How To Search Fast Using First Principles

Steven Czerwinski | Co-founder & CTO | Scalyr

What Is Scalyr

- Operational Visibility Tool
- SaaS
- Performance @ scale

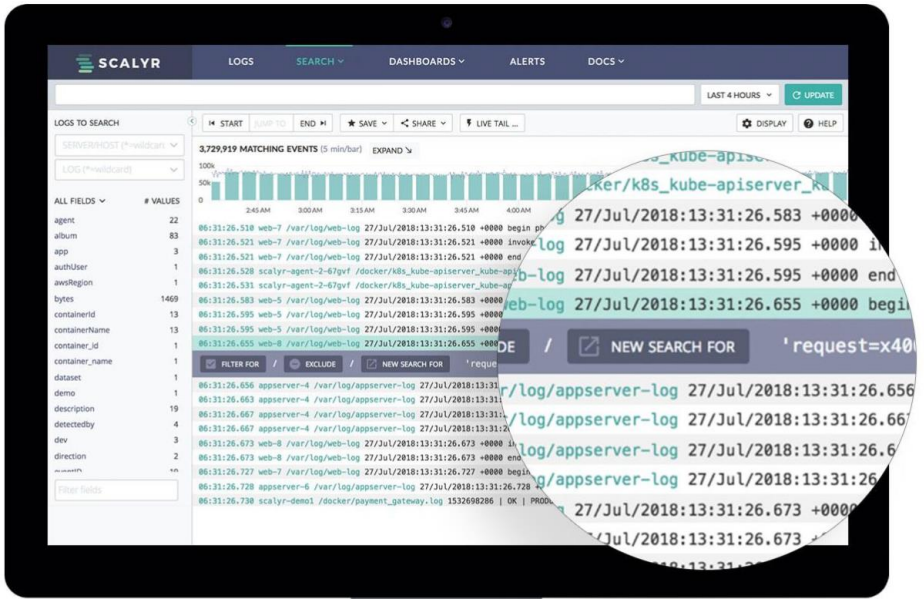

What is a log?

