

Bla Bla Car

Backend High Availability

OPEN R&DAY { } #1



Maxime Fouilleul

Data Architect

dude@blablacar.com


Agenda



| Facts & Figures



| 100% rkt powered carpooling



| MySQL as example

Facts and figures



40 million members



21 million mobile app downloaded
(iPhone + Android)



12 million travelers per quarter



Founded in 2006



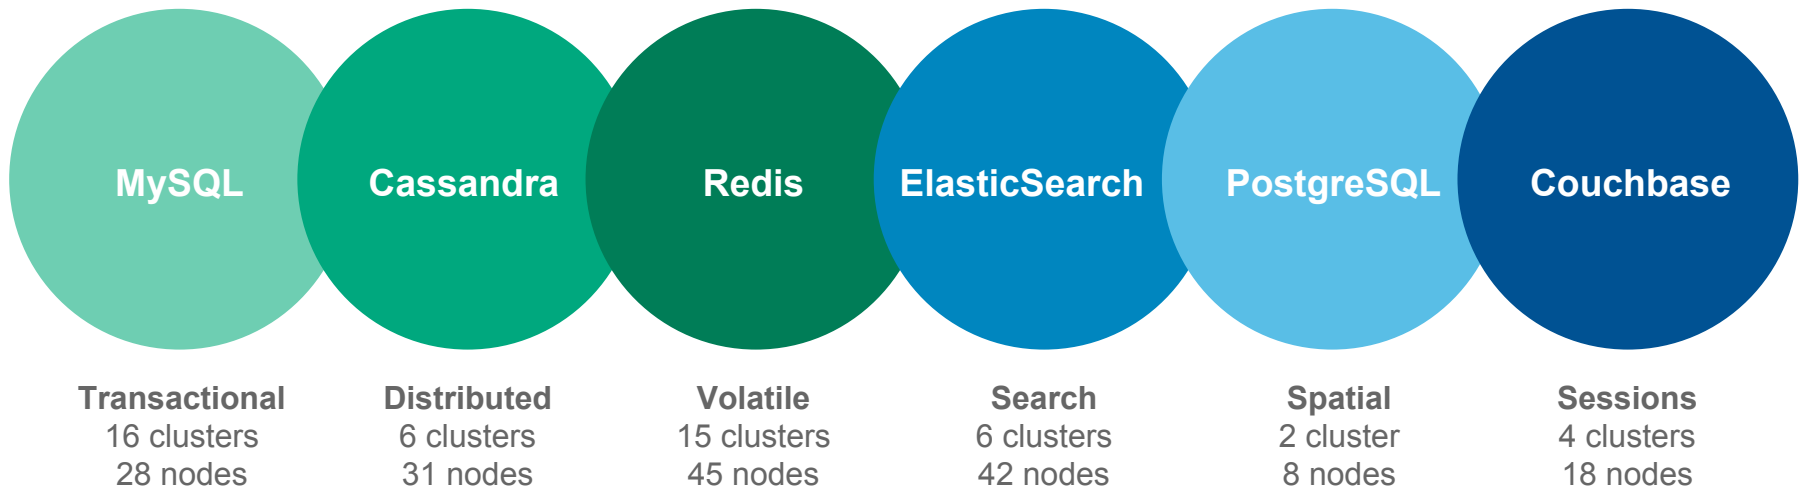
1 million tonnes less CO2 in the past year




Currently in 22 countries

France, Spain, UK, Italy, Poland, Hungary, Croatia, Serbia, Romania, Germany, Belgium, India, Mexico, The Netherlands, Luxembourg, Portugal, Ukraine, Czech Republic, Slovakia, Russia, Brazil and Turkey.

