

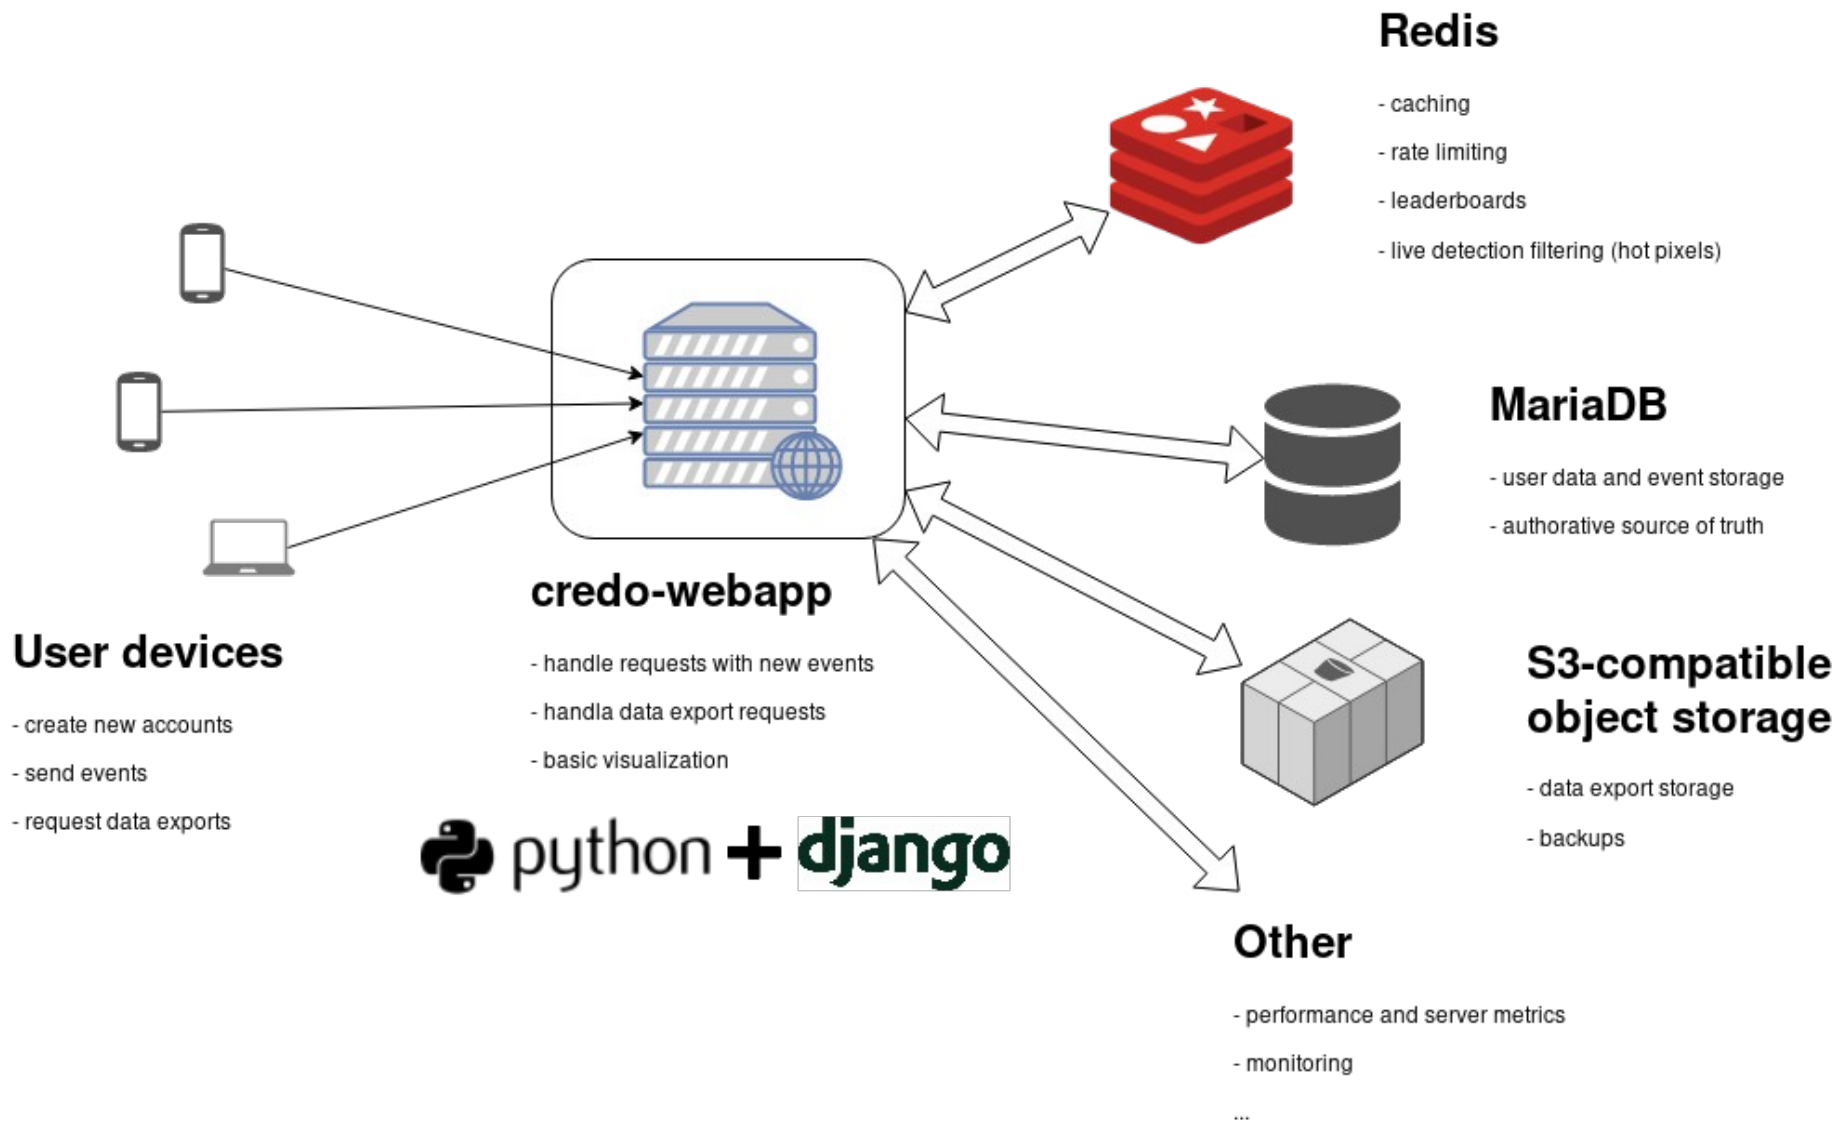


# CREDO backend infrastructure

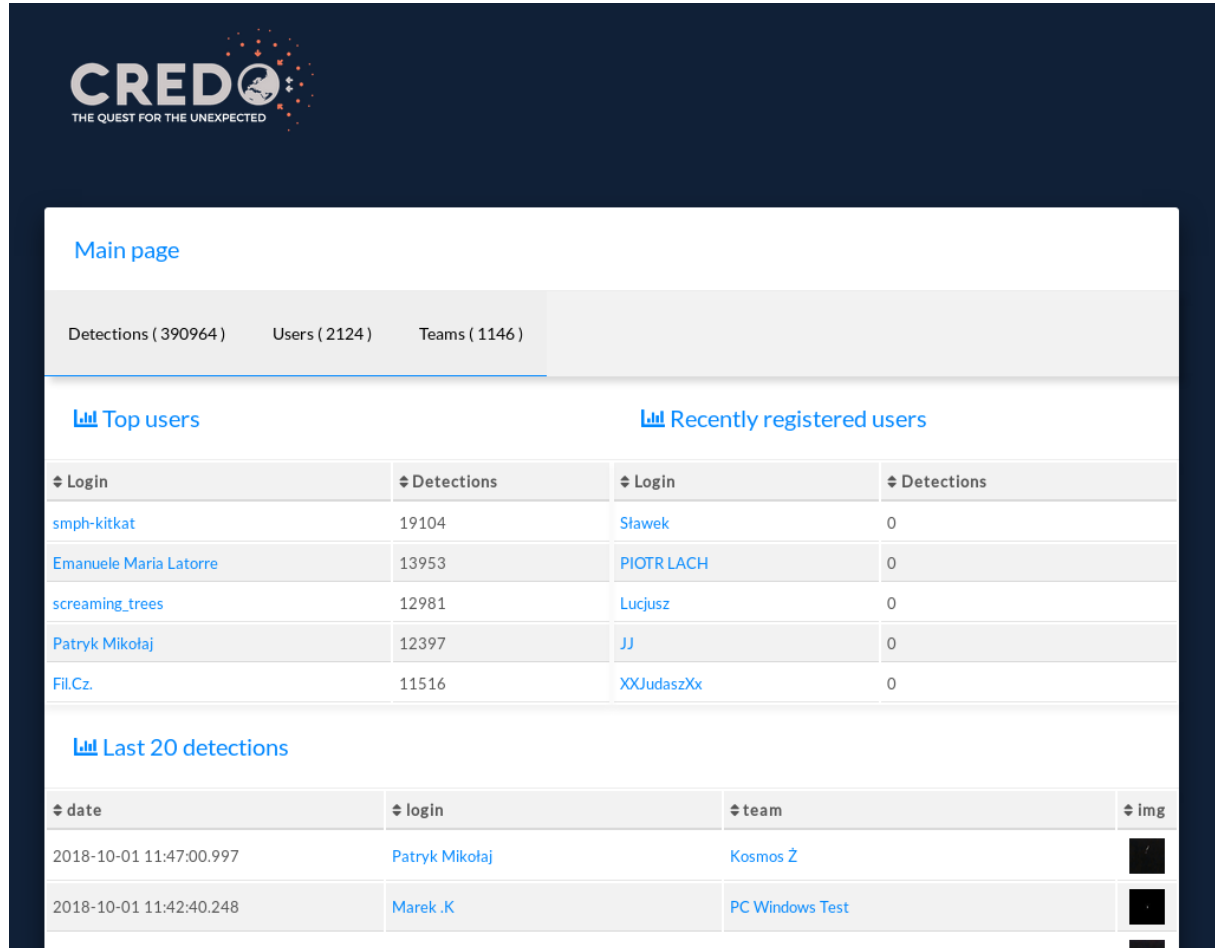
Marek Magryś, Krzysztof Oziomek, Maciej Pawlik

# Some statistics

- 390k+ visible detections (1.4m+ overall)
- 745k+ device pings (sums up to 48 years looking for particles)
- 2k+ users with at least 1 detection
- 4k+ devices
- 1k+ user teams
- 10s of GBs used for storage of data, metrics and backups



# credo-webapp (https://api.credo.science)



The screenshot displays the credo-webapp interface. At the top left is the Credo logo with the tagline "THE QUEST FOR THE UNEXPECTED". Below the logo, there are three tabs: "Detections (390964)", "Users (2124)", and "Teams (1146)". The "Users (2124)" tab is selected. Underneath, there are two sub-sections: "Top users" and "Recently registered users". Each sub-section contains a table with columns for "Login" and "Detections". Below these is a section for "Last 20 detections" with a table containing columns for "date", "login", "team", and "img".

