Behind the Magnifying Glass: How Search Works

Jeff Champagne

Staff Architect, Splunk

.conf2016

splunk>

Disclaimer

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

Who's This Dude?

Jeff Champagne

jchampagne@splunk.com
Staff Architect

- Started with Splunk in the fall of 2014
- Former Splunk customer in the Financial Services Industry
- Lived previous lives as a Systems Administrator, Engineer, and Architect
- Loves Skiing, traveling, photography, and a good Sazerac

Am I In The Right Place?

Some familiarity with…

- Splunk Components
  - Search Head, Indexer, Forwarder
- Splunk Search Interface
- Search Processing Language (SPL)

What Will I Learn?

1. What is going on when you click search
2. How to improve searches so they run faster

   - Splunk Architecture Overview
   - How Splunk stores events
   - Components of a search
   - Search tips and SPL command alternatives
   - Search command examples

Splunk Enterprise Architecture

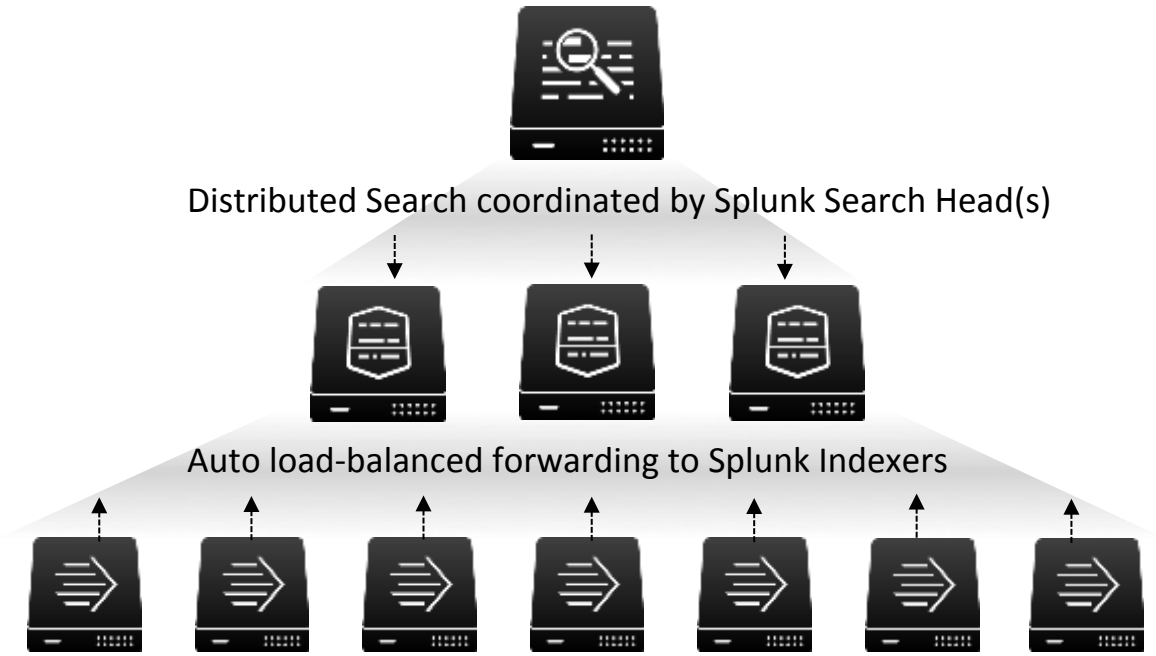

Index Vs. Index

An Overloaded Term

- Logical grouping for data
  - You or your Splunk admin create these
  - You reference these in your searches
    - Implicitly or explicitly

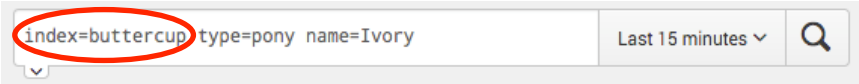

- TSIDX File
  - Time-series Index
  - Splunk's "secret sauce"
  - A logical Index is made up of many indexes/TSIDX files
  - This is how we search for your data
    - More on this later…

How Are Events Stored?

