elastic

Predictive Maintenance using Elastic Machine Learning

Search. Observe. Protect.

May 3rd, 2021

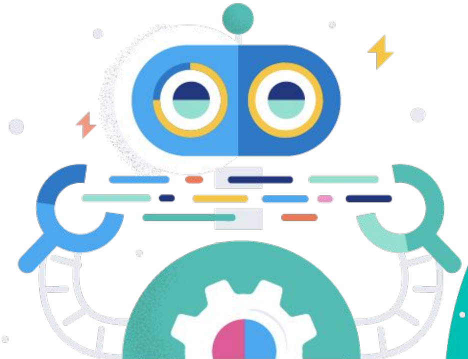

Felix Rössel

Principal Solutions Architect

Agenda

Intro predictive analytics

Real-time analytics and ML capabilities

Demo

Real world example

Q & A

Intro predictive analytics

World’s new most valuable resource

elastic

The Need for Relevant Real-Time Insights

Users are demanding more of applications.

Enterprise IT environments are becoming increasingly complex.

Digital threats are targeting any weakness in an evolving attack surface.

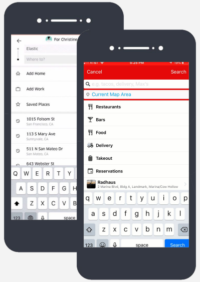

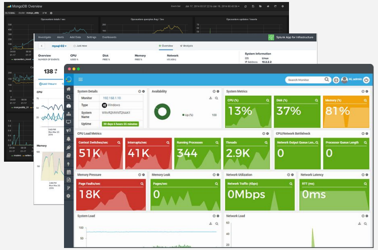

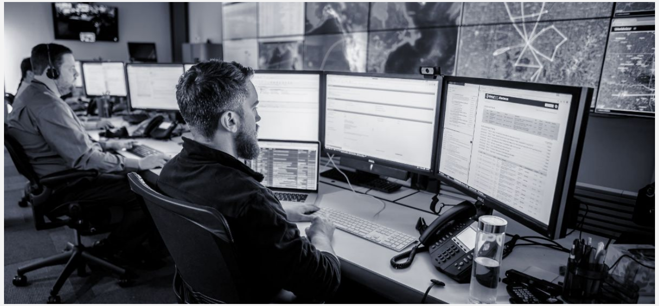

Business leaders are now, more than ever before, focused on using data to improve their business.

