


Distributed System Performance Troubleshooting Like You've Been Doing it for Twenty Years

Jon Haddad
Rustyrazorblade Consulting
rustyrazorblade.com



About me

- **20 Year Career**
- **30 Years of SW Development**
- **Apache Cassandra Committer**
- **Ex-Netflix**
- **Ex-Apple**
- **Ex-Lots of other stuff**
- **Now on my own!**








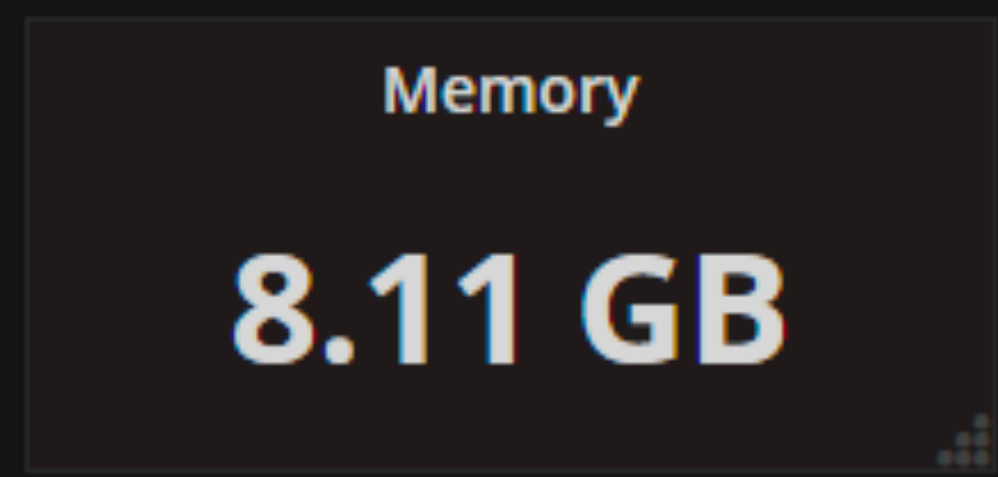
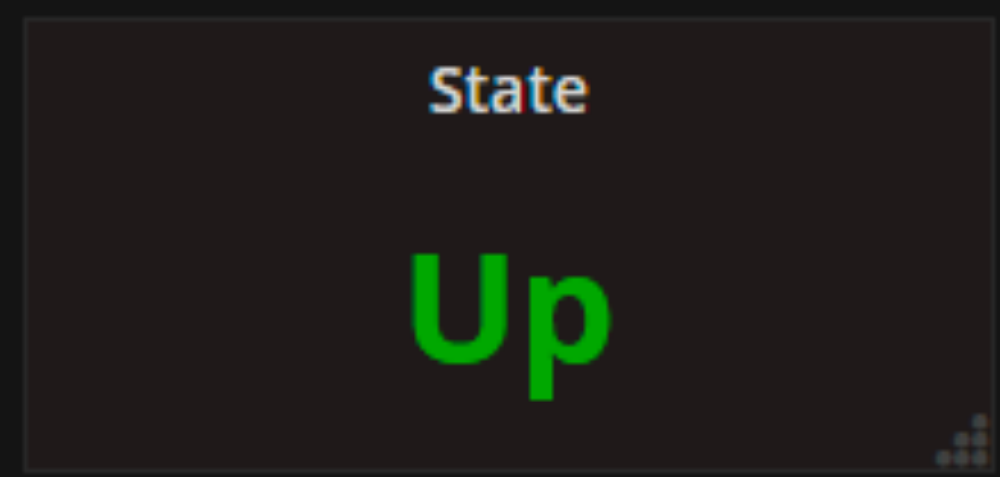