Our prod data ecosystem

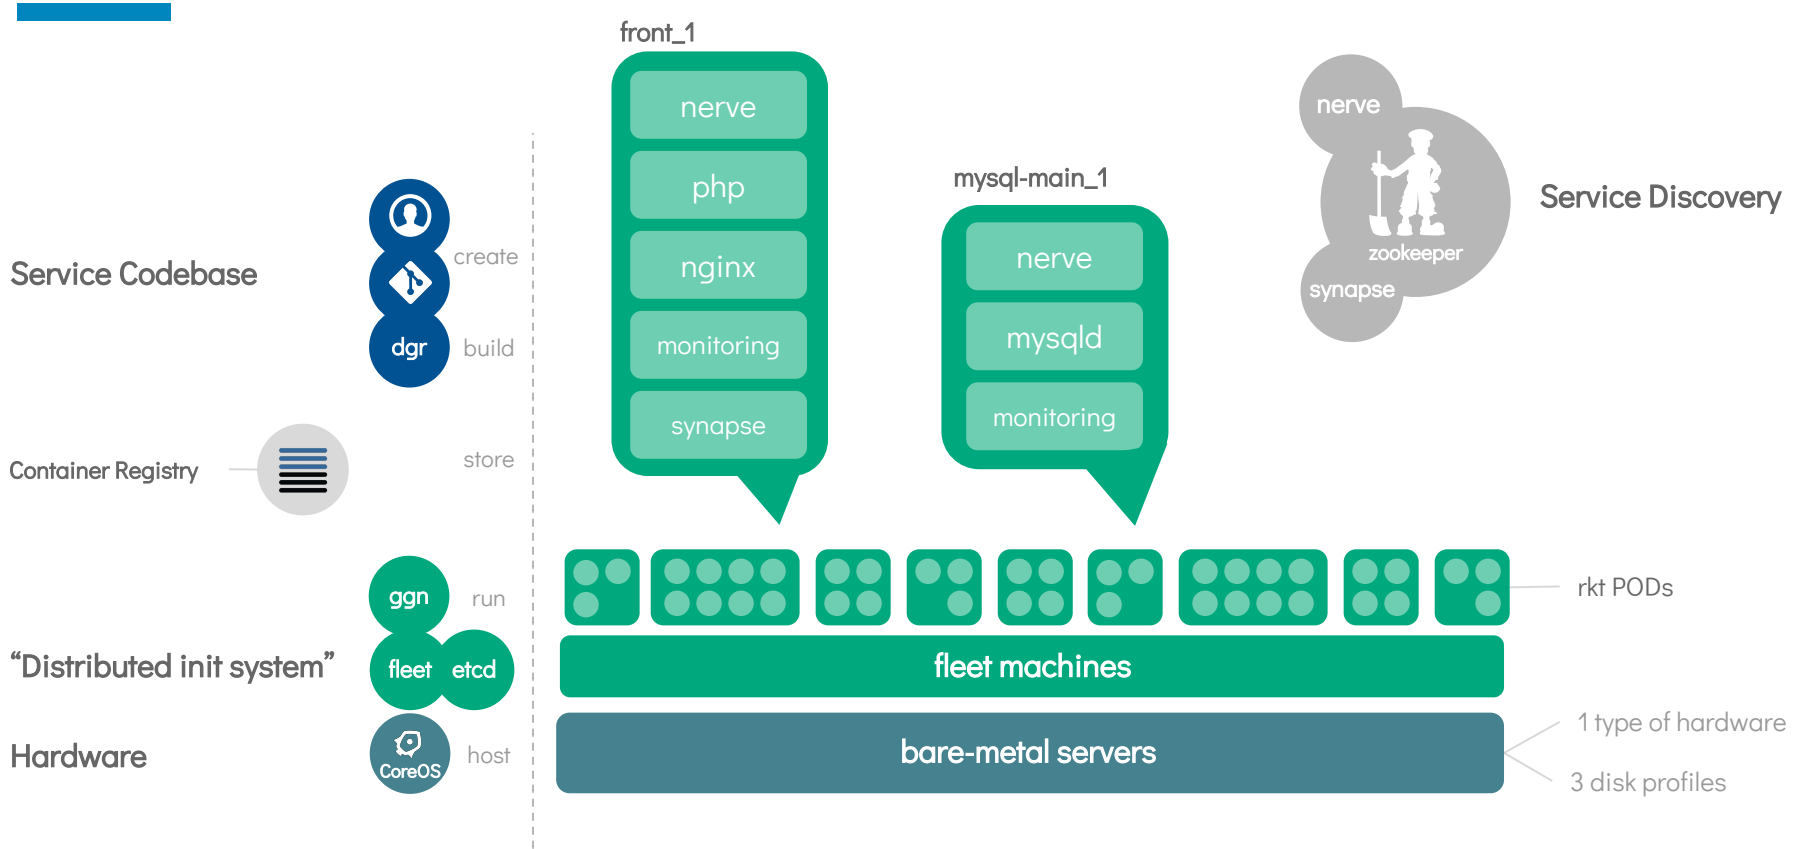




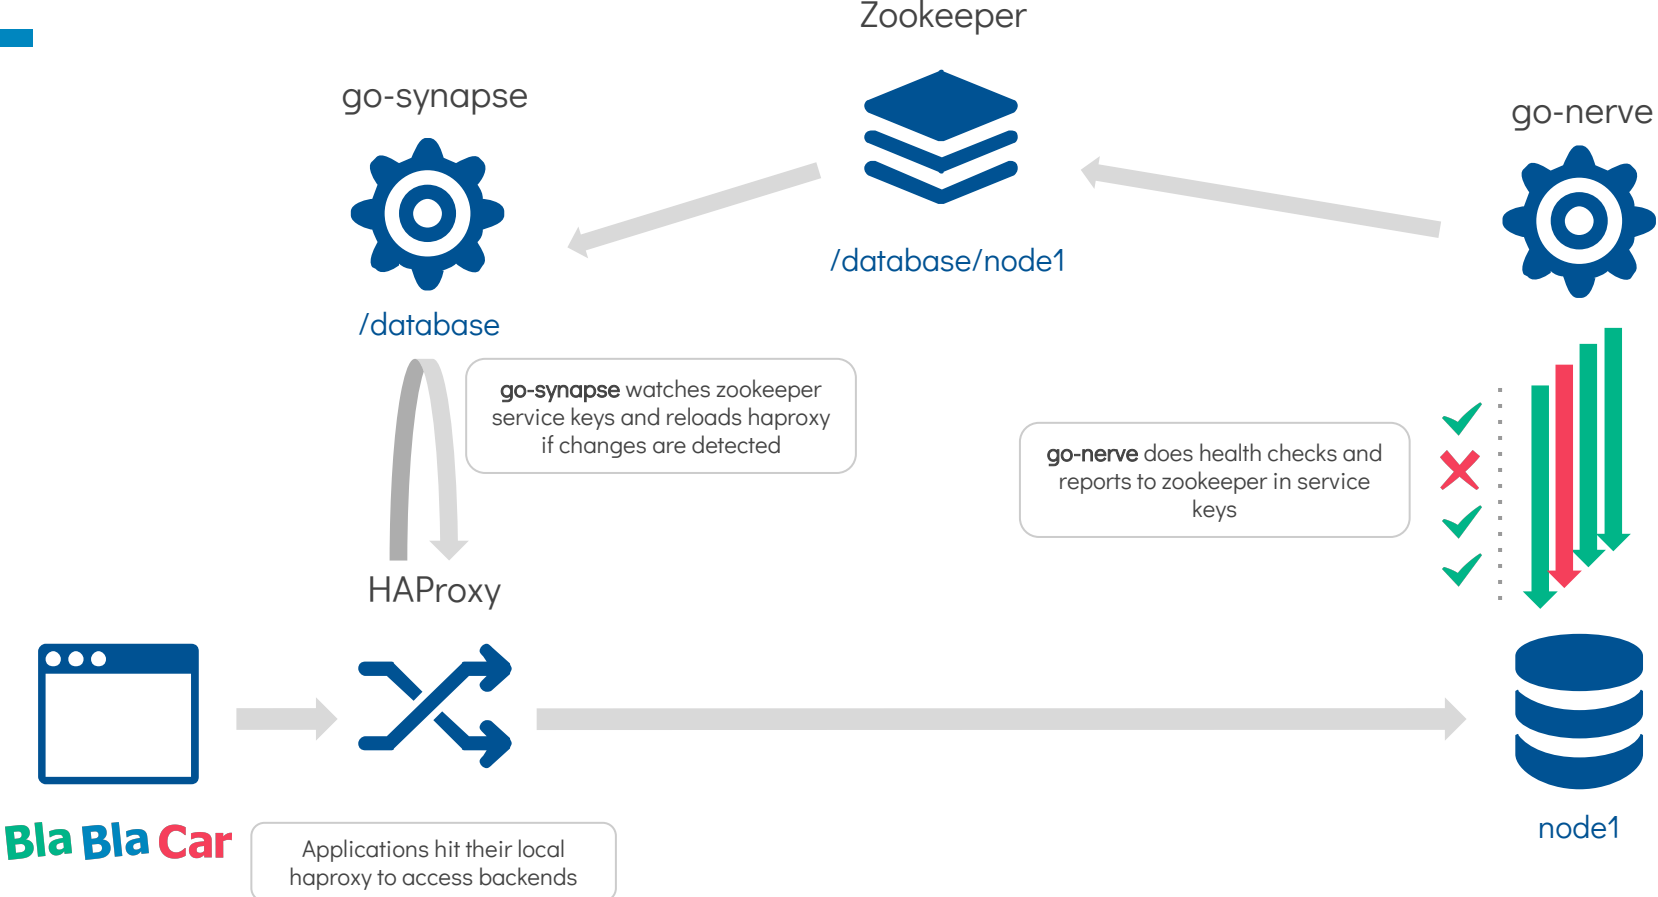
Infrastructure Ecosystem

100% rkt powered carpooling

Our infrastructure ecosystem



Service Discovery





Backend High Availability Pillars

MySQL as example



Abolish Slavery

Every node is the same


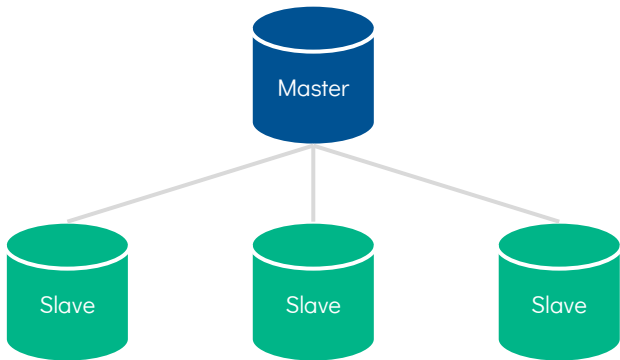
MySQL at BlaBlacar

Historical/main database

We use **sync** and **async** replication

Sync replication is ensured by **MariaDB Galera**

Galera = no master + no rep lag



Nerve - Track and report service status

```
# cat env/prod-dc1/services/mysql-main/attributes/nerve.yml
---
override:
  nerve:
    services:
      - name: "main-read"
        port: 3306
        reporters:
          - {type: zookeeper, path: /services/mysql/main_read}
        checks:
          - type: sql
            driver: mysql
            datasource: "local_monitor:local_monitor@tcp(127.0.0.1:3306)/"
            enableCheckStableCommand: ["/root/report_slow_queries.sh"]