**CREDO**  
THE QUEST FOR THE UNEXPECTED



Main page

Detections (390964) Users (2124) Teams (1146)

[Top users](#) [Recently registered users](#)

↕ Login	↕ Detections	↕ Login	↕ Detections
<a href="#">smph-kitkat</a>	19104	<a href="#">Slawek</a>	0
<a href="#">Emanuele Maria Latorre</a>	13953	<a href="#">PIOTR LACH</a>	0
<a href="#">screaming_trees</a>	12981	<a href="#">Lucjusz</a>	0
<a href="#">Patryk Mikołaj</a>	12397	<a href="#">JJ</a>	0
<a href="#">Fil.Cz.</a>	11516	<a href="#">XXJudaszXx</a>	0

[Last 20 detections](#)

↕ date	↕ login	↕ team	↕ img
2018-10-01 11:47:00.997	<a href="#">Patryk Mikołaj</a>	<a href="#">Kosmos Ż</a>	
2018-10-01 11:42:40.248	<a href="#">Marek .K</a>	<a href="#">PC Windows Test</a>	

# credo-webapp

- <https://github.com/credo-science/credo-webapp>
- We try to show filtered detections that are not noise
- Implements API for detectors to submit events
- Allows users to create accounts used for sending events and exporting data for analysis
- Django app running under Apache and mod\_wsgi

# Data stored with detection events

- Time (the the device thinks the event happened and time we received it)
- Device location (+ accuracy, location provider)
- Coordinates of the event on device sensor (used for hot pixel detection)
- Image of the event (if received)
- User and device the detection is coming from

# credo-api-tools

- <https://github.com/credo-science/credo-api-tools>
- Collection of tools designed to simplify process of exporting and analyzing CREDO data
- “data-exporter” handles making API calls to authenticate user, request and download data export incrementally
- “data-processor” handles incremental data processing and provides simple plugin interface for scientists to write their own code