**What do we do when we
have a production issue?**

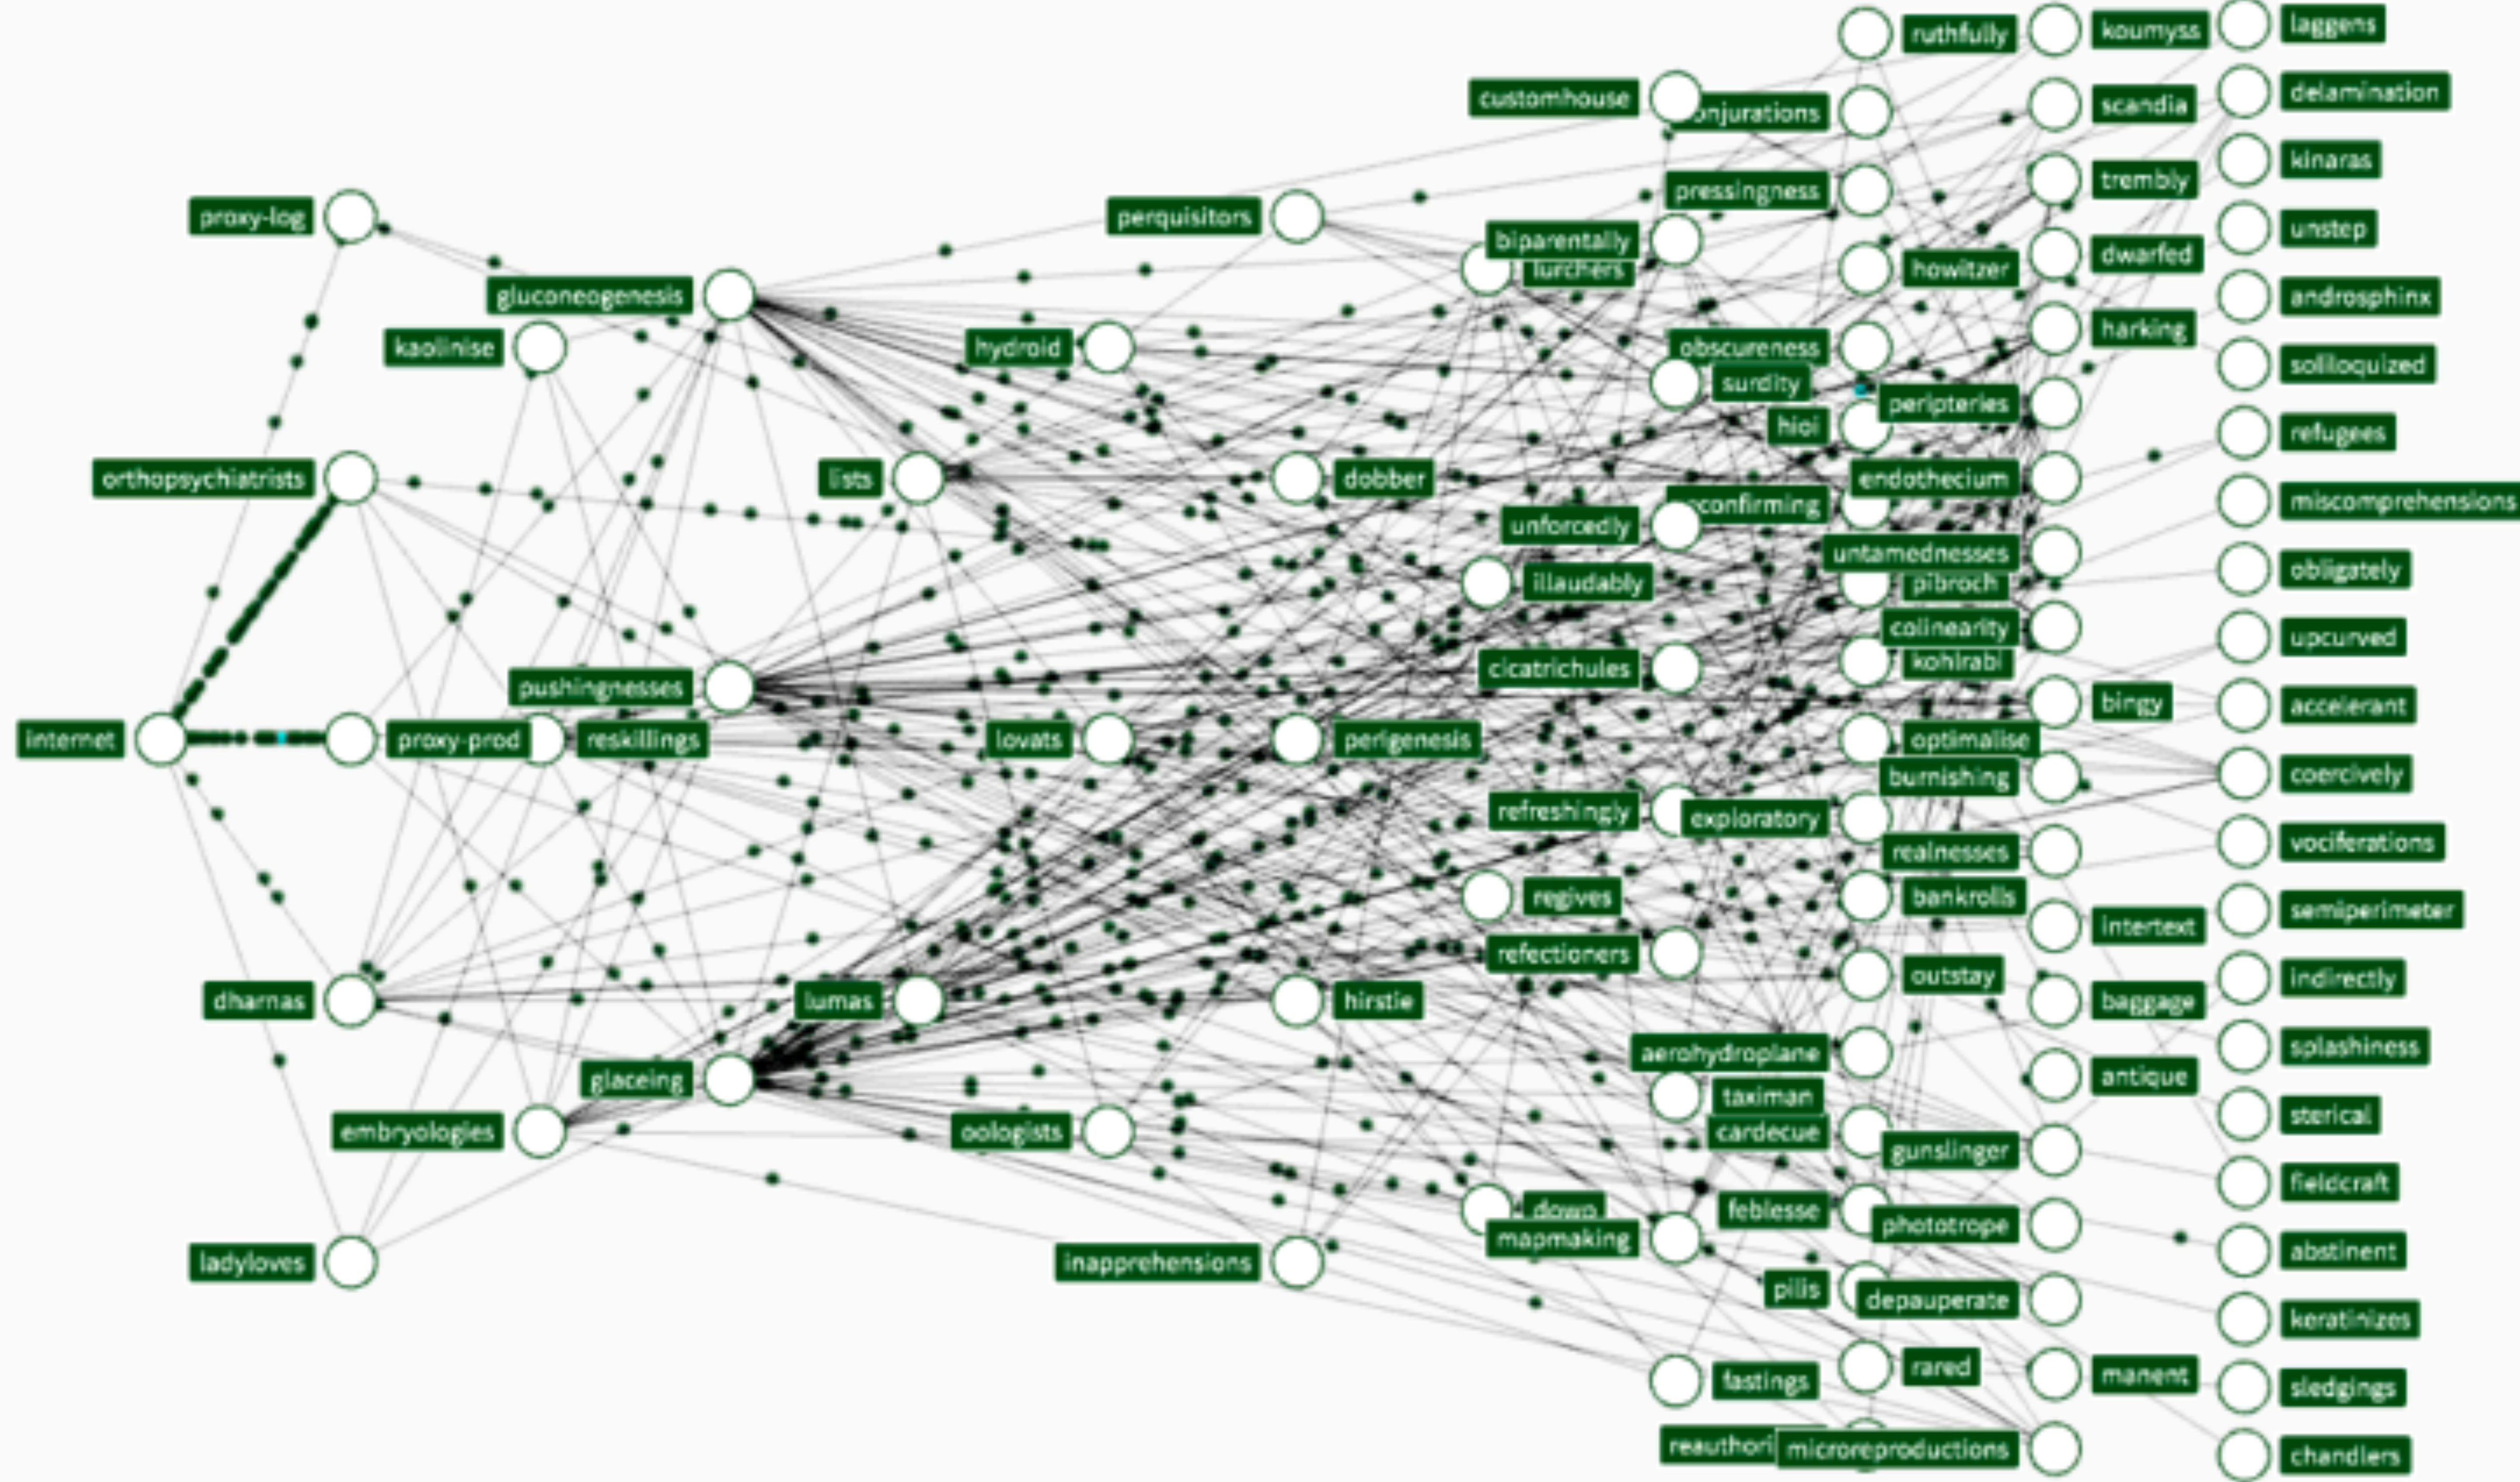
Blame The Database.


