello Bros. Tee	Mediocre Kingdoms	Orvil the Wolverine	Puppies vs. Zombies	SIM Cubicle	World of Cheese	World of Cheese Tee	ALL
STRATEGY	0	0	1316	1128	0	0	0	0	1319	0	965	0	0	0	4728
ARCADE	825	0	0	0	0	0	1264	0	0	887	0	0	0	0	2976
ACCESSORIES	0	0	0	0	1248	1038	0	0	0	0	0	0	0	0	2286
TEE	0	0	0	0	0	0	0	1180	0	0	0	0	0	913	2093
SHOOTER	0	0	0	0	0	0	0	0	0	0	0	0	1490	0	1490
SIMULATION	0	0	0	0	0	0	0	0	0	0	0	1380	0	0	1380
SPORTS	0	872	0	0	0	0	0	0	0	0	0	0	0	0	872
	825	872	1316	1128	1248	1038	1264	1180	1319	887	965	1380	1490	913	15825

The ALL column shows row totals by category

The bottom (bolded) row shows column totals by product name

Add Additional Filters

- You can refine a pivot by filtering on key/value pairs
  - Think of 'split by' as rows and columns as the fields to display
  - Think of filters as a field=value inclusion, exclusion or specific condition to apply to the search (=, <, >, !=, *)
- In the example, the pivot is filtered to exclude events from the ACCESSORIES category

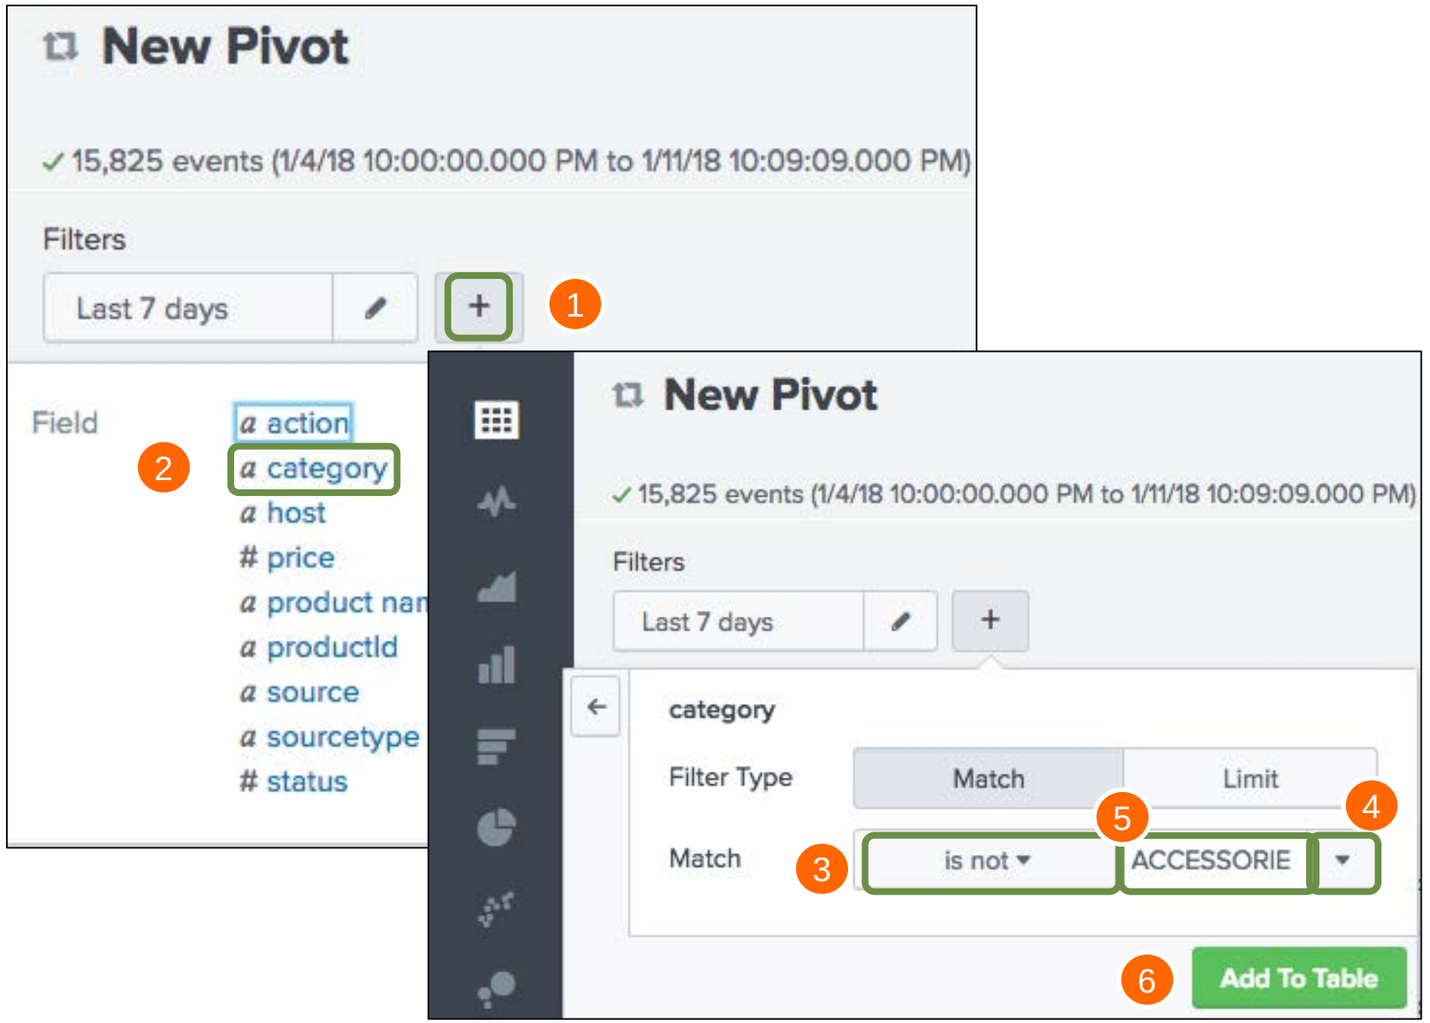

Filtered Pivot

- The ACCESSORIES category is filtered out
- All the other categories remain

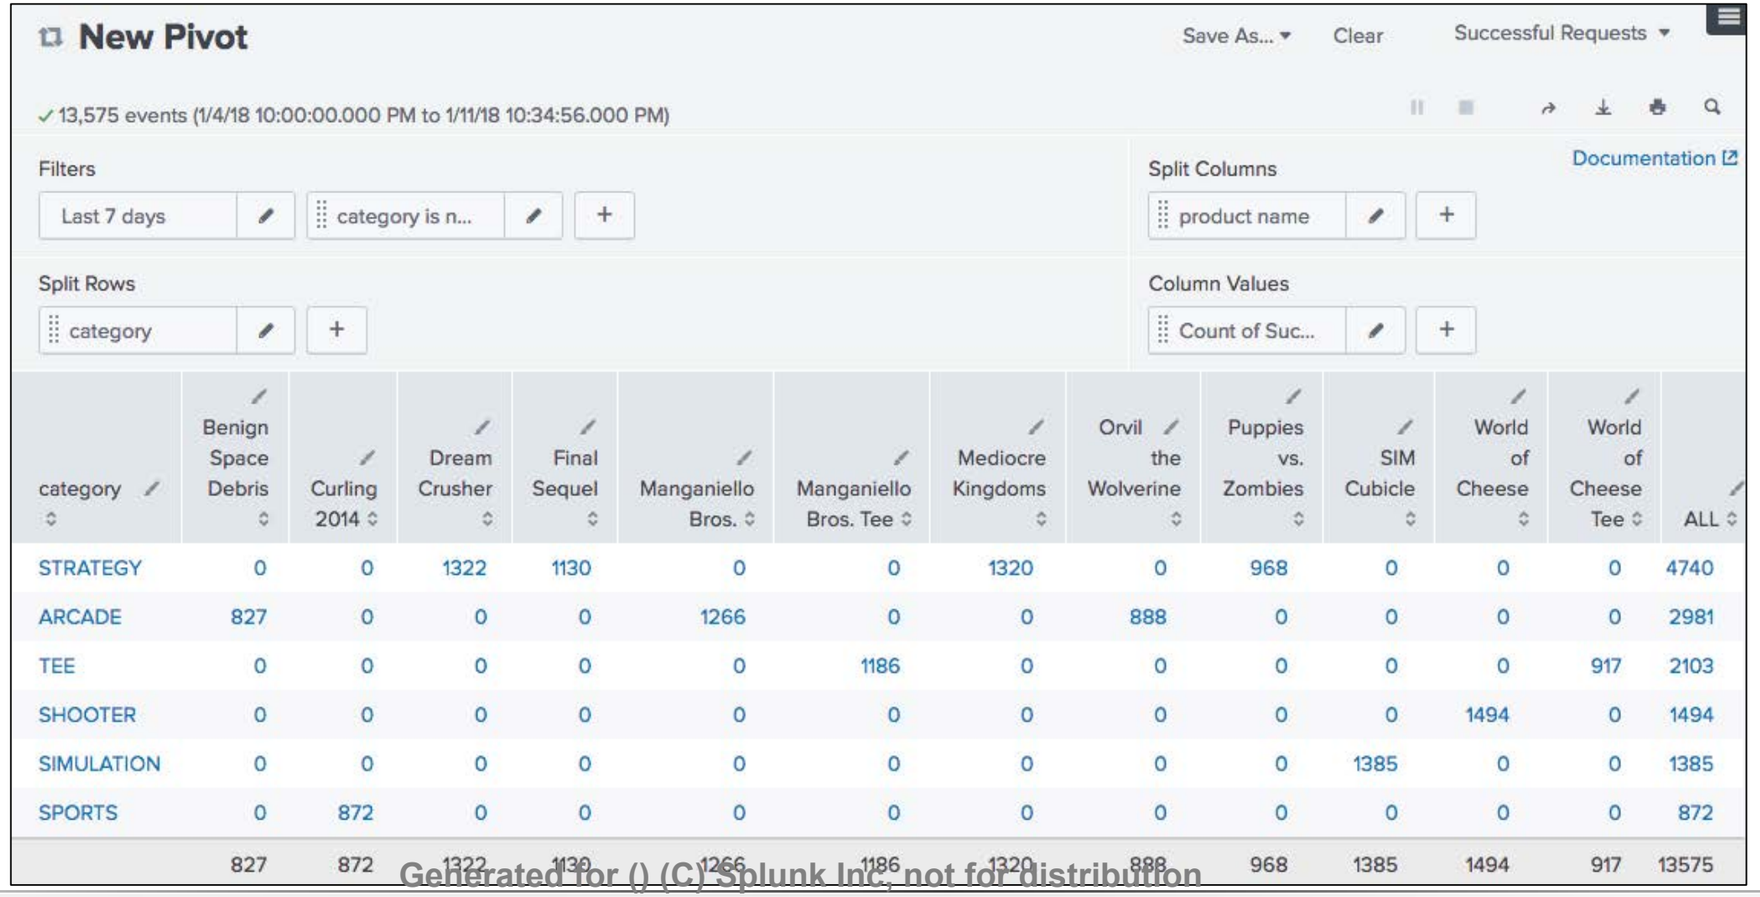

category	Benign Space Debris	Curling 2014	Dream Crusher	Final Sequel	Manganiello Bros.	Manganiello Bros. Tee	Mediocre Kingdoms	Orvil the Wolverine	Puppies vs. Zombies	SIM Cubicle	World of Cheese	World of Cheese Tee	ALL
STRATEGY	0	0	1322	1130	0	0	1320	0	968	0	0	0	4740
ARCADE	827	0	0	0	1266	0	0	888	0	0	0	0	2981
TEE	0	0	0	0	0	1186	0	0	0	0	0	917	2103
SHOOTER	0	0	0	0	0	0	0	0	0	0	1494	0	1494
SIMULATION	0	0	0	0	0	0	0	0	0	1385	0	0	1385
SPORTS	0	872	0	0	0	0	0	0	0	0	0	0	872
	827	872	1322	1130	1266	1186	1320	888	968	1385	1494	917	13575

Select a Visualization Format

You can display your pivot as a table or a visualization, such as a column chart

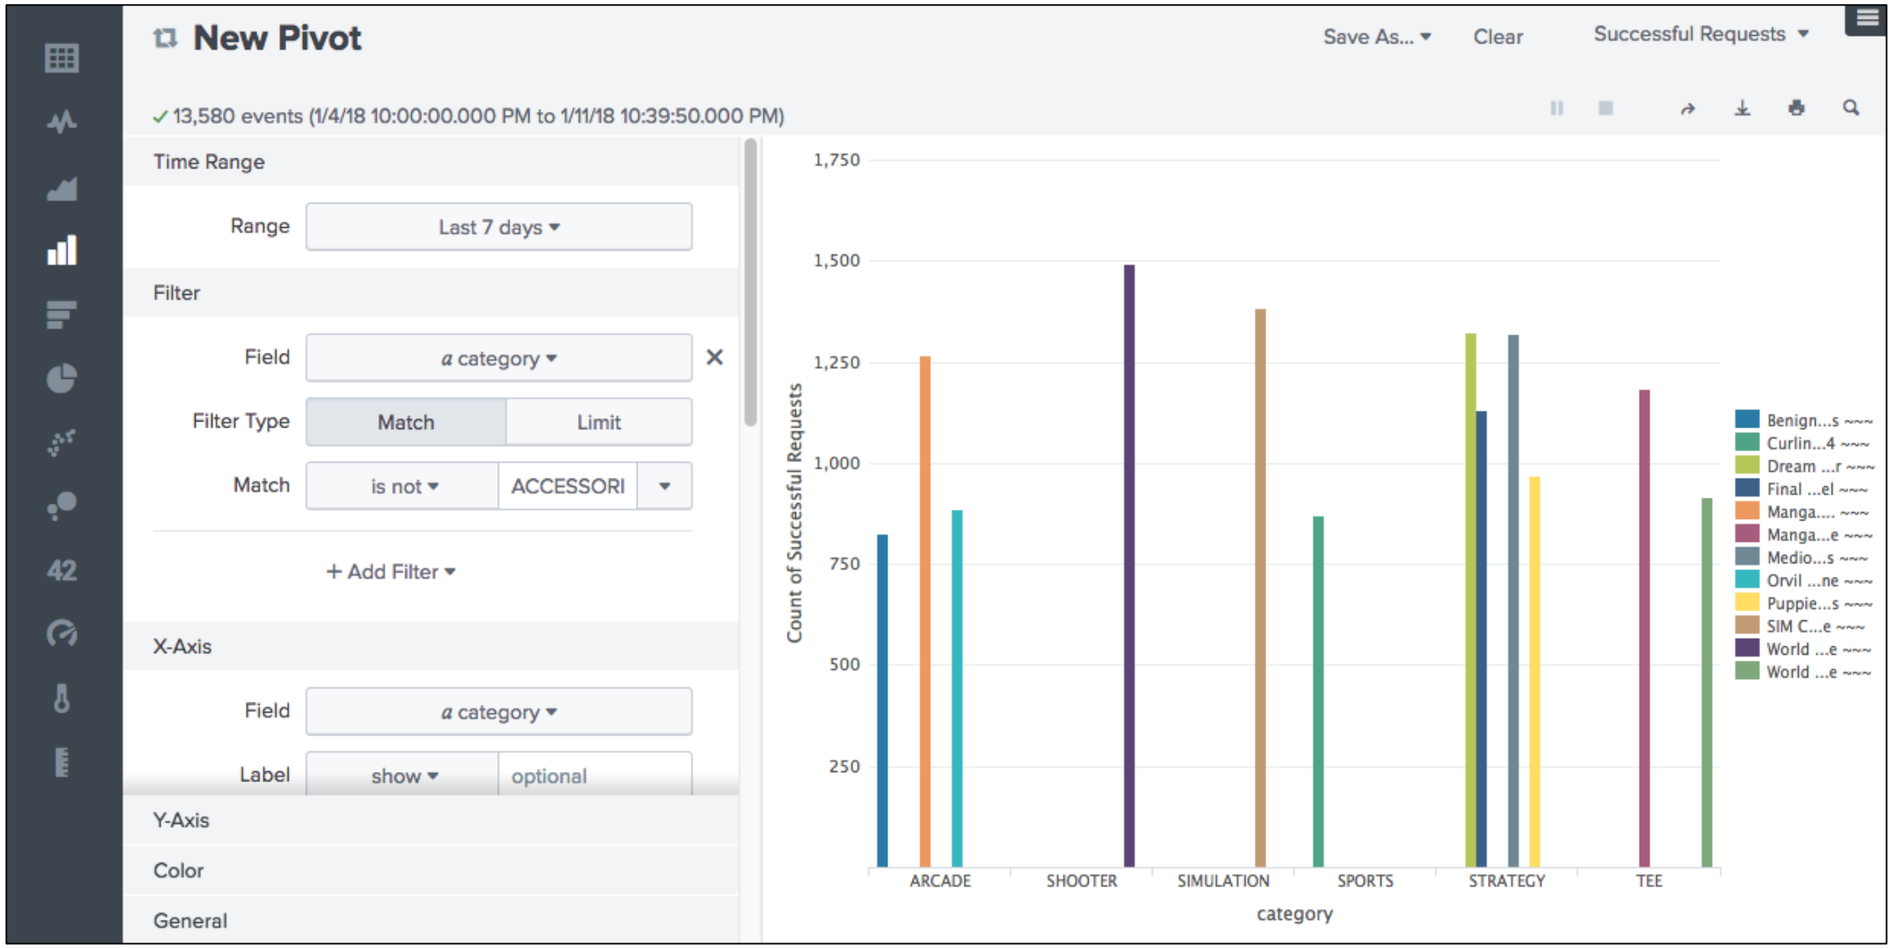

Modify Visualization Settings

- When a visualization control is selected, panels appear that let you configure its settings
- In this example:
  - The results for each category are broken down by **product_name**
  - The stack mode is set to stacked 

Some of the settings for the column chart include:
- Field to use to further breakdown the data
- Limit on the number of series to be charted
- Specify labels
- Set sort order
- Set stack mode

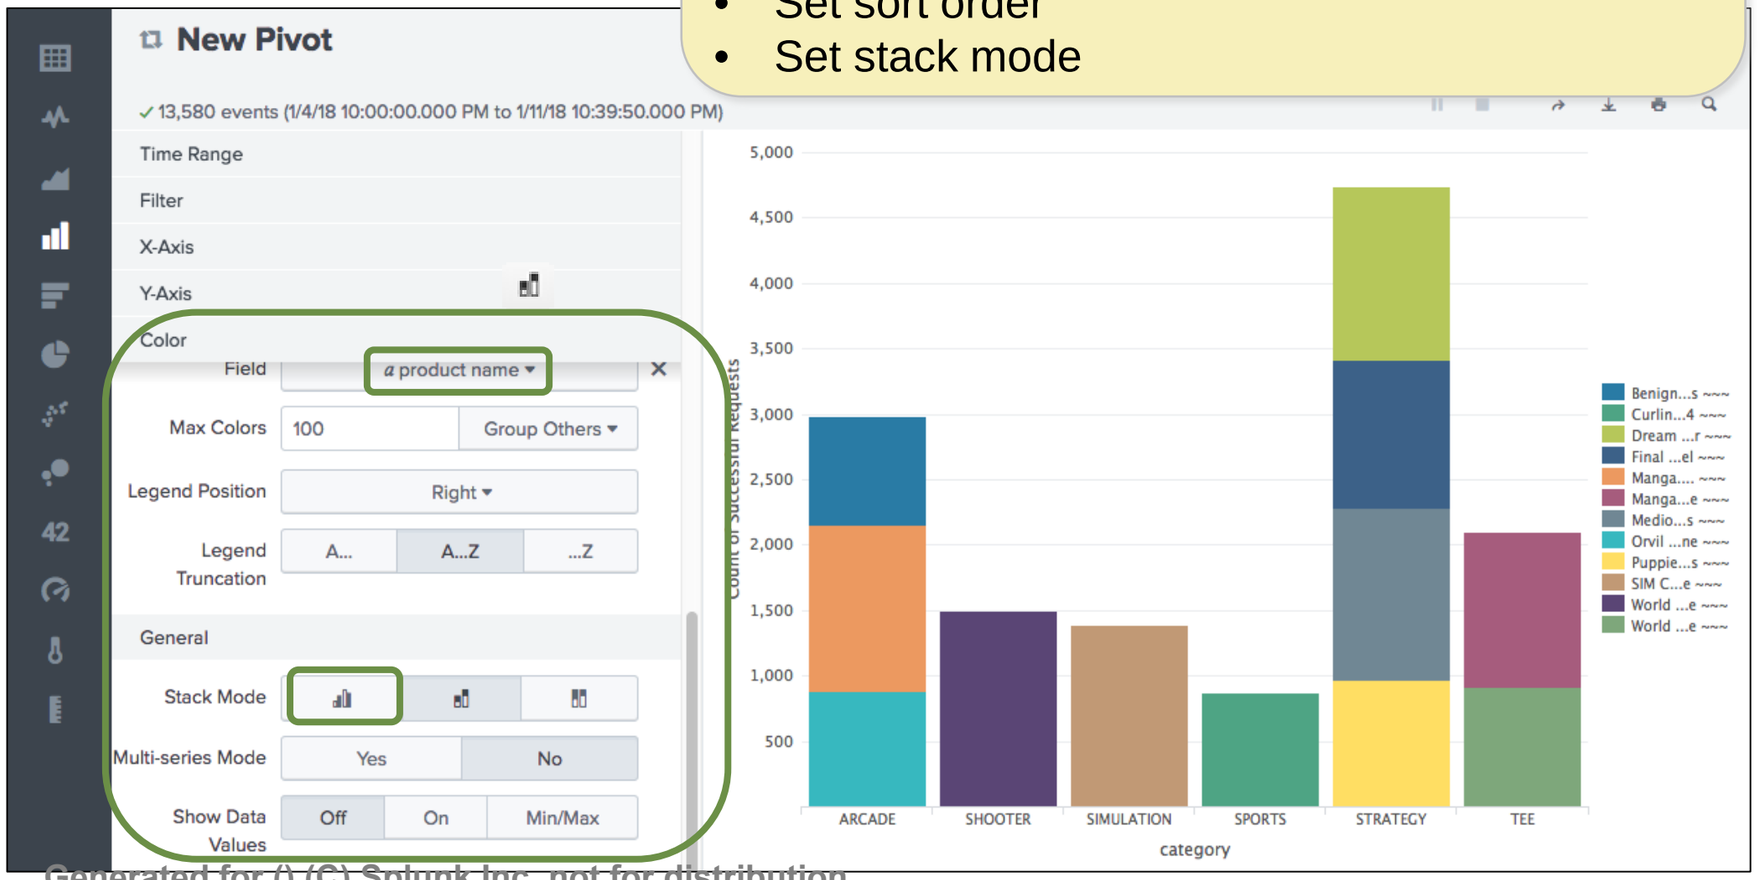

Saving a Pivot

- Pivots can be saved as reports
  - You can choose to include a Time Range Picker in the report to allow people who run it to change the time range (default is Yes)

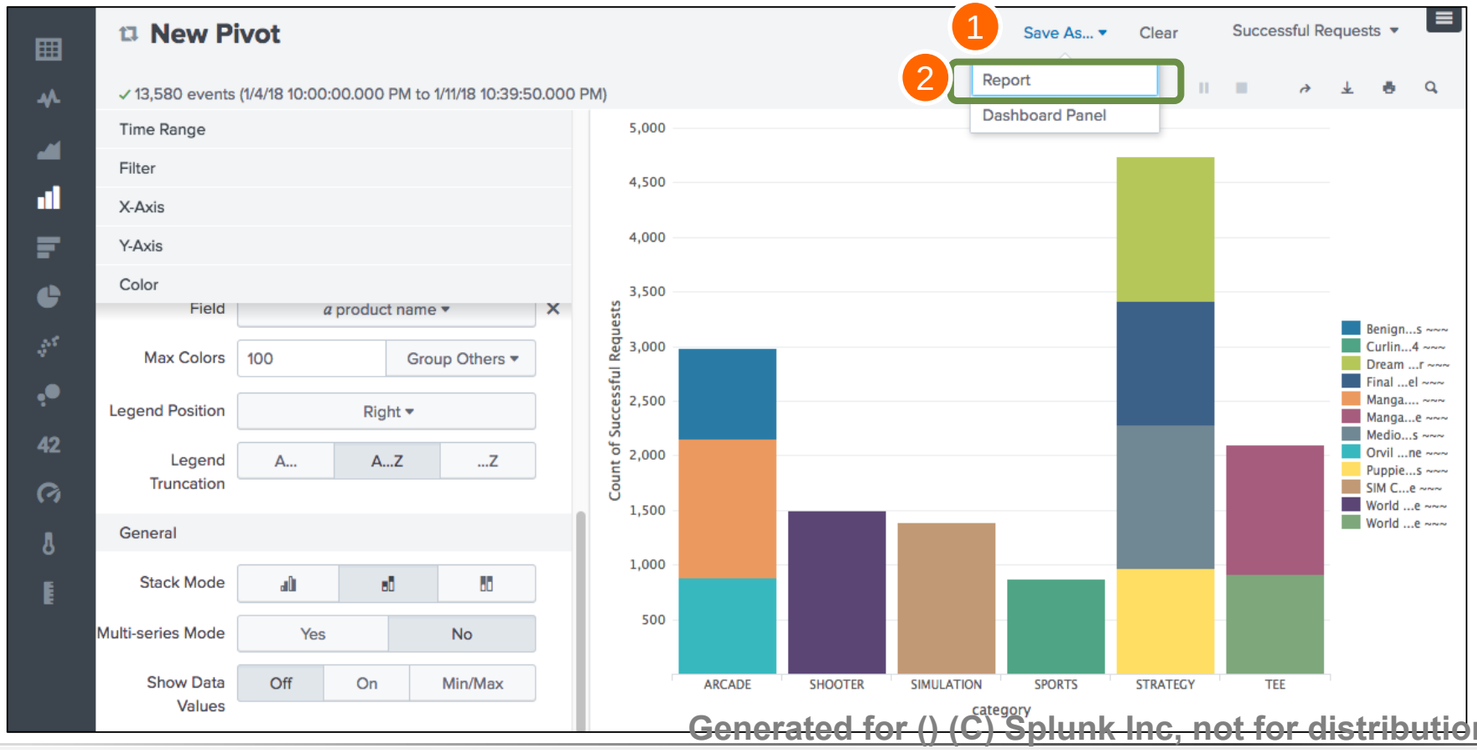

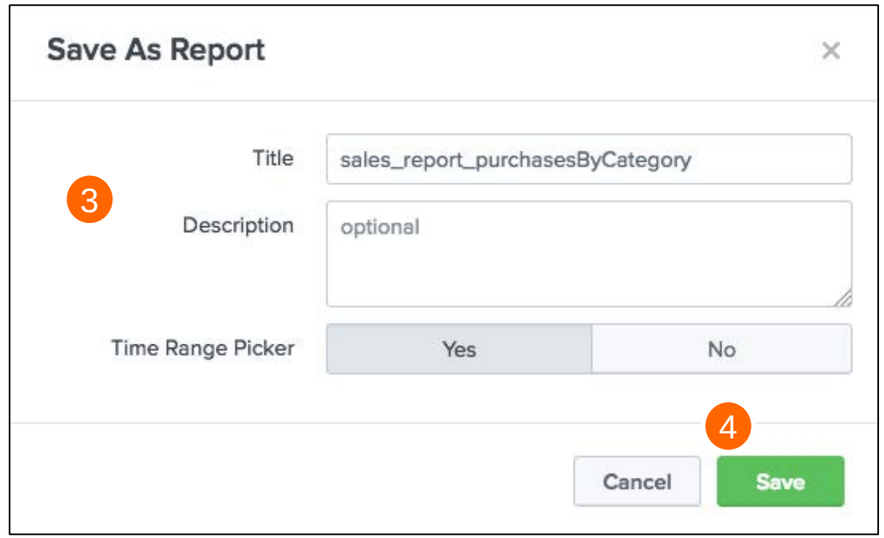

Saving a Pivot (cont.)

