A Practical Guide to Dashboarding

Jessica White
Software Engineer

What Is A Dashboard & Why Should I Care?

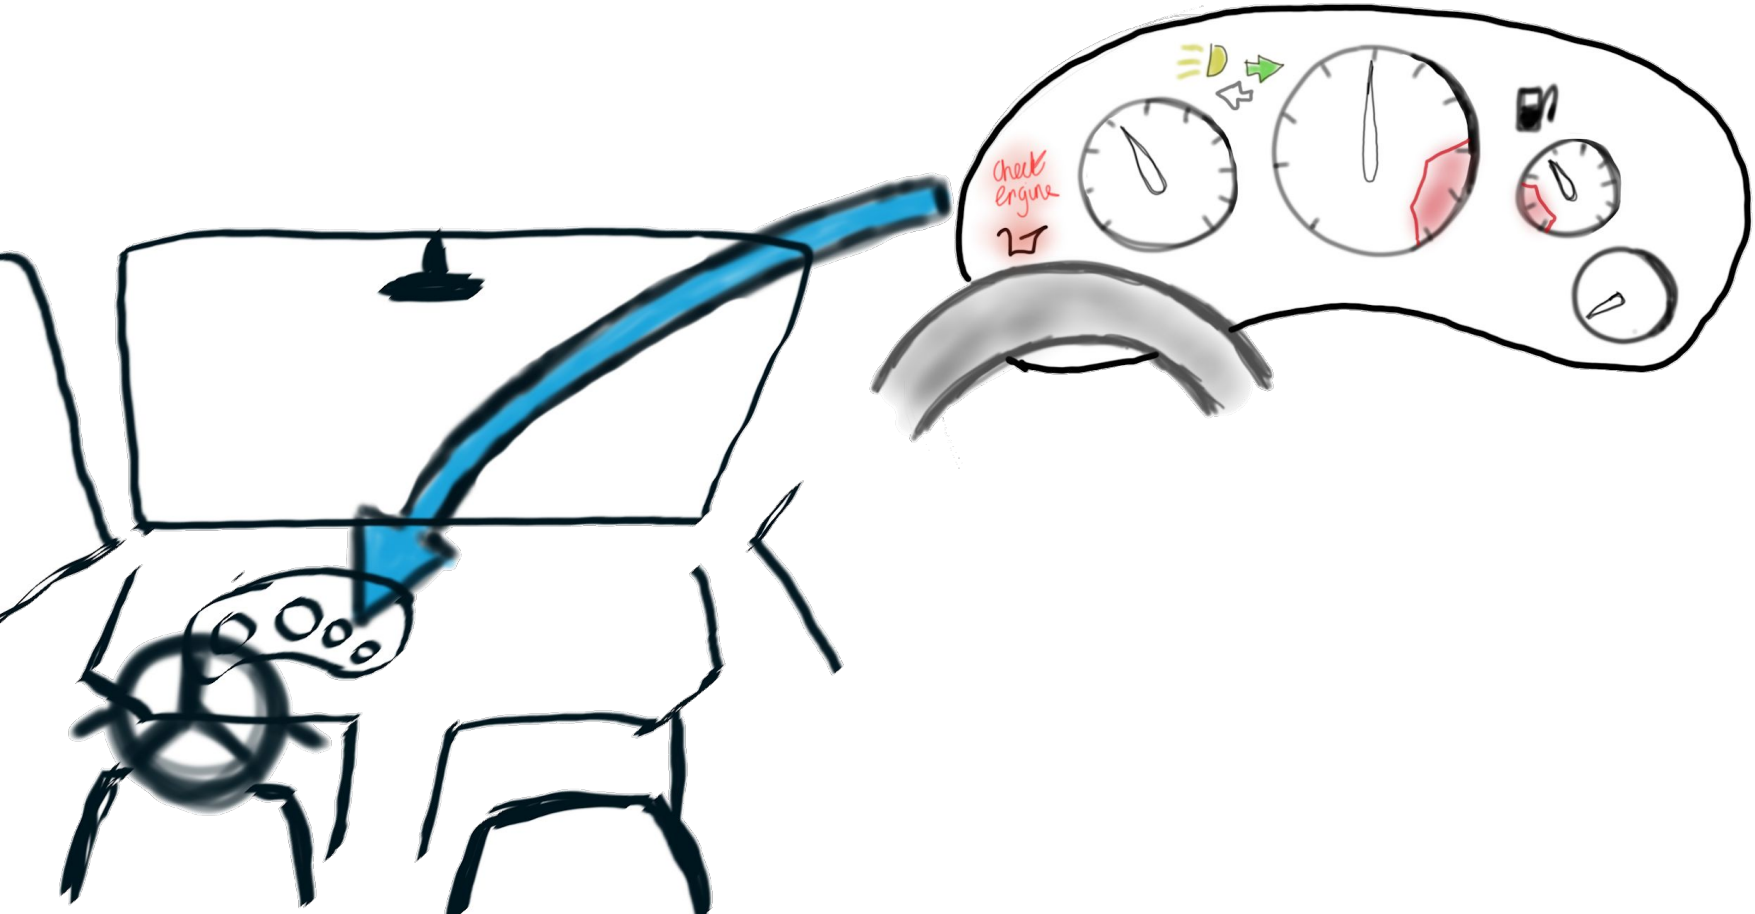
check
engine

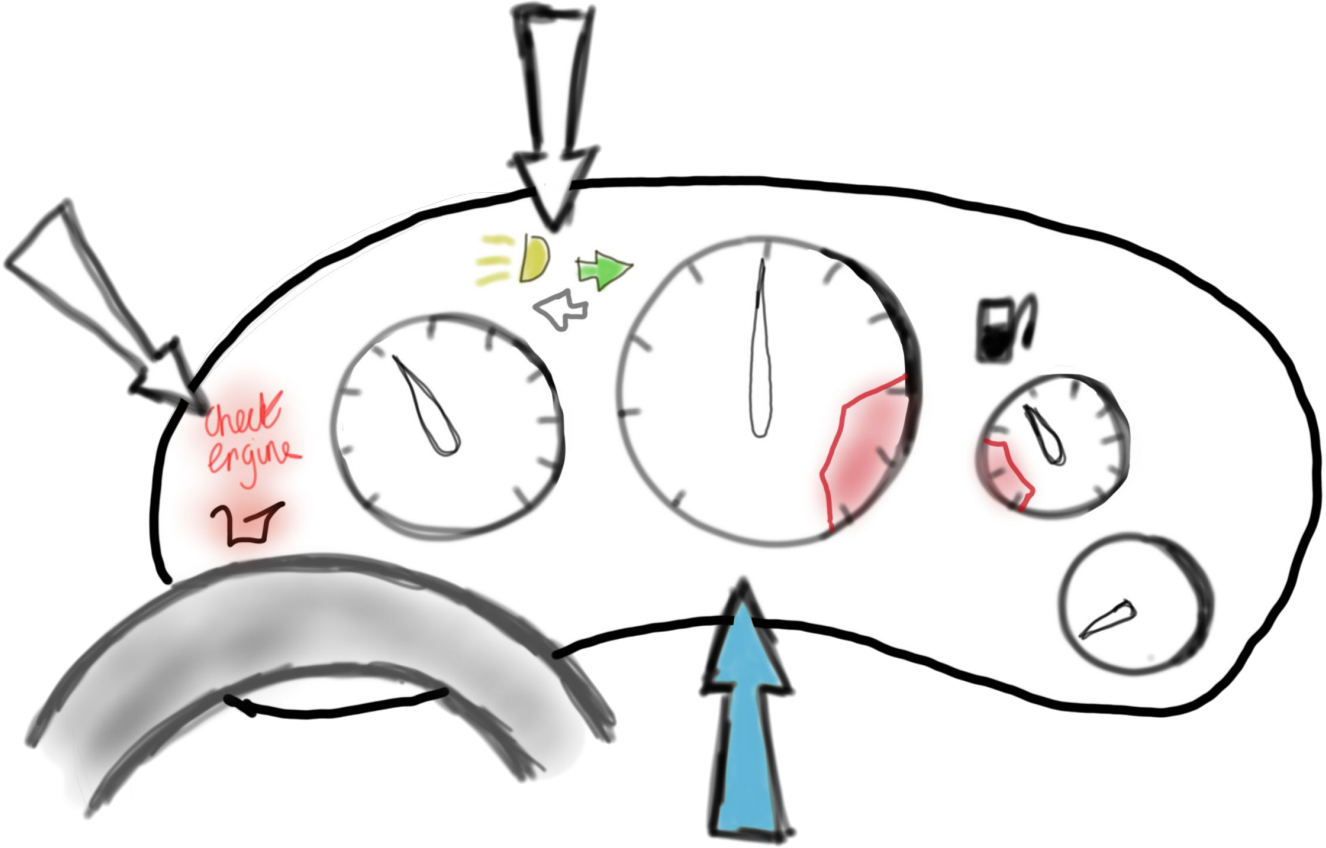
check
engine