elastic

Elastic Technology

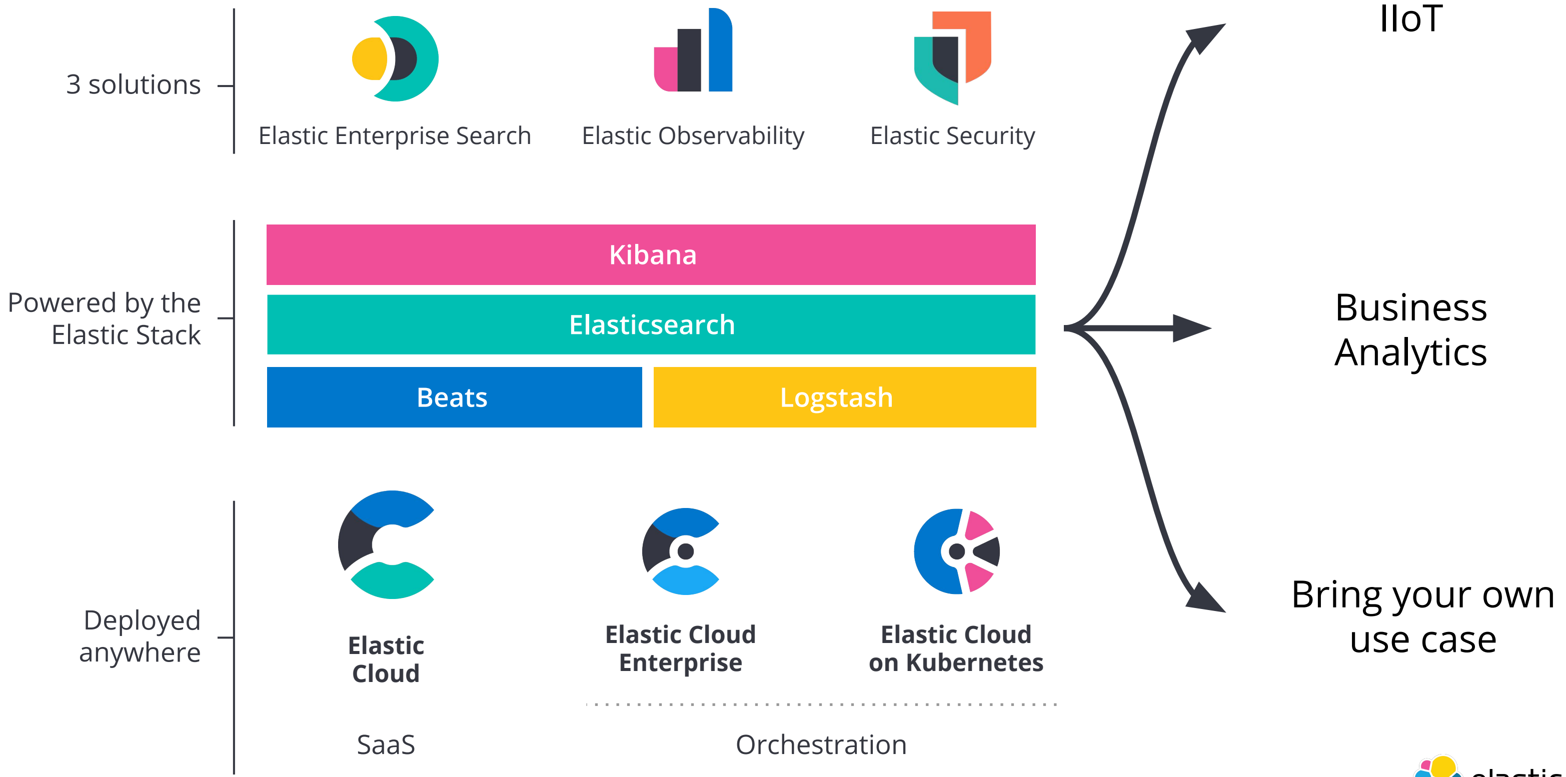

elastic

3 Steps of Optimizing your data use with Elastic Stack

Step 1: Storage	Step 2: Processing	Step 3: Automation

Step 1: Storage

Easy Integrations

Raw Data Processing

Enrichment

Granular level Security

Compliance (SOC2, CSA STAR, ISO/IEC 27001, ISO/IEC 27018, ISO IEC 27017, HIPAA, FeDRAMP)

Step 2: Processing

Fast, Scalable, and Relevant

Real-Time Aggregation

Data Lifecycle Management

Dynamic Visualisations

Step 3: Automation

Automation

Machine Learning Driven Analysis

- Anomaly detection
- Population Analysis
- Classification
- Prediction
- Forecasting

