Using Prometheus and Grafana for PostgreSQL monitoring in TomTom

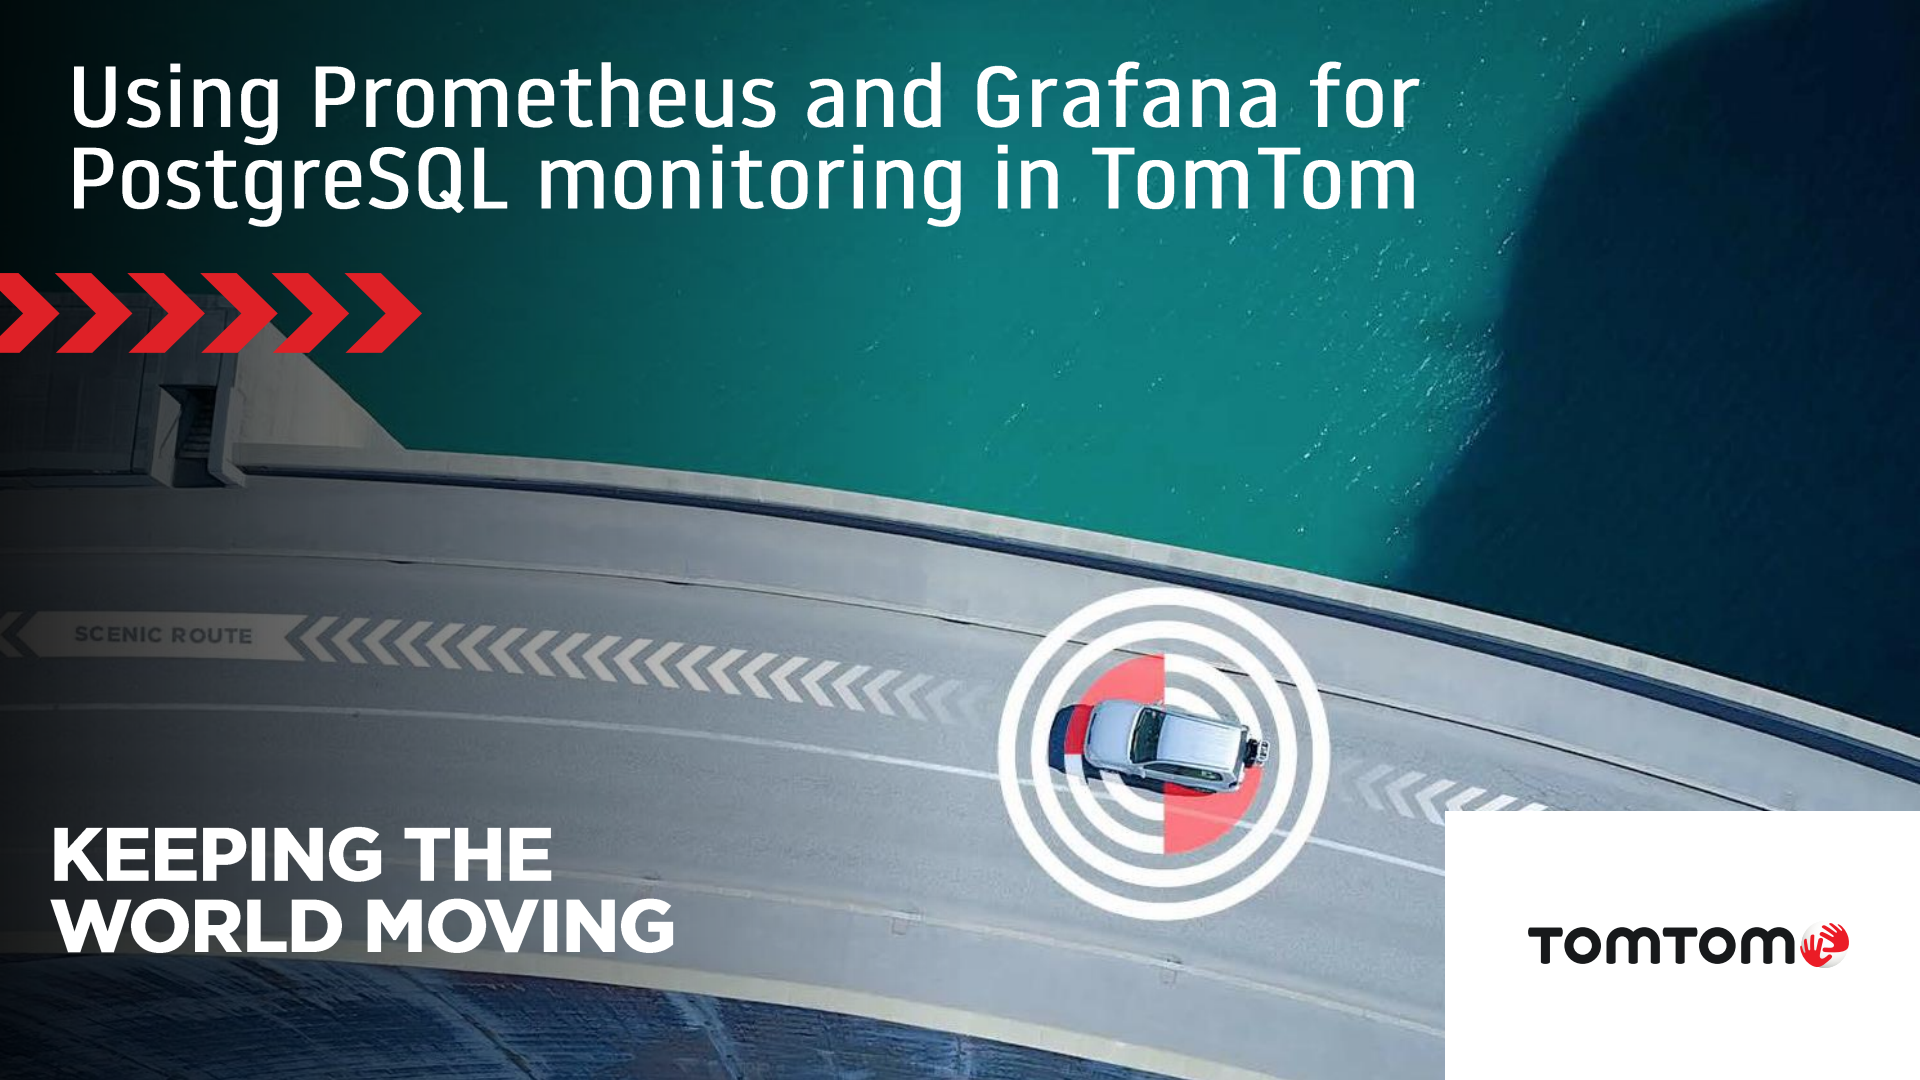

TOMTOM

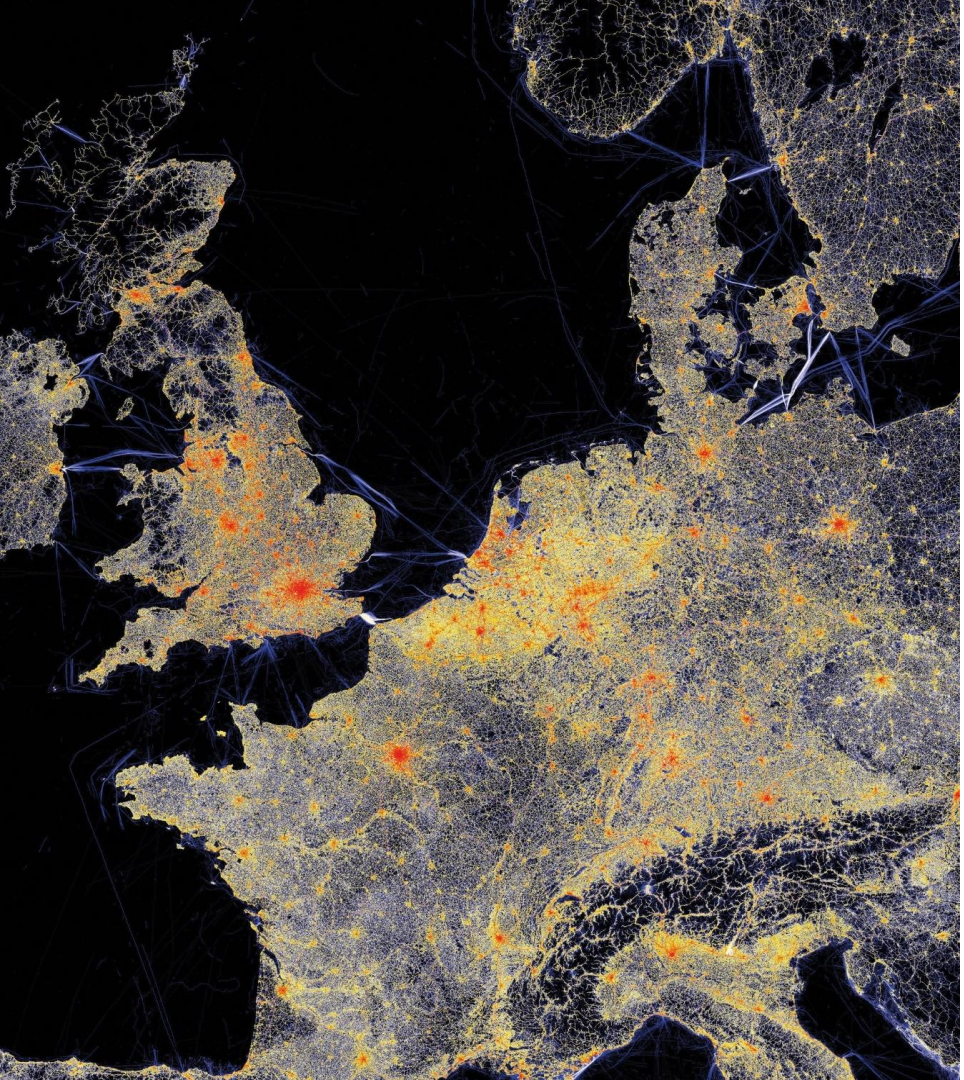

Agenda

- Who are we?
- Prometheus
  - Architecture
  - Collecting and Quering data
  - Alerting
  - High Availability
  - Installation and configuration
- Grafana
  - Data sources
  - Dashboards