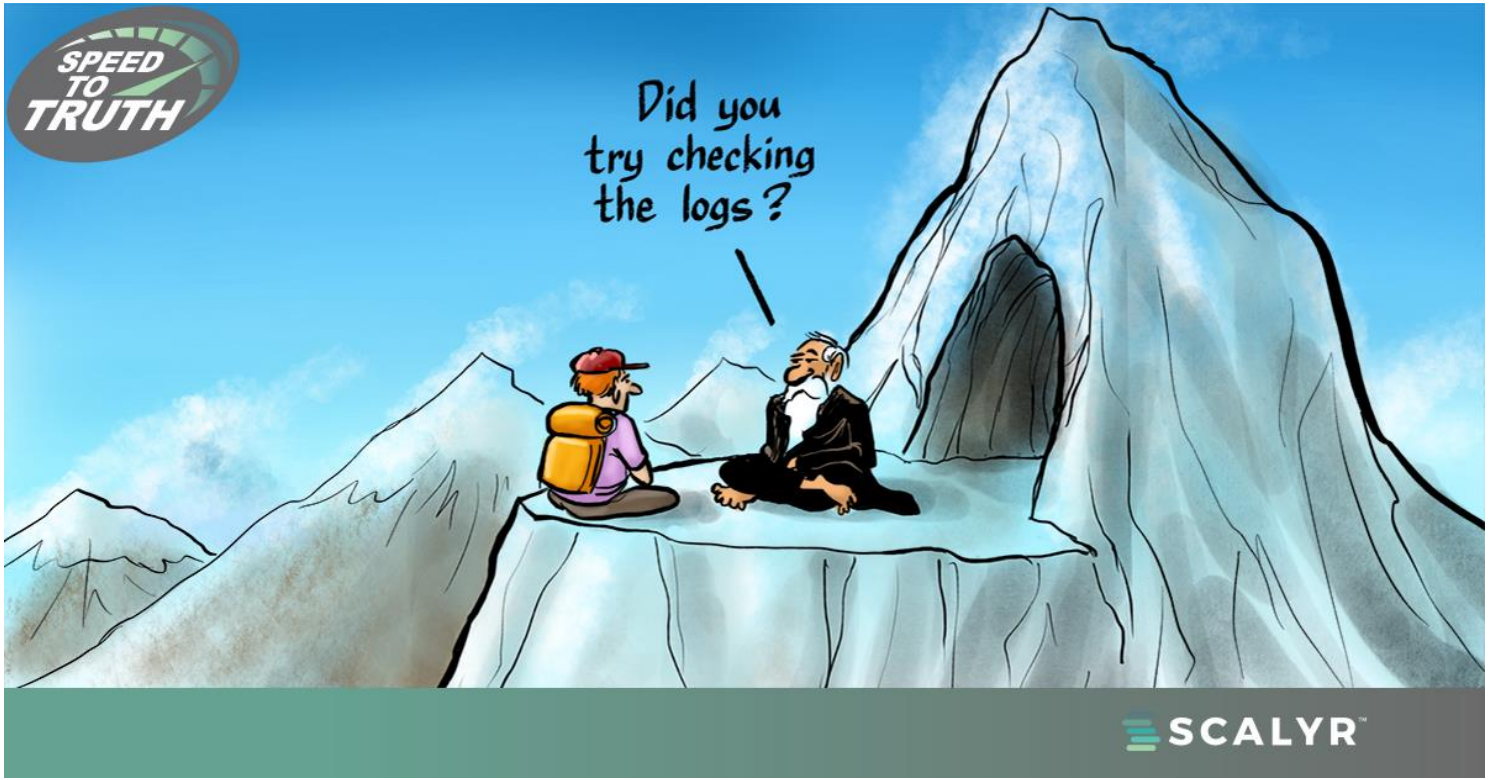

Log Sample

```
[17/Dec/2005:02:40:45 -0500] - - 85.226.238.xxx "GET /awstats/awstats.pl?configdir=|echo;echo%2
[17/Dec/2005:02:40:46 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats.pl?configdir=|echo;echo%2
[17/Dec/2005:02:40:47 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats/awstats.pl?configdir=|ech
[17/Dec/2005:02:40:48 -0500] - - 85.226.238.xxx "GET /index2.php?option=com_content&do_pdf=1&id
[17/Dec/2005:02:40:49 -0500] - - 85.226.238.xxx "GET /index.php?option=com_content&do_pdf=1&id=
[17/Dec/2005:02:40:50 -0500] - - 85.226.238.xxx "GET /mambo/index2.php?_REQUEST[option]=com_con
[17/Dec/2005:02:40:52 -0500] - - 85.226.238.xxx "GET /cvs/index2.php?_REQUEST[option]=com_conte
[17/Dec/2005:02:40:53 -0500] - - 85.226.238.xxx "GET /cvs/mambo/index2.php?_REQUEST[option]=con
[17/Dec/2005:03:07:04 -0500] - - 24.19.40.xxx "GET /awstats/awstats.pl?configdir=|echo;echo%20Y
[17/Dec/2005:04:34:04 -0500] - - 64.105.82.xxx "GET /blogtest/xmlsrv/xmlrpc.php HTTP/1.1" "-" 1
[17/Dec/2005:05:18:40 -0500] - - 195.82.6.xxx "GET /modules/Forums/admin/admin_styles.phpadmin_
[17/Dec/2005:05:18:41 -0500] - - 195.82.6.xxx "GET /Forums/admin/admin_styles.phpadmin_styles.p
```