When you click **View**, the report is displayed with a Time Range Picker (if that's the option you chose)

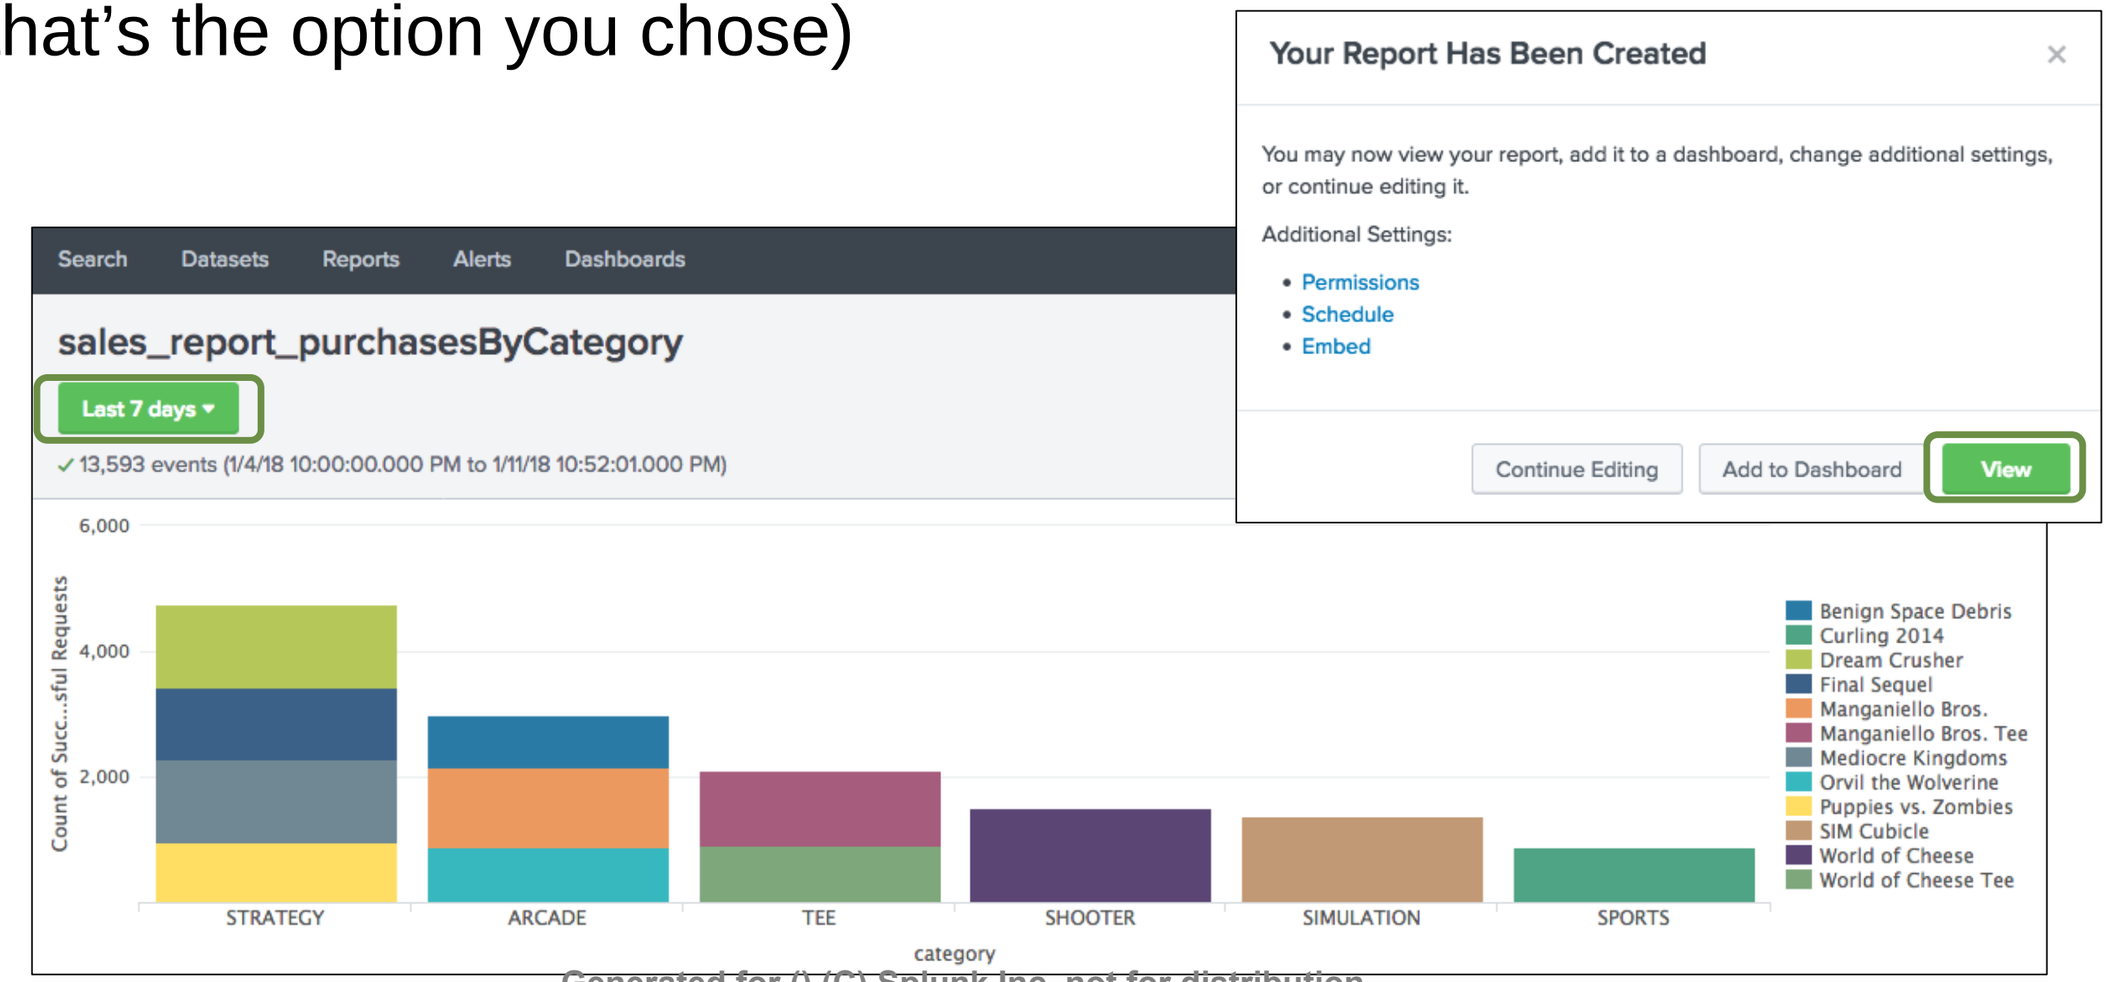

Mouse Actions

- Mouse over an object to reveal its details
- If drilldown is enabled, it is possible to click on the object to expose the underlying search

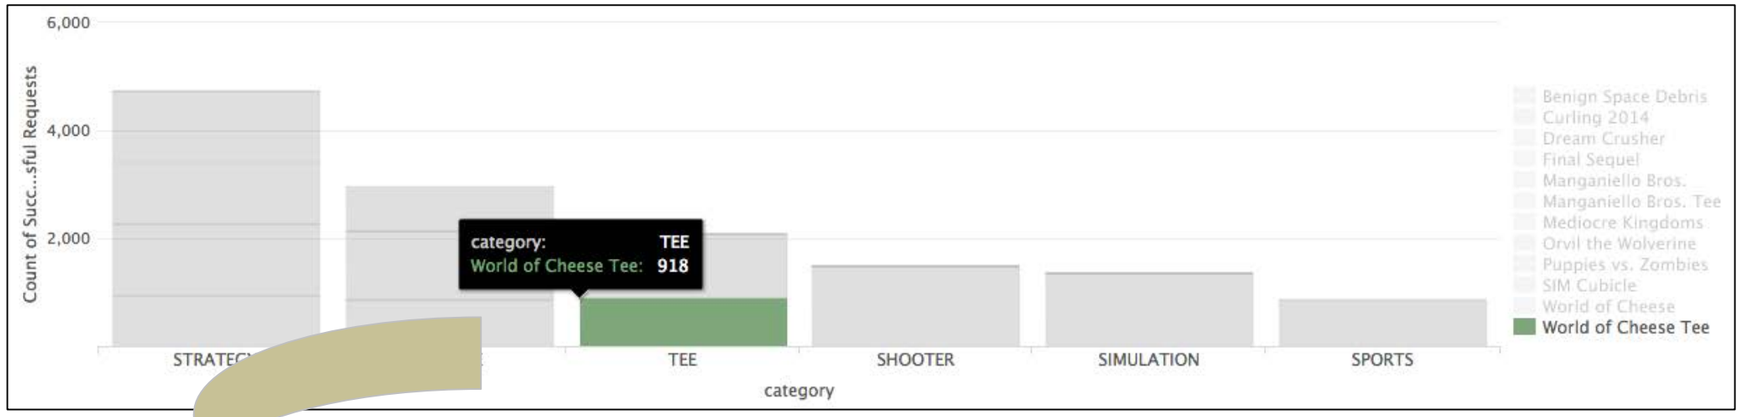

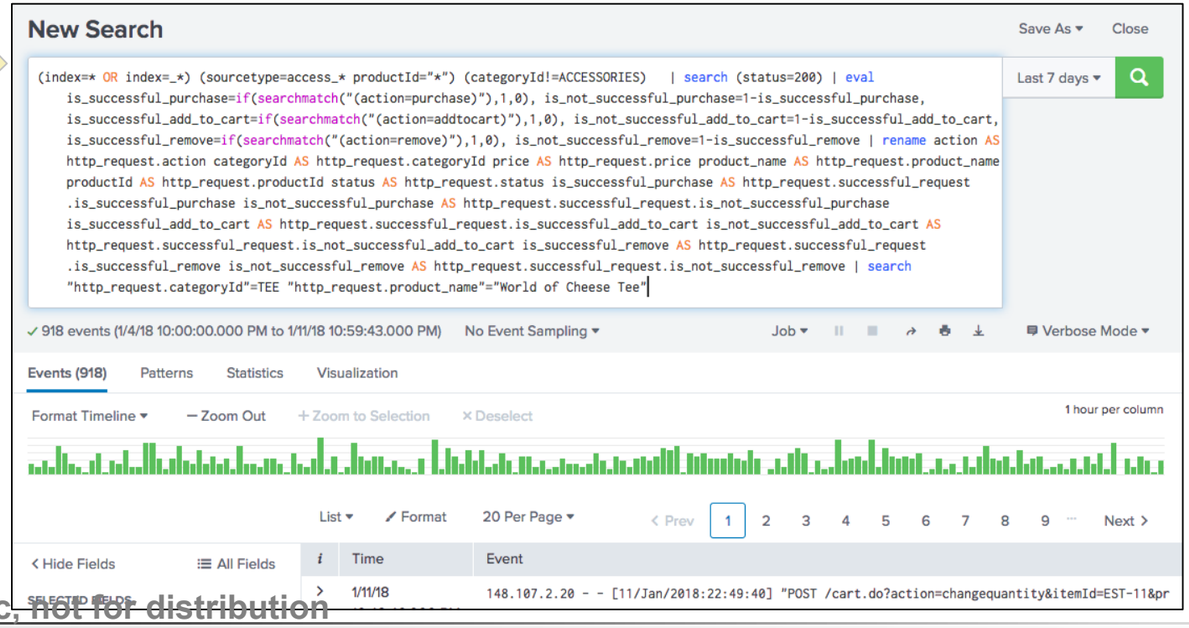

Note

The search generated by drilldown may be more detailed than your original search. However, it produces the same results.

Instant Pivot Overview

- Instant pivot allows you to utilize the pivot tool without a preexisting data model
  - Instant pivot creates an underlying data model utilizing the search criteria entered during the initial search
- To create an Instant Pivot
  1. Execute a search (search criteria only, no search commands)
  2. Click the **Statistics** or **Visualization** tab
  3. Click the **Pivot** icon
  4. Select the fields to be included in the data model object
  5. Create the pivot (table or chart)

Open Instant Pivot

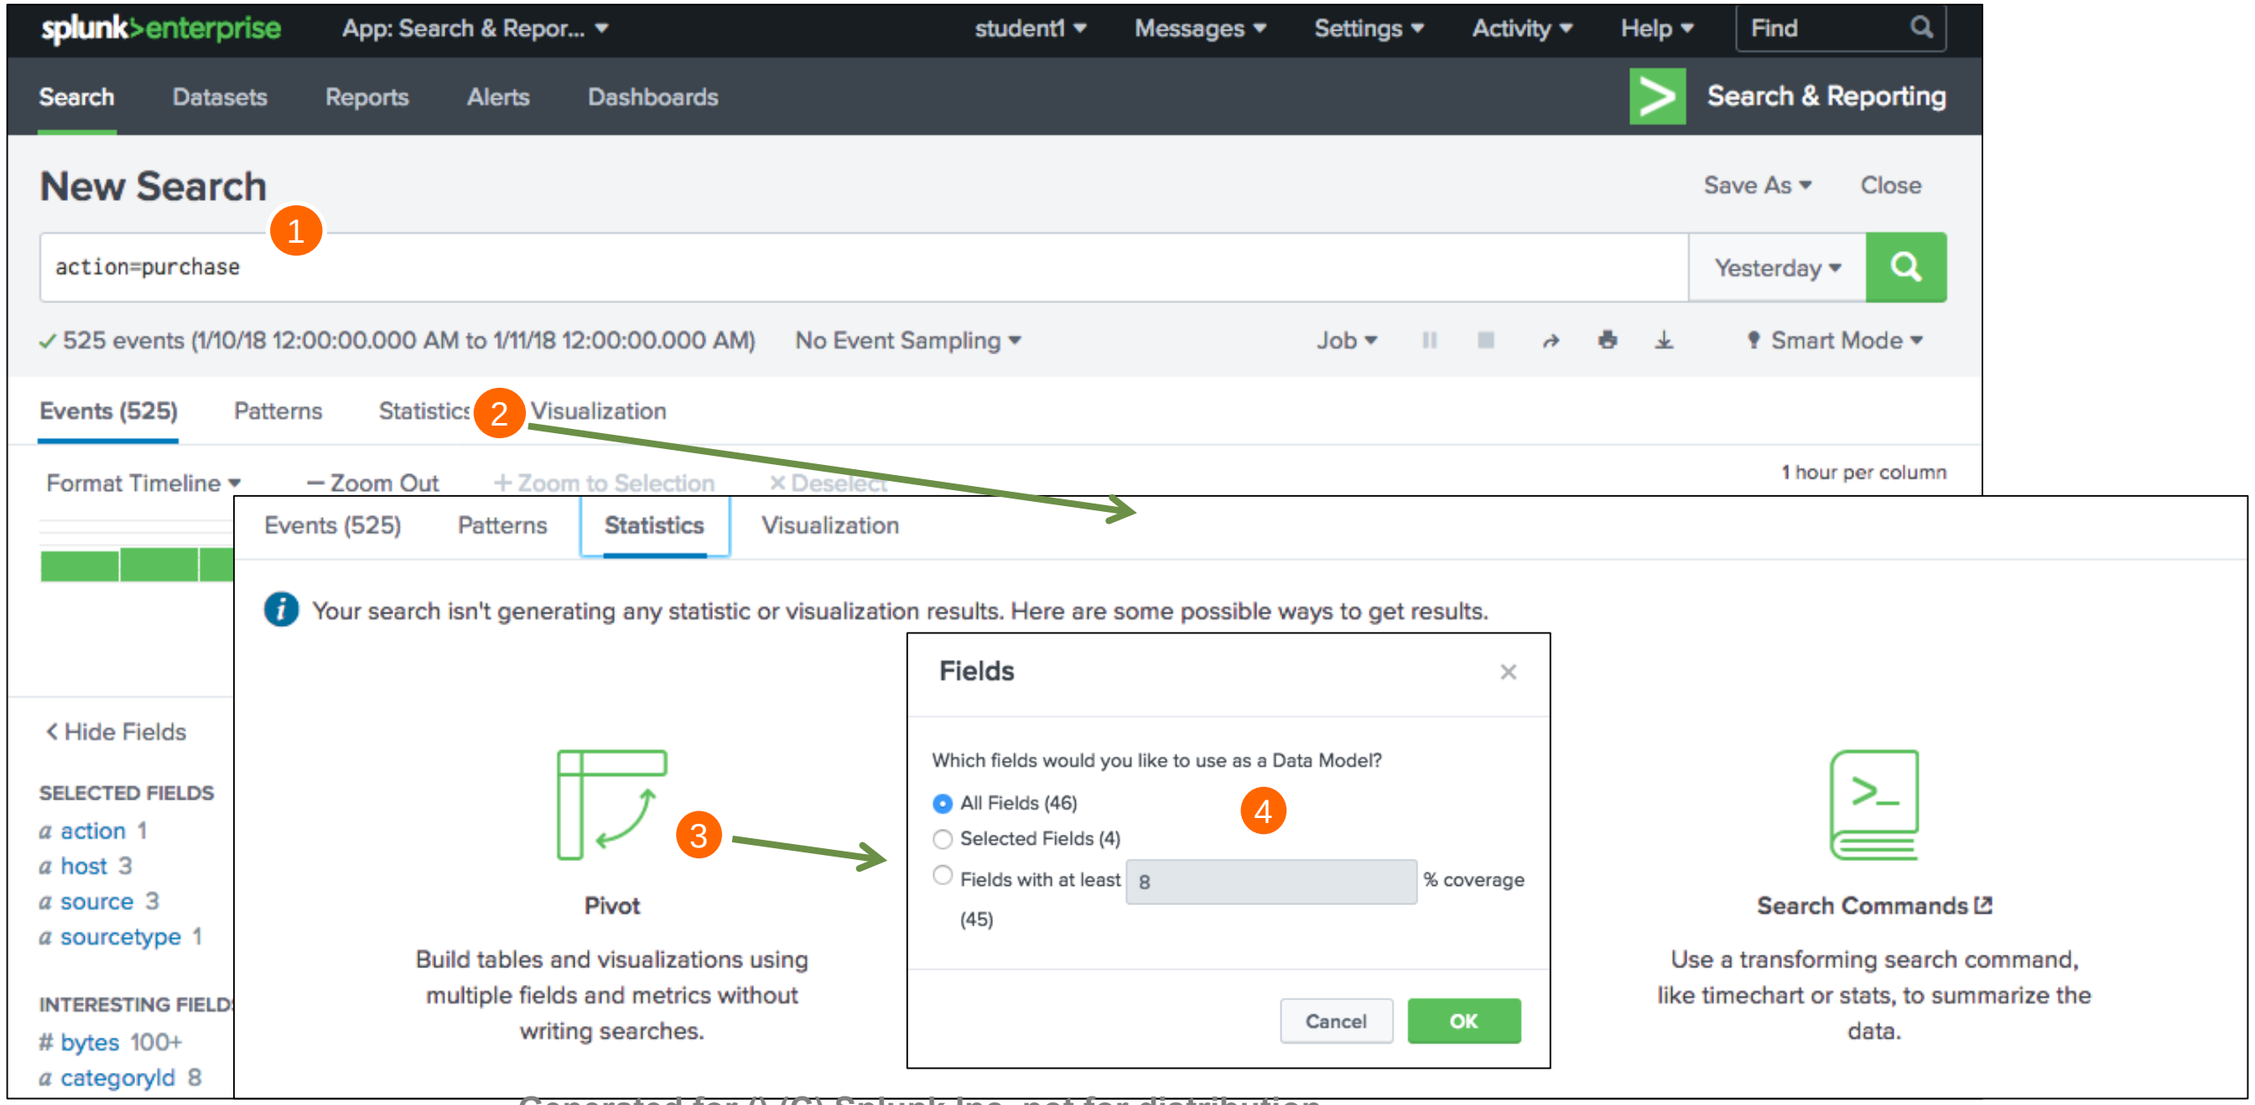

Saving a Pivot as a Report

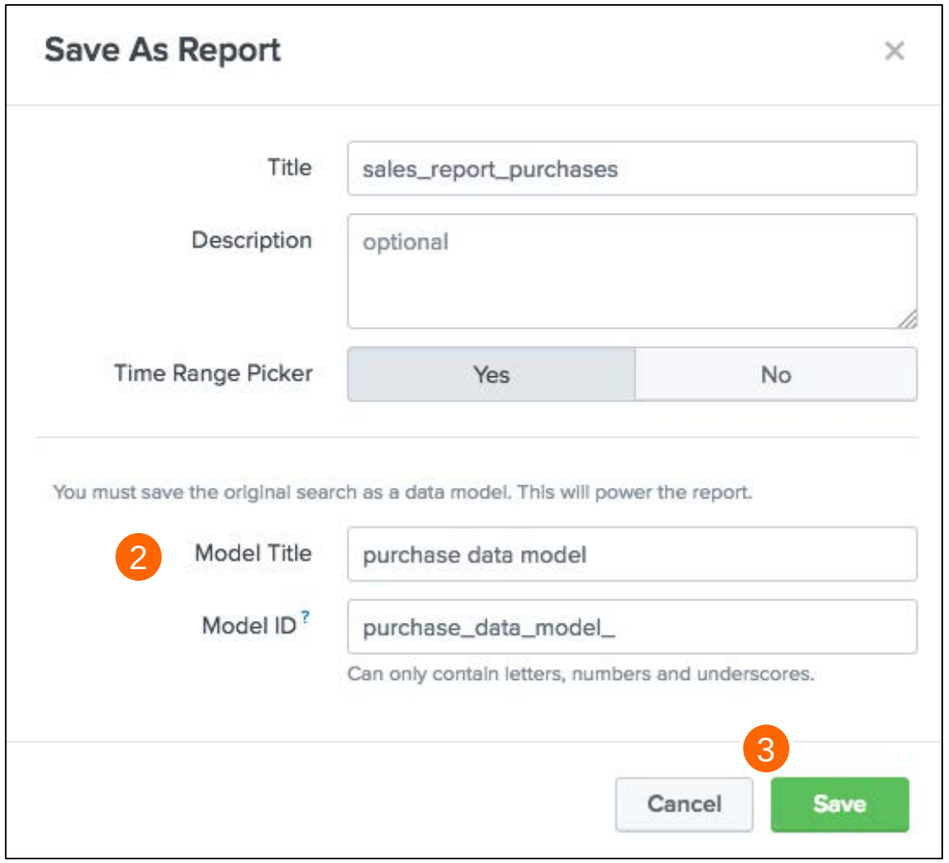

- When saving as a report, the **Model Title** is required
  - This is used to create a data model, which is required by the pivot report
- The **Model ID** is automatically generated based on the **Model Title**

Note

Manually changing the Model ID is not recommended.

Add a Pivot to a Dashboard

Similarly, you can save any pivot to a new or existing dashboard

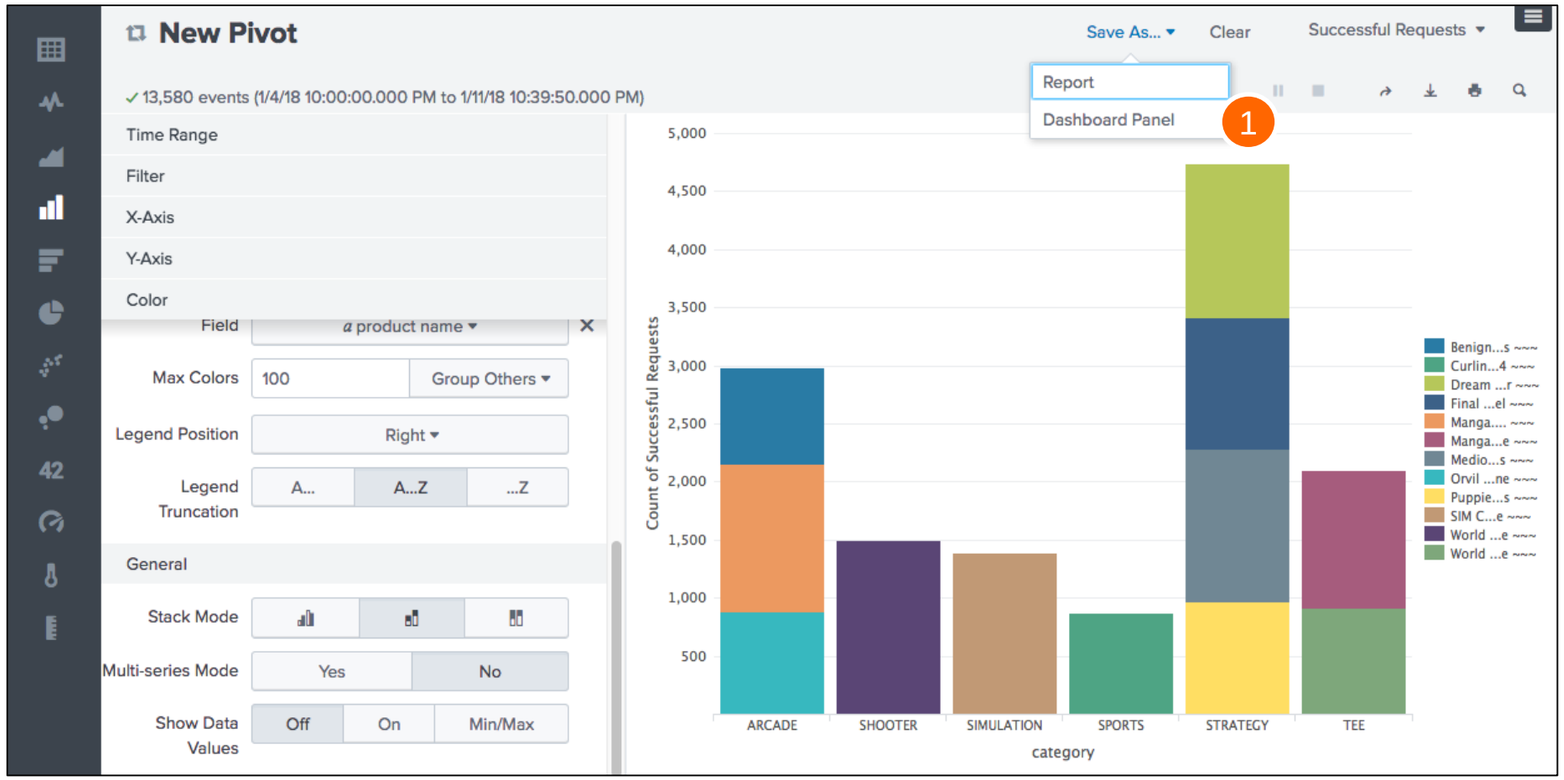

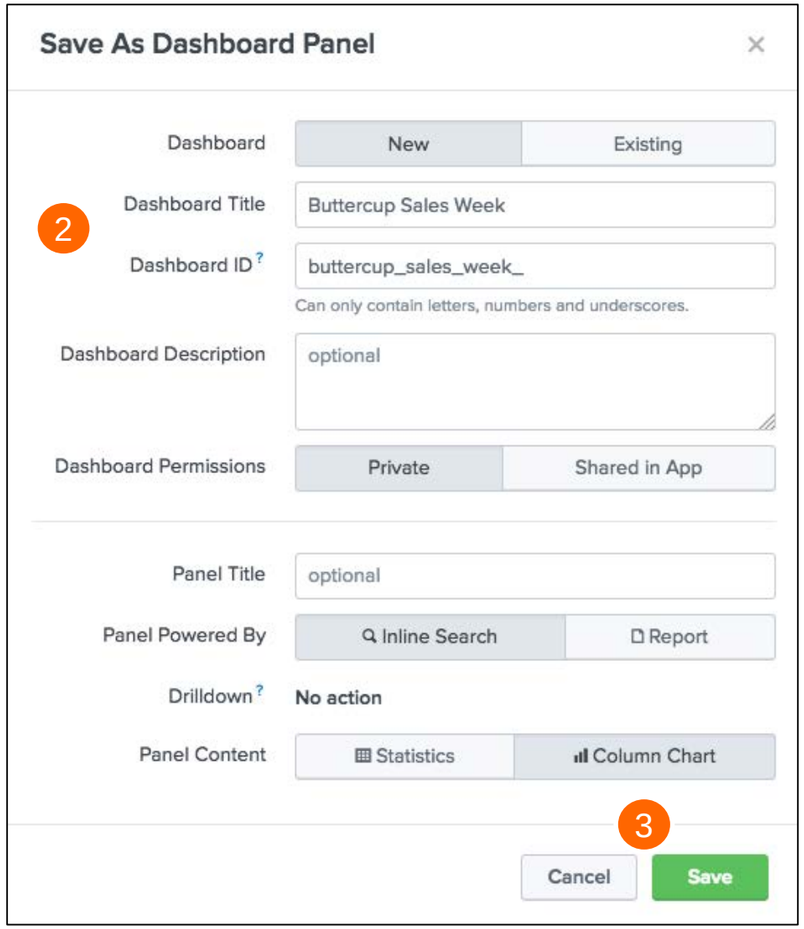

Module 12: Creating and Using Lookups

What Is a Lookup?