The End.











Blym

Bosloc.

C+ /

Fatic

Bimnr

Ficru



[Handwritten notes and diagrams in the first column of the chalkboard]

[Handwritten notes and diagrams in the second column of the chalkboard]

[Handwritten notes and diagrams in the third column of the chalkboard]

[Handwritten notes and diagrams in the fourth column of the chalkboard]

[Handwritten notes and diagrams in the fifth column of the chalkboard]

[Handwritten notes and diagrams in the sixth column of the chalkboard]



Guessing isn't effective.


We need a methodology!

We also need great tools.

Let's Start With The OODA Loop



John Boyd



Observe





Distributed Tracing



350ms



300ms



10ms



span



TraceId=<16 byte random>


frontend: HTTP GET /dispatch



Trace Start: July 20, 2018 2:48 PM | Duration: 732.11ms | Services: 6 | Depth: 5 | Total Spans: 51



Service & Operation






Choose service

payment












Last 2 hour

Aggregation of all "payment" metrics in selected timeframe. [View all traces](#)



Operations metrics under payment Over the last 2 hour

Search operation

Name <input type="button" value="v"/>	P95 Latency <input type="button" value="v"/>	Request rate <input type="button" value="v"/>	Error rate <input type="button" value="v"/>	Impact 1 <input type="button" value="v"/>
payment	 1.92ms	 1.56 req/s	 < 0.1%	
authorize payment	 1.9ms	 0.83 req/s	 < 0.1%	
health check	 1.9ms	 0.67 req/s	No Data	



OpenTelemetry

High-quality, ubiquitous, and portable telemetry to enable effective observability

Structured, Contextual Logging


```
127.0.0.1 alice Alice [06/May/2021:11:26:42  
+0200] "GET / HTTP/1.1" 200 3477
```



```
{
  "ip_address": "127.0.0.1",
  "user_identififier": "alice",
  "user_authentication": "Alice",
  "timestamp": "06/May/2021:11:26:42 +0200",
  "request_method": "GET",
  "request_url": "/",
  "protocol": "HTTP/1.1",
  "status_code": 200,
  "response_size": 3477
}
```


We've Narrowed Down the Problem.

