Observability for Software Teams

Shelby Spees Developer Advocate @ Honeycomb.io

 \bigotimes

Illustrations by @emilywithcurls!



Production is increasingly complex









NOW



Software teams

DevOps is more than "Dev" and "Ops":

- application engineers
- platform engineers
- infrastructure engineers
- SREs
- SDETs
- support engineers
- and more!



We've come a long way

80 6	Changes approved	Show all reviewers
_	1 approving review Learn more.	
~	1 approval	~
۶	1 pending reviewer	~
e	All checks have passed 1 successful check	Hide all checks
~	y 3 build-hound Successful in 7m — Workflow: build-hound	Details
C	This branch has no conflicts with the base branch Merging can be performed automatically.	
	Squash and merge You can also open this in GitHub Desktop or view command line instructions.	

21	# and private s	ubnet for that list index,
22	# then create an	n EIP and attach a nat_gateway for each one. and an aws rout
23	# table should I	be created for each private subnet, and add the correct nat_g
24		
25	resource "aws_si	ubnet" "private" {
26	vpc_id	= aws_vpc.mod.id
27	cidr_block	<pre>= var.private_ranges[count.index]</pre>
28	availability_	<pre>zone = var.azs[count.index]</pre>
29	count	<pre>= length(var.private_ranges)</pre>
30	tags = {	
31	Name = "\${va	ar.env}_private_\${count.index}"
32	}	
33	}	
34		
35	resource "aws_s	ubnet" "public" {
36	vpc_id	= aws_vpc.mod.id
37	cidr_block	<pre>= var.public_ranges[count.index]</pre>
38	availability_	<pre>zone = var.azs[count.index]</pre>
39	count	<pre>= length(var.public_ranges)</pre>
40	tags = {	
41	Name = "\${va	ar.env}_public_\${count.index}"
42	}	
43	<pre>map_public_ip</pre>	_on_launch = true
44	}	
45		
46	# refactor to ta	ake all the route {} sections out of routing tables,
47	# and turn them	into associated aws_route resources
48	# so we can add	vpc peering routes from specific environments.
49	resource "aws_re	<pre>pute_table" "public" {</pre>
50	<pre>vpc_id = aws_v</pre>	vpc.mod.id
51	tags = {	
52	Name = "\${va	ar.env}_public_subnet_route_table"

But we still encounter problems

emergent failure modes

small failures cascading together to degrade or take down a system

see also: how.complexsystems.fail





Getting paged can be scary





Our tools don't answer our questions



We don't make forward progress



of developer time is spent on dealing with bad code and tech debt



Source: The Developer Coefficient, Stripe, 2018



It doesn't have to be this way

We have the technology!

How can we better interact with production?

Our applications have all the answers at runtime!

We want to

- capture those answers
- make that data available to interrogate

Big companies have done this

- at Facebook: <u>Scuba</u>
- at Google: <u>Dapper</u>





We need observability!

What is observability?

in software, the ability to understand and explain any state a system can get into, no matter how novel or bizarre, without deploying new code



How to gain observability?

Observability requires

- capturing telemetry data with lots of runtime context
- interacting with that telemetry in near-real time



Instrumentation generates telemetry

instrumentation: code or tooling that captures data about the state of your running system

telemetry: data generated by your system that documents its state

```
"timestamp": "2018-03-20T00:47:25.339Z",
"app.interesting_thing": "banana",
"duration_ms": 772.446625,
"handler.name": "main.hello",
"handler.pattern": "/hello/",
"handler.type": "http.HandlerFunc",
"meta.beeline_version": "0.2.0",
"meta.local_hostname": "cobbler.local",
"meta.type": "http_request",
"name": "main.hello",
"request.header.user_agent": "curl/7.54.0",
"request.http_version": "HTTP/1.1",
"request.method": "GET",
"request.path": "/hello/",
"request.remote_addr": "127.0.0.1:60379",
"service_name": "sample app",
"trace.span_id": "9e4fe697-3ea9-48c9-b673-72d7ddf118a6",
"trace.trace_id": "b64c89a9-7671-4732-bef1-9ef75ab831f6"
```

You're already collecting data

You've probably instrumented your system to generate telemetry like

- metrics
- logs
- traces

That's great!





But traditional approaches aren't enough





Observability is for hard-to-debug problems

Distributed Systems

small change causing downstream effects?

Poor Performance

what part of the app is worth optimizing?

Subset of Traffic

only some users are complaining?