- Sometimes static (or relatively unchanging) data is required for searches, but isn't available in the index
- Lookups pull such data from standalone files at search time and add it to search results

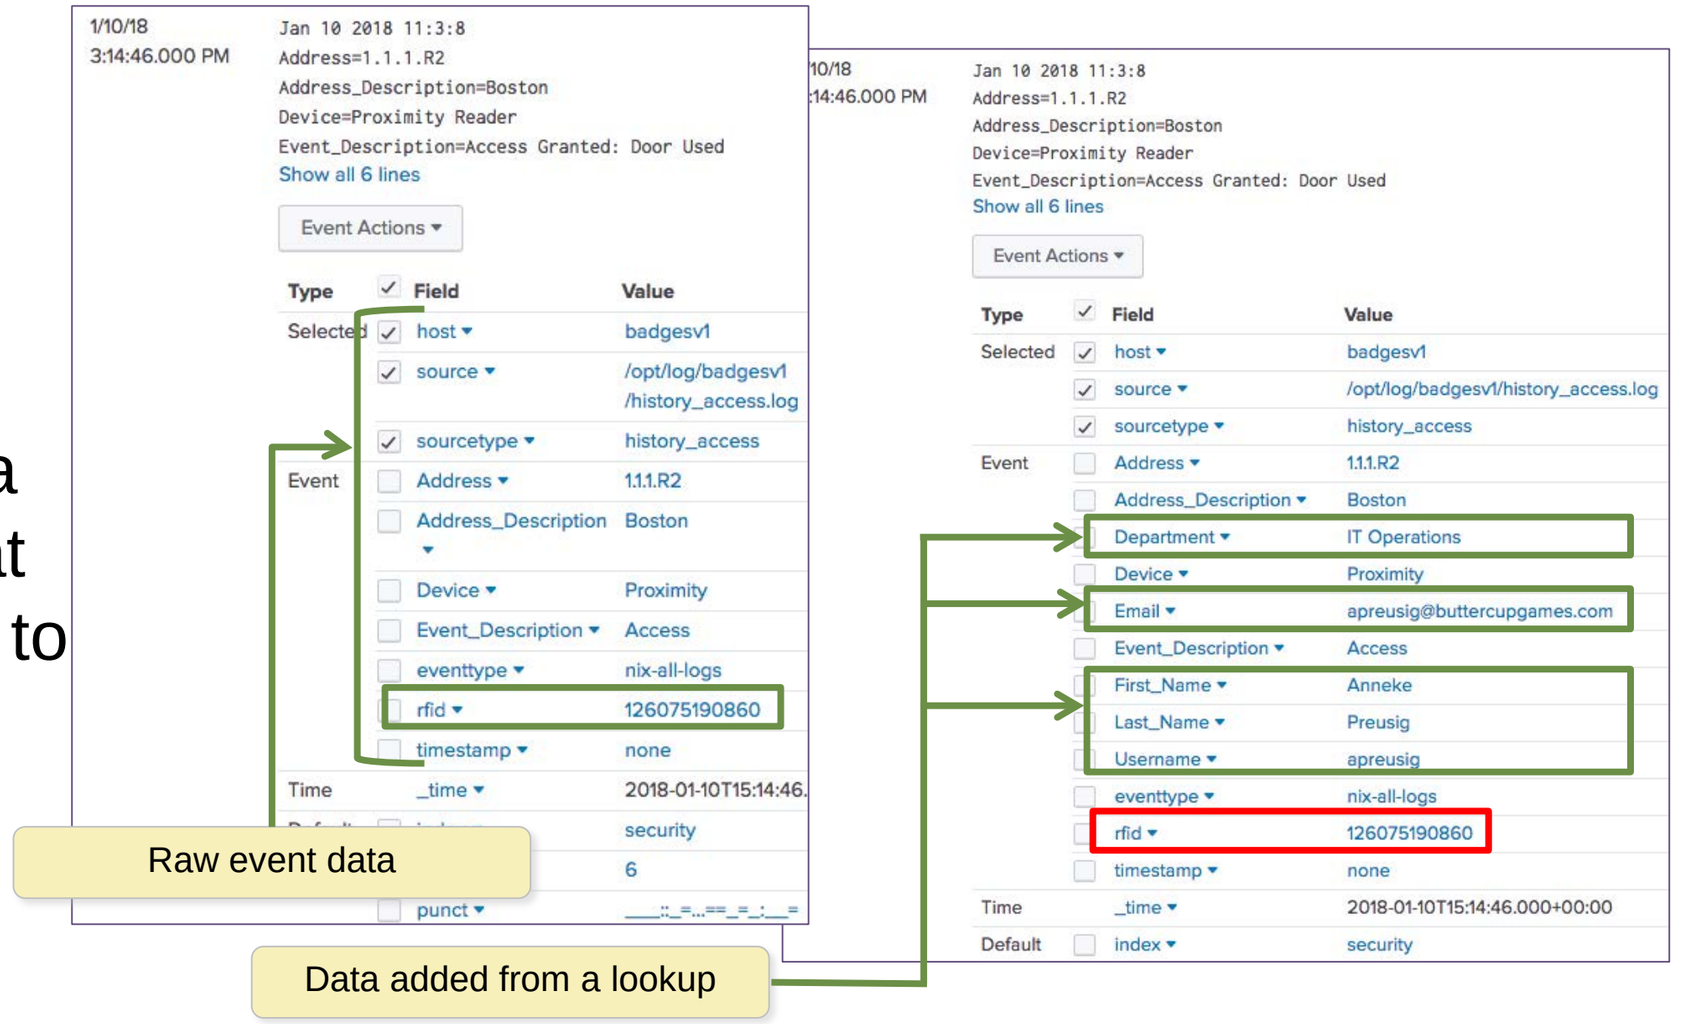

What Is a Lookup? (cont.)

- Lookups allow you to add more fields to your events, such as:
  - Descriptions for HTTP status codes (“File Not Found”, “Service Unavailable”)
  - Sale prices for products
  - User names, IP addresses, and workstation IDs associated with RFIDs
- After a lookup is configured, you can use the lookup fields in searches
- The lookup fields also appear in the Fields sidebar
- Lookup field values are case sensitive by default

Note

Admins can change the `case_sensitive_match` option to `false` in `transforms.conf`

A Sample Lookup File

- This example displays a lookup `.csv` file used to associate product information with productId
- First row represents field names (header)

  `productId, product_name, categoryId, price, sale_price, Code`

- The `productId` field exists in the `access_combined` events
  - This is the **input** field
- All of the fields listed above are available to search after the lookup is defined
  - These are the **output** fields

```
 GNU nano 2.3.1                File: products.csv

productId,product_name,categoryId,price,sale_price,Code
DB-SG-G01,Mediocre Kingdoms,STRATEGY,24.99,19.99,A
DC-SG-G02,Dream Crusher,STRATEGY,39.99,24.99,B
FS-SG-G03,Final Sequel,STRATEGY,24.99,16.99,C
WC-SH-G04,World of Cheese,SHOOTER,24.99,19.99,D
WC-SH-T02,World of Cheese Tee,TEE,9.99,6.99,E
PZ-SG-G05,Puppies vs. Zombies,STRATEGY,4.99,1.99,F
CU-PG-G06,Curling 2014,SPORTS,19.99,16.99,G
MB-AG-G07,Manganiello Bros.,ARCADE,39.99,24.99,H
MB-AG-T01,Manganiello Bros. Tee,TEE,9.99,6.99,I
FI-AG-G08,Orvil the Wolverine,ARCADE,39.99,24.99,J
BS-AG-G09,Benign Space Debris,ARCADE,24.99,19.99,K
SC-MG-G10,SIM Cubicle,SIMULATION,19.99,16.99,L
WC-SH-A01,Holy Blade of Gouda,ACCESSORIES,5.99,2.99,M
WC-SH-A02,Fire Resistance Suit of Provolone,ACCESSORIES,3.99,1.99,N
```

Creating a Lookup

1. Upload the file required for the lookup
2. Define the lookup type
3. Optionally, configure the lookup to run automatically

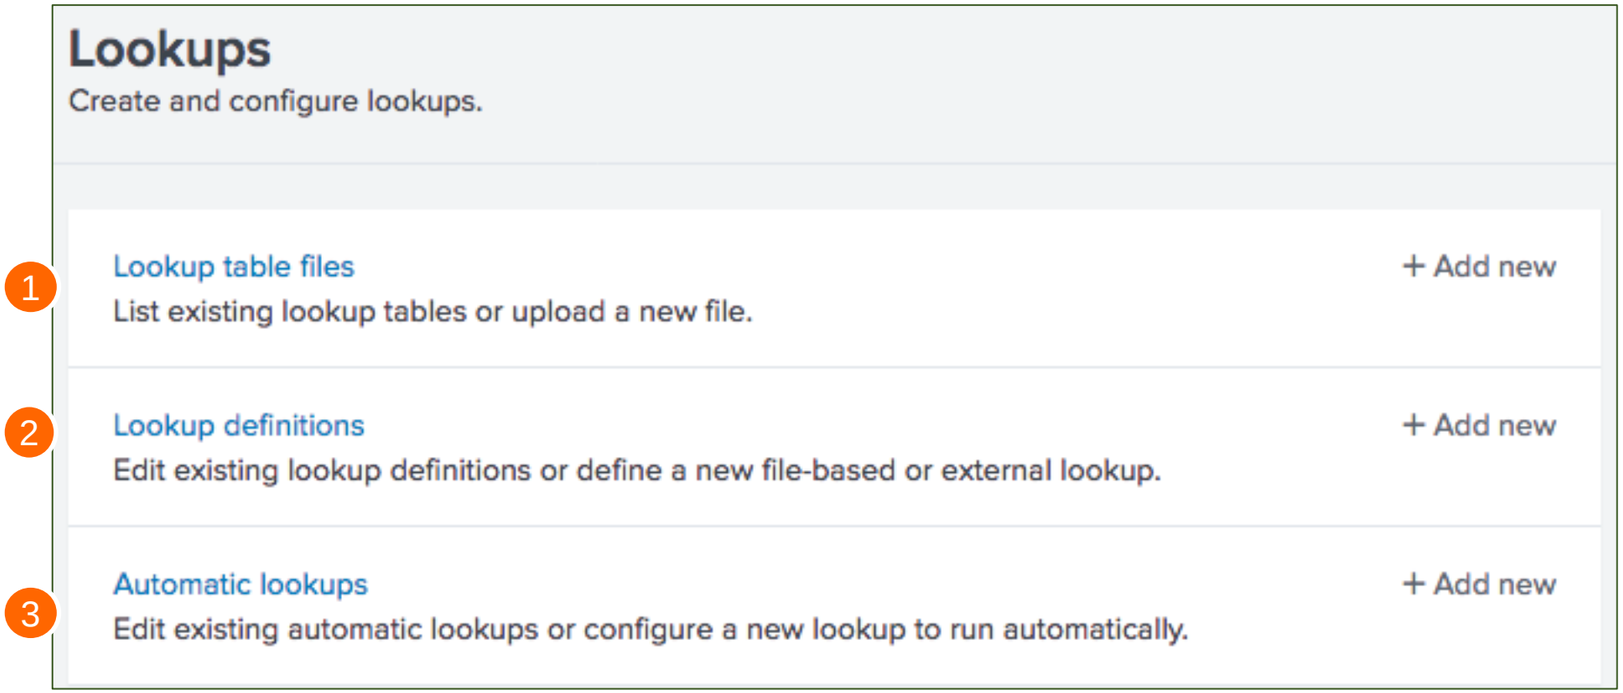

Adding a New Lookup Table File

Settings > Lookups > Lookup table files

1. Click **New Lookup Table File**
2. Select a destination app
3. Browse and select the `.csv` file to use for the lookup table
4. Enter a name for the lookup file
5. Save

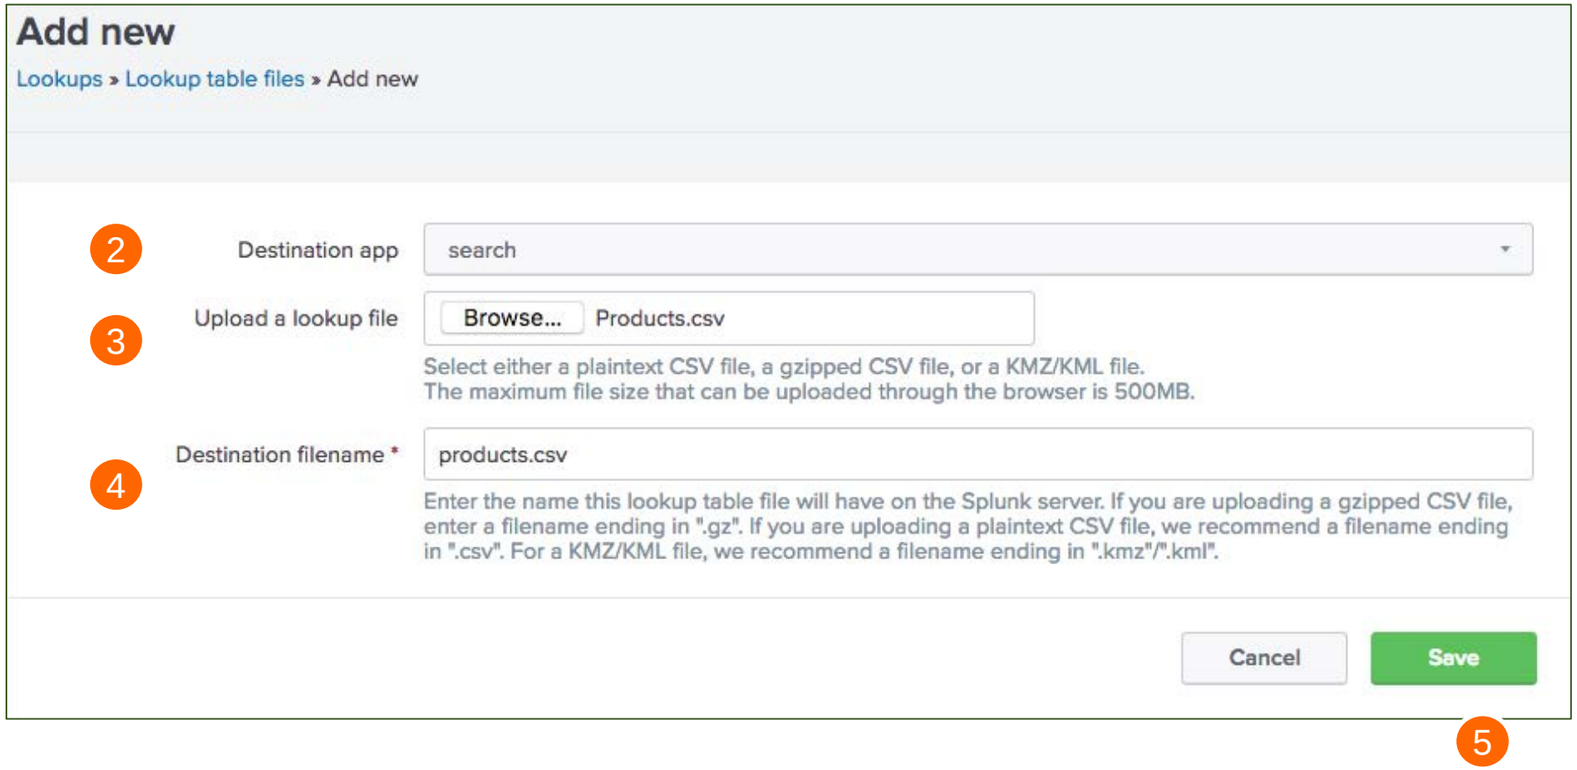

inputlookup Command

- Use the `inputlookup` command to load the results from a specified static lookup
- Useful to:
  - Review the data in the .csv file
  - Validate the lookup

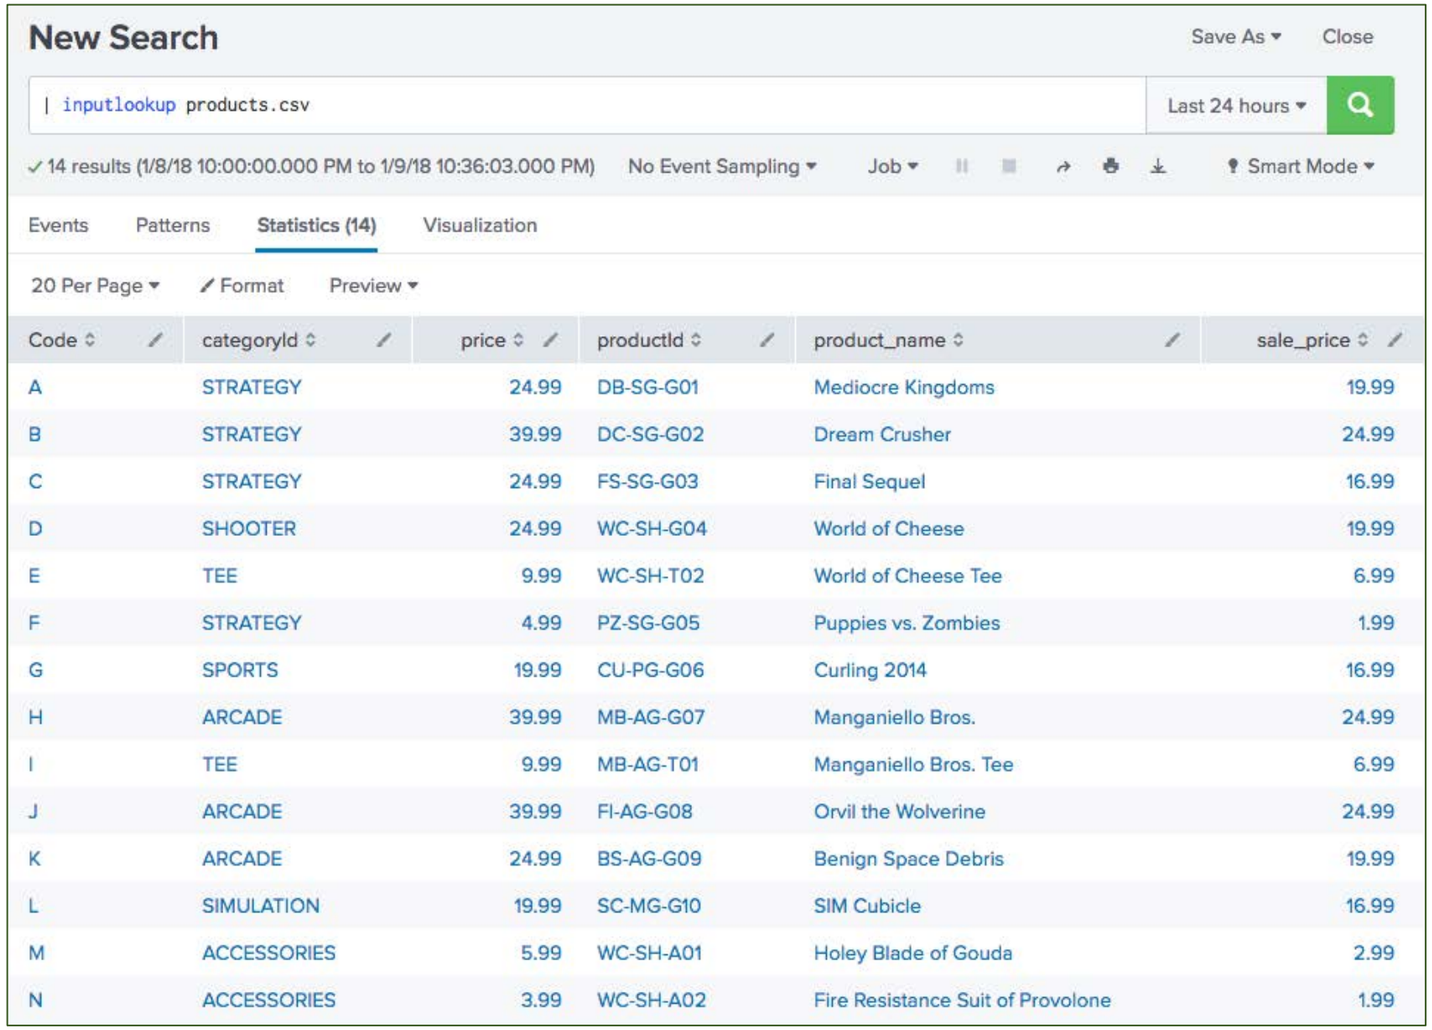

Code	categoryId	price	productId	product_name	sale_price
A	STRATEGY	24.99	DB-SG-G01	Mediocre Kingdoms	19.99
B	STRATEGY	39.99	DC-SG-G02	Dream Crusher	24.99
C	STRATEGY	24.99	FS-SG-G03	Final Sequel	16.99
D	SHOOTER	24.99	WC-SH-G04	World of Cheese	19.99
E	TEE	9.99	WC-SH-T02	World of Cheese Tee	6.99
F	STRATEGY	4.99	PZ-SG-G05	Puppies vs. Zombies	1.99
G	SPORTS	19.99	CU-PG-G06	Curling 2014	16.99
H	ARCADE	39.99	MB-AG-G07	Manganiello Bros.	24.99
I	TEE	9.99	MB-AG-T01	Manganiello Bros. Tee	6.99
J	ARCADE	39.99	FI-AG-G08	Orvil the Wolverine	24.99
K	ARCADE	24.99	BS-AG-G09	Benign Space Debris	19.99
L	SIMULATION	19.99	SC-MG-G10	SIM Cubicle	16.99
M	ACCESSORIES	5.99	WC-SH-A01	Holey Blade of Gouda	2.99
N	ACCESSORIES	3.99	WC-SH-A02	Fire Resistance Suit of Provolone	1.99

Note

When using the `inputlookup` command, you can specify the filename ending with .csv or the lookup definition name.

Creating a Lookup Definition

Settings > Lookups > Lookup definitions

1. Click **New Lookup Definition**
2. Select a destination app
3. Name the lookup definition
4. Select the lookup type, either File-based or External
5. From the drop-down, select a lookup file
6. Save

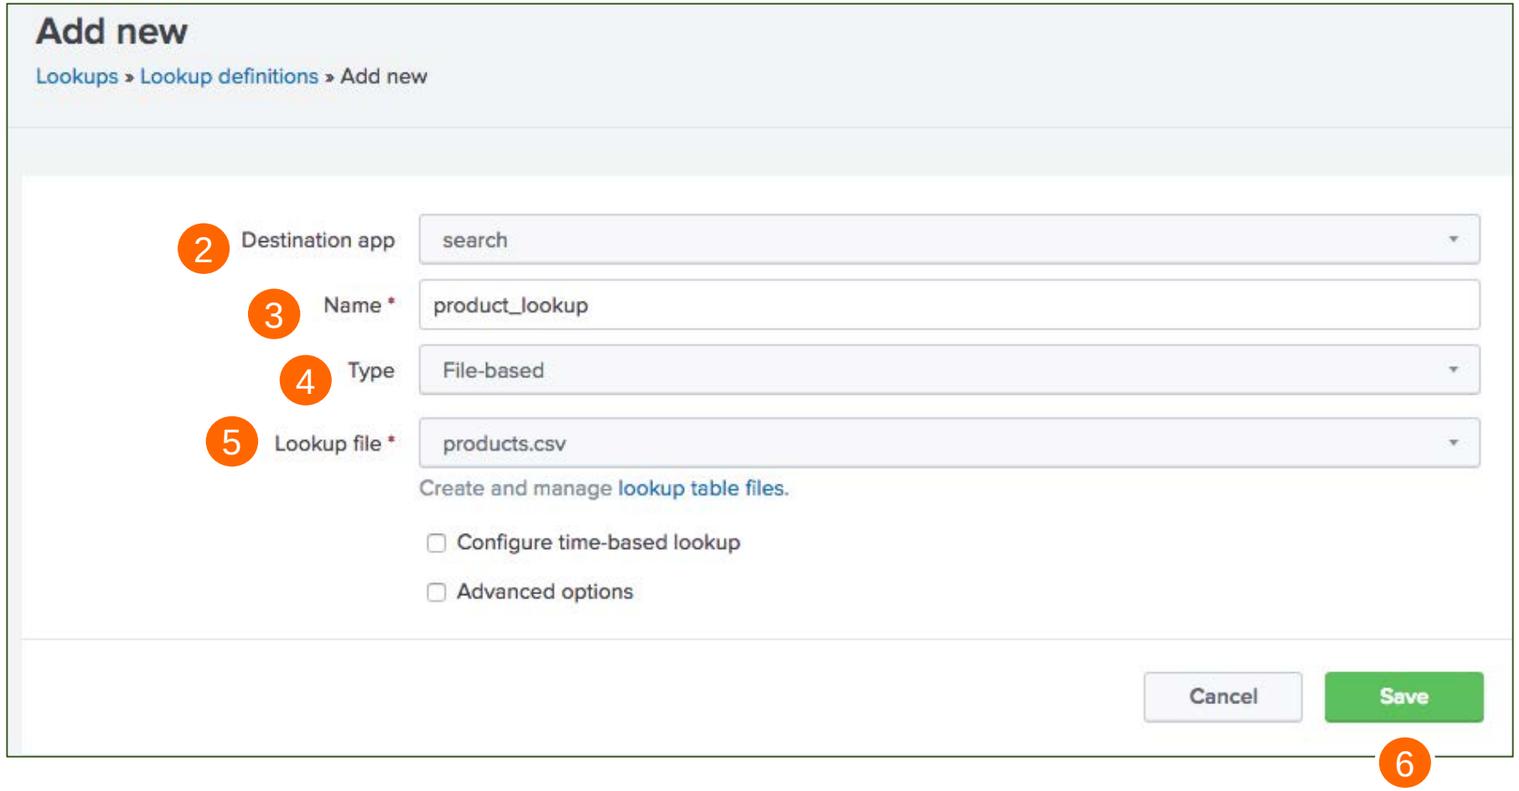

Applying Advanced Lookup Options

- Min/max # of matches for each input lookup value
- Default value to output (when fewer than the min # of matches present for a given input)
- Case sensitivity match on/off
- Batch index query: improves performance for large lookup files
- Match type: supplies format for non-exact matching
- Filter lookup: filters results before returning data

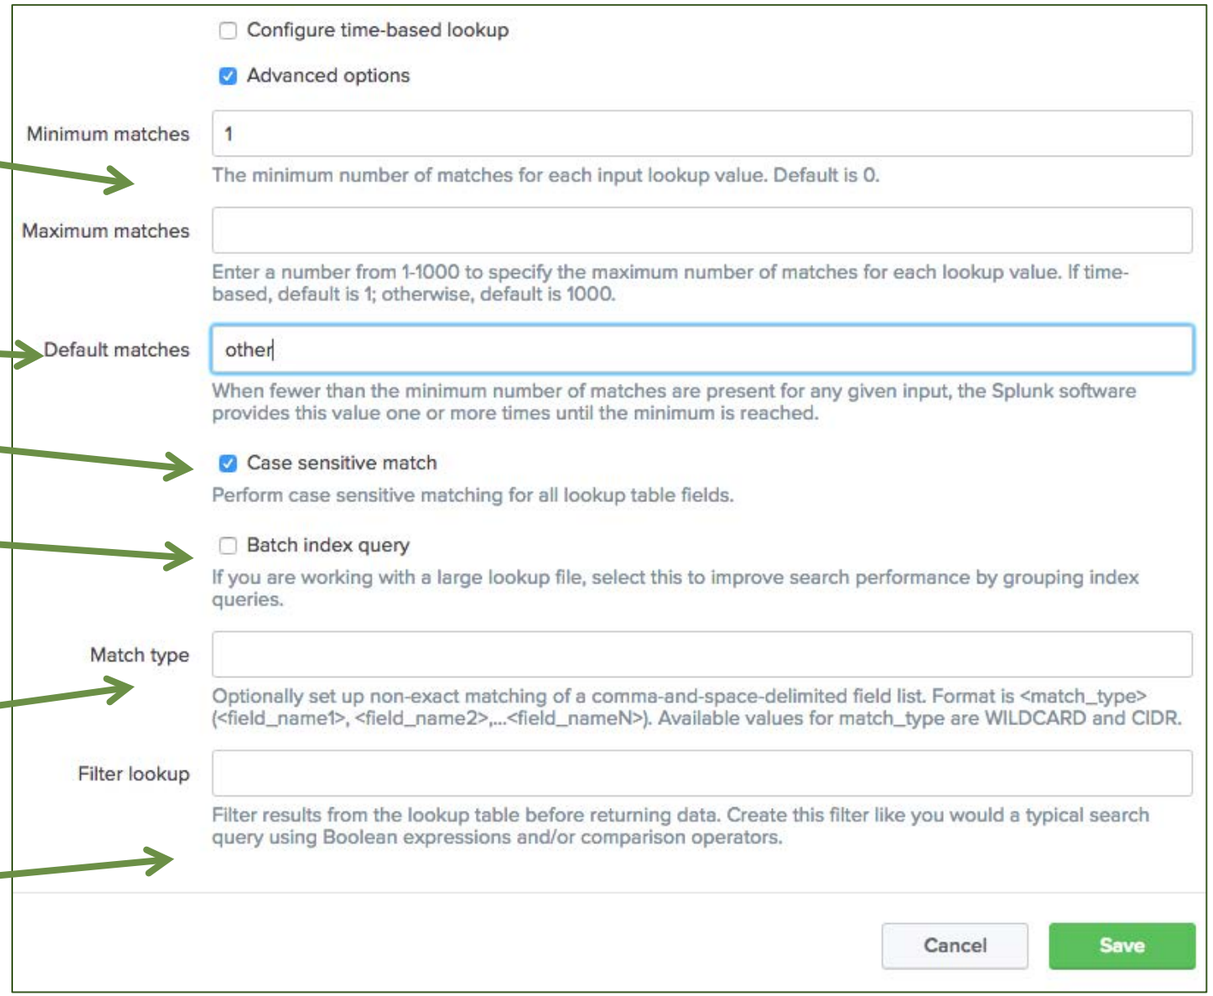

lookup Command

- If a lookup is not configured to run automatically, use the lookup command in your search to use the lookup fields
- The OUTPUT argument is optional
  - If OUTPUT not specified, lookup returns all the fields from the lookup table except the match fields
  - If OUTPUT is specified, the fields overwrite existing fields
- The output lookup fields exist only for the current search
- Use OUTPUTNEW when you do not want to overwrite existing fields

lookup Help More »
Explicitly invokes field value lookups.

Examples

There is a lookup table specified in a stanza name 'usertogroup' in transform.conf. This lookup table contains (at least) two fields, 'user' and 'group'. For each event, we look up the value of the field 'local_user' in the table and for any entries that matches, the value of the 'group' field in the lookup table will be written to the field 'user_group' in the event.