Real-Time Alerting

elastic

Single Stack Architecture

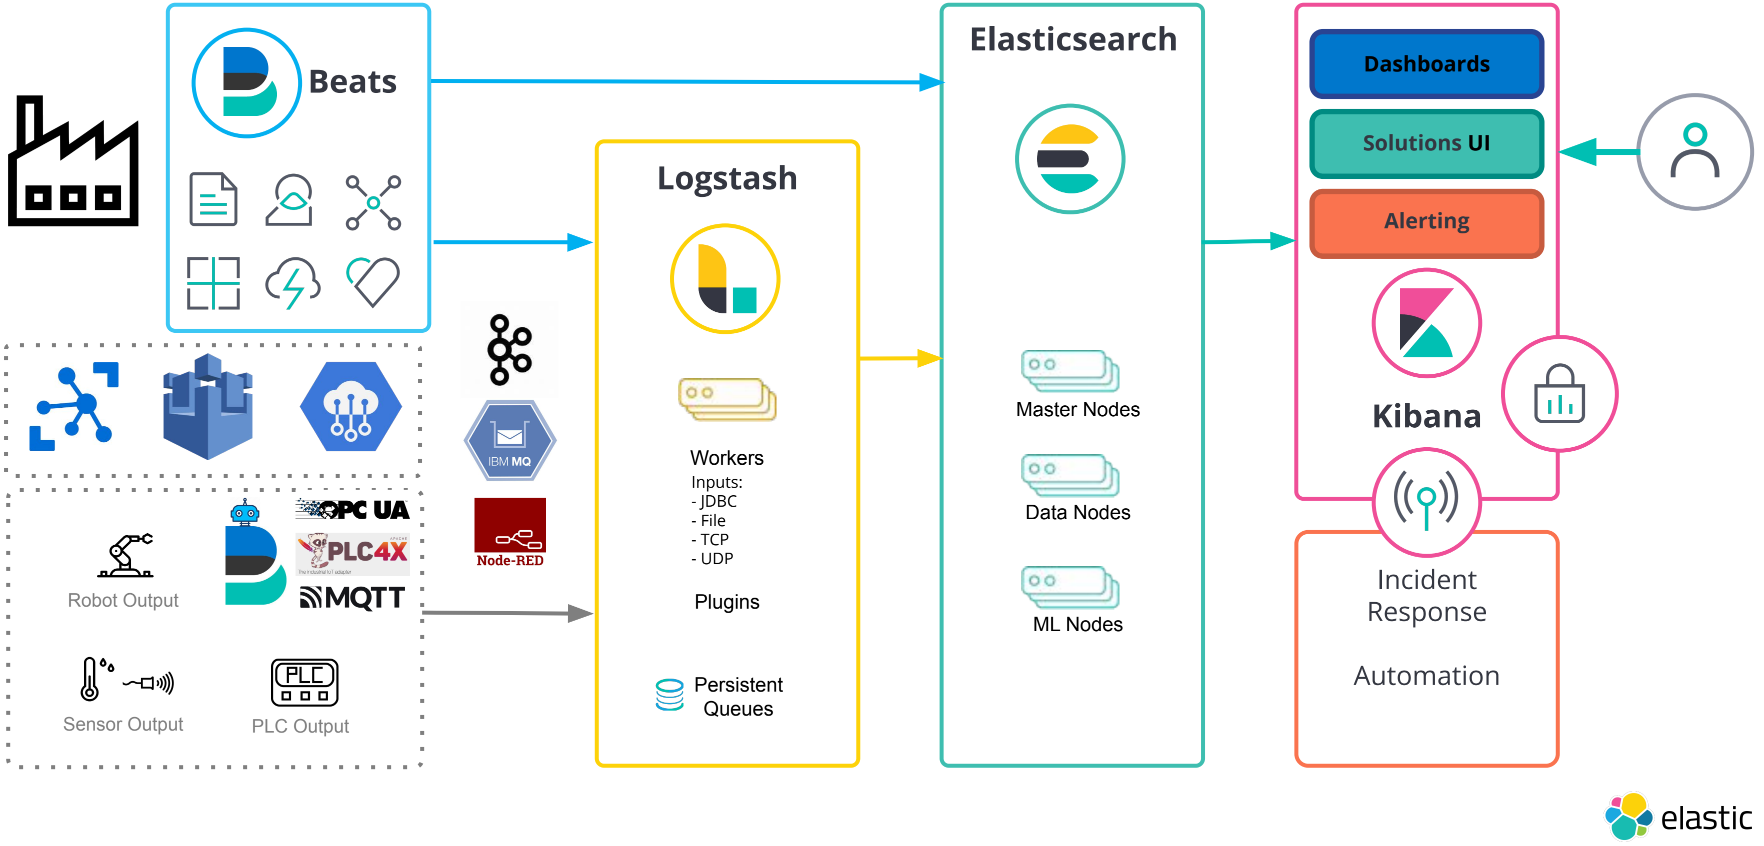

Elastic Machine Learning

<u>Operationalize</u> data science for <u>everyone</u>

Elastic Machine Learning

A tour of the Elastic ML stack

