Clara-fication: Job Inspector

Clara-fy your jobs

Clara Merriman
Senior Splunk Engineer | Splunk

Martin Müller
Principal Consultant | Consist Software Solutions

Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or plans 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 may differ materially. 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, it may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements made herein.

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 functionalities described or to include any such feature or functionality in a future release.

Clara Merriman

Senior Splunk Engineer | Splunk

Martin Müller

Principal Consultant | Consist Software Solutions

splunk> .conf20

Agenda

| rest splunk_server=local services/search/scheduler

1) What is the job inspector?

- Available both in the UI and REST
- Tool that offers insights into various aspects of search jobs

2) Why is the job inspector important?

- Troubleshoot a broken search
- Optimize an inefficient search

3) What can the job inspector tell me?

- Search execution costs
- Search properties (and saved search job properties)
- search.log, which is not currently logged in the internal indexes
- Errors, warnings, and debug messages

/job_inspector

BaDaDaDaDa Inspector Gadget

What is the job inspector?

localhost:8000/manager/search/job_inspector?sid=<sid>

- Tool offered in Splunk web allowing users to troubleshoot their searches
- Contains a wealth of data to make searches more efficient
- Provides helpful information such as:
  - job status
  - errors
  - events per second
  - execution costs
  - search job properties
  - saved search properties, if job uses saved search
  - search logs

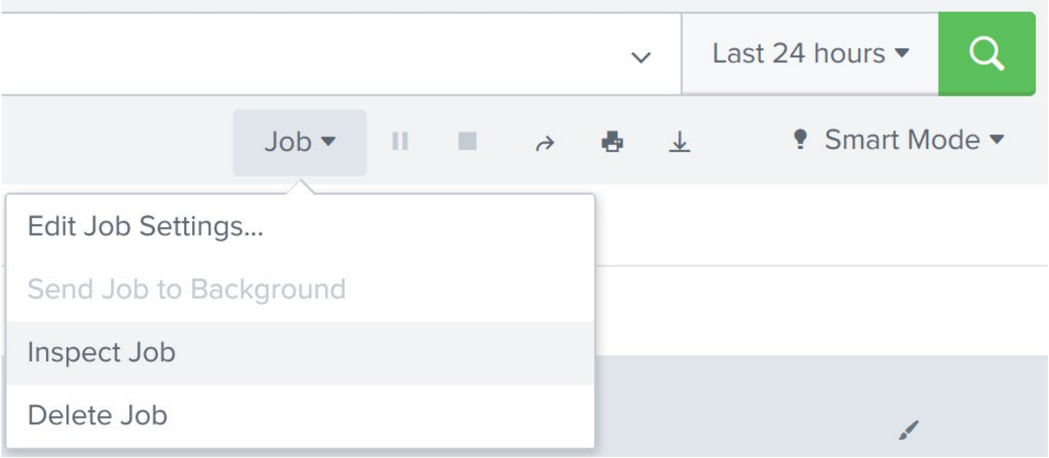

Basics

Status, errors, eps

Status

- Find the status of a job at the top of the job inspector or in job properties

Errors

- Find errors in search.log or at the top of the job inspector

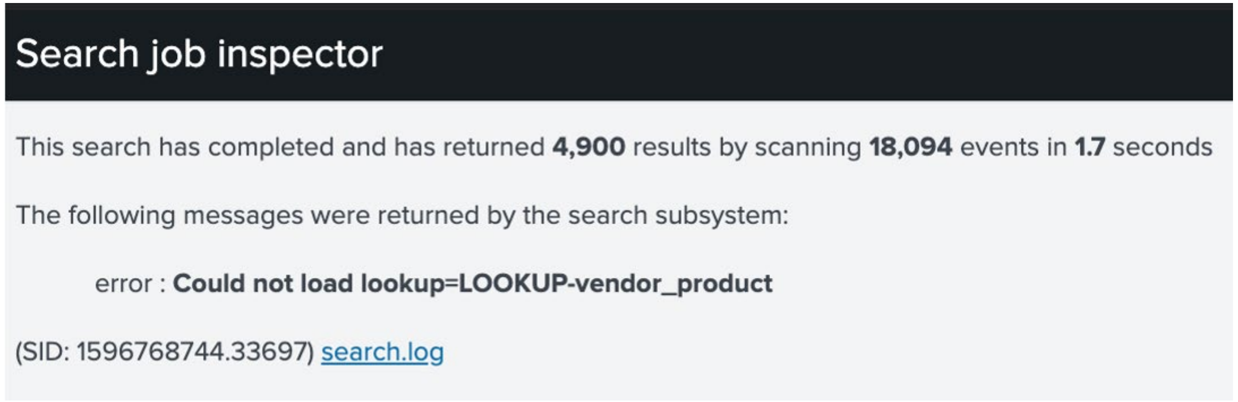

Search job inspector

This search has completed and has returned **4,900** results by scanning **18,094** events in **1.7** seconds

The following messages were returned by the search subsystem:

error : **Could not load lookup=LOOKUP-vendor_product**

(SID: 1596768744.33697) search.log

Measure Performance

- eps (events per second) = scanCount / runDuration
  eps > 10k / indexer = 👍

Example: Error Message

Lookup does not exist, but props reference does

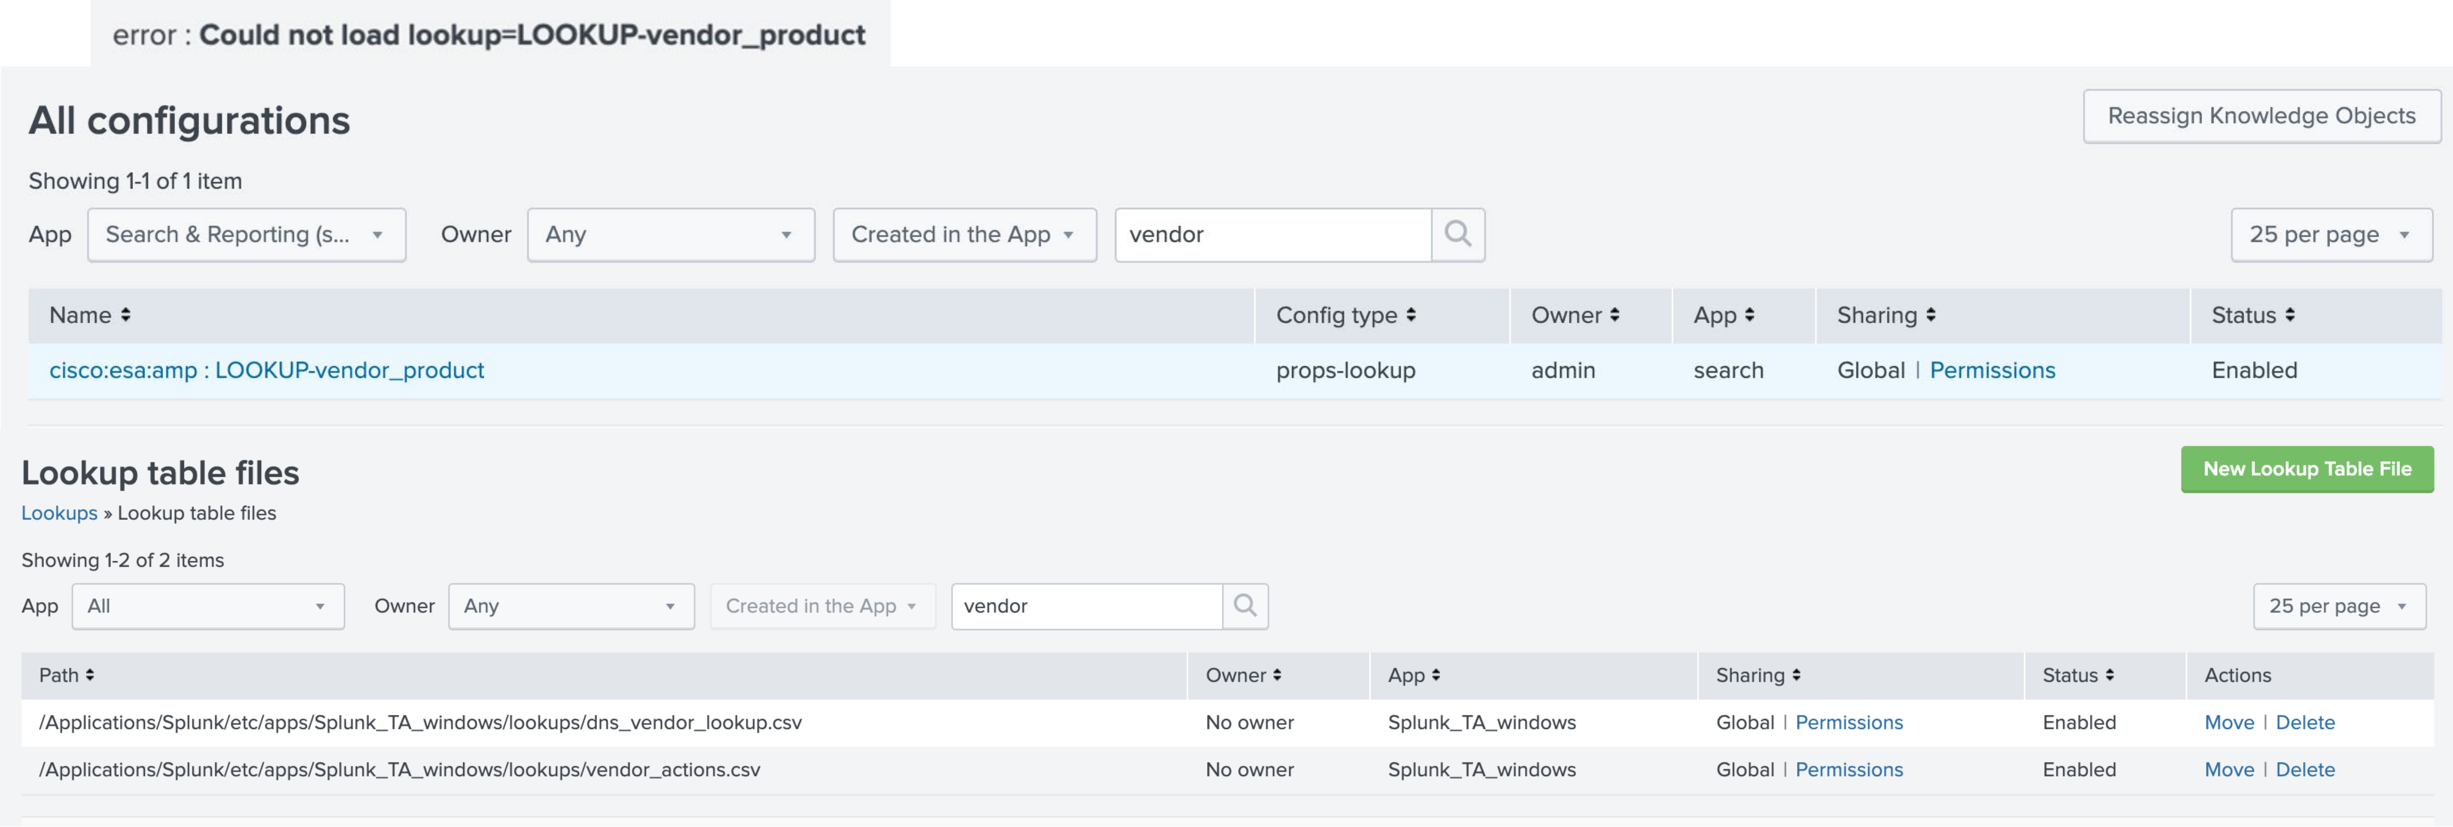

error : **Could not load lookup=LOOKUP-vendor_product**

All configurations

Showing 1-1 of 1 item

App Search & Reporting (s... Owner Any Created in the App vendor 25 per page

Name	Config type	Owner	App	Sharing	Status
cisco:esa:amp : LOOKUP-vendor_product	props-lookup	admin	search	Global | Permissions	Enabled

Lookup table files

Lookups » Lookup table files

Showing 1-2 of 2 items

App All Owner Any Created in the App vendor 25 per page

Path	Owner	App	Sharing	Status	Actions
/Applications/Splunk/etc/apps/Splunk_TA_windows/lookups/dns_vendor_lookup.csv	No owner	Splunk_TA_windows	Global | Permissions	Enabled	Move | Delete
/Applications/Splunk/etc/apps/Splunk_TA_windows/lookups/vendor_actions.csv	No owner	Splunk_TA_windows	Global | Permissions	Enabled	Move | Delete

*This could also be a problem if the lookup existed but was private or not shared globally

Execution Costs

KaPow!

Execution Cost Properties

| rest services/search/jobs/<sid>

1) Duration in seconds

Cumulative run time over all phases (SH + IDXs)

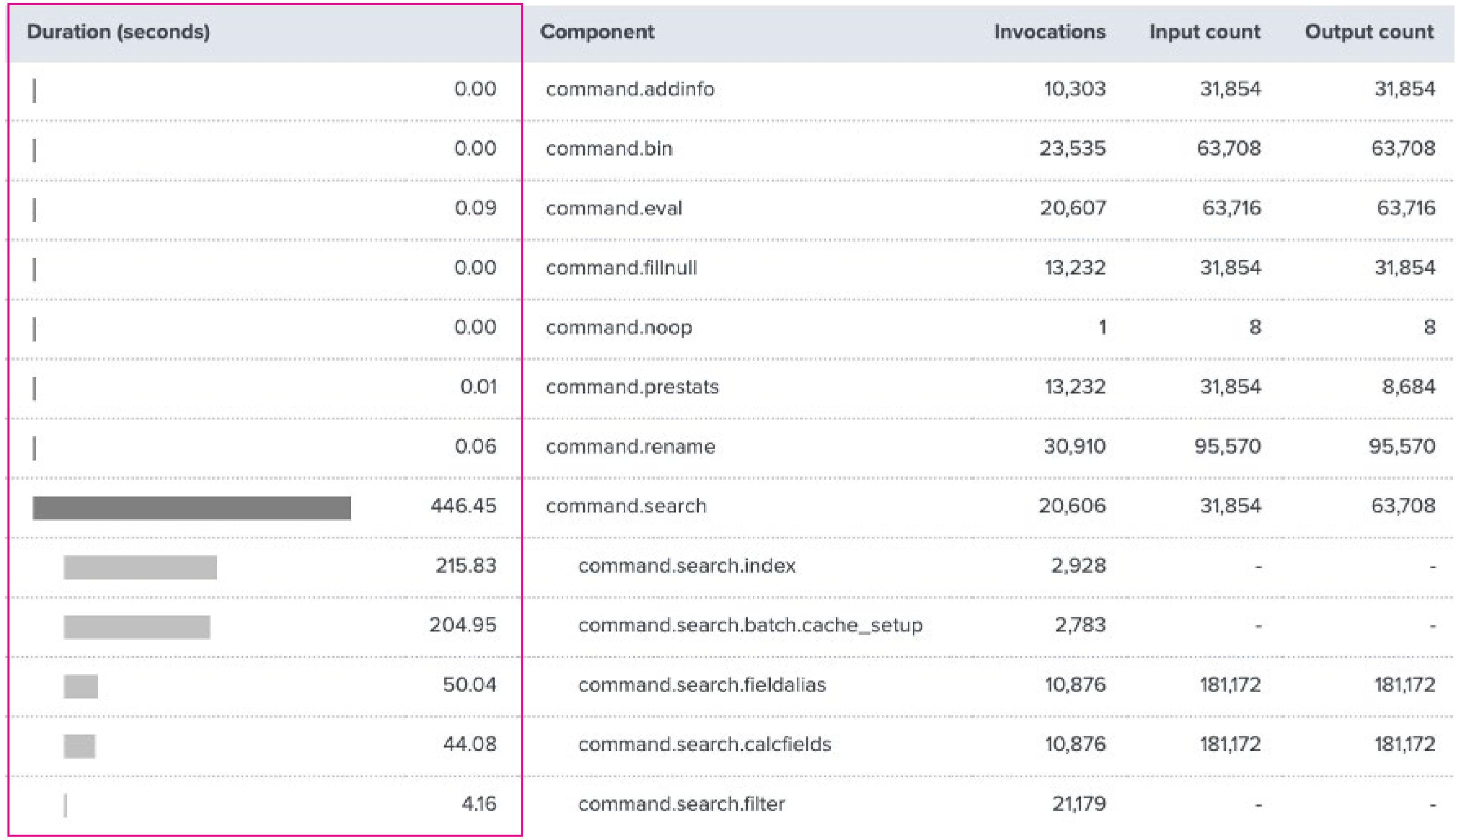

Duration (seconds)	Component	Invocations	Input count	Output count
0.00	command.addinfo	10,303	31,854	31,854
0.00	command.bin	23,535	63,708	63,708
0.09	command.eval	20,607	63,716	63,716
0.00	command.fillnull	13,232	31,854	31,854
0.00	command.noop	1	8	8
0.01	command.prestats	13,232	31,854	8,684
0.06	command.rename	30,910	95,570	95,570
446.45	command.search	20,606	31,854	63,708
215.83	command.search.index	2,928	-	-
204.95	command.search.batch.cache_setup	2,783	-	-
50.04	command.search.fieldalias	10,876	181,172	181,172
44.08	command.search.calcfields	10,876	181,172	181,172
4.16	command.search.filter	21,179	-	-

Execution Cost Properties

| rest services/search/jobs/<sid>

1) **Duration in seconds**

   Cumulative run time over all phases (SH + IDXs)

2) **Component**

   Hierarchical structure of commands and their subcomponents

   Note: Durations of indented subcomponents are included in the parent component's duration

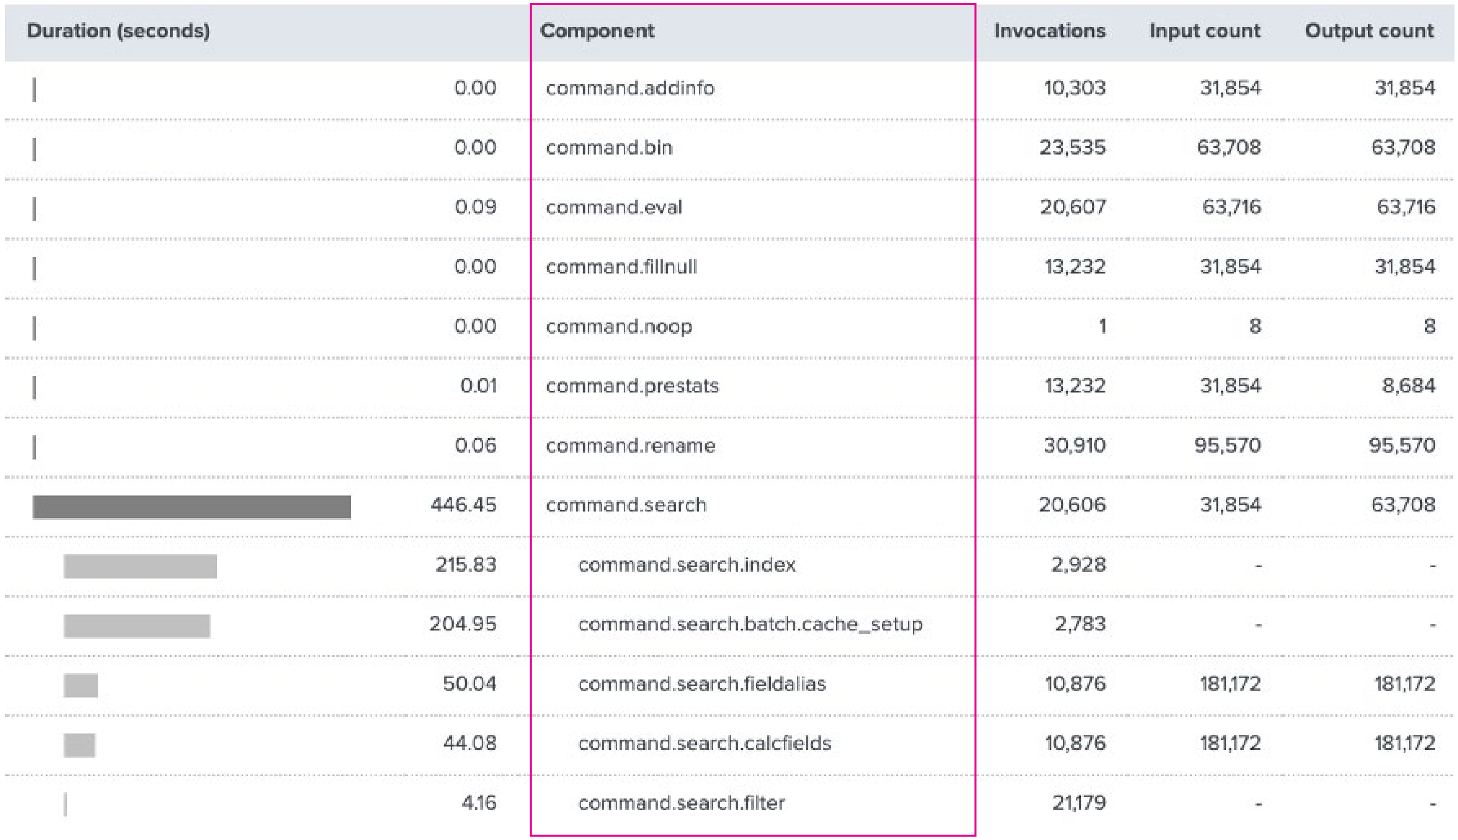

Duration (seconds)	Component	Invocations	Input count	Output count
0.00	command.addinfo	10,303	31,854	31,854
0.00	command.bin	23,535	63,708	63,708
0.09	command.eval	20,607	63,716	63,716
0.00	command.fillnull	13,232	31,854	31,854
0.00	command.noop	1	8	8
0.01	command.prestats	13,232	31,854	8,684
0.06	command.rename	30,910	95,570	95,570
446.45	command.search	20,606	31,854	63,708
215.83	command.search.index	2,928	-	-
204.95	command.search.batch.cache_setup	2,783	-	-
50.04	command.search.fieldalias	10,876	181,172	181,172
44.08	command.search.calcfields	10,876	181,172	181,172
4.16	command.search.filter	21,179	-	-