- Summary

TOMTOM

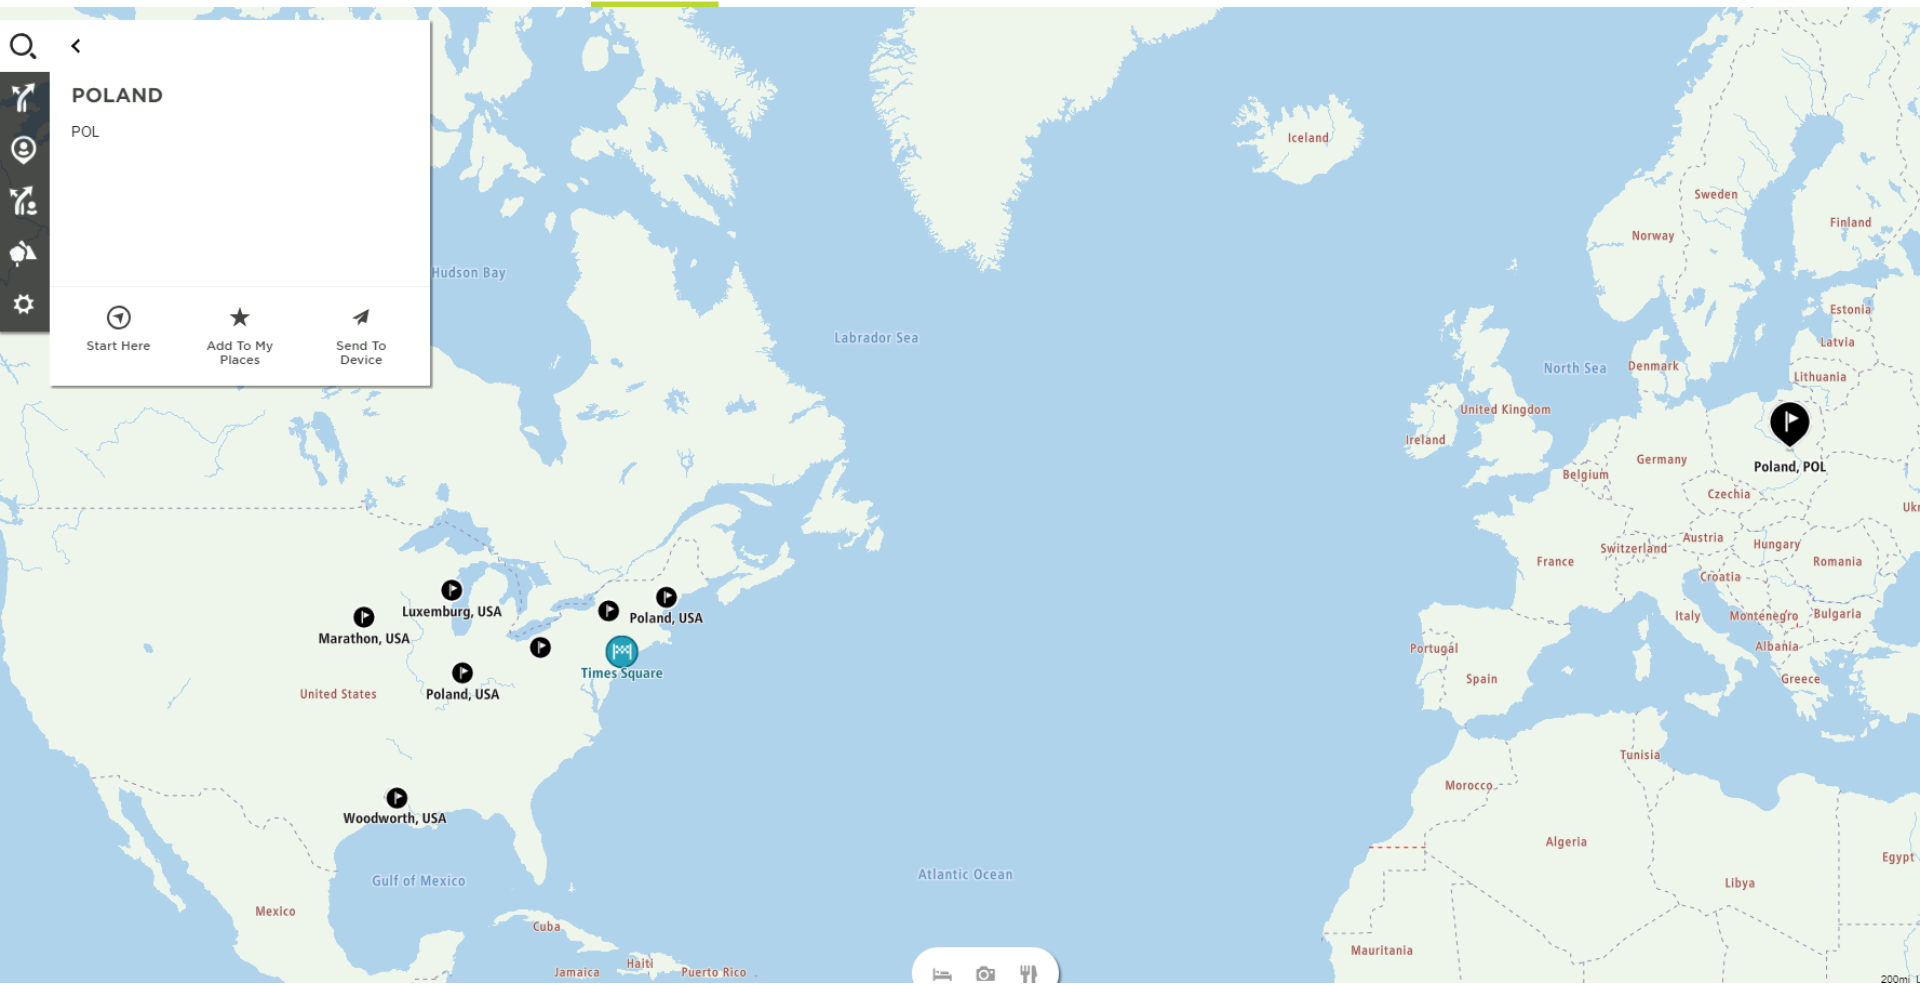
POLAND
POL
Start Here
Add To My Places
Send To Device
Iceland
Sweden
Finland
Norway
Hudson Bay
Labrador Sea
Estonia
Latvia
North Sea
Denmark
Lithuania
United Kingdom
Ireland
Germany
Belgium
Poland, POL
Czechia
Austria
Hungary
France
Switzerland
Romania
Croatia
Italy
Montenegro
Bulgaria
Albania
Portugal
Spain
Greece
Tunisia
Morocco
Algeria
Egypt
Libya
Mauritania
Luxemburg, USA
Marathon, USA
Poland, USA
Times Square
Poland, USA
United States
Woodworth, USA
Gulf of Mexico
Atlantic Ocean
Mexico
Cuba
Jamaica
Haiti
Puerto Rico

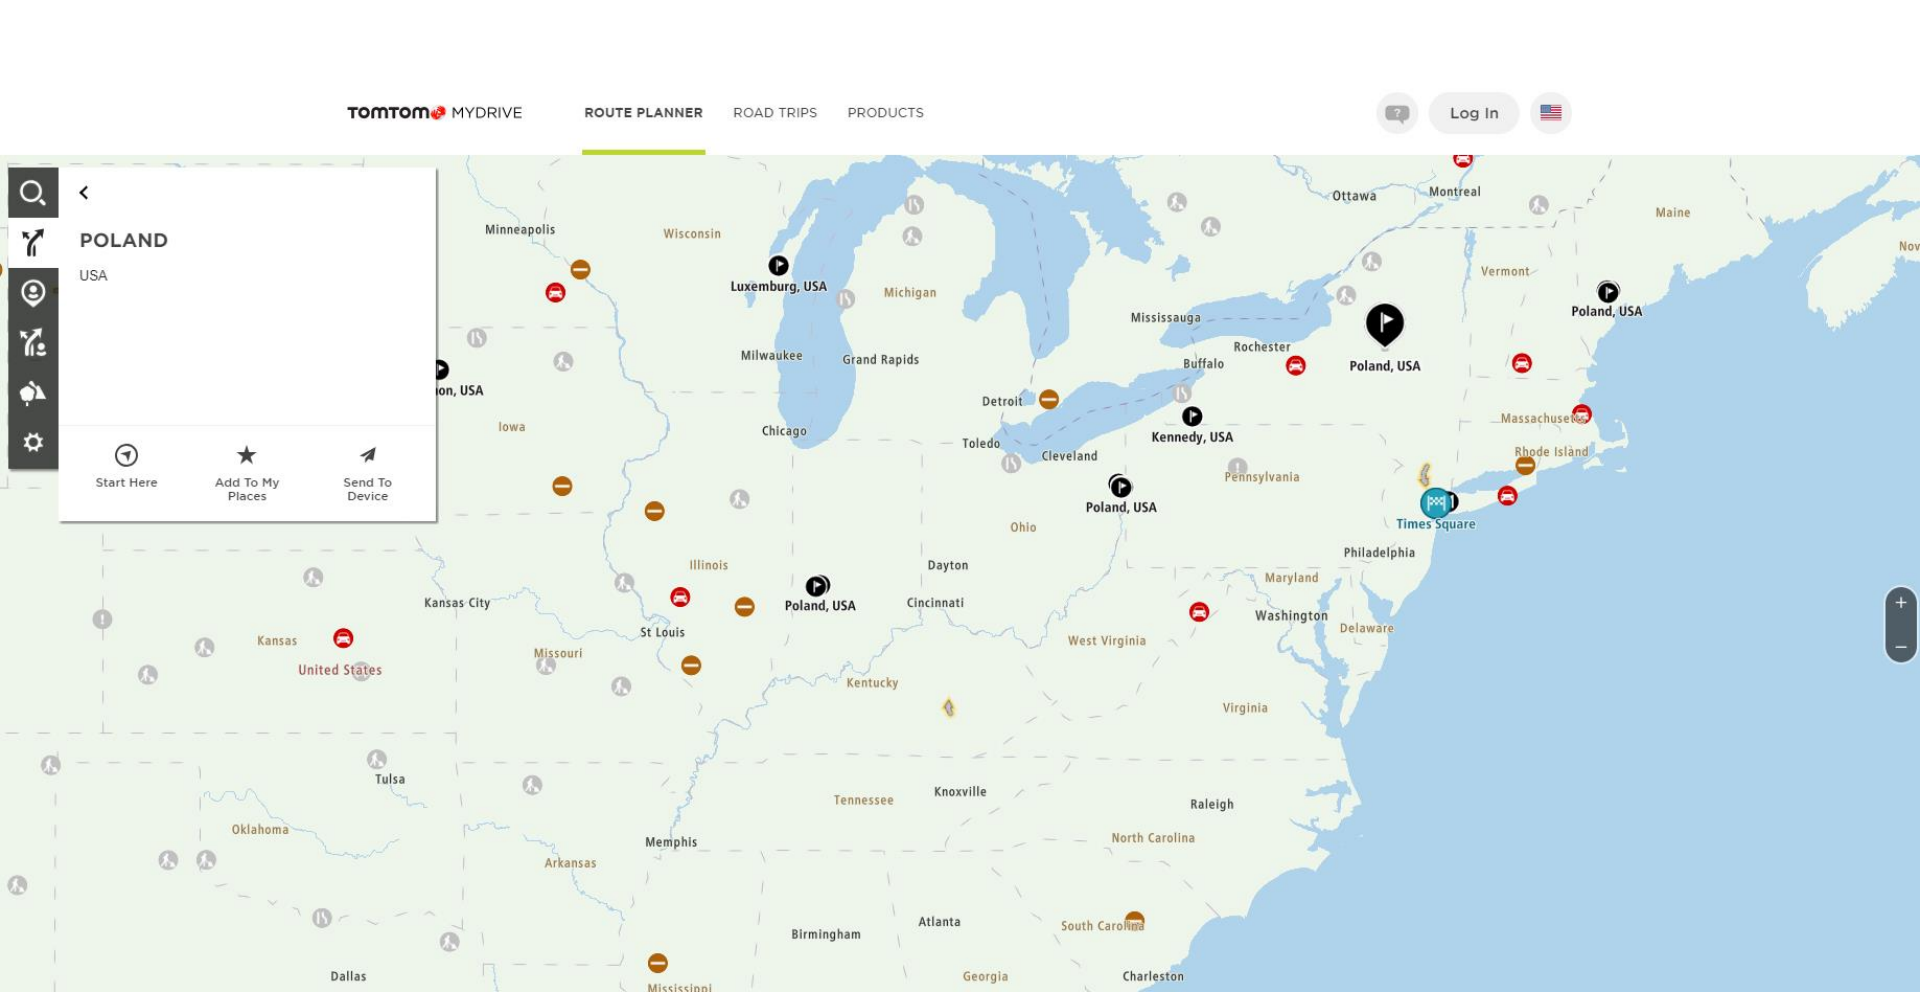
TOMTOM MYDRIVE
ROUTE PLANNER
ROAD TRIPS
PRODUCTS
Log In
POLAND
USA
Start Here
Add To My Places
Send To Device
Minneapolis
Wisconsin
Luxemburg, USA
Michigan
Milwaukee
Grand Rapids
Chicago
Iowa
Ottawa
Montreal
Maine
Vermont
Mississauga
Buffalo
Rochester
Poland, USA
Detroit
Toledo
Cleveland
Kennedy, USA
Massachusetts
Rhode Island
Pennsylvania
Ohio
Times Square
Philadelphia
Illinois
Dayton
Cincinnati
Maryland
Washington
Delaware
Kansas City
St Louis
West Virginia
Kansas
United States
Missouri
Kentucky
Virginia
Tulsa
Tennessee
Knoxville
Raleigh
Oklahoma
Memphis
North Carolina
Arkansas
Birmingham
Atlanta
South Carolina
Dallas
Georgia
Charleston