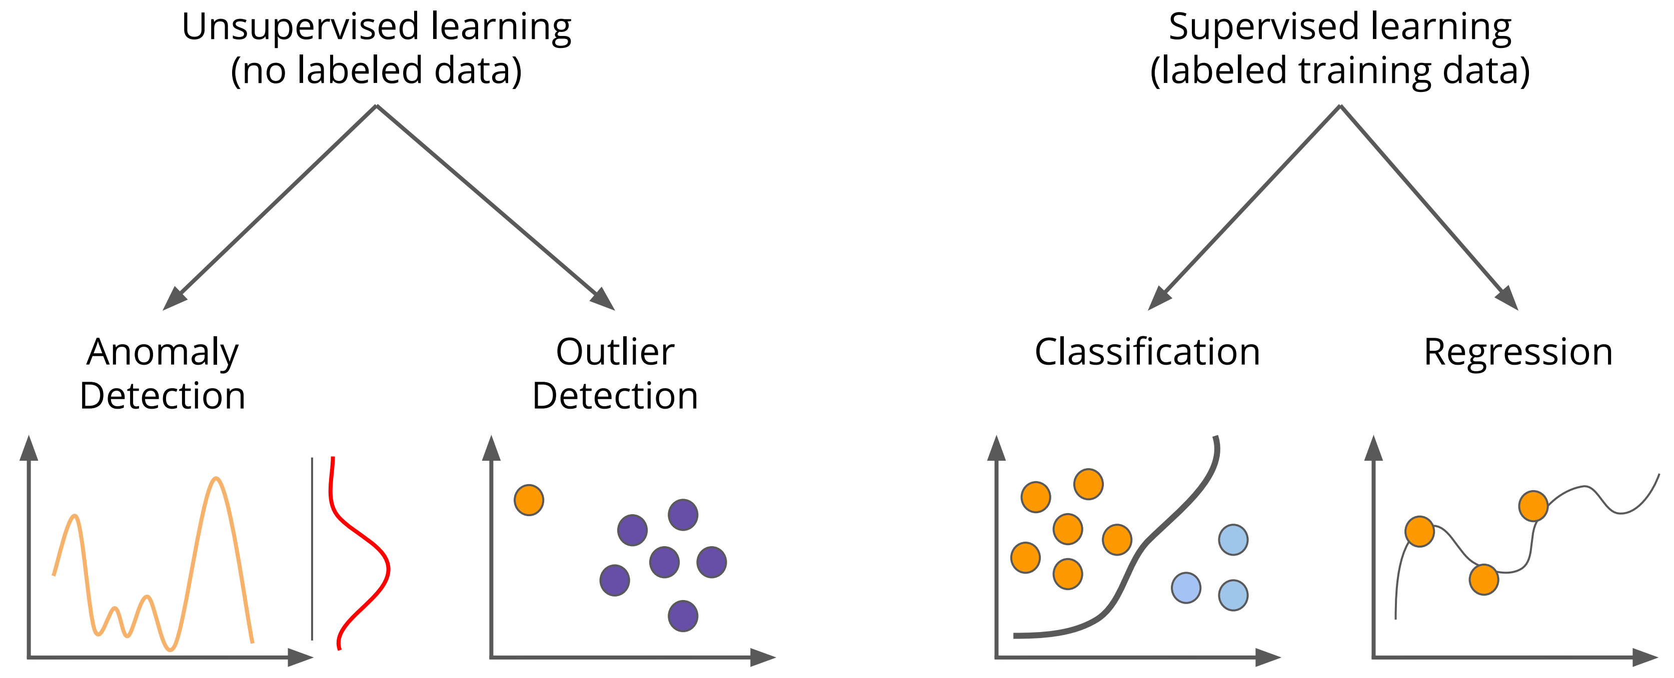

How to benefit from Elastic ML

Anomaly detection

- Is there any uncommon problem in my production process?
- What about the current traffic situation? Anything exceptional?
- Is there suspicious activity in my network environment?

