19/10/2017, Usergroup 2017

Next Generation Performance Monitoring
Machine Learning Algorithms for Anomaly Detection

Susanne Greiner

How to monitor performance?

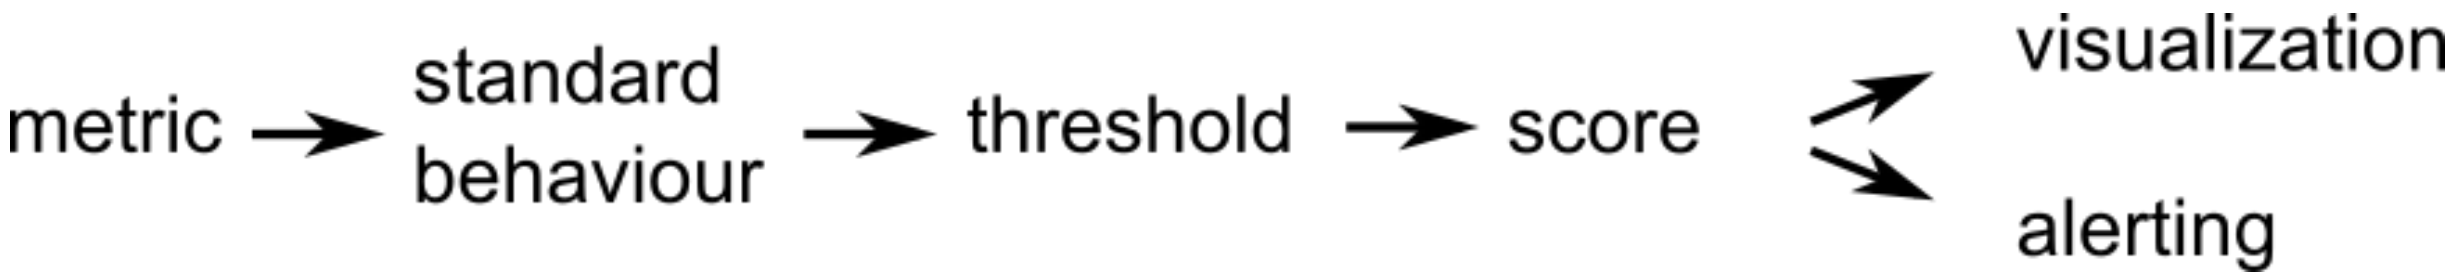

The right decision at each step is not trivial!

Data collection ≠ Problem solution

Performance Monitoring

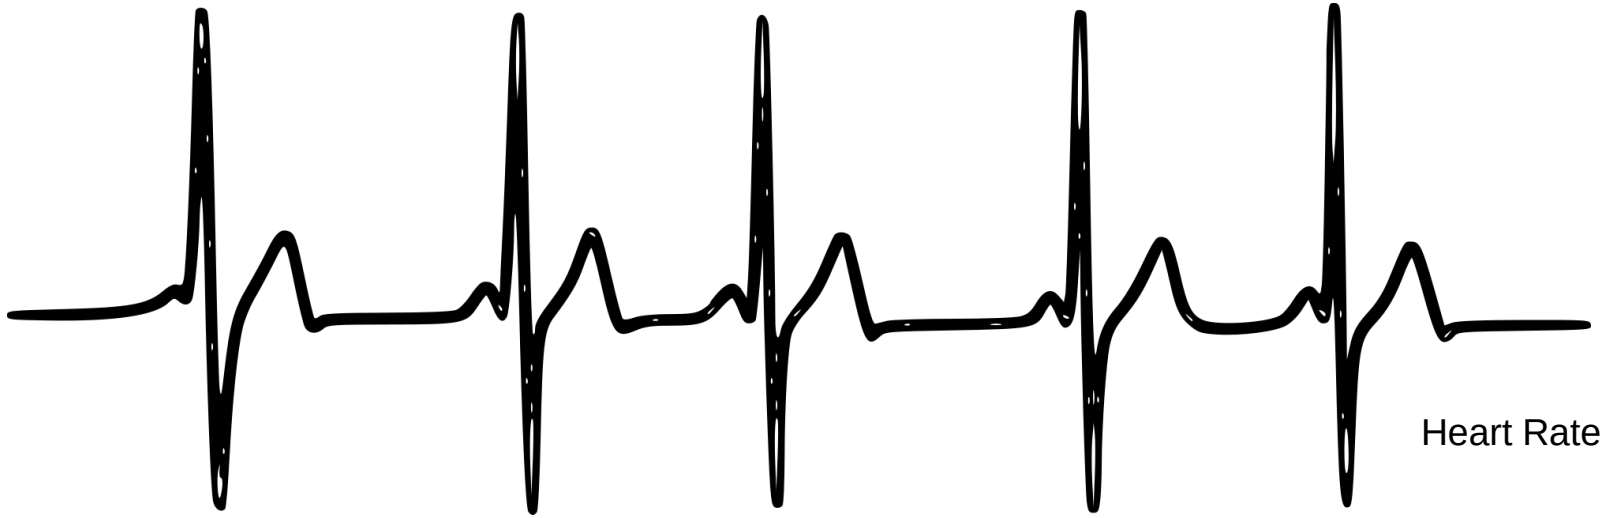

Can be influenced by

- Pathology
- Sport
- Breathing
- Drugs
- Temperature
- Dehydration
- Pressure
- Etc.

Monitoring & Alarms

Subject specific historical data

Expectation

- Time series
- Alarm thresholds

Population data

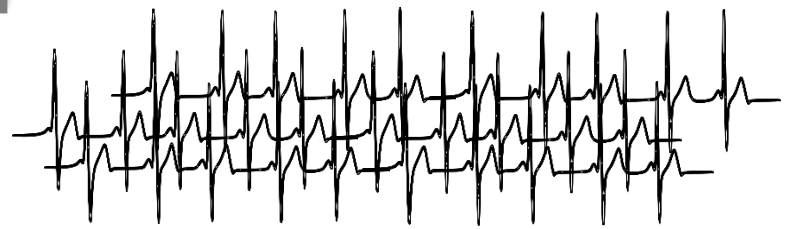

Performance Monitoring

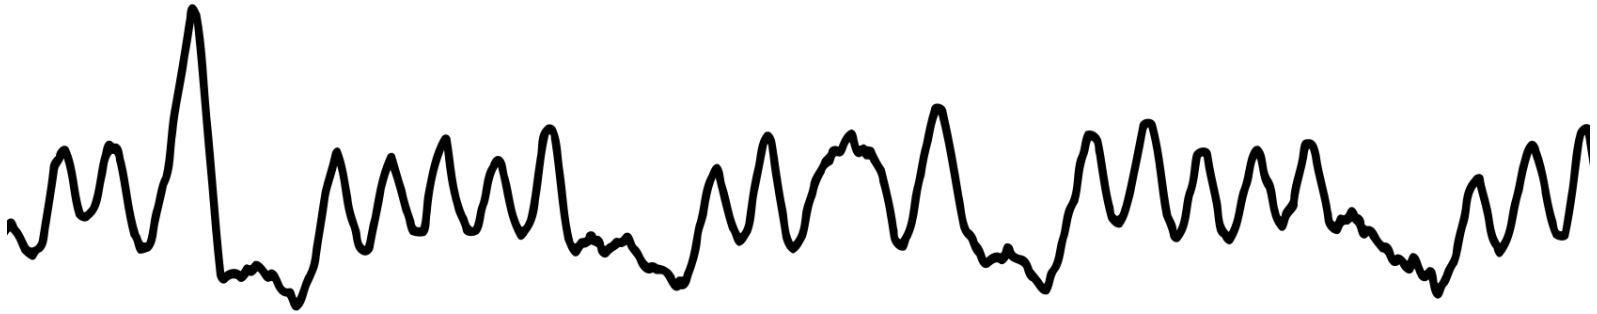

Percentage Processor Time

Can be influenced by

- Batch requests
- Transactions
- Memory
- SAN
- Network
- Side Processes
- Etc.