Buckets, Indexes, and Indexers

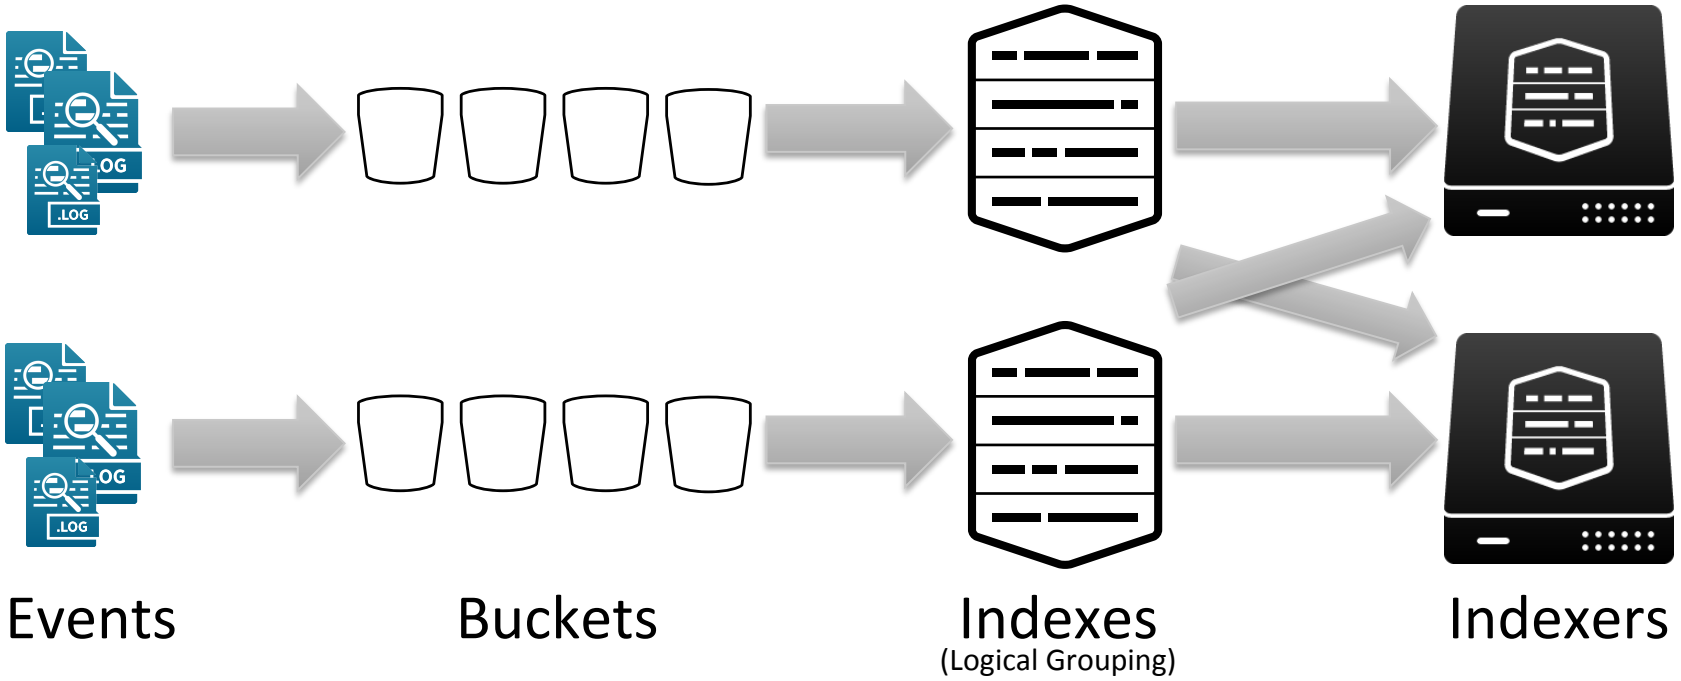

How Are Events Stored?

Bucket Aging Process

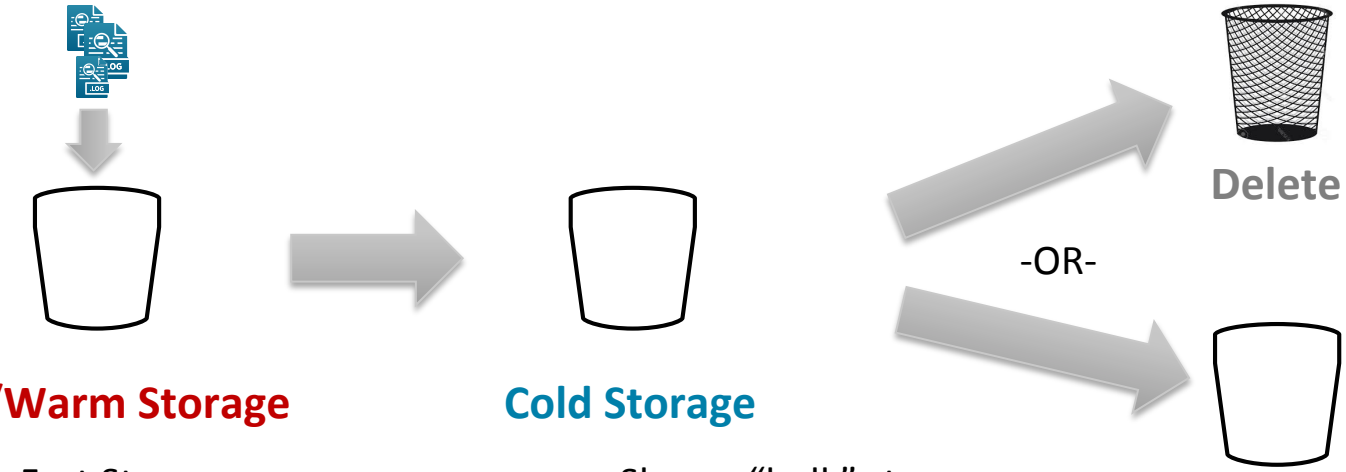

Hot/Warm Storage

- Fast Storage
- Recent data

Cold Storage

- Slower “bulk” storage
- Older data

Archive Storage

- Historical/Compliance data
- Online (searchable)/Offline

What's In A Bucket?

