Joint 10th Working IEEE/IFIP Conference on Software Architecture &
6th European Conference on Software Architecture

WORKLOAD-AWARE SYSTEM MONITORING USING PERFORMANCE PREDICTIONS APPLIED TO A LARGE-SCALE E-MAIL SYSTEM

Christoph Rathfelder (rathfelder@fzi.de)
FZI Research Center for Information Technology Karlsruhe, Germany

Stefan Becker (stefan.becker@1und1.de)
1&1 Internet AG, Germany

Klaus Krogmann (krogmann@fzi.de)
FZI Research Center for Information Technology Karlsruhe, Germany

Ralf Reussner (reussner@kit.edu)
Karlsruhe Institute of Technology, Germany

Agenda

1. Quick Facts about the 1&1 Mail System
2. Motivation and Approach
3. Foundation
    1. Model-based Performance Predictions
    2. Palladio Approach
4. Case Study
5. Validation and Applicability of our Approach
6. Lessons Learned
7. Conclusion and Future Work

The 1&1 Mail System

Quick Facts

- Basis for several brands like

GMX® WEB.DE mail.com india.com

- > 40 Million users
- > 200 Million emails per day

- Build up as service oriented system
  - More than 100 software services
  - Distributed over 2,000 hosts

Motivation

- Ensuring high availability with guaranteed quality of service
  - Detailed monitoring of the system
- Potential reasons for deviating resource utilisation
  - Misbehaving loadbalancer
  - Hardware failures
  - Service unavailability
  - Connection problems
- Identified reason was

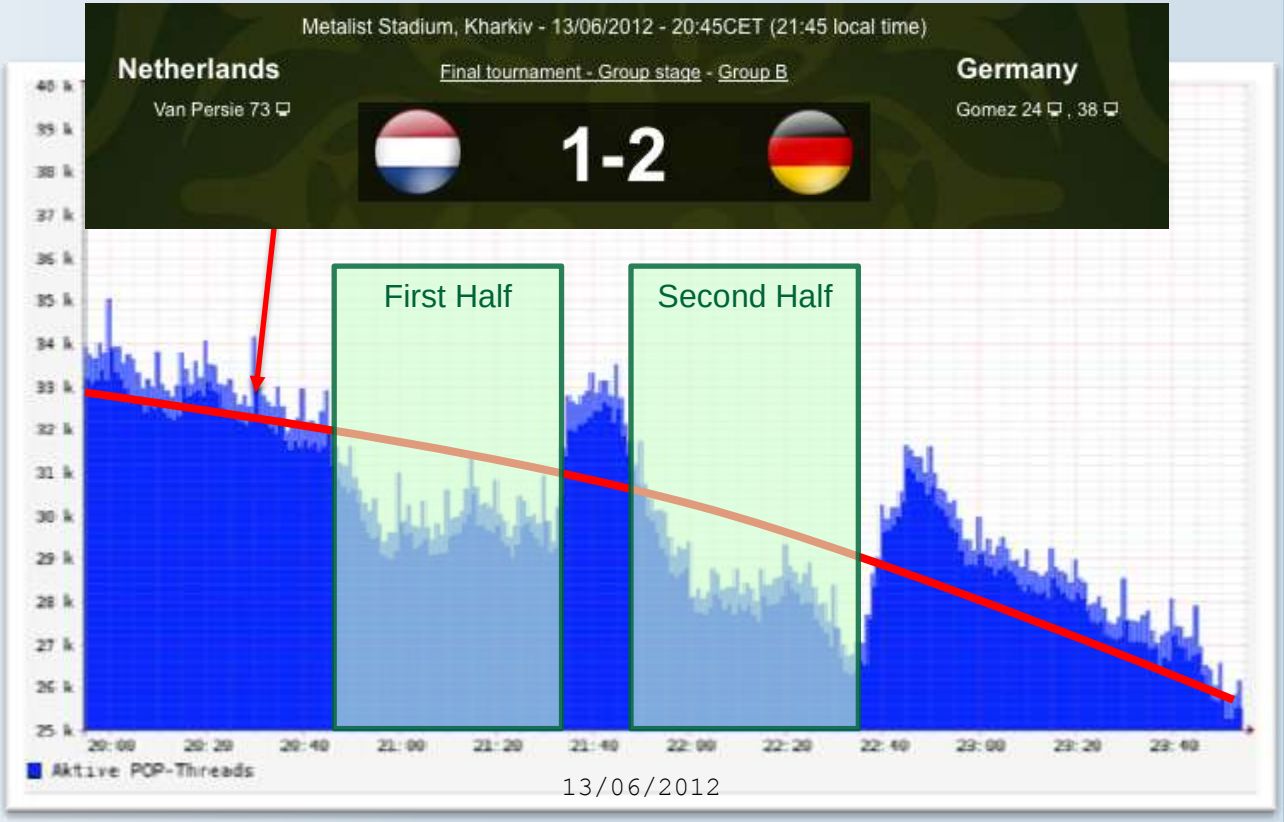

Motivation

- Ensuring high availability with guaranteed quality of service
  - Detailed monitoring of the system
- Potential reasons for deviating resource utilisation
  - Misbe… loadb…
  - Hardv…
  - Servic…
  - Conn…
- Identifiec…