Log Sample (2)

```
Dec 16 18:02:04 asterisk1 asterisk[31774]: NOTICE[31787]: chan_sip.c:11242 in
handle_request_register: Registration from '"503"<sip:503@192.168.1.107>' failed for
'192.168.1.137' - Wrong password

Dec 16 18:03:13 asterisk1 asterisk[31774]: NOTICE[31787]: chan_sip.c:11242 in
handle_request_register: Registration from '"502"<sip:502@192.168.1.107>' failed for
'192.168.1.137' - Wrong password

Dec 16 18:04:49 asterisk1 asterisk[31774]: NOTICE[31787]: chan_sip.c:11242 in
handle_request_register: Registration from '"1737245082"<sip:1737245082@192.168.1.107>'
failed for '192.168.1.137' - Username/auth name mismatch

Dec 16 18:04:49 asterisk1 asterisk[31774]: NOTICE[31787]: chan_sip.c:11242 in
handle_request_register: Registration from '"100"<sip:100@192.168.1.107>' failed for
'192.168.1.137' - Username/auth name mismatch

Jun 27 18:09:47 host asterisk[31774]: ERROR[27910]: chan_zap.c:10314 setup_zap: Unable to
register channel '1-2'
```

What our users want...

Keyword Indexing

```
(1) [17/Dec/2005:02:40:45 -0500] - - 85.226.238.xxx "GET /awstats/awstats.pl?configdir=|echo;e
(2) [17/Dec/2005:02:40:46 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats.pl?configdir=|echo;e
(3) [17/Dec/2005:02:40:47 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats/awstats.pl?configdi
(4) [17/Dec/2005:02:40:48 -0500] - - 85.226.238.xxx "GET /index2.php?option=com_content&do_pdf
(5) [17/Dec/2005:02:40:49 -0500] - - 85.226.238.xxx "GET /index.php?option=com_content&do_pdf=
(6) [17/Dec/2005:02:40:50 -0500] - - 85.226.238.xxx "GET /mambo/index2.php?_REQUEST[option]=co
(7) [17/Dec/2005:02:40:52 -0500] - - 85.226.238.xxx "GET /cvs/index2.php?_REQUEST[option]=com_
(8) [17/Dec/2005:02:40:53 -0500] - - 85.226.238.xxx "GET /cvs/mambo/index2.php?_REQUEST[optior
(9) [17/Dec/2005:03:07:04 -0500] - - 24.19.40.xxx "GET /awstats/awstats.pl?configdir=|echo;ech
(10)[17/Dec/2005:04:34:04 -0500] - - 64.105.82.xxx "GET /blogtest/xmlsrv/xmlrpc.php HTTP/1.1"
(11)[17/Dec/2005:05:18:40 -0500] - - 195.82.6.xxx "GET /modules/Forums/admin/admin_styles.php
(12)[17/Dec/2005:05:18:41 -0500] - - 195.82.6.xxx "GET /Forums/admin/admin_styles.phpadmin_st
```

Keyword Indexing

```
(1) [17/Dec/2005:02:40:45 -0500] - - 85.226.238.xxx "GET /awstats/awstats.pl?configdir=|echo;e
(2) [17/Dec/2005:02:40:46 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats.pl?configdir=|echo;e
(3) [17/Dec/2005:02:40:47 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats/awstats.pl?configdi
(4) [17/Dec/2005:02:40:48 -0500] - - 85.226.238.xxx "GET /index2.php?option=com_content&do_pdf
(5) [17/Dec/2005:02:40:49 -0500] - - 85.226.238.xxx "GET /index.php?option=com_content&do_pdf=
(6) [17/Dec/2005:02:40:50 -0500] - - 85.226.238.xxx "GET /mambo/index2.php?_REQUEST[option]=co
(7) [17/Dec/2005:02:40:52 -0500] - - 85.226.238.xxx "GET /cvs/index2.php?_REQUEST[option]=com_
(8) [17/Dec/2005:02:40:53 -0500] - - 85.226.238.xxx "GET /cvs/mambo/index2.php?_REQUEST[optior
(9) [17/Dec/2005:03:07:04 -0500] - - 24.19.40.xxx "GET /awstats/awstats.pl?configdir=|echo;ech
(10) [17/Dec/2005:04:34:04 -0500] - - 64.105.82.xxx "GET /blogtest/xmlsrv/xmlrpc.php HTTP/1.1"
(11) [17/Dec/2005:05:18:40 -0500] - - 195.82.6.xxx "GET /modules/Forums/admin/admin_styles.php
(12) [17/Dec/2005:05:18:41 -0500] - - 195.82.6.xxx "GET /Forums/admin/admin_styles.phpadmin_st
```

awstats: (1), (2), (3), (9)