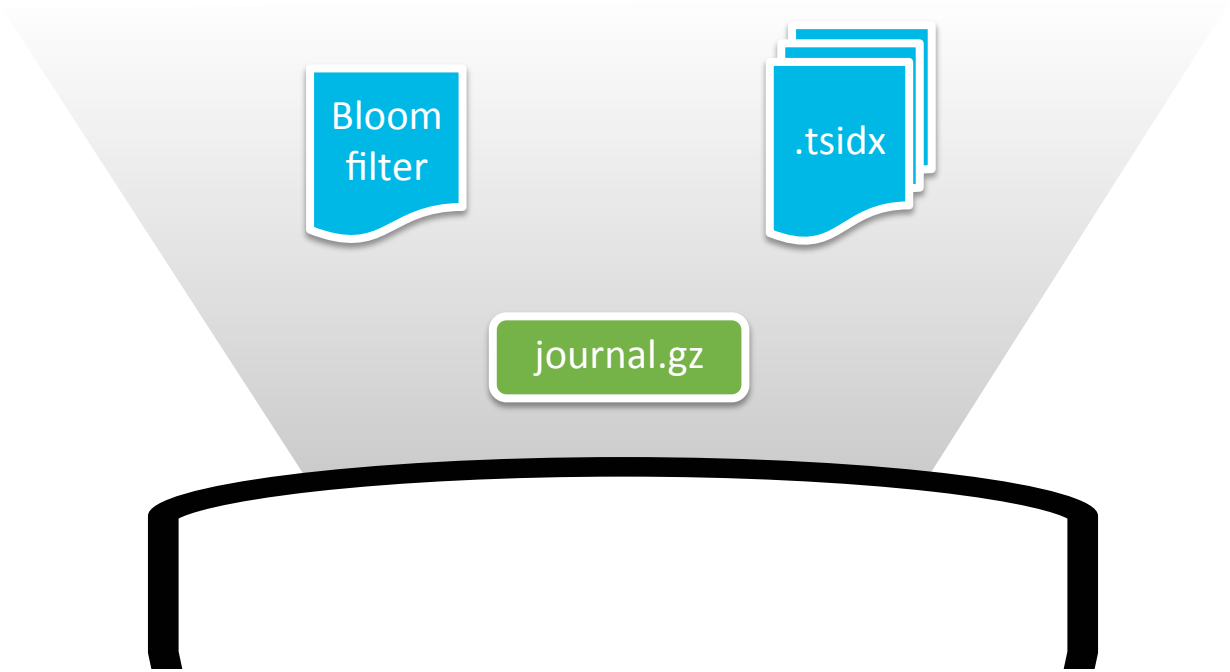

What's In A Bucket?

Journal.gz

- Your events go here
- Journal.gz is made up of many smaller compressed slices
- Raw data is collected and saved into slices
  - ~128KB of uncompressed data make up a slice

What's In A Bucket?

TSIDX

Raw Events
Jim likes Mickey
Suzie likes Donald
Pat likes Pluto

Unique terms from the raw events are written to the lexicon

Lexicon

Term	Postings List
Donald	1
Jim	0
likes	0,1,2
Mickey	0
Pat	2
Pluto	2
Suzie	1

The postings list tells us where we can find a specific term in the values array

Values Array

Posting Value	Seek Address
0	34
1	87
2	132

The seek address tells us where we can find the matching event(s) in the journal.gz slices

*The overall structure of a TSIDX file has been simplified for illustrative purposes

What's In A Bucket?

Bloom Filter

- Determines whether a term is likely to exist in the TSIDX of a bucket
  - False positives are possible, false negatives are not