Execution Cost Properties

| rest services/search/jobs/<sid>

1) **Duration in seconds**

   Cumulative run time over all phases (SH + IDXs)

2) **Component**

   Hierarchical structure of commands and their subcomponents

   Note: Durations of indented subcomponents are included in the parent component's duration

3) **Input/Output counts**

   Number of events going in and out of that component

   Exception: For dispatch.stream.remote this is bytes sent from the IDXs to the SH

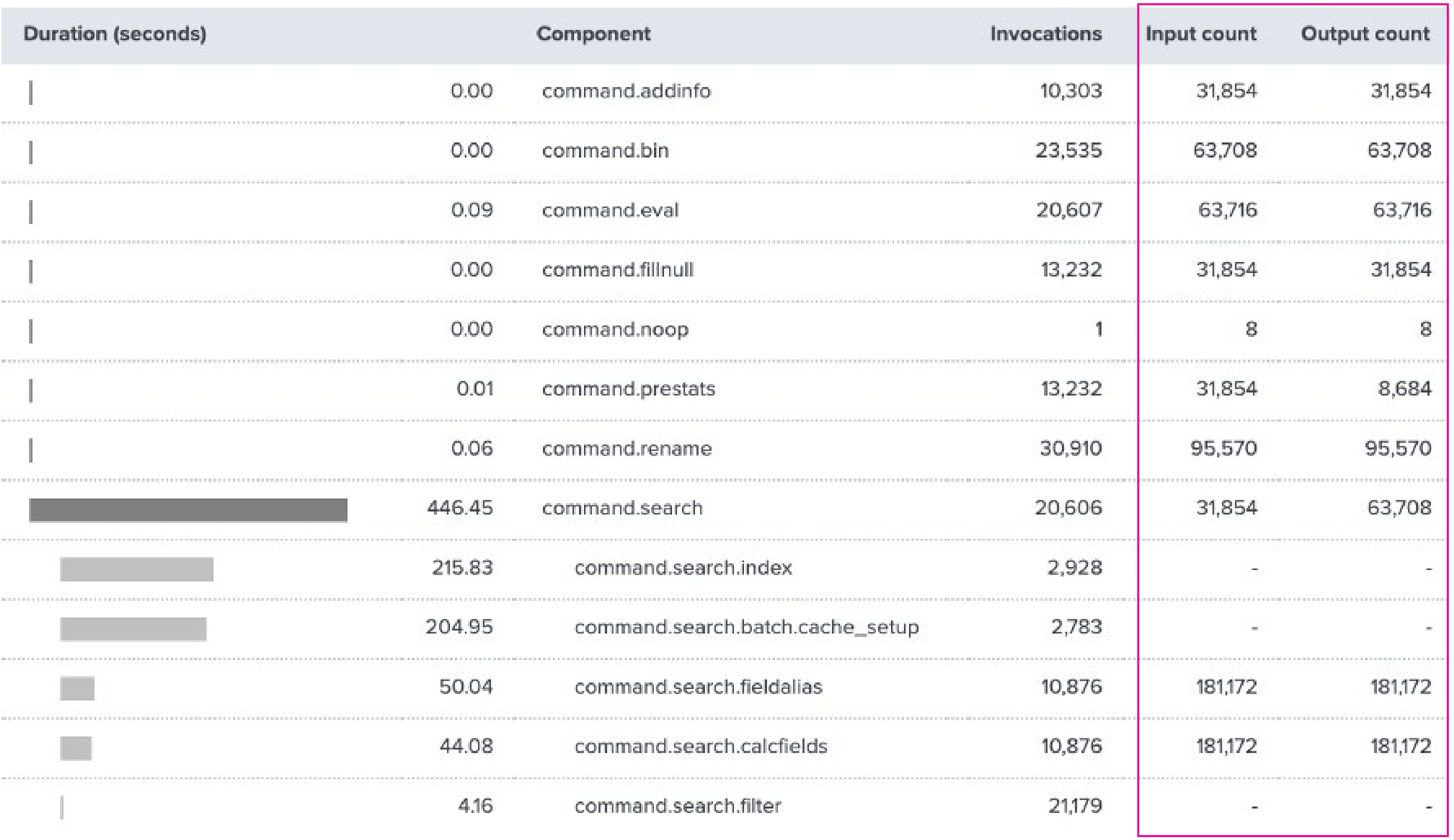

Duration (seconds)	Component	Invocations	Input count	Output count
0.00	command.addinfo	10,303	31,854	31,854
0.00	command.bin	23,535	63,708	63,708
0.09	command.eval	20,607	63,716	63,716
0.00	command.fillnull	13,232	31,854	31,854
0.00	command.noop	1	8	8
0.01	command.prestats	13,232	31,854	8,684
0.06	command.rename	30,910	95,570	95,570
446.45	command.search	20,606	31,854	63,708
215.83	command.search.index	2,928	-	-
204.95	command.search.batch.cache_setup	2,783	-	-
50.04	command.search.fieldalias	10,876	181,172	181,172
44.08	command.search.calcfields	10,876	181,172	181,172
4.16	command.search.filter	21,179	-	-

splunk> .conf20

Execution Costs

| rest services/search/jobs/<sid>

1) **command.<command>**

   Running <command> throughout search

   Focus optimization on large durations first

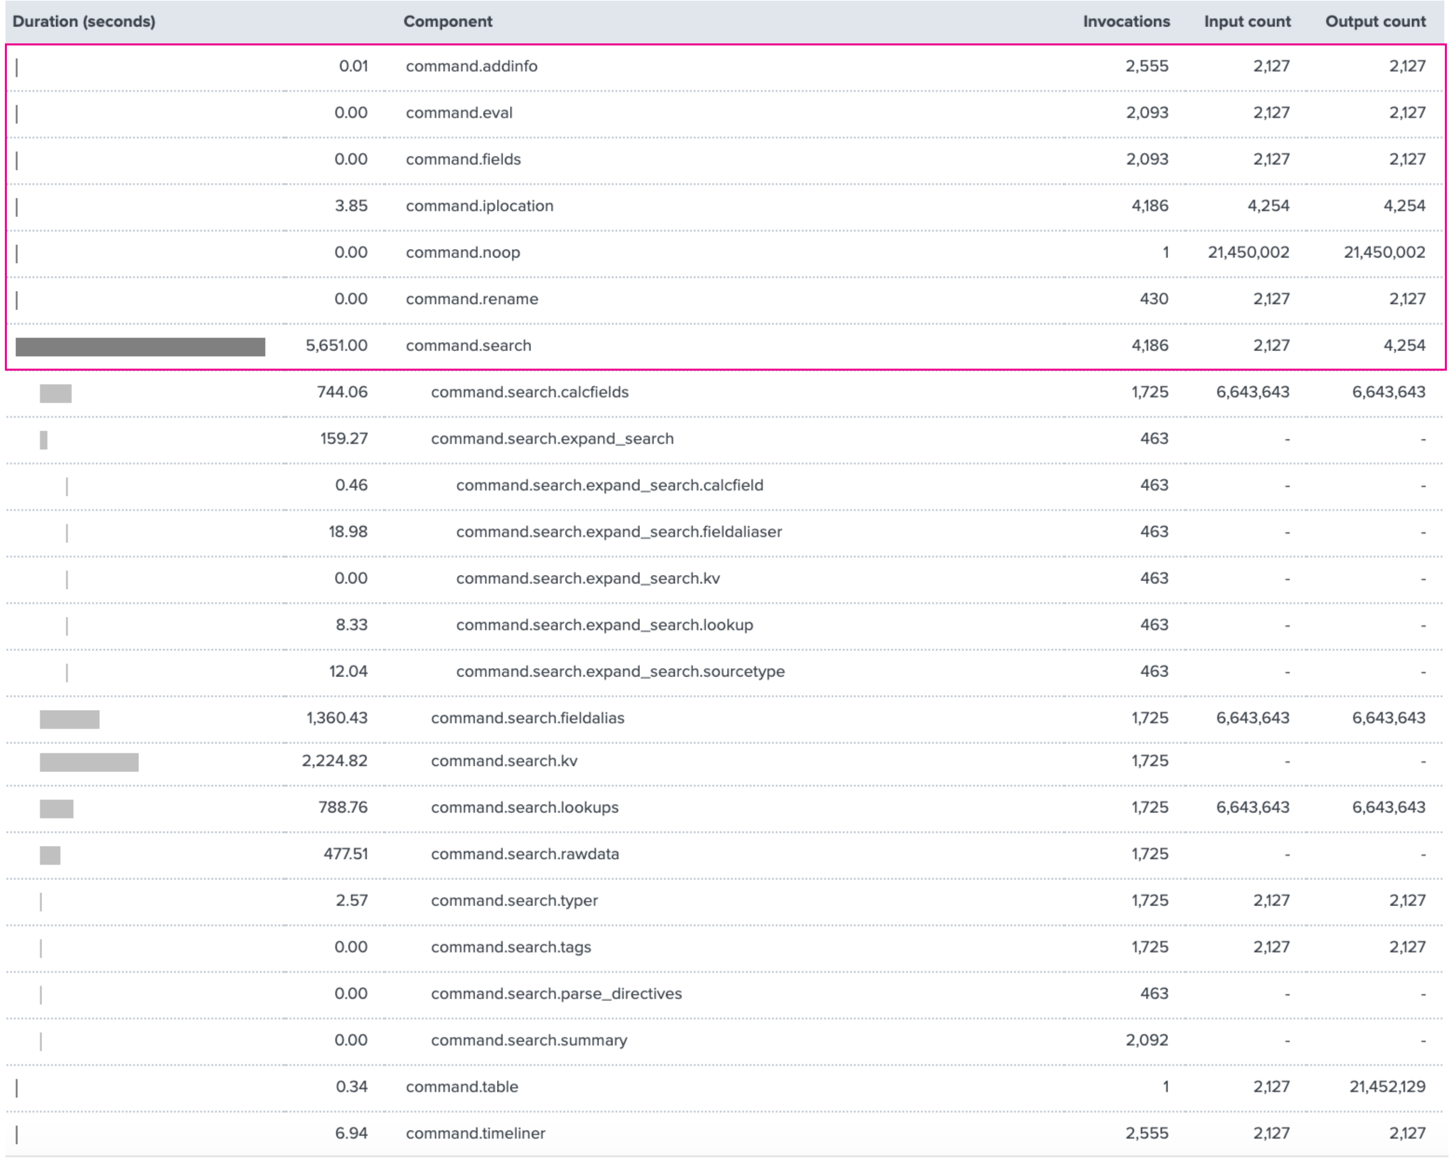

Duration (seconds)	Component	Invocations	Input count	Output count
0.01	command.addinfo	2,555	2,127	2,127
0.00	command.eval	2,093	2,127	2,127
0.00	command.fields	2,093	2,127	2,127
3.85	command.iplocation	4,186	4,254	4,254
0.00	command.noop	1	21,450,002	21,450,002
0.00	command.rename	430	2,127	2,127
5,651.00	command.search	4,186	2,127	4,254
744.06	command.search.calcfields	1,725	6,643,643	6,643,643
159.27	command.search.expand_search	463	-	-
0.46	command.search.expand_search.calcfield	463	-	-
18.98	command.search.expand_search.fieldaliaser	463	-	-
0.00	command.search.expand_search.kv	463	-	-
8.33	command.search.expand_search.lookup	463	-	-
12.04	command.search.expand_search.sourcetype	463	-	-
1,360.43	command.search.fieldalias	1,725	6,643,643	6,643,643
2,224.82	command.search.kv	1,725	-	-
788.76	command.search.lookups	1,725	6,643,643	6,643,643
477.51	command.search.rawdata	1,725	-	-
2.57	command.search.typer	1,725	2,127	2,127
0.00	command.search.tags	1,725	2,127	2,127
0.00	command.search.parse_directives	463	-	-
0.00	command.search.summary	2,092	-	-
0.34	command.table	1	2,127	21,452,129
6.94	command.timeliner	2,555	2,127	2,127

Execution Costs

| rest services/search/jobs/<sid>

1) **command.<command>**

   Running <command> throughout search

   Focus optimization on large durations first

2) **command.search.fieldalias**

   Renaming fields defined in props.conf

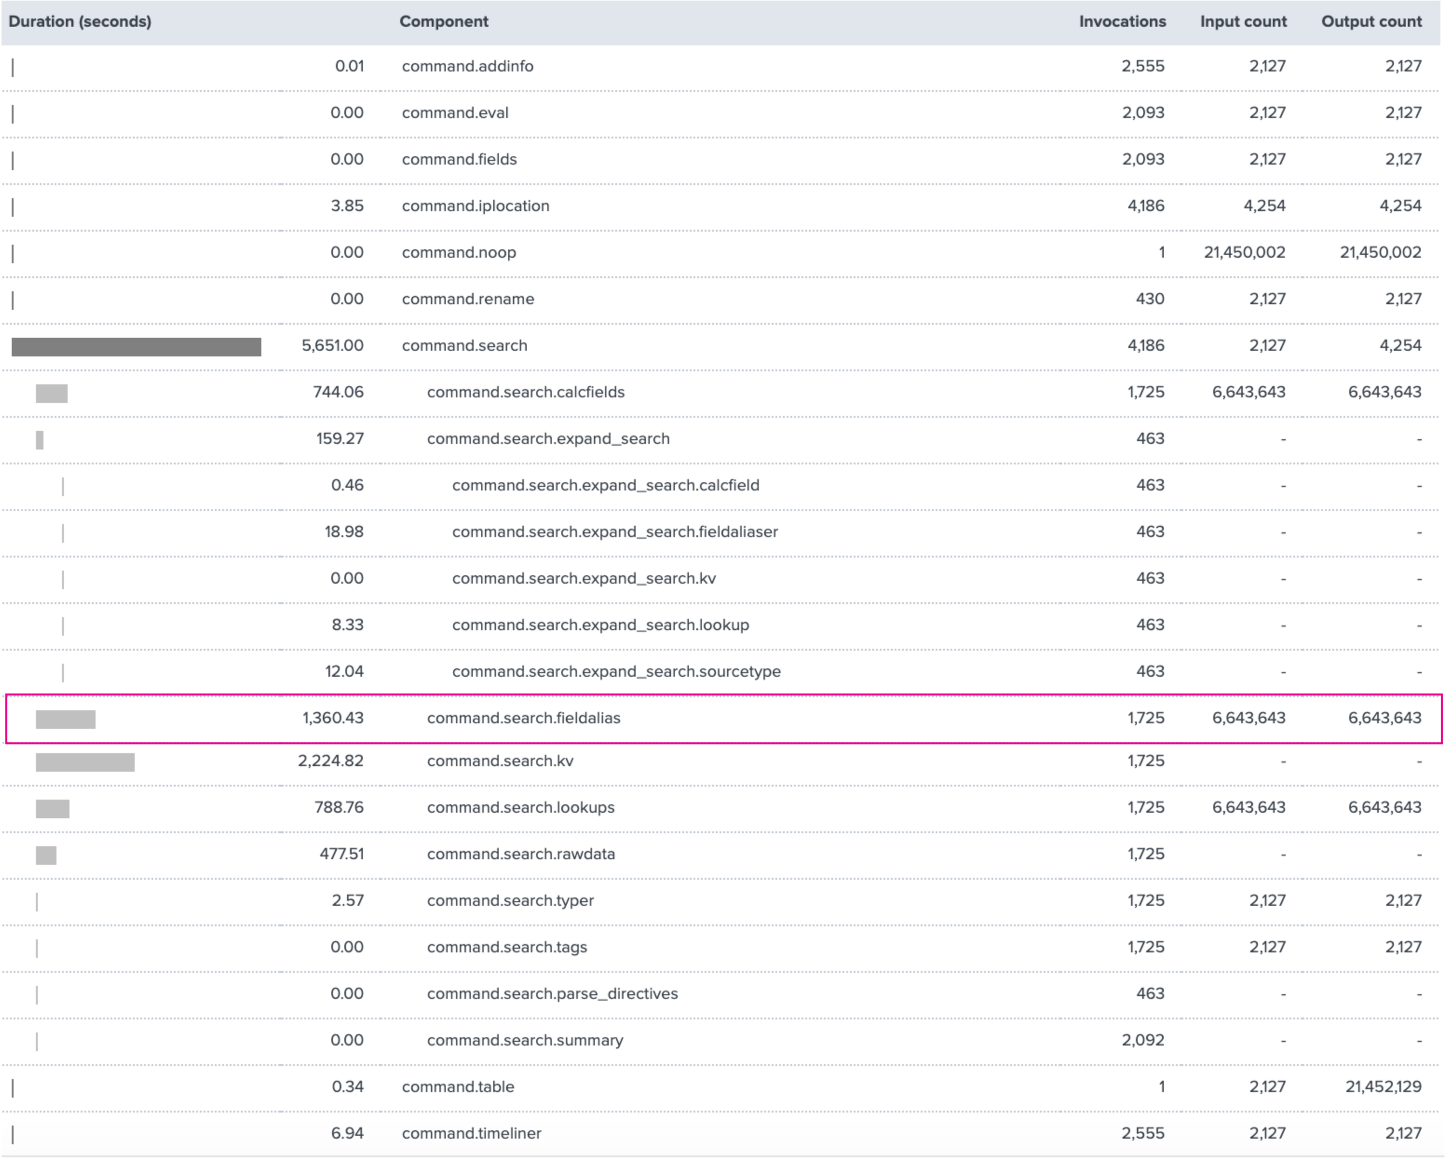

Duration (seconds)	Component	Invocations	Input count	Output count
0.01	command.addinfo	2,555	2,127	2,127
0.00	command.eval	2,093	2,127	2,127
0.00	command.fields	2,093	2,127	2,127
3.85	command.iplocation	4,186	4,254	4,254
0.00	command.noop	1	21,450,002	21,450,002
0.00	command.rename	430	2,127	2,127
5,651.00	command.search	4,186	2,127	4,254
744.06	command.search.calcfields	1,725	6,643,643	6,643,643
159.27	command.search.expand_search	463	-	-
0.46	command.search.expand_search.calcfield	463	-	-
18.98	command.search.expand_search.fieldaliaser	463	-	-
0.00	command.search.expand_search.kv	463	-	-
8.33	command.search.expand_search.lookup	463	-	-
12.04	command.search.expand_search.sourcetype	463	-	-
1,360.43	command.search.fieldalias	1,725	6,643,643	6,643,643
2,224.82	command.search.kv	1,725	-	-
788.76	command.search.lookups	1,725	6,643,643	6,643,643
477.51	command.search.rawdata	1,725	-	-
2.57	command.search.typer	1,725	2,127	2,127
0.00	command.search.tags	1,725	2,127	2,127
0.00	command.search.parse_directives	463	-	-
0.00	command.search.summary	2,092	-	-
0.34	command.table	1	2,127	21,452,129
6.94	command.timeliner	2,555	2,127	2,127