Classification

- Does a train / car / truck / machine / robot needs maintenance?
- Does a device works correctly?

Outlier detection

- What are the products with the biggest change of being unusable based on production sensor data?
- Do we have cheater / fraudster in our system?

elastic

Elastic Machine Learning Flow for anomaly detection

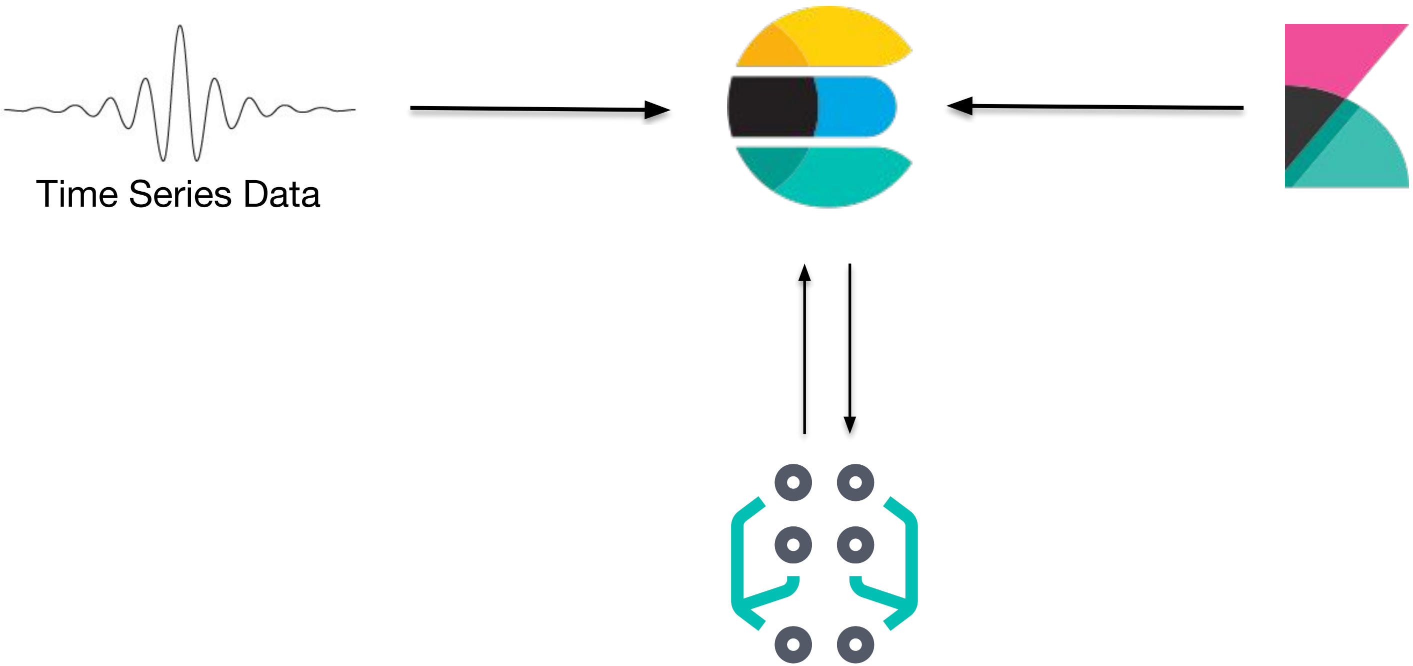

elastic

Let's build a model for anomaly detection