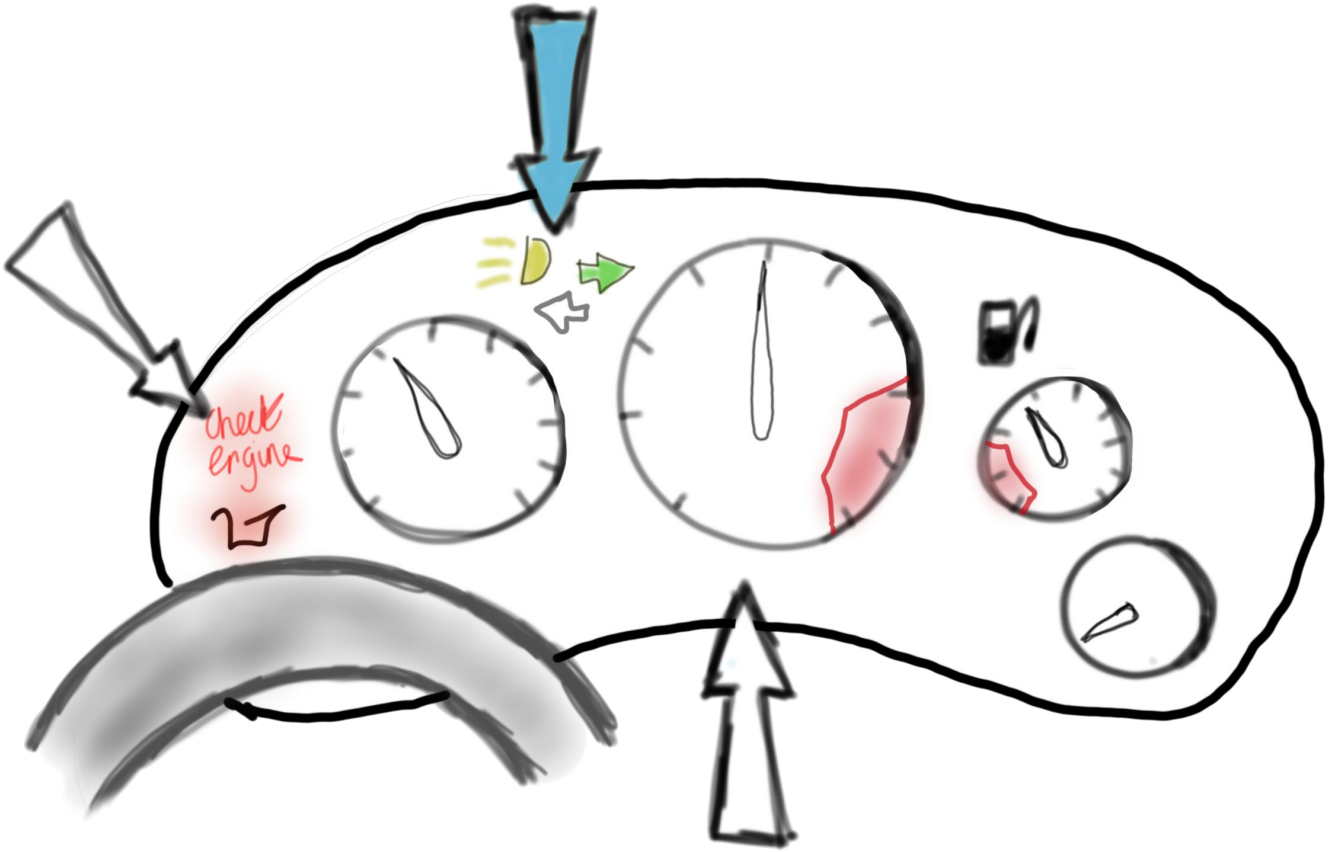
check
engine

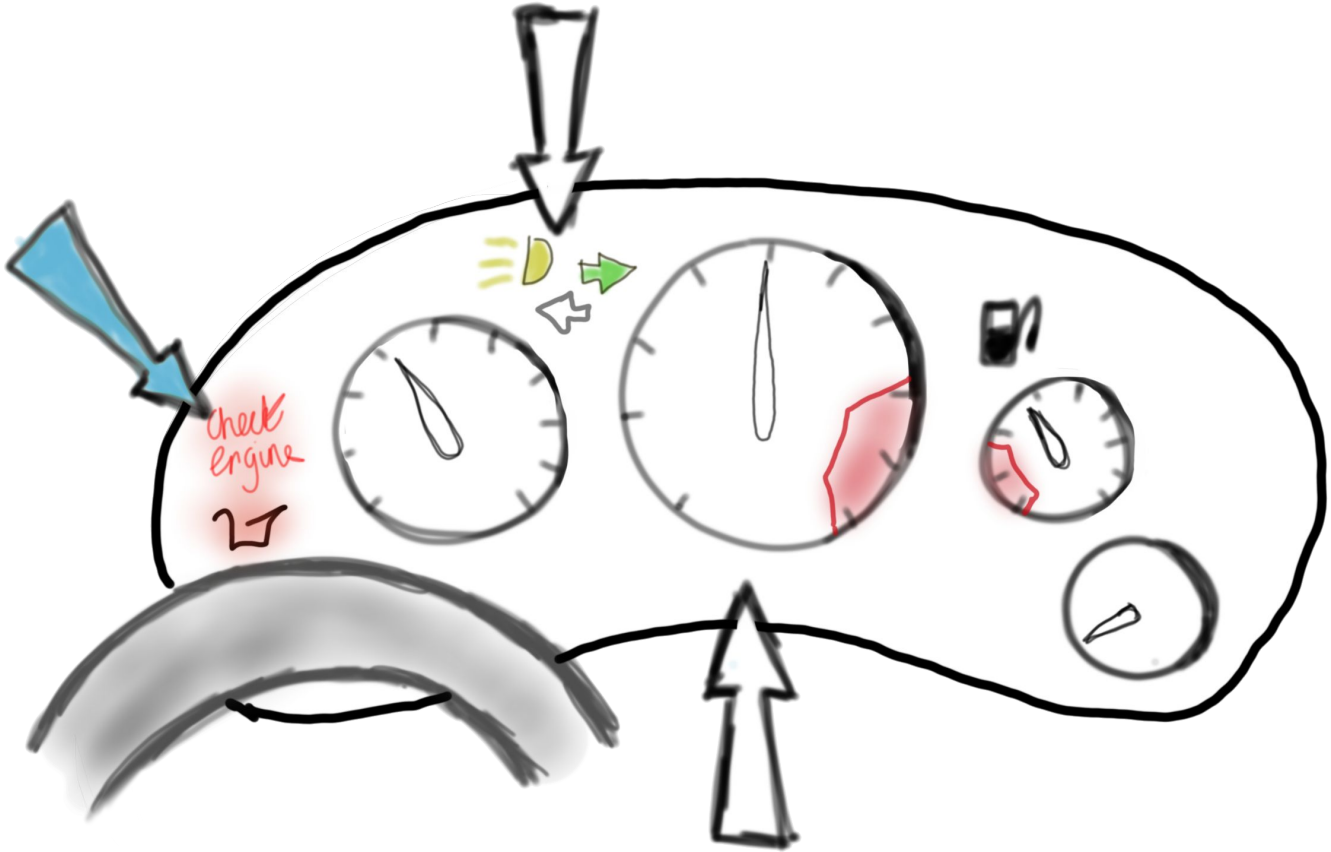
check
engine

Metrics VS Diagnostics

- High level view
- For assessing a system and spotting issues
- Kept in present view