Now What?

Narrow it further.

Dashboards!

Transactions

1-10 of 557 transactions

Last cleared Nov

TRANSACTIONS	STATEMENTS	RETRIES	EXECUTION COUNT	ROWS AFFECTED	LATENCY
SELECT FROM customer	5	0	13k	6	3.1 s
SELECT FROM customer	5	1	10k	6	3.7 s
SELECT FROM customer	5	0	10k	6	3.8 s
INSERT INTO history	4	1	9k	4	3.2 s
INSERT INTO history	4	1	7k	4	3.8 s
INSERT INTO history	4	1	7k	4	3.9 s
UPDATE order	24	0	2k	150	181.1 ms
SELECT FROM order_line	2	0	2k	2	178.8 ms
UPDATE order	24	0	2k	150	185.5 ms
SELECT FROM order_line	2	0	2k	2	201.0 ms

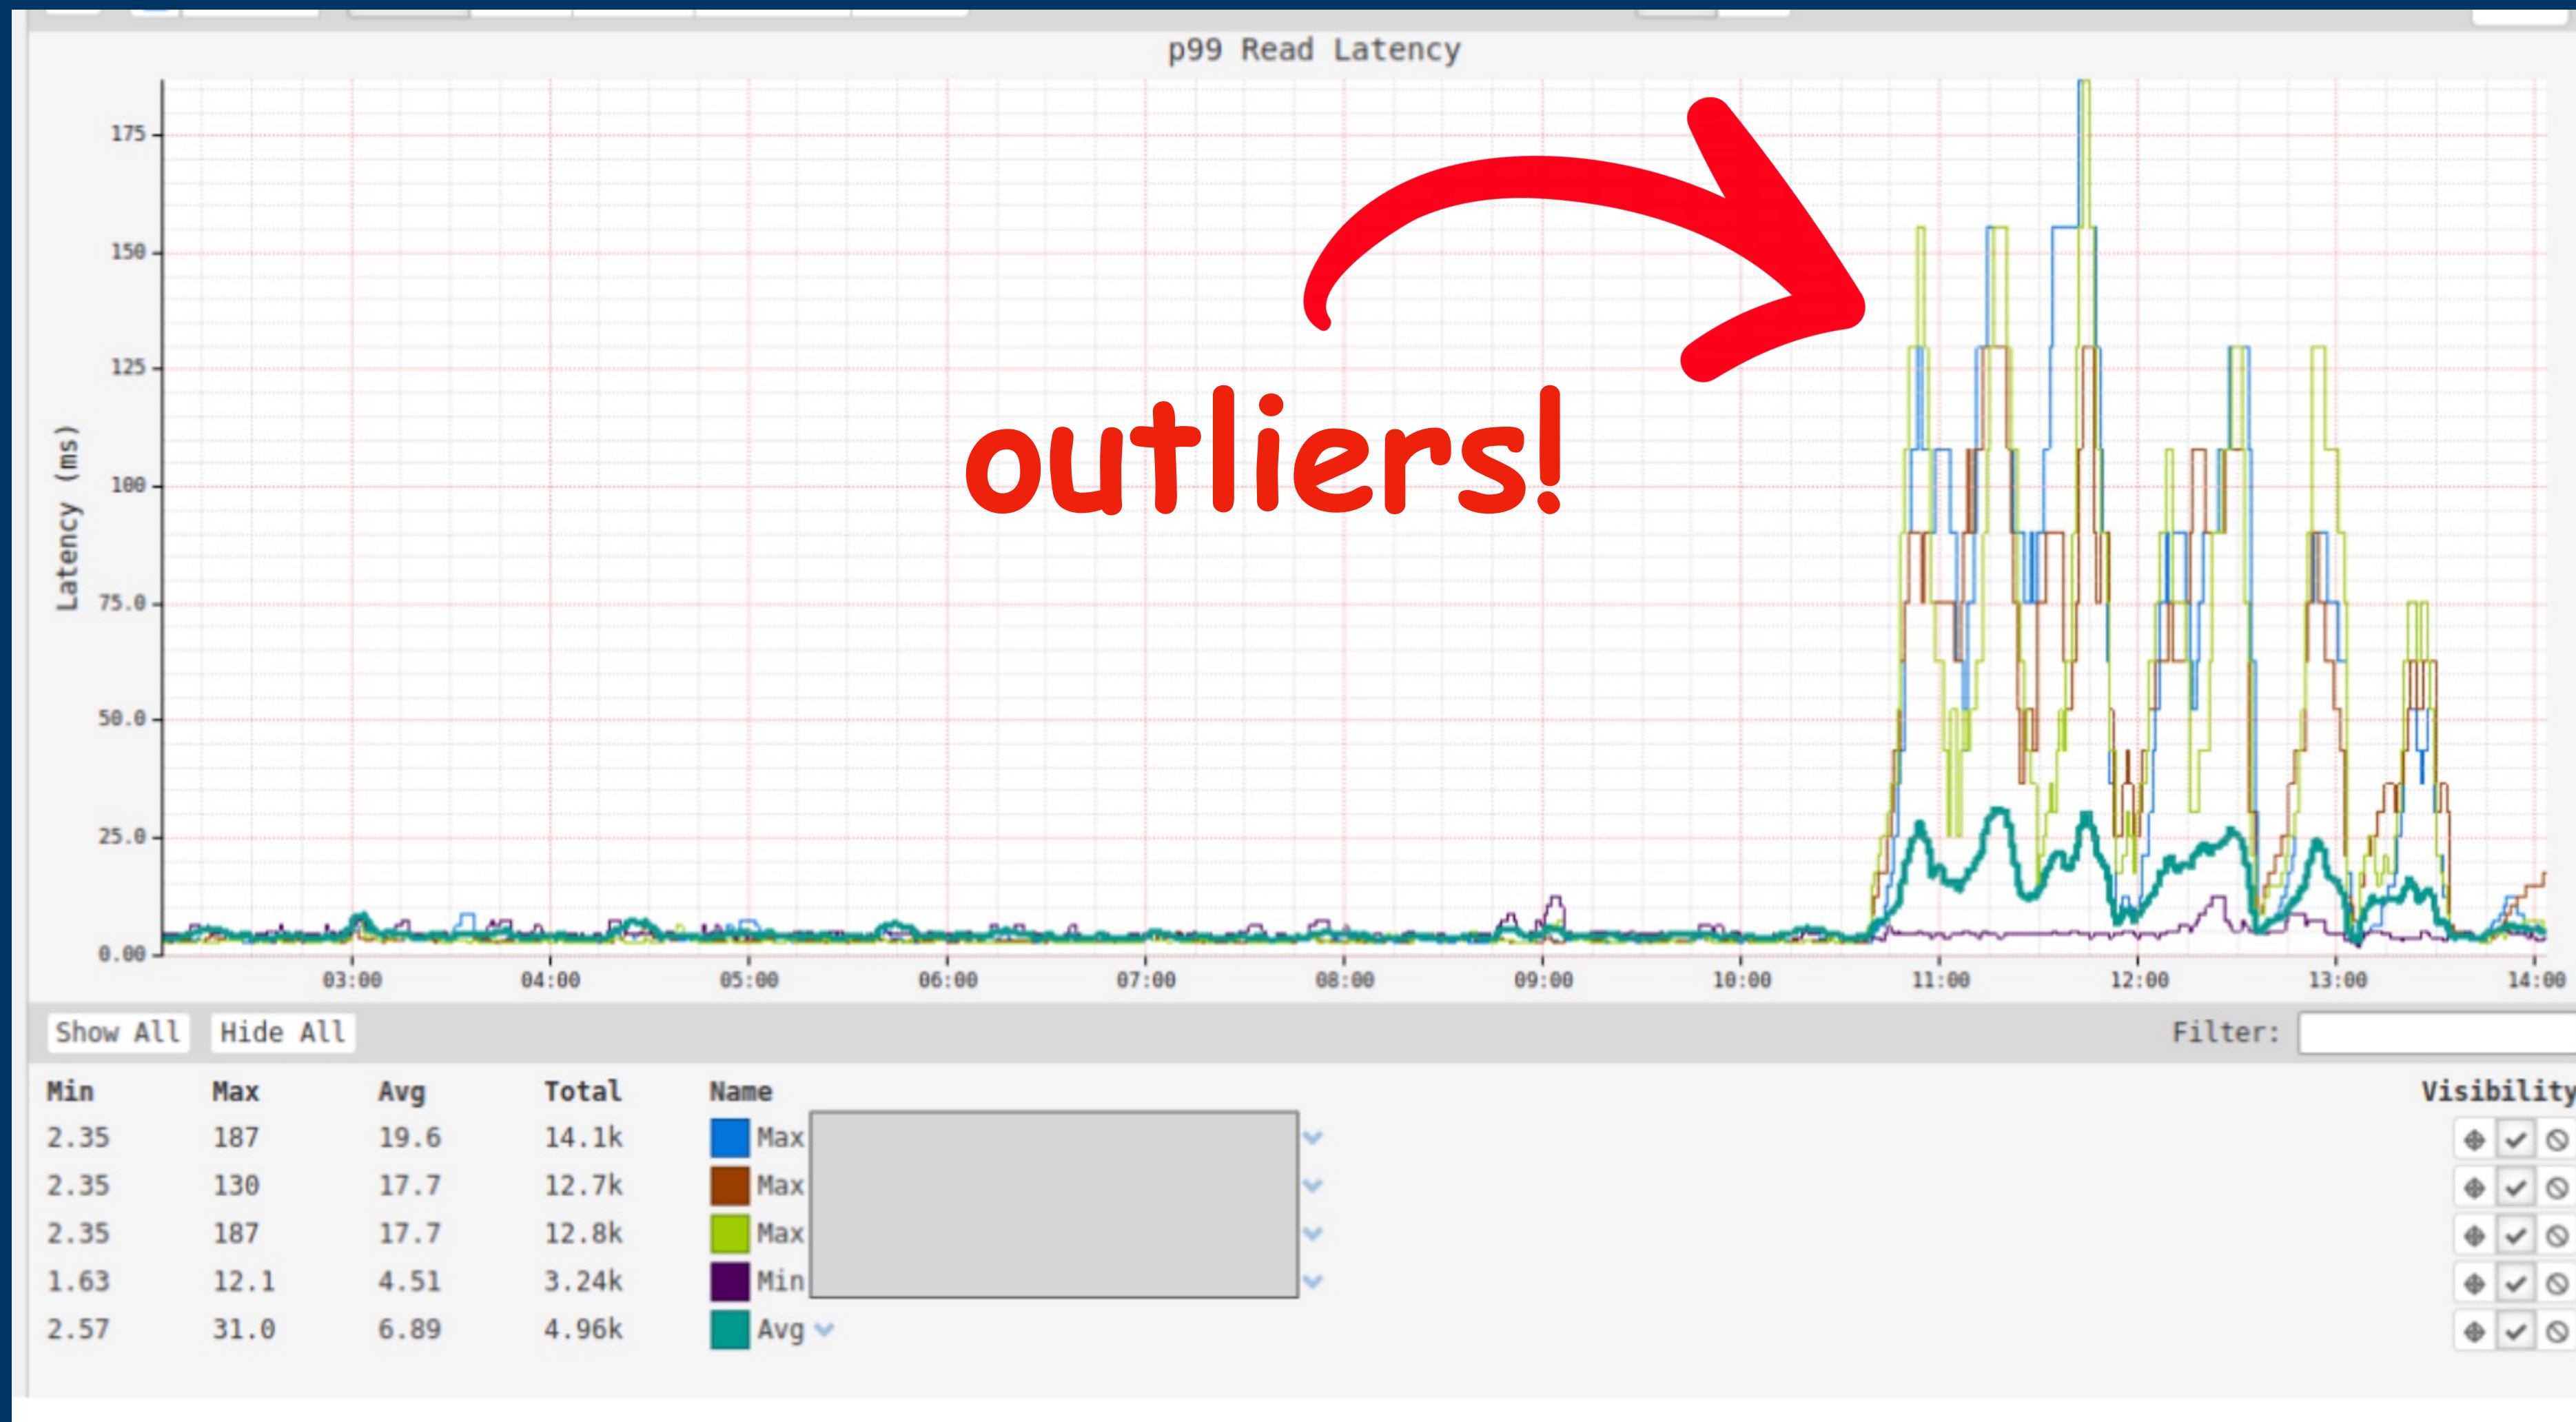
Narrow it further.