Metrics

great when thinking about infrastructure!

but they:

- only answer known questions
- don't capture application context
- don't support high-cardinality use cases



Metrics have little context

2021 Hound Technology, Inc. All Rights Reserved

Letters in Winning Word of Scripps National Spelling Bee correlates with Number of people killed by venomous spiders



- Number of people killed by venomous spidersSpelling Bee winning word

tylervigen.com

https://www.tylervigen.com/spurious-correlations



- capturing output from the code itself
- no standard log formatting across libraries or services
- requires string parsing--expensive at scale

Centralized logging services are expensive and slow to query

• great for compliance, not for debugging

Distributed traces

Visualization of the stack trace across your distributed system

Traces are formed by a directed graph made up of objects called *spans*

 each span has a start time and a duration, and points to its parent span



distributed trace visualization in Jaeger UI

h © 2021 Hound Technology, Inc. All Rights Reserved.

1.243s API call

What made it slow?



Structured events!

What are structured events?

structured logs + benchmarking + (optionally) tracing

```
{
    "Timestamp": "2018-07-03T04:57:12.517022Z",
    "duration_ms": 619.703,
    "request.method": "GET",
    "request.path": "/",
    "request.user_agent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_13_4) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/67.0.3396.99 Safari/537.36",
    "response.status_code": 200,
    "trace.span_id": "3eada0ce-934b-4ffd-bb72-2c9f57b02bf1",
    "trace.parent_id": "aa23-3becd17d726-607d2150c-5397-4e8e",
    "trace.trace_id": "07d2150c-5397-4e8e-aa23-3becd17d7266",
    ...
}
```

Keep all your runtime context together

No predefined index or schema

add fields as needed

Key-value pairs

less expensive to parse and query



Flat logs vs. structured events

Flat logs get written eagerly

You can't keep track of state changes, even within a specific context



```
@PostMapping("flip")
public ResponseEntity flipImage(@ReguestParam("image") MultipartFile file,
                                @RequestParam(value = "vertical") Boolean vertical,
                                @RequestParam(value = "horizontal") Boolean horizontal) {
     if (file.getContentType() != null) {
                                                                                 write as you go
        LOGGER.warn("Wrong content type uploaded: {}", file.getContentType());
        return new ResponseEntity<>("Wrong content type uploaded: " + file.getContentType());
                                                                                 write as you go
    LOGGER.info("Receiving {} image to flip.", file.getContentType());
    byte[] flippedImage = imageService.flip(file, vertical, horizontal);
   if (flippedImage == null) {
        return new ResponseEntity<>("Failed to flip image", HttpStatus.INTERNAL_SERVER_ERROR);
                                                                                 write as you go
    LOGGER.info("Successfully flipped image id: {}", file.getId());
    return new ResponseEntity<>(flippedImage, headers, HttpStatus.OK);
```

Flat logs vs. structured events

Structured events store all the data within your context, from beginning to end



```
@PostMapping("flip")
public ResponseEntity flipImage(...) {
    EVENT.addField("content.type", file.getContentType());
    EVENT.addField("action", "flip");
    EVENT.addField("image_id", file.getId());
    EVENT.addField("flip_vertical", vertical);
    EVENT.addField("flip horizontal", horizontal);
. . .
    LOGGER.info("Receiving {} image to flip.", file.getContentType());
    bvte[] flippedImage = imageService.flip(file, vertical, horizontal);
. . .
    LOGGER.info("Successfully flipped image id: {}", file.getId());
    EVENT.addField("action.success", "true");
    return new ResponseEntity<>(flippedImage, headers, HttpStatus.OK);
```

write at the end

From events: time-series graphs



elapsed query time: 3.404466517s rows examined: 216,747,687 nodes reporting: 100%





From events: error logging

response.status_code	app.error		COUNT		
500	something went wrong!				
500	We couldn't fetch the completed query results! Please reload in a bit.				
500	rpc error: code = Canceled desc = context canceled				
500	context canceled		10		
500	<pre>context canceled /home/circleci/project/types/dataset.go:1297 /home/circleci/project/cmd/poodle/handlers/queries.go:498 /home/circleci/project/cmd/poodle/handlers/datasets.go:195 /home/circleci/project/cmd/poodle/handlers/teams.go:193 /home/circleci/project/cmd/poodle/handlers/teams.go:193 /home/circleci/project/cmd/poodle/handlers/datasets.go:83 /home/circleci/project/cmd/poodle/handlers/datasets.go:118 /home/circleci/project/cmd/poodle/handlers/queries.go:468 /home/circleci/project/cmd/poodle/handlers/queries.go:89 /home/circleci/project/cmd/poodle/handlers/util.go:809 /home/circleci/project/cmd/poodle/handlers/util.go:300 /usr/local/go/src/net/http/server.go:2041 /home/circleci/project/cmd/poodle/handlers/util.go:627 /usr/local/go/src/net/http/server.go:2041 /home/circleci/project/cmd/poodle/handlers/util.go:531 /usr/local/go/src/net/http/server.go:2041 /home/circleci/project/cmd/poodle/handlers/util.go:455 /home/circleci/project/cmd/poodle/handlers/util.go:455</pre>	<pre>(*Dataset).FindSchemaWithoutEditors (*Dataset).FindSchema (*QueriesHandler).Samples.func2 (*Handler).withDataset.func1 (*Handler).withTeam (*Handler).withDataset (*Handler).withDataset (*Handler).withDatasetMaybePublic (*QueriesHandler).Samples AppHandler.ServeHTTP LogRequest.func1 HandlerFunc.ServeHTTP AttachRequestID.func1 HandlerFunc.ServeHTTP AttachReiderFunc.ServeHTTP AttachBuildInfo.func1.1 HandlerFunc.ServeHTTP AttachBuildInfo.func1.1 HandlerFunc.ServeHTTP AttachBuildInfo.func1.1</pre>	4		