- Log analytics
- For tracing and debugging
- Used when fixing issues

Suggested Baselines

Services

U.S.E

Utilisation

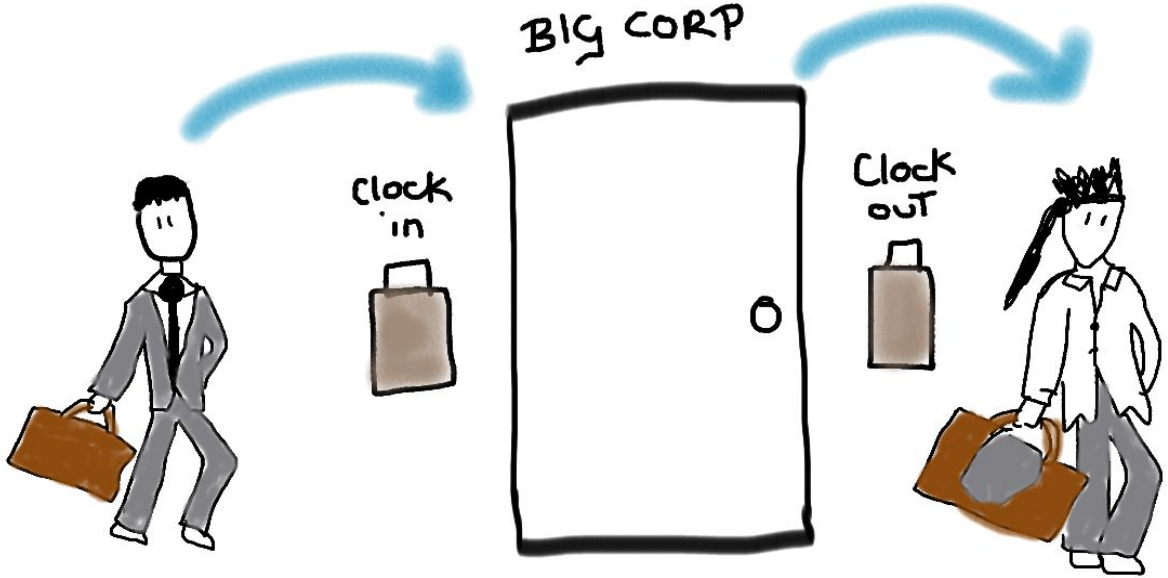

Saturation

Error

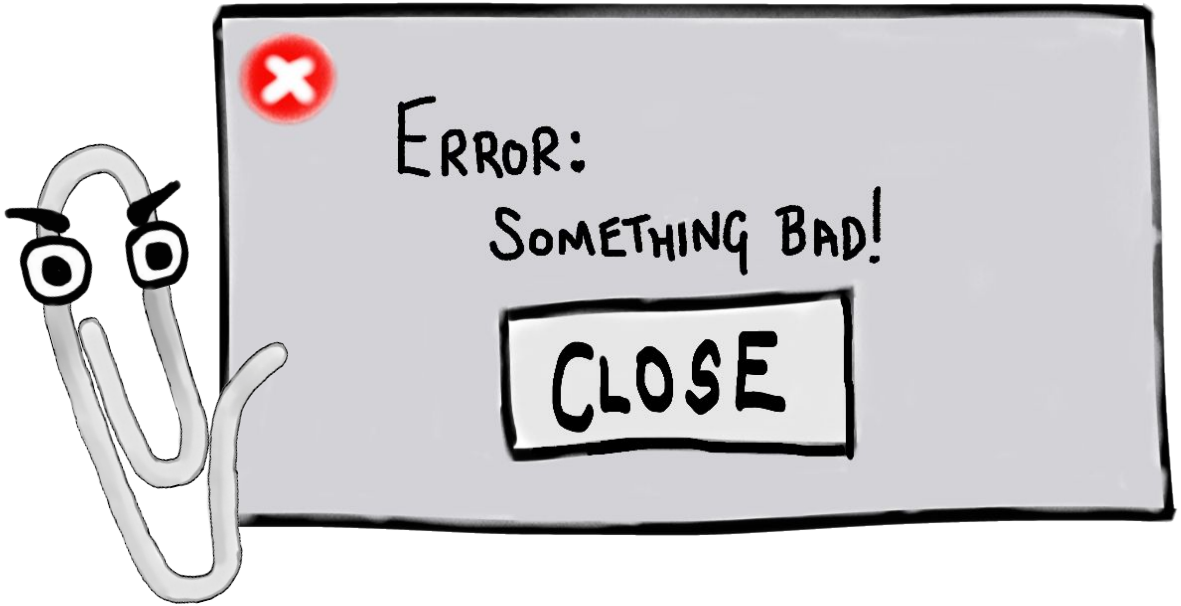

APIs

Four Golden Signals - Google

Traffic

Error

Latency

Saturation

Four Golden Signals - Google

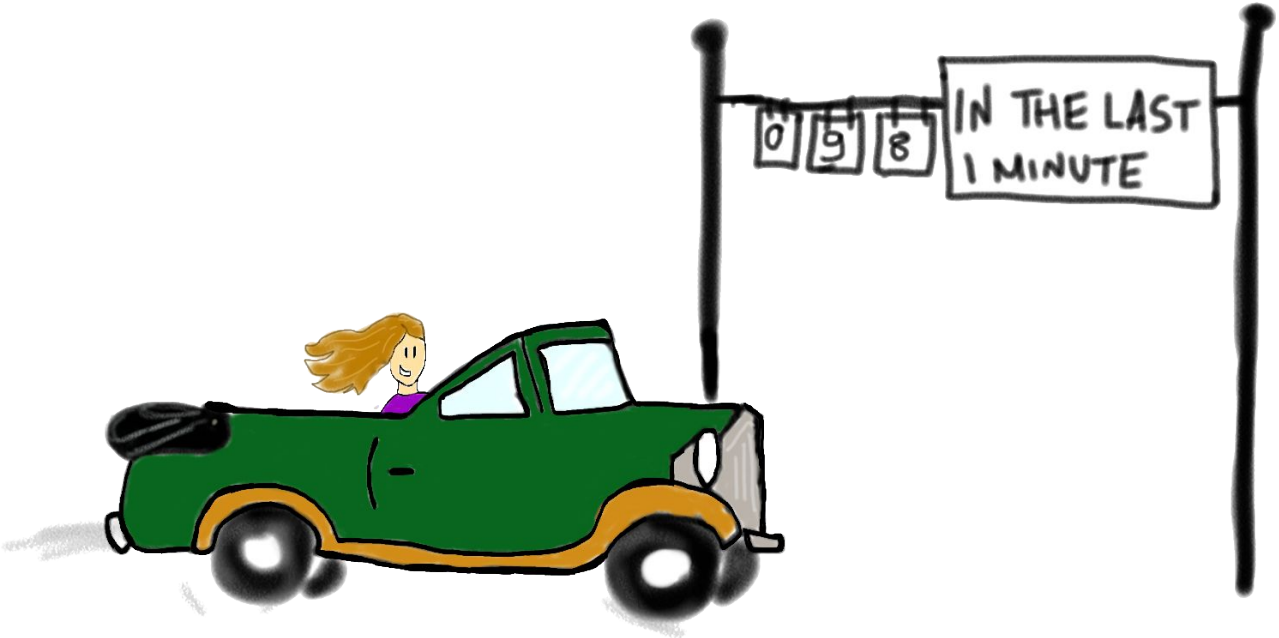

Traffic

Four Golden Signals - Google

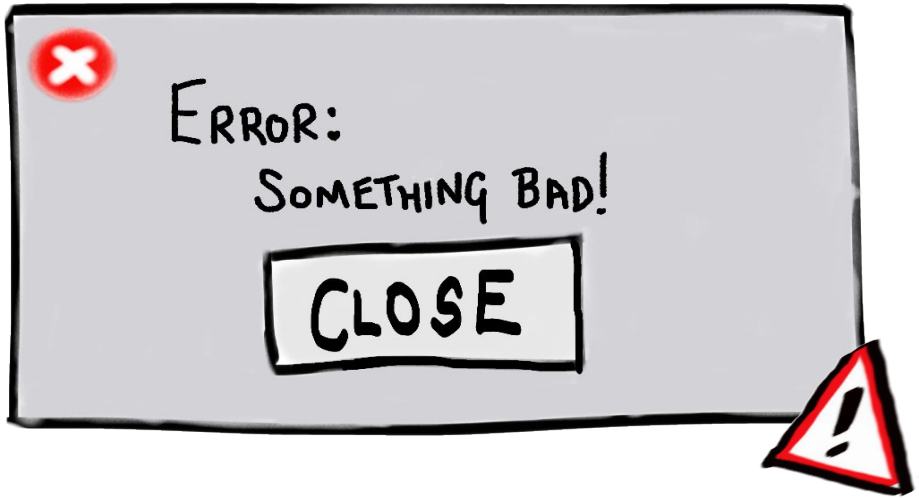

Error

Four Golden Signals - Google

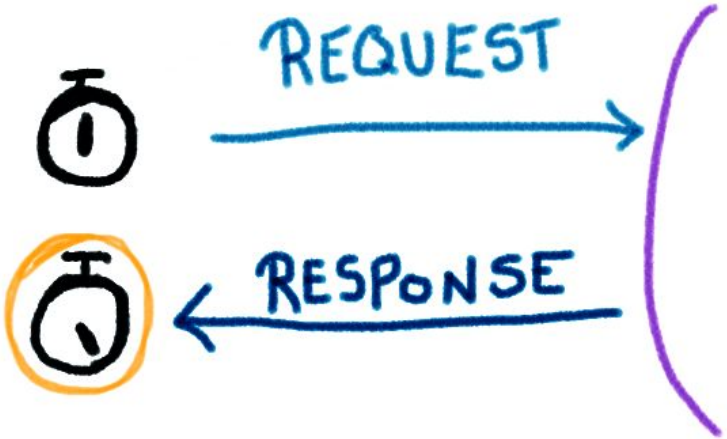

Latency

Four Golden Signals - Google

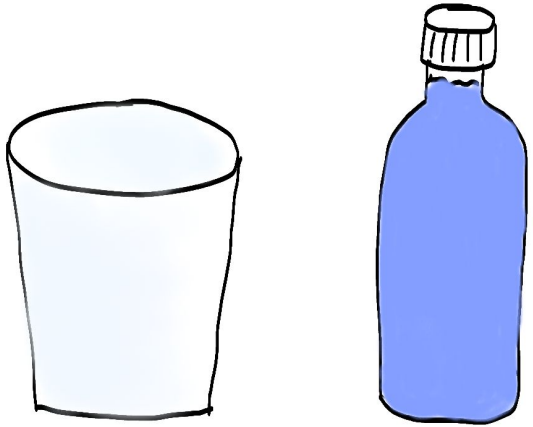

Saturation

Four Golden Signals - Google

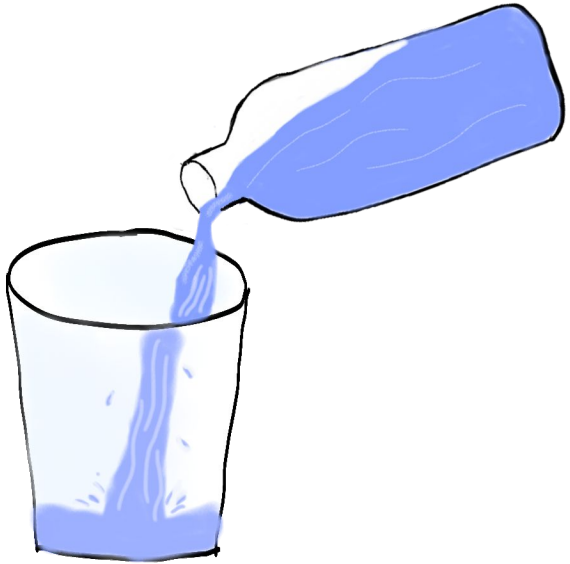

Saturation

Four Golden Signals - Google

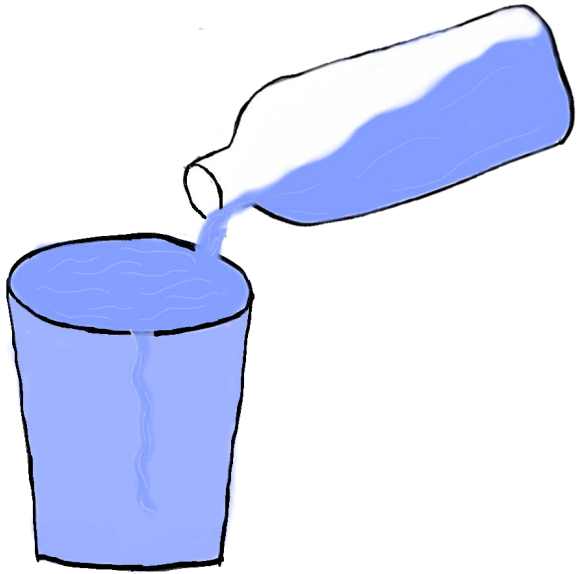

Saturation

Four Golden Signals - Google

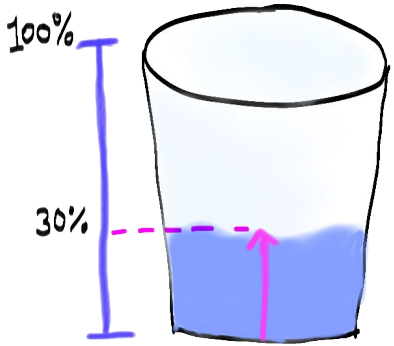

Saturation

R.E.D

Rate

Error

Duration

R.E.D

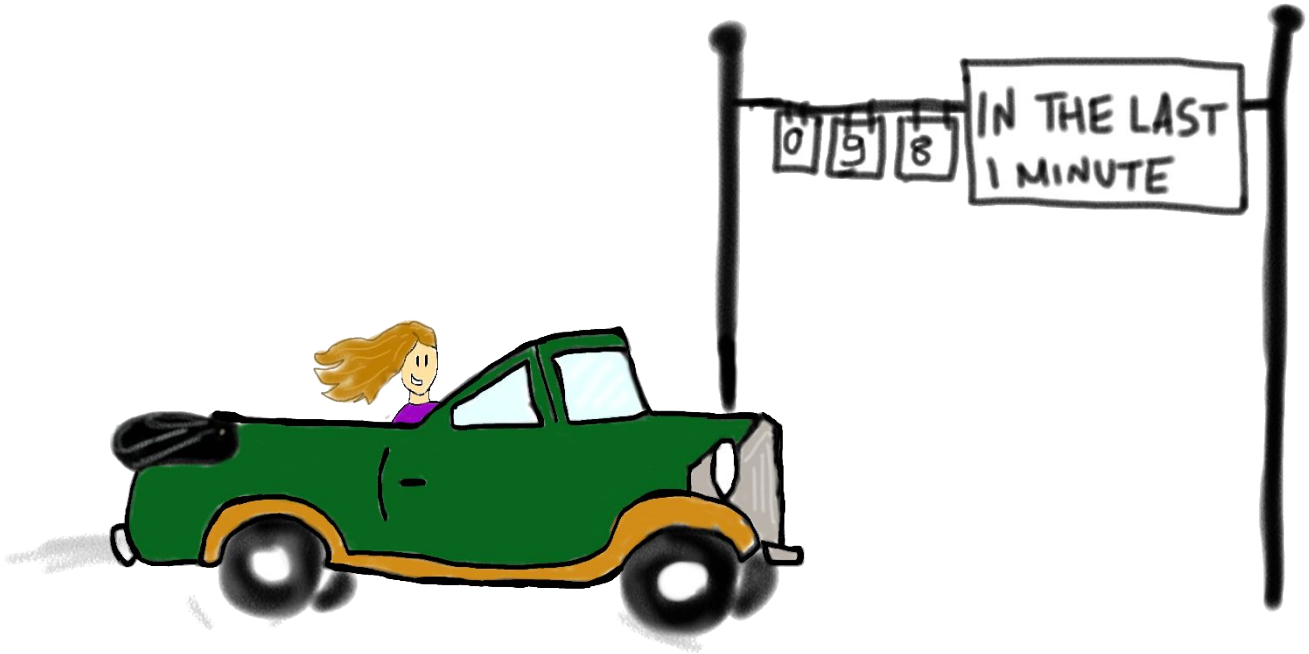

Rate

R.E.D

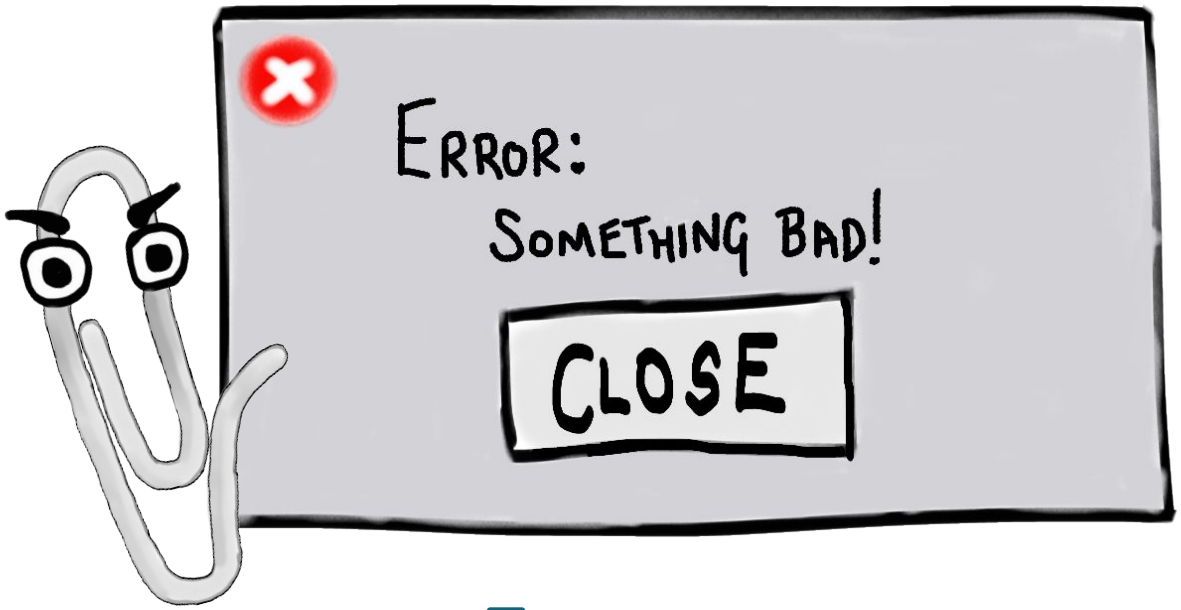

Error

R.E.D

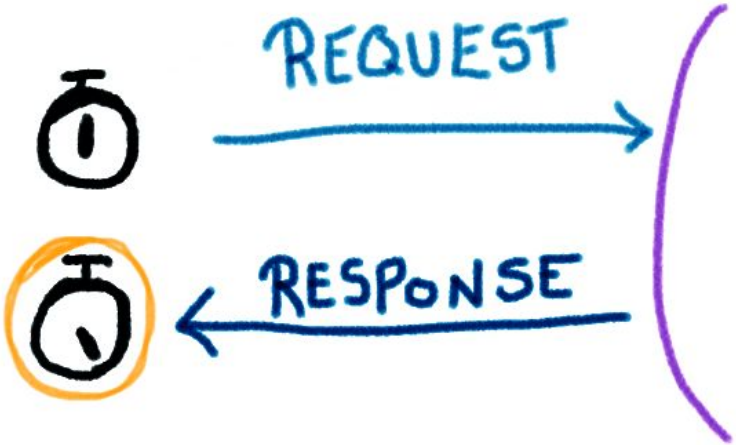

Duration

Rate	Traffic
Error	Error
Duration	Latency
	~~Saturation~~

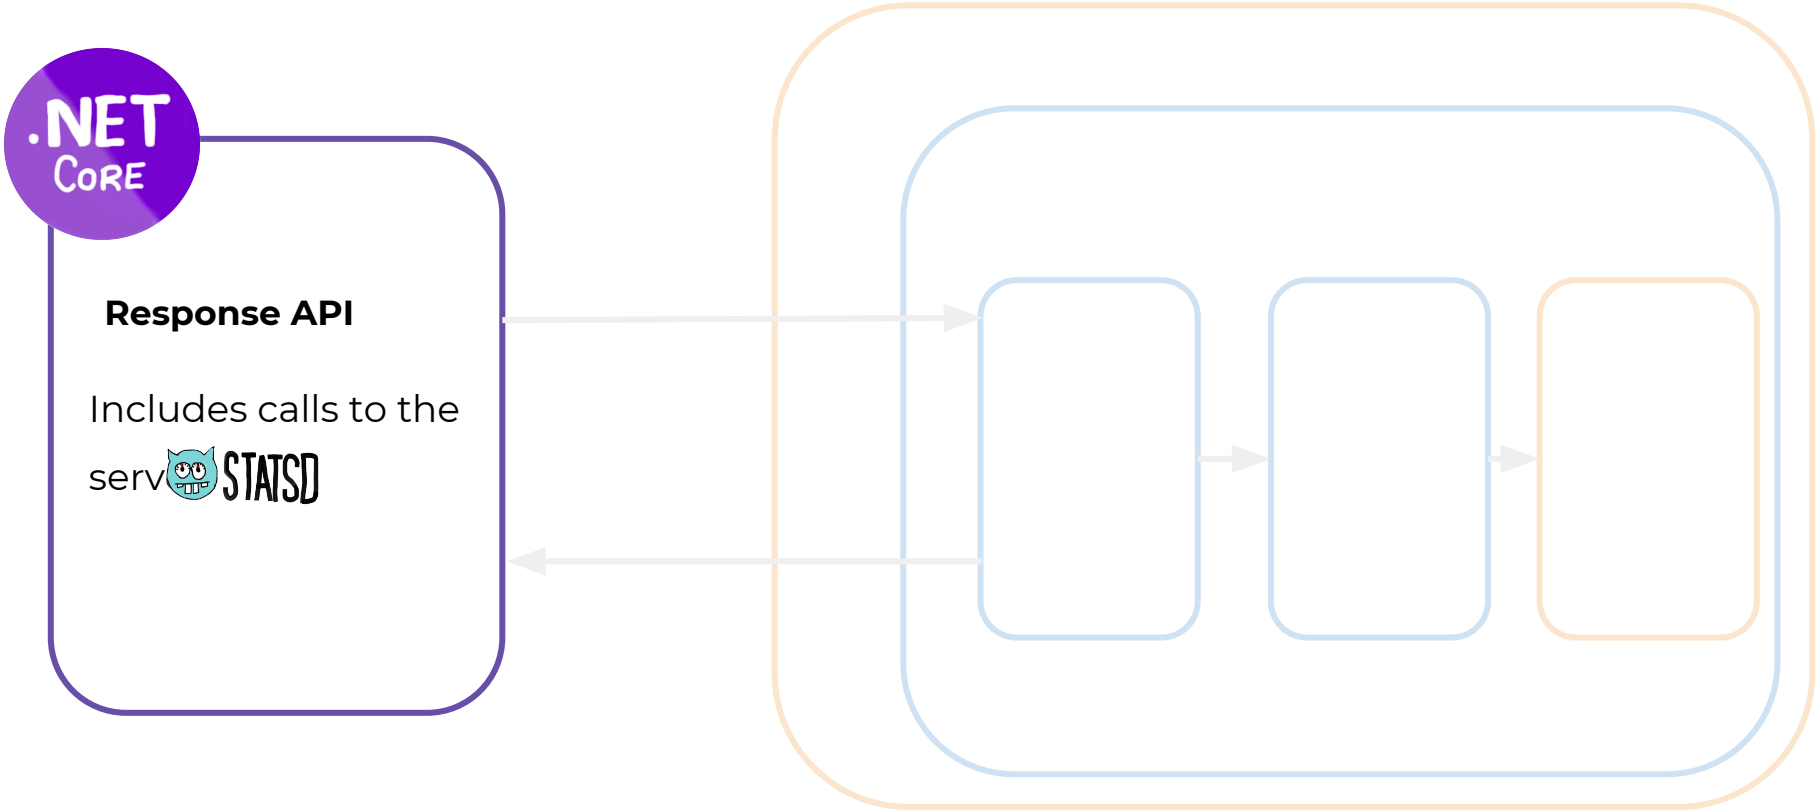
.NET
Core
Response API
Includes calls to the
serv
STATSD

46 lines (40 sloc) | 1.3 KB

Raw | Blame | History

```
using JustEat.StatsD;
using Microsoft.AspNetCore.Builder;
using Microsoft.AspNetCore.Hosting;
using Microsoft.Extensions.Configuration;
using Microsoft.Extensions.DependencyInjection;

namespace response_api
{
    public class Startup
    {
        public Startup(IConfiguration configuration)
        {
            Configuration = configuration;
        }

        public IConfiguration Configuration { get; }

        public void ConfigureServices(IServiceCollection services)
        {
            services.AddMvc();
        }

        public void Configure(IApplicationBuilder app, IHostingEnvironment env)
        {
            if (env.IsDevelopment())
            {
                app.UseDeveloperExceptionPage();
            }
            var statsDConfig = new StatsDConfiguration{
                Host = "graphite"
            };
            var statsDPublisher = new StatsDPublisher(statsDConfig);

            app.Use(async (context, next) =>
            {
                using (var timer = statsDPublisher.StartTimer("thing"))
                {
                    await next();
                    timer.StatName = $"response-api.code.{context.Response.StatusCode}";
                }
            });

            app.UseMvc();
        }
    }
}
```