# Performance monitoring and metrics


django-influxdb-metrics + InfluxDB + Grafana

Raven + Sentry


Metricbeat + Elasticsearch + Kibana




Requests per 10m




Detection count



Ping count

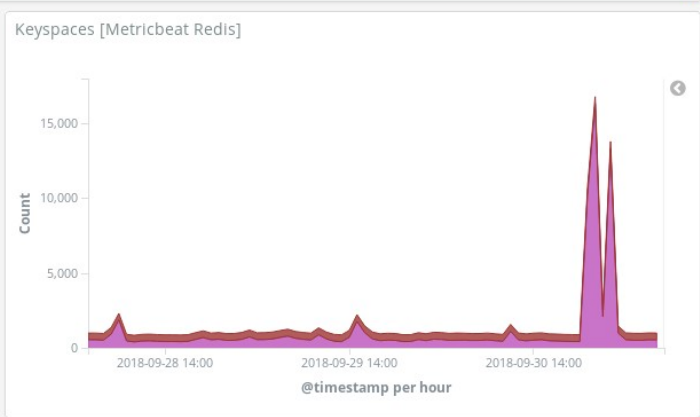

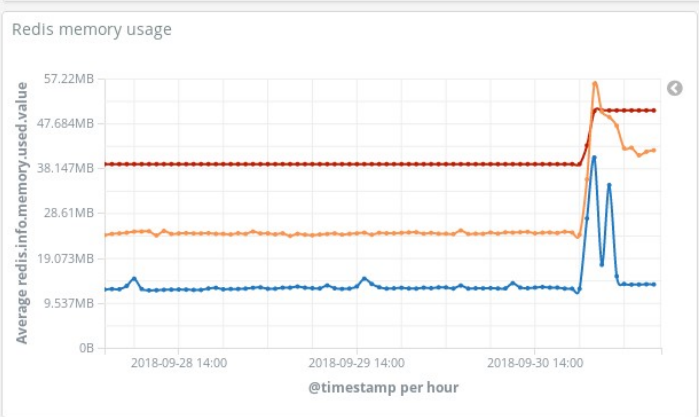


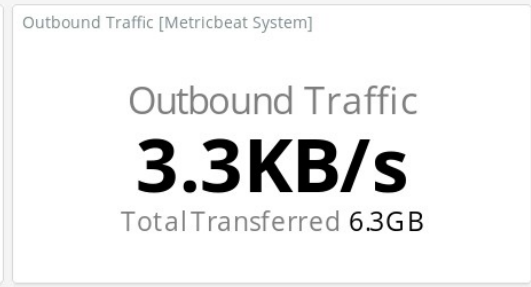
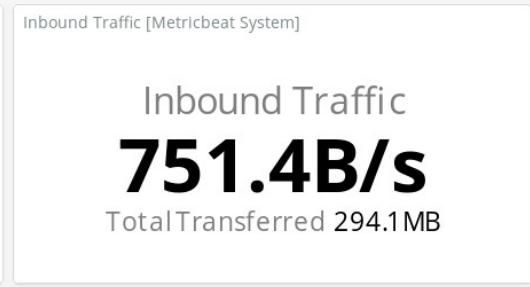
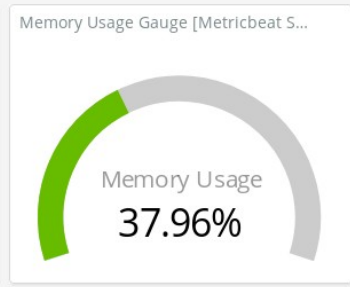
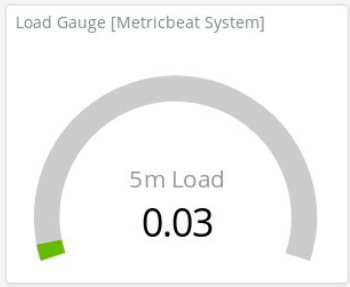
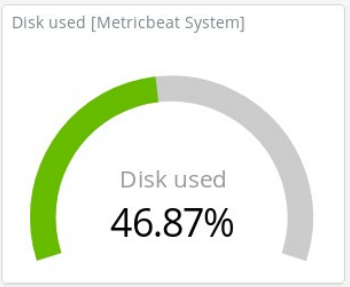