About us

- Rafał Hawrylak

  rafal.hawrylak@tomtom.com

  Software developer and database expert

- Michał Gutkowski

  michal.gutkowski@tomtom.com

  Software engineer solving problems with Java, Python, Bash... and SQL

TOMTOM

TOMTOM

Leading independent location technology specialist, shaping mobility with highly accurate maps, navigation software, real-time traffic information and services.

Sophisticated & detailed maps

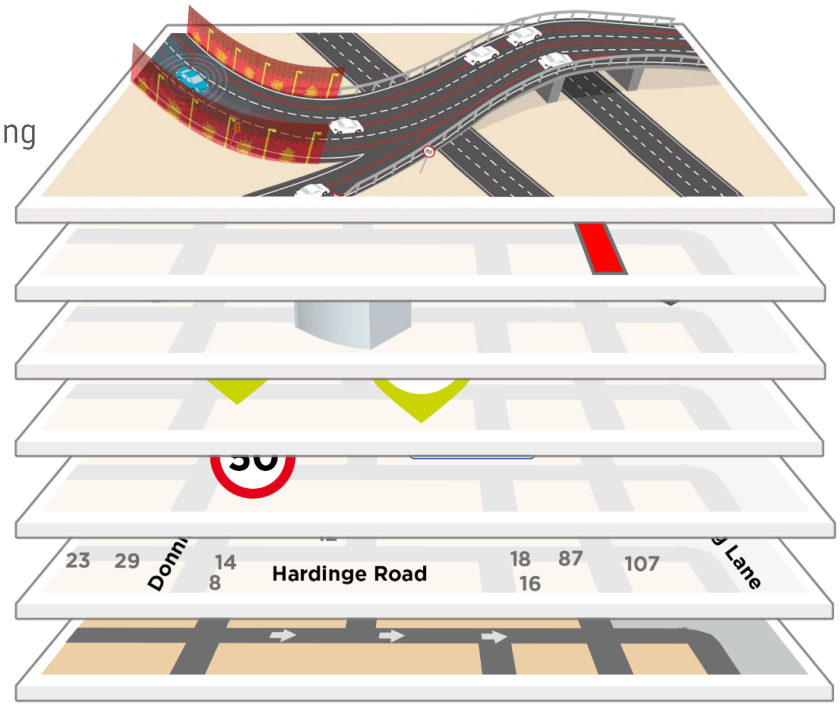

Map-making platform in 2019

- PostgreSQL + Postgis
- Data sharding and scalable reads
- 150TB of live data
- Daily db size increase: 400GB – up to 15k rows / sec
- Daily db transfers: 200TB – up to 500k queries / sec

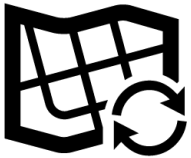

TOMTOM

Map users

Sensor input

Intelligent mapmaking

Transactional mapmaking engine

Continuously releasable map database

Incremental updates

Why monitoring is important?

- System health-check and maintenance
- Alerting and reliable notification system
- Detect performance regression
- Measure optimizations – software or business process
- Best value for money – maximum utilization
- Adjust business processes – self healing system

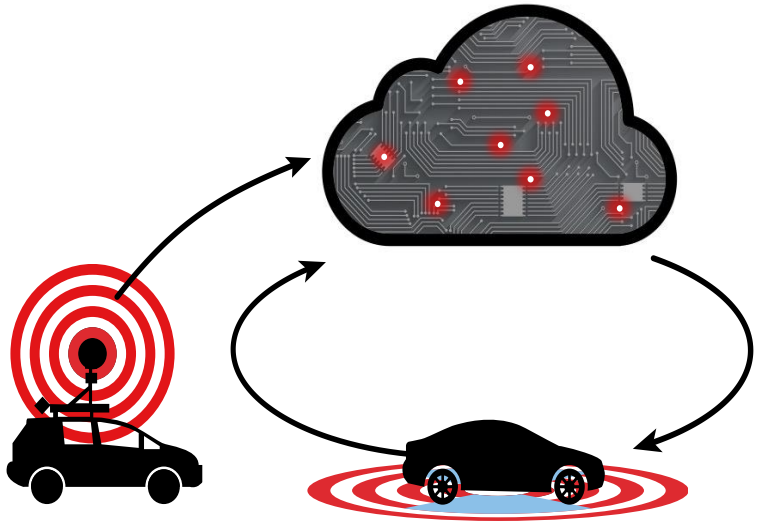

TOMTOM

Prometheus

Prometheus

What is it?

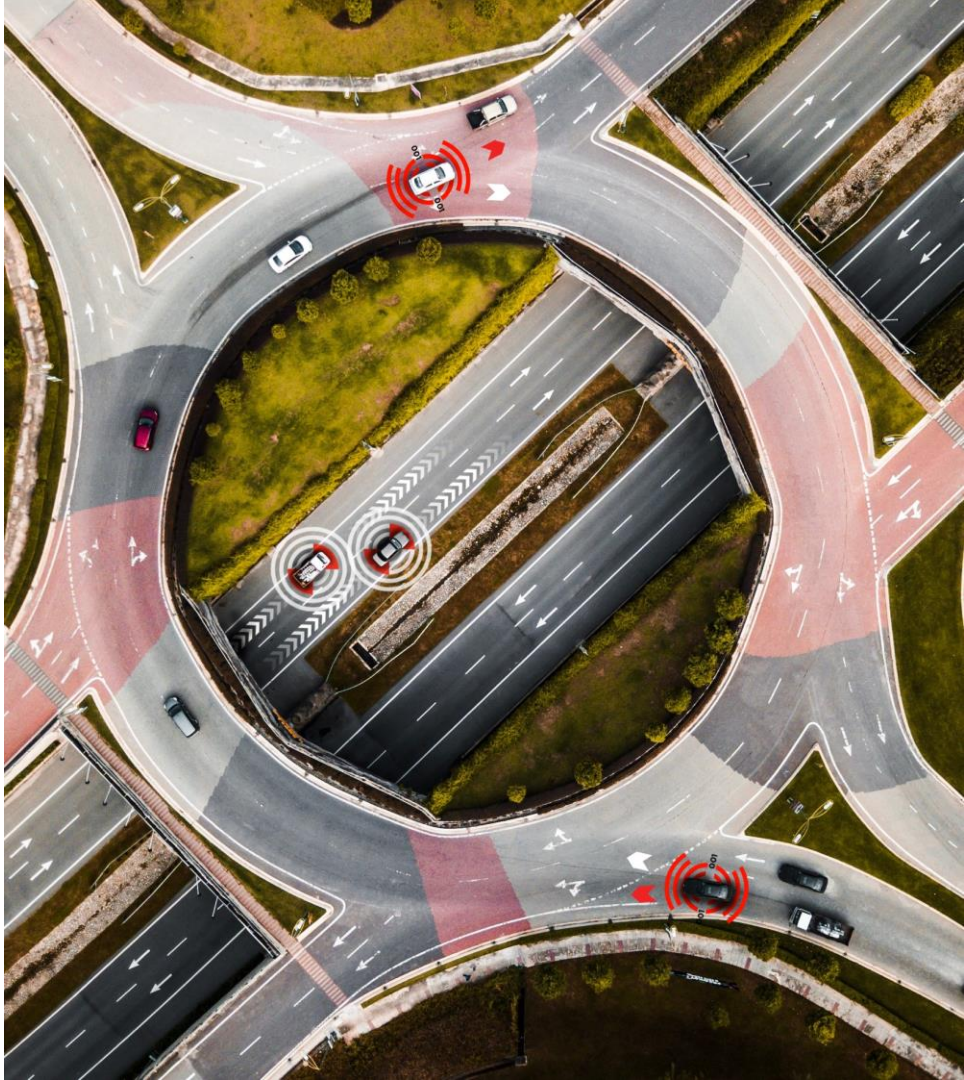

Prometheus

What is it?