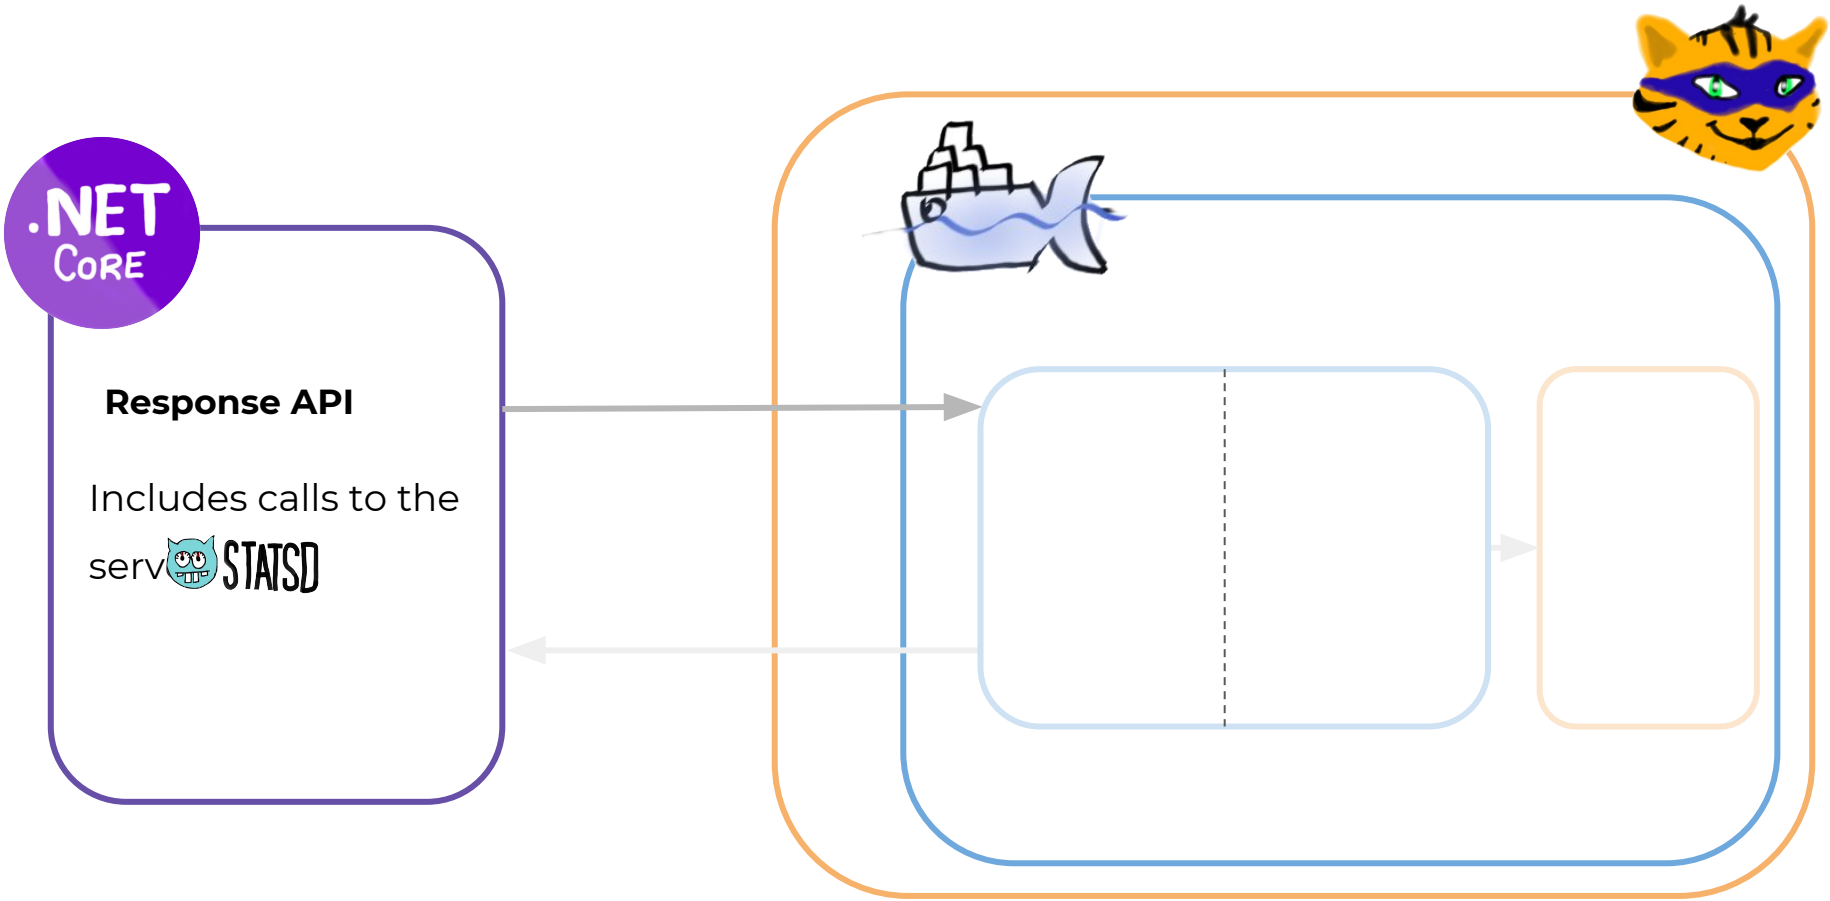
.NET
CORE
Response API
Includes calls to the
serv
STATSD

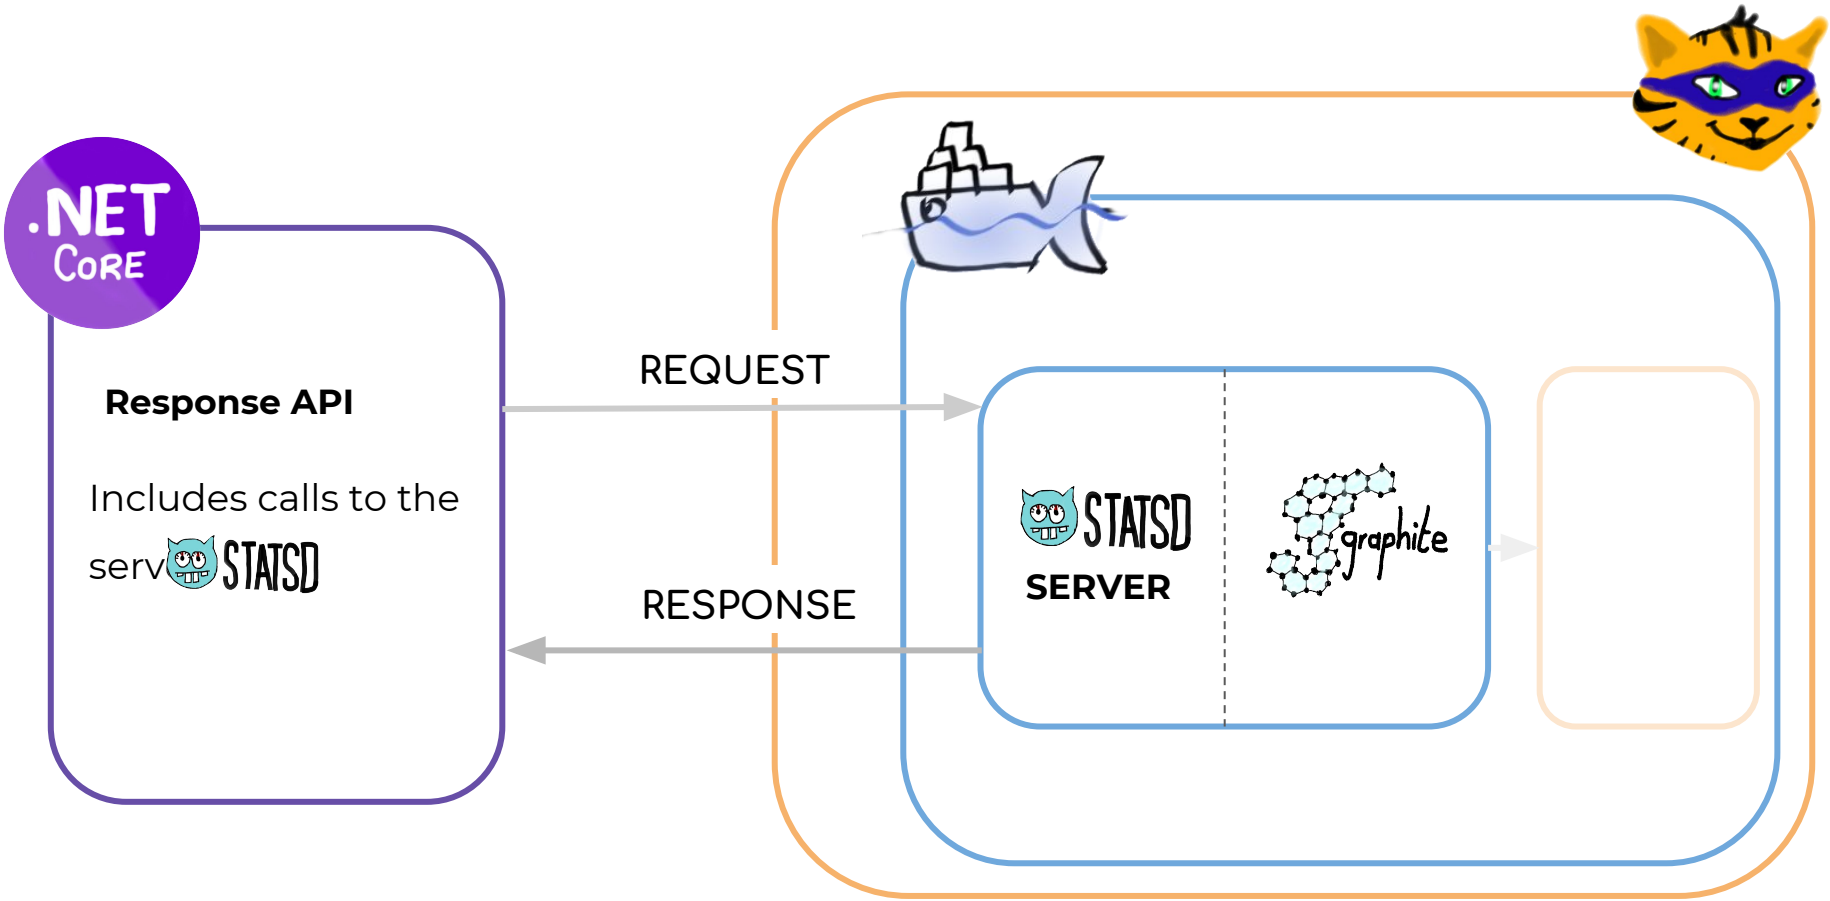
.NET
CORE
Response API
Includes calls to the
serv STATSD
REQUEST
RESPONSE
STATSD
SERVER
graphite