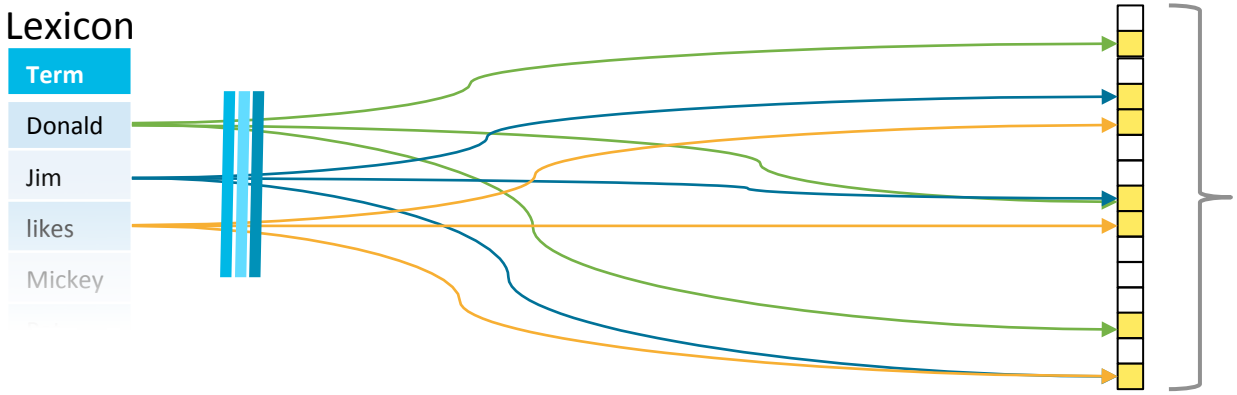

Each term from the lex is run through a set of hashing algorithms

The output of each hash sets a bit in the array to ON

- Regardless of the # of terms, bit array size remains fixed
- Binary format
- Fast to read vs. TSIDX, which grows with more unique terms

How Search Works…
An Example

How Search Works

Components of a Search String

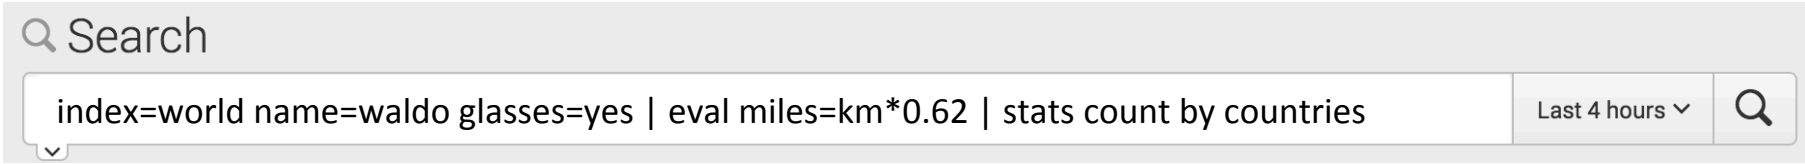

Base Search

Retrieves & filters events

SPL Commands

Evaluate, transform, and format events

Events are retrieved → Results move linearly through SPL commands

How Search Works

Where's Waldo?

index=world name=waldo 

How Search Works

Where's Waldo?

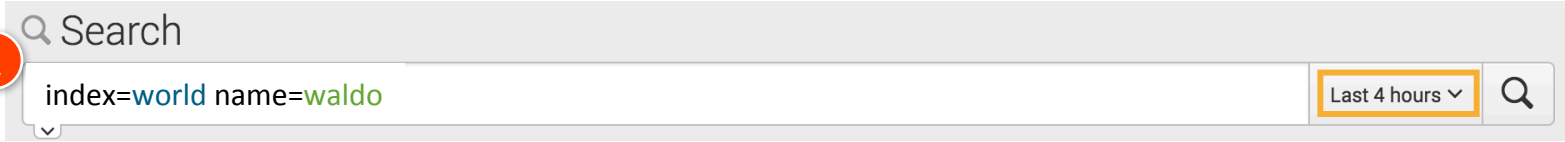

1

2 Hash the value waldo to create a bloom filter for our search

01010101001001

3 Begin searching world buckets containing events from the Last 4 hours

4 Compare our filter to the one in each bucket

Bloom filter

5 Locate the value waldo in the TSIDX

.tsidx

6 Retrieve events with waldo using the seek address in the TSIDX

journal.gz

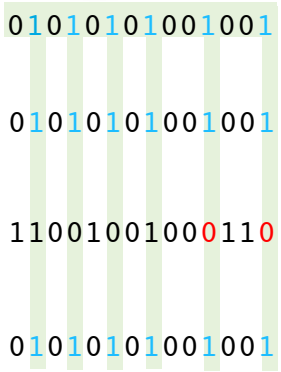

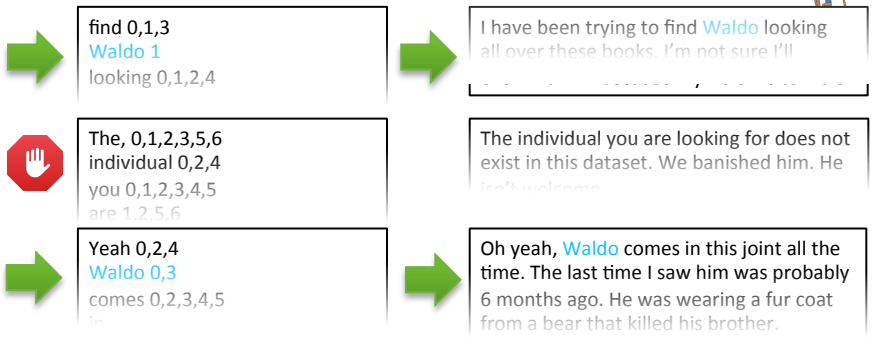

*The internal structure of Bloom filters, TSIDX, and Journal files has been simplified for illustrative purposes