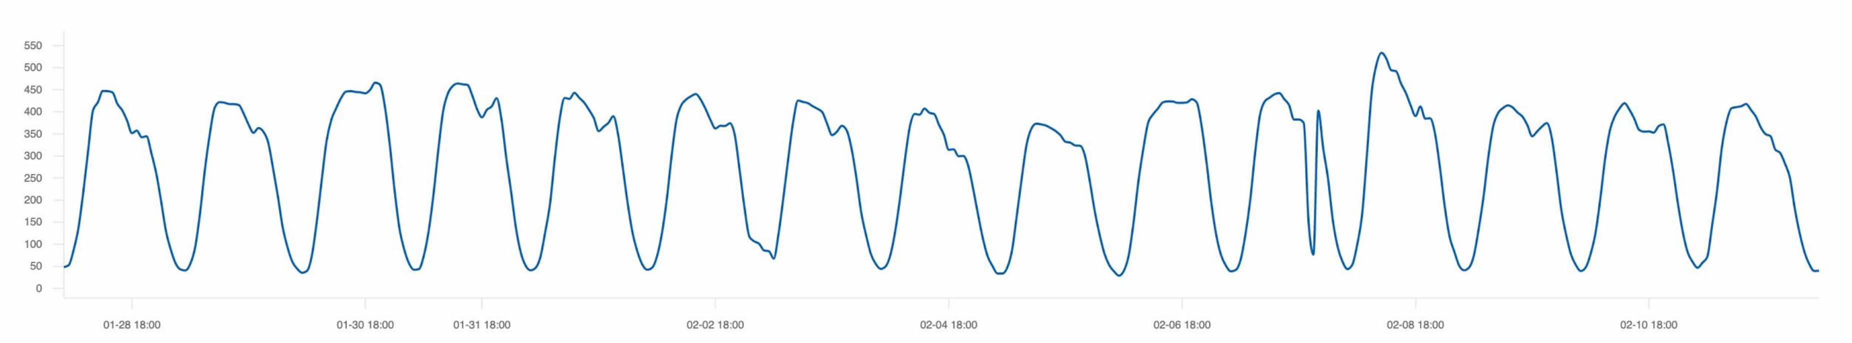

What kinds of patterns can we find in time series data?

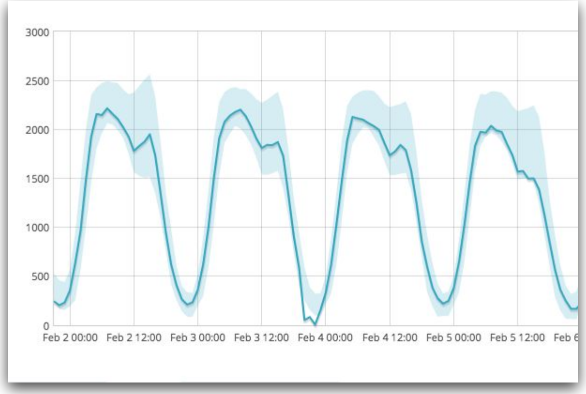

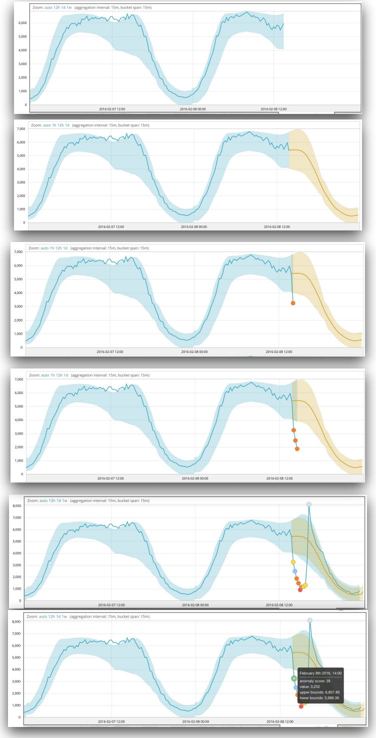

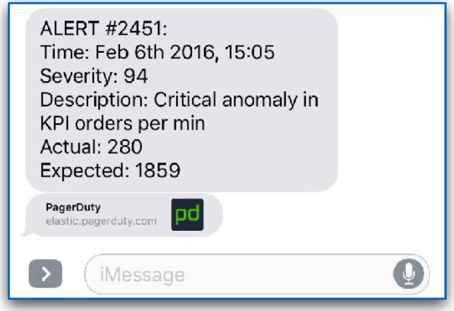

Learn

Predict

Operationalize

elastic

ML Data Frame Analytics | Predict customer churn

End to end methodology

Define a ML problem and propose a solution → Construct your dataset → Transform data → Train a model → Use the model to make predictions