- Open-source (Apache 2.0 License) monitoring toolkit https://prometheus.io
- Metrics collection
- Metrics storage in built-in time series database
- Flexible query language (promQL)
- Alerting with alert-manager

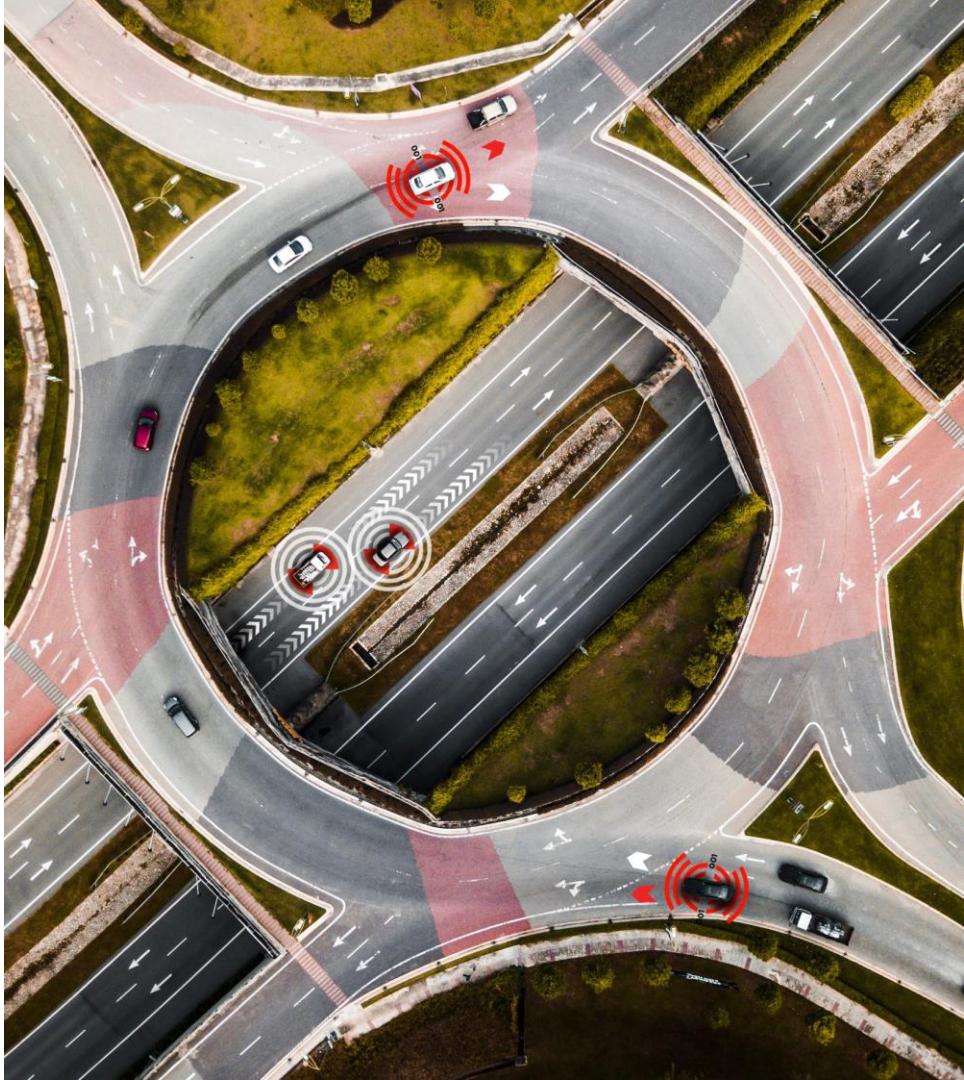

Prometheus

What is it?

Metric examples:

```
#TYPE metric_name gauge
426689100 metric_name{label1="value1",label2="value2"} 89
426689100 metric_name{label1="value3",label2="value4"} 110
426689160 metric_name{label1="value3",label2="value4"} 32
```

Prometheus

What is it?

Metric examples:

```
#TYPE metric_name gauge
426689100 metric_name{label1="value1",label2="value2"} 89
426689100 metric_name{label1="value3",label2="value4"} 110
426689160 metric_name{label1="value3",label2="value4"} 32
```

Timestamp
Metric name
Labels
Metric value

Prometheus

Prometheus

What is it not?

- Logging
- Tracing
- Anomaly detection
- Durable storage (infinitely)

Prometheus

Architecture

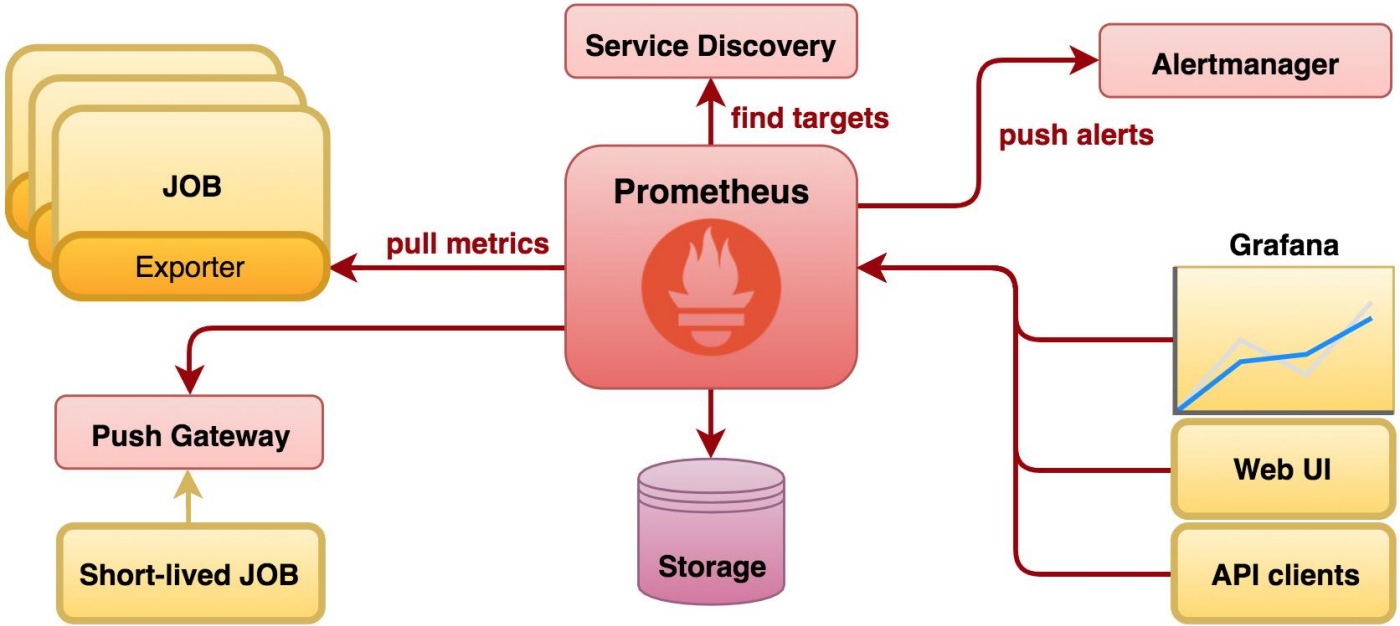

Prometheus

Prometheus

Collecting data – pulling

- Static configuration (list of IPs)
- Service discovery (Consul, AWS EC2)
- Fetching via HTTP
- Instrumentation

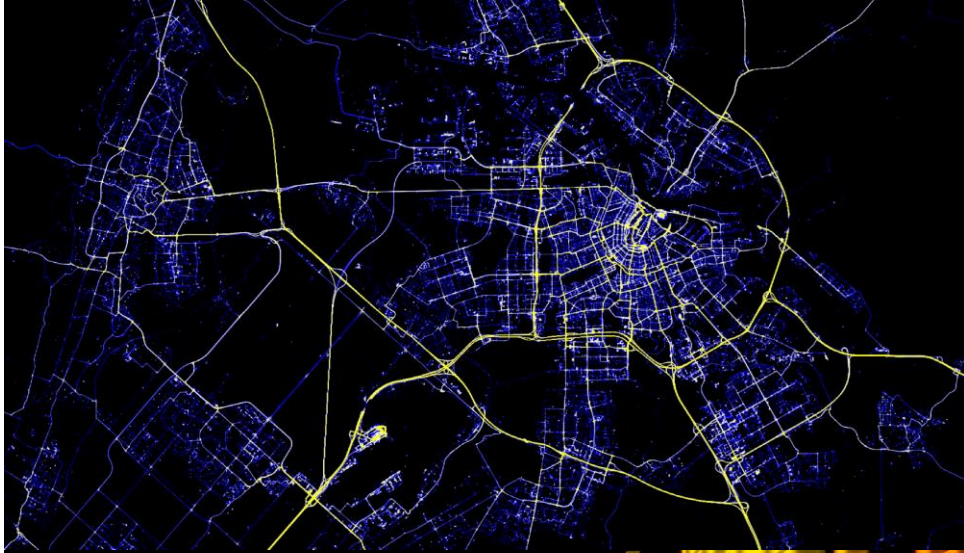

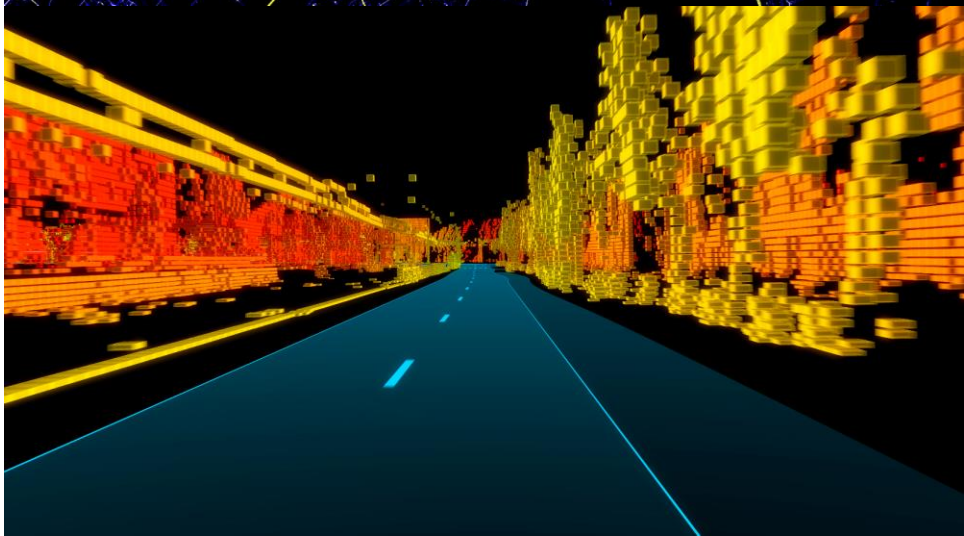

Prometheus

Collecting data - exporters

- Lots of third party exporters:
  - Hardware
  - Database
  - Messaging
  - Storage
  - HTTP
  - API
  - Other monitoring systems