How Search Works…
Distributed Search

How Search Works

Distributed Search

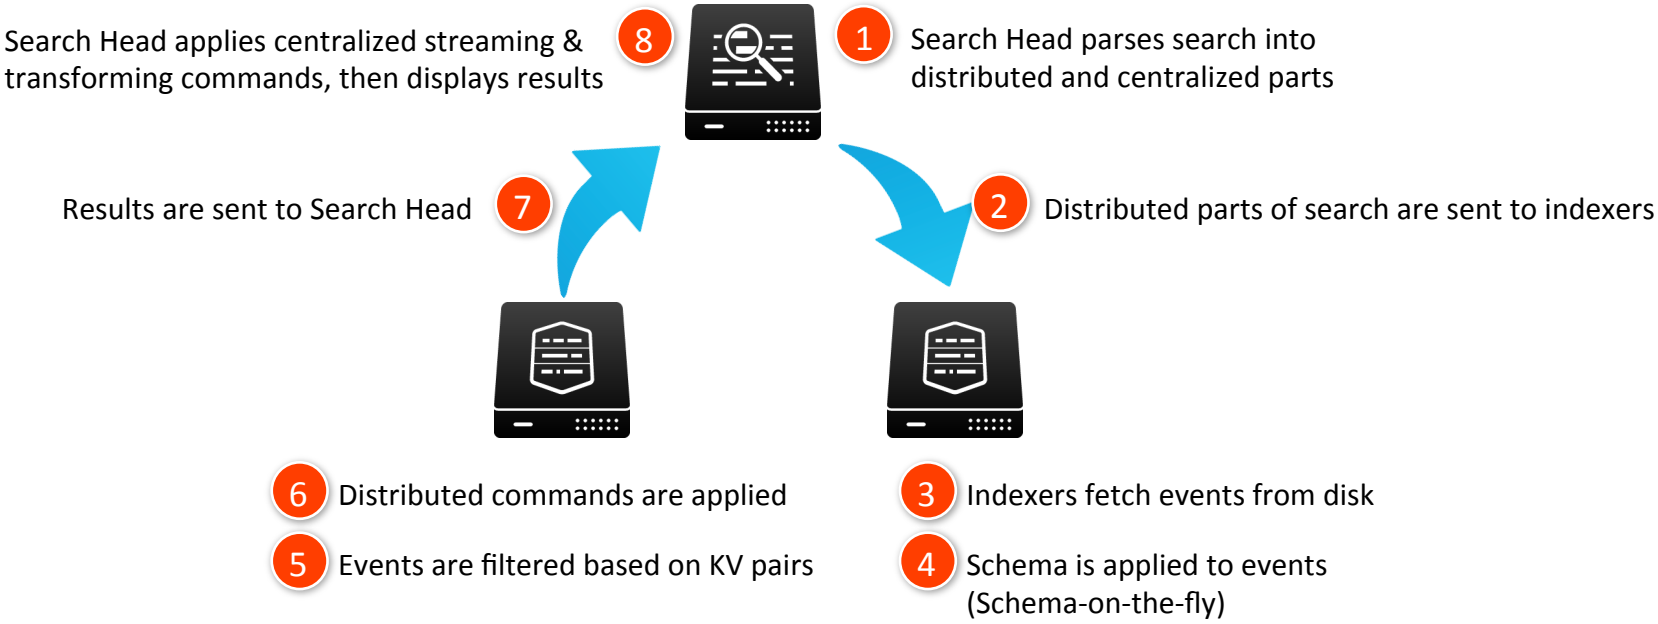

How Search Works

Types of Search Commands

- **Streaming Commands**
  - Distributable (Remote Streaming)
    - Operate on individual events
    - Run on indexers (distributed)
    - Ex: eval, rex, where, rename, fields…
  - Centralized (Stateful Streaming)
    - Operate on at least a sub-set of the entire result set
    - Run on Search Head (centralized)
    - Ex: head, streamstats
- **Transforming Commands**
  - Create a reporting data structure
  - Operate on the entire event set
    - Non-streaming
    - Typically run on the search head
  - Ex: transaction, stats, top, timechart…