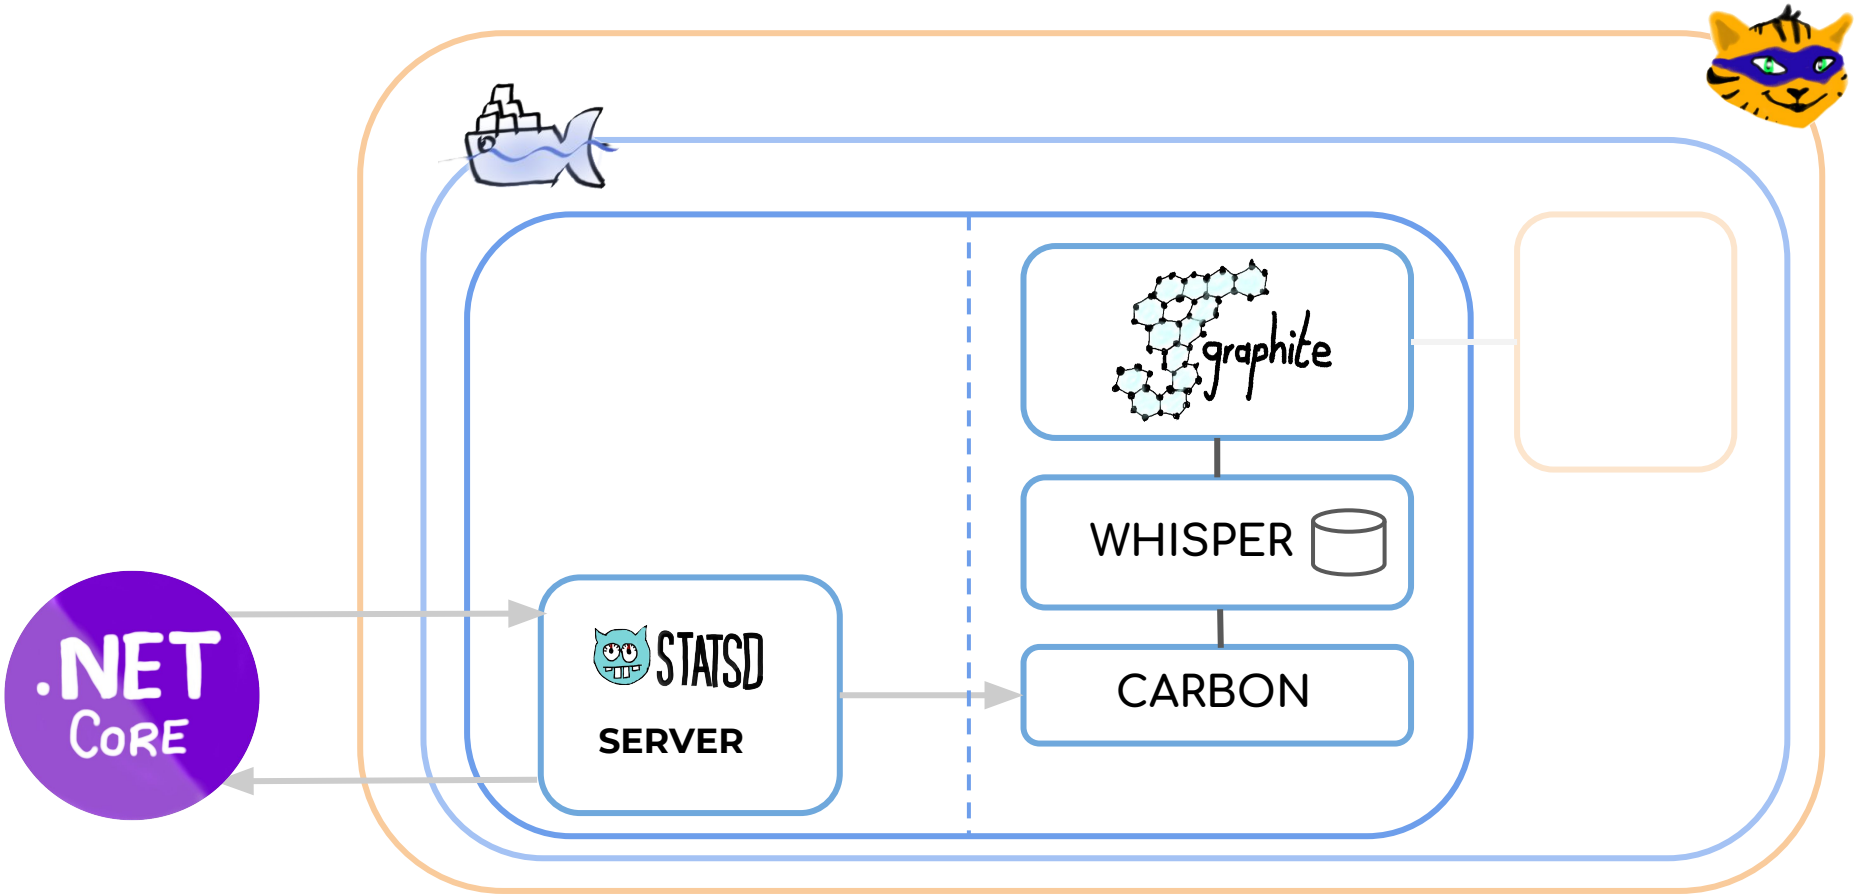
.NET
CORE
STATSD
SERVER
graphite
WHISPER
CARBON

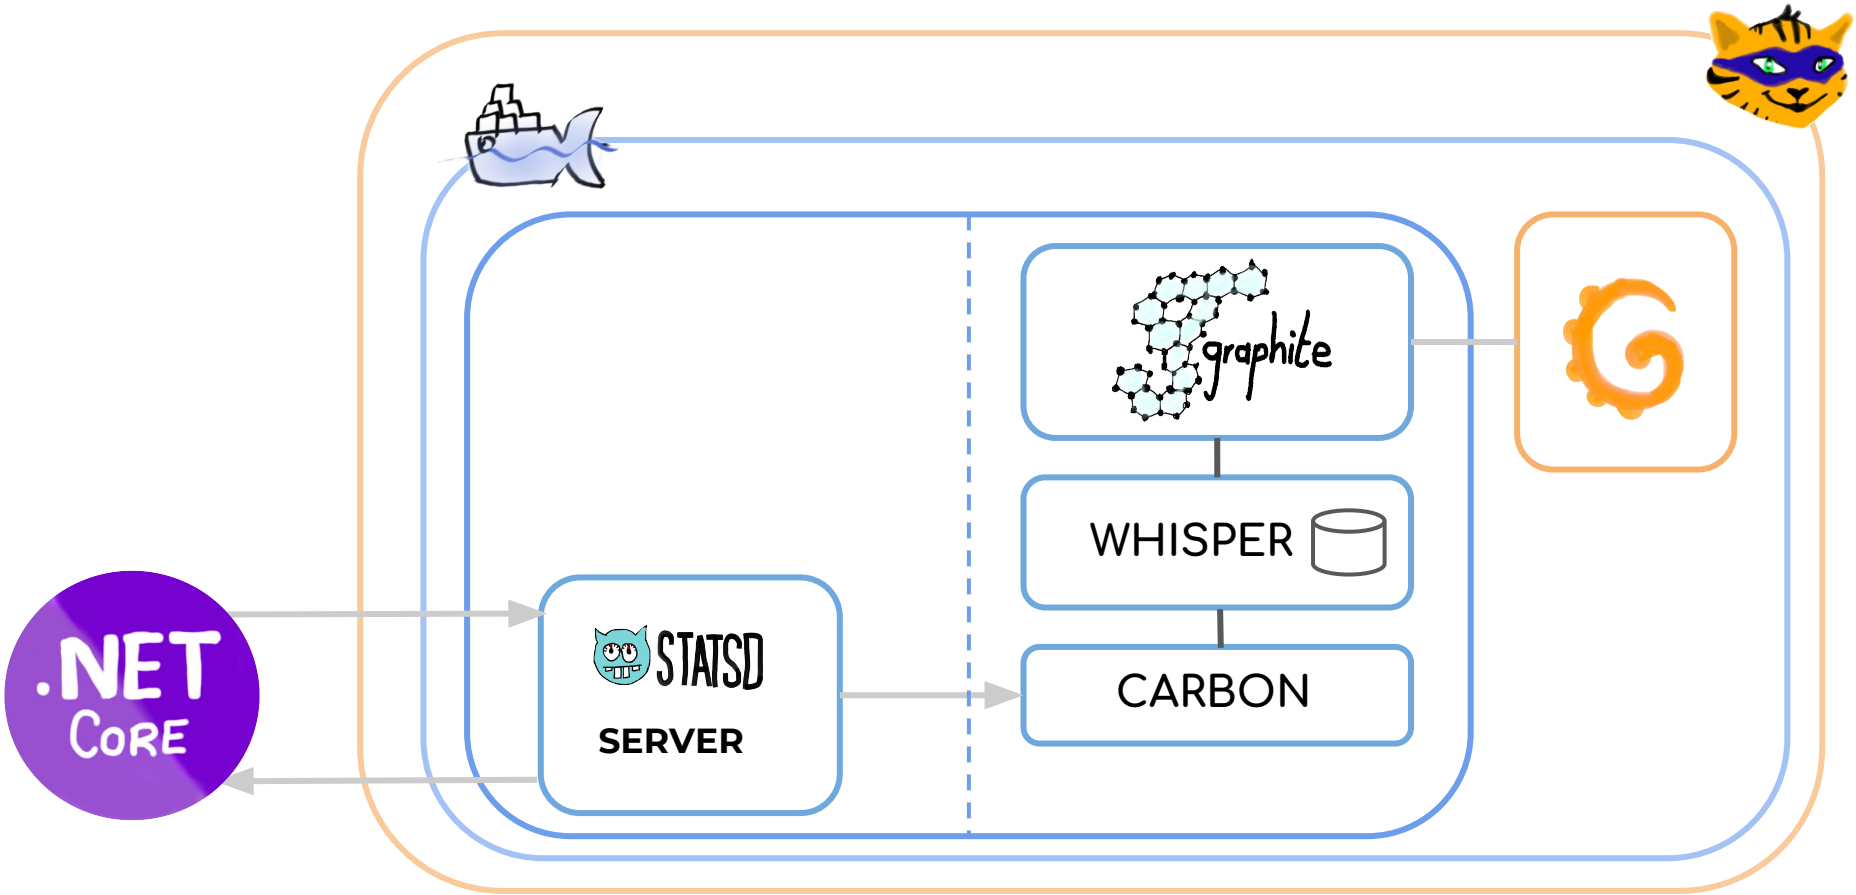
.NET
Core
STATSD
SERVER
graphite
WHISPER
CARBON