https://prometheus.io/docs/instrumenting/exporters/

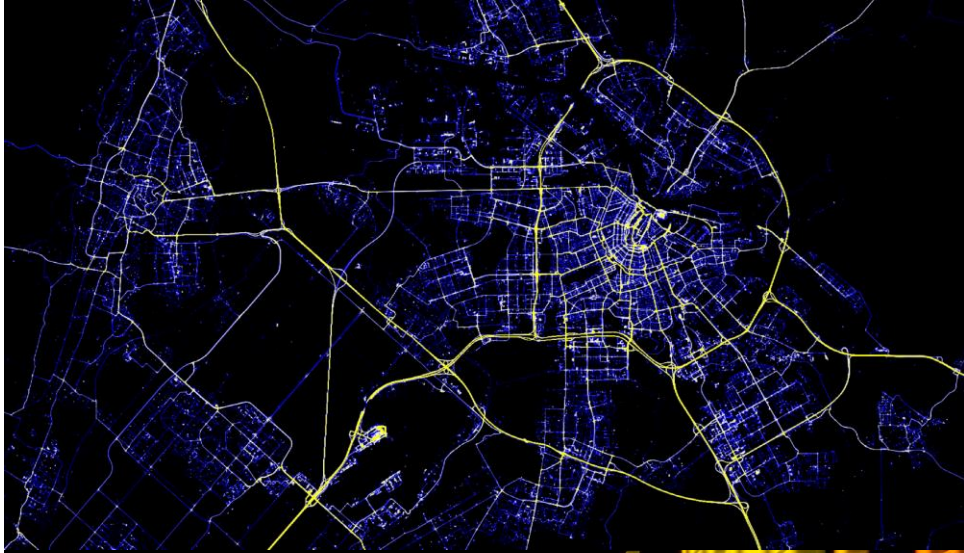

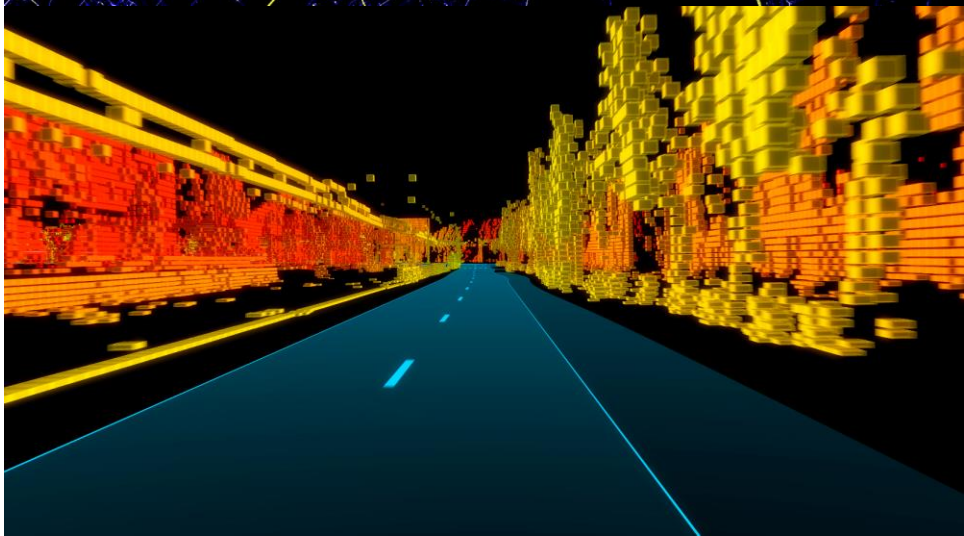

Prometheus

Collecting data - exporters

Hardware related

- apcupsd exporter
- Collins exporter
- IBM Z HMC exporter
- IoT Edison exporter
- IPMI exporter
- knxd exporter
- Netgear Cable Modem Exporter
- Netgear Router exporter
- Node/system metrics exporter **(official)**
- NVIDIA GPU exporter
- ProSAFE exporter
- Ubiquiti UniFi exporter

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

Messaging systems

- Beanstalkd exporter
- EMQ exporter
- Gearman exporter
- Kafka exporter
- NATS exporter
- NSQ exporter
- Mirth Connect exporter
- MQTT blackbox exporter
- RabbitMQ exporter
- RabbitMQ Management Plugin exporter

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

Storage

- Ceph exporter
- Ceph RADOSGW exporter
- Gluster exporter
- Hadoop HDFS FSImage exporter
- Lustre exporter
- ScaleIO exporter

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

HTTP

- Apache exporter
- HAProxy exporter (**official**)
- Nginx metric library
- Nginx VTS exporter
- Passenger exporter
- Squid exporter
- Tinyproxy exporter
- Varnish exporter
- WebDriver exporter

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

APIs

- AWS ECS exporter
- AWS Health exporter
- AWS SQS exporter
- Cloudflare exporter
- DigitalOcean exporter
- Docker Cloud exporter
- Docker Hub exporter
- GitHub exporter
- InstaClustr exporter
- Mozilla Observatory exporter
- OpenWeatherMap exporter
- Pagespeed exporter
- Rancher exporter
- Speedtest exporter

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

Logging

- Fluentd exporter
- Google's mtail log data extractor
- Grok exporter

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

Databases

- Aerospike exporter
- ClickHouse exporter
- Consul exporter **(official)**
- Couchbase exporter
- CouchDB exporter
- ElasticSearch exporter
- EventStore exporter
- Memcached exporter **(official)**
- MongoDB exporter
- MSSQL server exporter
- MySQL server exporter **(official)**
- OpenTSDB Exporter
- Oracle DB Exporter
- PgBouncer exporter
- PostgreSQL exporter
- ProxySQL exporter
- RavenDB exporter
- Redis exporter
- RethinkDB exporter
- SQL exporter
- Tarantool metric library
- Twemproxy

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data - exporters

Other monitoring systems

- Akamai Cloudmonitor exporter
- Alibaba Cloudmonitor exporter
- AWS CloudWatch exporter (**official**)
- Cloud Foundry Firehose exporter
- Collectd exporter (**official**)
- Google Stackdriver exporter
- Graphite exporter (**official**)
- Heka dashboard exporter
- Heka exporter
- InfluxDB exporter (**official**)
- JavaMelody exporter
- JMX exporter (**official**)
- Munin exporter
- Nagios / Naemon exporter
- New Relic exporter
- NRPE exporter
- Osquery exporter
- OTC CloudEye exporter
- Pingdom exporter
- scollector exporter
- Sensu exporter
- SNMP exporter (**official**)
- StatsD exporter (**official**)

https://prometheus.io/docs/instrumenting/exporters/

Prometheus

Collecting data – postgres exporter

- Out of the box
  - Replication lag
  - Uptime
  - pg_stat_user_tables
  - pg_statio_user_tables
  - pg_database
- You may add anything you need to monitor
  - Connections
  - Schema/tables/indexes sizes
  - Vacuums
  - Query Times
  - Bloat

https://github.com/wrouesnel/postgres_exporter

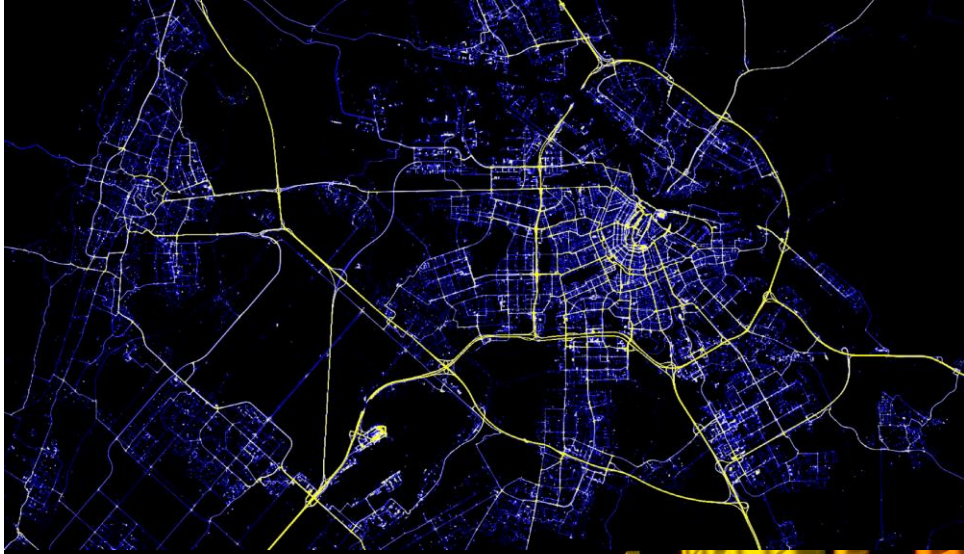

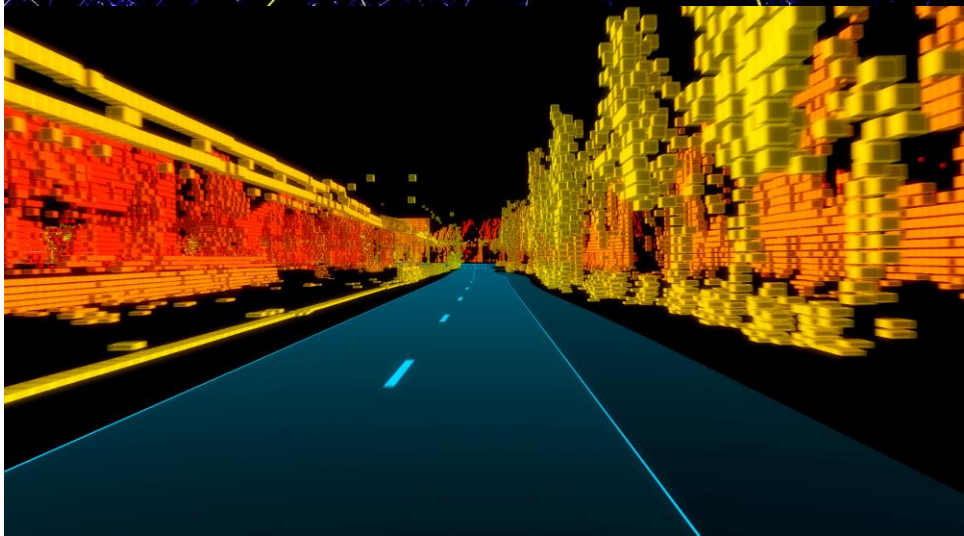

Prometheus

Custom metrics in postgres exporter

Definitions in queries.yml. Example:

```
pg_user_active_per_core:
  query: "select usename, datname, count(*) as active_userfrom pg_stat_activity where state =
'active' group by datname, usename"
  metrics:
  - usename:
    usage: "LABEL"
    description: "User"
  - datname:
    usage: "LABEL"
    description: "DB name"
  - active_user:
    usage: "GAUGE"
    description: "Active users"
```