Keyword Indexing

```
(1) [17/Dec/2005:02:40:45 -0500] - - 85.226.238.xxx "GET /awstats/awstats.pl?configdir=|echo;e
(2) [17/Dec/2005:02:40:46 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats.pl?configdir=|echo;e
(3) [17/Dec/2005:02:40:47 -0500] - - 85.226.238.xxx "GET /cgi-bin/awstats/awstats.pl?configdi
(4) [17/Dec/2005:02:40:48 -0500] - - 85.226.238.xxx "GET /index2.php?option=com_content&do_pdf
(5) [17/Dec/2005:02:40:49 -0500] - - 85.226.238.xxx "GET /index.php?option=com_content&do_pdf=
(6) [17/Dec/2005:02:40:50 -0500] - - 85.226.238.xxx "GET /mambo/index2.php?_REQUEST[option]=co
(7) [17/Dec/2005:02:40:52 -0500] - - 85.226.238.xxx "GET /cvs/index2.php?_REQUEST[option]=com_
(8) [17/Dec/2005:02:40:53 -0500] - - 85.226.238.xxx "GET /cvs/mambo/index2.php?_REQUEST[option
(9) [17/Dec/2005:03:07:04 -0500] - - 24.19.40.xxx "GET /awstats/awstats.pl?configdir=|echo;ech
(10) [17/Dec/2005:04:34:04 -0500] - - 64.105.82.xxx "GET /blogtest/xmlsrv/xmlrpc.php HTTP/1.1"
(11) [17/Dec/2005:05:18:40 -0500] - - 195.82.6.xxx "GET /modules/Forums/admin/admin_styles.php
(12) [17/Dec/2005:05:18:41 -0500] - - 195.82.6.xxx "GET /Forums/admin/admin_styles.phpadmin_st
```

awstats: (1), (2), (3), (9)

option: (4), (5), (6), (7), (8)

Keyword Index Challenges

- Expensive to build + maintain
  - Especially for machine data (huge vocabulary)
- Not good at sub-word “wildcard” matches (`*error*`)
- ...or regular expressions
- ...or numeric ranges (`latency > 1000`)
- ...or summarizing billions of matches

What About Grep?

Time For Math

Goal: search 1 day's logs in 1 second

- Competitor's pricing: $100 / daily GB / month
- ~= 1 CPU core in EC2
- ~= 2 GB/sec search performance
- → 2x better than we need!

Time For Math

Goal: search 1 day's logs in 1 second

- Competitor's pricing: $100 / daily GB / month
- ~= 1 CPU core in EC2
- ~= 2 GB/sec search performance
- → 2x better than we need!

BUT: other costs, profit margin, etc.