Monitoring & Alarms

Machine/ setting specific historical data

Expectation

- Time series
- Alarm thresholds

Experience,
Data from similar machines/ settings

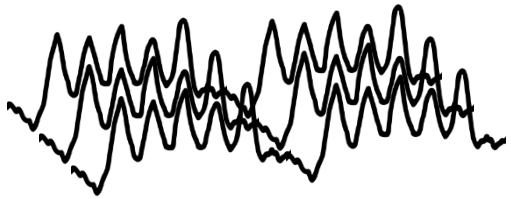

Recent Trends

Reactivity → Proactivity

Standard Stats → Advanced Stats & Machine Learning

Combination of Performance Monitoring and User Experience

Agile Implementation of Solutions

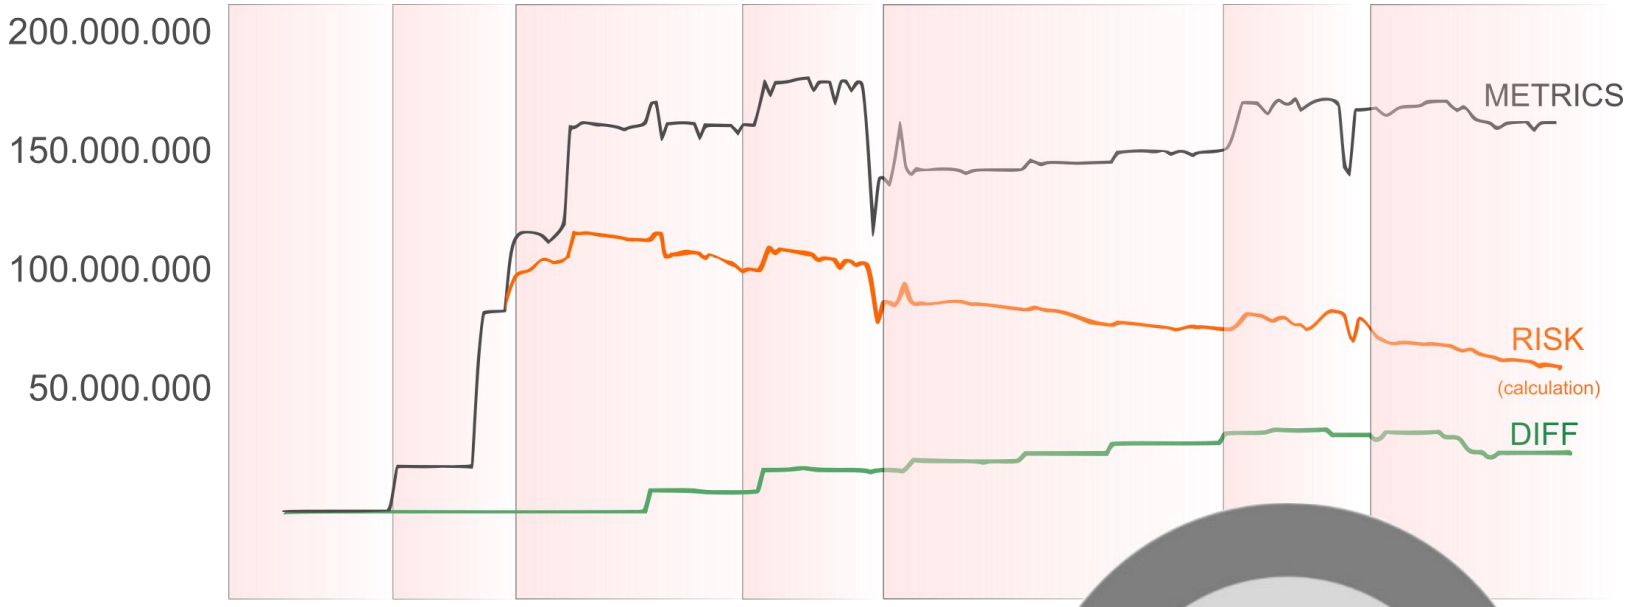

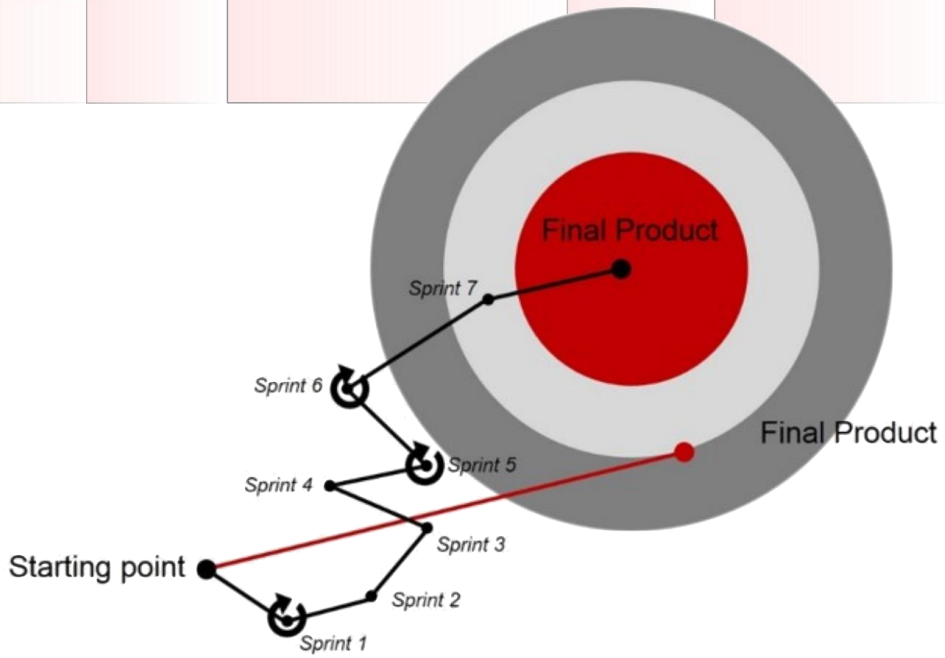

Customer specific solutions with agile implementation

REACTIVITY → PROACTIVITY

Visualization: Trend Detection

VORHER

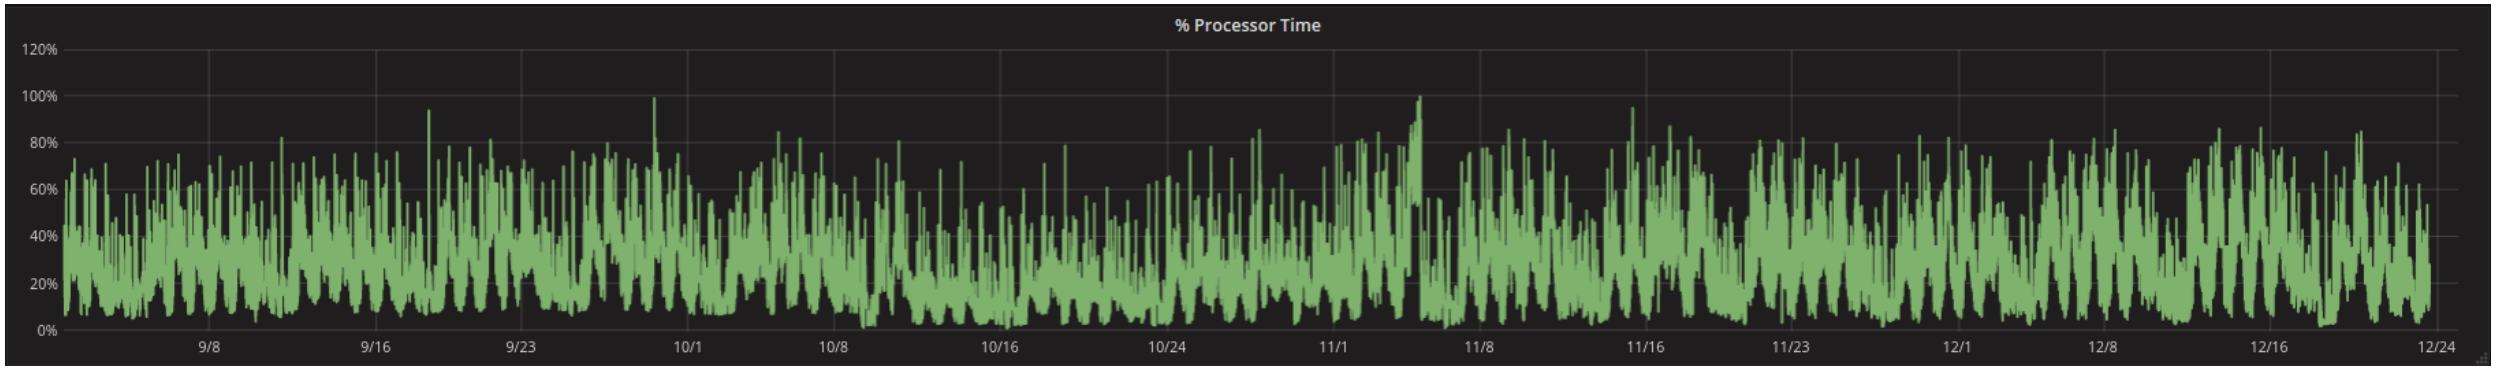

NACHHER

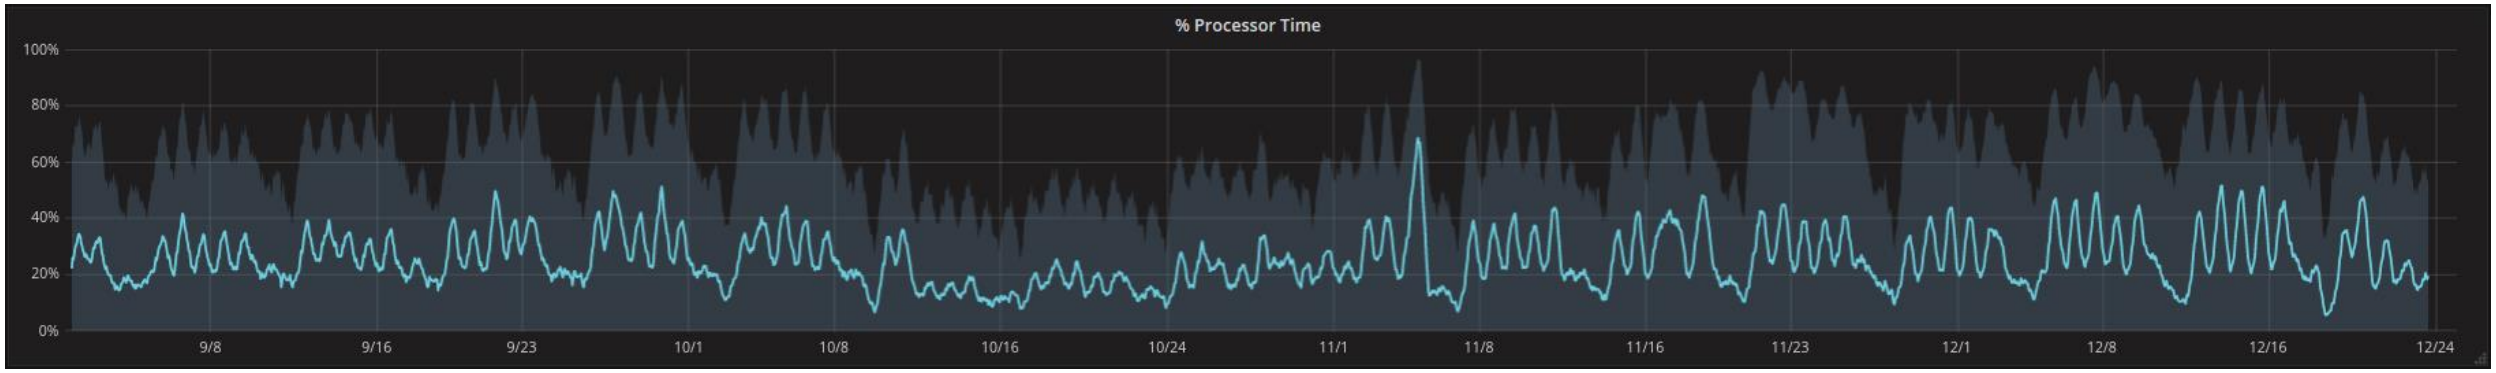

- Areas not to many points
- Smoothed signal

Historic Data

VORHER