Prometheus

Custom metrics in postgres exporter

What we get from HTTP request from Prometheus:

```
#HELP pg_user_active_per_core_active_user Active users
#TYPE pg_user_active_per_core_active_user gauge
pg_user_active_per_core_active_user{datname="core103",usename="cpp_ro"} 89
pg_user_active_per_core_active_user{datname="core104",usename="cpp_wr"} 34
pg_user_active_per_core_active_user{datname="postgres",usename="postgres"} 1
```

Prometheus

Prometheus

Quering data

It uses its own query language – PromQL

- Designed for time series data
- Not SQL style
- Functional

Prometheus

Quering data

Example PromQL query:

```
sum by (datname, usename)
(pg_user_active_per_core_active_user{datname=
~"^core.*?"}[1m])
```

Prometheus

Quering data

Example PromQL query:

```
sum by (datname, usename)
(pg_user_active_per_core_active_user{datname=
~"^core.*?"}[1m])
```

Aggregating function
Metric name
Filters
Time window for aggregation

Prometheus

Prometheus

Quering data

- Data types:
  - Instant vector
  - Range vector
  - Scalar
- Operators
  - Arithmetic
  - Comparision
  - Logical
  - Aggregation
- Functions

Prometheus

Quering data

Built in UI query browser

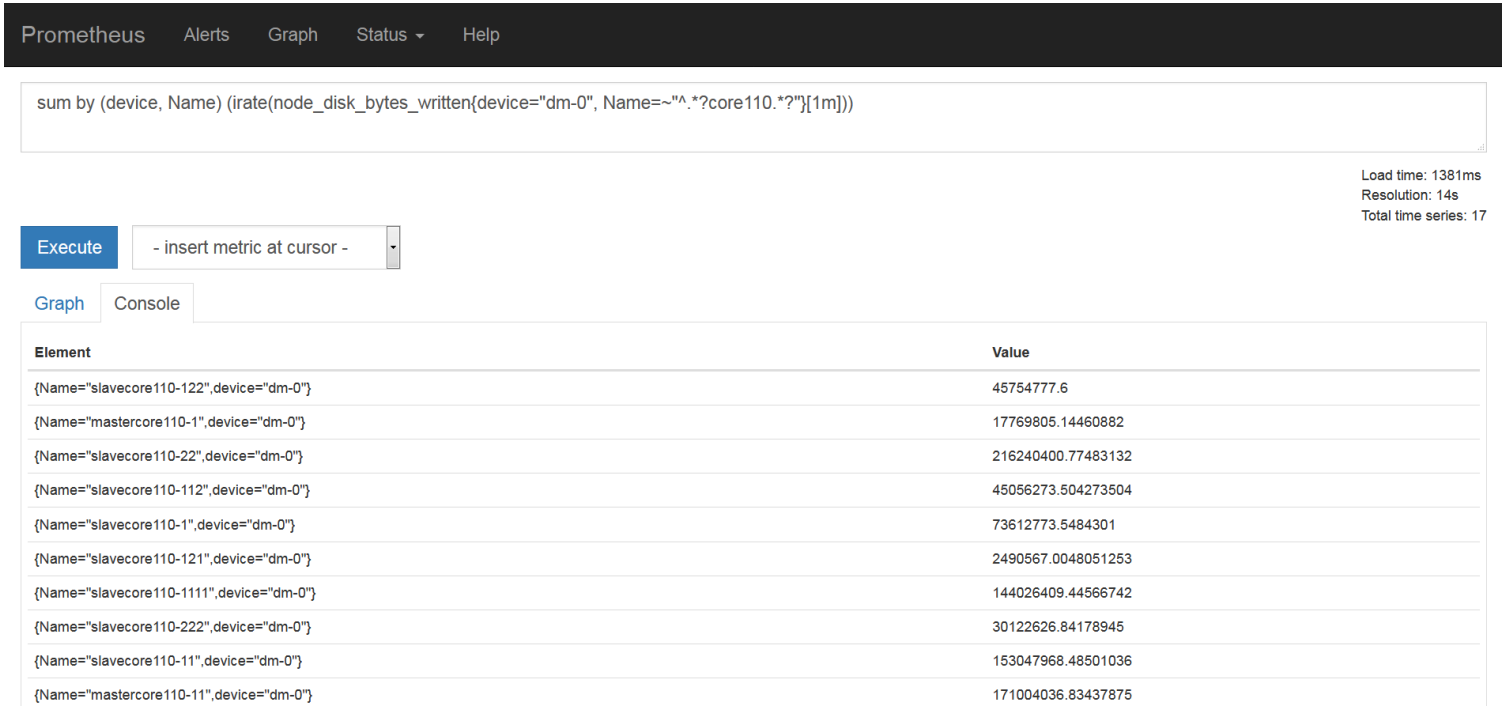

Prometheus

Prometheus

Alerting

- Alerting rules
  - Use PromQL
  - Templating
- Alerting manager
  - Inhibition
  - Aggregation
  - Routing

Prometheus

Prometheus

Alerting

CORESUP_pgdata_disc_left (0 active)

```
ALERT CORESUP_pgdata_disc_left
  IF node_filesystem_avail{device="/dev/md0",job="ec2-node-coresup",rolegroup="coresup"} < 300000000000
  FOR 10m
  LABELS {description="{{ $labels.instance }}  : {{ $value }} is less than 300GB", service="monitoring-vd
b-service", severity="critical", summary="Instance {{ $labels.instance }} vmds pgsql volume space"}
```

Prometheus

High Availability

- Autonomous single server nodes
- Local storage, no clustering
- Minimal network dependency
- Preferred metrics collection by pulling
- To avoid single point of failure run two identical but independent Prometheus servers that collect the same metrics
- To backup your metrics use decoupled remote storage (starting from 1.6)

Prometheus

Prometheus

Installation and configuration

- Precompiled binaries
- Makefile
- Docker
- Ansible
- Chef
- Puppet
- Salt

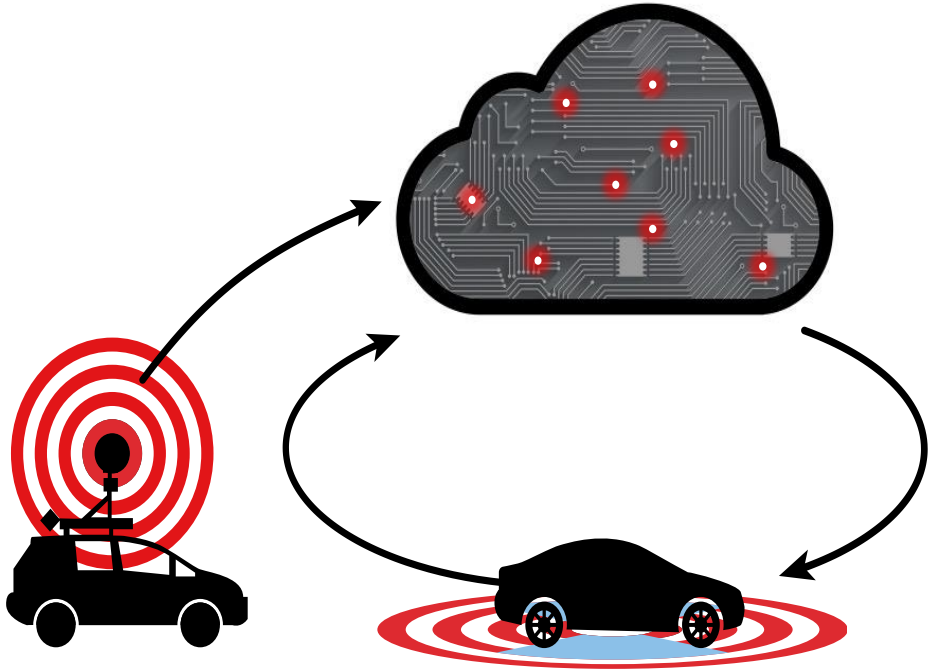

TOMTOM

Prometheus

Installation and configuration

- YAML format
- Reload at runtime

- Sections:
  - global
  - rule_files
  - scrape_configs
  - alerting
  - remote_write
  - remote_read

```
# A list of scrape configurations.
scrape_configs:
  - job_name: 'prometheus'
    scrape_interval: 10s
    scrape_timeout:  10s
    static_configs:
      - targets: ['localhost:9090']

  - job_name: "node"
    file_sd_configs:
      - files:
        - '/etc/prometheus/tgroups/*.json'
        - '/etc/prometheus/tgroups/*.yml'
        - '/etc/prometheus/tgroups/*.yaml'
  - job_name: 'ec2-node-vmds'
    ec2_sd_configs:
      - region: eu-west-1
        port: 9100
        refresh_interval: 180s
    relabel_configs:
        - source_labels: [__meta_ec2_tag_rolegroup]
          regex: vmds.*
          action: keep
        - source_labels: [__meta_ec2_instance_id]
          target_label: instance
        - action: labelmap
          regex: __meta_ec2_tag_(.+)
  - job_name: 'postgres-vmds'
    ec2_sd_configs:
      - region: eu-west-1
        port: 9187
        refresh_interval: 180s
    relabel_configs:
        - source_labels: [__meta_ec2_tag_rolegroup]
          regex: vmds.*
          action: keep
        - source_labels: [__meta_ec2_instance_id]
          target_label: instance
        - action: labelmap
          regex: __meta_ec2_tag_(.+)
```

Grafana

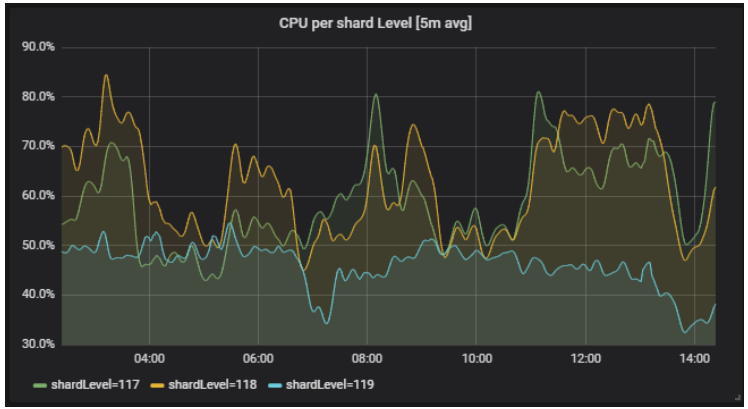
CPU per shard Level [5m avg]
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
04:00
06:00
08:00
10:00
12:00
14:00
shardLevel=117
shardLevel=118
shardLevel=119