Execution Costs

| rest services/search/jobs/<sid>

1) **command.<command>**
   Running <command> throughout search
   Focus optimization on large durations first
2) **command.search.fieldalias**
   Renaming fields defined in props.conf
3) **command.search.kv**
   Applying search-time field extractions

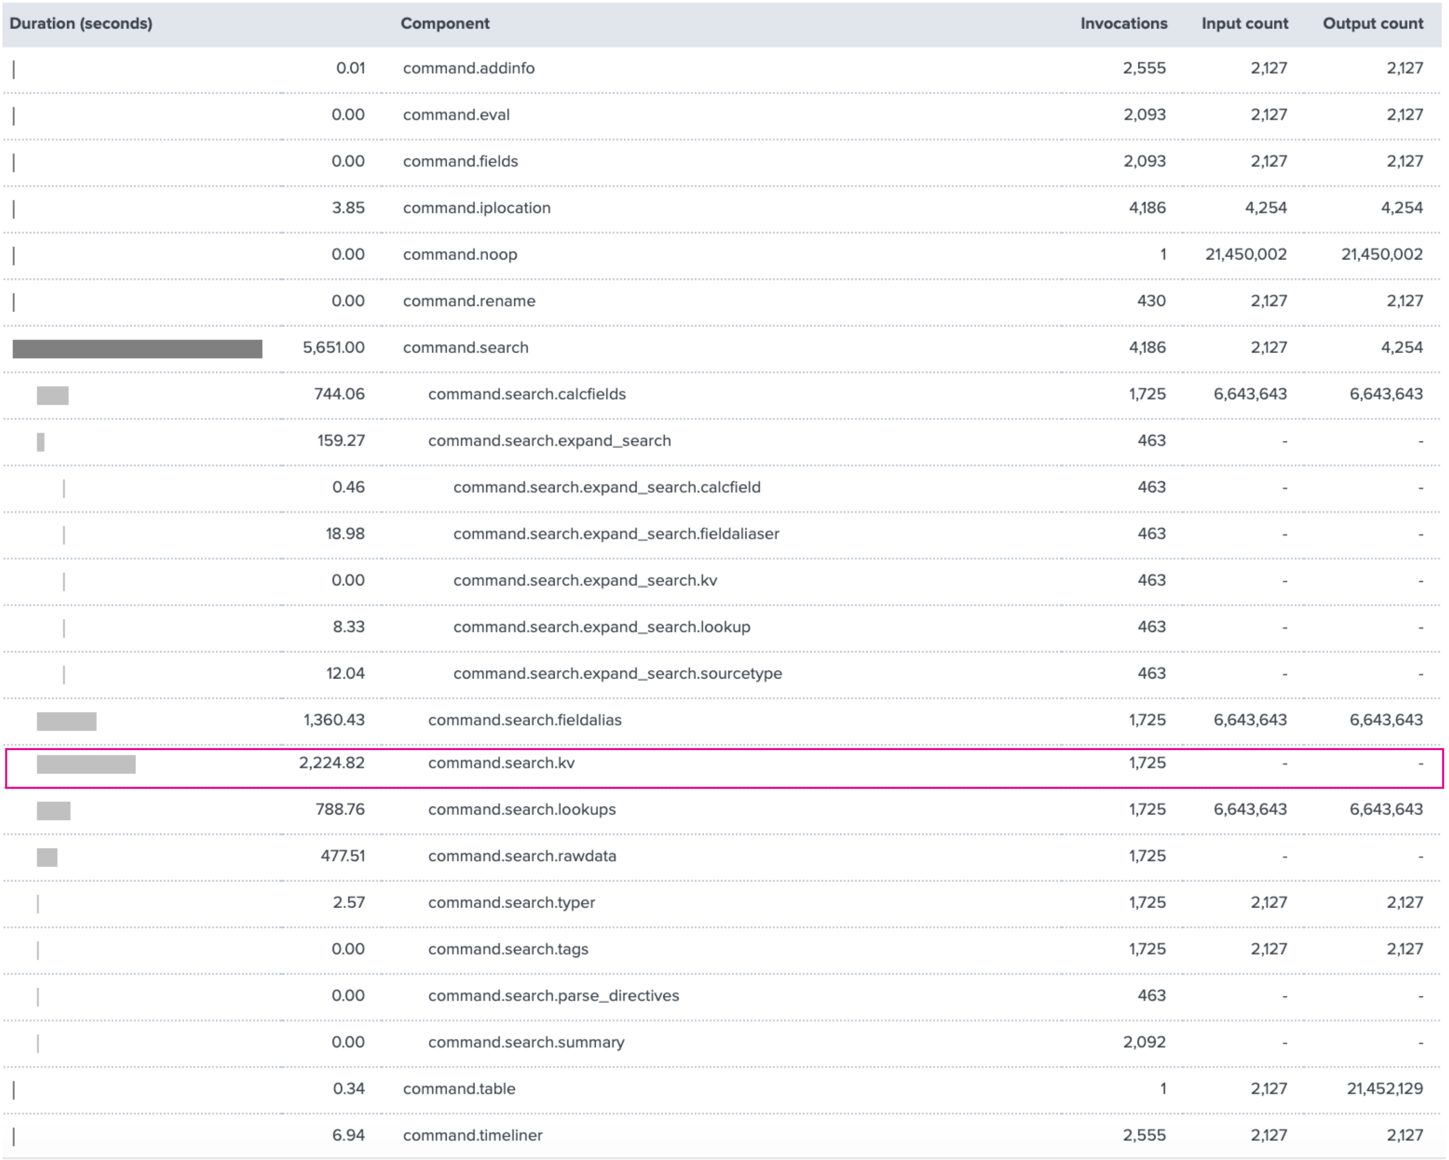

Duration (seconds)	Component	Invocations	Input count	Output count
0.01	command.addinfo	2,555	2,127	2,127
0.00	command.eval	2,093	2,127	2,127
0.00	command.fields	2,093	2,127	2,127
3.85	command.iplocation	4,186	4,254	4,254
0.00	command.noop	1	21,450,002	21,450,002
0.00	command.rename	430	2,127	2,127
5,651.00	command.search	4,186	2,127	4,254
744.06	command.search.calcfields	1,725	6,643,643	6,643,643
159.27	command.search.expand_search	463	-	-
0.46	command.search.expand_search.calcfield	463	-	-
18.98	command.search.expand_search.fieldaliaser	463	-	-
0.00	command.search.expand_search.kv	463	-	-
8.33	command.search.expand_search.lookup	463	-	-
12.04	command.search.expand_search.sourcetype	463	-	-
1,360.43	command.search.fieldalias	1,725	6,643,643	6,643,643
2,224.82	command.search.kv	1,725	-	-
788.76	command.search.lookups	1,725	6,643,643	6,643,643
477.51	command.search.rawdata	1,725	-	-
2.57	command.search.typer	1,725	2,127	2,127
0.00	command.search.tags	1,725	2,127	2,127
0.00	command.search.parse_directives	463	-	-
0.00	command.search.summary	2,092	-	-
0.34	command.table	1	2,127	21,452,129
6.94	command.timeliner	2,555	2,127	2,127

Execution Costs

| rest services/search/jobs/<sid>

1) **command.<command>**
   Running <command> throughout search
   Focus optimization on large durations first
2) **command.search.fieldalias**
   Renaming fields defined in props.conf
3) **command.search.kv**
   Applying search-time field extractions
4) **command.search.lookups**
   Creating new fields from automatic lookups

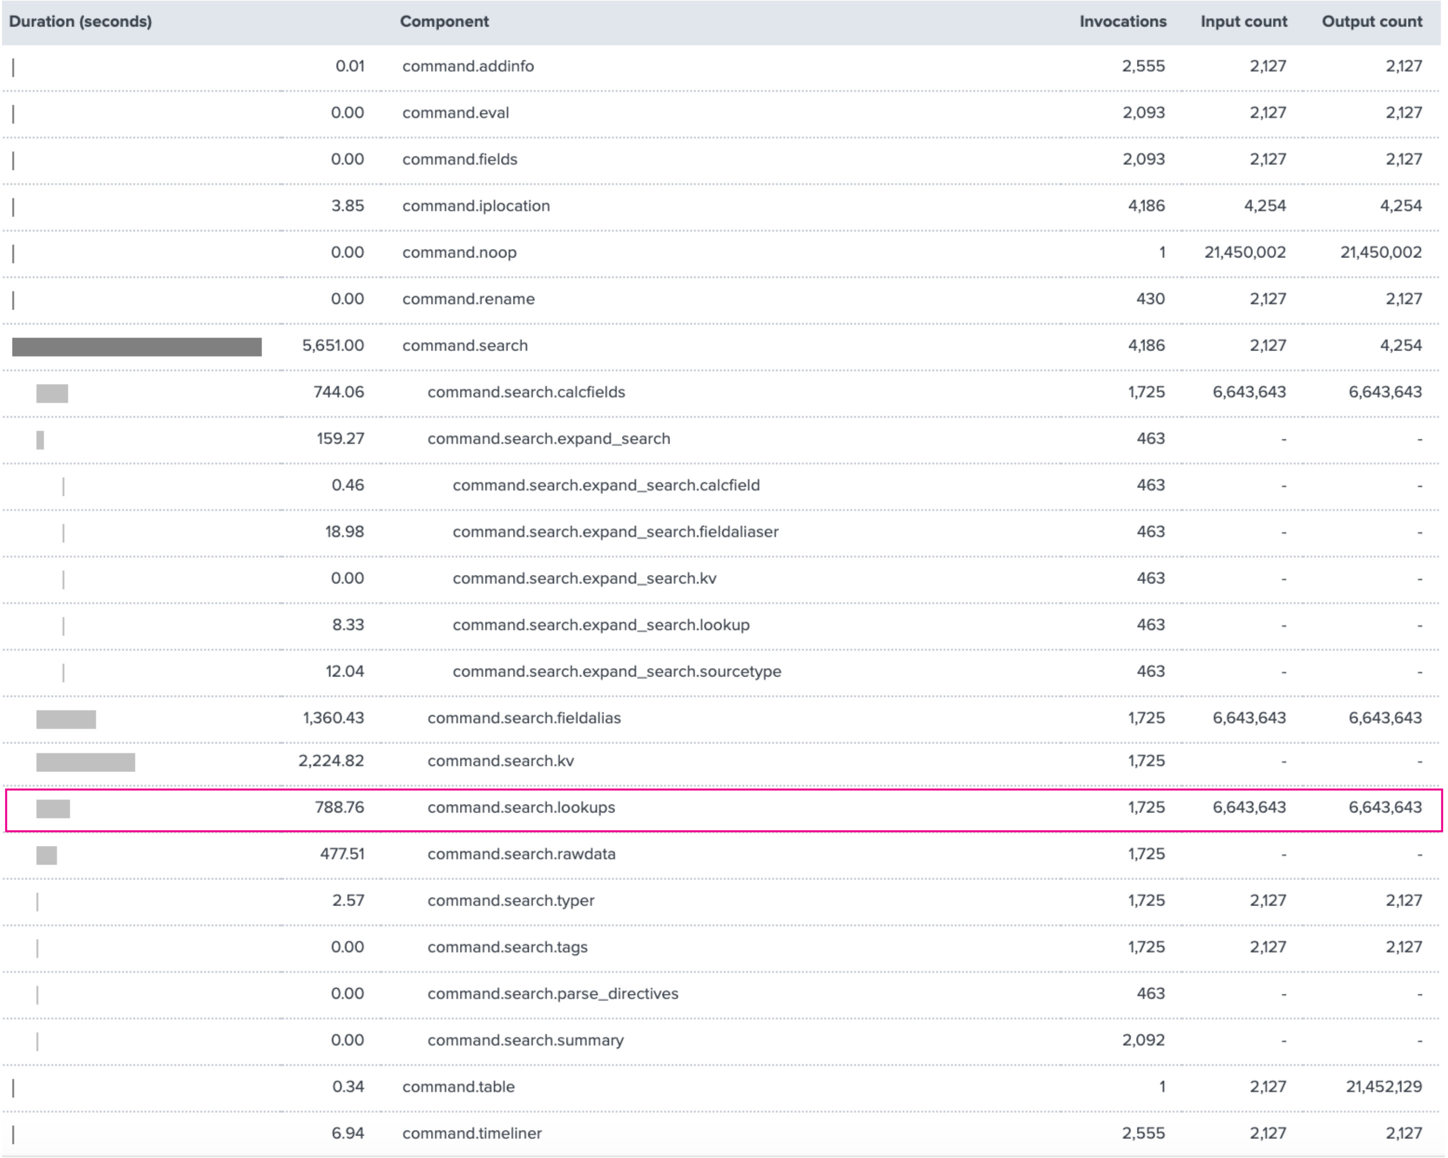

Duration (seconds)	Component	Invocations	Input count	Output count
0.01	command.addinfo	2,555	2,127	2,127
0.00	command.eval	2,093	2,127	2,127
0.00	command.fields	2,093	2,127	2,127
3.85	command.iplocation	4,186	4,254	4,254
0.00	command.noop	1	21,450,002	21,450,002
0.00	command.rename	430	2,127	2,127
5,651.00	command.search	4,186	2,127	4,254
744.06	command.search.calcfields	1,725	6,643,643	6,643,643
159.27	command.search.expand_search	463	-	-
0.46	command.search.expand_search.calcfield	463	-	-
18.98	command.search.expand_search.fieldaliaser	463	-	-
0.00	command.search.expand_search.kv	463	-	-
8.33	command.search.expand_search.lookup	463	-	-
12.04	command.search.expand_search.sourcetype	463	-	-
1,360.43	command.search.fieldalias	1,725	6,643,643	6,643,643
2,224.82	command.search.kv	1,725	-	-
788.76	command.search.lookups	1,725	6,643,643	6,643,643
477.51	command.search.rawdata	1,725	-	-
2.57	command.search.typer	1,725	2,127	2,127
0.00	command.search.tags	1,725	2,127	2,127
0.00	command.search.parse_directives	463	-	-
0.00	command.search.summary	2,092	-	-
0.34	command.table	1	2,127	21,452,129
6.94	command.timeliner	2,555	2,127	2,127

Execution Costs

| rest services/search/jobs/<sid>

1) **command.<command>**
   Running <command> throughout search
   Focus optimization on large durations first
2) **command.search.fieldalias**
   Renaming fields defined in props.conf
3) **command.search.kv**
   Applying search-time field extractions
4) **command.search.lookups**
   Creating new fields from automatic lookups
5) **command.search.typer**
   Assigning event types

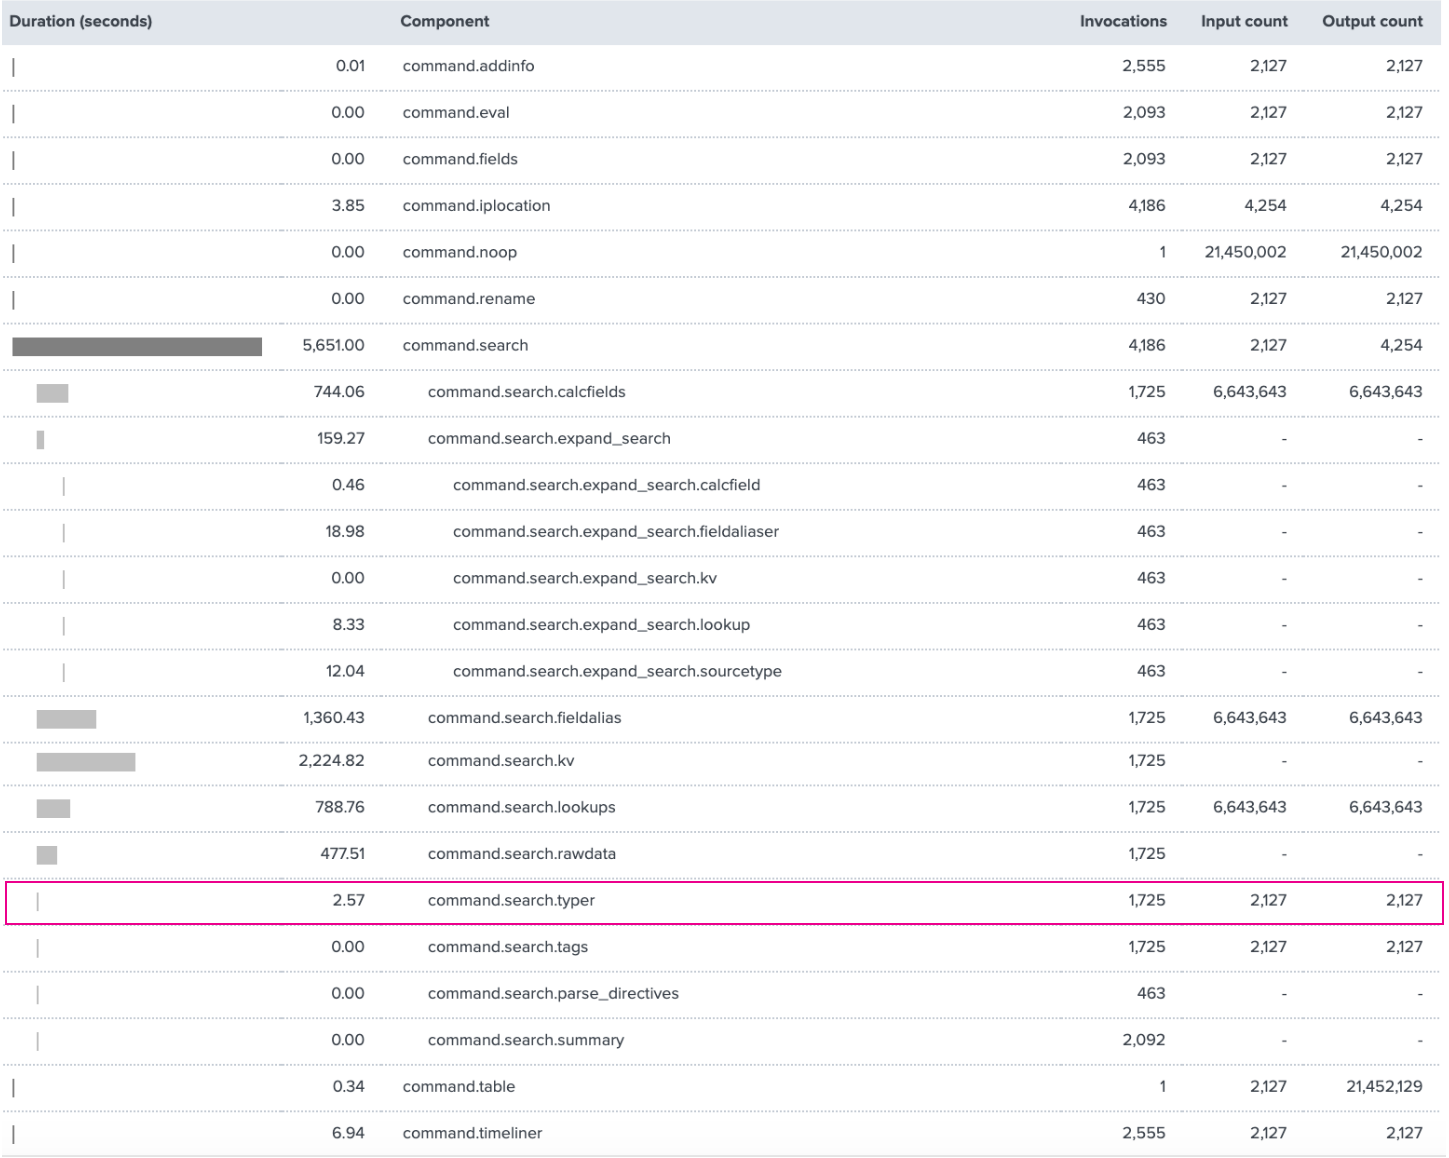

Duration (seconds)	Component	Invocations	Input count	Output count
0.01	command.addinfo	2,555	2,127	2,127
0.00	command.eval	2,093	2,127	2,127
0.00	command.fields	2,093	2,127	2,127
3.85	command.iplocation	4,186	4,254	4,254
0.00	command.noop	1	21,450,002	21,450,002
0.00	command.rename	430	2,127	2,127
5,651.00	command.search	4,186	2,127	4,254
744.06	command.search.calcfields	1,725	6,643,643	6,643,643
159.27	command.search.expand_search	463	-	-
0.46	command.search.expand_search.calcfield	463	-	-
18.98	command.search.expand_search.fieldaliaser	463	-	-
0.00	command.search.expand_search.kv	463	-	-
8.33	command.search.expand_search.lookup	463	-	-
12.04	command.search.expand_search.sourcetype	463	-	-
1,360.43	command.search.fieldalias	1,725	6,643,643	6,643,643
2,224.82	command.search.kv	1,725	-	-
788.76	command.search.lookups	1,725	6,643,643	6,643,643
477.51	command.search.rawdata	1,725	-	-
2.57	command.search.typer	1,725	2,127	2,127
0.00	command.search.tags	1,725	2,127	2,127
0.00	command.search.parse_directives	463	-	-
0.00	command.search.summary	2,092	-	-
0.34	command.table	1	2,127	21,452,129
6.94	command.timeliner	2,555	2,127	2,127