How Search Works

Command Ordering

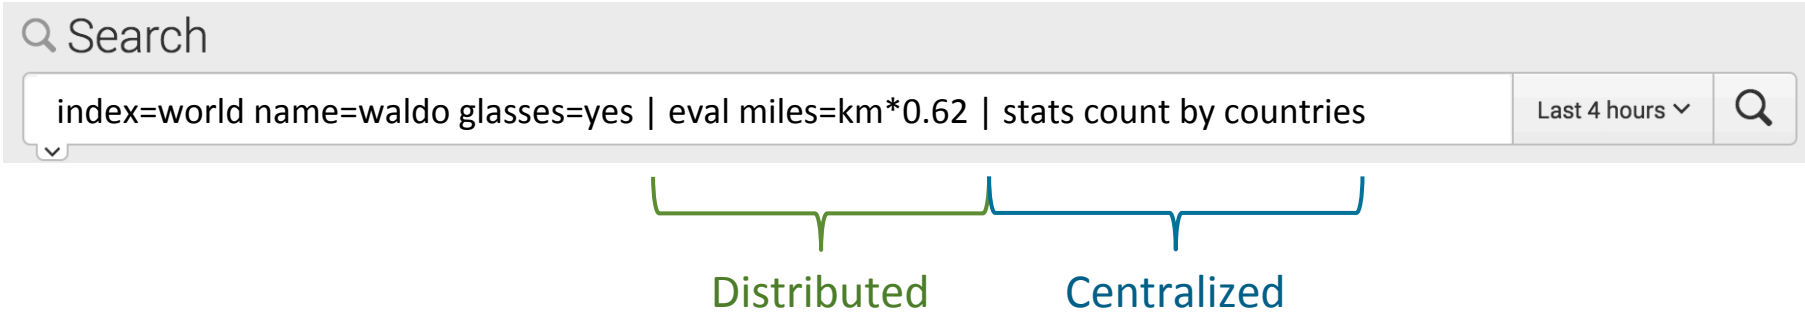

- Commands are processed in the order you write them
- Placing centralized or transforming commands before distributable commands may force unnecessary data and/or processing to the Search Head

Want To Know More?

Search: Under the Hood by Chris Pride

– Wednesday, Sept. 28th 4:35PM – 5:20PM

Search Tips

Just Because You Can…doesn't Mean You Should

Plan your search to leverage the power of Splunk!

Search Tips

- Reduce the amount of data Splunk has to Search
  - Specify and limit the index(es)
  - Limit the time range
  - Search for values that are unique to your events where possible
    - Reduce the number of events filtered after schema-on-the-fly
- Distributed Search
  - Ensure events are well distributed
  - Place distributed commands before centralized commands

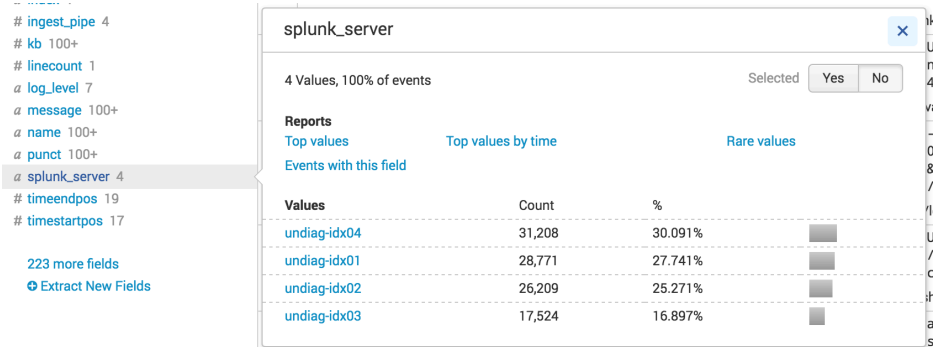

Thou shalt not use index= or All Time*

-Moses

Search Tips

Avoid	Explanation	Suggested Alternative
All Time	• Events are grouped by time • Reduce searched buckets by being specific about time	• Use a specific time range • Narrow the time range as much as possible
index=*	• Events are grouped into indexes • Reduce searched buckets by specifying an index	• Always specify an index in your search
Wildcards	• Wildcards are not compatible with Bloom Filters • Wildcard matching of terms in the index takes time	• Varying levels of suck-itude > myterm* → Not great > *myterm → Bad > *myterm* → Death • Use the OR operator i.e.: MyTerm1 OR MyTerm2

Search Tips

Avoid	Explanation	Suggested Alternative
NOT !=	• Bloom filters & indexes are designed to quickly locate terms that exist • Searching for terms that don't exist takes longer	• Use the OR/AND operators (host=c OR host=d) (host=f AND host=h) vs. (host!=a host!=b) NOT host=a host=b
Verbose Search Mode	• Verbose search mode causes full event data to be sent to the search head, even if it isn't needed	• Use Smart Mode or Fast Mode
Real-time Searches	• RT Searches put an increased load on search head and indexers • The same effect can typically be accomplished with a 1 min. or 5 min. scheduled search	• Use a scheduled search that occurs more frequently • Use Indexed-Realtime searches (Set by Splunk admin)

Search Tips

Avoid	Explanation	Suggested Alternative
Transaction	• Not distributed to indexers • Typically only needed if using additional parameters (maxSpan, startsWith, etc…)	• Use the stats command to link events where possible
Joins/Sub-searches	• Joins can be used to link events by a common field value, but this is an intensive search command	• Use the stats (preferred) or transaction command to link events
Search after first |	• Filtering search results using a second "| search" command in your query is inefficient	• As much as possible, add all filtering criteria before the first | i.e.: >index=main foo bar vs. >index=main foo | search bar

The TERM Directive

Why does it matter?