Monitoring - Screens - engineering
100.00%
3
100.00%
7
96.94%
PROCESSING
9526
80
56.96 s
570 ms
112.58 s
570 ms
1.87 min
getFeaturesByIds - AVG SYNC
23 ms
21 ms
getFeaturesByIds - AVG PROC
86 ms
114 ms
getObjectsByIds - AVG SYNC
3 ms
2 ms
getObjectsByIds - AVG PROC
19 ms
38 ms
VMDS SLAVES LAG
VMDS SLAVES LOAD
EDITING
Success Rate - Edit manual
100.00%
Success Rate - Edit AUTO
99.99%
Edits Count [15m]
manual failure count
automatic failure count
40.00

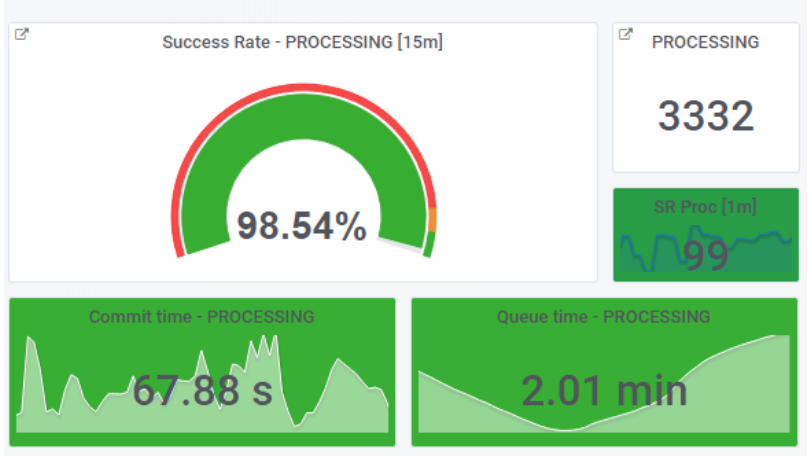
Success Rate - PROCESSING [15m]
98.54%
PROCESSING
3332
SR Proc [1m]
99
Commit time - PROCESSING
67.88 s
Queue time - PROCESSING
2.01 min

Grafana

- Rich visualization toolkit https://grafana.com/
- Open-source (Apache 2.0 License) with enterprise support
- Flexible dashboards helps understand data
- Time series data-sources:
  - Elasticsearch,
  - Promethues,
  - AWS CloudWatch,
  - Graphite

Grafana

Data-sources

- Installed as plugins
- Query editors to get metrics

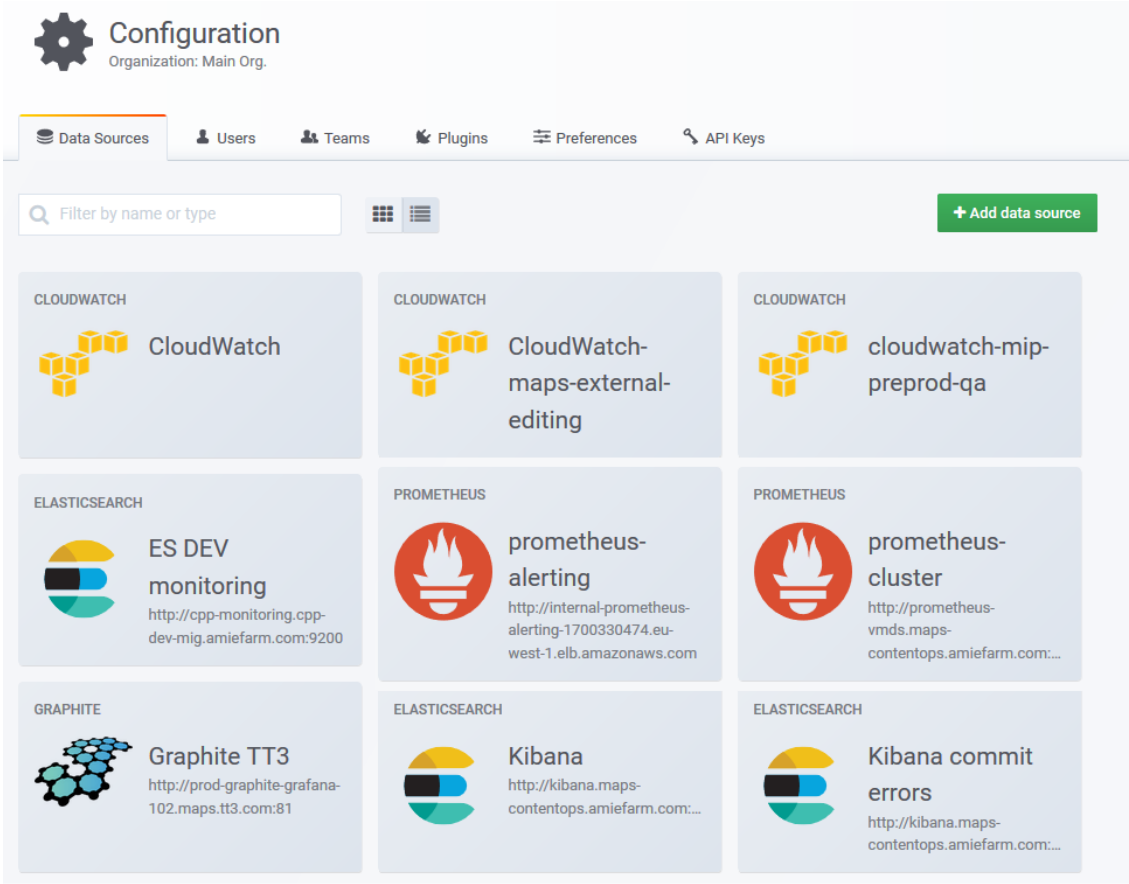

Grafana

Dashboards

- Rows
- Panels
- Query editors

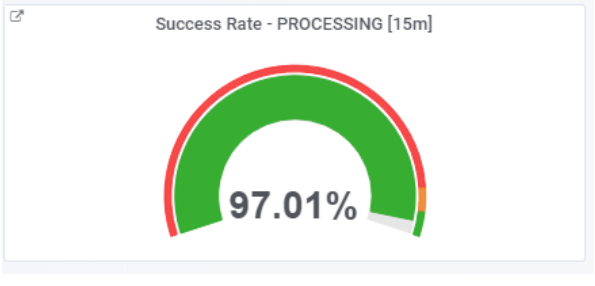

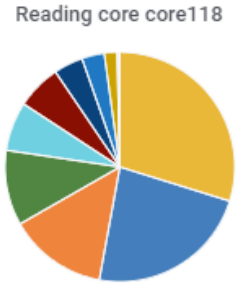

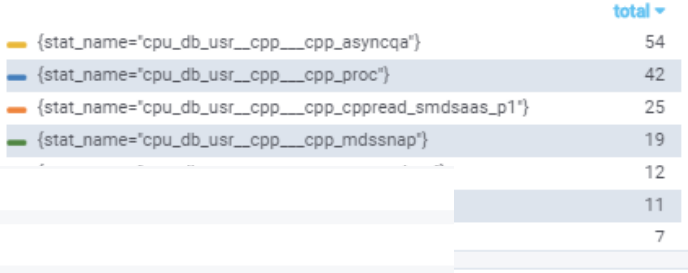

TomTom

Grafana

Query editors

- Difference in data-source settings
- Native queries

- AWS CloudWatch exporter
- Elasticsearch exporter

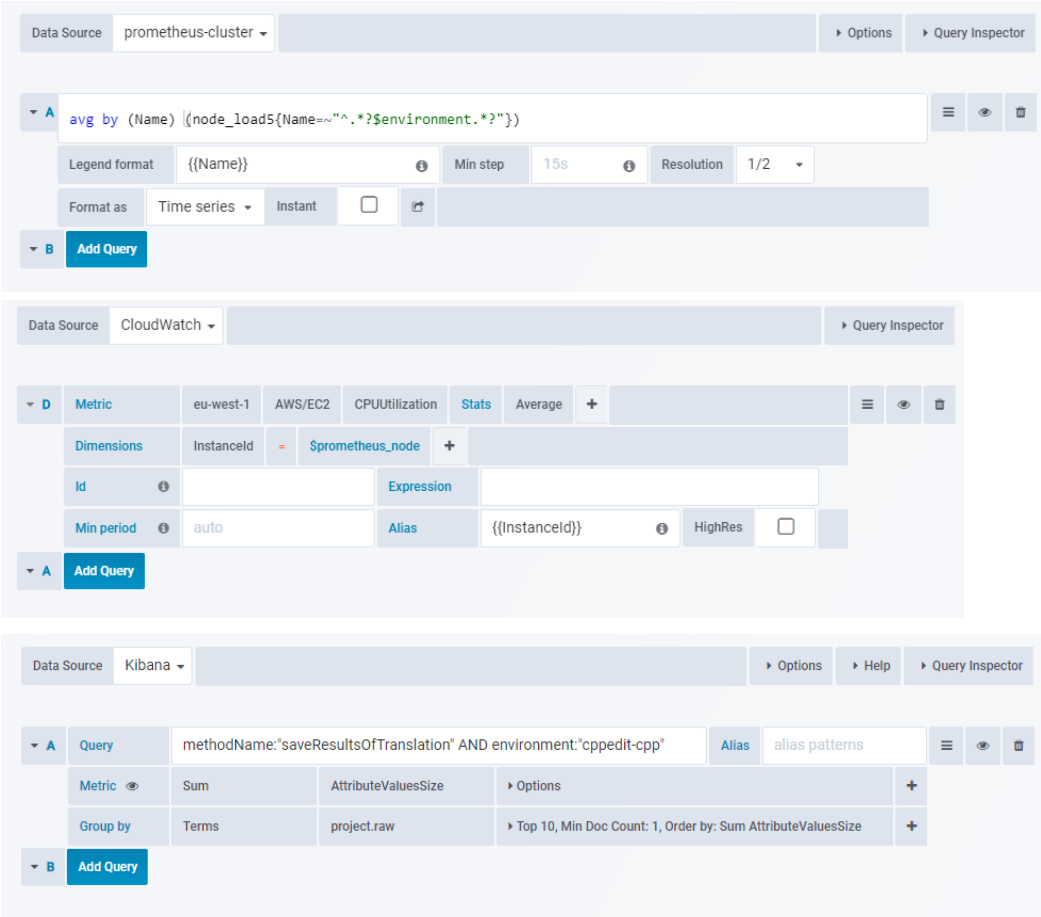

Grafana

Advanced dashboards

- Templating with dynamic variables
- Annotations for marking events
- ACL per user per dashboard