Execution Costs

| rest services/search/jobs/<sid>

1) **dispatch.stream.remote**

Time spent by the Indexers: Shows whether IDXs or the SH is doing the heavy lifting

Total bytes returned to SH: Lower numbers indicate more efficient search and first reporting command

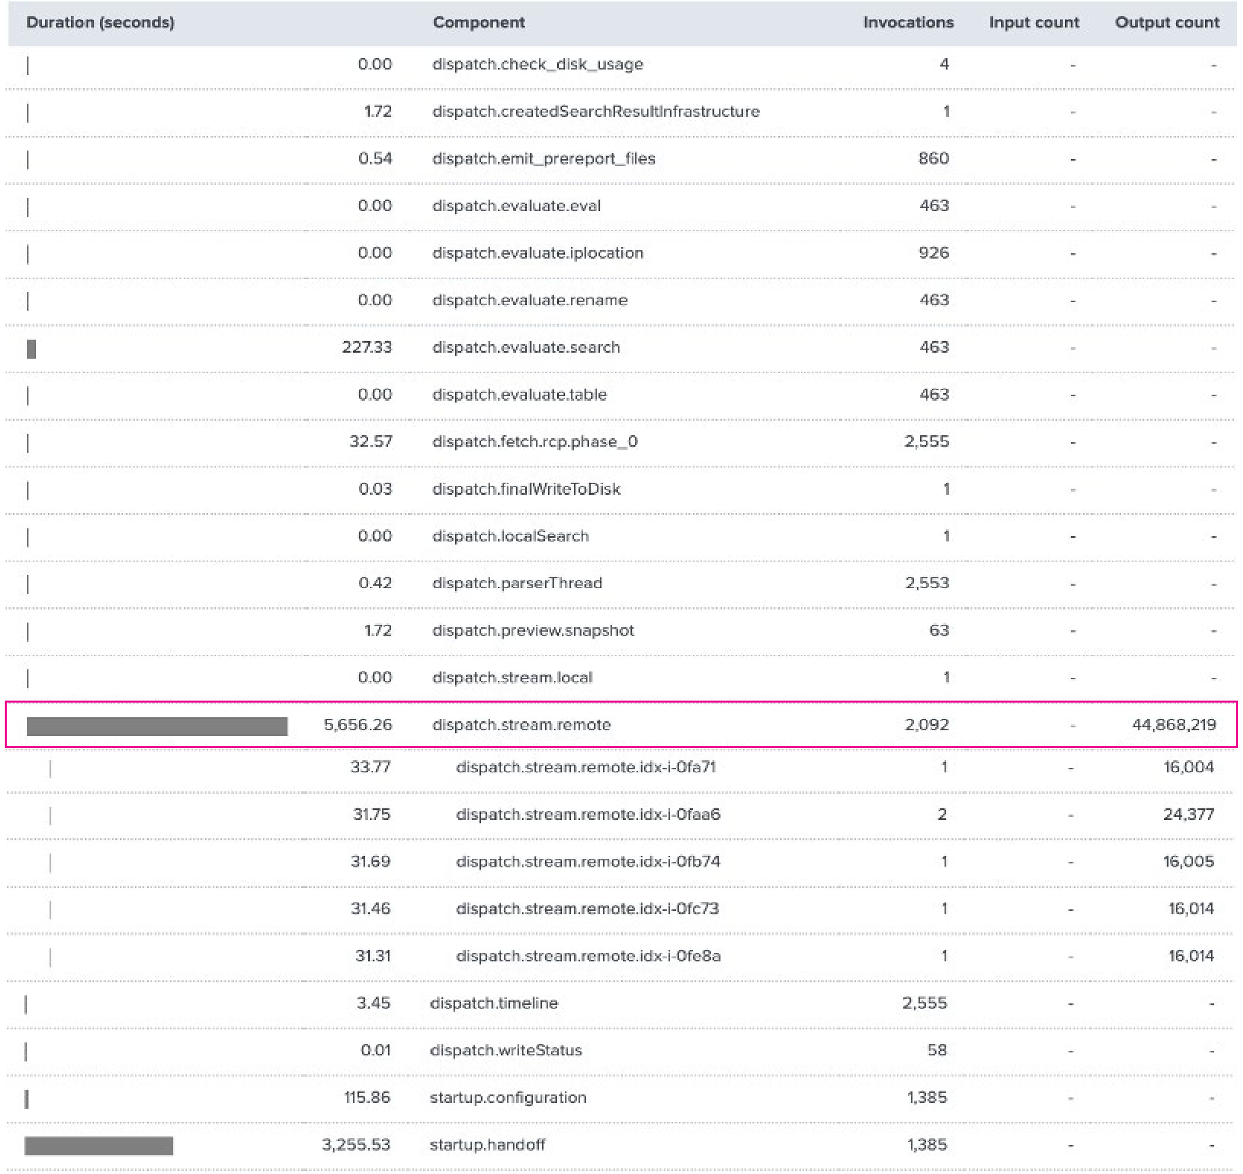

Duration (seconds)	Component	Invocations	Input count	Output count
0.00	dispatch.check_disk_usage	4	-	-
1.72	dispatch.createdSearchResultInfrastructure	1	-	-
0.54	dispatch.emit_prereport_files	860	-	-
0.00	dispatch.evaluate.eval	463	-	-
0.00	dispatch.evaluate.iplocation	926	-	-
0.00	dispatch.evaluate.rename	463	-	-
227.33	dispatch.evaluate.search	463	-	-
0.00	dispatch.evaluate.table	463	-	-
32.57	dispatch.fetch.rcp.phase_0	2,555	-	-
0.03	dispatch.finalWriteToDisk	1	-	-
0.00	dispatch.localSearch	1	-	-
0.42	dispatch.parserThread	2,553	-	-
1.72	dispatch.preview.snapshot	63	-	-
0.00	dispatch.stream.local	1	-	-
5,656.26	dispatch.stream.remote	2,092	-	44,868,219
33.77	dispatch.stream.remote.idx-i-0fa71	1	-	16,004
31.75	dispatch.stream.remote.idx-i-0faa6	2	-	24,377
31.69	dispatch.stream.remote.idx-i-0fb74	1	-	16,005
31.46	dispatch.stream.remote.idx-i-0fc73	1	-	16,014
31.31	dispatch.stream.remote.idx-i-0fe8a	1	-	16,014
3.45	dispatch.timeline	2,555	-	-
0.01	dispatch.writeStatus	58	-	-
115.86	startup.configuration	1,385	-	-
3,255.53	startup.handoff	1,385	-	-

Execution Costs

| rest services/search/jobs/<sid>

1) **dispatch.stream.remote**

   Time spent by the Indexers: Shows whether IDXs or the SH is doing the heavy lifting
   Total bytes returned to SH: Lower numbers indicate more efficient search and first reporting command

2) **dispatch.stream.remote.<IDX>**

   Large differences between Indexers can indicate poor data balance (different bytes returned) or lower hardware performance (similar bytes but slower)

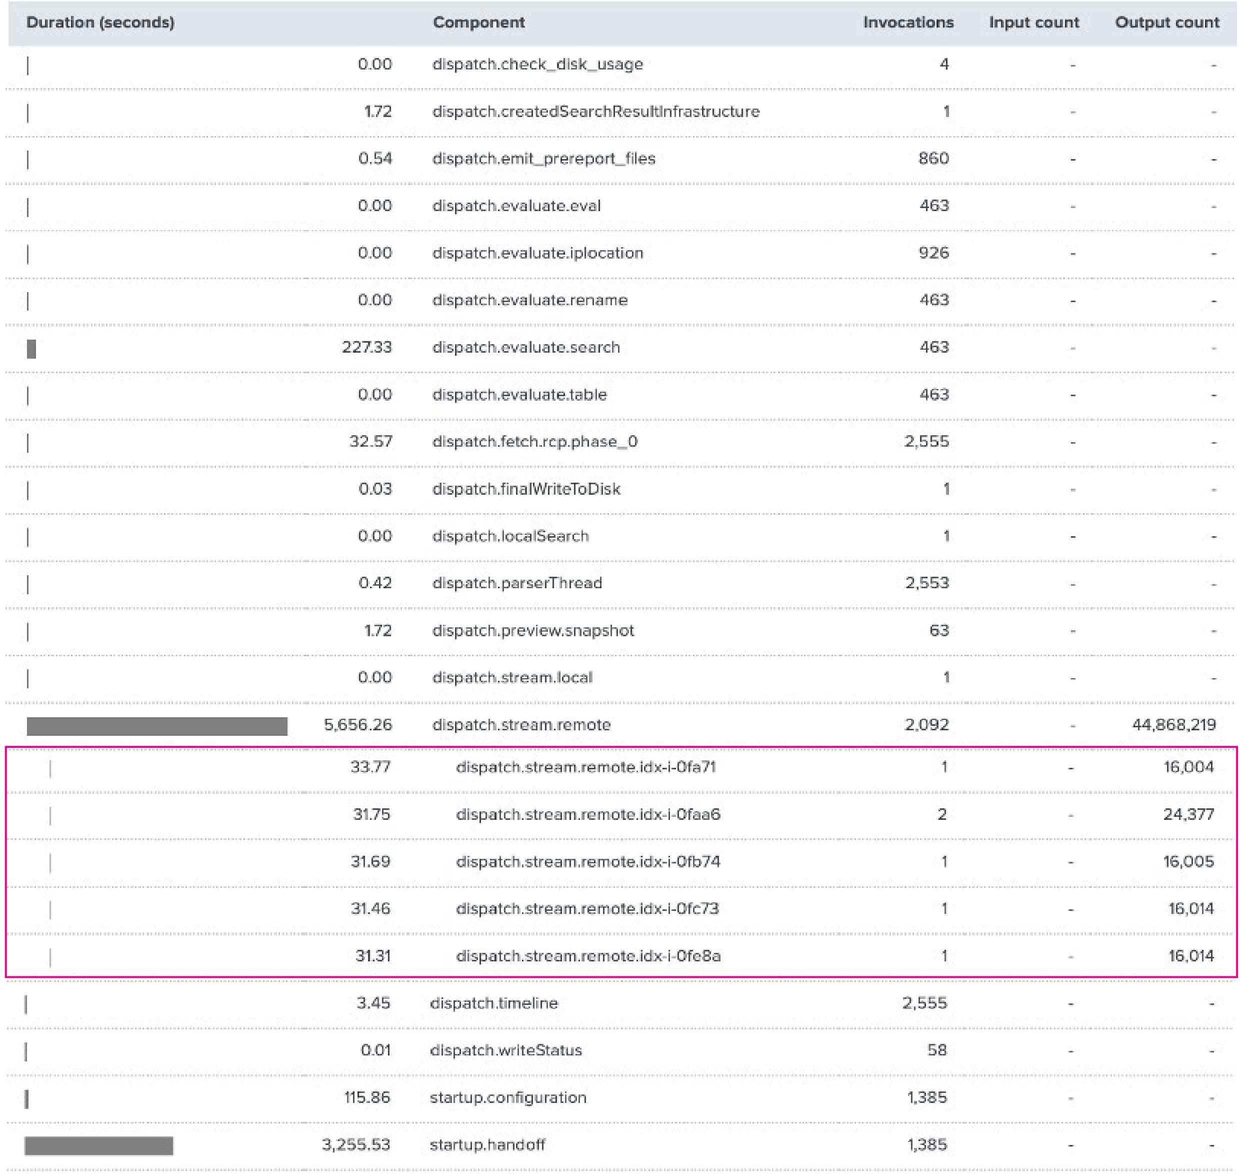

Duration (seconds)	Component	Invocations	Input count	Output count
0.00	dispatch.check_disk_usage	4	-	-
1.72	dispatch.createdSearchResultInfrastructure	1	-	-
0.54	dispatch.emit_prereport_files	860	-	-
0.00	dispatch.evaluate.eval	463	-	-
0.00	dispatch.evaluate.iplocation	926	-	-
0.00	dispatch.evaluate.rename	463	-	-
227.33	dispatch.evaluate.search	463	-	-
0.00	dispatch.evaluate.table	463	-	-
32.57	dispatch.fetch.rcp.phase_0	2,555	-	-
0.03	dispatch.finalWriteToDisk	1	-	-
0.00	dispatch.localSearch	1	-	-
0.42	dispatch.parserThread	2,553	-	-
1.72	dispatch.preview.snapshot	63	-	-
0.00	dispatch.stream.local	1	-	-
5,656.26	dispatch.stream.remote	2,092	-	44,868,219
33.77	dispatch.stream.remote.idx-i-0fa71	1	-	16,004
31.75	dispatch.stream.remote.idx-i-0faa6	2	-	24,377
31.69	dispatch.stream.remote.idx-i-0fb74	1	-	16,005
31.46	dispatch.stream.remote.idx-i-0fc73	1	-	16,014
31.31	dispatch.stream.remote.idx-i-0fe8a	1	-	16,014
3.45	dispatch.timeline	2,555	-	-
0.01	dispatch.writeStatus	58	-	-
115.86	startup.configuration	1,385	-	-
3,255.53	startup.handoff	1,385	-	-

Execution Costs

| rest services/search/jobs/<sid>

1) **dispatch.stream.remote**

   Time spent by the Indexers: Shows whether IDXs or the SH is doing the heavy lifting
   Total bytes returned to SH: Lower numbers indicate more efficient search and first reporting command

2) **dispatch.stream.remote.<IDX>**

   Large differences between Indexers can indicate poor data balance (different bytes returned) or lower hardware performance (similar bytes but slower)

3) **startup.handoff**

   Cumulative time spent communicating to all Indexers and setting up the search processes
   Large durations relative to Indexer count can indicate network issues or overloaded Indexers

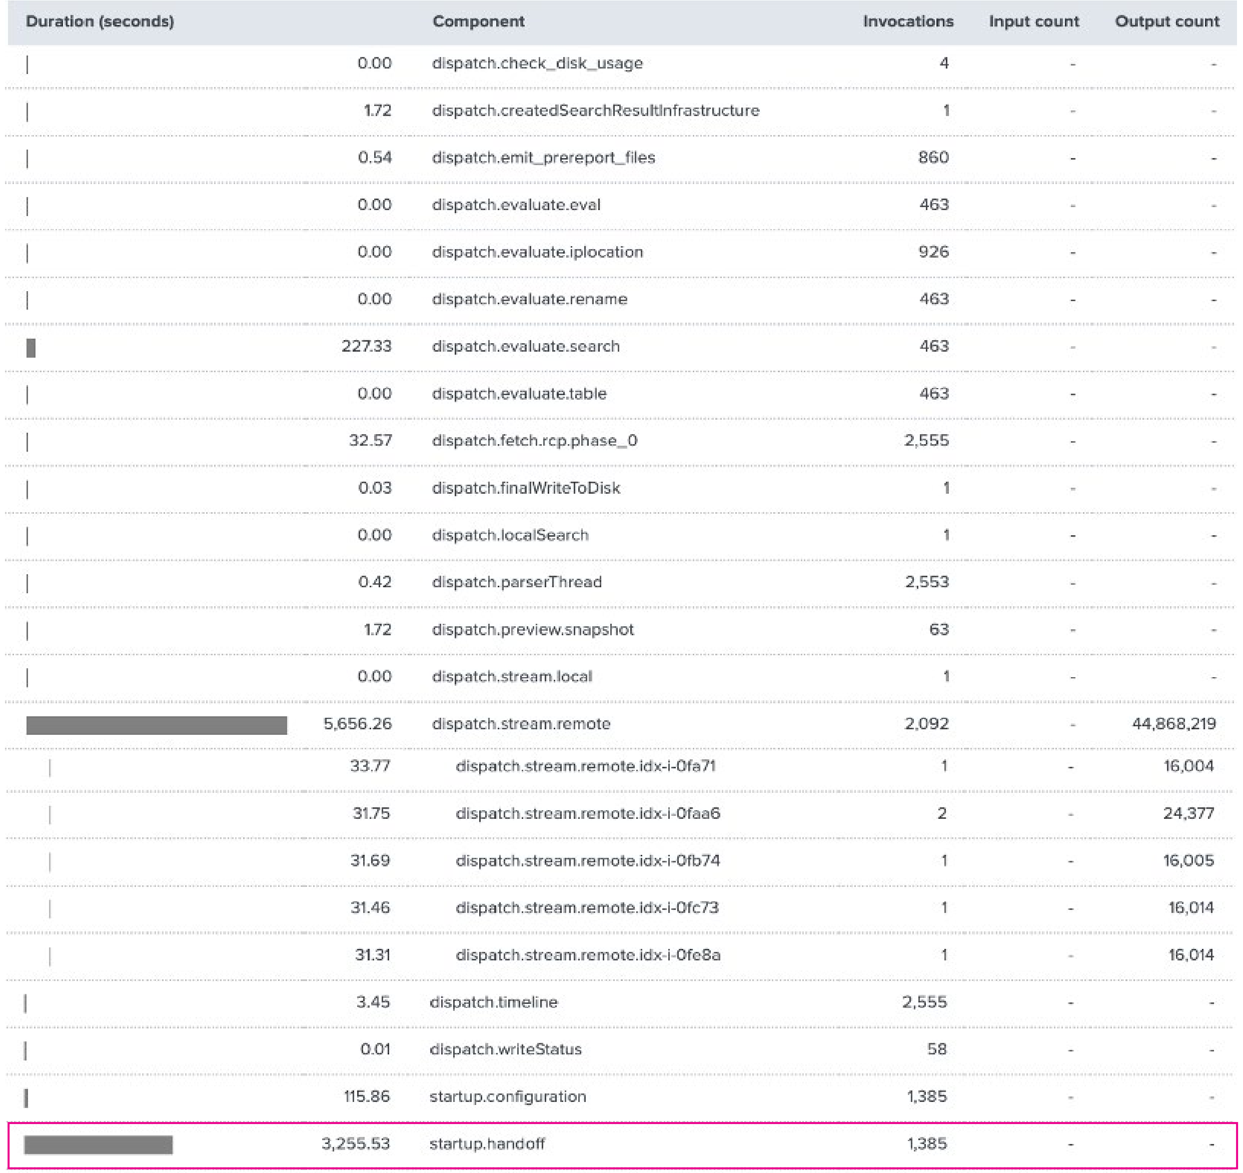

Duration (seconds)	Component	Invocations	Input count	Output count
0.00	dispatch.check_disk_usage	4	-	-
1.72	dispatch.createdSearchResultInfrastructure	1	-	-
0.54	dispatch.emit_prereport_files	860	-	-
0.00	dispatch.evaluate.eval	463	-	-
0.00	dispatch.evaluate.iplocation	926	-	-
0.00	dispatch.evaluate.rename	463	-	-
227.33	dispatch.evaluate.search	463	-	-
0.00	dispatch.evaluate.table	463	-	-
32.57	dispatch.fetch.rcp.phase_0	2,555	-	-
0.03	dispatch.finalWriteToDisk	1	-	-
0.00	dispatch.localSearch	1	-	-
0.42	dispatch.parserThread	2,553	-	-
1.72	dispatch.preview.snapshot	63	-	-
0.00	dispatch.stream.local	1	-	-
5,656.26	dispatch.stream.remote	2,092	-	44,868,219
33.77	dispatch.stream.remote.idx-i-0fa71	1	-	16,004
31.75	dispatch.stream.remote.idx-i-0faa6	2	-	24,377
31.69	dispatch.stream.remote.idx-i-0fb74	1	-	16,005
31.46	dispatch.stream.remote.idx-i-0fc73	1	-	16,014
31.31	dispatch.stream.remote.idx-i-0fe8a	1	-	16,014
3.45	dispatch.timeline	2,555	-	-
0.01	dispatch.writeStatus	58	-	-
115.86	startup.configuration	1,385	-	-
3,255.53	startup.handoff	1,385	-	-