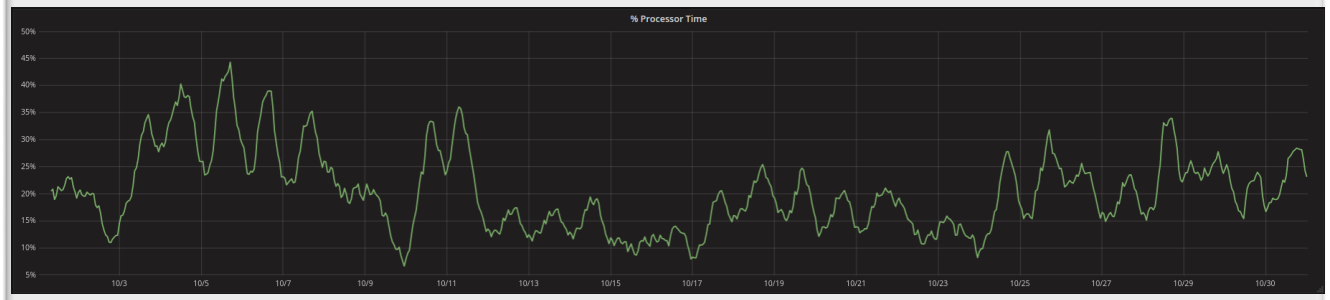

- Comparison time consuming
- Quantification complicated

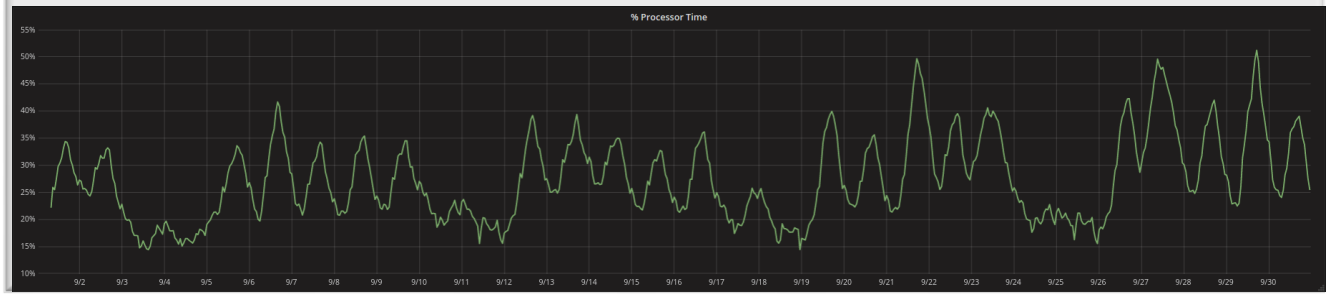

NACHHER

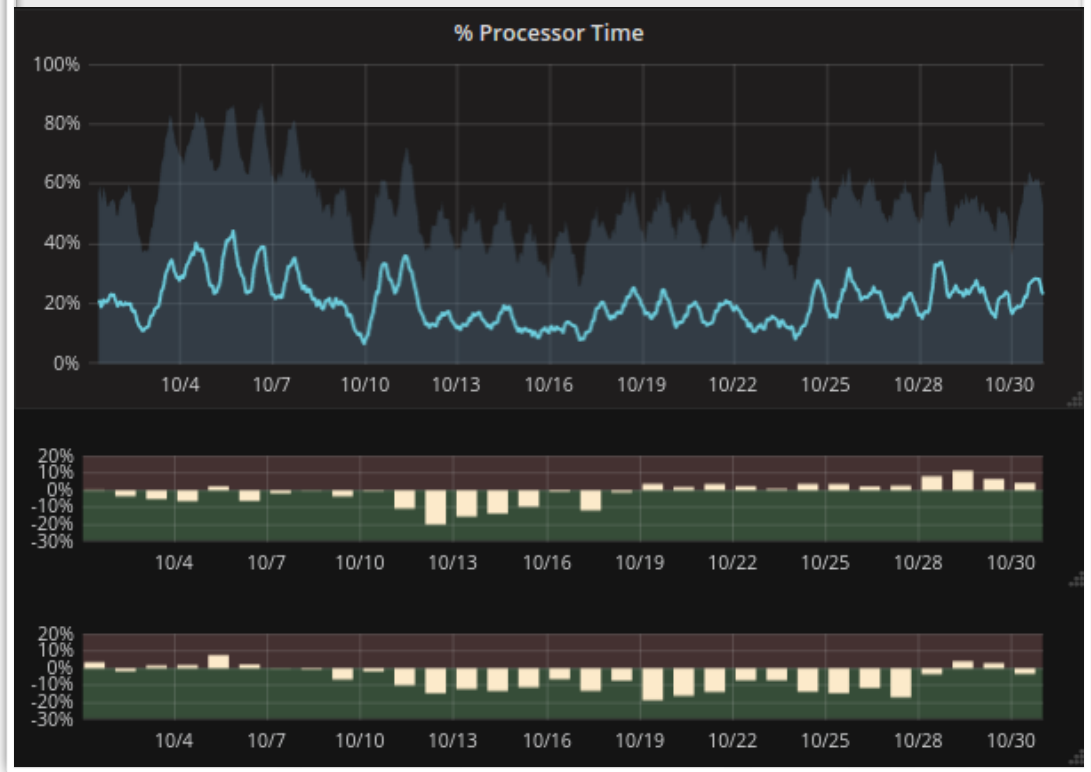

- Historic data at hand
- Visualization of differences
  => Easy trend detection

STANDARD STATS → ADVANCED STATS & ML

Univariate Data Analysis

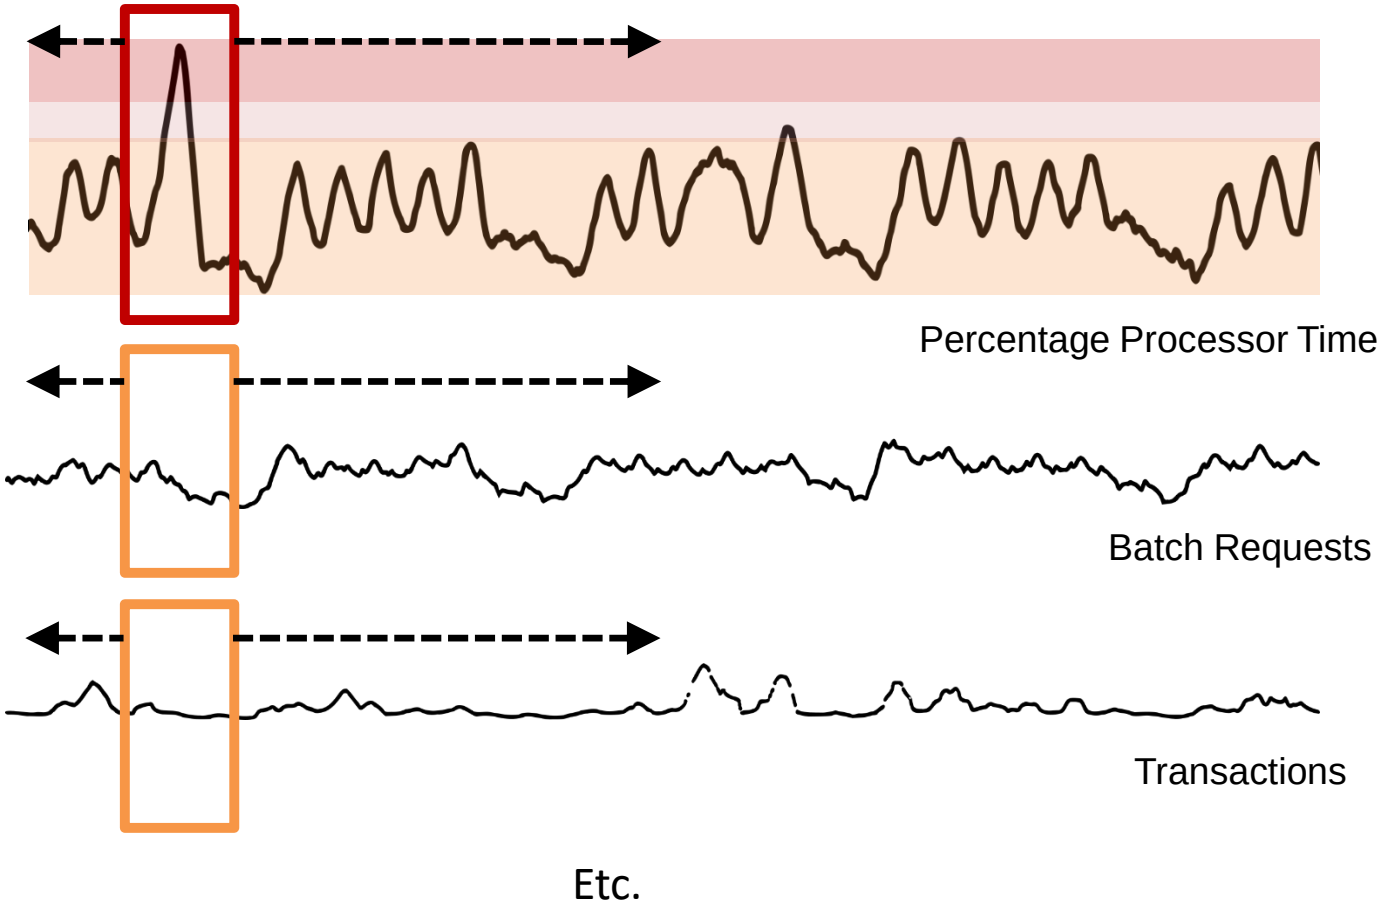

- Every time series is analyzed on **separately**
- Thresholds are calculated on via **baselining**
- Alarms from separate time series are combined into a **global alarm**
- Relationships between time series are **ignored**
- Shape of time series is **ignored**

Motivations

- Separate Data Sources
- Different Precision
- Evolution of networks (complexity)
- Common practice was enough