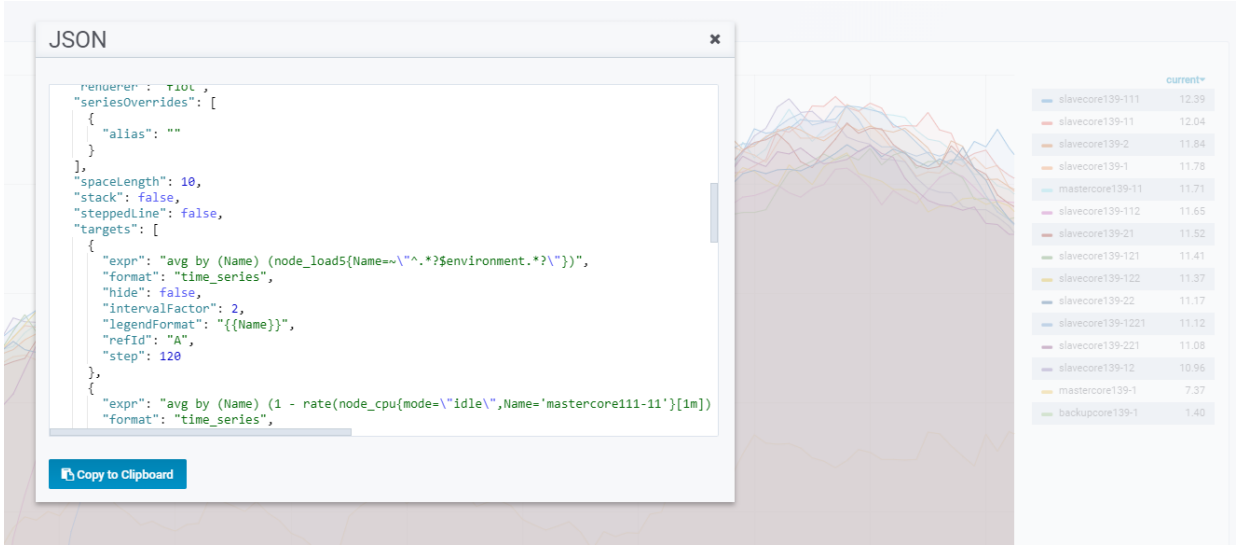

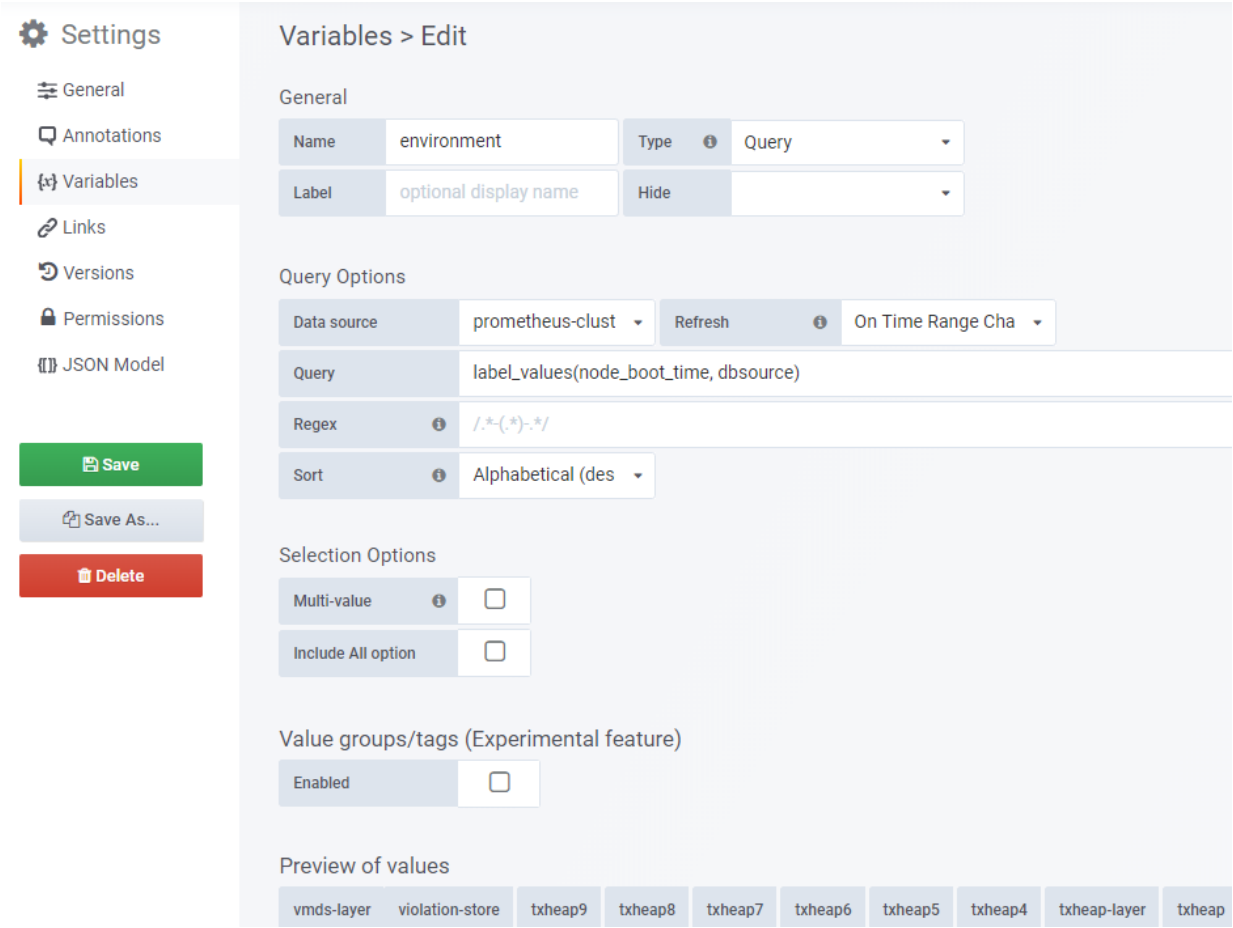
Settings
General
Annotations
Variables
Links
Versions
Permissions
JSON Model
Save
Save As...
Delete
Variables > Edit
General
Name
environment
Type
Query
Label
optional display name
Hide
Query Options
Data source
prometheus-clust
Refresh
On Time Range Cha
Query
label_values(node_boot_time, dbsource)
Regex
/.*-(.*)-.*/
Sort
Alphabetical (des
Selection Options
Multi-value
Include All option
Value groups/tags (Experimental feature)
Enabled
Preview of values
vmds-layer
violation-store
txheap9
txheap8
txheap7
txheap6
txheap5
txheap4
txheap-layer
txheap

TOMTOM

Grafana

Advanced dashboards

- HTTP API – manage Grafana and dashboards
- JSON model: properties, variables, panels and queries
- Versioning and changes tracking
- Benefit from automation - limit manual work

```
POST /api/dashboards/db HTTP/1.1
Accept: application/json
Content-Type: application/json
Authorization: Bearer eyJrIjoiT0tTcG1pUl

{
  "dashboard": {
    "id": null,
    "uid": null,
    "title": "Production Overview",
    "tags": [ "templated" ],
    "timezone": "browser",
    "schemaVersion": 16,
    "version": 0
  },
  "folderId": 0,
  "overwrite": false
}
```

TOMTOM

Layer Baseline All | API Endpoint | Process Id genesis_NAVRE_L74 | CPPLyr Cluster All

Status

- Selected (1)
- All
- liers
- liers-density
- liers-monitoring
- liers-streetname
- liers-genesis
- liers-development
- liers-openlr
- liers-openlr-genesis

Sync Lag

Layer	itDelay	Average versionDelay	Max commitDelay	Max versionDelay	Commits
		20 Mil	59.97 day	20 Mil	
		20 Mil	59.83 day	20 Mil	
		15	50.78 s	135	
		8	52.67 s	121	

Sync Lag *(2 panels)*

Commit Success Rate *(4 panels)*

Synchronization *(6 panels)*

DB

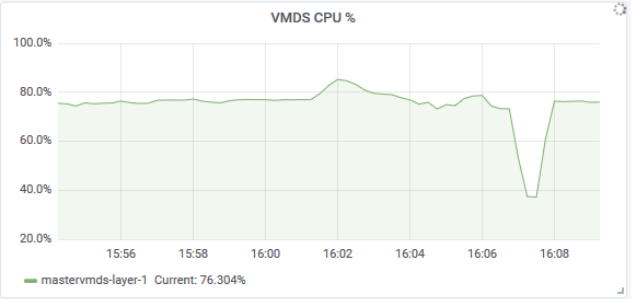

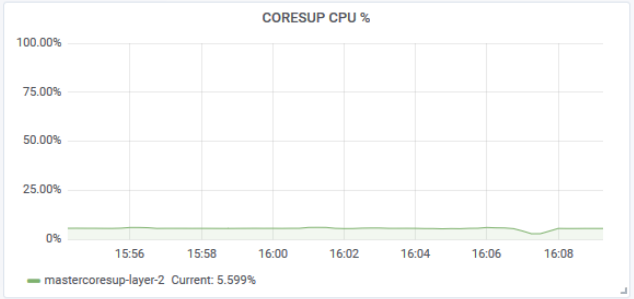

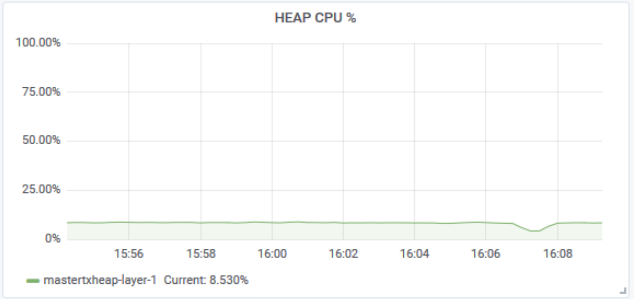

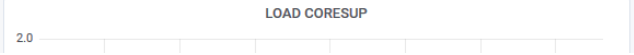

Summary

- Prometheus can collect and store time-series data with fast lookups.
- External datastores for long-term storage and HA (InfluxDB or PostgreSQL).
- Grafana enables visualization of data from different sources.
- Both tools allow automation in building monitoring infrastructure.
- Dashboards are important, but an alerting system should notify about crossing treshholds.

What's next?

- Start playing with it!

- wget https://github.com/prometheus/prometheus/releases/download/v2.8.0/prometheus-2.8.0.linux-amd64.tar.gz

  tar –zxvf prometheus-...

- wget https://dl.grafana.com/oss/release/grafana-6.0.2.linux-amd64.tar.gz
  tar -zxvf grafana-..

TOMTOM

Our offices

5000 employees worldwide

We are hiring!

https://www.tomtom.com/careers/

TOMTOM

THANK YOU

TOMTOM