Example: NTP Inspector

Indexer time synchronization is broken → startup.handoff shows implausible number

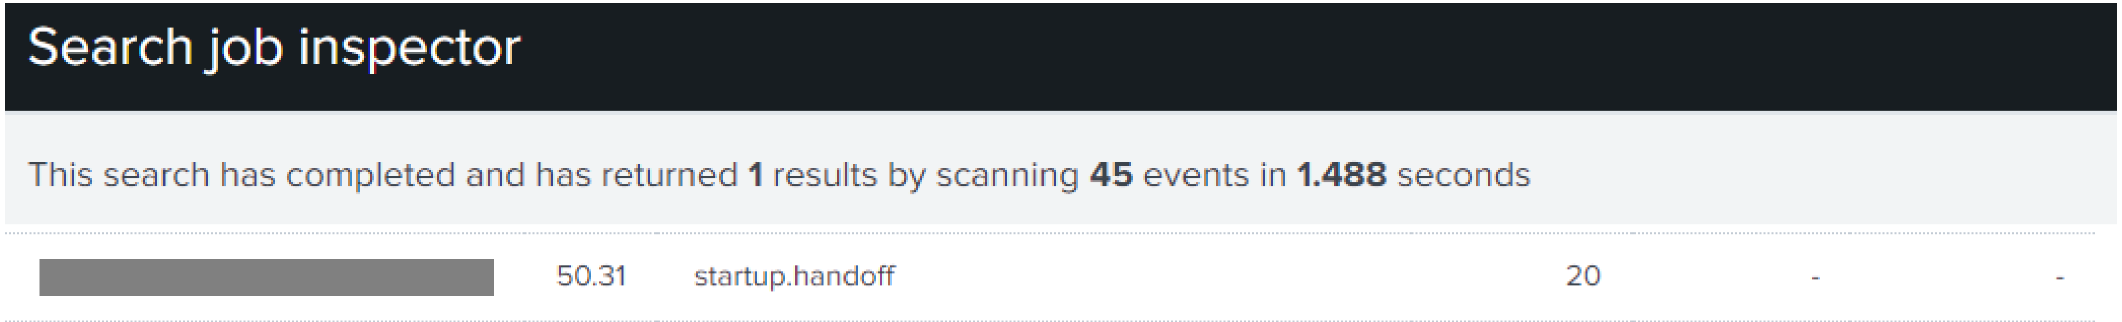

Search job inspector

This search has completed and has returned **1** results by scanning **45** events in **1.488** seconds

50.31	startup.handoff	20	-	-

search.log to the rescue:

```
IDX0: 08-04-2020 20:40:32.888 INFO  dispatchRunner - Search process [...]
IDX1: 08-04-2020 20:40:32.887 INFO  dispatchRunner - Search process [...]
IDX2: 08-04-2020 20:40:40.579 INFO  dispatchRunner - Search process [...]
IDX3: 08-04-2020 20:40:47.969 INFO  dispatchRunner - Search process [...]
```

→ IDX2+3 need their NTP config fixed

Execution Costs

| rest services/search/jobs/<sid>

1) **command.search.index**

   Time spent crawling the tsidx files

2) **command.search.index.bucketcache**

   Number of SmartStore buckets hit in the cache or missed and downloaded from S3

3) **command.search.index.usec_N_M**

   Number of IO operations that took between N and M microseconds

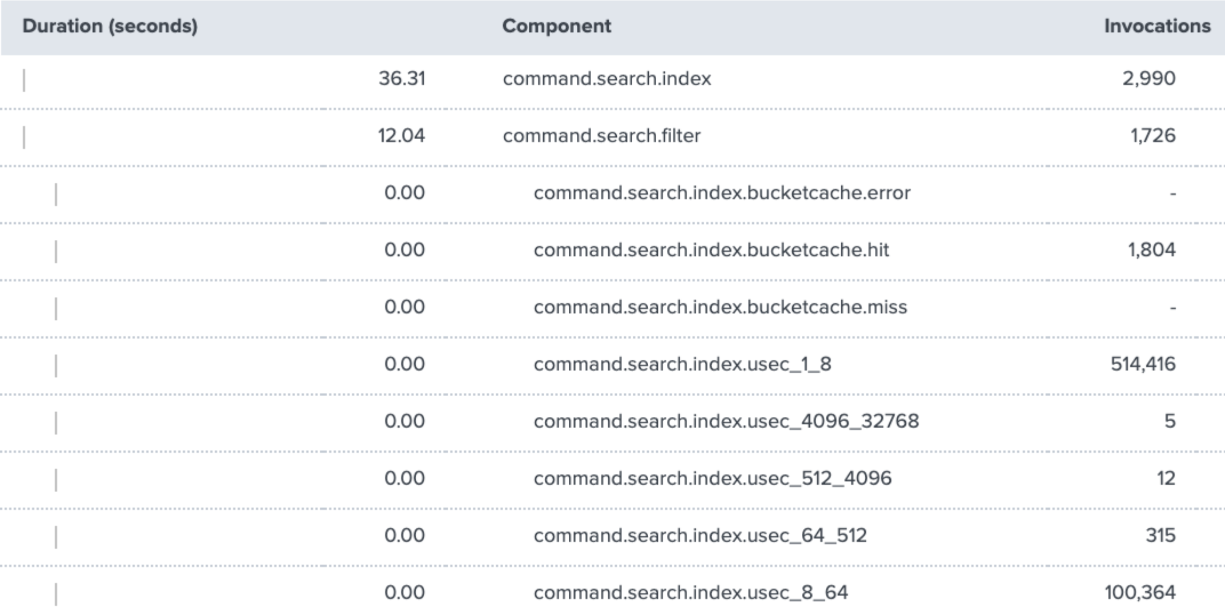

Duration (seconds)	Component	Invocations
36.31	command.search.index	2,990
12.04	command.search.filter	1,726
0.00	command.search.index.bucketcache.error	-
0.00	command.search.index.bucketcache.hit	1,804
0.00	command.search.index.bucketcache.miss	-
0.00	command.search.index.usec_1_8	514,416
0.00	command.search.index.usec_4096_32768	5
0.00	command.search.index.usec_512_4096	12
0.00	command.search.index.usec_64_512	315
0.00	command.search.index.usec_8_64	100,364

.conf20

splunk>

Search Job Properties

Wham!

Search Job Properties

| rest services/search/jobs/<sid>

1) **scanCount**

   Events loaded off disk and ”field extracted” (props.conf applied)

2) **eventCount**

   Number of events returned by the generating search command

3) **resultCount**

   Number of results returned from last command

4) **runDuration**

   Duration (in seconds) search took to run

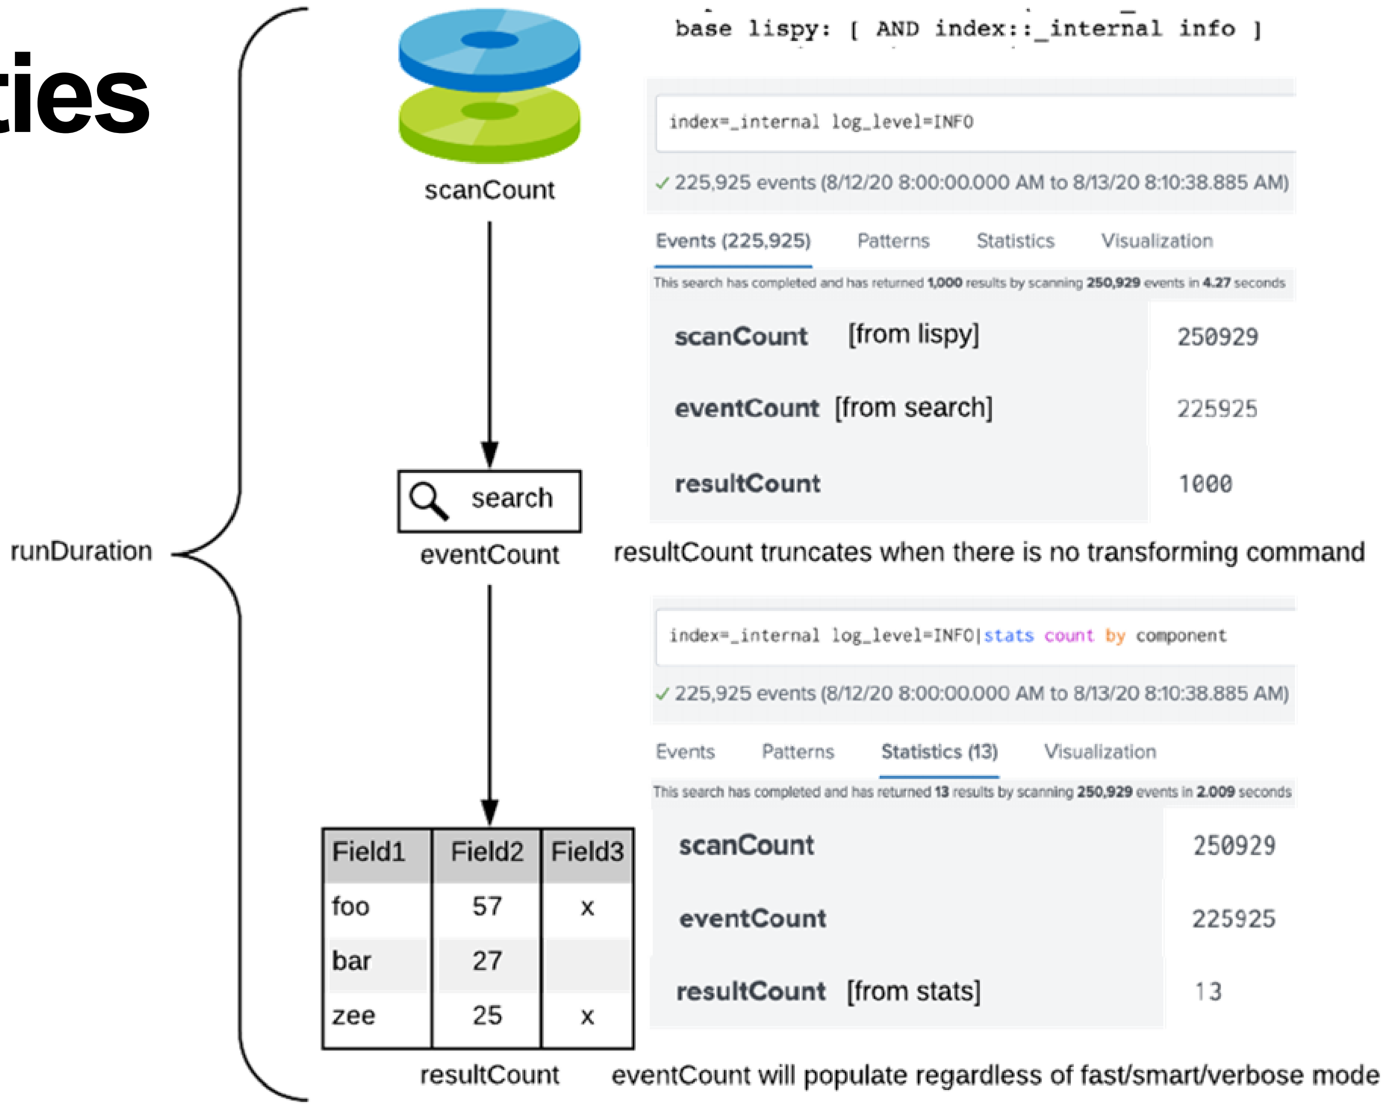

```
base lispy: [ AND index::_internal info ]
```

```
index=_internal log_level=INFO
```

✓ 225,925 events (8/12/20 8:00:00.000 AM to 8/13/20 8:10:38.885 AM)

Events (225,925) Patterns Statistics Visualization

This search has completed and has returned **1,000** results by scanning **250,929** events in **4.27** seconds

scanCount [from lispy]	250929
eventCount [from search]	225925
resultCount	1000

resultCount truncates when there is no transforming command

```
index=_internal log_level=INFO|stats count by component
```

✓ 225,925 events (8/12/20 8:00:00.000 AM to 8/13/20 8:10:38.885 AM)

Events Patterns Statistics (13) Visualization

This search has completed and has returned **13** results by scanning **250,929** events in **2.009** seconds

scanCount	250929
eventCount	225925
resultCount [from stats]	13

Field1	Field2	Field3
foo	57	x
bar	27	
zee	25	x

eventCount will populate regardless of fast/smart/verbose mode

Search Job Properties

| rest services/search/jobs/<sid>

1) **optimizedSearch**

   Search ran by Splunk after optimizing SPL
   (none if optimizer doesn't run:
   |noop search_optimization=false)

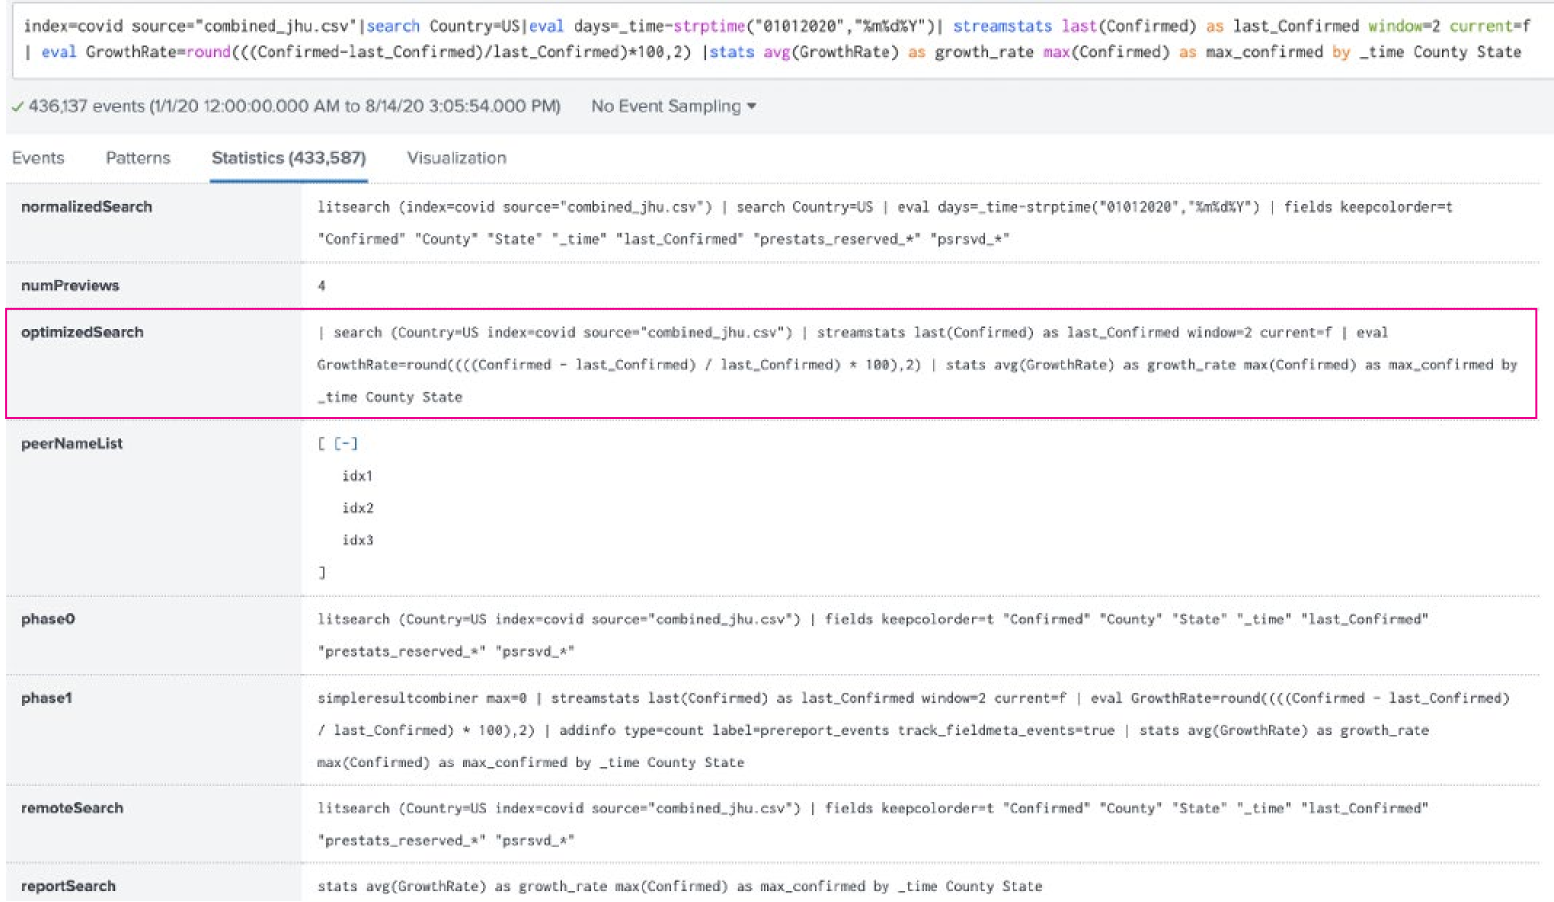

```
index=covid source="combined_jhu.csv"|search Country=US|eval days=_time-strptime("01012020","%m%d%Y")| streamstats last(Confirmed) as last_Confirmed window=2 current=f
| eval GrowthRate=round(((Confirmed-last_Confirmed)/last_Confirmed)*100,2) |stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State
```

✓ 436,137 events (1/1/20 12:00:00.000 AM to 8/14/20 3:05:54.000 PM)

normalizedSearch	`litsearch (index=covid source="combined_jhu.csv") | search Country=US | eval days=_time-strptime("01012020","%m%d%Y") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"`
numPreviews	4
optimizedSearch	`| search (Country=US index=covid source="combined_jhu.csv") | streamstats last(Confirmed) as last_Confirmed window=2 current=f | eval GrowthRate=round((((Confirmed - last_Confirmed) / last_Confirmed) * 100),2) | stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State`
peerNameList	[[-] idx1 idx2 idx3]
phase0	`litsearch (Country=US index=covid source="combined_jhu.csv") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"`
phase1	`simpleresultcombiner max=0 | streamstats last(Confirmed) as last_Confirmed window=2 current=f | eval GrowthRate=round((((Confirmed - last_Confirmed) / last_Confirmed) * 100),2) | addinfo type=count label=prereport_events track_fieldmeta_events=true | stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State`
remoteSearch	`litsearch (Country=US index=covid source="combined_jhu.csv") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"`
reportSearch	`stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State`

Search Job Properties

| rest services/search/jobs/<sid>

1) **optimizedSearch**

   Search ran by Splunk after optimizing SPL
   (none if optimizer doesn't run: |noop search_optimization=false)

2) **phase0**

   Search ran by indexers – expands field aliases, calculated fields, etc.