Demo

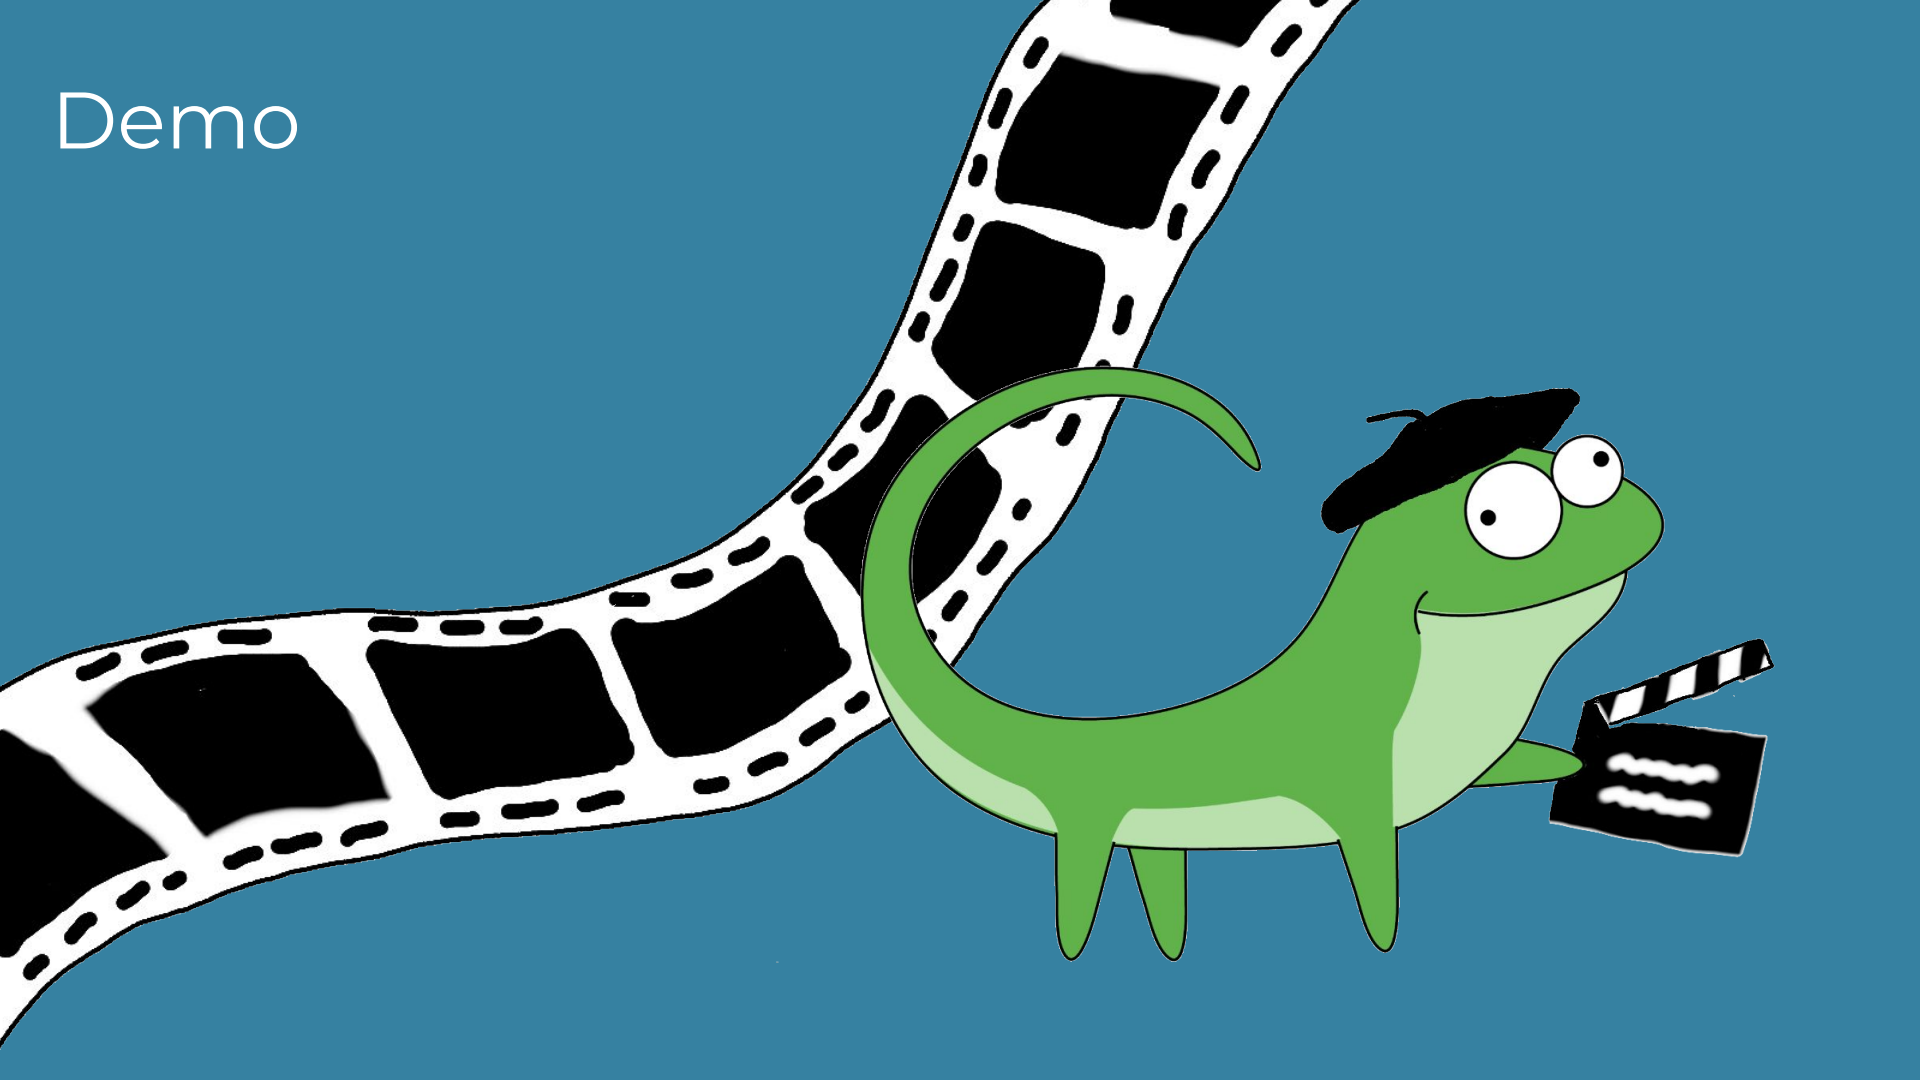

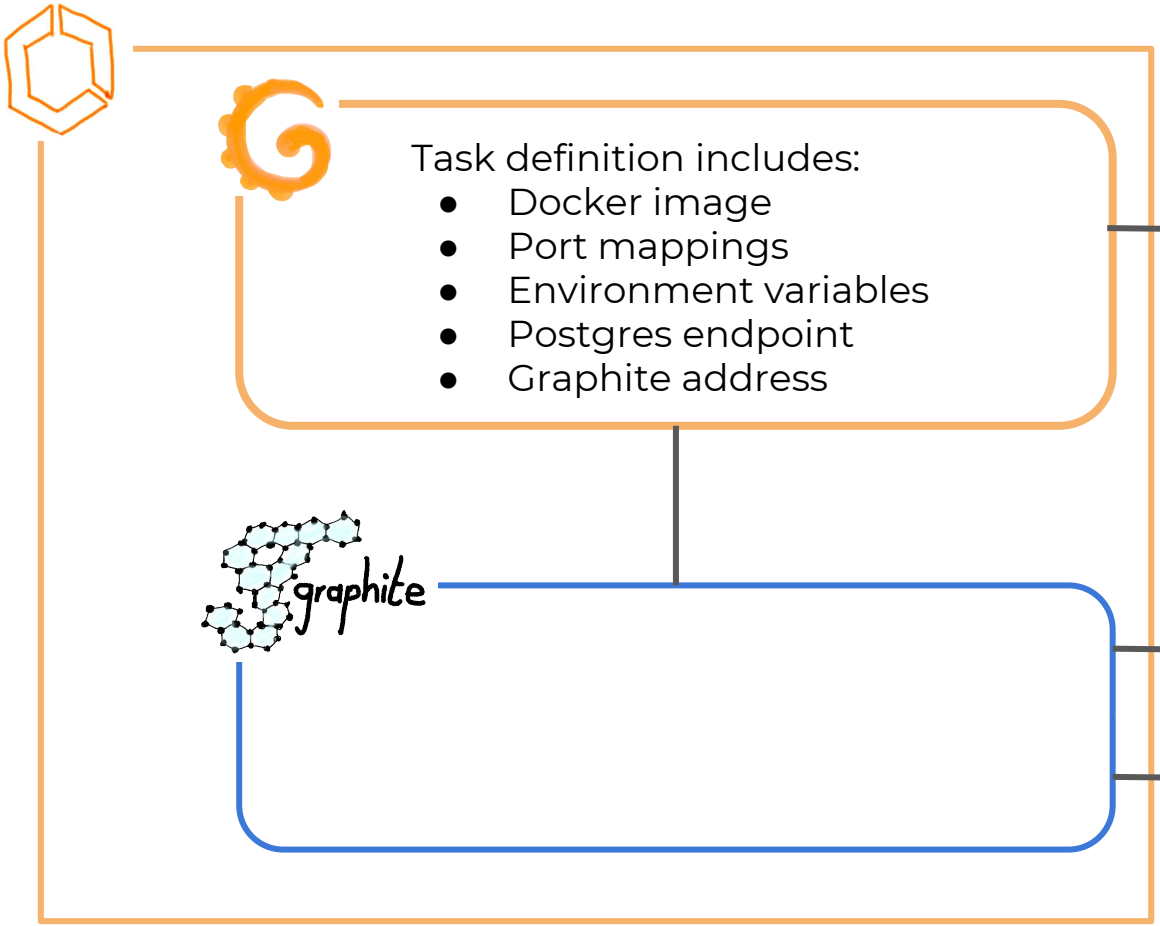
Task definition includes:
Docker image
Port mappings
Environment variables
Postgres endpoint
Graphite address
graphite

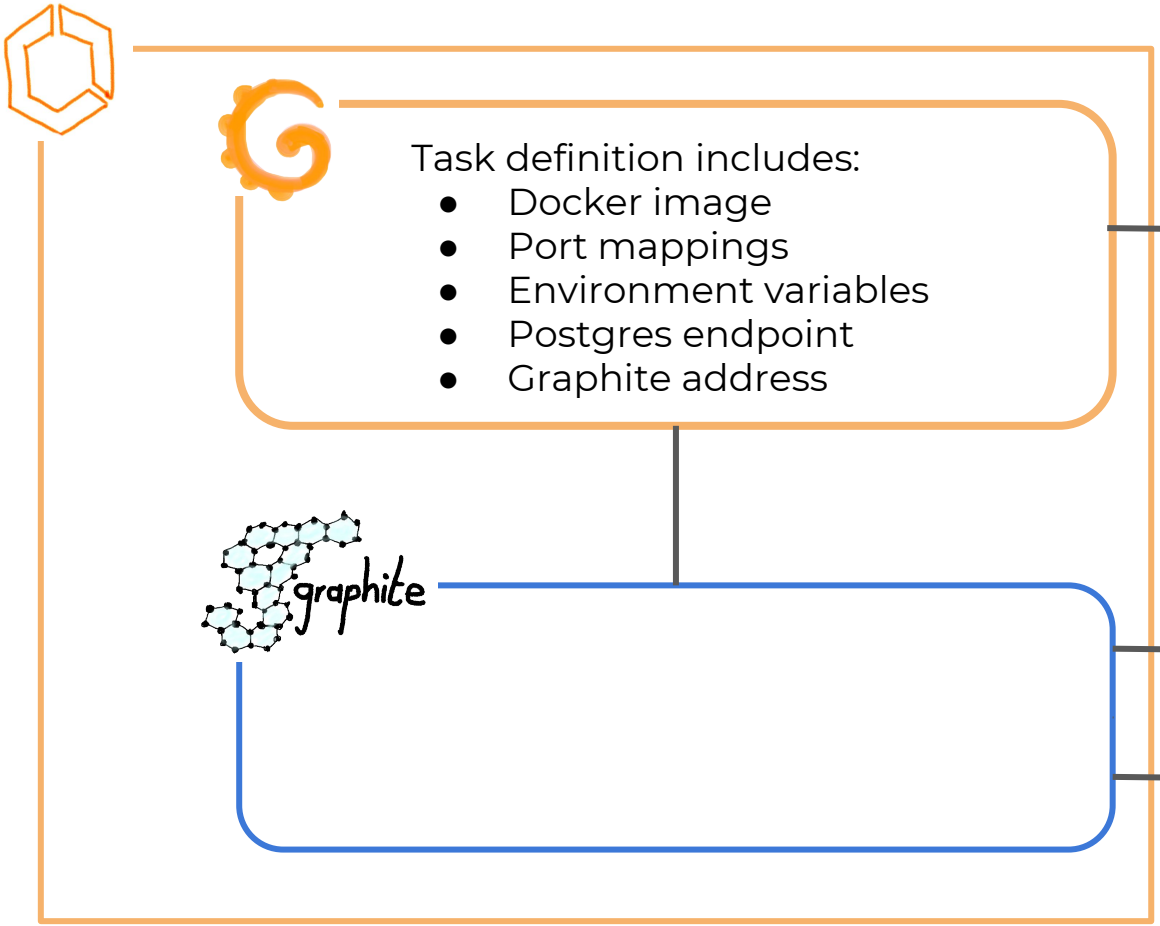
Task definition includes:
Docker image
Port mappings
Environment variables
Postgres endpoint
Graphite address
graphite
Postgres
Whisper
STATSD
Config

Do these metrics help you in real life?