... | lookup usertogroup user as local_user OUTPUT group as user_group

Using the lookup Command

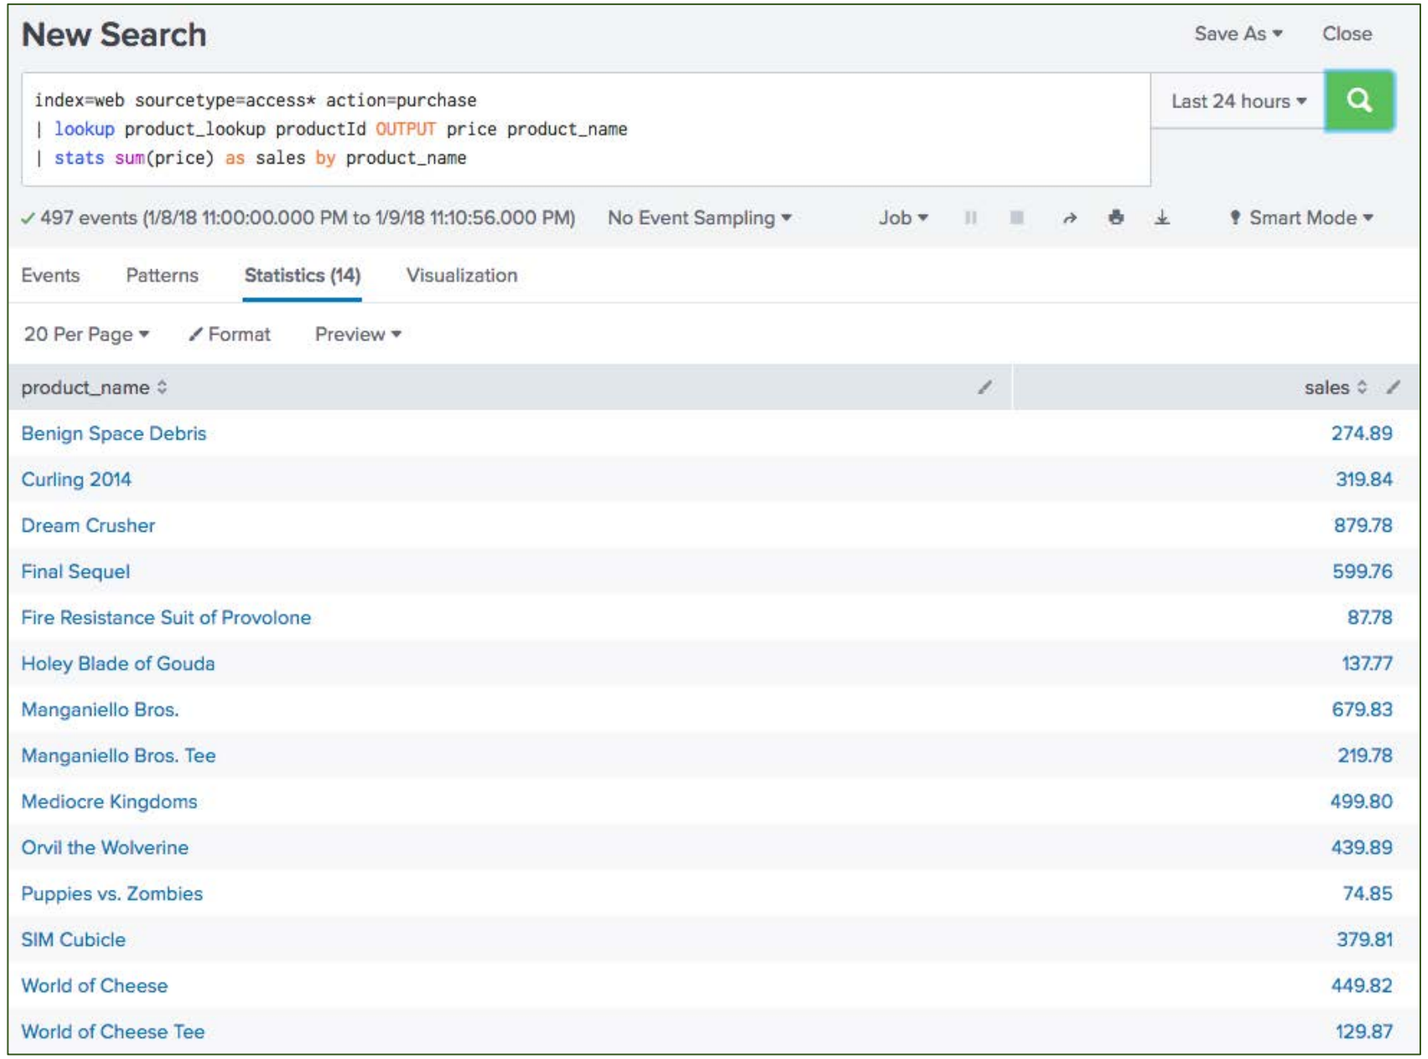

Scenario

Calculate the sales for each product in the last 24 hours.

Creating an Automatic Lookup

Settings > Lookups > Automatic lookups

1. Click **New Automatic Lookup**
2. Select the Destination app
3. Enter a Name for the lookup
3. Select the Lookup table definition
4. Select host, source, or sourcetype to apply to the lookup and specify the name

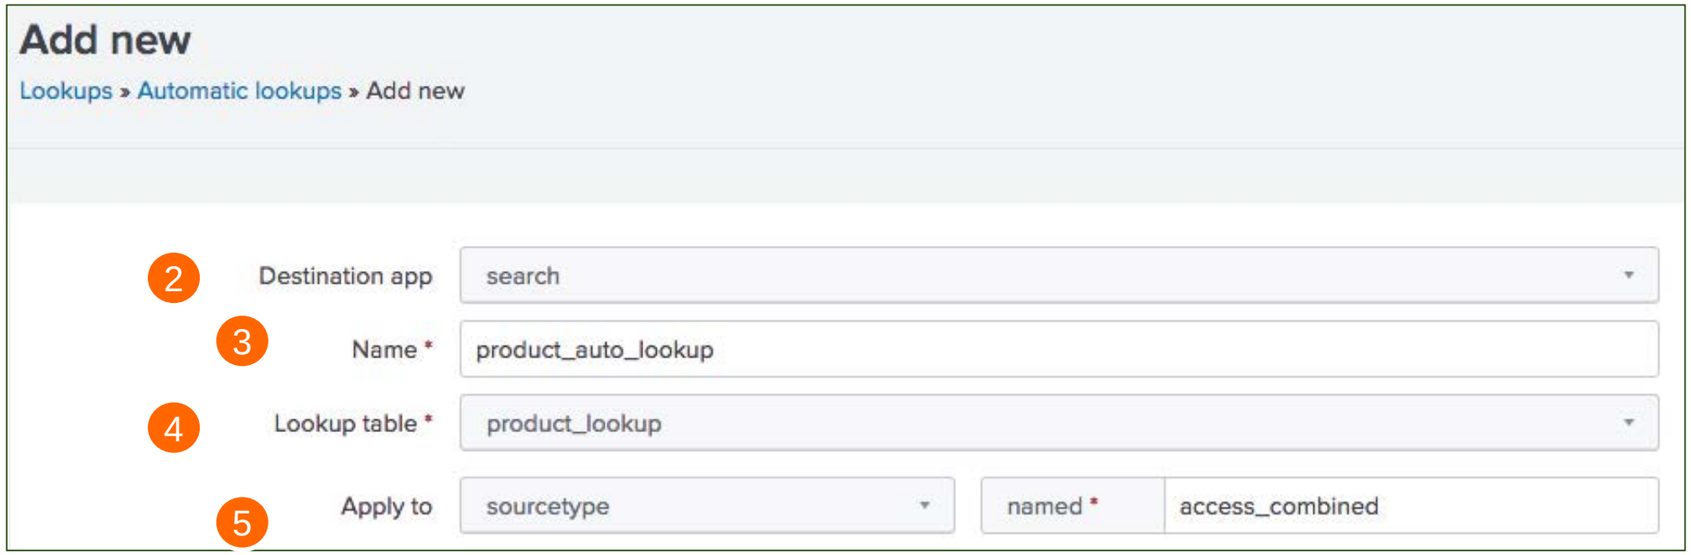

Creating an Automatic Lookup (cont.)

5. Define the Lookup input fields

   Field(s) that exist in your events that you are relating to the lookup table

   A. Column name in CSV

   B. Field name in Splunk, if different from column name

6. Define the Lookup output fields

   Field(s) from your lookup table that are added to the events

   C. Field name in lookup table

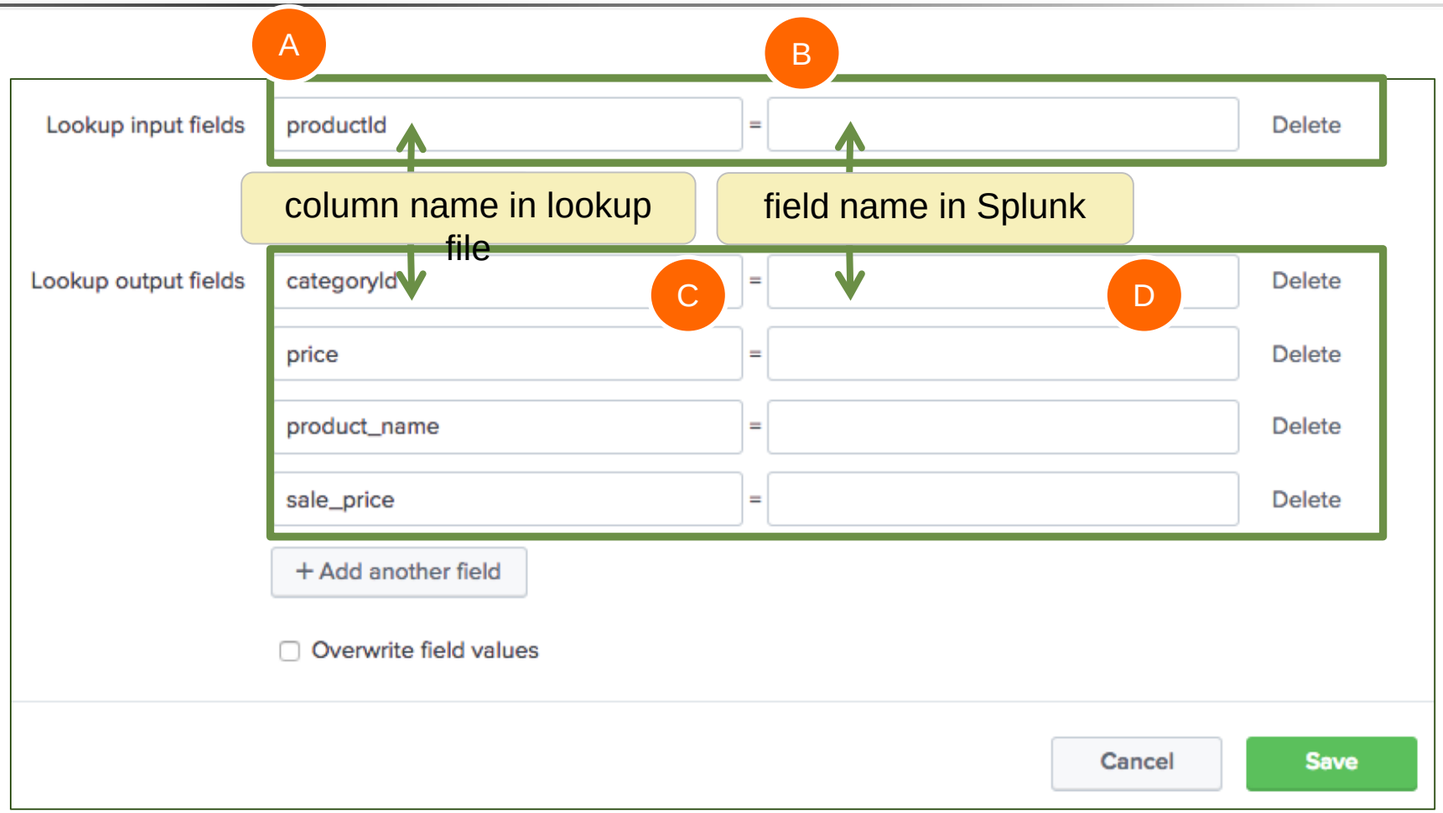

   D. Name you want displayed in Splunk; otherwise it inherits the column name

7. Save

Using the Automatic Lookup

To use an automatic lookup, specify the output fields in your search

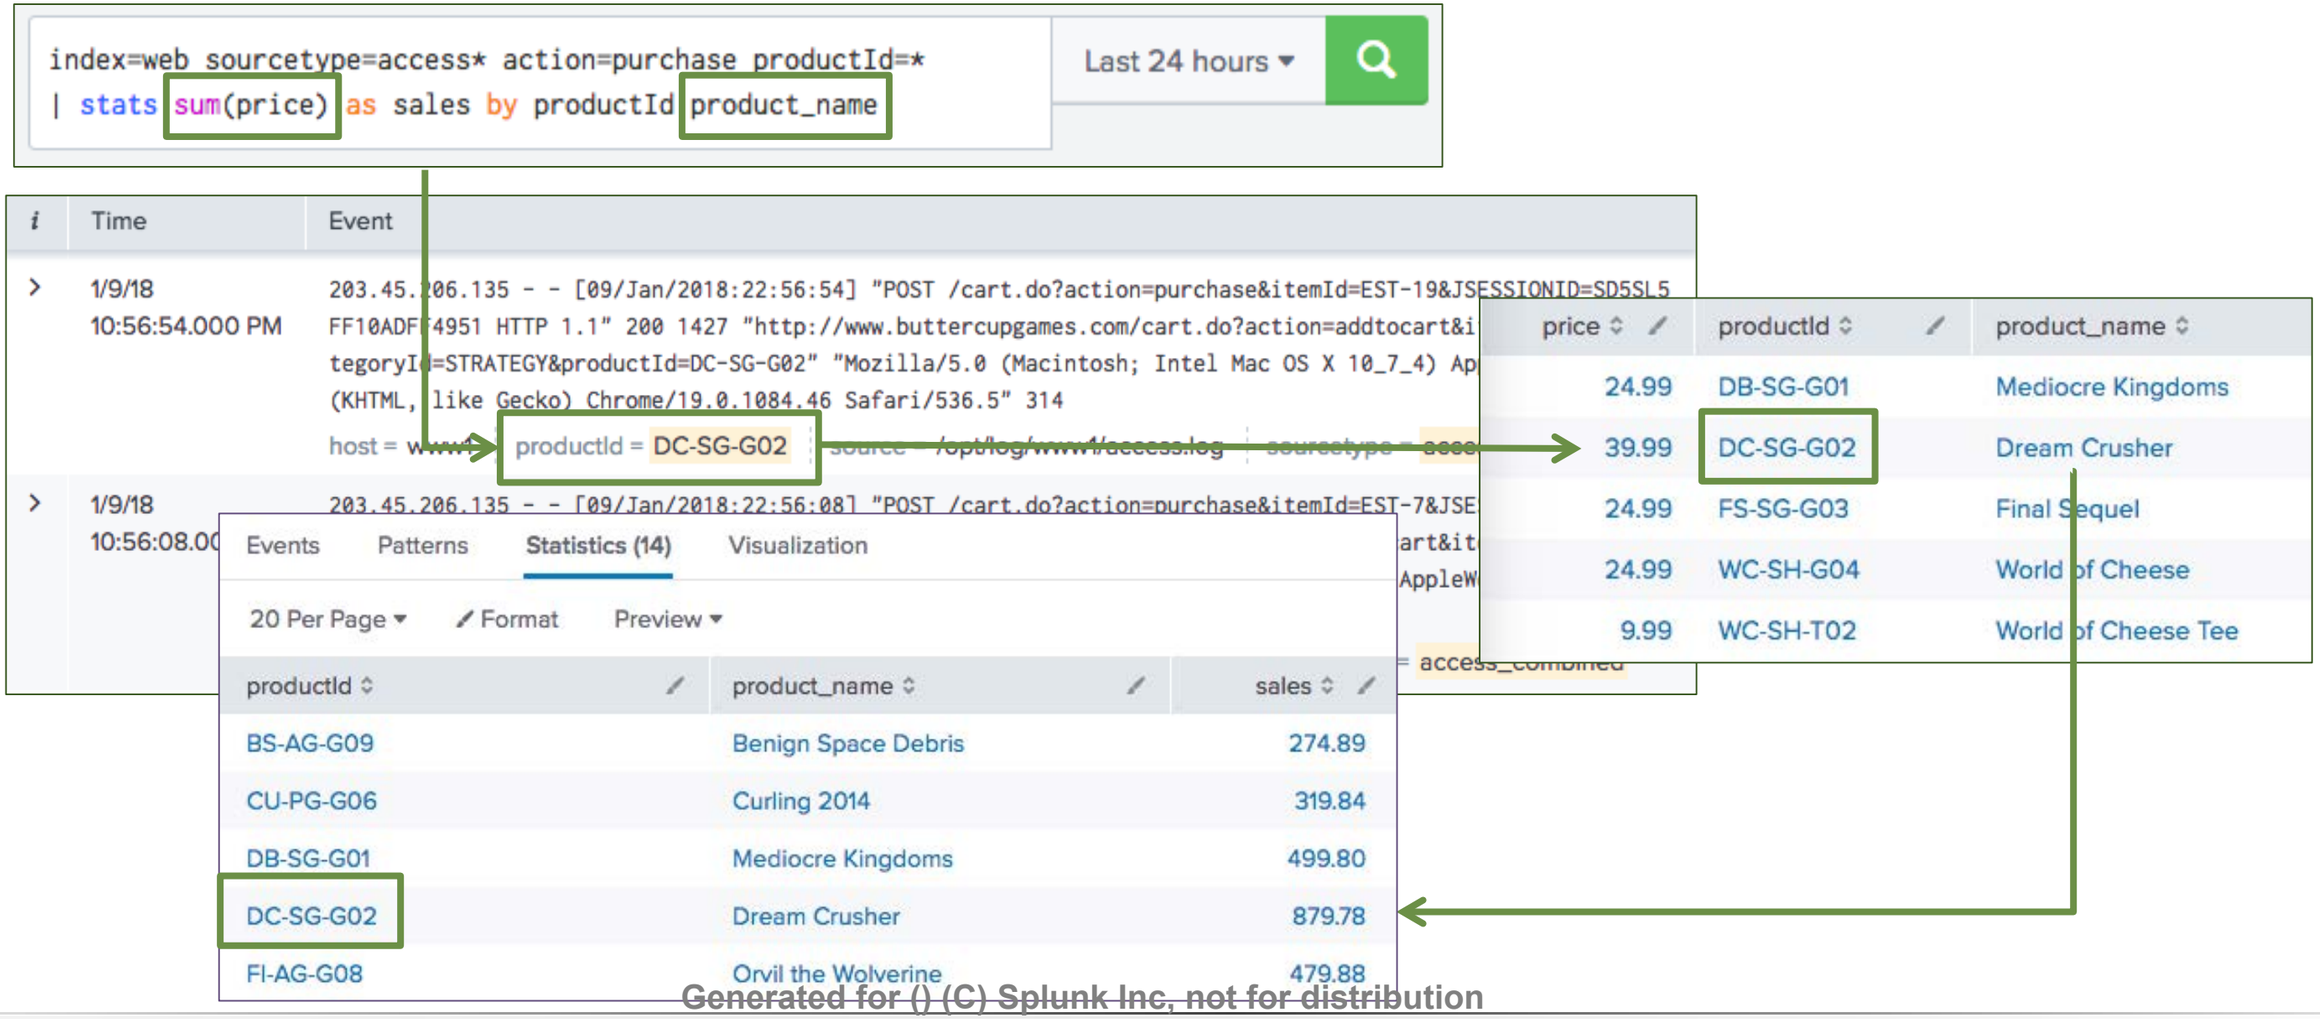

Time-based Lookups

- If a field in the lookup table represents a timestamp, you can create a time-based lookup
- In this example, the search retrieved events for December and January and calculated the sales based on the correct unit price for those dates

product_time	productId	product_name	categoryId	price	sale_price
12/1/17	DB-SG-G01	Mediocre Kingdoms	STRATEGY	24.99	19.99
12/1/17	DC-SG-G02	Dream Crusher	STRATEGY	39.99	24.99
12/1/17	FS-SG-G03	Final Sequel	STRATEGY	24.99	16.99
12/1/17	WC-SH-G04	World of Cheese	SHOOTER	24.99	19.99
12/1/17	WC-SH-T02	World of Cheese Tee	TEE	9.99	6.99
12/1/17	PZ-SG-G05	Puppies vs. Zombies	STRATEGY	4.99	1.99
12/1/17	CU-PG-G06	Curling 2014	SPORTS	19.99	16.99
12/1/17	MB-AG-G07	Manganiello Bros.	ARCADE	39.99	24.99
12/1/17	MB-AG-T01	Manganiello Bros. Tee	TEE	9.99	6.99
12/1/17	FI-AG-G08	Orvil the Wolverine	ARCADE	39.99	24.99
12/1/17	BS-AG-G09	Benign Space Debris	ARCADE	24.99	19.99
12/1/17	SC-MG-G10	SIM Cubicle	SIMULATION	19.99	16.99
12/1/17	WC-SH-A01	Holey Blade of Gouda	ACCESSORIES	5.99	2.99
12/1/17	WC-SH-A02	Fire Resistance Suit of Provolone	ACCESSORIES	3.99	1.99
1/1/18	DB-SG-G01	Mediocre Kingdoms	STRATEGY	24.99	19.99
1/1/18	DC-SG-G02	Dream Crusher	STRATEGY	40.99	24.99
1/1/18	FS-SG-G03	Final Sequel	STRATEGY	25.99	16.99
1/1/18	WC-SH-G04	World of Cheese	SHOOTER	24.99	19.99
1/1/18	WC-SH-T02	World of Cheese Tee	TEE	9.99	6.99
1/1/18	PZ-SG-G05	Puppies vs. Zombies	STRATEGY	4.99	1.99
1/1/18	CU-PG-G06	Curling 2014	SPORTS	20.99	16.99
1/1/18	MB-AG-G07	Manganiello Bros.	ARCADE	39.99	24.99
1/1/18	MB-AG-T01	Manganiello Bros. Tee	TEE	9.99	6.99
1/1/18	FI-AG-G08	Orvil the Wolverine	ARCADE	39.99	24.99
1/1/18	BS-AG-G09	Benign Space Debris	ARCADE	25.99	19.99
1/1/18	SC-MG-G10	SIM Cubicle	SIMULATION	19.99	16.99
1/1/18	WC-SH-A01	Holey Blade of Gouda	ACCESSORIES	6.99	2.99
1/1/18	WC-SH-A02	Fire Resistance Suit of Provolone	ACCESSORIES	4.99	1.99

product_name	Month	price	count	sales	SubTotal Sales
Benign Space Debris	Dec	24.99	828	42,211.44	
Benign Space Debris	Jan	25.99	646	32,933.08	
Benign Space Debris Subtotal					150,289.04
Curling 2014	Dec	19.99	868	35,570.64	
Curling 2014	Jan	20.99	620	25,407.60	
Curling 2014 Subtotal					121,956.48
Dream Crusher	Dec	39.99	1,360	110,132.80	
Dream Crusher	Jan	40.99	1,004	81,303.92	
Dream Crusher Subtotal					382,873.44
Final Sequel	Dec	24.99	1,200	61,176.00	
Final Sequel	Jan	25.99	812	41,395.76	
Final Sequel Subtotal					205,143.52

Module 13
Creating Scheduled Reports and Alerts

Why Scheduled Reports?

Scheduled Reports are useful for:

- Monthly, weekly, daily executive/managerial roll up reports
- Dashboard performance
- Automatically sending reports via email

Creating a Scheduled Report

1. Create your search
2. From the Save As menu, select Report

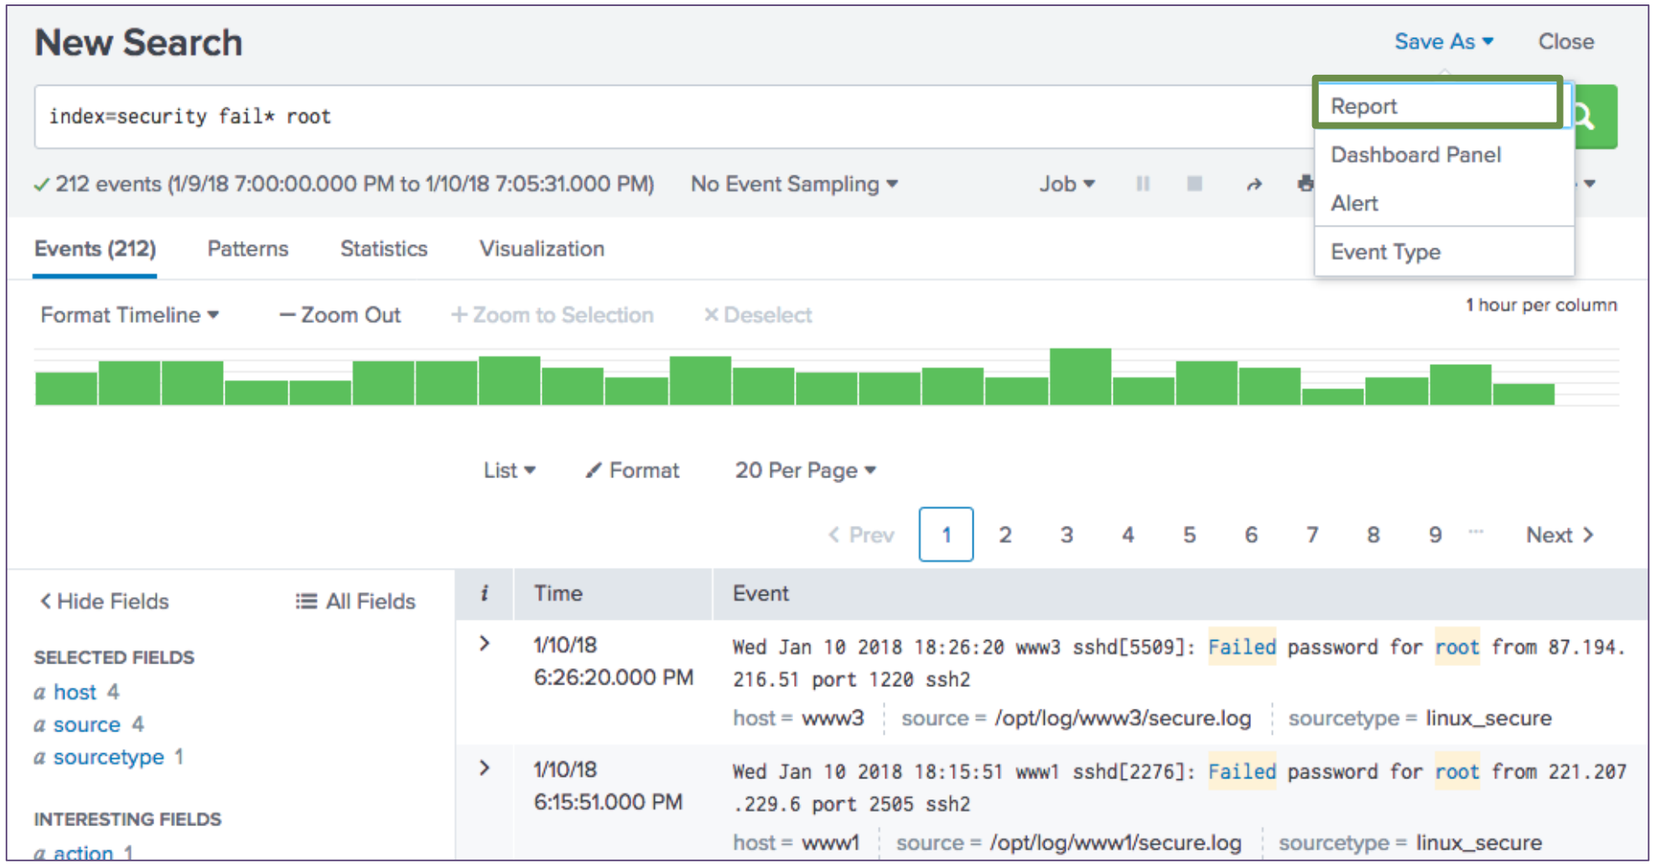

Creating a Scheduled Report (cont.)

3. Enter Title
4. Enter Description
5. Set Time Range Picker to No
6. Click Save

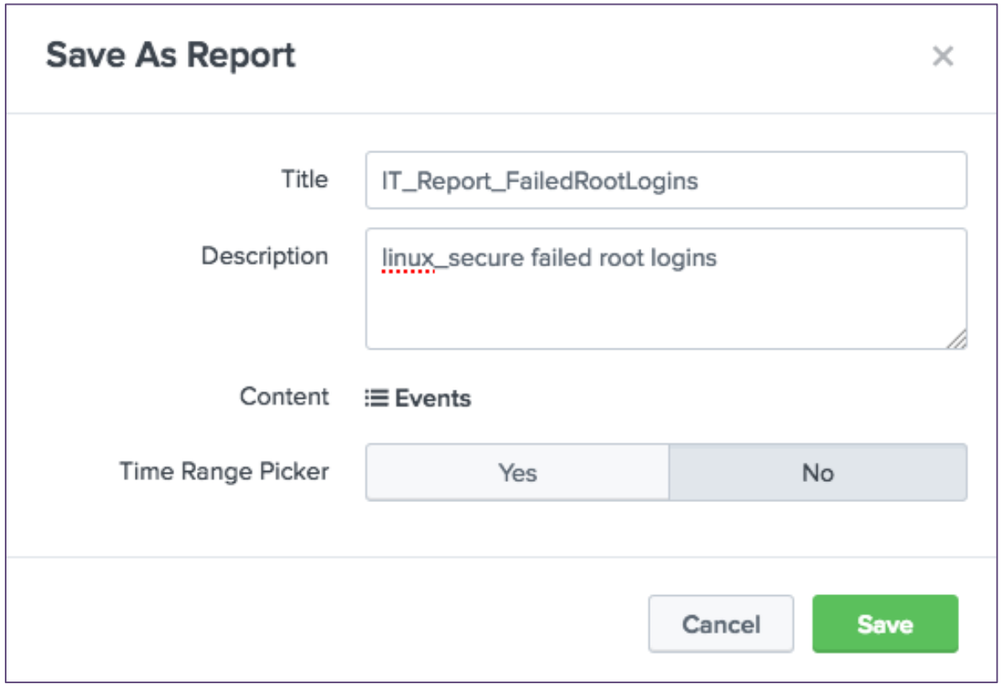

Note

Time Range Picker cannot be used with scheduled reports.

Creating a Scheduled Report (cont.)

- After the report is created, click Schedule
- If you inadvertently set Time Range Picker to Yes on previous screen, a warning displays and time picker is disabled

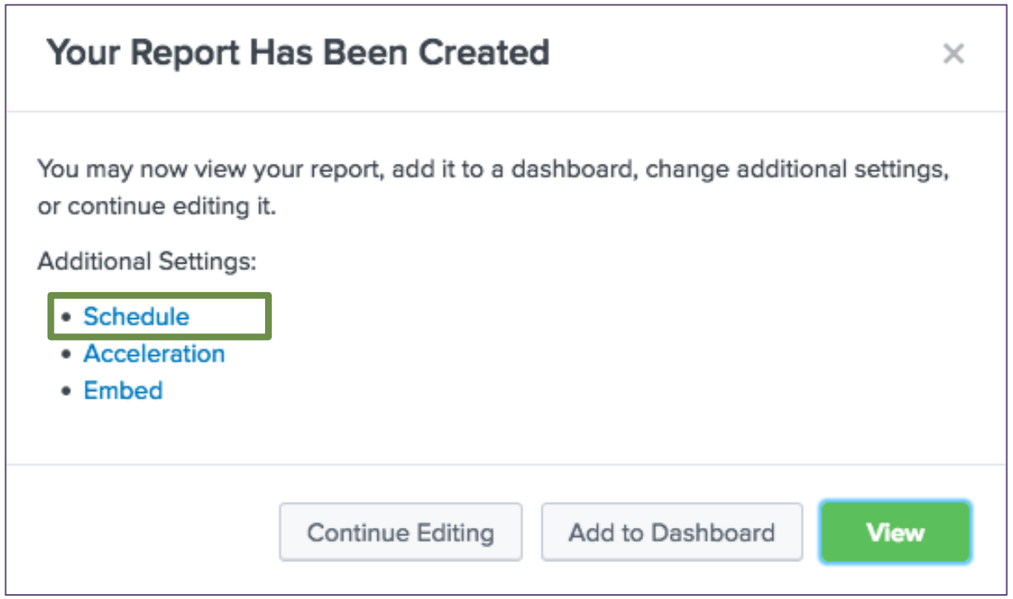

Scheduling this report results in removal of the time picker from the report display.