Even More Dashboards!


Common Mistakes

- Averaging Percentiles
 - "average p99 latency is X"
- Showing feel good metrics
 - p75, average

Make It Easy



USE Method



Utilization

the average time that the resource was busy servicing work

a percent over a time interval. eg, "one cpu is running at 90% utilization".

Saturation

the degree to which the resource has extra work which it can't service, often queued

as a queue length. eg, "the CPUs have an average run queue length of four".

Errors

the count of error events

scalar counts. eg, "this network interface has had fifty late collisions".


```
root@ubuntu-vm:~# iostat -dmc 2
```

```
Linux 5.15.0-84-generic (ubuntu-vm)
```

```
09/24/2023
```

```
_aarch64_ (2 CPU)
```

```
avg-cpu:  %user  %nice %system %iowait  %steal   %idle  
          0.77   0.00   5.38  42.82   0.00  51.03
```

Device	tps	MB_read/s	MB_wrtn/s	MB_dscd/s	MB_read	MB_wrtn	MB_dscd
dm-0	0.00	0.00	0.00	0.00	0	0	0
dm-1	12165.50	47.52	0.00	0.00	95	0	0
loop0	0.00	0.00	0.00	0.00	0	0	0
loop1	0.00	0.00	0.00	0.00	0	0	0
loop2	0.00	0.00	0.00	0.00	0	0	0
loop3	0.00	0.00	0.00	0.00	0	0	0
sr0	0.00	0.00	0.00	0.00	0	0	0
vda	12165.50	47.52	0.00	0.00	95	0	0