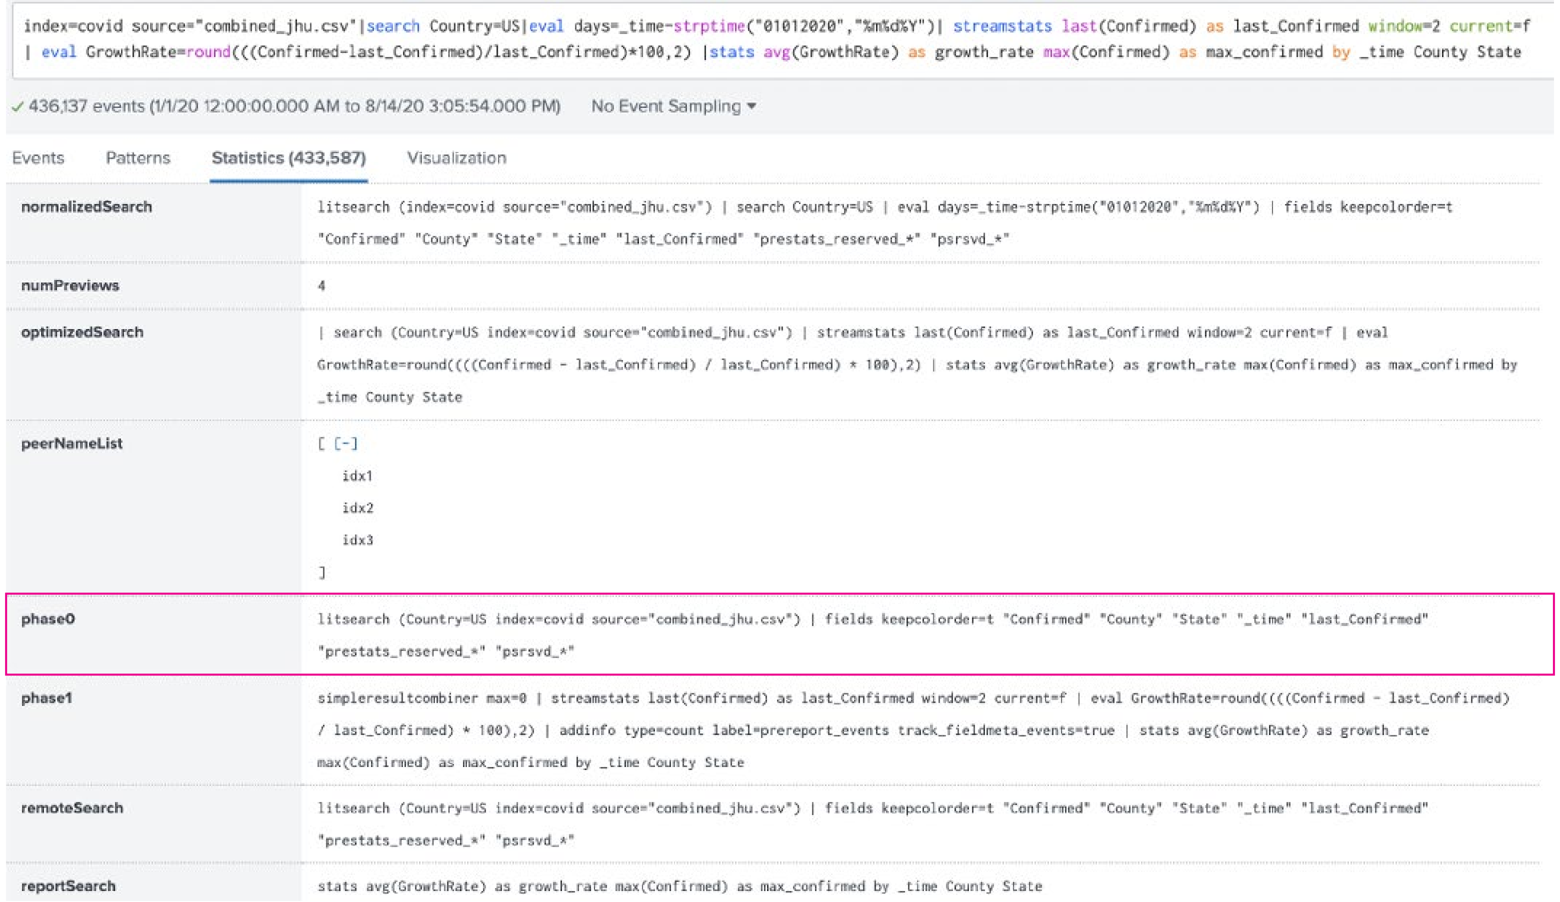

```
index=covid source="combined_jhu.csv"|search Country=US|eval days=_time-strptime("01012020","%m%d%Y")| streamstats last(Confirmed) as last_Confirmed window=2 current=f
| eval GrowthRate=round(((Confirmed-last_Confirmed)/last_Confirmed)*100,2) |stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State
```

✓ 436,137 events (1/1/20 12:00:00.000 AM to 8/14/20 3:05:54.000 PM)

<table>
<tr><td>normalizedSearch</td><td>litsearch (index=covid source="combined_jhu.csv") | search Country=US | eval days=_time-strptime("01012020","%m%d%Y") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"</td></tr>
<tr><td>numPreviews</td><td>4</td></tr>
<tr><td>optimizedSearch</td><td>| search (Country=US index=covid source="combined_jhu.csv") | streamstats last(Confirmed) as last_Confirmed window=2 current=f | eval GrowthRate=round((((Confirmed - last_Confirmed) / last_Confirmed) * 100),2) | stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State</td></tr>
<tr><td>peerNameList</td><td>[[-]
idx1
idx2
idx3
]</td></tr>
<tr><td>phase0</td><td>litsearch (Country=US index=covid source="combined_jhu.csv") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"</td></tr>
<tr><td>phase1</td><td>simpleresultcombiner max=0 | streamstats last(Confirmed) as last_Confirmed window=2 current=f | eval GrowthRate=round((((Confirmed - last_Confirmed) / last_Confirmed) * 100),2) | addinfo type=count label=prereport_events track_fieldmeta_events=true | stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State</td></tr>
<tr><td>remoteSearch</td><td>litsearch (Country=US index=covid source="combined_jhu.csv") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"</td></tr>
<tr><td>reportSearch</td><td>stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State</td></tr>
</table>

Search Job Properties

| rest services/search/jobs/<sid>

1) **optimizedSearch**
   Search ran by Splunk after optimizing SPL
   (none if optimizer doesn't run:
   |noop search_optimization=false)
2) **phase0**
   Search ran by indexers – expands field aliases, calculated fields, etc.
3) **phase1**
   Executing the streaming and transformatin commands

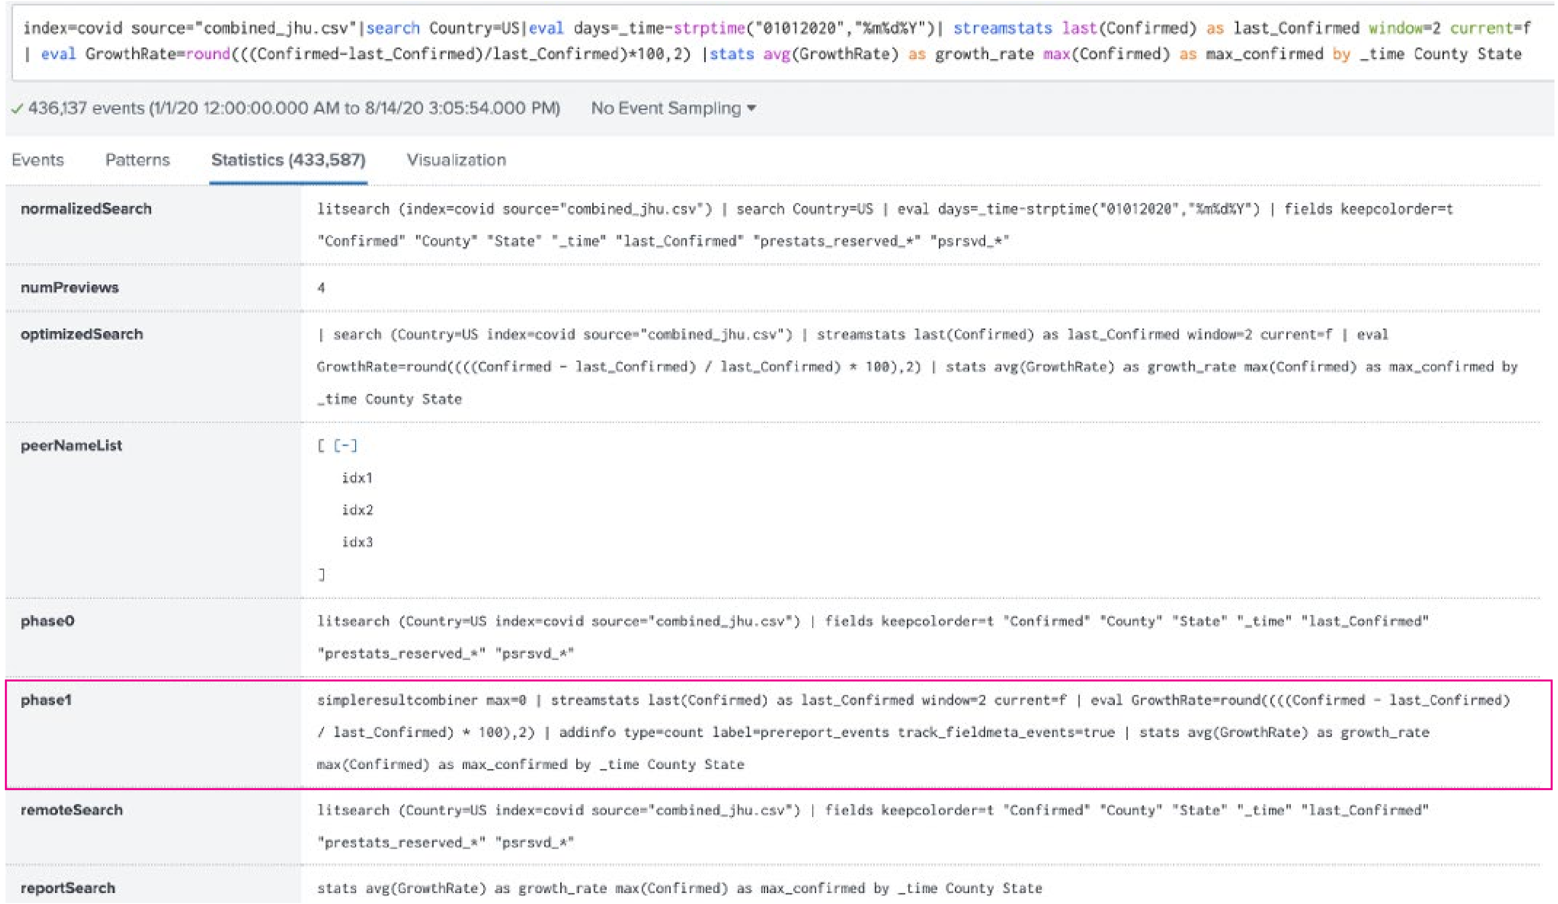

```
index=covid source="combined_jhu.csv"|search Country=US|eval days=_time-strptime("01012020","%m%d%Y")| streamstats last(Confirmed) as last_Confirmed window=2 current=f
| eval GrowthRate=round(((Confirmed-last_Confirmed)/last_Confirmed)*100,2) |stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State
```

✓ 436,137 events (1/1/20 12:00:00.000 AM to 8/14/20 3:05:54.000 PM)

normalizedSearch	litsearch (index=covid source="combined_jhu.csv") | search Country=US | eval days=_time-strptime("01012020","%m%d%Y") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"
numPreviews	4
optimizedSearch	| search (Country=US index=covid source="combined_jhu.csv") | streamstats last(Confirmed) as last_Confirmed window=2 current=f | eval GrowthRate=round((((Confirmed - last_Confirmed) / last_Confirmed) * 100),2) | stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State
peerNameList	[[-] idx1 idx2 idx3]
phase0	litsearch (Country=US index=covid source="combined_jhu.csv") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"
phase1	simpleresultcombiner max=0 | streamstats last(Confirmed) as last_Confirmed window=2 current=f | eval GrowthRate=round((((Confirmed - last_Confirmed) / last_Confirmed) * 100),2) | addinfo type=count label=prereport_events track_fieldmeta_events=true | stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State
remoteSearch	litsearch (Country=US index=covid source="combined_jhu.csv") | fields keepcolorder=t "Confirmed" "County" "State" "_time" "last_Confirmed" "prestats_reserved_*" "psrsvd_*"
reportSearch	stats avg(GrowthRate) as growth_rate max(Confirmed) as max_confirmed by _time County State

Improve Job Inspector

Add debugging

DEBUG logging to job inspector to troubleshoot

`In limits.conf:`

```
[search_info]
infocsv_log_level = DEBUG
```

- Pulls lispy to the top of the job inspector
- Displays subsearch output in the top of the job inspector

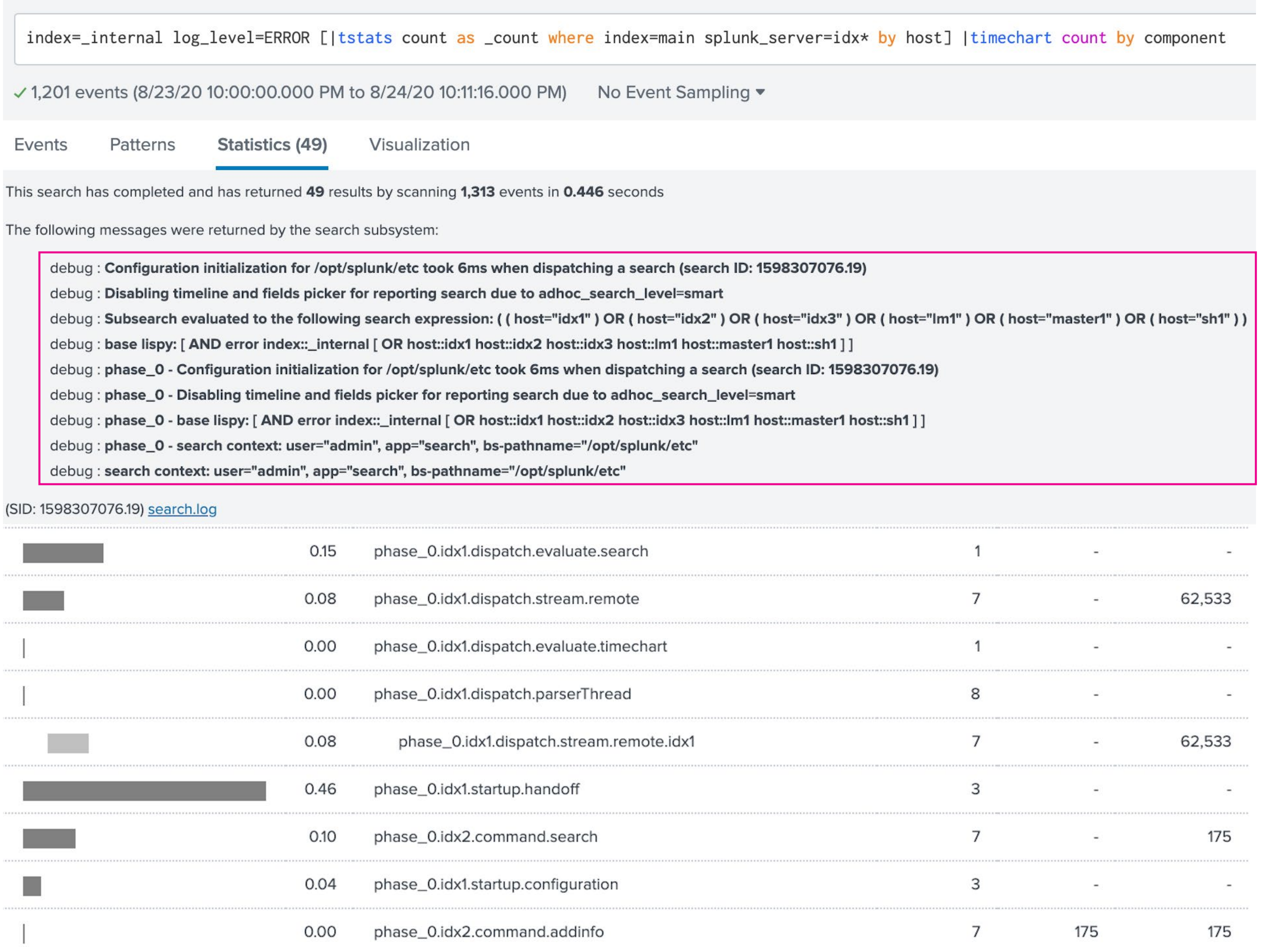

```
index=_internal log_level=ERROR [|tstats count as _count where index=main splunk_server=idx* by host] |timechart count by component
```

✓ 1,201 events (8/23/20 10:00:00.000 PM to 8/24/20 10:11:16.000 PM)   No Event Sampling ▾

This search has completed and has returned **49** results by scanning **1,313** events in **0.446** seconds

The following messages were returned by the search subsystem:

```
debug : Configuration initialization for /opt/splunk/etc took 6ms when dispatching a search (search ID: 1598307076.19)
debug : Disabling timeline and fields picker for reporting search due to adhoc_search_level=smart
debug : Subsearch evaluated to the following search expression: ( ( host="idx1" ) OR ( host="idx2" ) OR ( host="idx3" ) OR ( host="lm1" ) OR ( host="master1" ) OR ( host="sh1" ) )
debug : base lispy: [ AND error index::_internal [ OR host::idx1 host::idx2 host::idx3 host::lm1 host::master1 host::sh1 ] ]
debug : phase_0 - Configuration initialization for /opt/splunk/etc took 6ms when dispatching a search (search ID: 1598307076.19)
debug : phase_0 - Disabling timeline and fields picker for reporting search due to adhoc_search_level=smart
debug : phase_0 - base lispy: [ AND error index::_internal [ OR host::idx1 host::idx2 host::idx3 host::lm1 host::master1 host::sh1 ] ]
debug : phase_0 - search context: user="admin", app="search", bs-pathname="/opt/splunk/etc"
debug : search context: user="admin", app="search", bs-pathname="/opt/splunk/etc"
```

(SID: 1598307076.19) search.log

0.15	phase_0.idx1.dispatch.evaluate.search	1	-	-
0.08	phase_0.idx1.dispatch.stream.remote	7	-	62,533
0.00	phase_0.idx1.dispatch.evaluate.timechart	1	-	-
0.00	phase_0.idx1.dispatch.parserThread	8	-	-
0.08	phase_0.idx1.dispatch.stream.remote.idx1	7	-	62,533
0.46	phase_0.idx1.startup.handoff	3	-	-
0.10	phase_0.idx2.command.search	7	-	175
0.04	phase_0.idx1.startup.configuration	3	-	-
0.00	phase_0.idx2.command.addinfo	7	175	175

Improve Job Inspector

Add debugging

DEBUG logging to job inspector to troubleshoot

In limits.conf:

```
[search_info]
infocsv_log_level = DEBUG
```

- Pulls lispy to the top of the job inspector
- Displays subsearch output in the top of the job inspector

```
[search_metrics]
debug_metrics = true
```

- Outputs detailed search metrics, such as command time spent on each peer
- Not related to metrics.log.