- Splunk breaks terms by Major and Minor Segmenters
  - When writing to the TSIDX and searching
  - Default minor segmenters:
    `/ : = @ . - $ # % \\ _`
- TERM prevents breaking on Minor segmenters

`[ AND 0 10 6 index::myindex ]`

Raw Events
10.0.0.6
9/28/2016
jeff@splunk.com

Lexicon

Term	Postings List
0	0
6	0
9	1
10	0
28	1
2016	1
10.0.0.6	0
9/28/2016	1
com	2
jeff	2
splunk	2
jeff@splunk.com	2

The TERM Directive

What about quotes?

- TERM controls how we search the lexicon and which events are retrieved from disk
- Quotes can help filter after the events are retrieved from disk
- Use quotes when the *value* in your key-value pair has major breakers

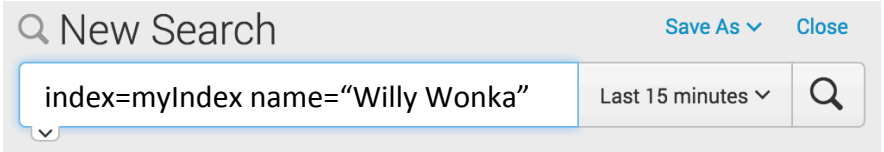


```
[ AND wonka willy index::myindex ]
```

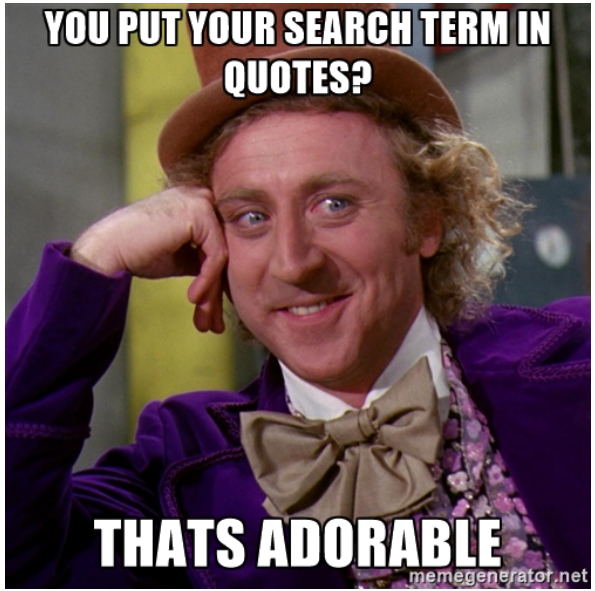

The TERM Directive

How do I use it?

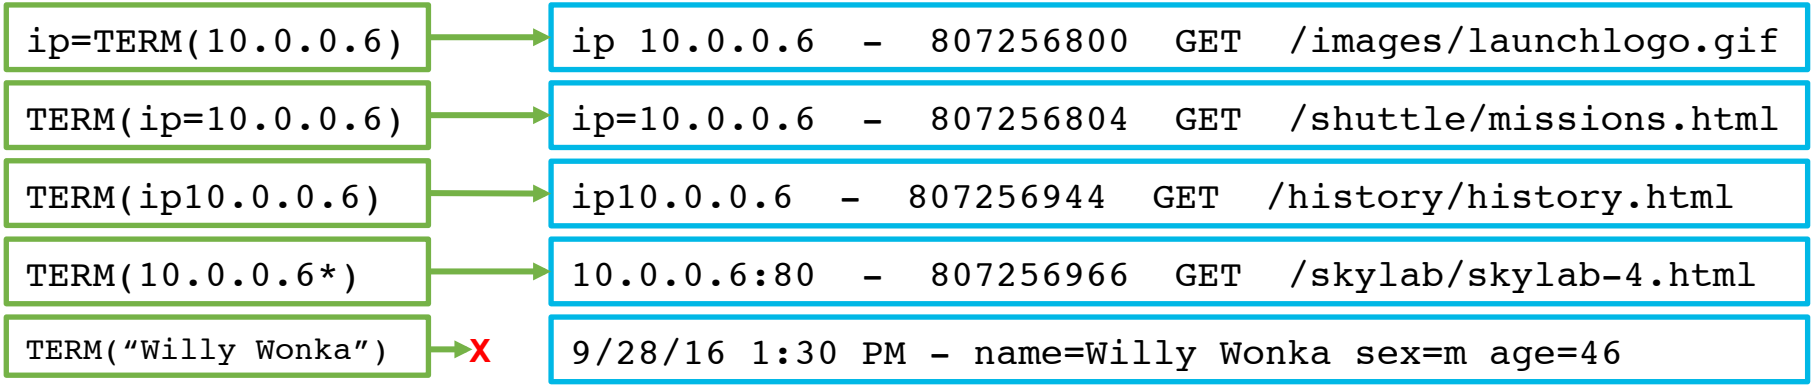

- Your term MUST be bounded by major segmenters
  - Example: Spaces, tabs, carriage returns
    - See Segmenters.conf spec for full details
  - Your term cannot contain major segmenters