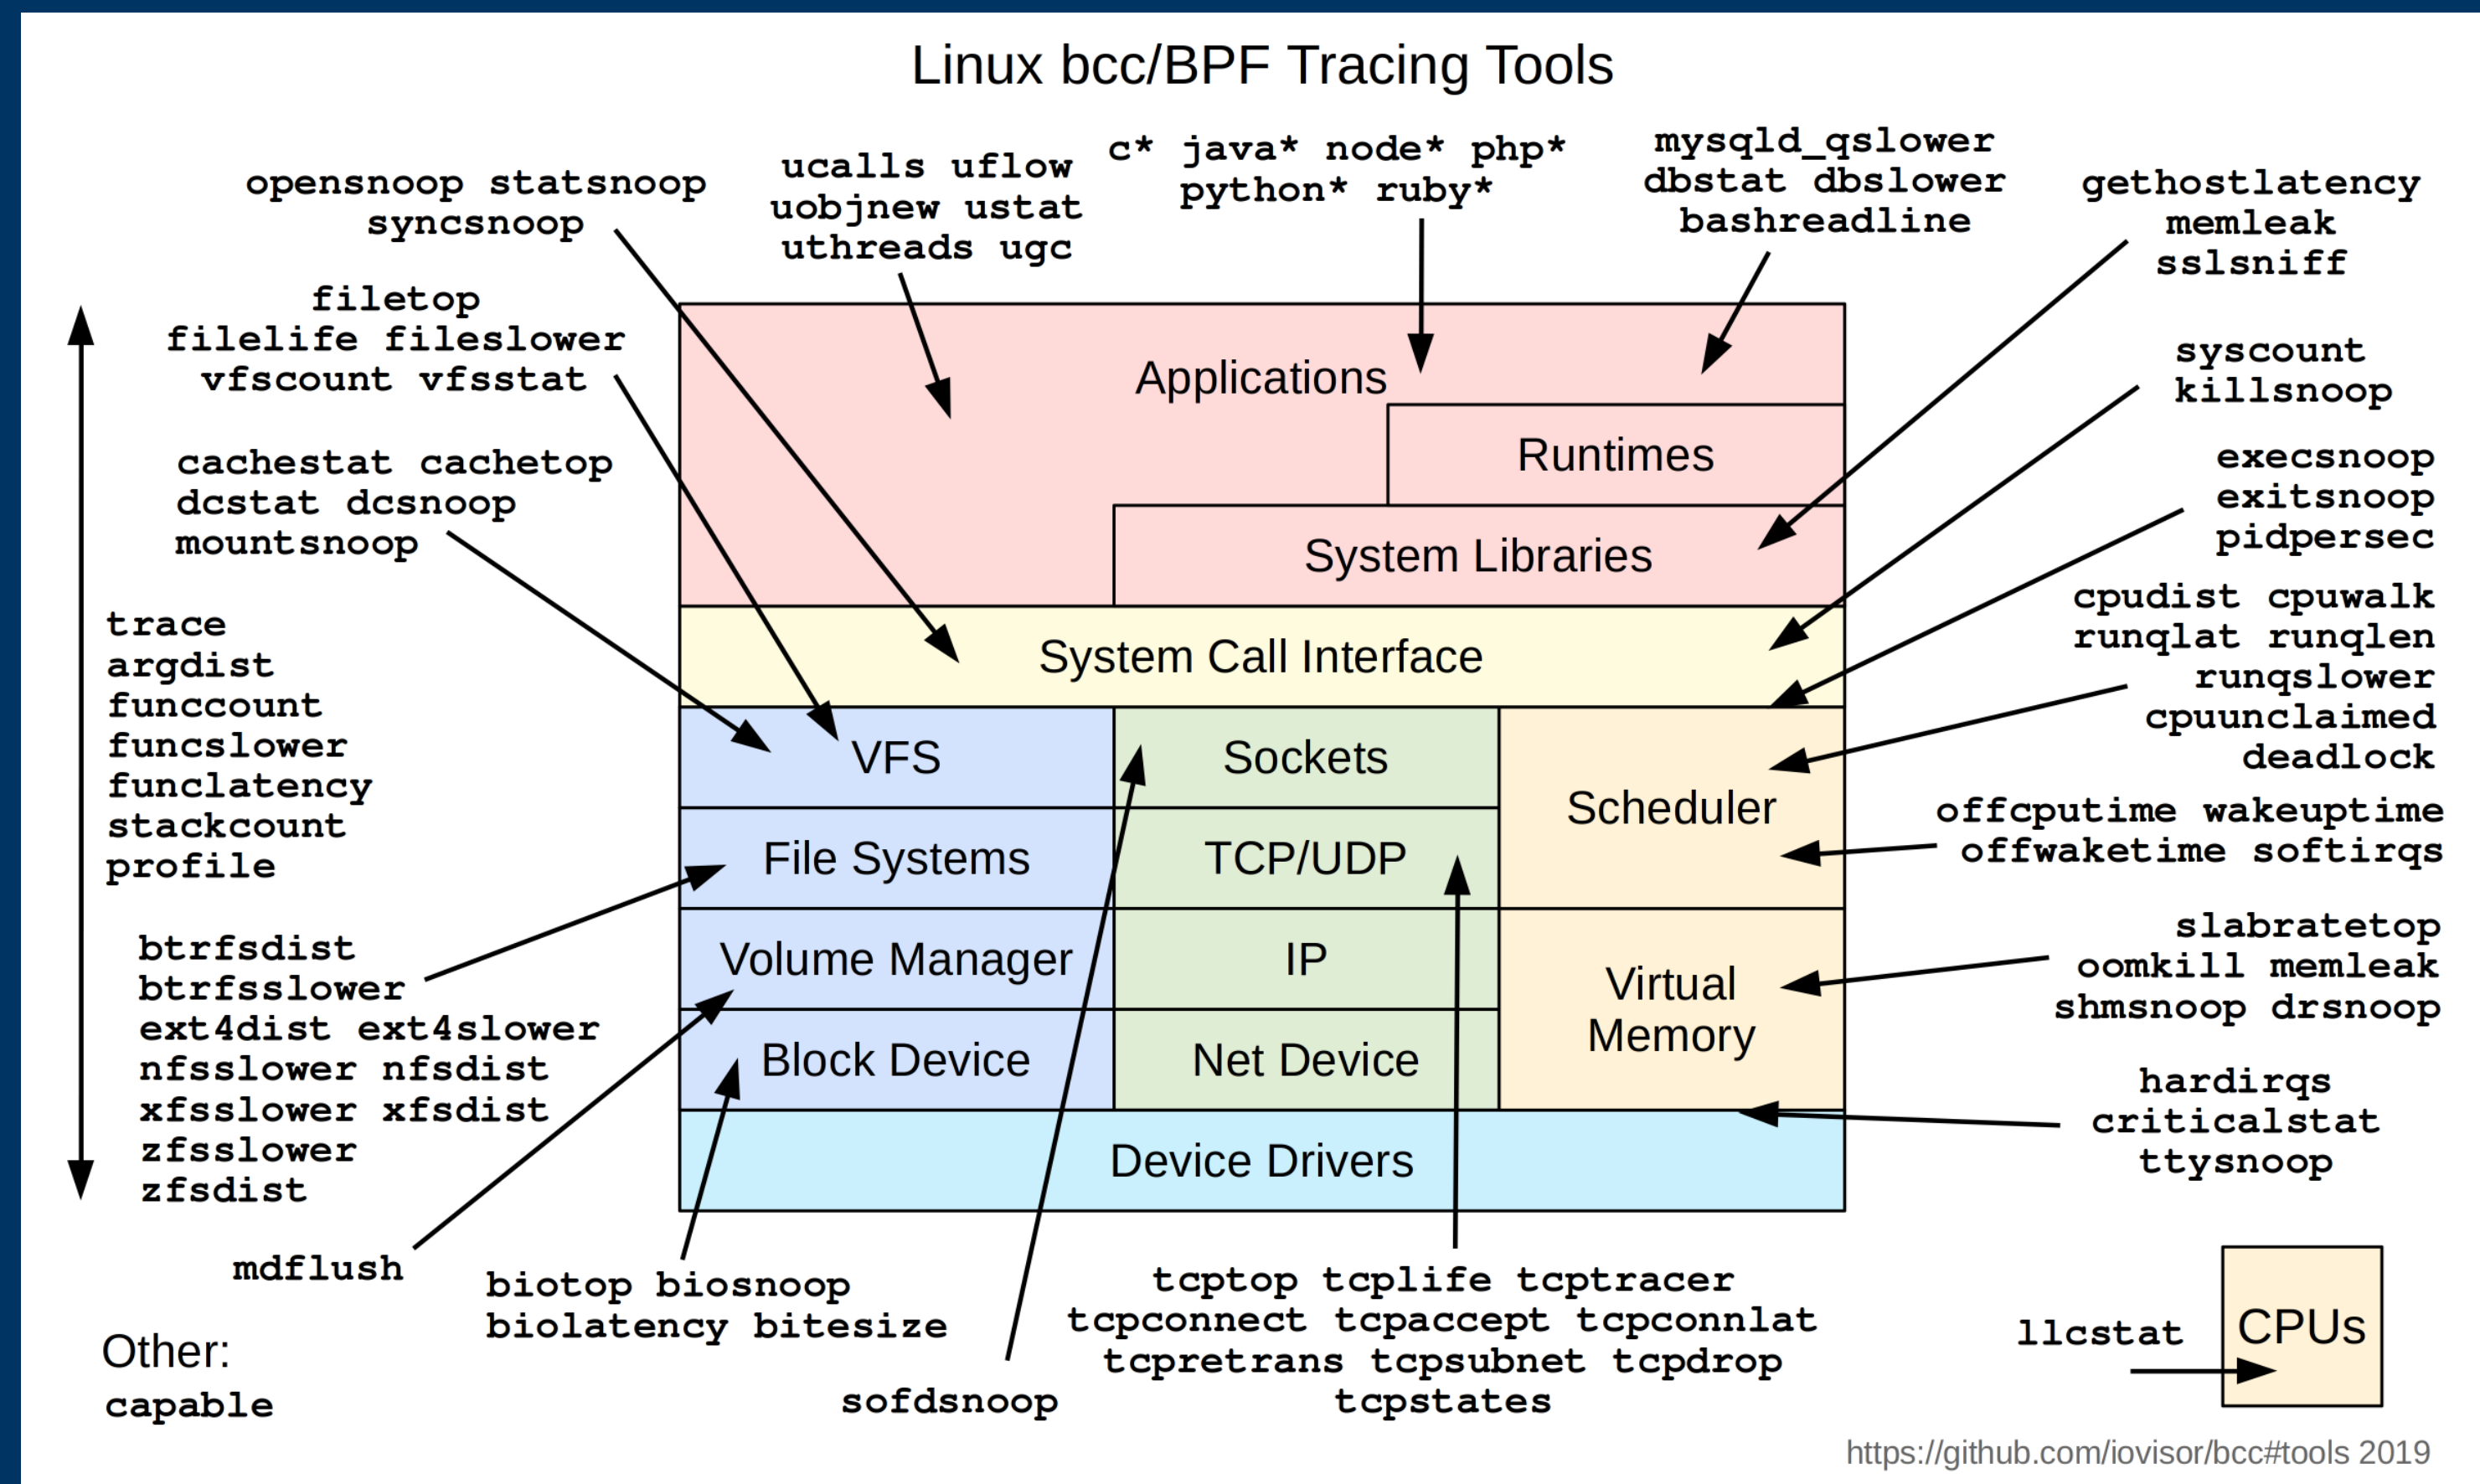
```
root@ubuntu-vm:~# mpstat -P ALL 2
Linux 5.15.0-84-generic (ubuntu-vm)
```