Note

Depending on the permissions granted to you by your Splunk administrator, you may be able to set permissions to share your scheduled report.

Creating a Scheduled Report – Define Schedule

- Schedule Report – select this checkbox
- Schedule – select the frequency to run the report
  - Run every hour
  - Run every day
  - Run every week
  - Run every month
  - Run on Cron Schedule

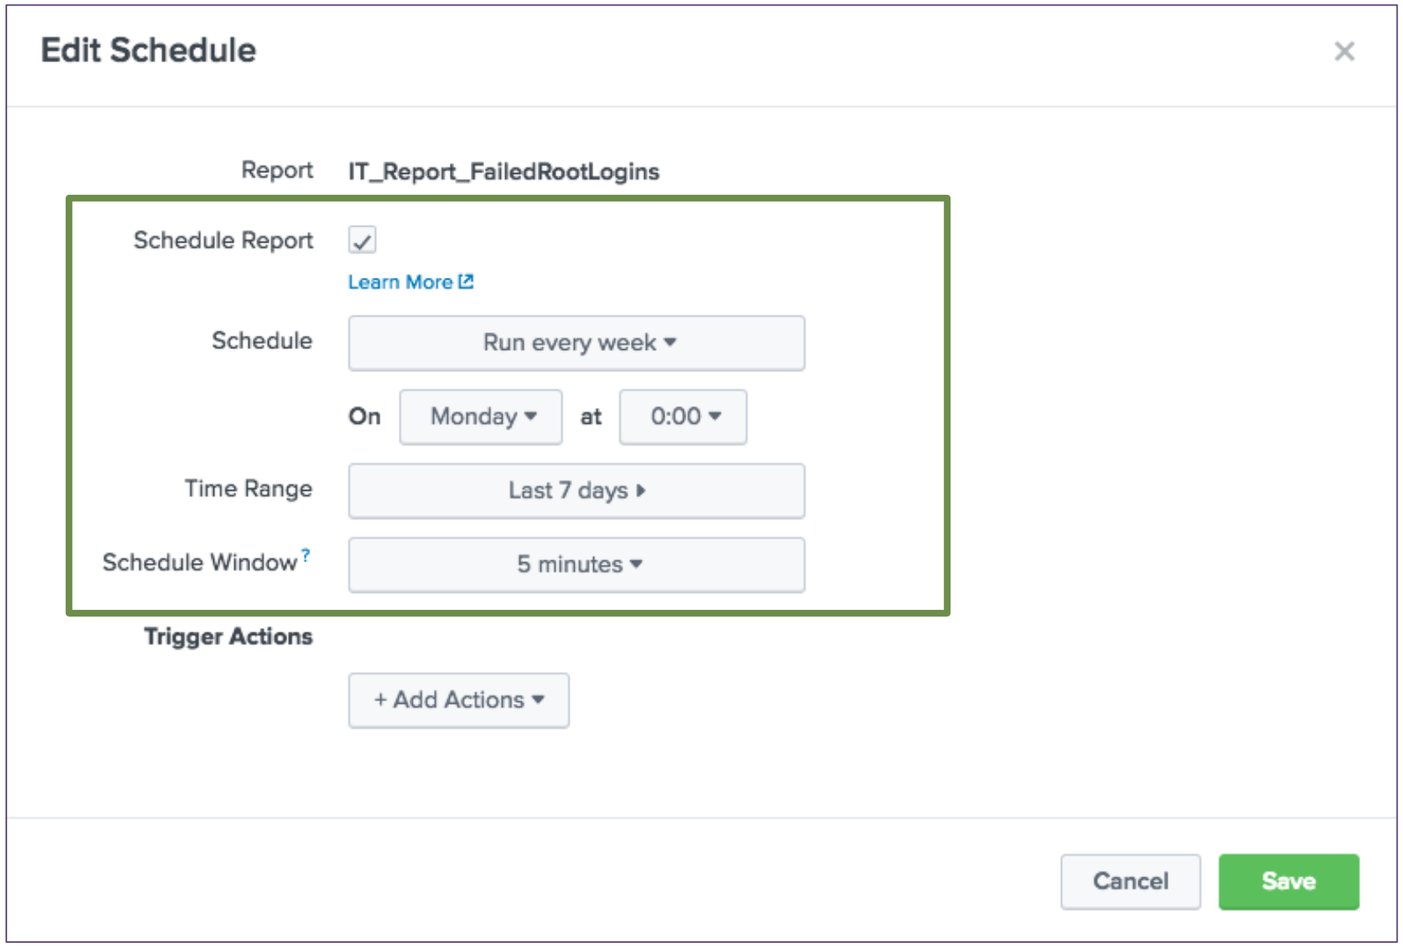

Creating a Scheduled Report – Select Time Range

- Time Range – By default, search time range used
  - Click the Time Range button to change the time range
  - You can select a time range from Presets, Relative, or Advanced
  - Typically, the time range is relative to the Schedule

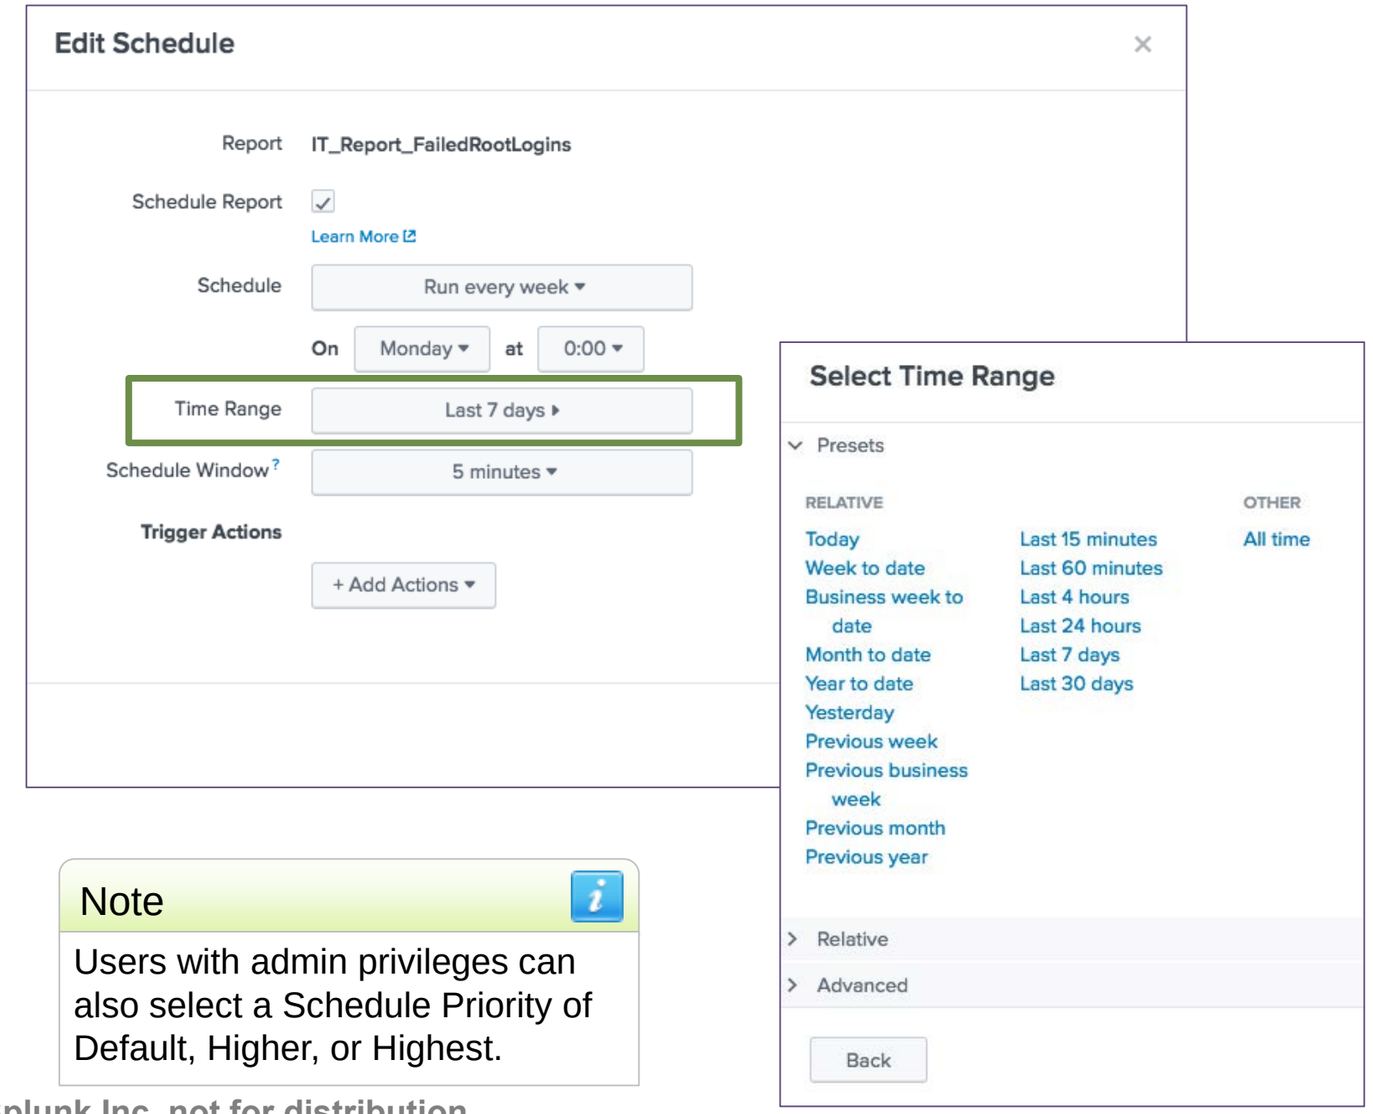

Note

Users with admin privileges can also select a Schedule Priority of Default, Higher, or Highest.

Creating a Scheduled Report – Schedule Window

- Schedule Window – this setting determines a time frame to run the report
  - If there are other reports scheduled to run at the same time, you can provide a window in which to run the report
  - This setting provides efficiency when scheduling several reports to run
- After you configure the report schedule, click Next

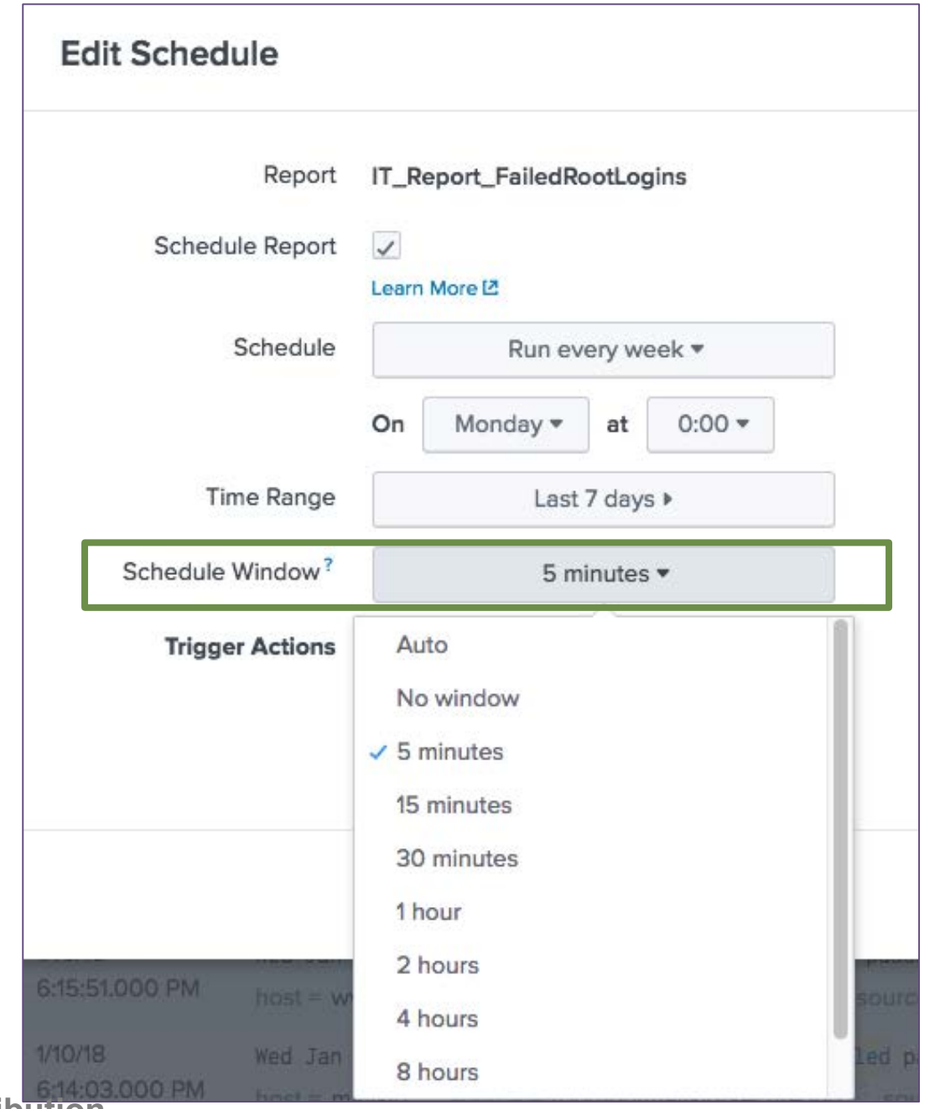

Creating a Scheduled Report – Add Actions

- **Log Event** – creates an indexed, searchable log event
- **Output results to lookup** – sends results of search to CSV lookup file
- **Output results to telemetry endpoint** – sends usage metrics back to Splunk (if your company has opted-in to program)
- **Run a script** – runs a previously created script
- **Send email** – sends an email with results to specified recipients
- **Webhook** – sends an HTTP POST request to specified URL

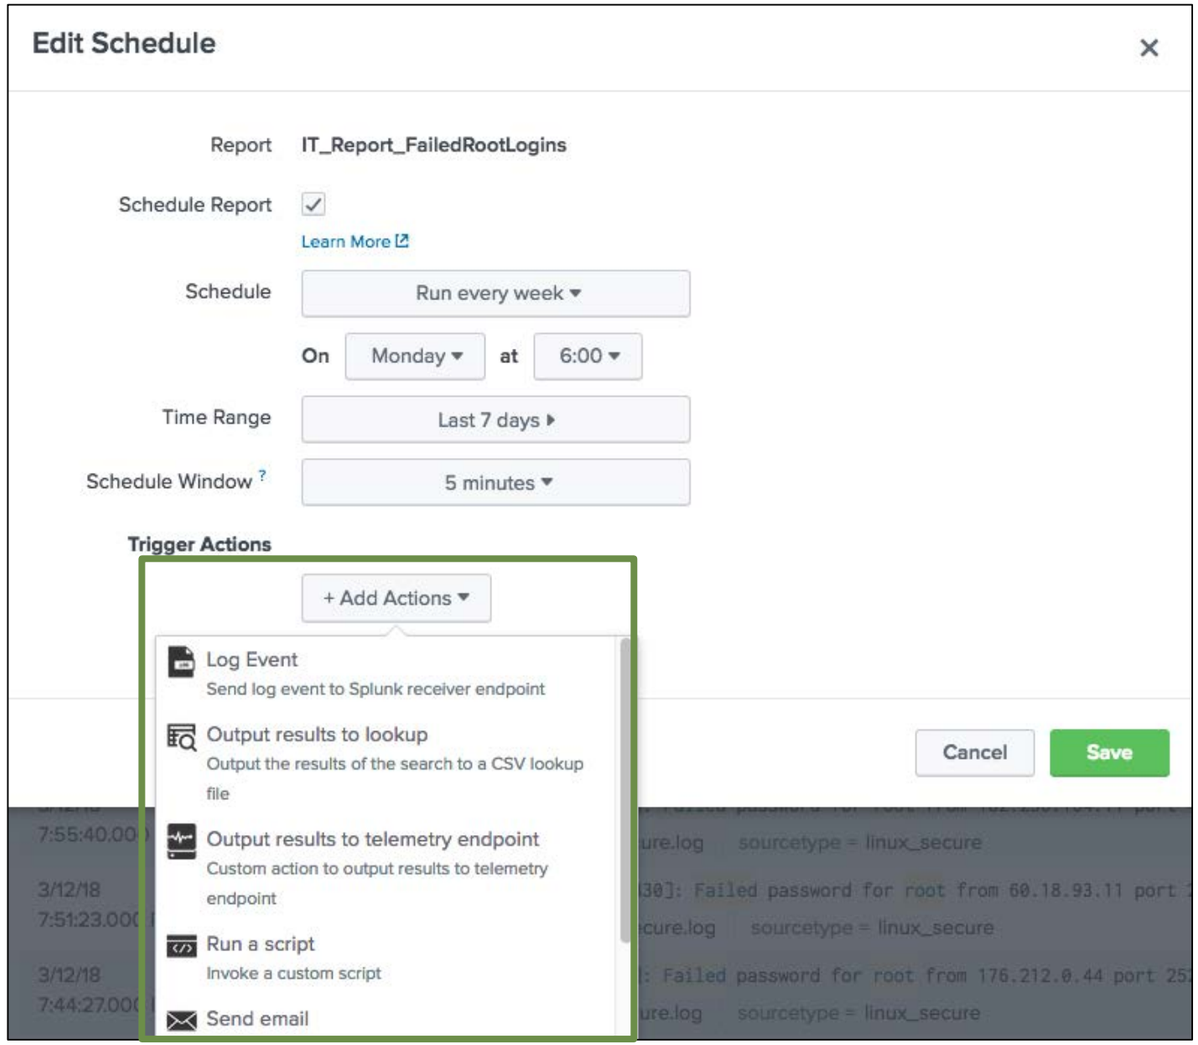

Creating a Scheduled Report – Send Email

1. Enter addresses in the To field, separated by commas
2. Set the priority
3. Edit or keep the default subject

   The $name$ variable includes the name of the report
4. If desired, include other options, such as an inline table of results
5. Define the email text type
6. Click Save

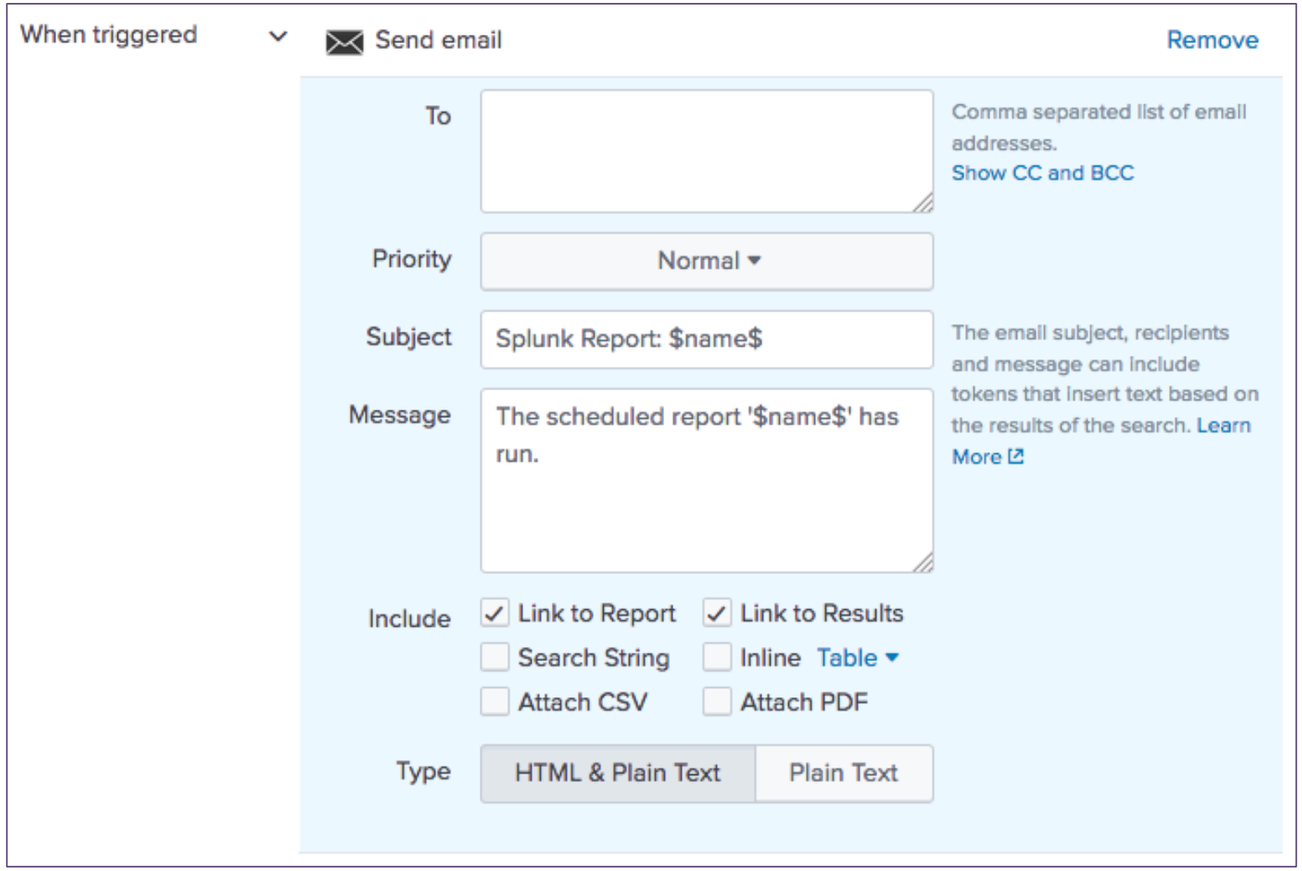

Managing Reports – Edit Permissions

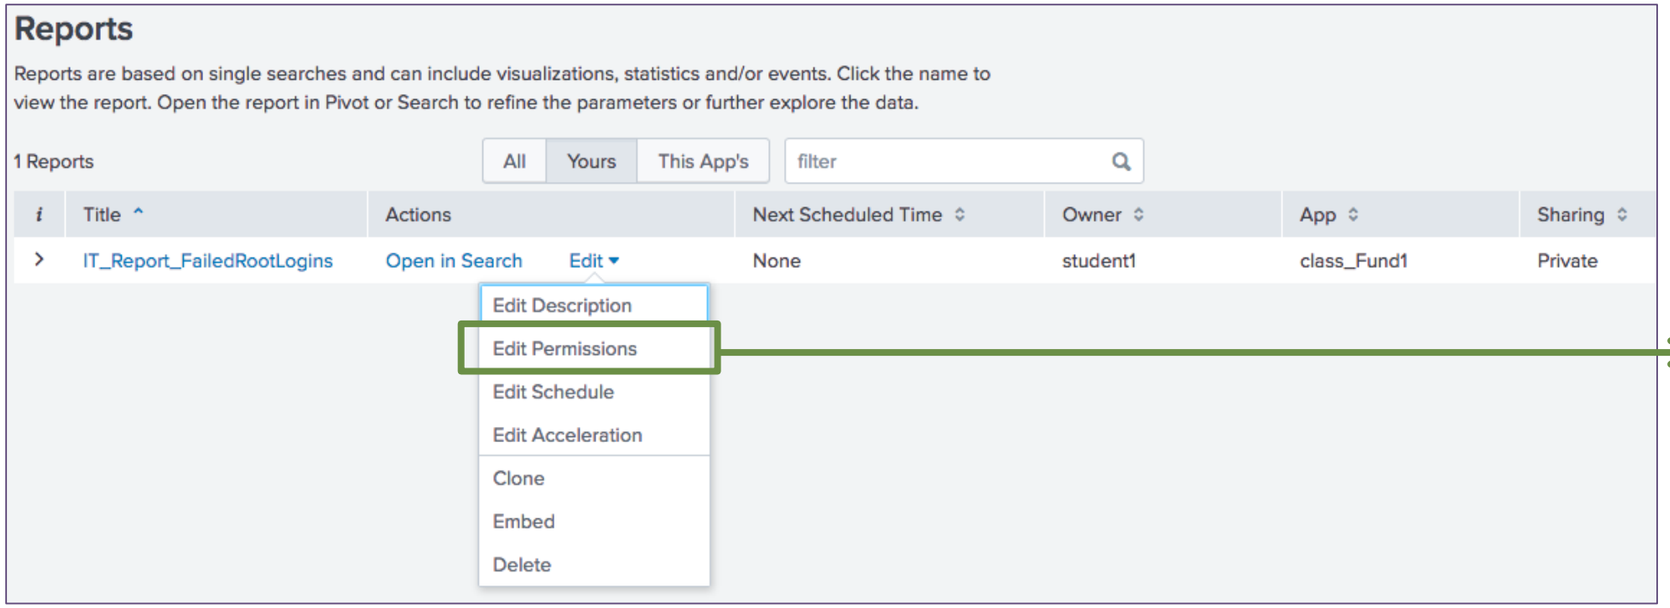

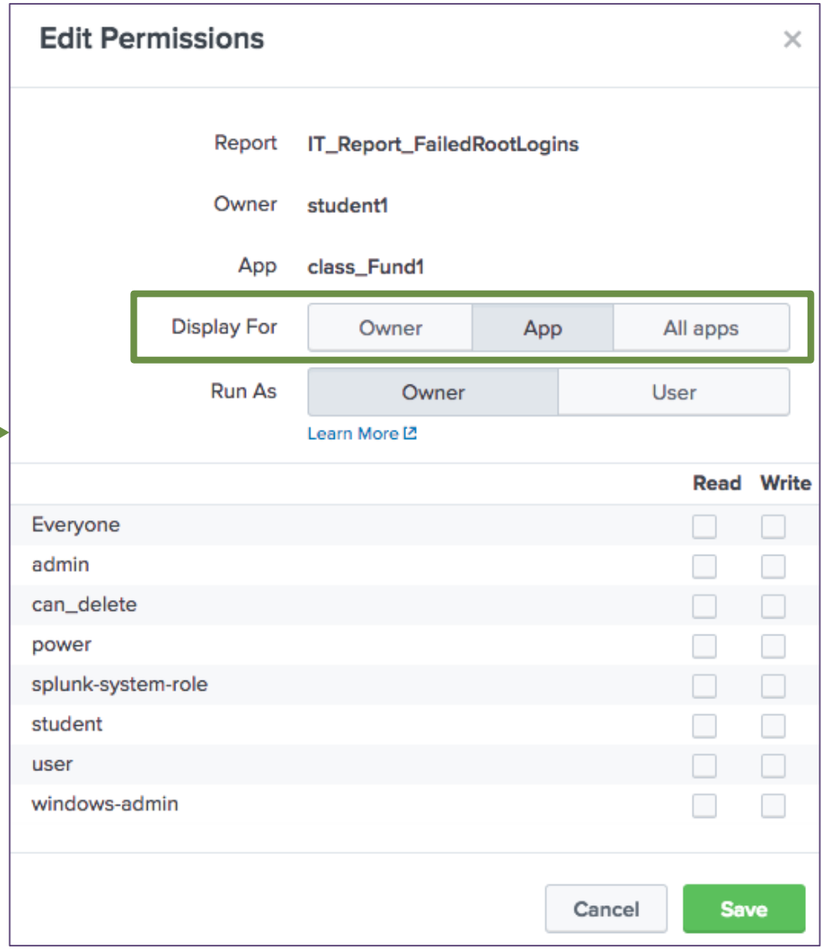

Note

The proper permissions from your Splunk administrator are required to edit the permissions on a scheduled report.

Managing Reports – Edit Permissions (cont.)