Multi-variate Data Analysis

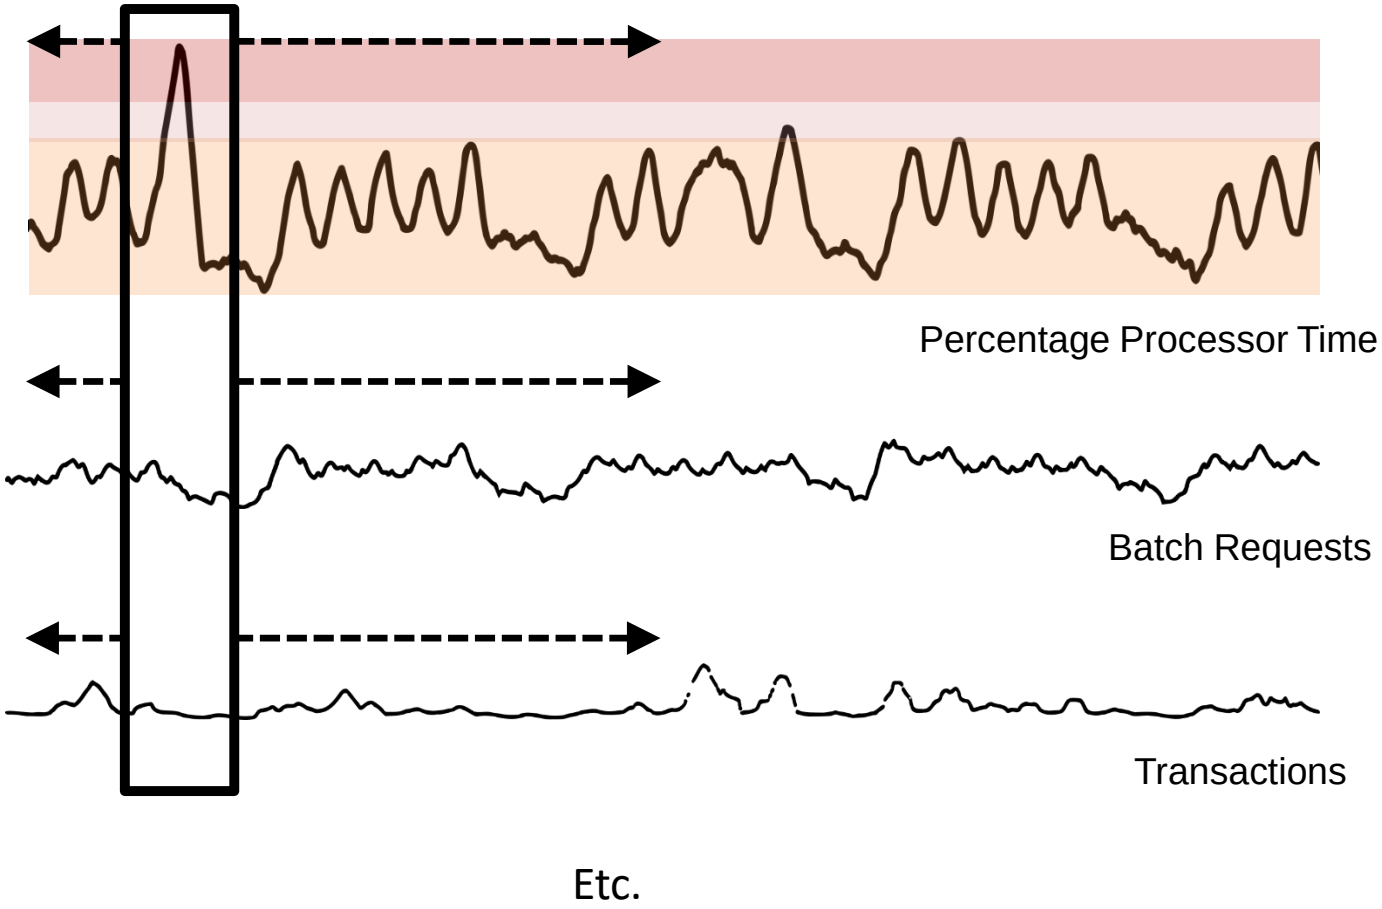

- All time series are analyzed on **together**
- Thresholds are calculated dynamically via **baselining and anomaly detection**
- **Risk estimation** in addition to **global alarm and specific alarms**
- Relationships between time series are used to create more reliable alarms and risks
- Shape of time series is **considered**

Motivations

- Common Data Source
- Grafana & InfluxDB
- Today we need more than common practice
- **Proactivity**