            disableGracefullyDoneCommand: ["/root/report_remaining_processes.sh"]

      - name: "main-write"
        port: 3306
        reporters:
          - {type: zookeeper, path: /services/mysql/main_write}
        checks:
          - type: sql
            driver: mysql
            datasource: "local_monitor:local_monitor@tcp(127.0.0.1:3306)/"
            haproxyServerOptions: "backup"
```

```
# zookeepercli -c lsr /services/mysql/main_read
mysql-main_read1_192.168.1.2_ba0f1f8b3
mysql-main_read2_192.168.1.3_734d63da
mysql-main_read3_192.168.1.4_dde45787

# zookeepercli -c get /services/mysql/mysql-main_read1_192.168.1.2_ba0f1f8b3
{
  "available":true,
  "host":"192.168.1.2",
  "port":3306,
  "name":"mysql-main1",
  "weight":255,
  "labels":{
    "host":"r10-srv4"
  }
}

# zookeepercli -c get /services/mysql/mysql-main_write1_192.168.1.2_ba0f1f8b3
{
  "available":true,
  "host":"192.168.1.2",
  "port":3306,
  "name":"mysql-main1",
  "haproxy_server_options":"backup",
  "weight":255,
  "labels":{
    "host":"r10-srv4"
  }
}
```

Synapse - Service discovery router

```
# cat env/prod-dc1/services/tripsearch/attributes/synapse.yml
```

```
---
```

```
override:
```

```
synapse:
```

```
services:
```

```
- name: mysql-main_read
```

```
  path: /services/mysql/main_read
```

```
  port: 3307
```

```
  serverCorrelation:
```

```
    type: excludeServer
```

```
    otherServiceName: mysql-main_write
```

```
    scope: first
```

```
- name: mysql-main_write
```

```
  path: /services/mysql/main_write
```

```
  port: 3308
```

```
  serverSort: date
```

```
# cat env/prod-dc1/services/tripsearch/attributes/tripsearch.yml
```

```
---
```

```
override:
```

```
tripsearch:
```

```
  database:
```

```
    read:
```

```
      host: localhost
```

```
      database: tripsearch
```

```
      user: tripsearch_rd
```

```
      port: 3307
```

```
    write:
```

```
      host: localhost
```

```
      database: tripsearch
```

```
      user: tripsearch_wr
```

```
      port: 3308
```

Synapse - haproxy

What is looks like in the local haproxy