- Run As determines which user profile is used at run time
  - Owner – all data accessible by the owner appears in the report
  - User – only data allowed to be accessed by the user role appears

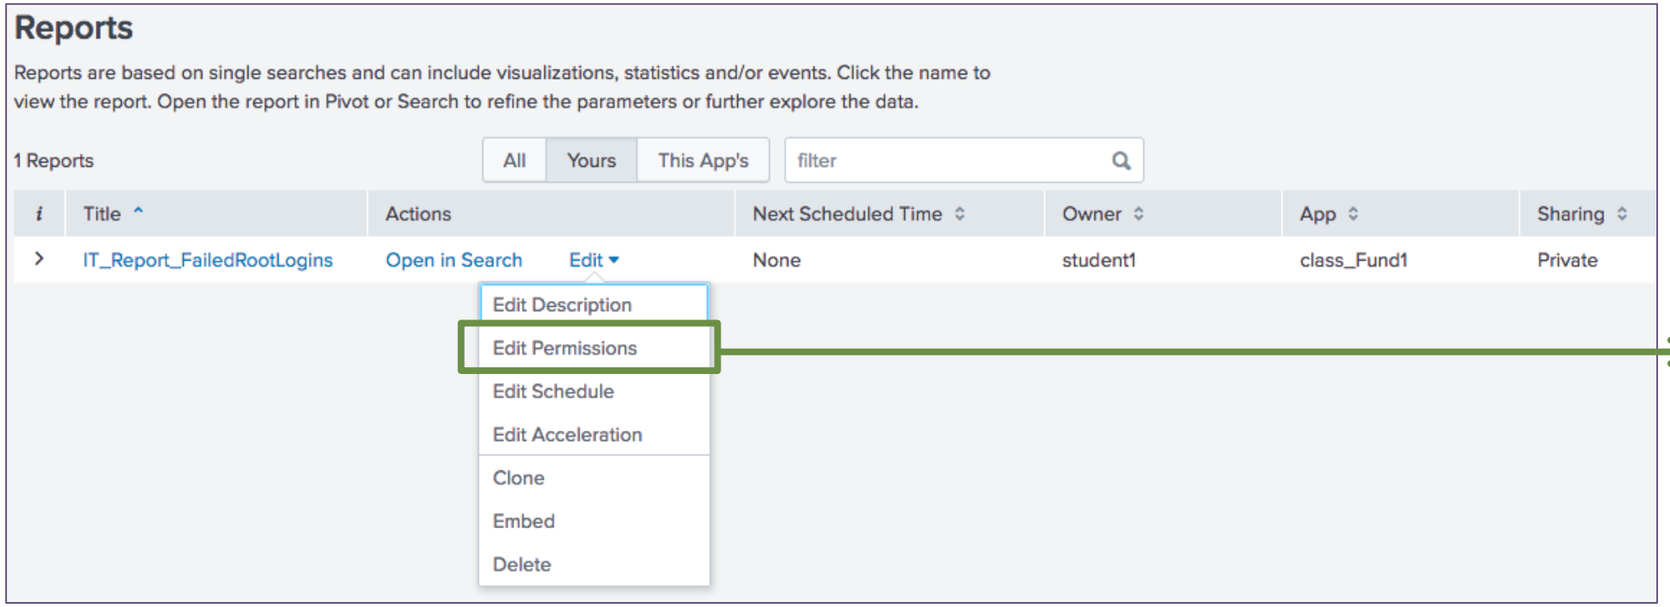

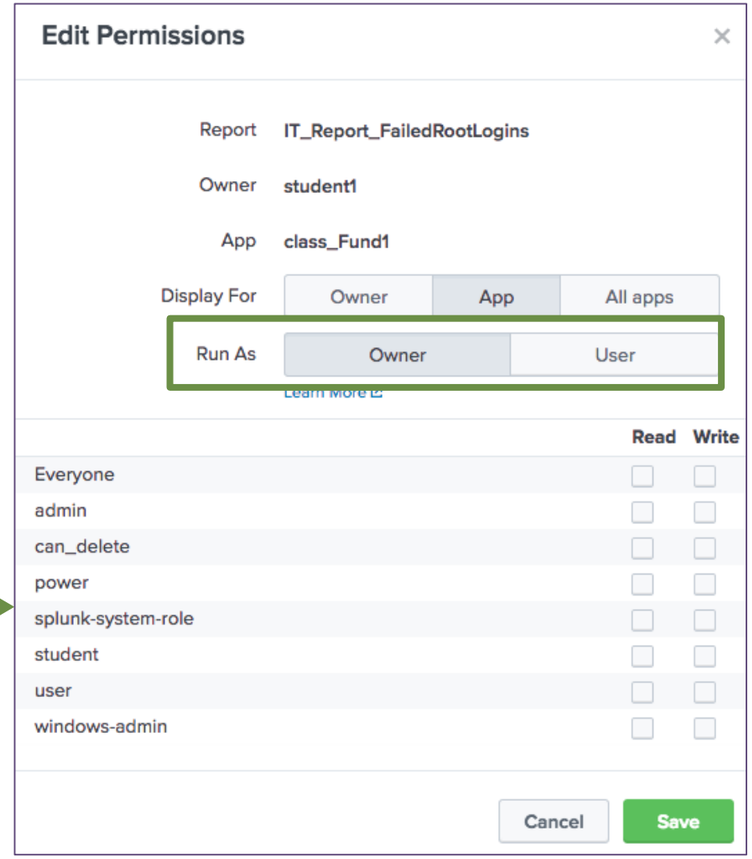

Managing Reports – Embed

- To access the report results from a webpage, click Edit > Embed
  - Before a report can be embedded, it must be scheduled

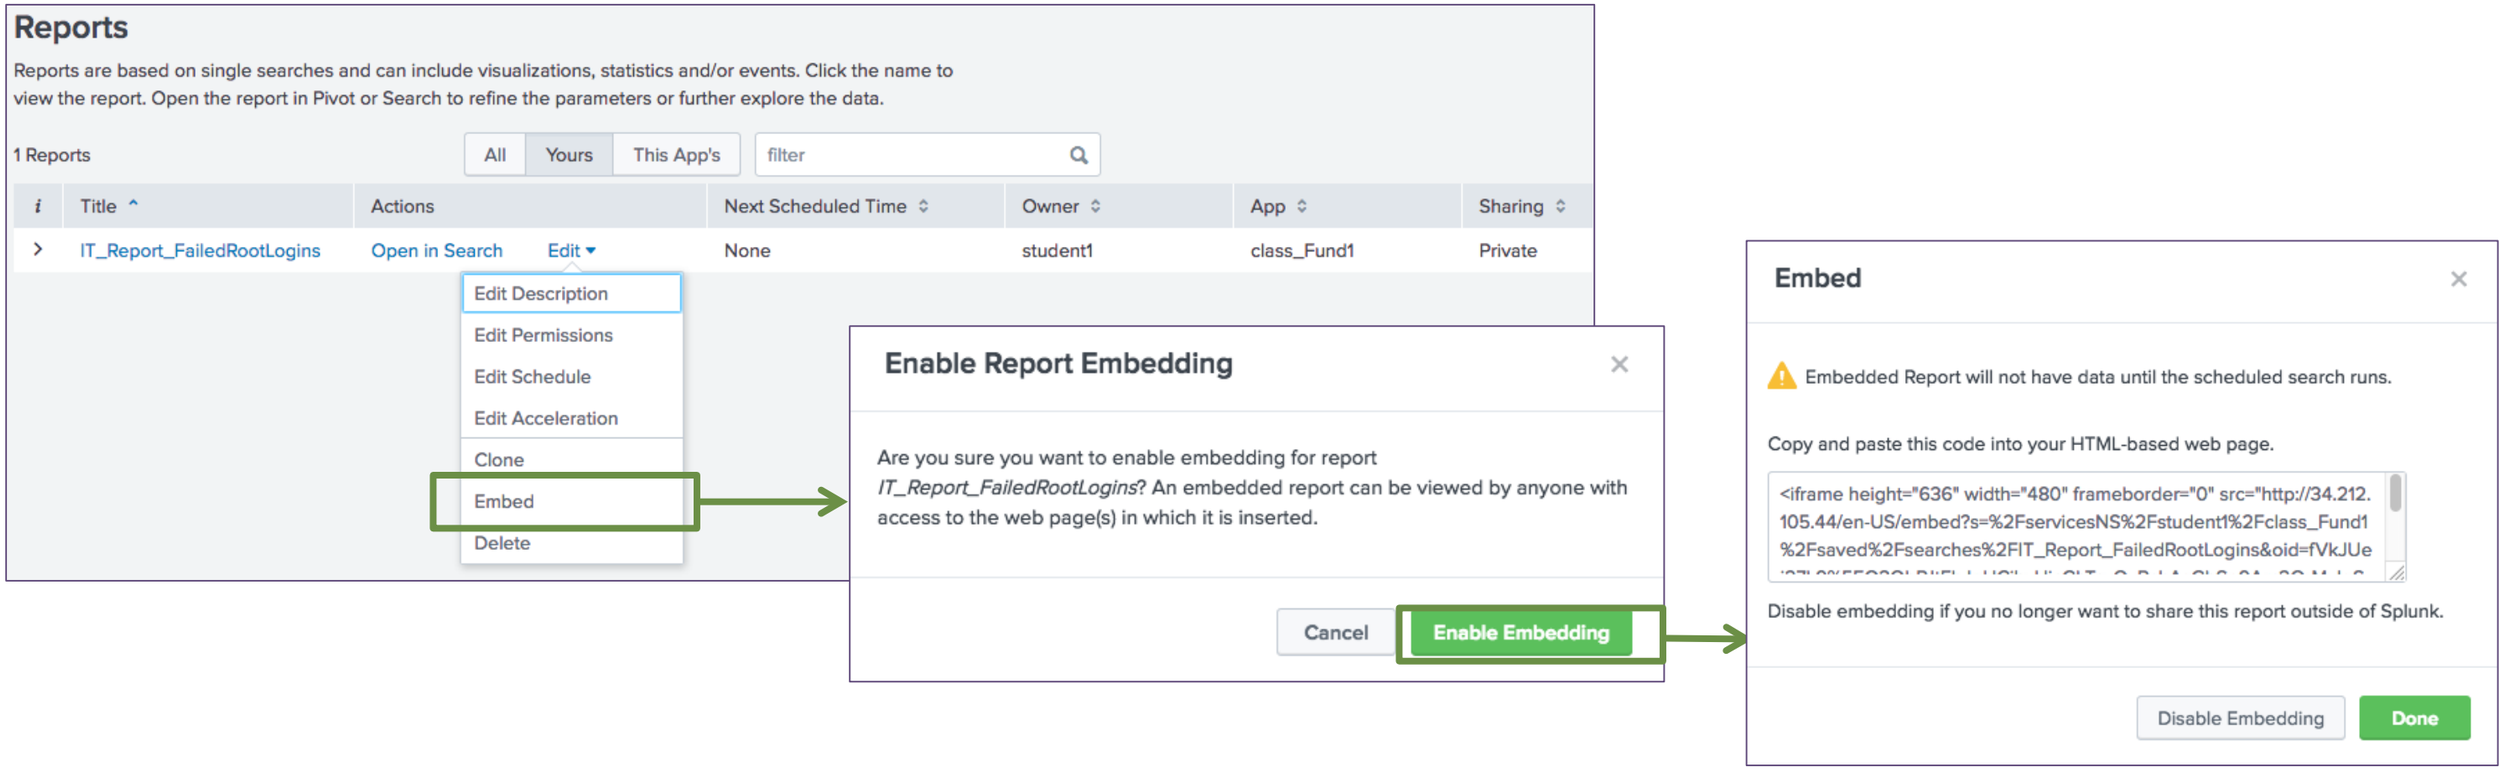

What Are Alerts?

- Splunk alerts are based on searches that can run either:
  - On a regular scheduled interval
  - In real-time
- Alerts are triggered when the results of the search meet a specific condition that you define
- Based on your needs, alerts can:
  - Create an entry in Triggered Alerts
  - Log an event
  - Output results to a lookup file
  - Send emails
  - Use a webhook
  - Perform a custom action

Creating an Alert

- Run a search
  - In this example, you're searching for server errors—any HTTP request status that begins with 50 over the last 5 minutes
- Select Save As > Alert
- Give the alert a Title and Description

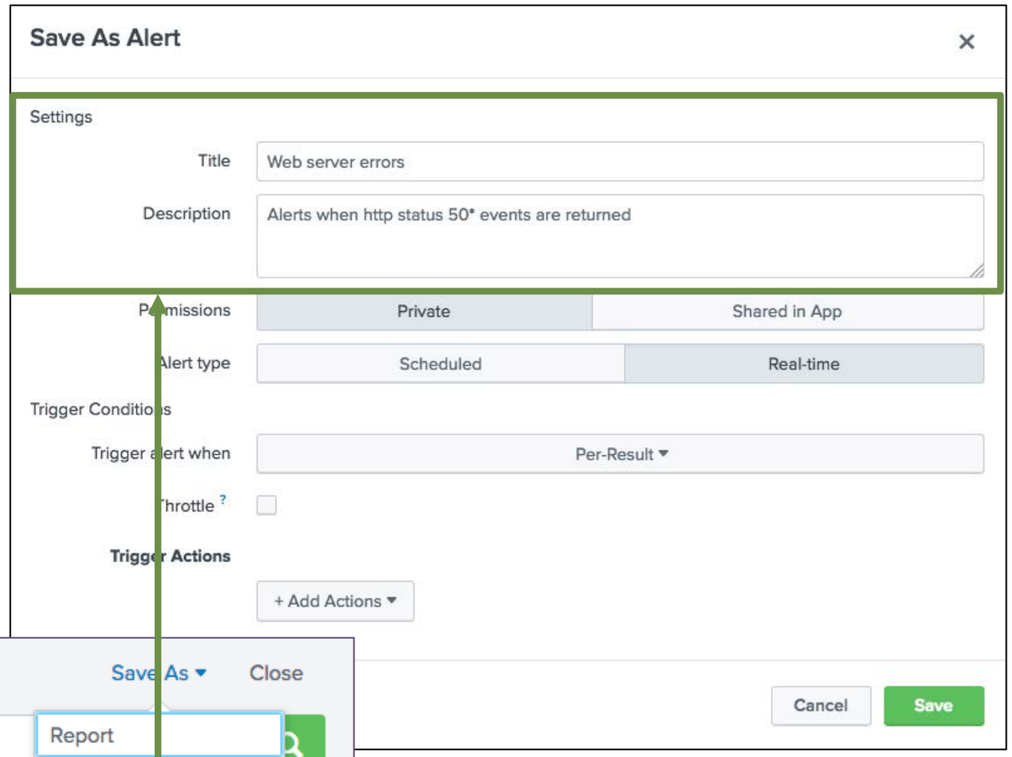

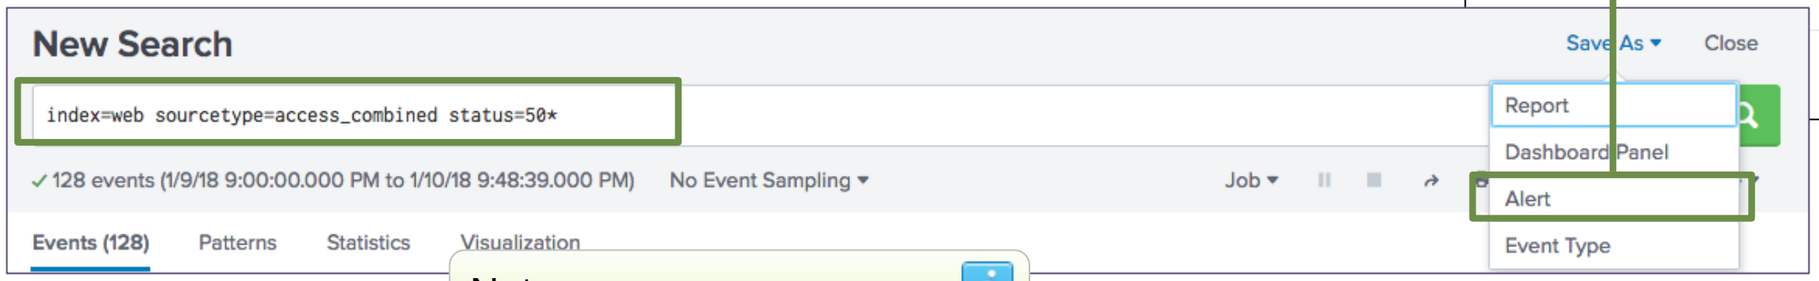

Note

This is the underlying search on which all the subsequent Alert slides are based.

Setting Alert Permissions

- Private – only you can access, edit, and view triggered alerts
- Shared in app
  - All users of the app can view triggered alerts
  - By default, everyone has read access and power has write access to the alert

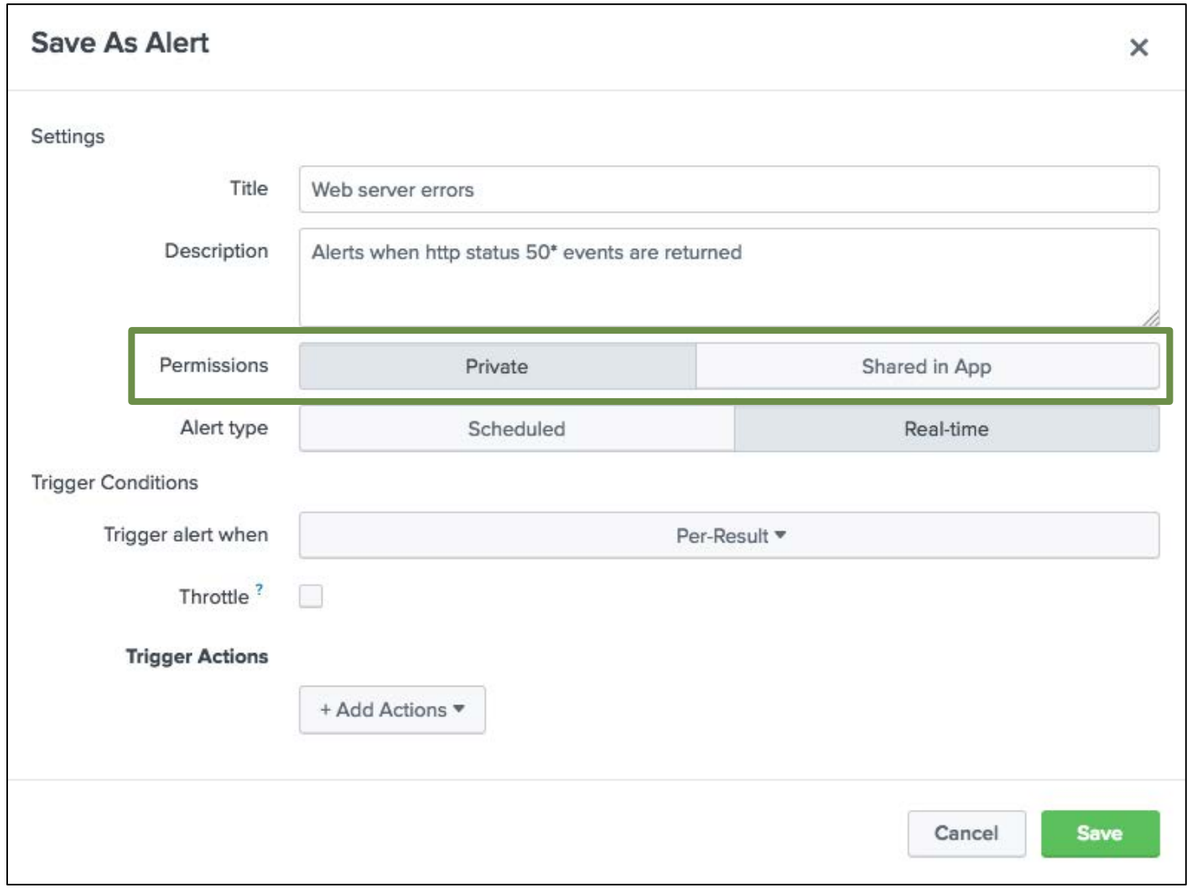

Note

The proper permissions from your Splunk administrator are required to set the permissions on an alert.

Choosing Real-time or Scheduled Alert Type

Choose an **Alert type** to determine how Splunk searches for events that match your alert

- **Scheduled** alerts
  - Search runs at a defined interval
  - Evaluates trigger condition when the search completes
- **Real-time** alerts
  - Search runs constantly in the background
  - Evaluates trigger conditions within a window of time based on the conditions you define

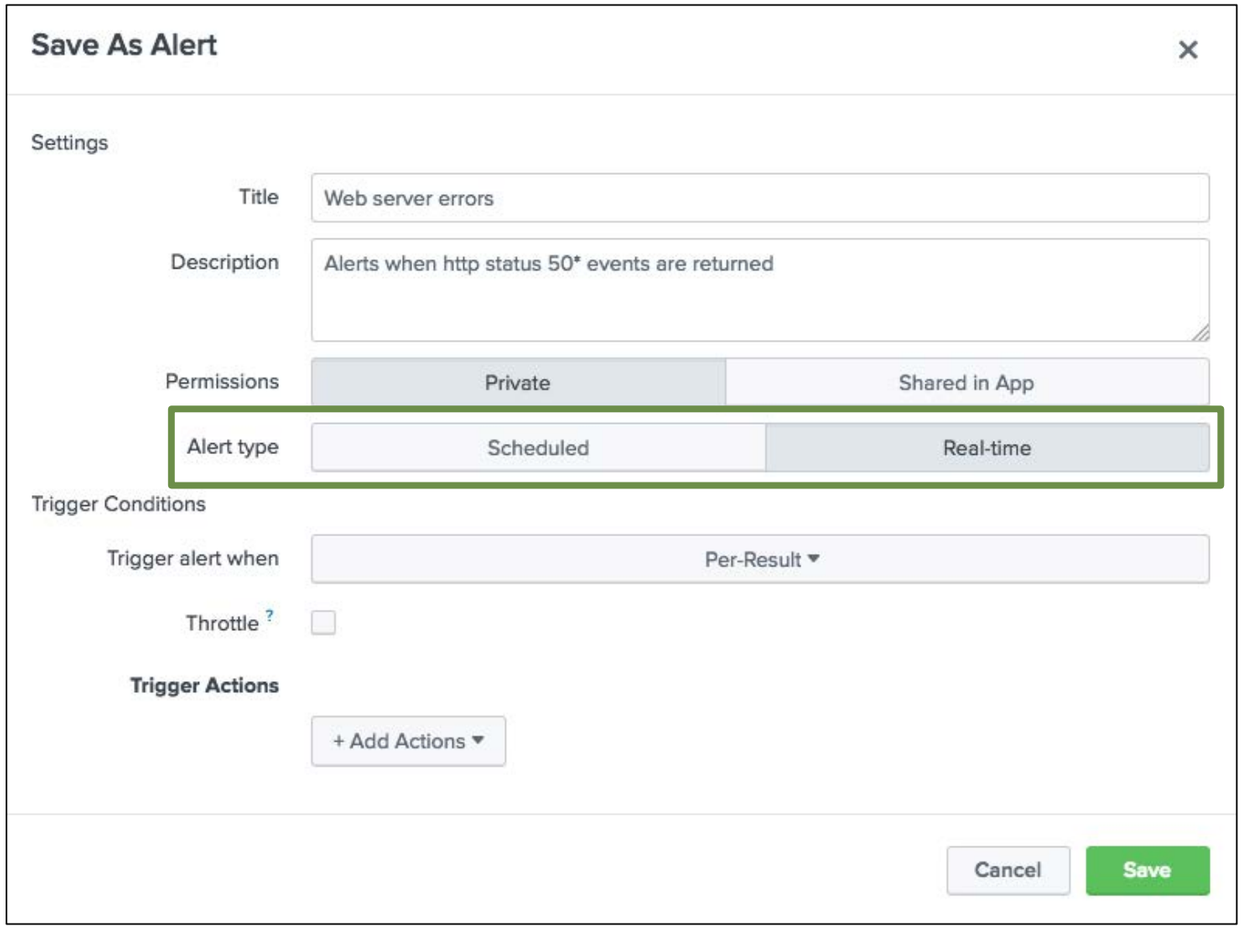

Setting the Alert Type – Scheduled

- From the frequency menu, choose to run the search every hour, day, week, month, or on a cron schedule
  - For the scheduled interval options, select the time the search will run
  - For cron schedule, define the cron expression

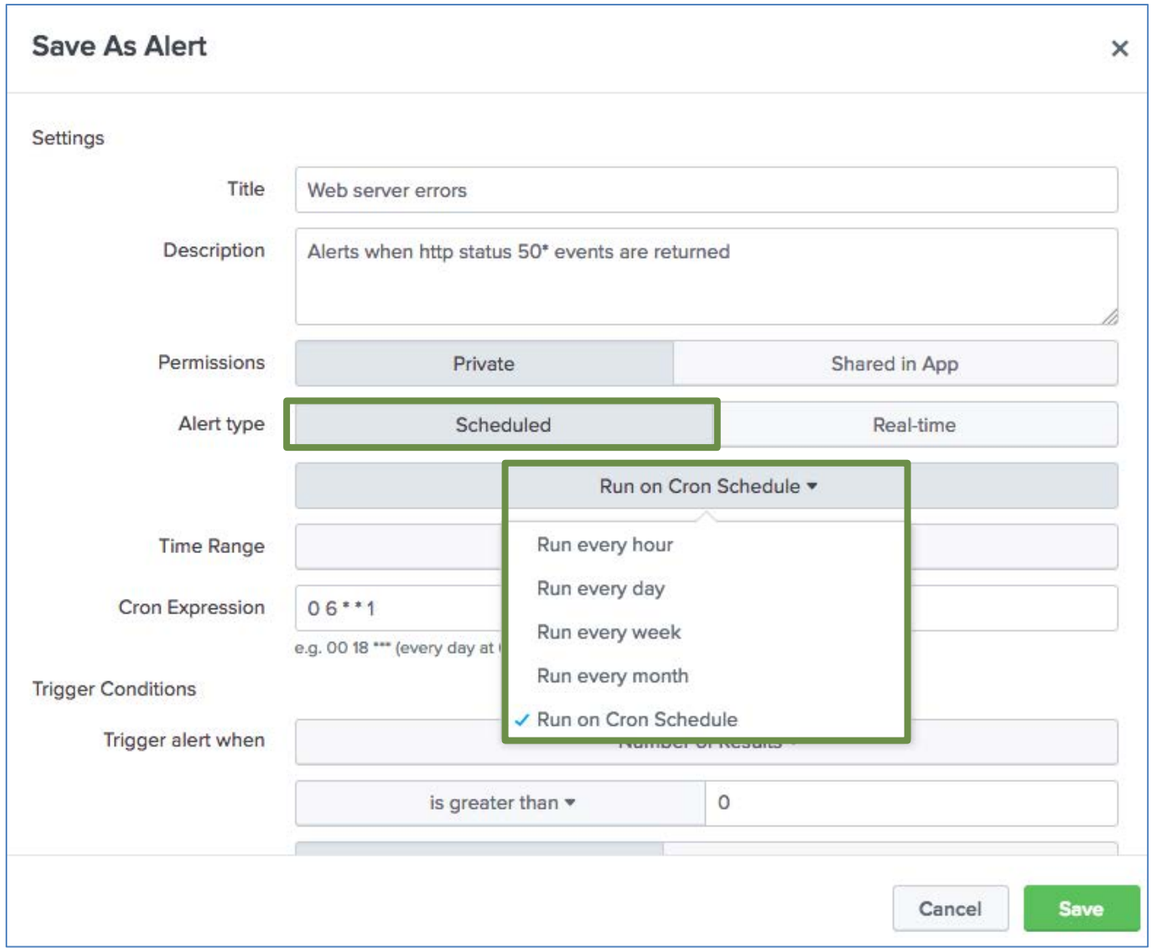

Setting Trigger Conditions – Scheduled

- For the cron schedule, choose a Time Range and enter a Cron Expression
- Set trigger conditions for scheduled alerts (same steps outlined for real-time alerts)
  - The alert examines the complete results set after the search is run

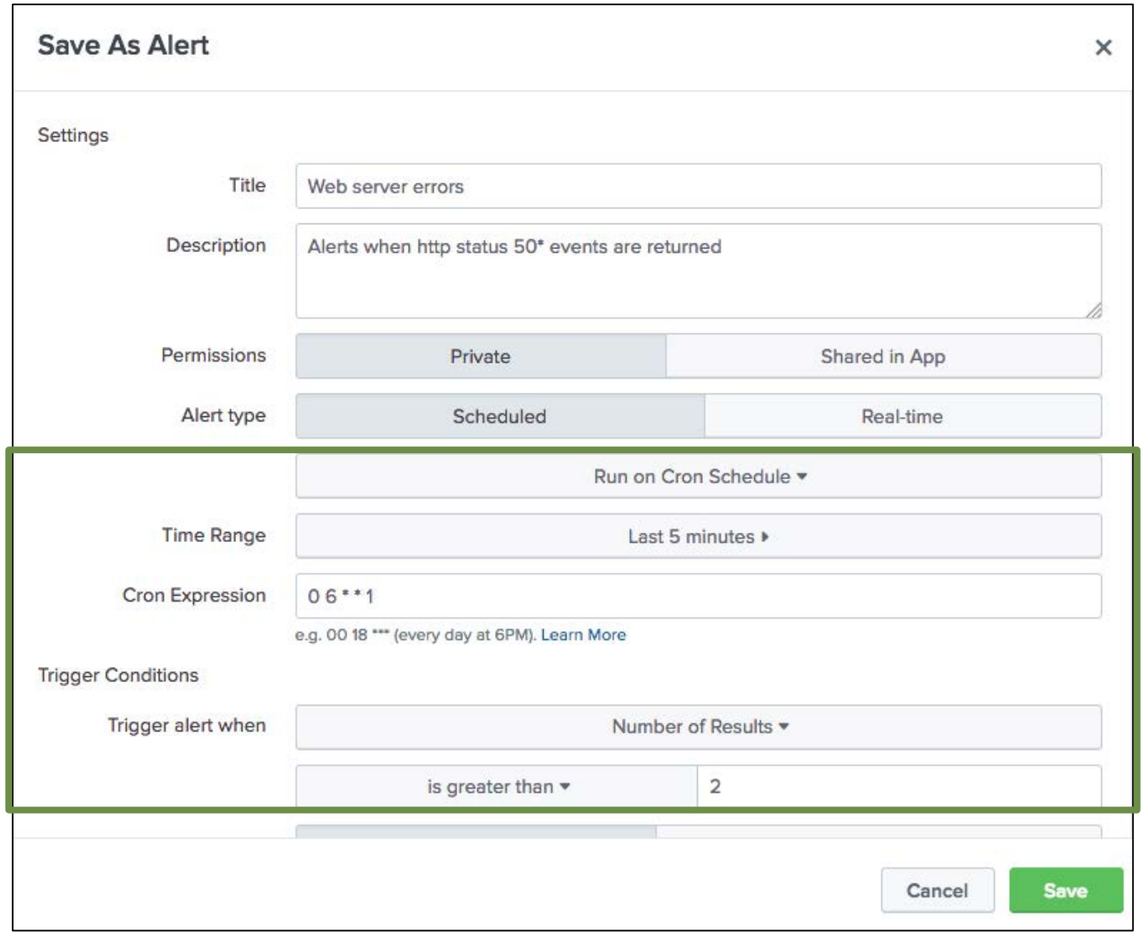

Scenario

In this example, a scheduled search will run every 5 minutes.