Anomaly vs. Threshold

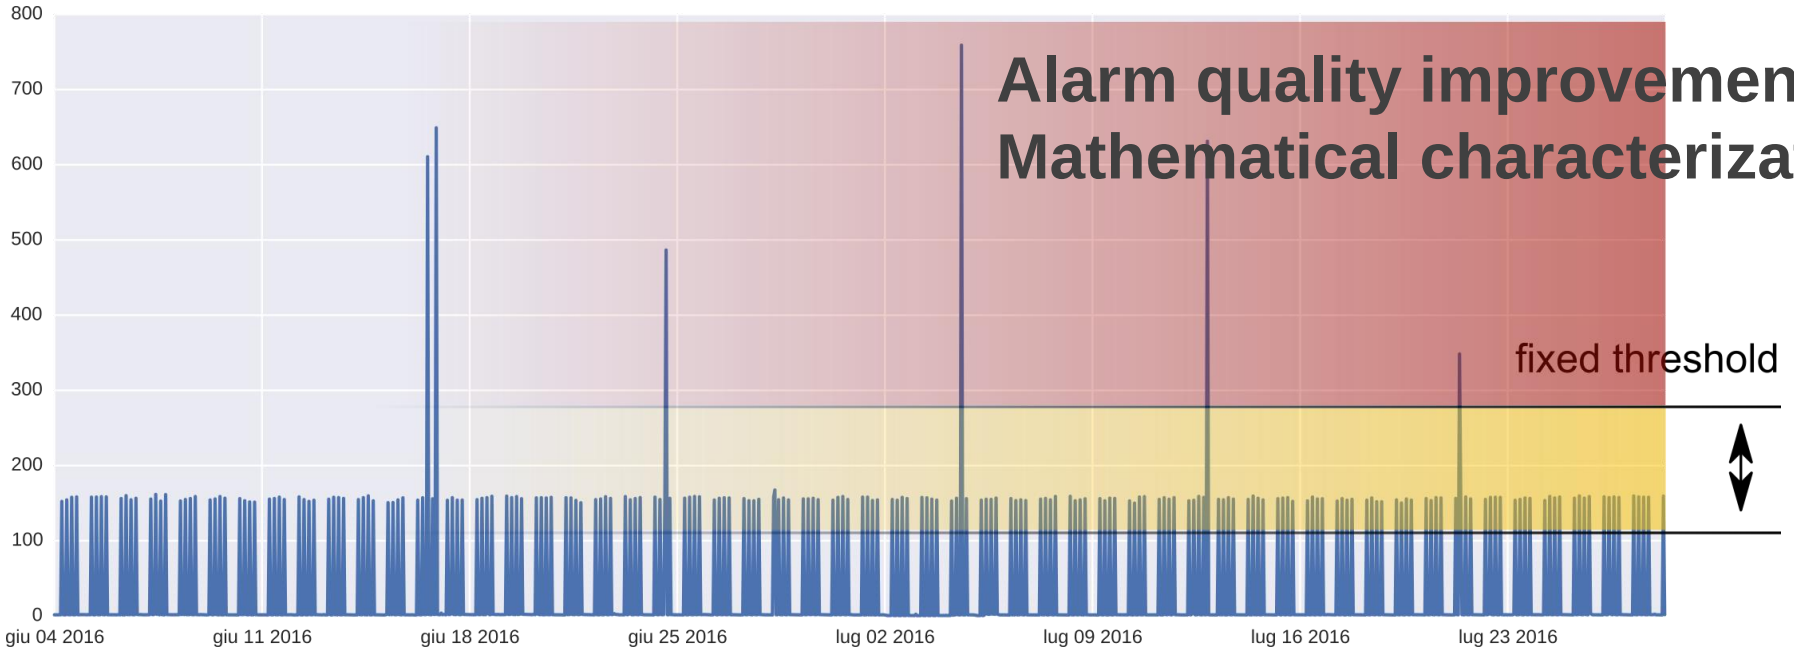

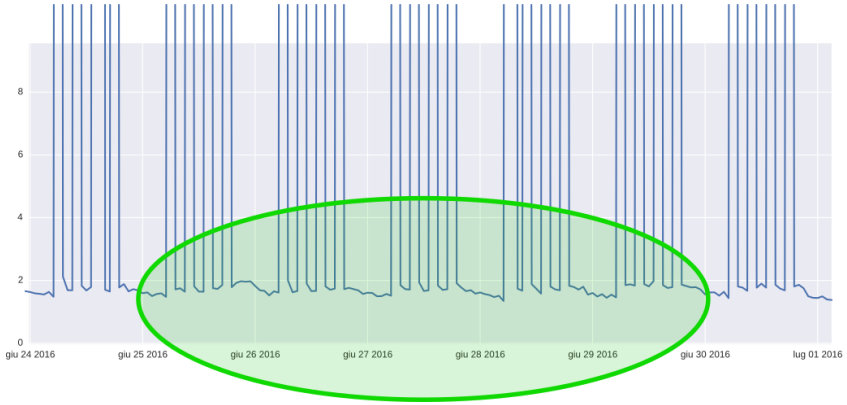

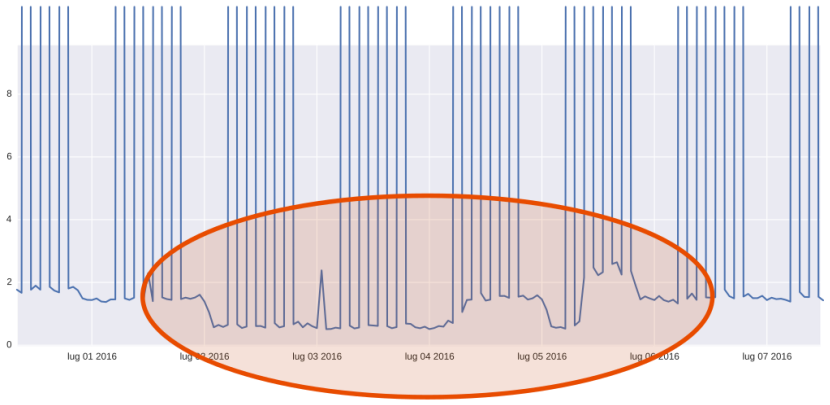

Anomaly vs. Threshold

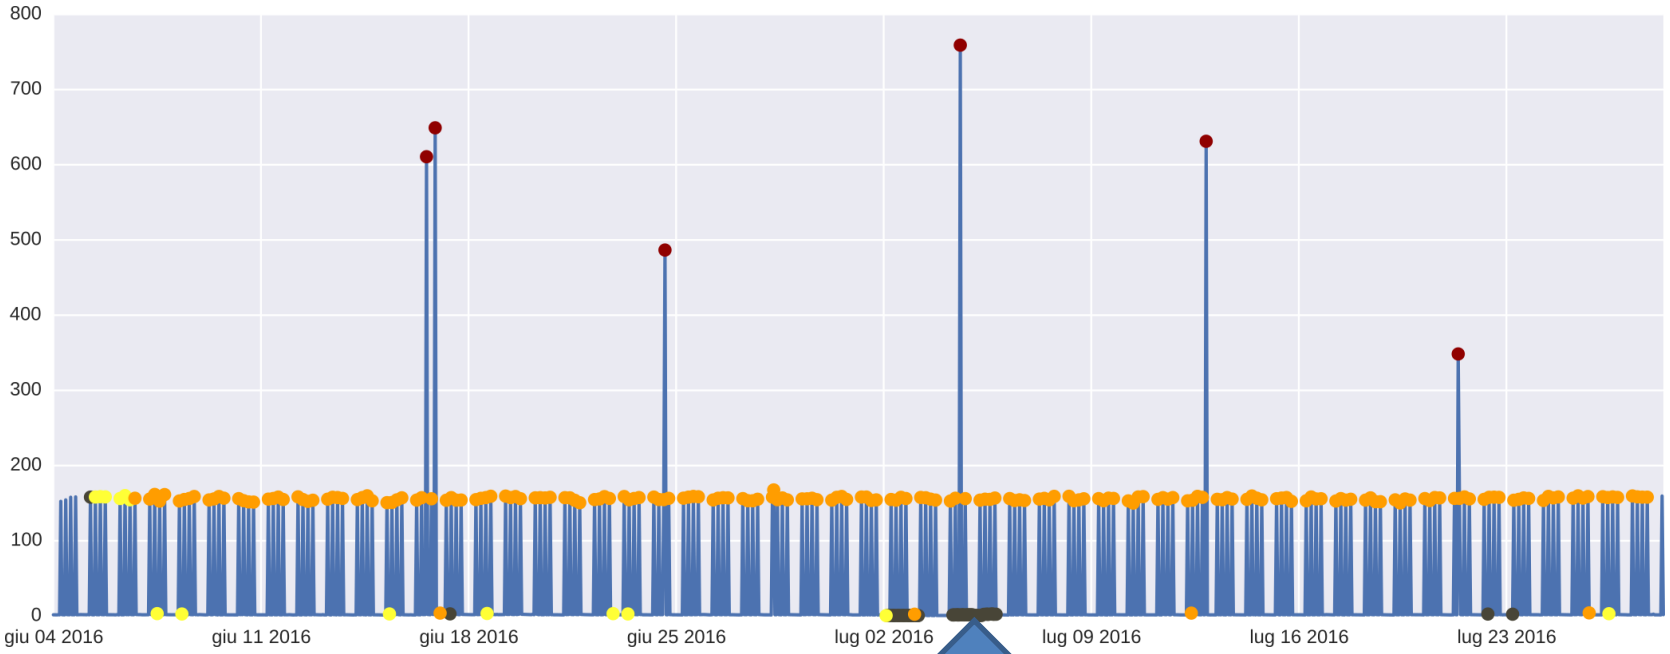

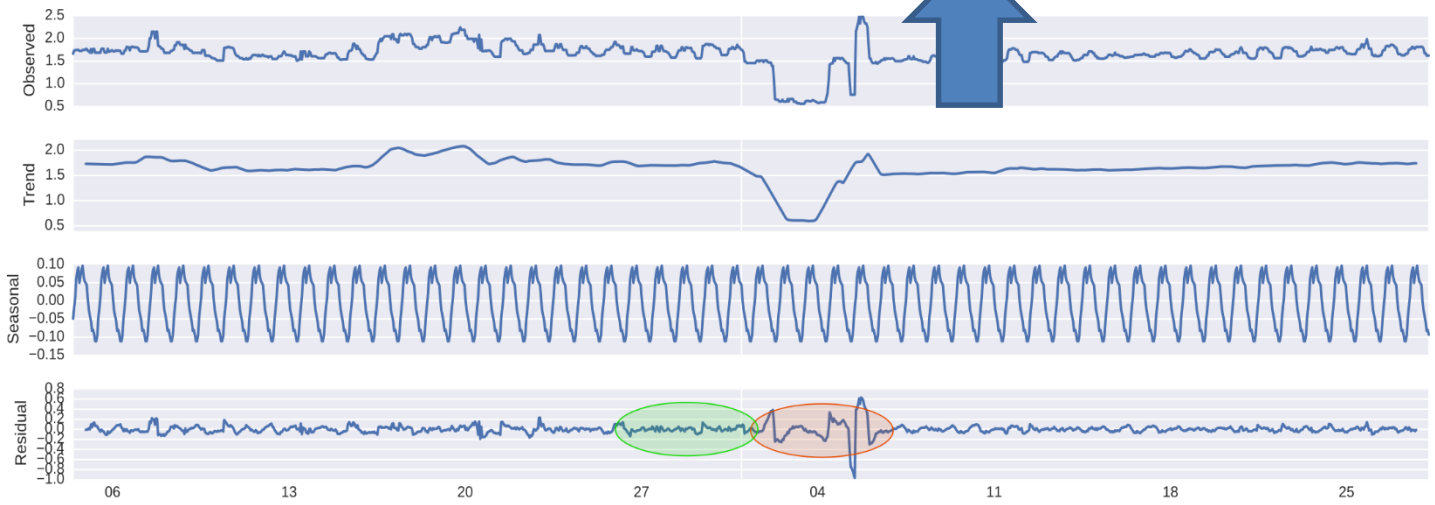

Automatic detection of relevant changes

Risk: Anomaly Detection via Multivariate ML Analysis

Metrics