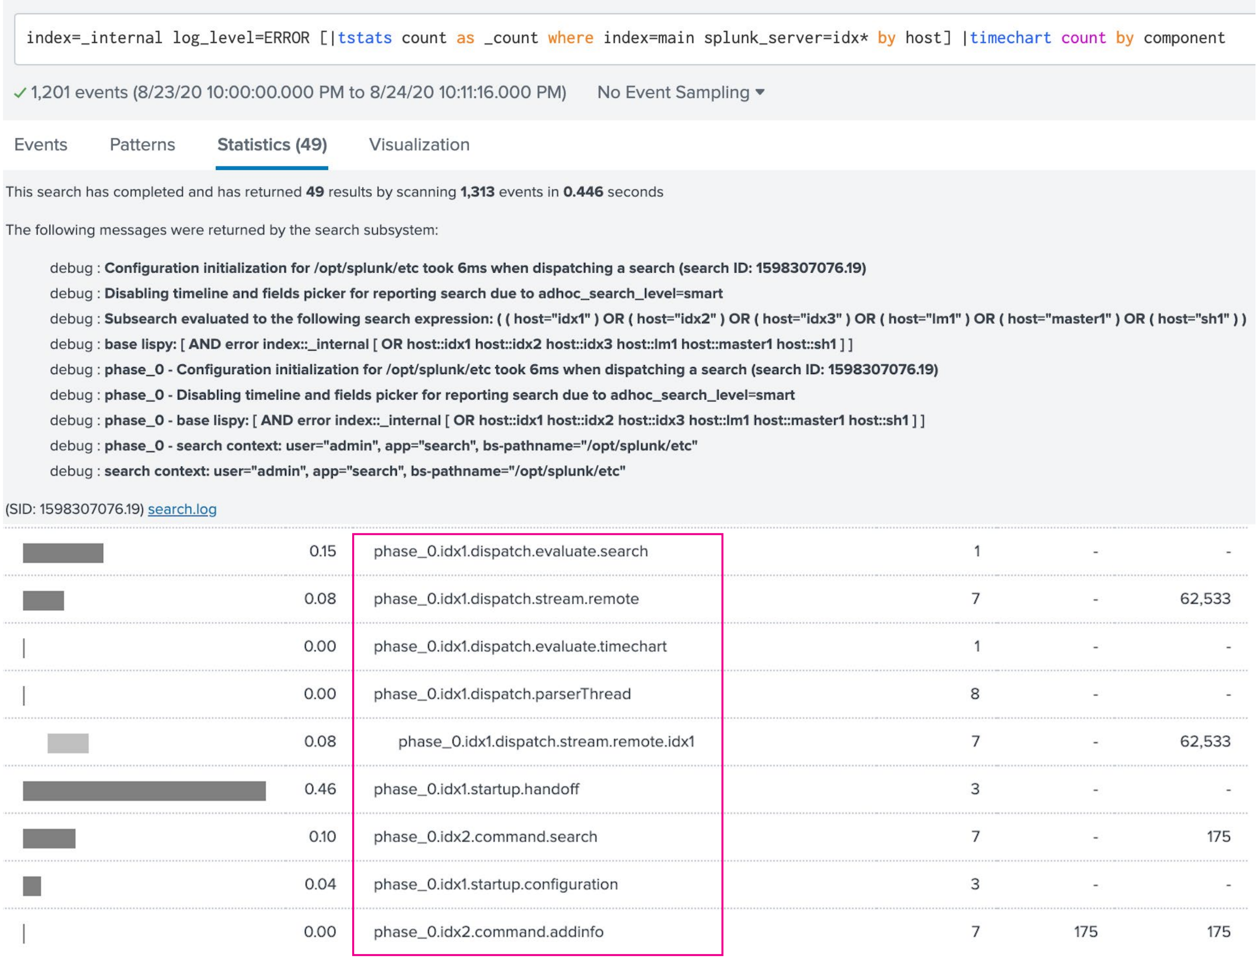

```
index=_internal log_level=ERROR [|tstats count as _count where index=main splunk_server=idx* by host] |timechart count by component
```

✓ 1,201 events (8/23/20 10:00:00.000 PM to 8/24/20 10:11:16.000 PM) No Event Sampling ▾

This search has completed and has returned **49** results by scanning **1,313** events in **0.446** seconds

The following messages were returned by the search subsystem:

```
debug : Configuration initialization for /opt/splunk/etc took 6ms when dispatching a search (search ID: 1598307076.19)
debug : Disabling timeline and fields picker for reporting search due to adhoc_search_level=smart
debug : Subsearch evaluated to the following search expression: ( ( host="idx1" ) OR ( host="idx2" ) OR ( host="idx3" ) OR ( host="lm1" ) OR ( host="master1" ) OR ( host="sh1" ) )
debug : base lispy: [ AND error index::_internal [ OR host::idx1 host::idx2 host::idx3 host::lm1 host::master1 host::sh1 ] ]
debug : phase_0 - Configuration initialization for /opt/splunk/etc took 6ms when dispatching a search (search ID: 1598307076.19)
debug : phase_0 - Disabling timeline and fields picker for reporting search due to adhoc_search_level=smart
debug : phase_0 - base lispy: [ AND error index::_internal [ OR host::idx1 host::idx2 host::idx3 host::lm1 host::master1 host::sh1 ] ]
debug : phase_0 - search context: user="admin", app="search", bs-pathname="/opt/splunk/etc"
debug : search context: user="admin", app="search", bs-pathname="/opt/splunk/etc"
```

(SID: 1598307076.19) search.log

0.15	phase_0.idx1.dispatch.evaluate.search	1	-	-
0.08	phase_0.idx1.dispatch.stream.remote	7	-	62,533
0.00	phase_0.idx1.dispatch.evaluate.timechart	1	-	-
0.00	phase_0.idx1.dispatch.parserThread	8	-	-
0.08	phase_0.idx1.dispatch.stream.remote.idx1	7	-	62,533
0.46	phase_0.idx1.startup.handoff	3	-	-
0.10	phase_0.idx2.command.search	7	-	175
0.04	phase_0.idx1.startup.configuration	3	-	-
0.00	phase_0.idx2.command.addinfo	7	175	175

.conf20 splunk>

Search Logs

Log, I am your father

What’s In Search Logs?

| rest services/search/jobs/<sid>/search.log

1) **ERROR/WARN logs**

   Find all ERROR and WARN messages, even ones not displayed in the top level of the job inspector

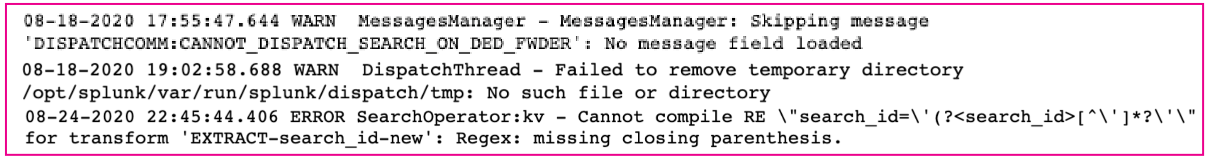

```
08-18-2020 17:55:47.644 WARN  MessagesManager - MessagesManager: Skipping message
'DISPATCHCOMM:CANNOT_DISPATCH_SEARCH_ON_DED_FWDER': No message field loaded
08-18-2020 19:02:58.688 WARN  DispatchThread - Failed to remove temporary directory
/opt/splunk/var/run/splunk/dispatch/tmp: No such file or directory
08-24-2020 22:45:44.406 ERROR SearchOperator:kv - Cannot compile RE \"search_id=\'(?<search_id>[^\']*?\'\"
for transform 'EXTRACT-search_id-new': Regex: missing closing parenthesis.
```

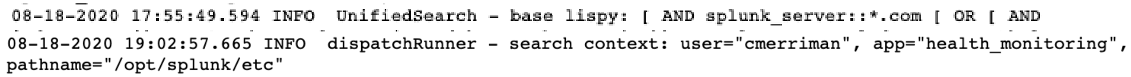

```
08-18-2020 17:55:49.594 INFO  UnifiedSearch - base lispy: [ AND splunk_server::*.com [ OR [ AND
08-18-2020 19:02:57.665 INFO  dispatchRunner - search context: user="cmerriman", app="health_monitoring",
pathname="/opt/splunk/etc"
```

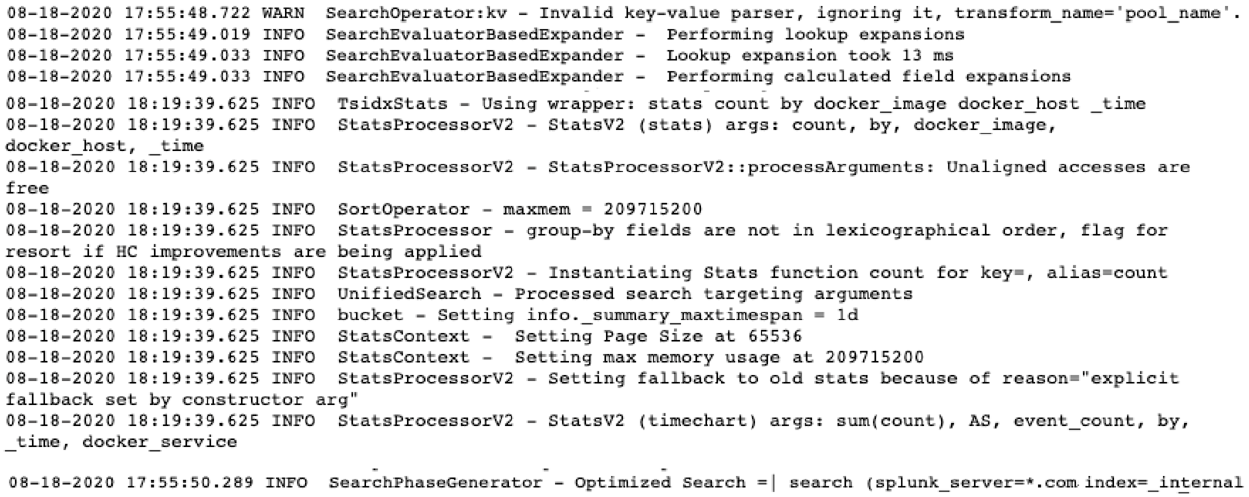

```
08-18-2020 17:55:48.722 WARN  SearchOperator:kv - Invalid key-value parser, ignoring it, transform_name='pool_name'.
08-18-2020 17:55:49.019 INFO  SearchEvaluatorBasedExpander -  Performing lookup expansions
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Lookup expansion took 13 ms
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Performing calculated field expansions
08-18-2020 18:19:39.625 INFO  TsidxStats - Using wrapper: stats count by docker_image docker_host _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (stats) args: count, by, docker_image,
docker_host, _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsProcessorV2::processArguments: Unaligned accesses are
free
08-18-2020 18:19:39.625 INFO  SortOperator - maxmem = 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessor - group-by fields are not in lexicographical order, flag for
resort if HC improvements are being applied
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Instantiating Stats function count for key=, alias=count
08-18-2020 18:19:39.625 INFO  UnifiedSearch - Processed search targeting arguments
08-18-2020 18:19:39.625 INFO  bucket - Setting info._summary_maxtimespan = 1d
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting Page Size at 65536
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting max memory usage at 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Setting fallback to old stats because of reason="explicit
fallback set by constructor arg"
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (timechart) args: sum(count), AS, event_count, by,
_time, docker_service

08-18-2020 17:55:50.289 INFO  SearchPhaseGenerator - Optimized Search =| search (splunk_server=*.com index=_internal
```

What’s In Search Logs?

| rest services/search/jobs/<sid>/search.log

1) **ERROR/WARN logs**
   Find all ERROR and WARN messages, even ones not displayed in the top level of the job inspector
2) **Search Context**
   dispatchRunner, UnifiedSearch, SearchParser

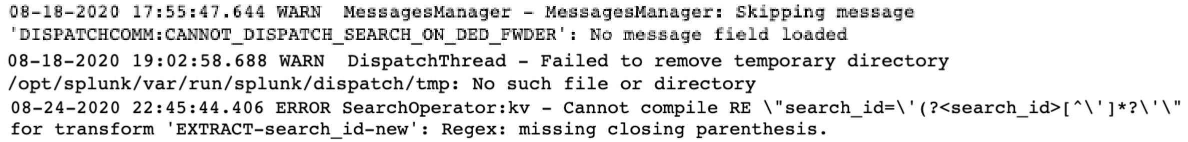

```
08-18-2020 17:55:47.644 WARN  MessagesManager - MessagesManager: Skipping message
'DISPATCHCOMM:CANNOT_DISPATCH_SEARCH_ON_DED_FWDER': No message field loaded
08-18-2020 19:02:58.688 WARN  DispatchThread - Failed to remove temporary directory
/opt/splunk/var/run/splunk/dispatch/tmp: No such file or directory
08-24-2020 22:45:44.406 ERROR SearchOperator:kv - Cannot compile RE \"search_id=\'(?<search_id>[^\']*?\'\"
for transform 'EXTRACT-search_id-new': Regex: missing closing parenthesis.
```

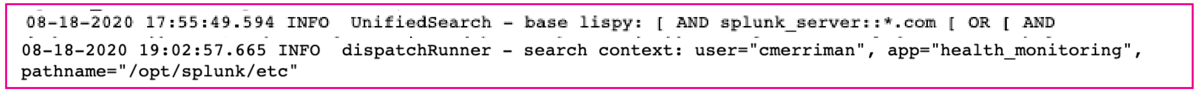

```
08-18-2020 17:55:49.594 INFO  UnifiedSearch - base lispy: [ AND splunk_server::*.com [ OR [ AND
08-18-2020 19:02:57.665 INFO  dispatchRunner - search context: user="cmerriman", app="health_monitoring",
pathname="/opt/splunk/etc"
```

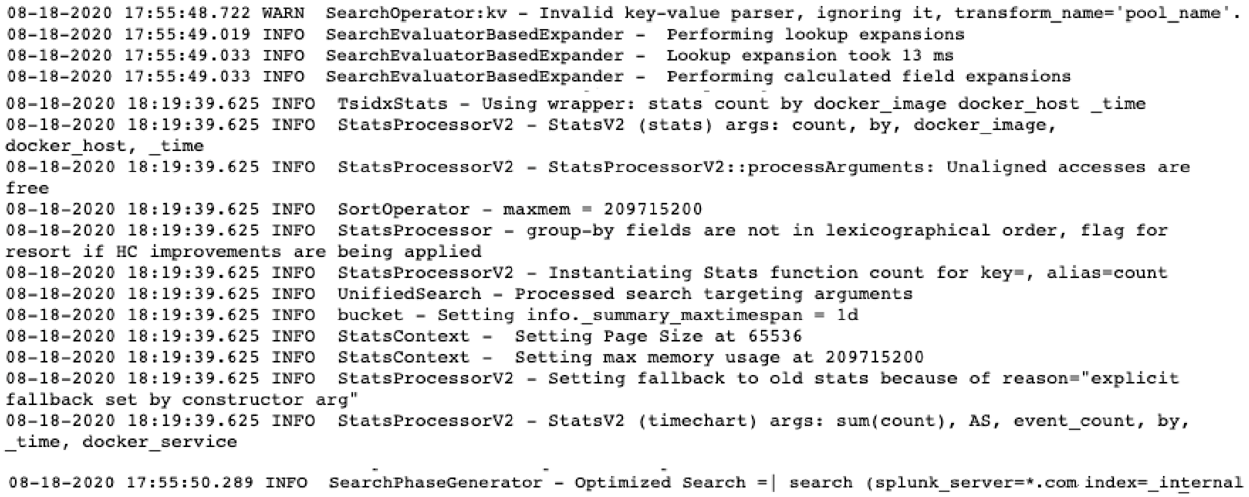

```
08-18-2020 17:55:48.722 WARN  SearchOperator:kv - Invalid key-value parser, ignoring it, transform_name='pool_name'.
08-18-2020 17:55:49.019 INFO  SearchEvaluatorBasedExpander -  Performing lookup expansions
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Lookup expansion took 13 ms
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Performing calculated field expansions
08-18-2020 18:19:39.625 INFO  TsidxStats - Using wrapper: stats count by docker_image docker_host _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (stats) args: count, by, docker_image,
docker_host, _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsProcessorV2::processArguments: Unaligned accesses are
free
08-18-2020 18:19:39.625 INFO  SortOperator - maxmem = 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessor - group-by fields are not in lexicographical order, flag for
resort if HC improvements are being applied
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Instantiating Stats function count for key=, alias=count
08-18-2020 18:19:39.625 INFO  UnifiedSearch - Processed search targeting arguments
08-18-2020 18:19:39.625 INFO  bucket - Setting info._summary_maxtimespan = 1d
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting Page Size at 65536
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting max memory usage at 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Setting fallback to old stats because of reason="explicit
fallback set by constructor arg"
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (timechart) args: sum(count), AS, event_count, by,
_time, docker_service

08-18-2020 17:55:50.289 INFO  SearchPhaseGenerator - Optimized Search =| search (splunk_server=*.com index=_internal
```

What’s In Search Logs?

| rest services/search/jobs/<sid>/search.log

1) **ERROR/WARN logs**
   Find all ERROR and WARN messages, even ones not displayed in the top level of the job inspector
2) **Search Context**
   dispatchRunner, UnifiedSearch, SearchParser
3) **Lookups**
   CsvDataProvider, SearchOperator:kv, SearchEvaluatorBasedExpander