Setting Trigger Conditions – Real-time

- Trigger conditions allow you to capture a larger data set, then apply more stringent criteria to results before executing the alert
- You can set alerts to trigger:
  - **Per-Result –** triggers when a result is returned
  - **Number of Results** – define how many results are returned before the alert triggers
  - **Number of Hosts** – define how many unique hosts are returned before the alert triggers
  - **Number of Sources** – define how many unique sources are returned before the alert triggers
  - **Custom** – define custom conditions using the search language

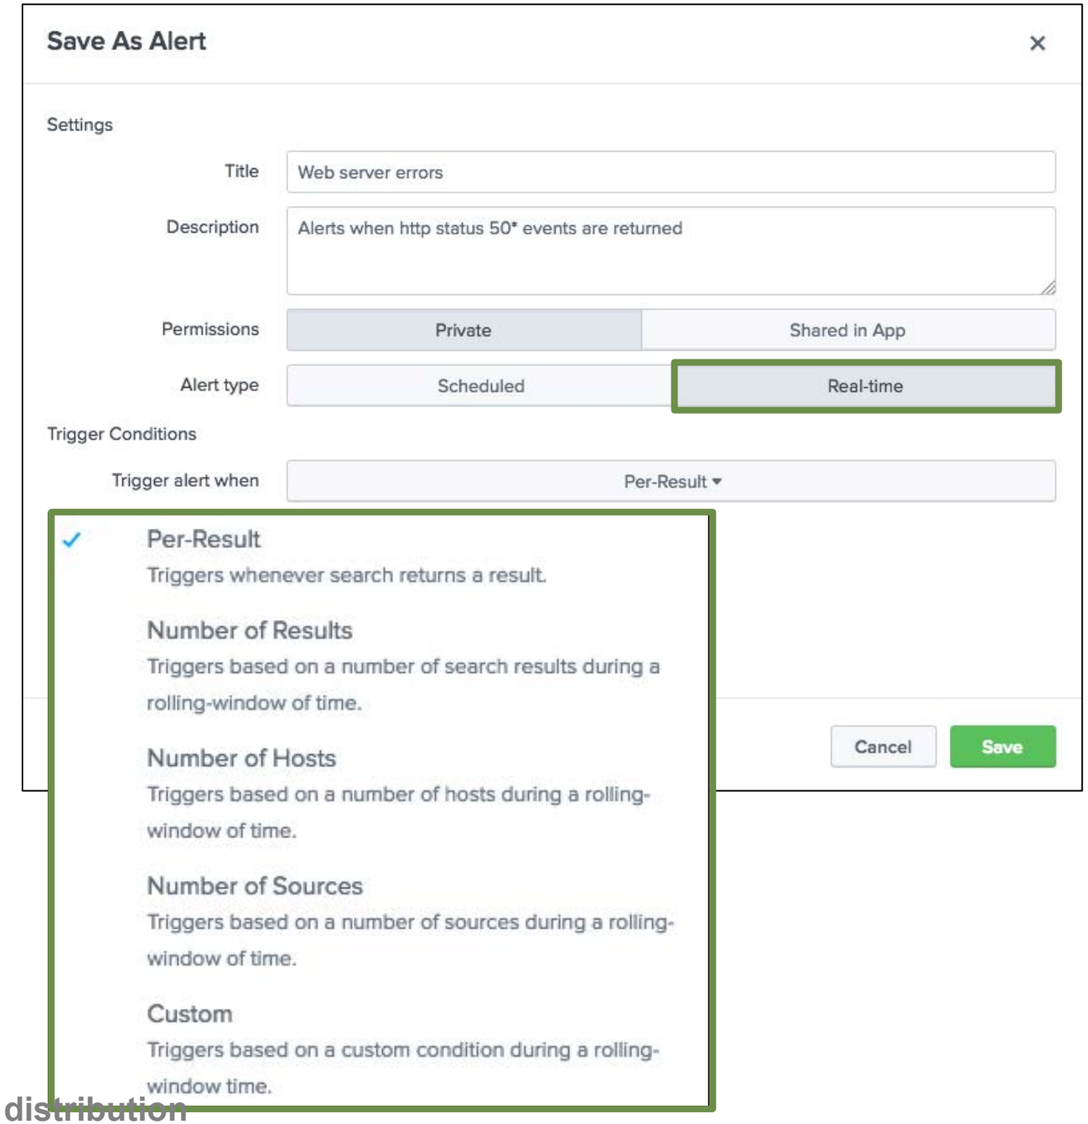

Setting Trigger Conditions – Real-time (cont.)

- In this example, the trigger condition is set to Number of Results
- In this Real-time alert example, if the number of results is greater than 2 within 1 minute, the alert triggers

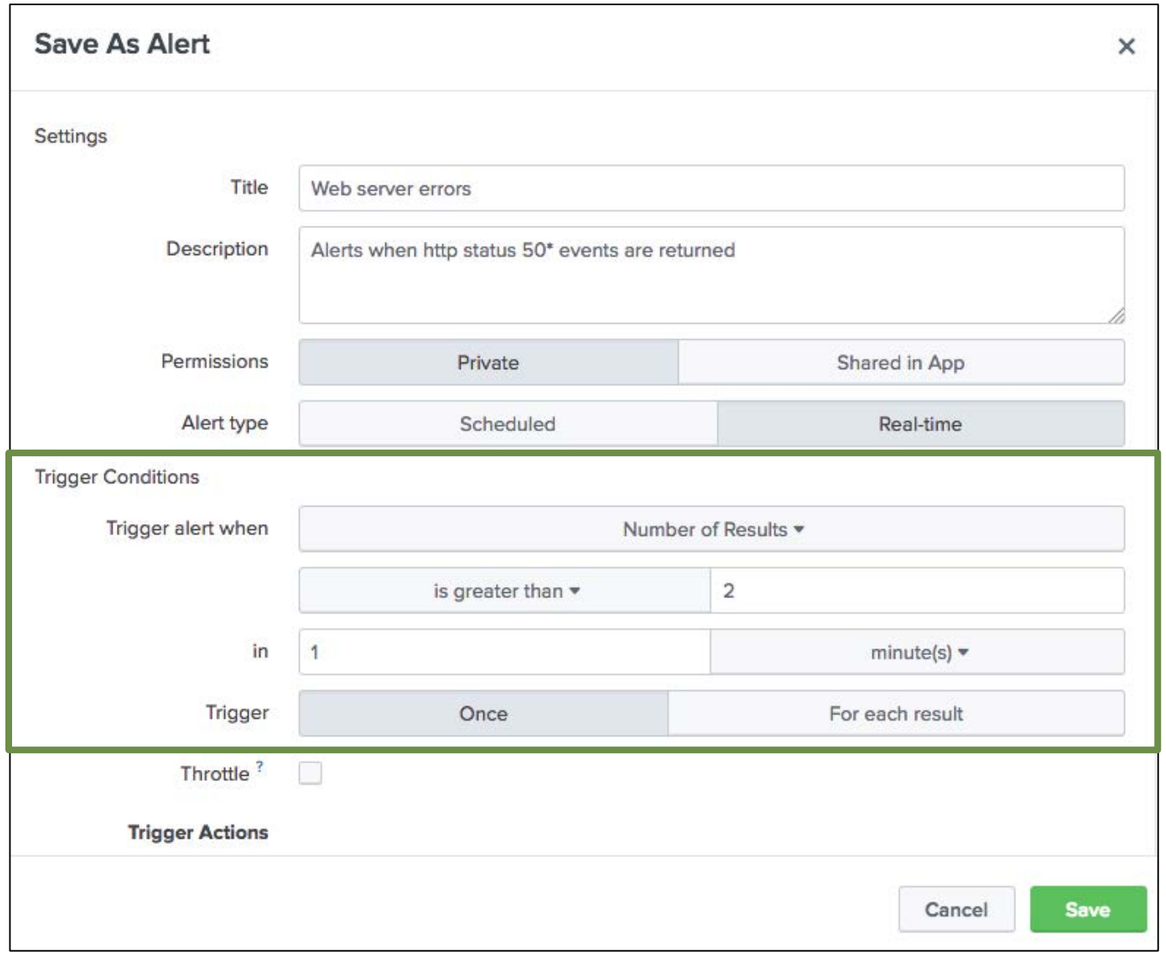

Note

The Number of Results setting does **not** determine how many actions associated with the alert are triggered. Rather, it sets a threshold to determine whether the alert is triggered in the first place.

Alert Actions – Trigger Conditions: Once

- **Once** executes actions *one time* for all matching events within the scheduled time and conditions
  - Example: If your alert is scheduled to run every 5 minutes, and 40 results are returned, the alert only triggers and executes actions one time
- Select the Throttle option to suppress the actions for results within a specified time range

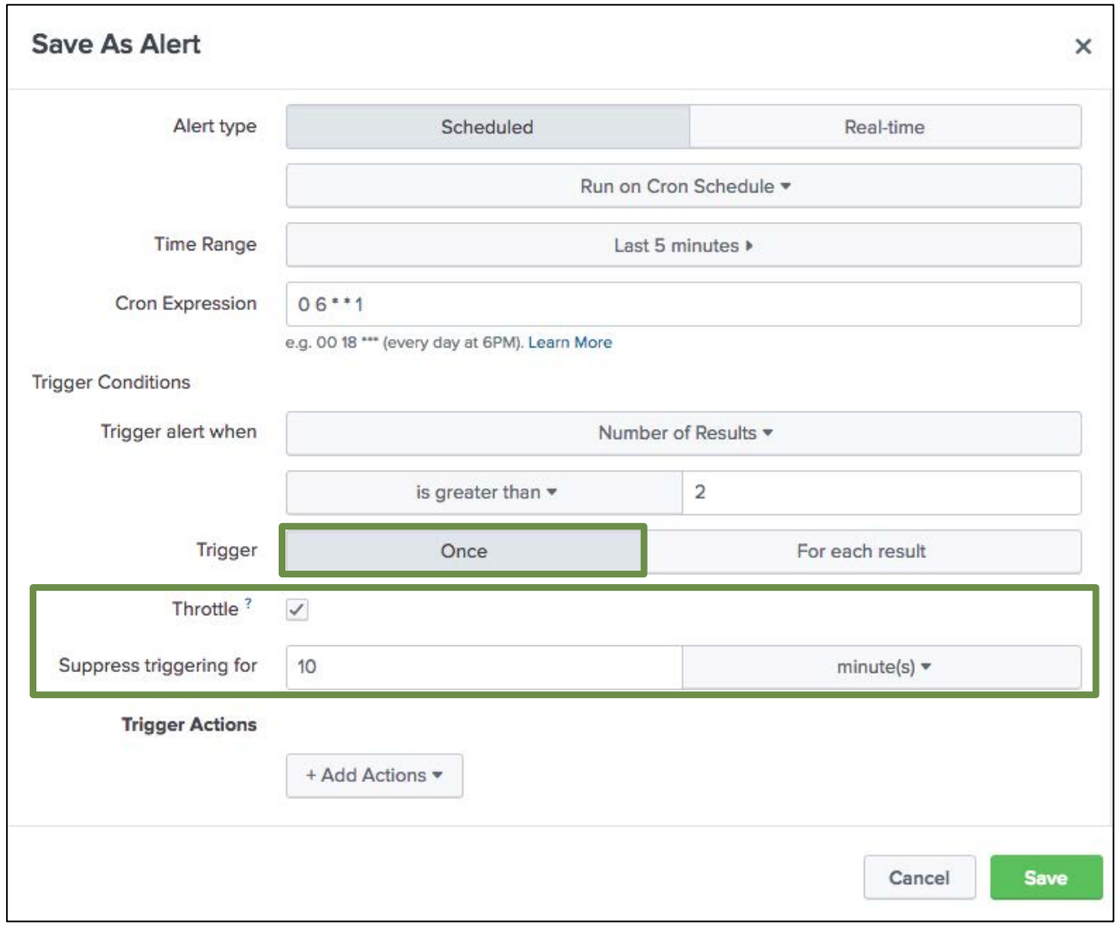

Alert Actions – Trigger Conditions: For Each Result

- **For each result** – executes the alert actions once *for each result* that matches the conditions
- Select the Throttle option to suppress the actions for results that have the same field value within a specified time range
  - Certain situations can cause a flood of alerts, when really you only want one
- In this example:
  - The search runs every 5 minutes
  - 70 events are returned in a 5 minute window—50 events with status=**500**, 20 with status=**503**
  - Since *For each result* is selected, **two actions** trigger—one for each status

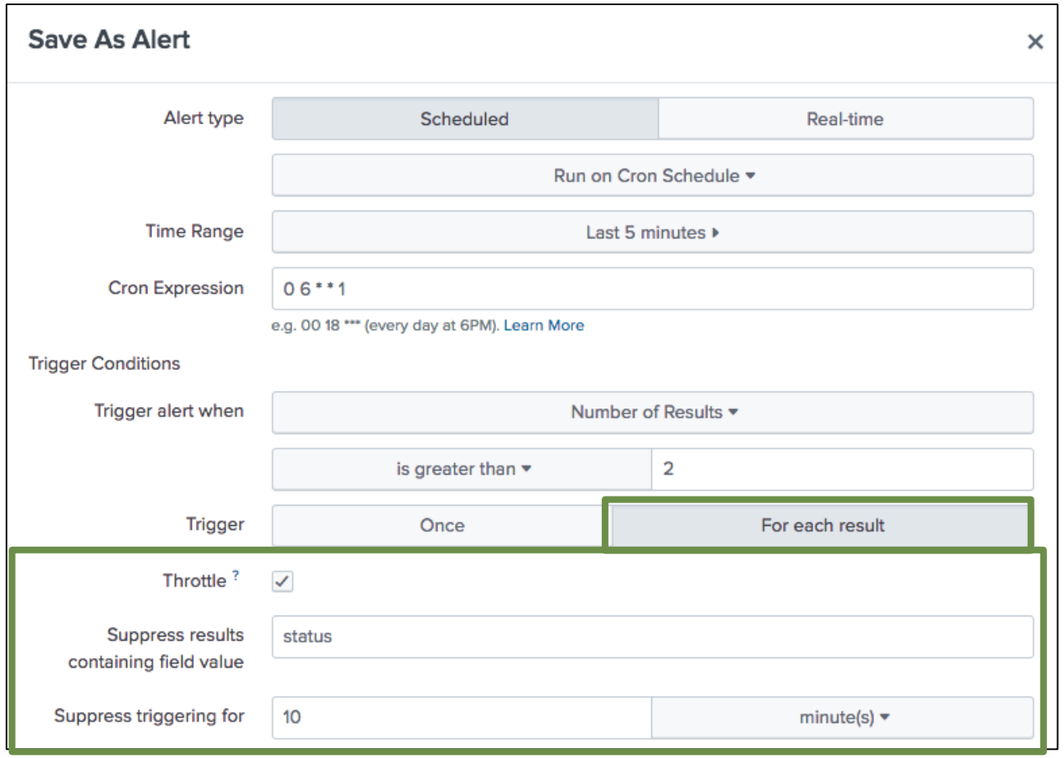

Add Trigger Actions

- Add to Triggered Alerts – adds the alert to the Activity > Triggered alerts list
- All actions available for scheduled reports are also available for alerts:
  - Log Event
  - Output results to lookup
  - Output results to telemetry endpoint
  - Run a script
  - Send email
  - Webhook

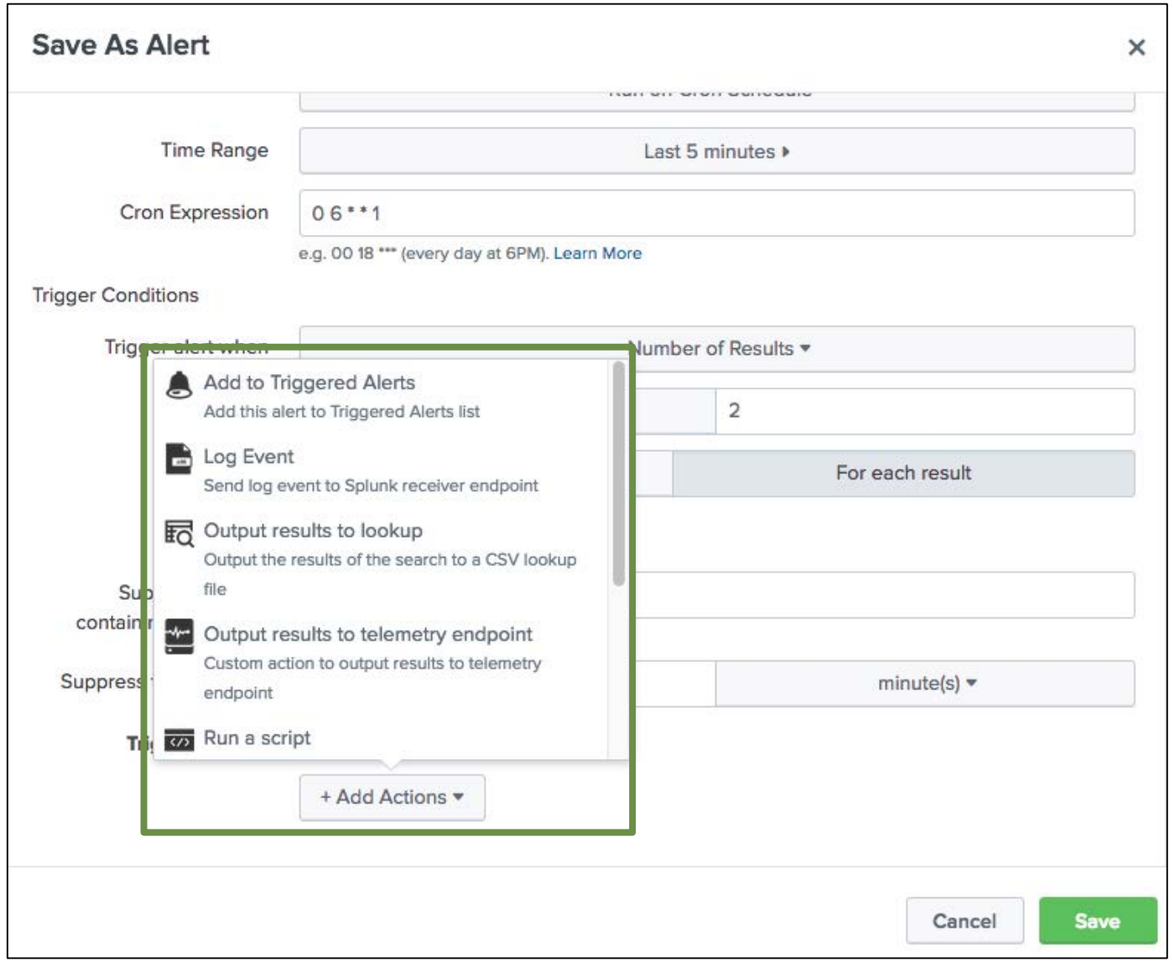

Alert Actions – Add to Triggered Alerts

Choose an appropriate severity for the alert

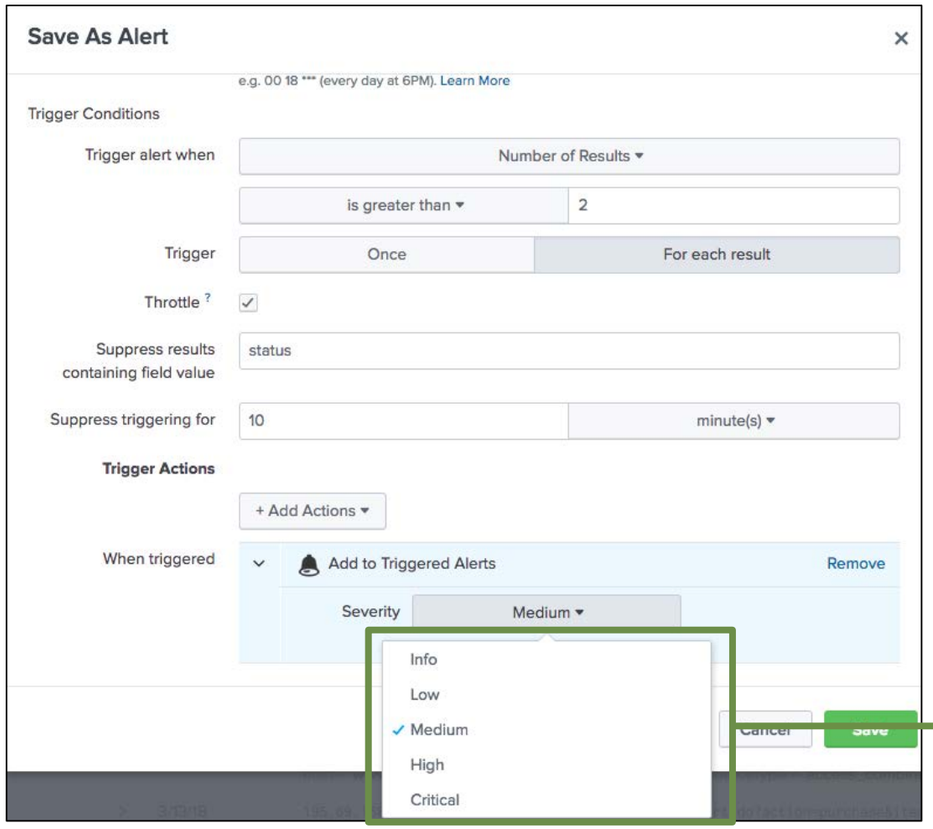

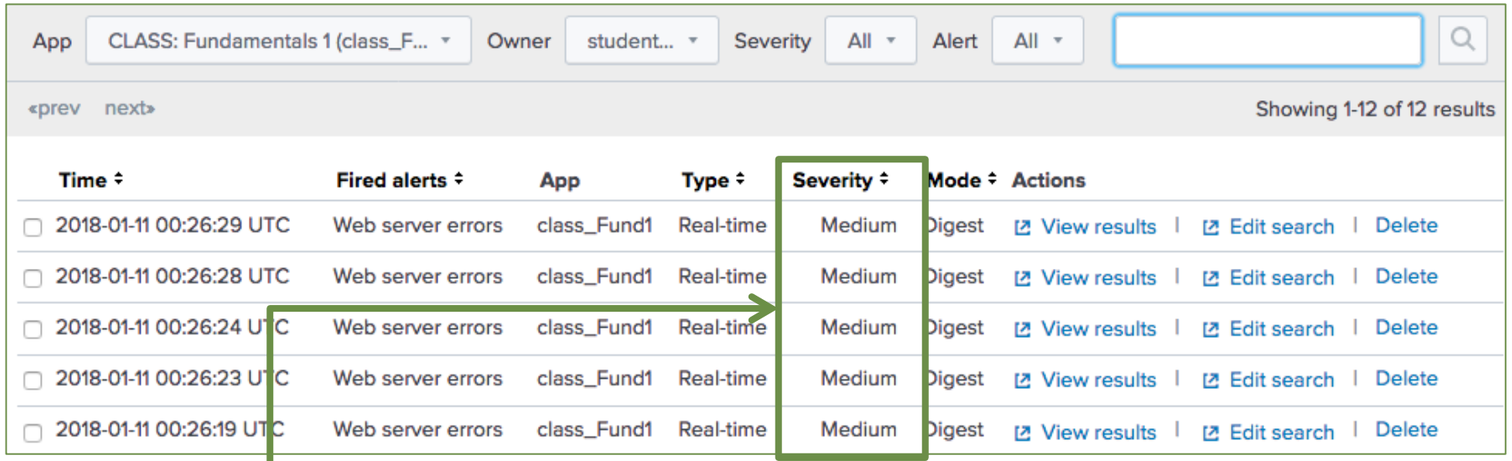

Alert Actions – Log Event

If you have administrator privileges, you can use a log event action

- Event – Enter the information that will be written to the new log event
- Source – Source of the new log event (by default, the alert name)
- Sourcetype – Sourcetype to which the new log event will be written
- Host – Host field value of the new log event (by default, IP address of the host of the alert)
- Index – Destination index for the new log event (default value is main)

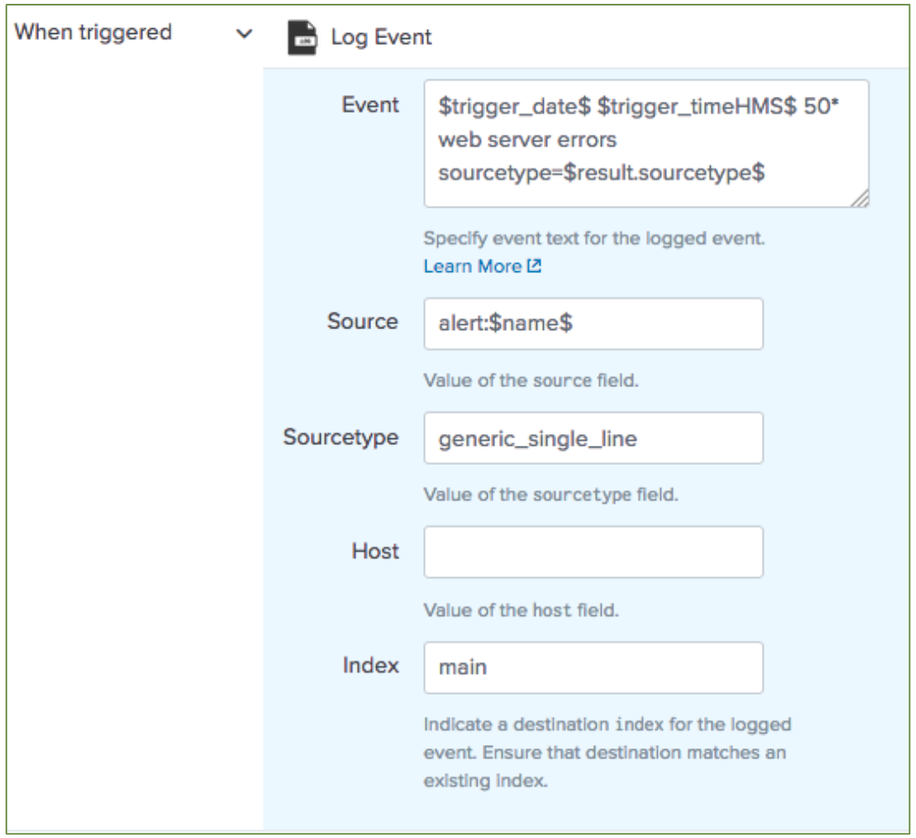

Note

For a complete list of available tokens, go to: http://docs.splunk.com/Documentation/Splunk/latest/Alert/EmailNotificationTokens

Alert Actions – Log Event (cont.)