raw logs

aggregated data


```
{ "customer_id": "028fa21e", "session_id": "MA0l6PC5", "@timestamp":
"2019-05-08T18:46:22", "request_type": "streaming_tv", "channel": "bbc",
"title": "Line of Duty" }
{ "customer_id": "a4ca7c7c", "session_id": "LMSXQXHg", "@timestamp":
"2019-05-08T18:49:34", "request_type": "streaming_film", "channel": "ziggo",
"title": "Glass" }
{ "customer_id": "avad97s3", "session_id": "LMSXQXHg", "@timestamp":
"2019-05-08T18:50:34", "request_type": "streaming_film", "channel": "ziggo",
"title": "Glass" }
{ "customer_id": "dce909a0", "session_id": "MA0l6PC5", "@timestamp":
"2019-05-08T18:51:23", "request_type": "streaming_film", "channel": "ziggo",
"title": "Glass" }
{ "customer_id": "vfva09a09", "session_id": "LMSXQXHg", "@timestamp":
"2019-05-08T18:52:14", "request_type": "streaming_film", "channel": "ziggo",
"title": "Glass" }
{ "customer_id": "sdfd9s90", "session_id": "MA0l6PC5", "@timestamp":
"2019-05-08T18:54:51", "request_type": "streaming_film", "channel": "ziggo",
"title": "Glass" }
...
```


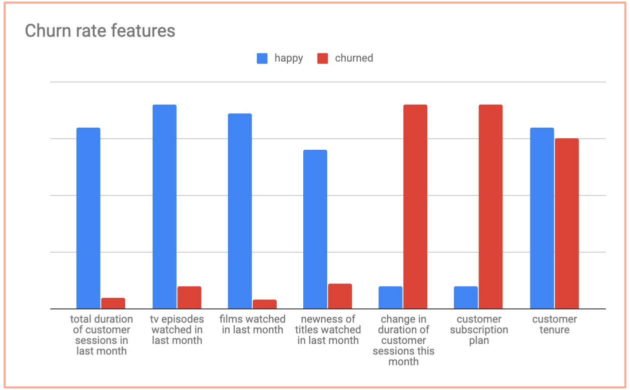

elastic

Machine Learning end-to-end methodology

Transform raw data to a feature index

Define a ML problem and propose a solution → Construct your dataset → **Transform data** → Train a model → Use the model to make predictions

RAW Data

```
{
"customer_id": "028fa21e",
"session_id": "MA0l6PC5",
"@timestamp": "2019-05-08T18:46:22",
"request_type": "streaming_tv",
"channel": "bbc",
"title": "Line of Duty"
},
{
"customer_id": "a4ca7c7c",
"session_id": "LMSXQXHg",
"@timestamp": "2019-05-08T18:49:34",
"request_type": "streaming_film",
"channel": "ziggo",
"title": "Glass"
},
...
```

→

```
PUT _transform/customer_behaviour
{
 "source": {
   "index": ["viewing_logs"]
 },
 "description": "Pivot viewing logs to customer-centric index",
 "dest": {"index": "customer_behaviour"},
 "pivot": {
   "group_by": {
     "customer_id": {"terms":{"field": "customer_id"}
     }
   },
   "aggregations": {
     "total_tv_shows": {...},
     "total_films": {...},
     ...
   }
 }
}
```