From events: tracing

If your events have start time, duration, parent: then you can generate a trace!



← Trace 30d094c6-7f39-4c2e-8f2b-3b843e051b8e at 2020-08-18 21:01:56

Rerun

name v @	service_name v	0s	1s	23	35	4s	5.015
4 http_request	rails	5.015s					
redis	rails	1.164ms					
4 process_action.action_controller	rails	5.002s					
sql.active_record	rails	2.350ms					
6 render_template.action_view	rails	4.982s					
3 render_partial.action_view	rails	7.617ms					
- sql.active_record	rails	1.544ms					
sql.active_record	rails	1.568ms					
+ sql.active_record	rails	1.539ms					
-4 render_partial.action_view	rails	6.888ms					
-+ sql.active_record	rails	1.411ms					
sql.active_record	rails	0.1093ms					
- sql.active_record	rails	0.1079ms					
sql.active_record	rails	0.1083ms					
sql.active_record	rails	1.474ms					
sql.active_record	rails	4.951s					
sql.active_record	rails						4.224ms
render_partial.action_view	rails						0.1686ms
1 render_partial.action_view	rails						8.621ms
1 render_partial.action_view	rails						1.818ms
redis	rails						1.016ms
render_partial.action_view	rails						1.931ms
redis	rails						1.399ms
- sql.active_record	rails						0.1265ms

sql.active_record: 4.951s

Trace 30d094c6-7f39-4c2e-8f2b-3b843e051b8e at 2020-08-18 21:01:56

arch spans < >							III Fields
name 🗸 🖨	service_name ~	Os	1s	2s	3s	4s	5.015
4 http_request	rails	5.015s					
redis	rails	1.164ms					
4 process_action.action_controller	rails	5.002s	1				
sql.active_record	rails	2.350ms					
6 render_template.action_view	rails	4.982s		0			
3 render_partial.action_view	rails	7.617ms					
- sql.active_record	rails	1.544ms					
sql.active_record	rails	1.568ms					
+ sql.active_record	rails	1.539ms					
4 render_partial.action_view	rails	6.888ms					
sql.active_record	rails	1.411ms					
sql.active_record	rails	0.1093ms					
- sql.active_record	rails	0.1079ms					
sql.active_record	rails	0.1083ms					
sql.active_record	rails	1.474ms					
sql.active_record	rails	4.9518					
sql.active_record	rails						4.224ms
render_partial.action_view	rails						0.1686ms
render_partial.action_view	rails						8.621ms
render_partial.action_view	rails						1.818ms
redis	rails						1.016ms
render_partial.action_view	rails						1.931ms
redis	rails						1.399ms
sql.active_record	rails						0.1265ms

Fields Filter fields and values in span name sql.active_record service_name rails sql.active_record.connection # <ActiveRecord::ConnectionAdapters::PostgreSQ LAdapter:0x00007f39eb11c868> sql.active_record.connection_id 1147120 sql.active_record.name Follow Load sql.active_record.sql SELECT "follows".* FROM "follows" WHERE "follows"."follower_id" = AND "follows"."follower_type" = 'User' AND "follows"."blocked" = FALSE AND "follows"."followable_type" = 'User' ORDER BY

"follows"."created_at" DESC LIMIT 80

© 2021 Hound Technology, Inc. All Rights Reserved.

@shelbyspees at #DevWeek2021

Rerun

000

Instrument your code

Instrument with OpenTelemetry!