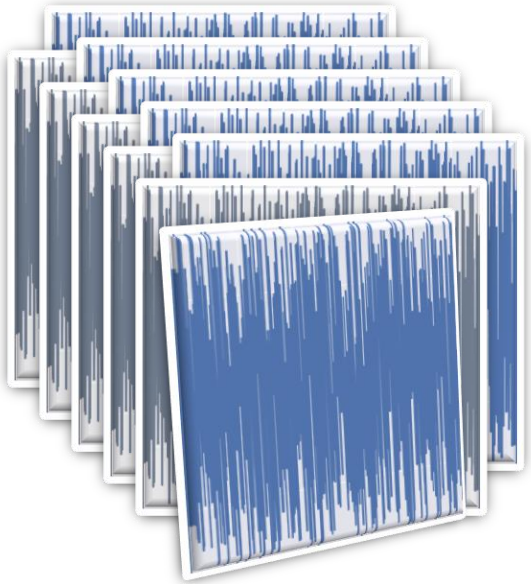

- The farer away from expectation the higher the RISK
- RISK: different and rare

HOW FAR ARE WE FROM OUR EXPECTATION?

MODEL
STANDARD
BEHAVIOUR

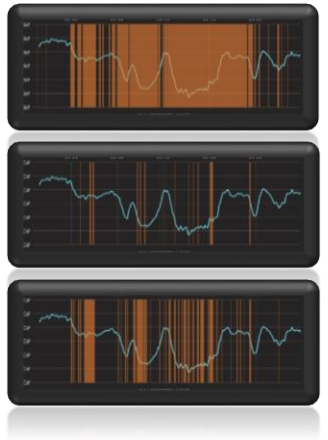

RISK
SCORE

- historical data of same metric
- historical data of similar metric
- historical data of similar machine

WHAT DO WE EXPECT?