service_mysql-main_read

| | Queue | | | Session rate | | | Sessions | | | | | |
|-------------|-------|-----|-------|--------------|-----|-------|----------|-----|-------|-------|-------|------|
| | Cur | Max | Limit | Cur | Max | Limit | Cur | Max | Limit | Total | LbTot | Last |
| mysql-main2 | 0 | 0 | - | 0 | 0 | | 0 | 0 | - | 0 | 0 | ? |
| mysql-main3 | 0 | 0 | - | 0 | 0 | | 0 | 0 | - | 0 | 0 | ? |
| Backend | 0 | 0 | | 0 | 1 | | 0 | 1 | 200 | 1 | 1 | 2s |

service_mysql-main_write

| | Queue | | | Session rate | | | Sessions | | | | | |
|-------------|-------|-----|-------|--------------|-----|-------|----------|-----|-------|-------|-------|------|
| | Cur | Max | Limit | Cur | Max | Limit | Cur | Max | Limit | Total | LbTot | Last |
| mysql-main1 | 0 | 0 | - | 0 | 1 | | 0 | 1 | - | 1 | 1 | 2s |
| mysql-main2 | 0 | 0 | - | 0 | 0 | | 0 | 0 | - | 0 | 0 | ? |
| mysql-main3 | 0 | 0 | - | 0 | 0 | | 0 | 0 | - | 0 | 0 | ? |
| Backend | 0 | 0 | | 0 | 1 | | 0 | 1 | 200 | 1 | 1 | 2s |



Be Quiet!