09/24/2023

aarch64 (2 CPU)

03:12:50 AM	CPU	%usr	%nice	%sys	%iowait	%irq	%soft	%steal	%guest	%gnice	%idle
03:12:52 AM	all	0.78	0.00	4.40	43.26	0.00	0.00	0.00	0.00	0.00	51.55
03:12:52 AM	0	1.03	0.00	4.62	39.49	0.00	0.00	0.00	0.00	0.00	54.87
03:12:52 AM	1	0.52	0.00	4.19	47.12	0.00	0.00	0.00	0.00	0.00	48.17
03:12:52 AM	CPU	%usr	%nice	%sys	%iowait	%irq	%soft	%steal	%guest	%gnice	%idle
03:12:54 AM	all	0.78	0.00	4.91	42.89	0.00	0.00	0.00	0.00	0.00	51.42
03:12:54 AM	0	1.03	0.00	5.15	44.33	0.00	0.00	0.00	0.00	0.00	49.48
03:12:54 AM	1	0.52	0.00	4.66	41.45	0.00	0.00	0.00	0.00	0.00	53.37

bcc-tools & bpftrace



```
$ xfsdist 1 10 -m
Tracing XFS operation latency... Hit Ctrl-C to end.
```

```
23:51:18:
```

```
operation = 'read'
```

msecs	: count	distribution
0 -> 1	: 110925	*****

```
23:51:19:
```

```
operation = 'read'
```

msecs	: count	distribution
0 -> 1	: 30715	*****
2 -> 3	: 73	
4 -> 7	: 19	
8 -> 15	: 1	


```
root@ubuntu-vm:~# cachestat 2
```

HITS	MISSES	DIRTIES	HITRATIO	BUFFERS_MB	CACHED_MB
0	24016	0	0.00%	31	709
0	24288	0	0.00%	31	677
0	23686	0	0.00%	31	705
0	22041	0	0.00%	31	664
0	20342	0	0.00%	31	680
0	22785	0	0.00%	31	705
0	22714	0	0.00%	31	666
0	22904	0	0.00%	31	692
0	22805	0	0.00%	31	654
0	22782	0	0.00%	31	679
0	22999	0	0.00%	31	705
0	22851	0	0.00%	31	667
0	22758	0	0.00%	31	692

```
$ root@ubuntu-vm:~# biolatency 2
Tracing block device I/O... Hit Ctrl-C to end.
```

usecs	:	count	distribution
0 -> 1	:	0	
2 -> 3	:	0	
4 -> 7	:	0	
8 -> 15	:	0	
16 -> 31	:	0	
32 -> 63	:	4093	*****
64 -> 127	:	15175	*****
128 -> 255	:	250	
256 -> 511	:	108	
512 -> 1023	:	44	
1024 -> 2047	:	17	
2048 -> 4095	:	9	
4096 -> 8191	:	3	
8192 -> 16383	:	4	
16384 -> 32767	:	3	
32768 -> 65535	:	1	

Profiling is great.

Flame Graphs Are Amazing

Wrap it up buddy

- Keep narrowing down the problem
- Distributed Tracing For High Level
- Use Structured Logging To Build a Complete Picture Across Services
- Analyze Performance Problems With eBPF and Flame Graphs

Thank You!