Overview
00:00:14 to Dec 10, 2018 16:17:25 UTC
Region
eu-west-1
Aggregation
10m
Invocations
95.3 K
Errors
78
Duration vs. Throttles
Invocations
13.48 K
Errors
0
Duration vs. Throttles
Invocations
13.47 K
Errors
0
Duration vs. Throttles
Invocations
13.47 K
Errors
0
Duration vs. Throttles
Invocations
5.54 K
Errors
13
Duration vs. Throttles
Invocations
11.98 K
Errors
0
Duration vs. Throttles
04:00
08:00
12:00
16:00
30 s
20 s
10 s
0 ms

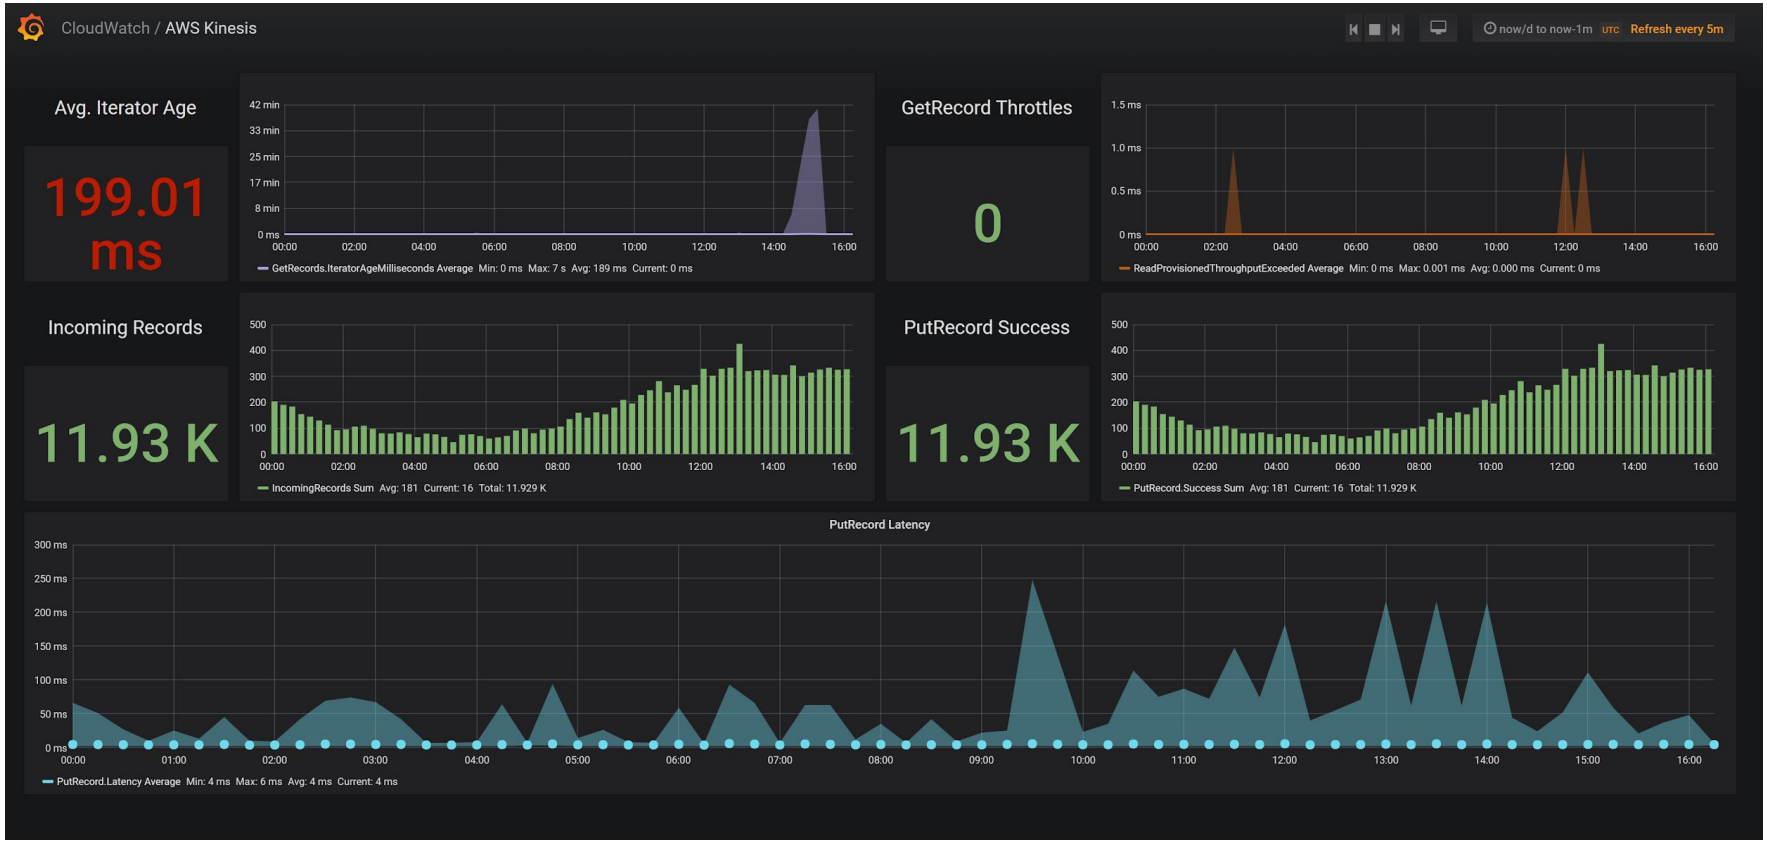
CloudWatch / AWS Kinesis
now/d to now-1m UTC Refresh every 5m
Avg. Iterator Age
199.01 ms
42 min
33 min
25 min
17 min
8 min
0 ms
00:00 02:00 04:00 06:00 08:00 10:00 12:00 14:00 16:00
GetRecords.IteratorAgeMilliseconds Average Min: 0 ms Max: 7 s Avg: 189 ms Current: 0 ms
GetRecord Throttles
0
1.5 ms
1.0 ms
0.5 ms
0 ms
ReadProvisionedThroughputExceeded Average Min: 0 ms Max: 0.001 ms Avg: 0.000 ms Current: 0 ms
Incoming Records
11.93 K
500
400
300
200
100
0
IncomingRecords Sum Avg: 181 Current: 16 Total: 11.929 K
PutRecord Success
11.93 K
PutRecord.Success Sum Avg: 181 Current: 16 Total: 11.929 K
PutRecord Latency
300 ms
250 ms
200 ms
150 ms
100 ms
50 ms
0 ms
00:00 01:00 02:00 03:00 04:00 05:00 06:00 07:00 08:00 09:00 10:00 11:00 12:00 13:00 14:00 15:00 16:00
PutRecord.Latency Average Min: 4 ms Max: 6 ms Avg: 4 ms Current: 4 ms

Where Can I Learn More?

These slides and resources will be made available on my be.notist page after this session:

https://noti.st/jesspwhite

The link to which will be shared on Twitter @JessPWhite

Tom Wilkie -The Red Method: How To Instrument Your Services

- Talk from Influx Days
- Slides From GrafanaCon EU

Brendan Greg -The USE Method

- Article on USE Method

Google SRE Book

- Book available online
- Other google SRE resources

Extra

- I also occasionally blog

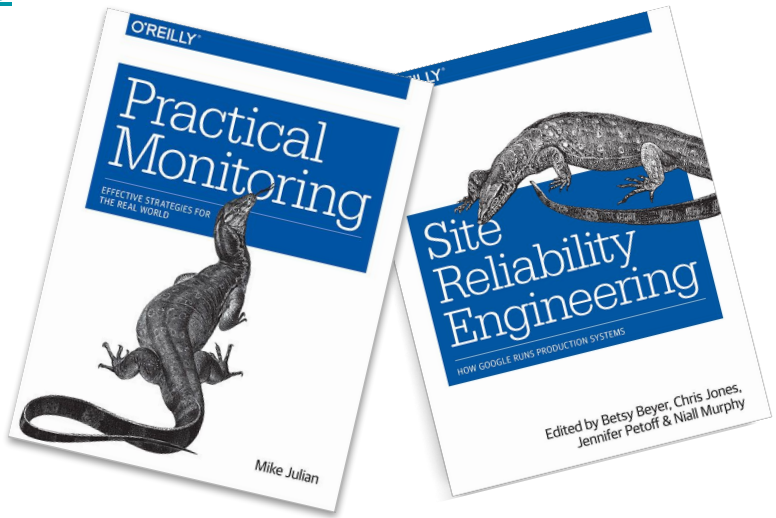

HEY! TECH CHATTER...
MORE EQUALLY AWESOME CHATTER!
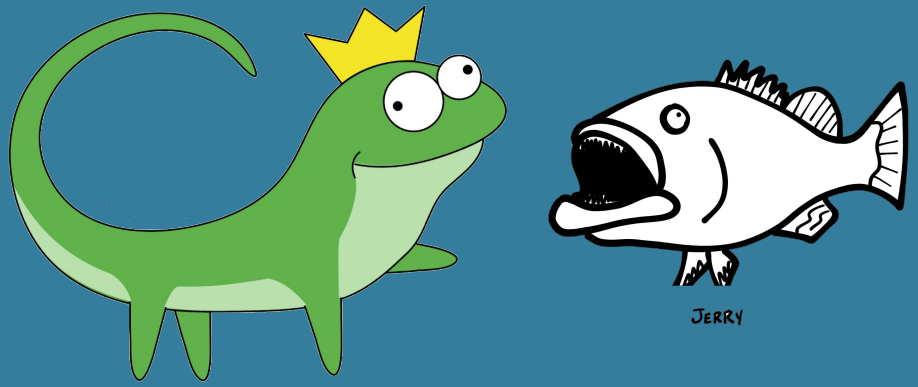
Thank you for listening