Column Store Reduces Scan Requirements

173.186.97.118
-- --
[15/Feb/2017.21
:54:06 +0000]
“POST /home
HTTP 1.1” 200
…

152.169.219.124
-- --
[15/Feb/2017.21
:54:06 +0000]
“POST /home
HTTP 1.1” 200
…

33.201.177.40
-- --
[15/Feb/2017.21
:54:07 +0000]
“POST /friends
HTTP 1.1” 200
…

173.186.97.118	15/Feb/2017:21:54:06 +0000	POST	/home	HTTP 1.1	200	13508	Mozilla/5.0
152.169.219.124	15/Feb/2017:21:54:06 +0000	POST	/home	HTTP 1.1	200	34815	Mozilla/5.0
33.201.177.40	15/Feb/2017:21:54:07 +0000	POST	/friends	HTTP 1.1	200	8517	Mozilla/5.0
249.93.144.201	15/Feb/2017:21:54:07 +0000	POST	/album	HTTP 1.1	200	10053	Mozilla/5.0
189.167.124.186	15/Feb/2017:21:54:07 +0000	POST	/photo	HTTP 1.1	200	3323	Mozilla/5.0
245.235.220.237	15/Feb/2017:21:54:08 +0000	POST	/profile?user=u1609726	HTTP 1.1	200	33284	Mozilla/5.0
39.54.222.112	15/Feb/2017:21:54:08 +0000	POST	/photo	HTTP 1.1	200	2723	Mozilla/5.0
90.94.230.51	15/Feb/2017:21:54:08 +0000	POST	/profile?user=p422479&width=320&height=320	HTTP 1.1	200	3642	Mozilla/5.0

Key Architectural Decisions

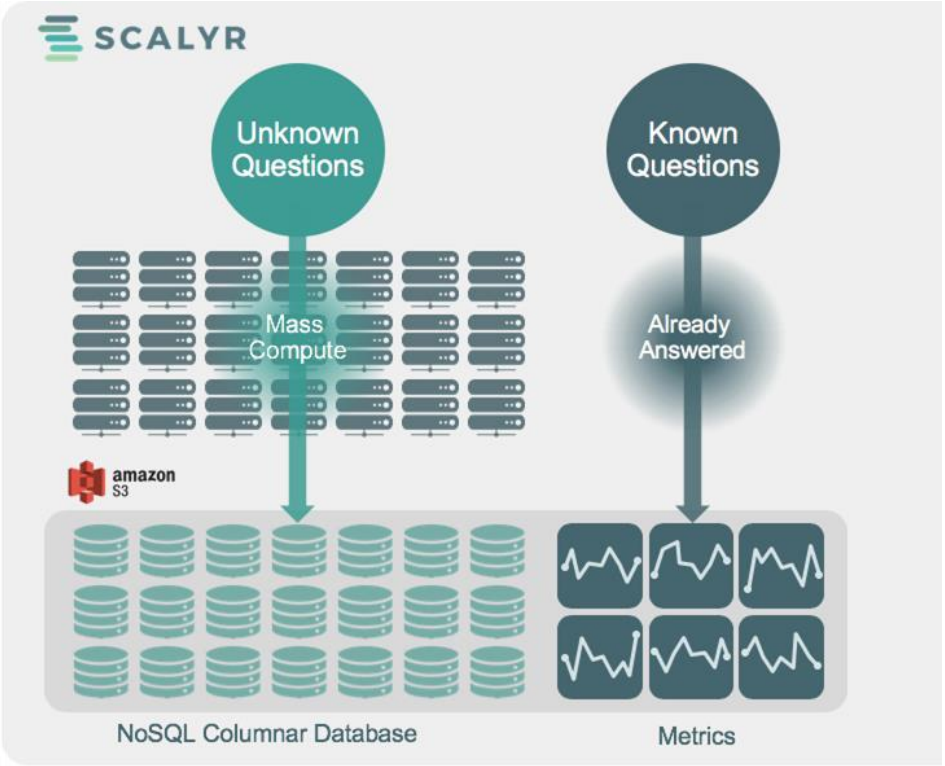

Time For Some More Math

Goal: support searching 1TB of daily logs using Lambda

- Competitor's pricing: $100K / daily TB / month
- ~= 2.1B Lambda seconds / month (Lambda invoke w/ 3GB RAM)
- ~= 46K Lambda seconds / search (100 users x 15 searches / day)
- ~= 9TB / search

  (200MB scanned / Lambda sec)
- → 9x better than we need!

Time For Some More Math

Goal: support searching 1TB of daily logs using Lambda

- Competitor's pricing: $100K / daily TB / month
- ~= 2.1B Lambda seconds / month (Lambda invoke w/ 3GB RAM)
- ~= 46K Lambda seconds / search (100 users x 15 searches / day)
- ~= 9TB / search

  (200MB scanned / Lambda sec)
- → 9x better than we need!