Vendor-neutral framework with auto-instrumentation

Learn more: **OpenTelemetry.io**



Start with auto-instrumentation

Rich context and distributed tracing out of the box:

- HTTP headers
- gRPC calls
- SQL queries



Trace 30d094c6-7f39-4c2e-8f2b-3b843e051b8e at 2020-08-18 21:01:56

name 🗸 o	rails	0s	15	70			
4 http://www.upped.	rans	E 01E-		23	38	45	5.015
4 nttp_request		5.0155					
redis	rails	1.164ms					
4 process_action.action_controller	rails	5.002s					
sql.active_record	rails	2.350ms					
6 render_template.action_view	rails	4.982s					
3 render_partial.action_view	rails	7.617ms					
- sql.active_record	rails	1.544ms					
sql.active_record	rails	1.568ms					
+ sql.active_record	rails	1.539ms					
-4 render_partial.action_view	rails	6.888ms					
-+ sql.active_record	rails	1.411ms					
- sql.active_record	rails	0.1093ms					
- sql.active_record	rails	0.1079ms					
- sql.active_record	rails	0.1083ms					
sql.active_record	rails	1.474ms					
sql.active_record	rails	4.9518					
sql.active_record	rails						4.224ms
render_partial.action_view	rails						0.1686ms
1 render_partial.action_view	rails						8.621ms
1 render_partial.action_view	rails						1.818ms
redis	rails						1.016ms
render_partial.action_view	rails						1.931ms
redis	rails						1.399ms
sql.active record	rails						0.1265ms

Fields Filter fields and values in span name sql.active_record service_name rails sql.active_record.connection # <ActiveRecord::ConnectionAdapters::PostgreSQ LAdapter:0x00007f39eb11c868> sql.active_record.connection_id 1147120 sql.active_record.name Follow Load sql.active_record.sql SELECT "follows".* FROM "follows" WHERE "follows"."follower_id" = AND "follows"."follower_type" = 'User' AND "follows"."blocked" = FALSE AND "follows"."followable_type" = 'User' ORDER BY

"follows"."created_at" DESC LIMIT 80

© 2021 Hound Technology, Inc. All Rights Reserved.

@shelbyspees at #DevWeek2021

Rerun

000

Custom instrumentation

What's important for the service that I'm providing?

- lib/ code consumed by your request endpoints
- 3rd-party API calls



.

```
def get_all_tweets(user)
  all = []
  options = { count: 200, include_rts: true }
  loop do
    tweets = client.user_timeline(user, options)
    return all if tweets.empty?
    all += tweets
    options[:max_id] = tweets.last.id - 1
  end
end
```

```
def get all tweets(user)
   Instrumentation.start_span(name: 'get_all_tweets') do
      Instrumentation.add field to trace('user', user)
     all = []
     options = { count: 200, include rts: true }
      loop do
         Instrumentation.start span(name: 'get batch') do
           Instrumentation.add field('options', options)
           tweets = client.user timeline(user, options)
           return all if tweets.empty?
           all += tweets
           options[:max_id] = tweets.last.id - 1
         end
                                    name 🗸 🗅
                                                      service_name
                                                                     0s
                                                                          0.2s
                                                                               0.4s
                                                                                     0.65
                                                                                          0.8s
      end
                                                                      1.678s
                                    8 get_all_tweets
                                                      twitter-example
                                                                      279.3ms

    get_batch

                                                      twitter-example
   end
                                                                              215.3ms
                                       get_batch
                                                      twitter-example
end
                                       get_batch
                                                      twitter-example
                                                                                    232.7ms
                                                                                          224.9ms
                                       get_batch
                                                      twitter-example
                                     · get_batch
                                                      twitter-example
                                       get_batch
                                                      twitter-example
                                       get batch
                                                      twitter-example
                                       get_batch
                                                      twitter-example
```



@shelbyspees at #DevWeek2021

15

201.3ms

1.25

235.3ms

152.2ms

1.4s

1.678s

How does this benefit teams?

Better data \rightarrow better conversations

Shared source of truth.

Improved knowledge transfer.

Shared ownership of production.



Observability-driven development

Before making a change, ask:

How will I know this code is working as intended in prod?

Observe the current state:

What am I trying to change?

Then implement + instrument: Look, it's working!

One tool for dev and prod

Fast feedback loop in dev

Strong debugging skills in prod

Read more: The Future of Developer Careers go.hny.co/dev-careers



Level up your team

Better data leads to better conversations

Custom instrumentation enables self-serve querying

Knowledge transfer creates production ownership



Where to start?

Start where you're at

Take inventory of the tools and data your team has access to

- How often are people actually using them?
- How often are you actually able to answer questions?
- How much of the time do you rely on experience (and guessing) vs. data?

Instrument your code

Structured events!

- How to capture them?
- Look into OpenTelemetry
- Is there auto-instrumentation for your tech stack?

Start small and grow

Start with one app or one chunk of code in your critical path

• Critical code helps you learn the fastest

Try out some tools

- Some vendors offer a free tier or trial
- Send from dev or deploy a canary



Reach out!

Get these slides: hny.co/shelby Twitter: @shelbyspees



Visit our Booth!

Get O'Reilly pre-release chapters! Get Honeycomb swag!

O'REILLY"

Observability Engineering

Achieving Production Excellence





Questions?