Come gently into prod

Manage the load on newcomers

```
# cat env/prod-dc1/services/mysql-main/attributes/nerve.yml
---
override:
  nerve:
    services:
      - name: "main-read"
        port: 3306
        reporters:
          - {type: zookeeper, path: /services/mysql/main_read}
        checks:
          - type: sql
            driver: mysql
            datasource: "local_monitor:local_monitor@tcp(127.0.0.1:3306)/"
            enableCheckStableCommand: ["/root/report_slow_queries.sh"]
```

Manage the load on newcomers

```
# bbc mysql prod-dc1 mysql-main overview
=== Service Overview 'prod-dc1 mysql-main' ===
mysql-main1 (192.168.1.2) PING, PORT, Synced
---
mysql-main1_read (3306) - enabled - weight = 255/255
mysql-main1_write (3306) - enabled - weight = 255/255

mysql-main2 (192.168.1.3) PING, PORT, Synced
---
mysql-main2_read (3306) - enabled - weight = 255/255
mysql-main2_write (3306) - enabled - weight = 255/255

mysql-main3 (192.168.1.4) PING, PORT, Synced
---
mysql-main3_read (3306) - enabled - weight = 255/255
mysql-main3_write (3306) - enabled - weight = 255/255
```

```
# bbc mysql prod-dc1 mysql-main mysql-main1 monitor
Weight: 0/255 Processes: 0 Slow: 0
Weight: 2/255 Processes: 0 Slow: 0
Weight: 3/255 Processes: 3 Slow: 0
Weight: 4/255 Processes: 7 Slow: 0
Weight: 6/255 Processes: 10 Slow: 0
Weight: 9/255 Processes: 12 Slow: 0
Weight: 15/255 Processes: 20 Slow: 1
Weight: 0/255 Processes: 20 Slow: 1
Weight: 2/255 Processes: 12 Slow: 0
Weight: 3/255 Processes: 4 Slow: 0
Weight: 4/255 Processes: 7 Slow: 0
Weight: 6/255 Processes: 10 Slow: 0
Weight: 9/255 Processes: 12 Slow: 0
Weight: 15/255 Processes: 20 Slow: 0
Weight: 23/255 Processes: 35 Slow: 0
Weight: 38/255 Processes: 40 Slow: 0
Weight: 38/255 Processes: 35 Slow: 0
Weight: 61/255 Processes: 36 Slow: 0
Weight: 61/255 Processes: 47 Slow: 0
Weight: 98/255 Processes: 44 Slow: 0
Weight: 98/255 Processes: 41 Slow: 0
Weight: 158/255 Processes: 38 Slow: 0
Weight: 158/255 Processes: 50 Slow: 0
Weight: 255/255 Processes: 46 Slow: 0
Weight: 255/255 Processes: 46 Slow: 0
```




Die in Peace...

Get out when you are ready

Gracefull shutdown

```
# cat env/prod-dc1/services/mysql-main/attributes/nerve.yml
---
override:
  nerve:
    services:
      - name: "main-read"
        port: 3306
        reporters:
          - {type: zookeeper, path: /services/mysql/main_read}
        checks:
          - type: sql
            driver: mysql
            datasource: "local_monitor:local_monitor@tcp(127.0.0.1:3306)/"
            enableCheckStableCommand: ["/root/report_slow_queries.sh"]
            disableGracefullyDoneCommand: ["/root/report_remaining_processes.sh"]
```

Backend High Availability Pillars



Abolish Slavery

Every node is the same

- 1 type of Hardware
- rkt Containers
- No master/slave
- Service Discovery (nerve/synapse)



Be Quiet!

Come gently into prod

- Graceful restart
- Service Discovery (nerve/synapse)
- Weight system
- Slow query tracking



Die in Peace...

Get out when you are ready

- Graceful restart
- Service Discovery (nerve/synapse)
- Weight system

Thanks!



blablatech.com
github.com/blablacar

OPEN R&DAY  #1