Search Tips

Indexed Extractions

- Special Key-Value pairs that are stored in the TSIDX file
- Default Extractions
  - source, host, sourcetype
  - Use these whenever possible
- TSTATS
  - Super-fast command
  - Doesn't search or return raw data
  - Can be used on report/data model accelerations AND indexed extractions

Want To Know More?

How to Scale: From _raw to tstats by David Veuve

- Wednesday, Sept 28th 2:15 PM – 3:00 PM

Previous Session:

- **Fields, Indexed Tokens and You** by Martin Müller

Commands In Action

Command Abuse

Fields vs. Table

Goal: Remove fields I don't need from results

BAD:

```
index=myIndex field1=value1 | table field1, field2, field4 | head 10000
| table field2, field4
```

GOOD:

```
index=myIndex field1=value1 | fields field1, field2, field4 | head 10000
| table field2, field4
```

- Table is a formatting command NOT a filtering command
  - If used improperly, it will cause unnecessary data to be transferred to the search head from search peers
- Fields tells Splunk to explicitly drop or retain fields from your results

Command Abuse

Fields vs. Table Example

Search Term	Status	Artifact Size	# of Events	Run Time
| table	Running (1%)	624.93MB	2,037,500	00:02:44
| fields	Done	9.95MB	10,000	00:00:13

Command Abuse

Stats vs. Transaction

Goal: Group multiple events by a common field value

NOT GREAT:

```
index=mail from=joe@schmoe.com| transaction message_id | table _time, to,
from, subject, message_id
```

GOOD:

```
index=mail from=joe@schmoe.com | stats latest(_time) AS mTime values(to)
AS to values(from) AS from values(subject) AS subject BY message_id
```

- If you're not using any of the Transaction command parameters, the same results can usually be accomplished using Stats
  - startswith, endswith, maxspan, maxpause, etc…

Command Abuse

Joins & Sub-searches

Goal: Return the latest JSESSIONID across two sourcetypes

NOT GREAT:

```
sourcetype=access_combined | join type=inner JSESSIONID [search
      sourcetype=applogs | dedup JSESSIONID  | table JSESSIONID,
      clienip, othervalue]
```

GOOD:

```
sourcetype=access_combined OR sourcetype=applogs | stats latest(*) AS *
BY JSESSIONID
```

Resources

- Splunk Docs
  - Write Better Searches
    http://docs.splunk.com/Documentation/Splunk/latest/Search/Writebettersearches
  - Wiki: How Distributed Search Works
    http://wiki.splunk.com/Community:HowDistSearchWorks
  - Splunk Search Types
    http://docs.splunk.com/Documentation/Splunk/6.2.3/Capacity/HowsearchtypesaffectSplunkEnterpriseperformance
  - Blog: When to use Transaction and when to use Stats
    http://blogs.splunk.com/2012/11/29/book-excerpt-when-to-use-transaction-and-when-to-use-stats/
  - Segmenters.conf Spec
    http://docs.splunk.com/Documentation/Splunk/latest/Admin/Segmentersconf
  - Splunk Book: Exploring Splunk
    http://www.splunk.com/goto/book
- How Bloom Filters Work: An Interactive Demo
  https://www.jasondavies.com/bloomfilter/

Resources

Training

- **eLearning**
  - What is Splunk (Intro to Splunk)
    - http://www.splunk.com/view/SP-CAAAH9U
- **Instructor Led Courses with Labs**
  - Using Splunk
    - http://www.splunk.com/view/SP-CAAAH9A
  - Searching & Reporting with Splunk
    - http://www.splunk.com/view/SP-CAAAH9C
  - Advanced Searching & Reporting
    - http://www.splunk.com/view/SP-CAAAH9D

What Now?

Related breakout sessions and activities…

- **How to Scale: From _raw to tstats** by David Veuve
  - Wednesday, Sept 28th 2:15 PM – 3:00 PM
- **Search: Under the Hood** by Chris Pride
  - Wednesday, Sept. 28th 4:35PM – 5:20PM
- **Best Practices and Better Practices for Users** by Burch Simon
  - Thursday, Sept. 29th 12:25PM – 1:10PM

Previous Sessions…

- **Fields, Indexed Tokens and You** by Martin Müller
- **Worst Practices…and How to Fix Them** by Jeff Champagne
- **Best and Better Practices for Admins** by Burch Simon
- **Observations and Recommendations on Splunk Performance** by Dritan Bitincka

Questions?

THANK YOU

splunk>