BUT: other costs, discounts, profit margin, yada, yada

BTW, Old Fashioned Coding Still Matters

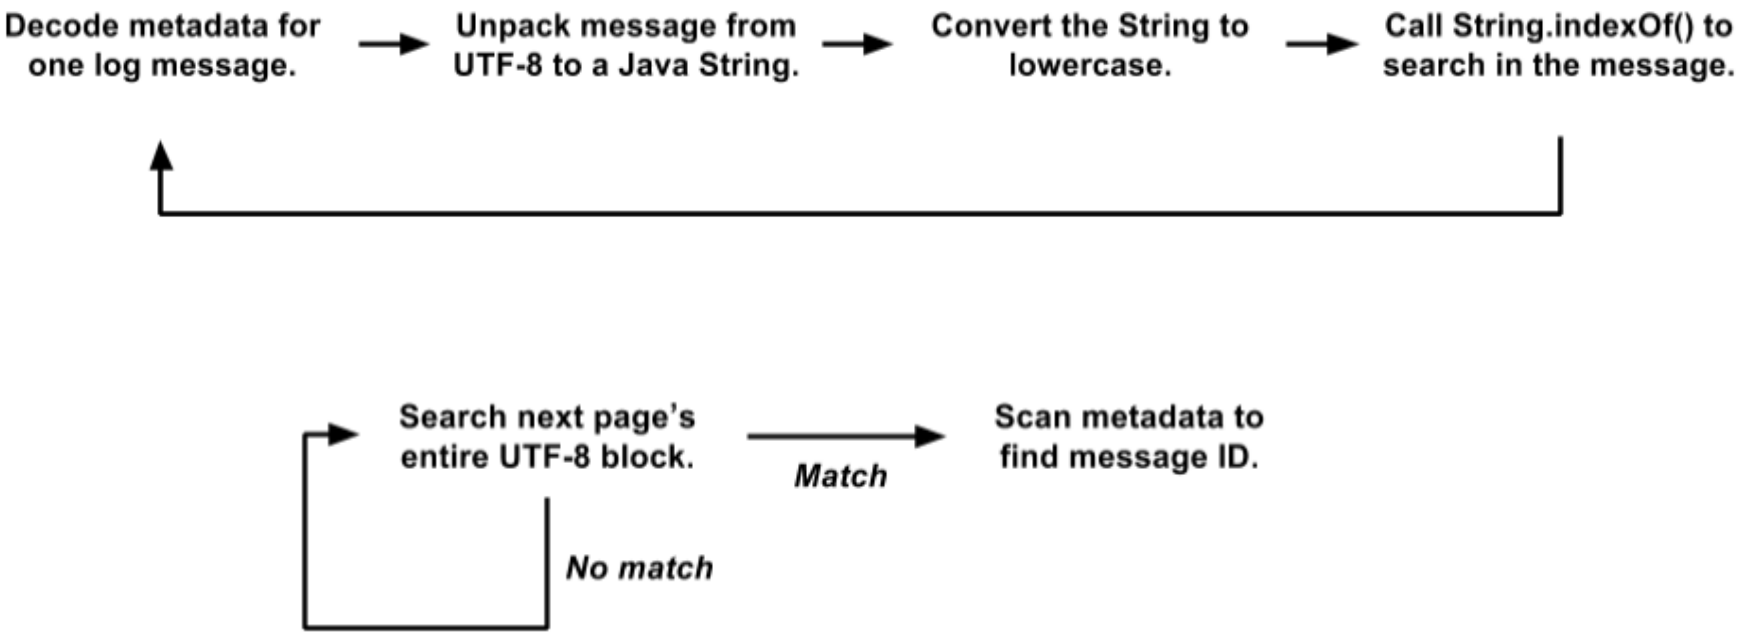

Lessons Learned

- Always good to analyze from first principles
- Anchor your analysis in a user / business goal
- Do be afraid to buck conventional wisdom
  - *But sometimes do it anyway*
- Systems engineering can be more important than fancy algorithms
  - *Constant factors matter*

Q + A