Detection/Ping latency percentiles

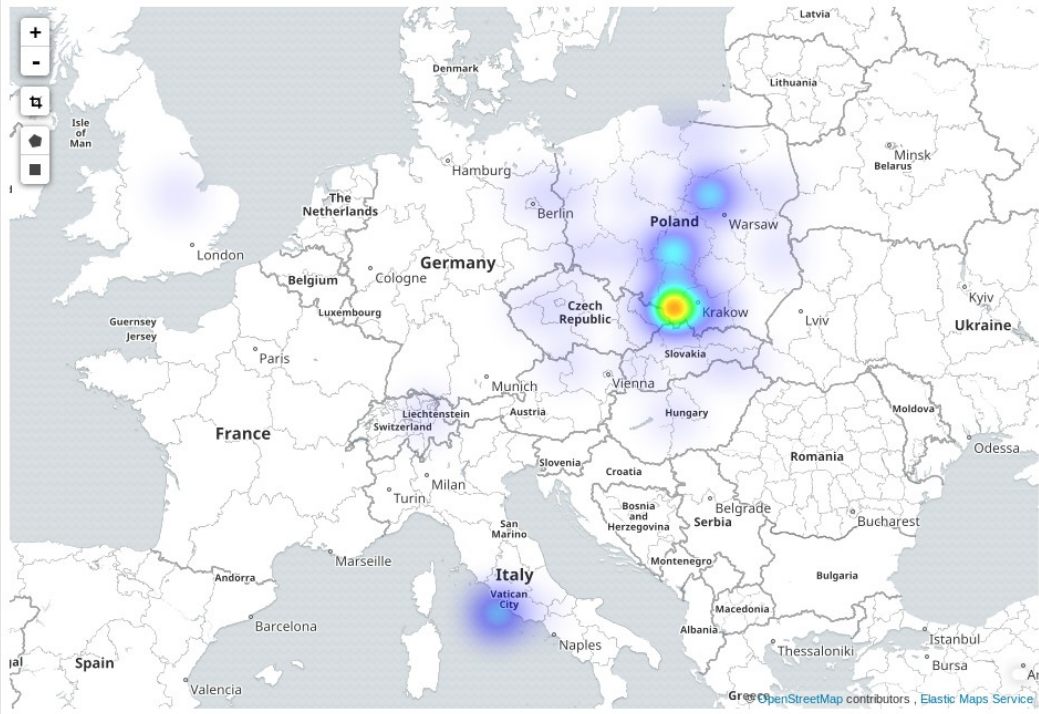


Slowest views (95th percentile)

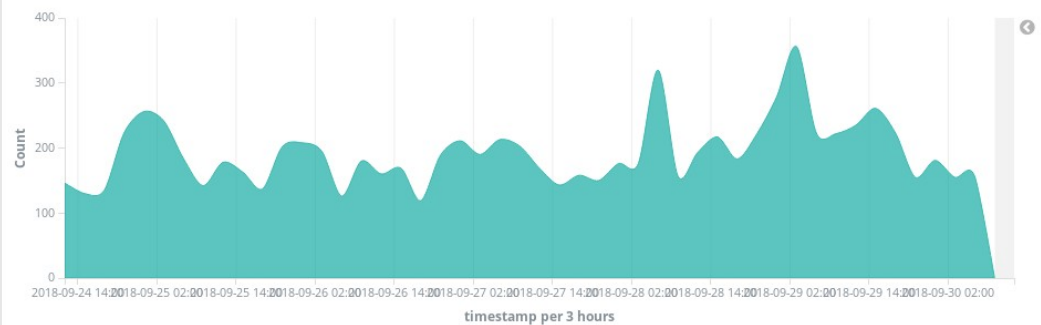
view	Time spent
changelist_view	2 seconds, 65 milliseconds
user_page	370 milliseconds
index	250 milliseconds
finished_jobs	225 milliseconds
DataExportView	197 milliseconds
detection_list	194 milliseconds
stats	118 milliseconds
team_page	108 milliseconds



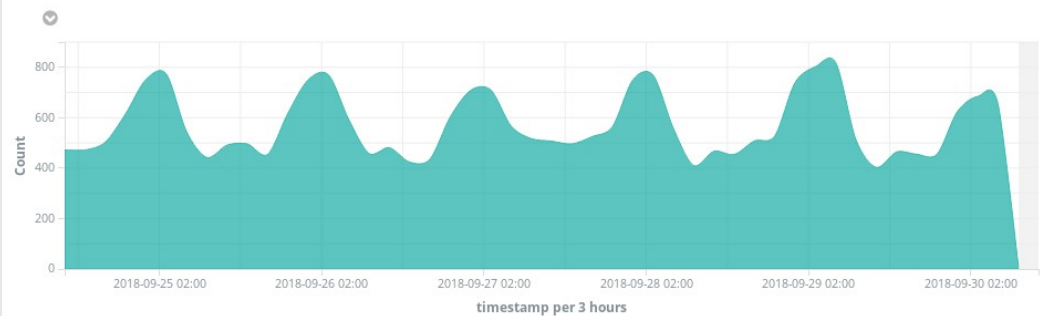
[CREDO] Detection heatmap



[CREDO] Visible detections rate



[CREDO] Ping rate



[CREDO] Top users (detection count)

User ID	Detection count
1510	1,432
4980	1,180
4125	886
157	747
5363	736
Other	4,033
	<b>9,014</b>

[CREDO] Top users (by on time)

User ID	On time	Ping count
4980	21 days	3,041
4125	14 days	2,161
157	10 days	1,438
1584	6 days	889
1849	6 days	847
Other	4 months	18,434
	<b>6 months</b>	<b>26,810</b>