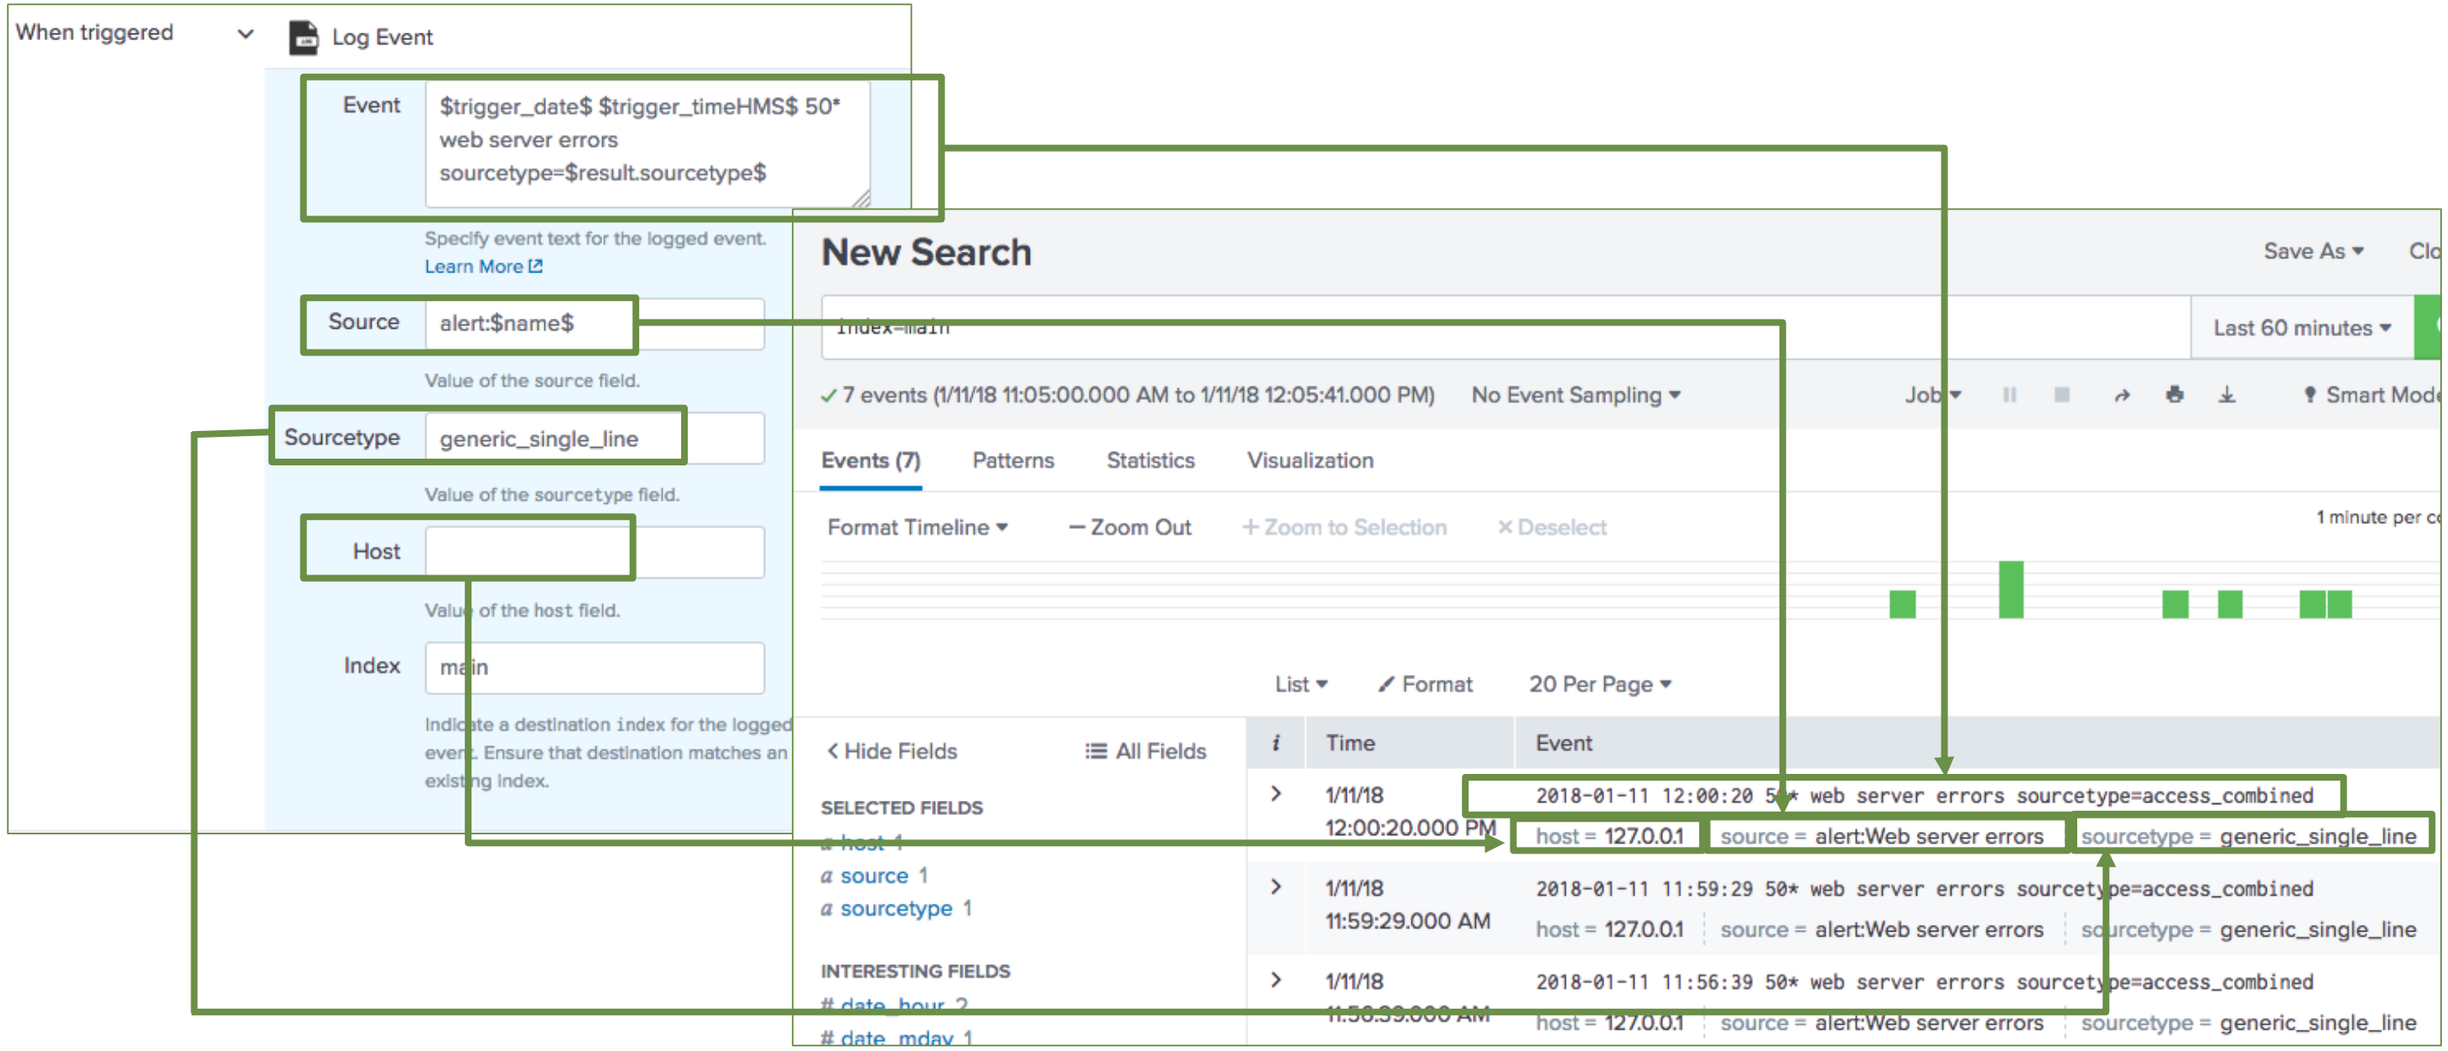

Alert Actions – Send Email

Customize the content of email alerts

- To - enter the email address(es) of the alert recipients
- Priority – select the priority
- Subject – edit the subject of the email (the $name$ token is the title of the alert)
- Message – provide the message body of the email
- Include – select the format of the alert
- Type – select the format of the text message

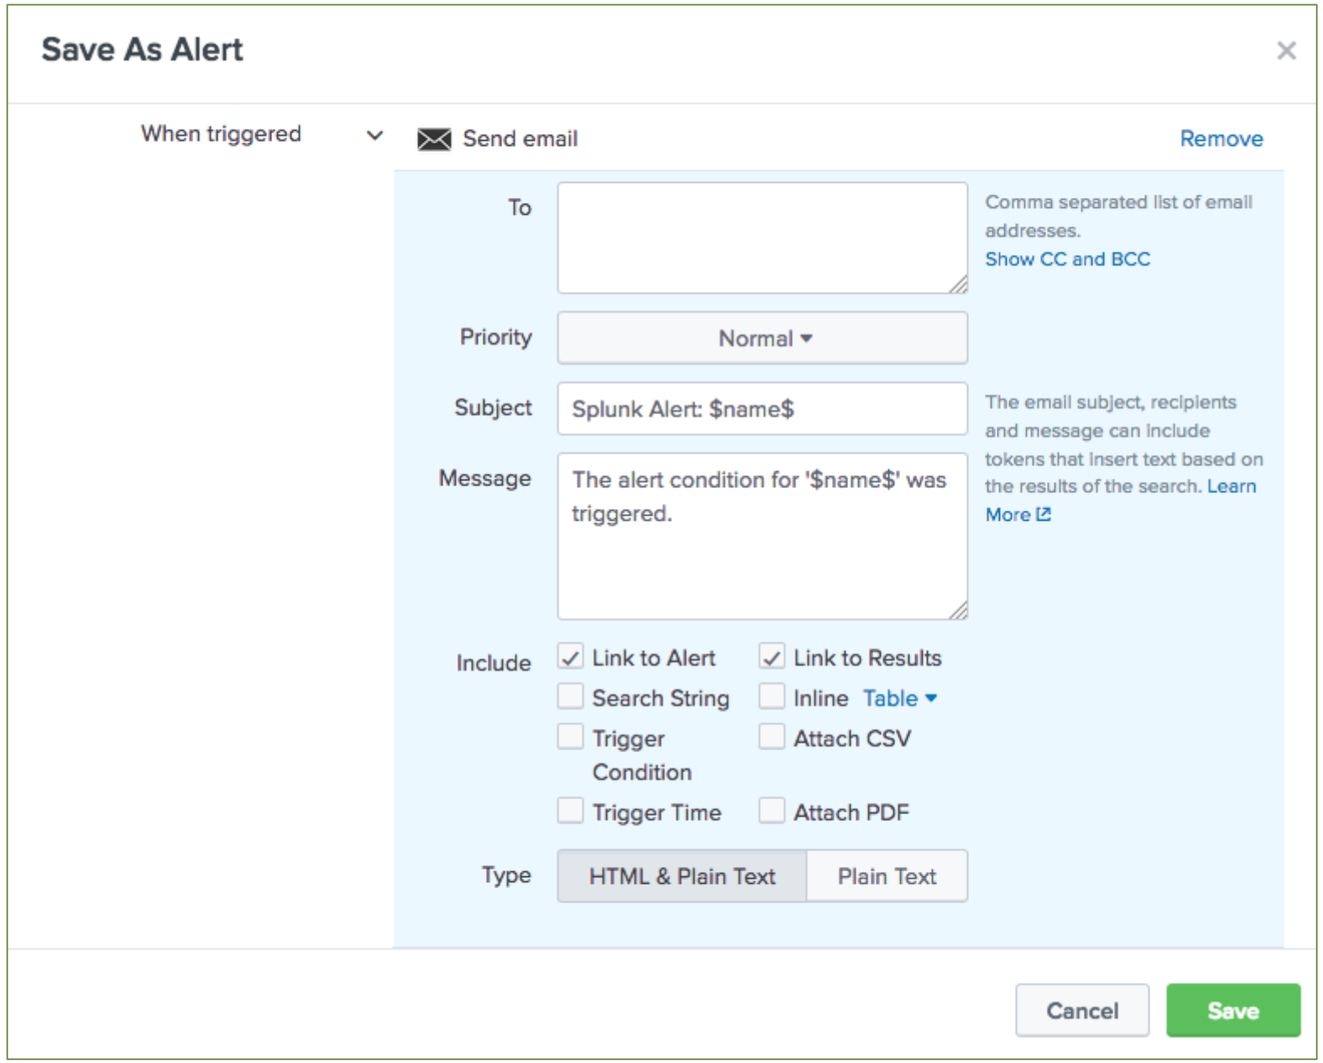

Viewing Triggered Alerts

- If you elected to list in triggered alerts, you can view the results by accessing **Activity > Triggered Alerts**
- Click **View results** to see the matching events that triggered the alert
- Click **Edit search** to modify the alert definition

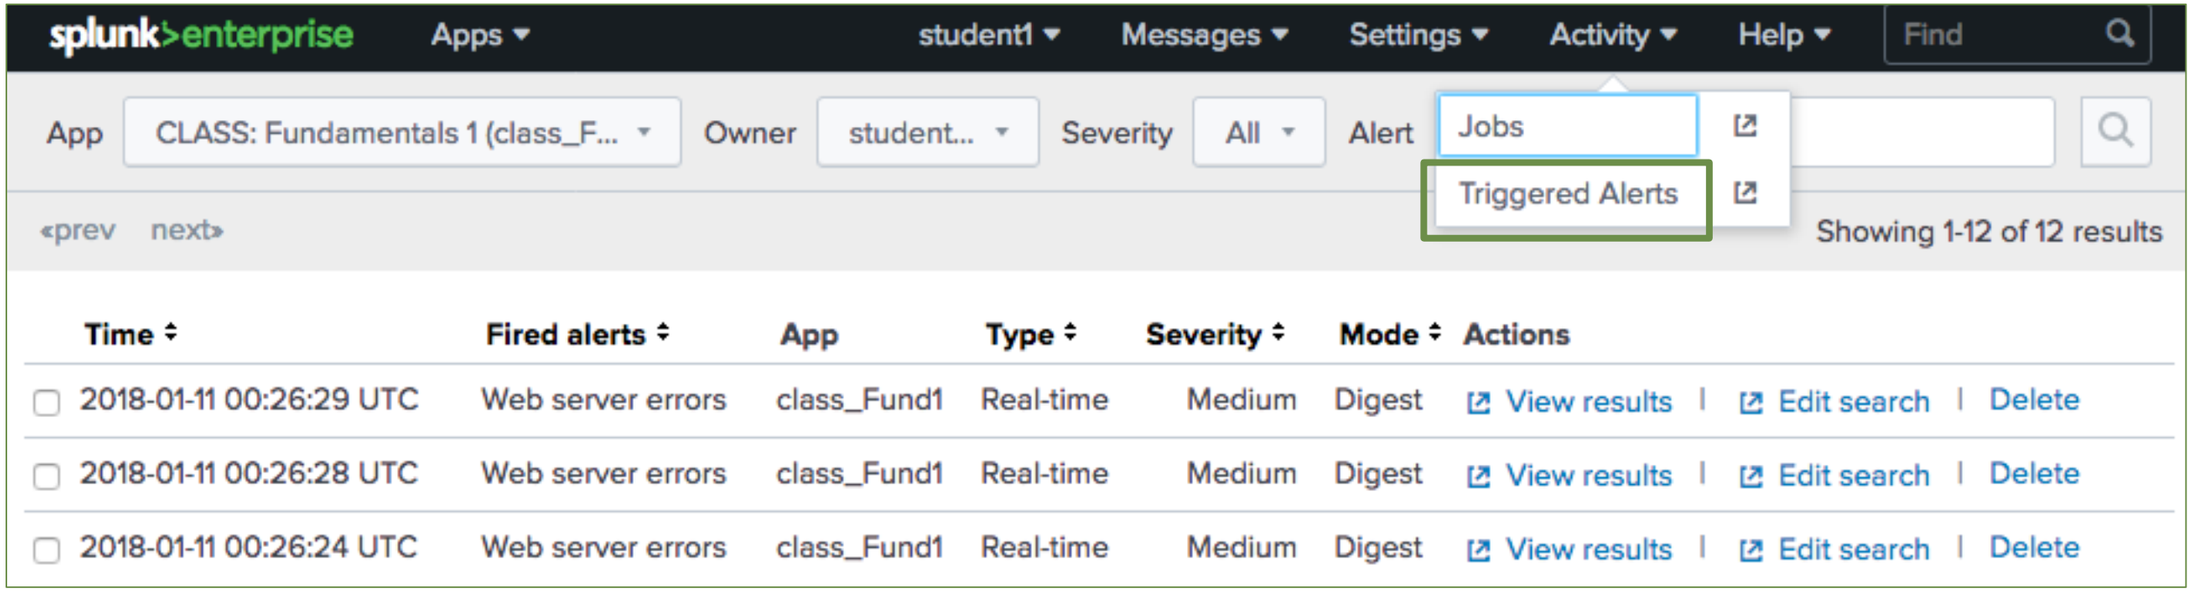

Editing Alerts

1. From the search bar, click **Alerts**
2. Select the alert and click **Edit**

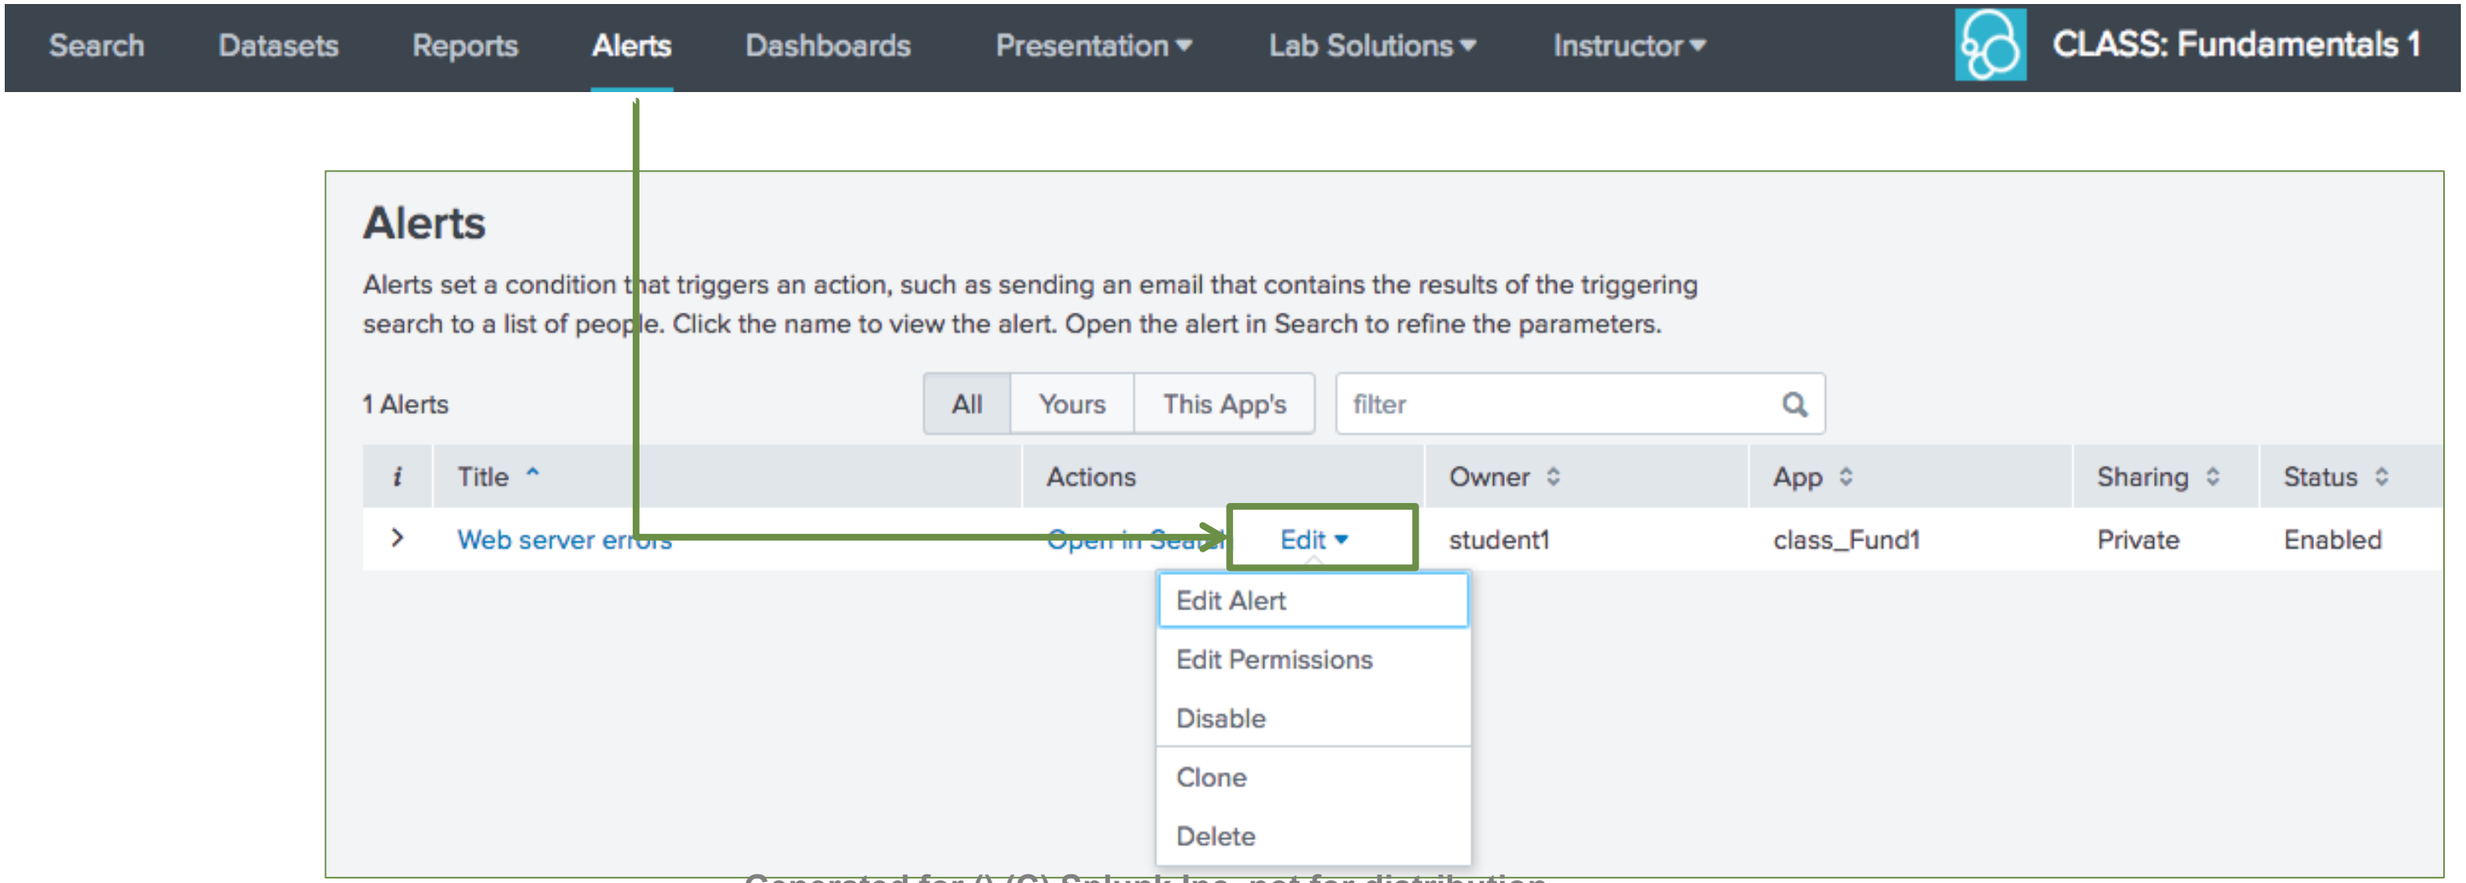

Editing Alert Permissions

- Edit permissions
  - Owner – only you can access, edit, and view triggered alerts
  - App – users of the app can access, edit, and view triggered alerts

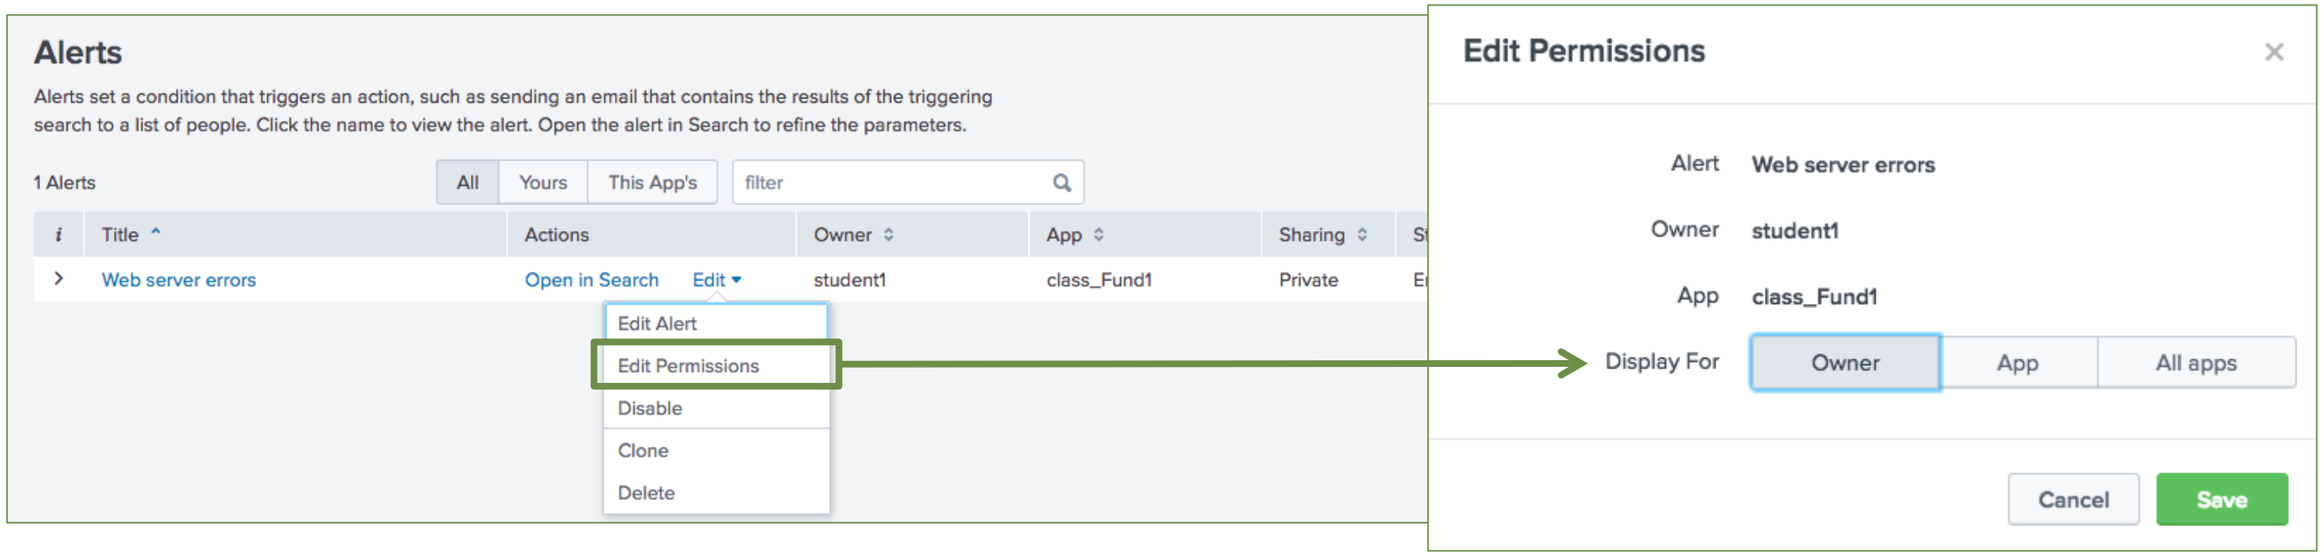

What's Next?: *Splunk Fundamentals 2*

- Use transforming commands and visualizations
- Filter and format the results of a search
- Correlate events into transactions
- Create and manage Knowledge Objects
- Create & manage extracted fields, field aliases, calculated fields
- Create tags and event types
- Create and use macros and workflow objects
- Create and manage data models
- Use the Splunk Common Information Model (CIM)

YouTube: The Splunk How-To Channel

- In addition to our roster of training courses, check out the Splunk Education How-To channel: http://www.youtube.com/c/SplunkHowTo
- This site provides useful, short videos on a variety of Splunk topics

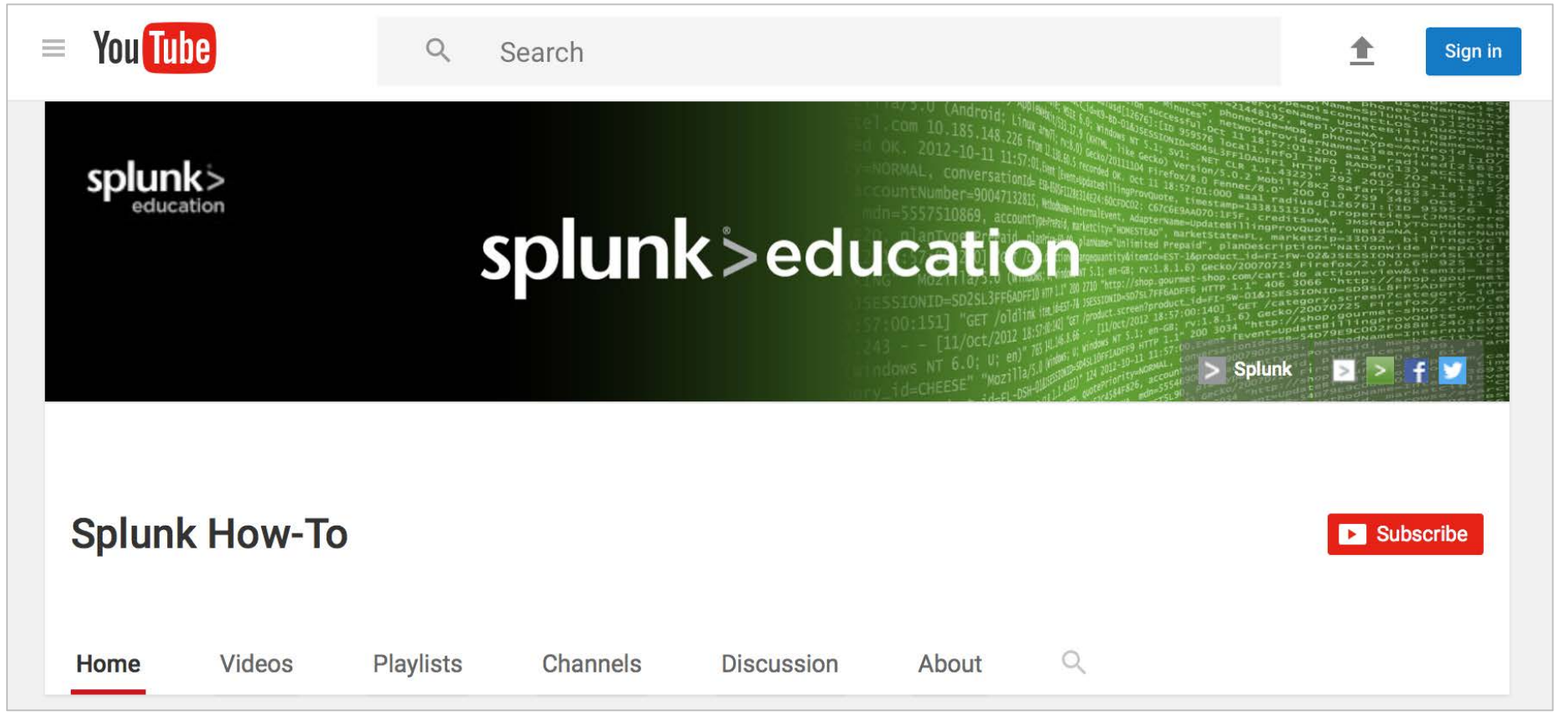

Support Programs

- **Community**
  - **Splunk Answers:** answers.splunk.com
    Post specific questions and get them answered by Splunk community experts.
  - **Splunk Docs:** docs.splunk.com
    These are constantly updated. Be sure to select the version of Splunk you are using.
  - **Wiki:** wiki.splunk.com
    A community space where you can share what you know with other Splunk users.
  - **IRC Channel:** #splunk on the EFNet IRC server Many well-informed Splunk users “hang out” here.
- **Global Support**
  Support for critical issues, a dedicated resource to manage your account – 24 x 7 x 365.
  - **Phone: (855) SPLUNK-S or (855) 775-8657**
  - **Web:** http://www.splunk.com/index.php/submit_issue
- **Enterprise Support**
  Access your customer support team by phone and manage your cases online 24 x 7 (depending on support contract.)

Other Resources

- Splunk App Repository
  https://splunkbase.splunk.com/
- Splunk Answers
  http://answers.splunk.com/
- Splunk Blogs
  http://blogs.splunk.com/
- Splunk Wiki
  http://wiki.splunk.com/
- Splunk Docs
  http://docs.splunk.com/Documentation/Splunk
- Splunk User Groups
  http://usergroups.splunk.com/