Risk: Culprit Detection with Risk

IOs

- historical data of same metric
- historical data of similar metric
- historical data of similar machine
- historical data of neighbours

HIGH RISK PERIODS

AUTOMATED ANALYSIS

- Which (set of) machine(s) is most probably causing the high risk
- Proactive analysis to prevent congestion

Proactive search for potential future culprits

Faster Troubleshooting

- Check machines with high risk first, there might be no need to control the others

COMBINATION OF PERFORMANCE MONITORING AND USER EXPERIENCE

3 Levels of Dashboards

Level 1: Overview Dashboard

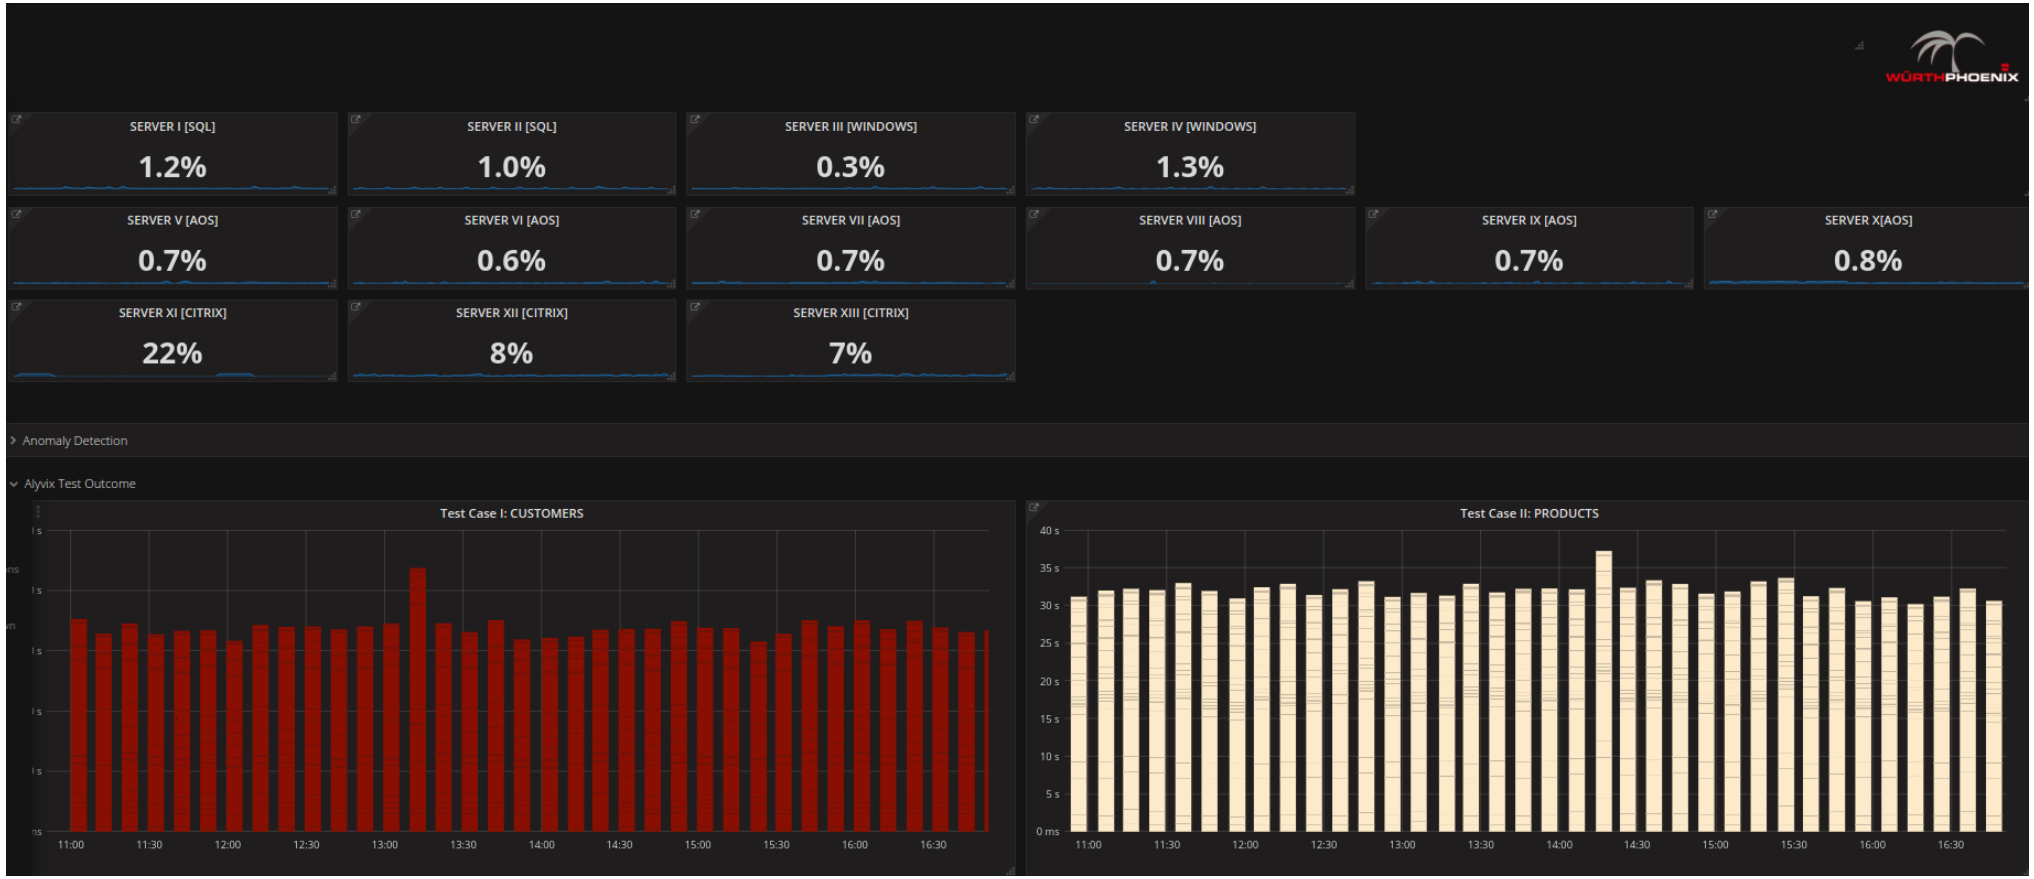

Server Overview

User Experience Overview

Level 2: Multimeasure - Troubleshooting

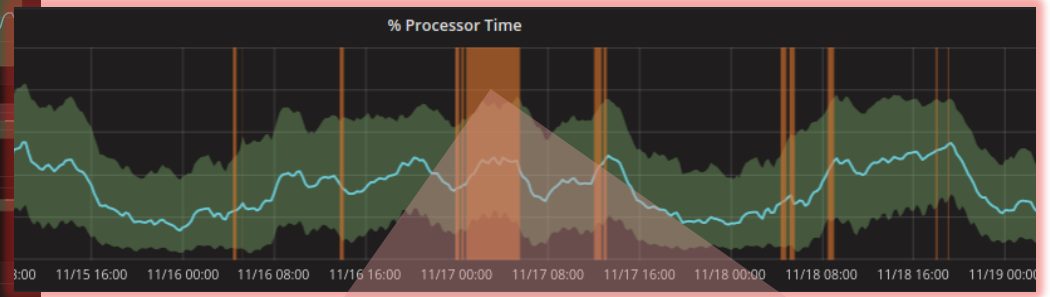

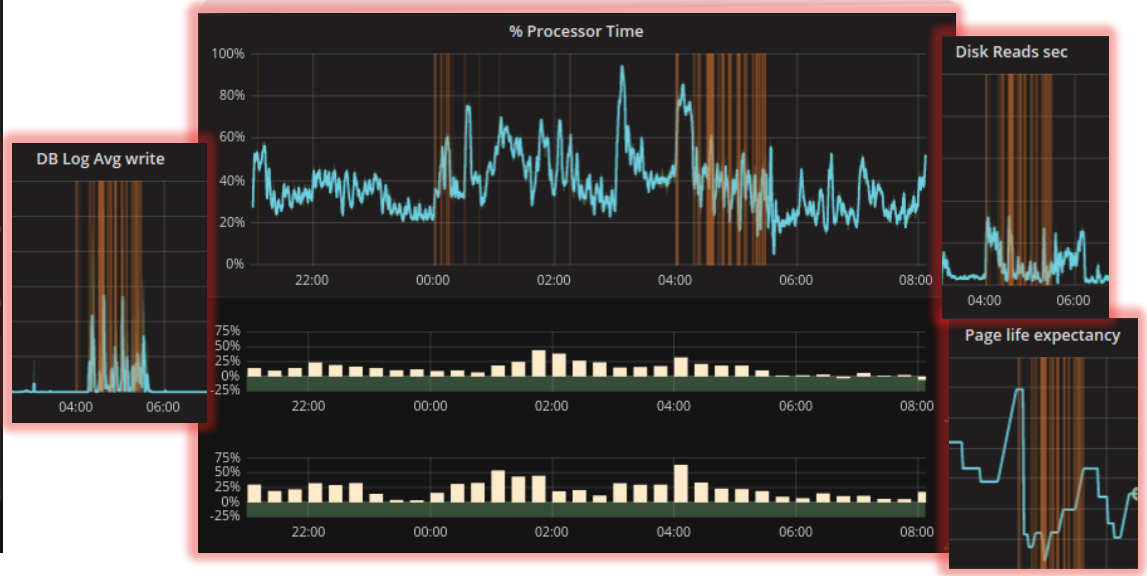

Detailed Dashboard (Alyvix)