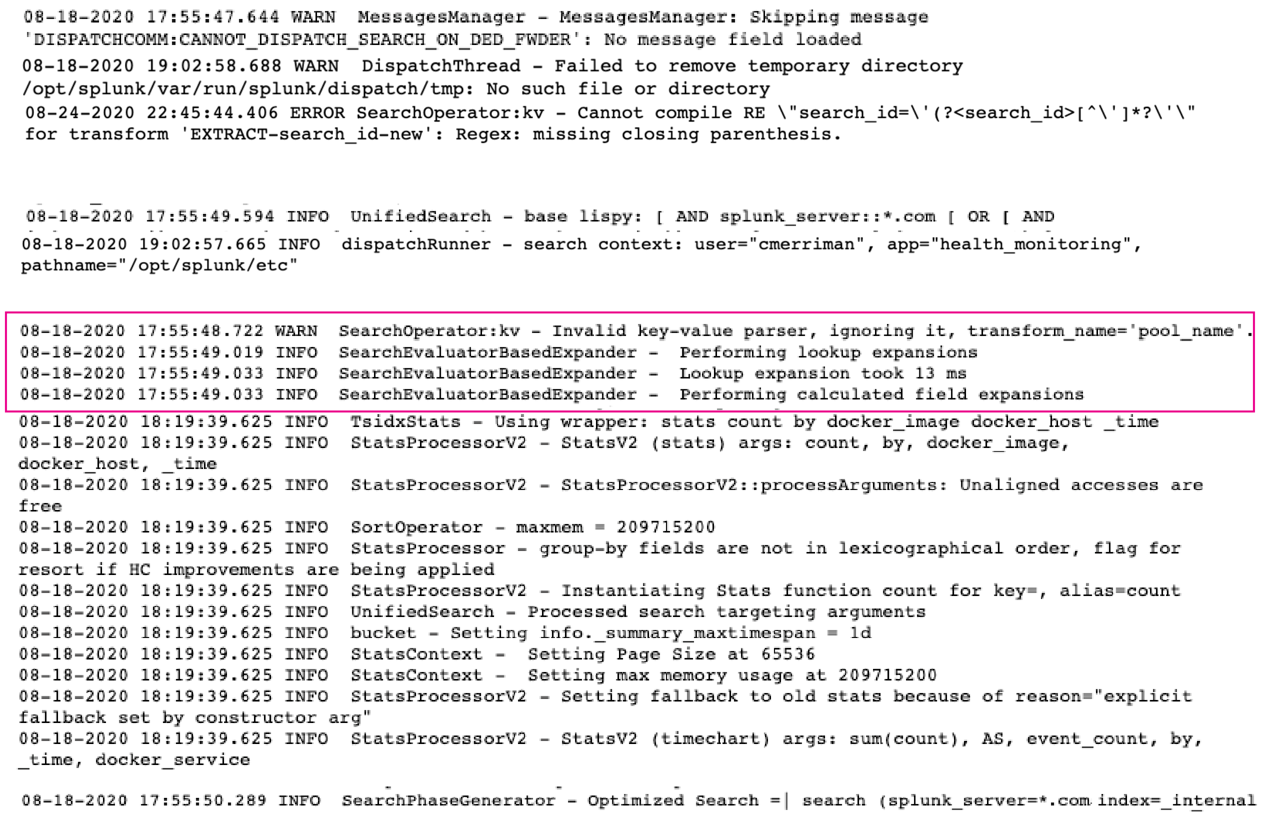

```
08-18-2020 17:55:47.644 WARN  MessagesManager - MessagesManager: Skipping message
'DISPATCHCOMM:CANNOT_DISPATCH_SEARCH_ON_DED_FWDER': No message field loaded
08-18-2020 19:02:58.688 WARN  DispatchThread - Failed to remove temporary directory
/opt/splunk/var/run/splunk/dispatch/tmp: No such file or directory
08-24-2020 22:45:44.406 ERROR SearchOperator:kv - Cannot compile RE \"search_id=\'(?<search_id>[^\']*?\'\"
for transform 'EXTRACT-search_id-new': Regex: missing closing parenthesis.
```

```
08-18-2020 17:55:49.594 INFO  UnifiedSearch - base lispy: [ AND splunk_server::*.com [ OR [ AND
08-18-2020 19:02:57.665 INFO  dispatchRunner - search context: user="cmerriman", app="health_monitoring",
pathname="/opt/splunk/etc"
```

```
08-18-2020 17:55:48.722 WARN  SearchOperator:kv - Invalid key-value parser, ignoring it, transform_name='pool_name'.
08-18-2020 17:55:49.019 INFO  SearchEvaluatorBasedExpander -  Performing lookup expansions
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Lookup expansion took 13 ms
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Performing calculated field expansions
08-18-2020 18:19:39.625 INFO  TsidxStats - Using wrapper: stats count by docker_image docker_host _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (stats) args: count, by, docker_image,
docker_host, _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsProcessorV2::processArguments: Unaligned accesses are
free
08-18-2020 18:19:39.625 INFO  SortOperator - maxmem = 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessor - group-by fields are not in lexicographical order, flag for
resort if HC improvements are being applied
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Instantiating Stats function count for key=, alias=count
08-18-2020 18:19:39.625 INFO  UnifiedSearch - Processed search targeting arguments
08-18-2020 18:19:39.625 INFO  bucket - Setting info._summary_maxtimespan = 1d
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting Page Size at 65536
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting max memory usage at 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Setting fallback to old stats because of reason="explicit
fallback set by constructor arg"
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (timechart) args: sum(count), AS, event_count, by,
_time, docker_service
```

```
08-18-2020 17:55:50.289 INFO  SearchPhaseGenerator - Optimized Search =| search (splunk_server=*.com index=_internal
```

What’s In Search Logs?

| rest services/search/jobs/<sid>/search.log

1) **ERROR/WARN logs**
   Find all ERROR and WARN messages, even ones not displayed in the top level of the job inspector
2) **Search Context**
   dispatchRunner, UnifiedSearch, SearchParser
3) **Lookups**
   CsvDataProvider, SearchOperator:kv, SearchEvaluatorBasedExpander
4) **Search Processes**
   StatsContext, StatsProcessor[V2], SortOperator, TsidxStats

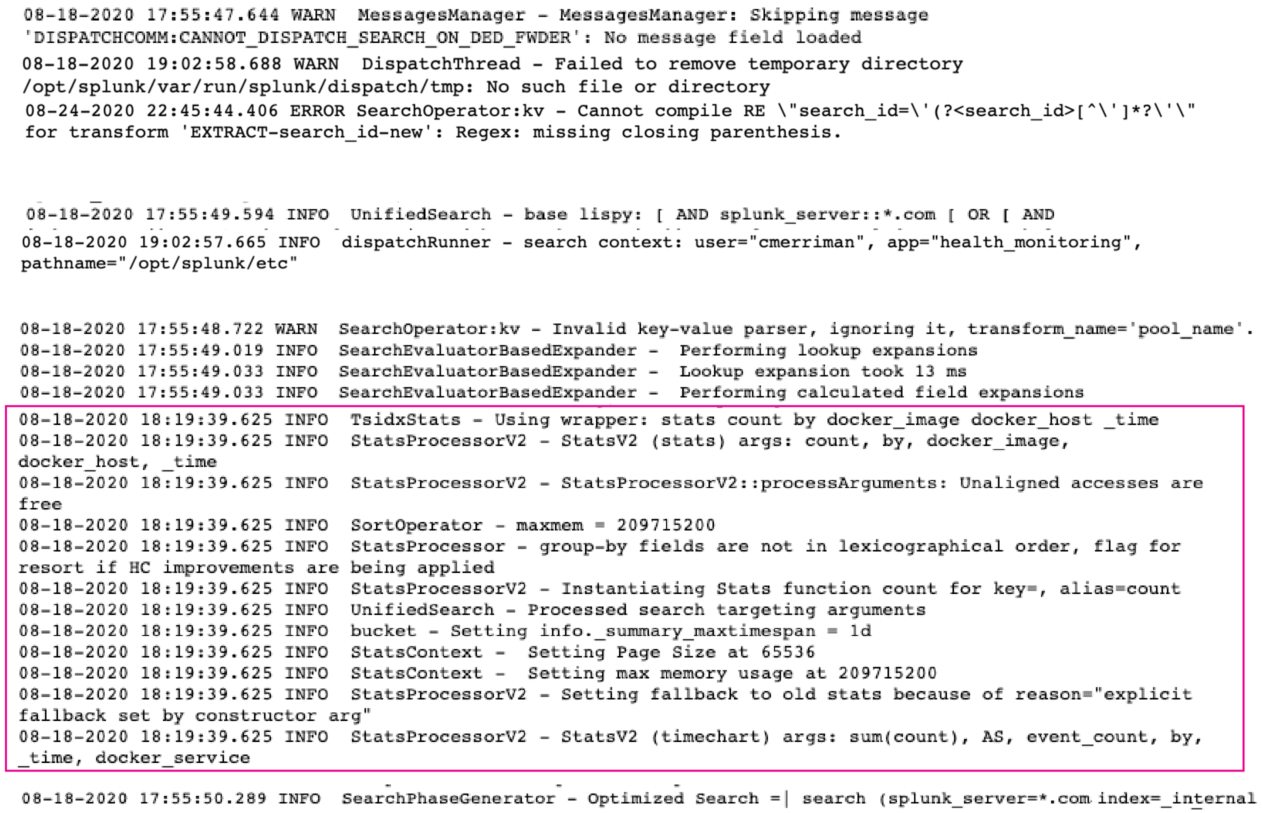

```
08-18-2020 17:55:47.644 WARN  MessagesManager - MessagesManager: Skipping message
'DISPATCHCOMM:CANNOT_DISPATCH_SEARCH_ON_DED_FWDER': No message field loaded
08-18-2020 19:02:58.688 WARN  DispatchThread - Failed to remove temporary directory
/opt/splunk/var/run/splunk/dispatch/tmp: No such file or directory
08-24-2020 22:45:44.406 ERROR SearchOperator:kv - Cannot compile RE \"search_id=\'(?<search_id>[^\']*?\'\"
for transform 'EXTRACT-search_id-new': Regex: missing closing parenthesis.
```

```
08-18-2020 17:55:49.594 INFO  UnifiedSearch - base lispy: [ AND splunk_server::*.com [ OR [ AND
08-18-2020 19:02:57.665 INFO  dispatchRunner - search context: user="cmerriman", app="health_monitoring",
pathname="/opt/splunk/etc"
```

```
08-18-2020 17:55:48.722 WARN  SearchOperator:kv - Invalid key-value parser, ignoring it, transform_name='pool_name'.
08-18-2020 17:55:49.019 INFO  SearchEvaluatorBasedExpander -  Performing lookup expansions
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Lookup expansion took 13 ms
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Performing calculated field expansions
```

```
08-18-2020 18:19:39.625 INFO  TsidxStats - Using wrapper: stats count by docker_image docker_host _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (stats) args: count, by, docker_image,
docker_host, _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsProcessorV2::processArguments: Unaligned accesses are
free
08-18-2020 18:19:39.625 INFO  SortOperator - maxmem = 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessor - group-by fields are not in lexicographical order, flag for
resort if HC improvements are being applied
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Instantiating Stats function count for key=, alias=count
08-18-2020 18:19:39.625 INFO  UnifiedSearch - Processed search targeting arguments
08-18-2020 18:19:39.625 INFO  bucket - Setting info._summary_maxtimespan = 1d
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting Page Size at 65536
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting max memory usage at 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Setting fallback to old stats because of reason="explicit
fallback set by constructor arg"
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (timechart) args: sum(count), AS, event_count, by,
_time, docker_service
```

```
08-18-2020 17:55:50.289 INFO  SearchPhaseGenerator - Optimized Search =| search (splunk_server=*.com index=_internal
```

What’s In Search Logs?

| rest services/search/jobs/<sid>/search.log

1) **ERROR/WARN logs**
   Find all ERROR and WARN messages, even ones not displayed in the top level of the job inspector
2) **Search Context**
   dispatchRunner, UnifiedSearch, SearchParser
3) **Lookups**
   CsvDataProvider, SearchOperator:kv, SearchEvaluatorBasedExpander
4) **Search Processes**
   StatsContext, StatsProcessor[V2], SortOperator, TsidxStats
5) **Search Optimization**
   AstOptimizer, SearchPhaseGenerator

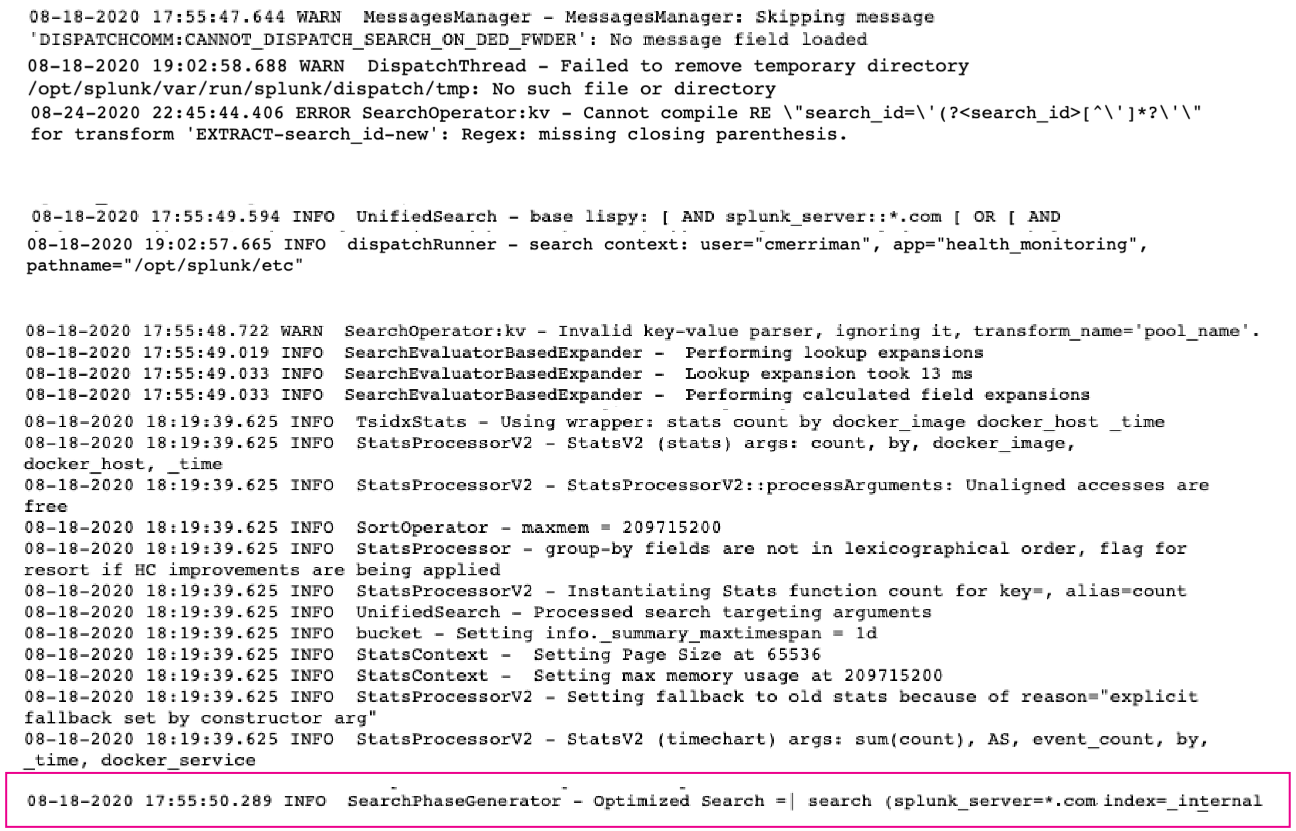

```
08-18-2020 17:55:47.644 WARN  MessagesManager - MessagesManager: Skipping message
'DISPATCHCOMM:CANNOT_DISPATCH_SEARCH_ON_DED_FWDER': No message field loaded
08-18-2020 19:02:58.688 WARN  DispatchThread - Failed to remove temporary directory
/opt/splunk/var/run/splunk/dispatch/tmp: No such file or directory
08-24-2020 22:45:44.406 ERROR SearchOperator:kv - Cannot compile RE \"search_id=\'(?<search_id>[^\']*?\'\"
for transform 'EXTRACT-search_id-new': Regex: missing closing parenthesis.
```

```
08-18-2020 17:55:49.594 INFO  UnifiedSearch - base lispy: [ AND splunk_server::*.com [ OR [ AND
08-18-2020 19:02:57.665 INFO  dispatchRunner - search context: user="cmerriman", app="health_monitoring",
pathname="/opt/splunk/etc"
```

```
08-18-2020 17:55:48.722 WARN  SearchOperator:kv - Invalid key-value parser, ignoring it, transform_name='pool_name'.
08-18-2020 17:55:49.019 INFO  SearchEvaluatorBasedExpander -  Performing lookup expansions
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Lookup expansion took 13 ms
08-18-2020 17:55:49.033 INFO  SearchEvaluatorBasedExpander -  Performing calculated field expansions
08-18-2020 18:19:39.625 INFO  TsidxStats - Using wrapper: stats count by docker_image docker_host _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (stats) args: count, by, docker_image,
docker_host, _time
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsProcessorV2::processArguments: Unaligned accesses are
free
08-18-2020 18:19:39.625 INFO  SortOperator - maxmem = 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessor - group-by fields are not in lexicographical order, flag for
resort if HC improvements are being applied
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Instantiating Stats function count for key=, alias=count
08-18-2020 18:19:39.625 INFO  UnifiedSearch - Processed search targeting arguments
08-18-2020 18:19:39.625 INFO  bucket - Setting info._summary_maxtimespan = 1d
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting Page Size at 65536
08-18-2020 18:19:39.625 INFO  StatsContext -  Setting max memory usage at 209715200
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - Setting fallback to old stats because of reason="explicit
fallback set by constructor arg"
08-18-2020 18:19:39.625 INFO  StatsProcessorV2 - StatsV2 (timechart) args: sum(count), AS, event_count, by,
_time, docker_service
```

```
08-18-2020 17:55:50.289 INFO  SearchPhaseGenerator - Optimized Search =| search (splunk_server=*.com index=_internal
```

What’s Next?

latest = +1d

1) Go check out the job inspector

Find one of the longest-running jobs in your environment and see what the job inspector tells you. Tweak the search a little and see what changes in the job inspector. Purposely hit a limit (that only affects you/that search, not the system and other users, something like a join limit) or syntax error to see what messages are created.

2) Have more questions?

Check out the resources on the next slide

Read the docs

Reach out on Answers

Reach out on Slack

Reach out to us

Ask now ☺

Resources

| search splunkdocs = true

Job Inspector Properties Docs
https://docs.splunk.com/Documentation/Splunk/latest/Search/ViewsearchjobpropertieswiththeJobInspector

REST Endpoint for developers/poking around
https://docs.splunk.com/Documentation/Splunk/latest/RESTREF/RESTsearch#search.2Fjobs

Splunk Optimizations Docs
https://docs.splunk.com/Documentation/Splunk/latest/Search/Built-inoptimization#Optimization_settings

Martin’s B-Sides Job Inspector talk
https://www.youtube.com/watch?v=1QCZ5klSptM

Martin’s Optimizing Knowledge Object’s talk
https://conf.splunk.com/session/2015/conf2015_MMueller_Consist_Deploying_OptimizingSplunkKnowledge.pdf / http://conf.splunk.com/session/2015/recordings/2015-splunk-134.mp4

Martin’s Lispy talk
https://conf.splunk.com/files/2019/slides/FN1003.pdf / https://conf.splunk.com/files/2019/summit/FN1003.mp4

Splunk Math: How Users Learn About Job Inspector

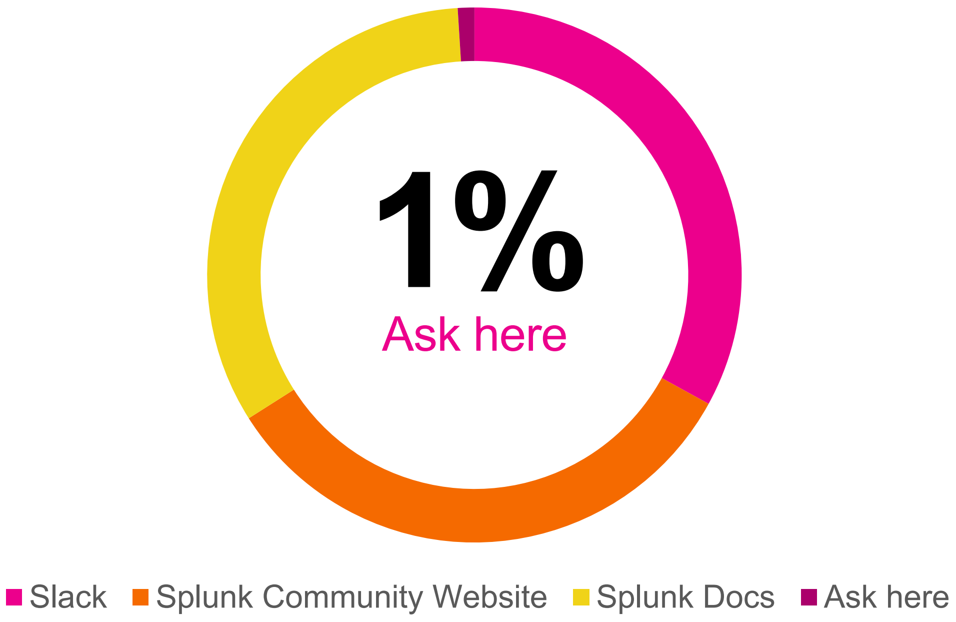

splunk> .conf20

Thank You

Please provide feedback via the

SESSION SURVEY