→

Customer Index

```
{
"customer_id": "028fa21e",
"total_tv_shows": 10,
"total_films": 2,
"total_watching_duration": 72123,
"last_active": "019-05-08T18:46:22",
...
},
{
"customer_id": "a4ca7c7c",
"total_tv_shows": 23,
"total_films": 8,
"total_watching_duration": 184212,
"last_active": "2019-05-08T18:49:34",
...
},
...
```

elastic

Machine Learning end-to-end methodology

Build a model on historical data that has a churn indicator

Define a ML problem and propose a solution → Construct your dataset → Transform data → Train a model → Use the model to make predictions

train/validate/test

	customer a	customer b
total duration of customer sessions	80:21:07	1:01:11
tv episodes watched	24	1
films watched in last month	5	0
newness of titles watched in last month	9.8	1.2
Change in duration	6:22:17	16:43:29
subscription plan	gold	platinum
customer tenure	32	26
has churned?	no	yes

ML Supervised Model

Model Name: churn_e2r21
Model Precision: 96.3%
Model Recall: 95.7%
Model F1 score: 96.0%

elastic

Machine Learning end-to-end methodology

Use model inference to make predictions on streaming data

Define a ML problem and propose a solution → Construct your dataset → Transform data → Train a model → Use the model to make predictions

predict

	customer c
total duration of customer sessions	10:10:06
tv episodes watched	2
films watched in last month	1
newness of titles watched in last month	1.6
change in duration this month	17:22:17
customer plan	gold
customer tenure	5

→ ML Supervised Model →

	customer c	Feature Influence
total duration of customer sessions	10:10:06	0.1
tv episodes watched	2	0.8
films watched in last month	1	0.8
newness of titles watched in last month	1.6	0.01
change in duration of this month	17:22:17	0.6
customer plan	gold	0.01
customer tenure	5	0.1
will churn?	p(churn) = 97%	

elastic

Real world example

elastic

Powering the Search for Real-Time Costing

"Being a market leader means being a tech leader. Our ability to quickly detect deviations in the manufacturing process and adjust on the fly is a competitive advantage."

Jürgen Kerner

Head of Operations, Corporate IT at Mayr-Melnhof Karton

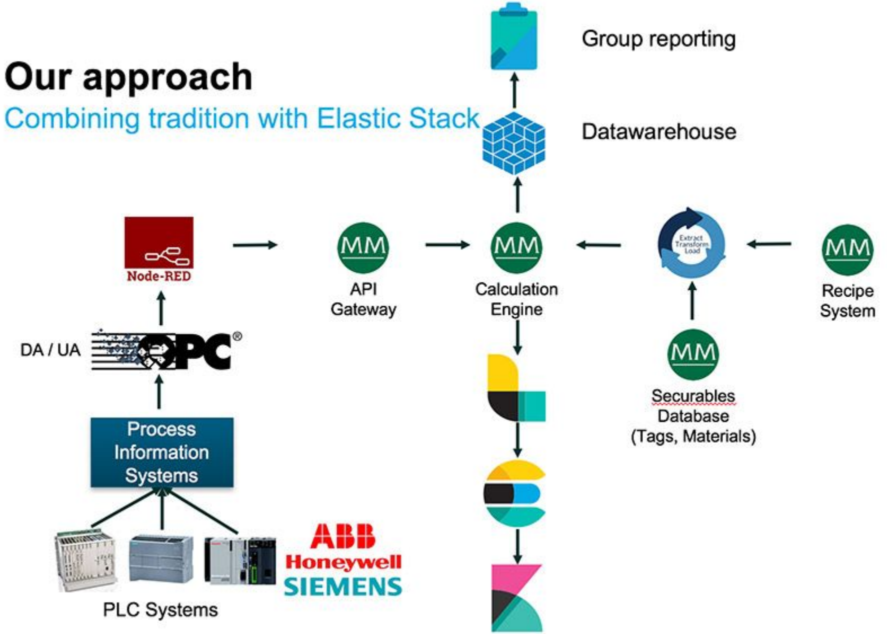

https://www.elastic.co/customers/mm-karton

elastic

Thank You

Elastic is a Search Company.

www.elastic.co