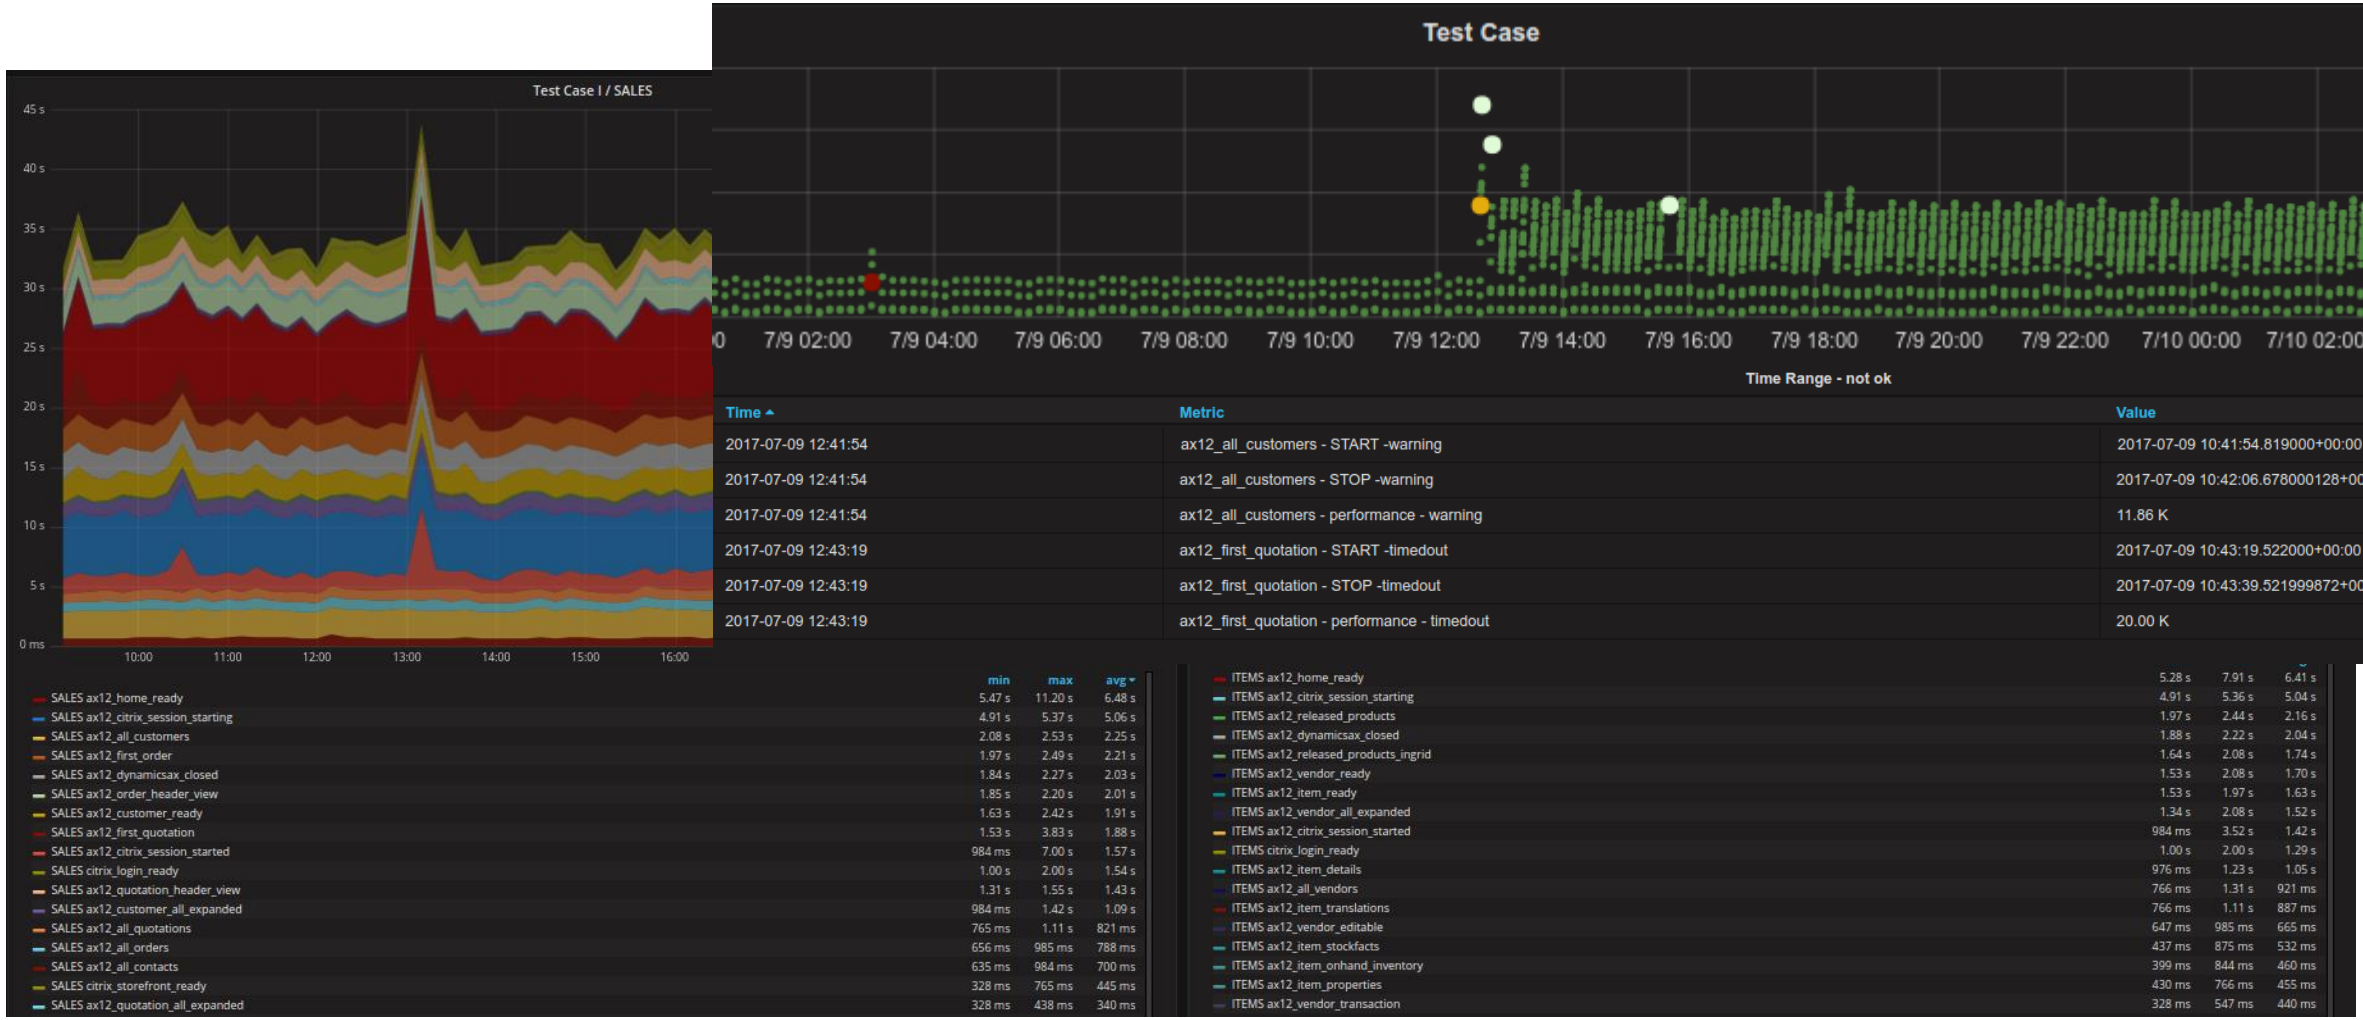

GRAZIE!