System monitoring should be aware of the current usage profile

Approach

- Applying **performance prediction** to derive the **expected system behaviour** depending on the **monitored user behaviour**

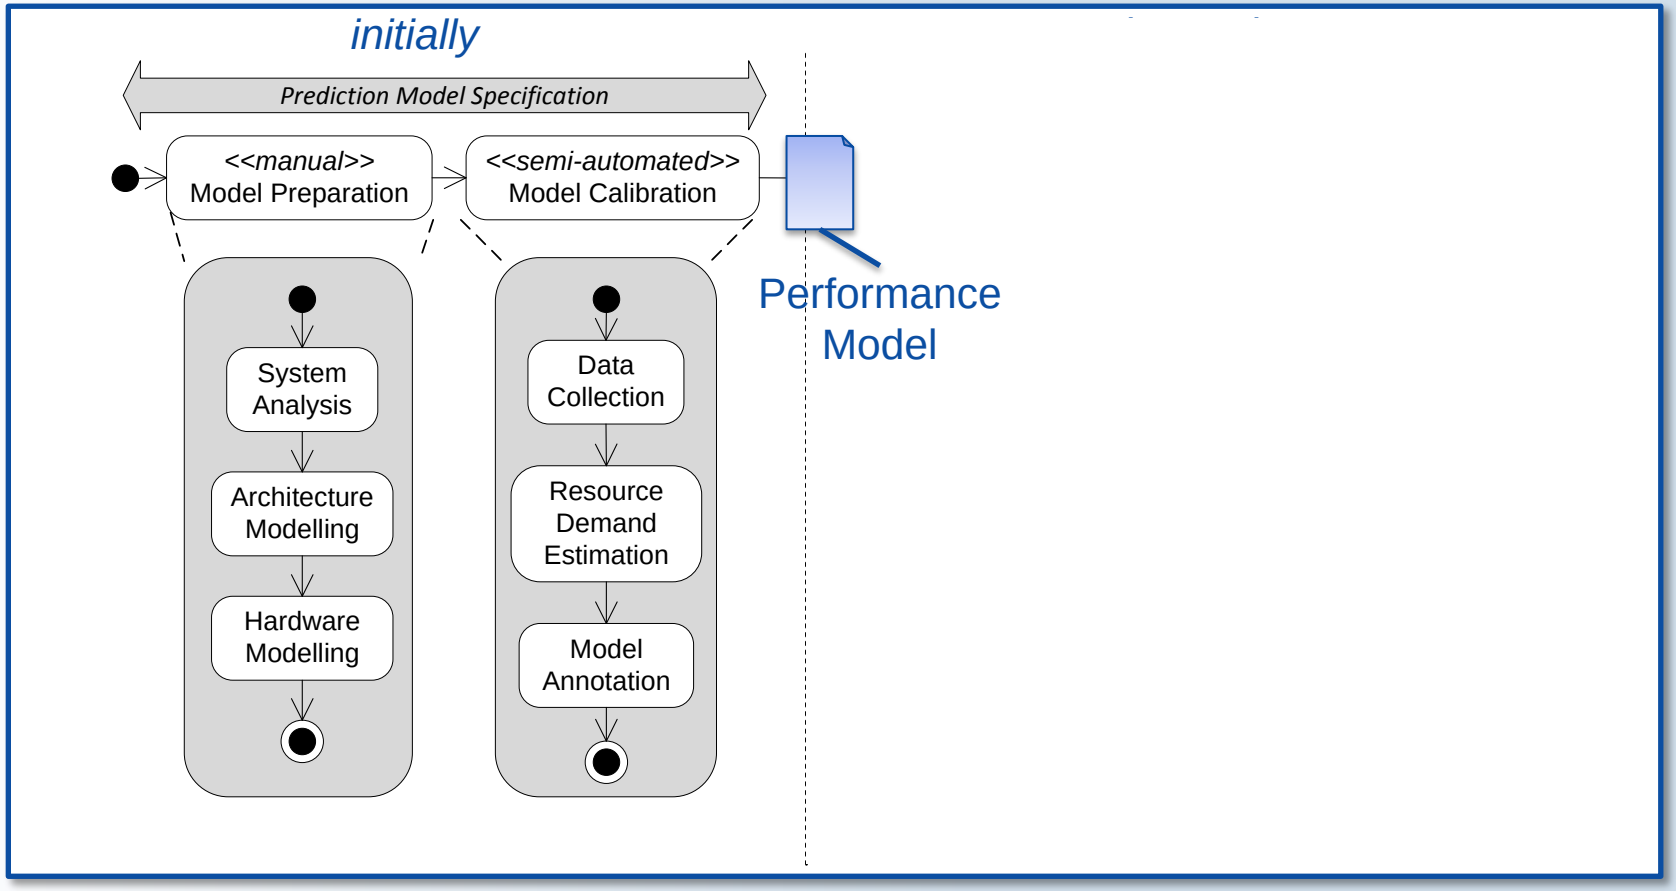

Approach

- Applying **performance prediction** to derive the **expected system behaviour** depending on the **monitored user behaviour**

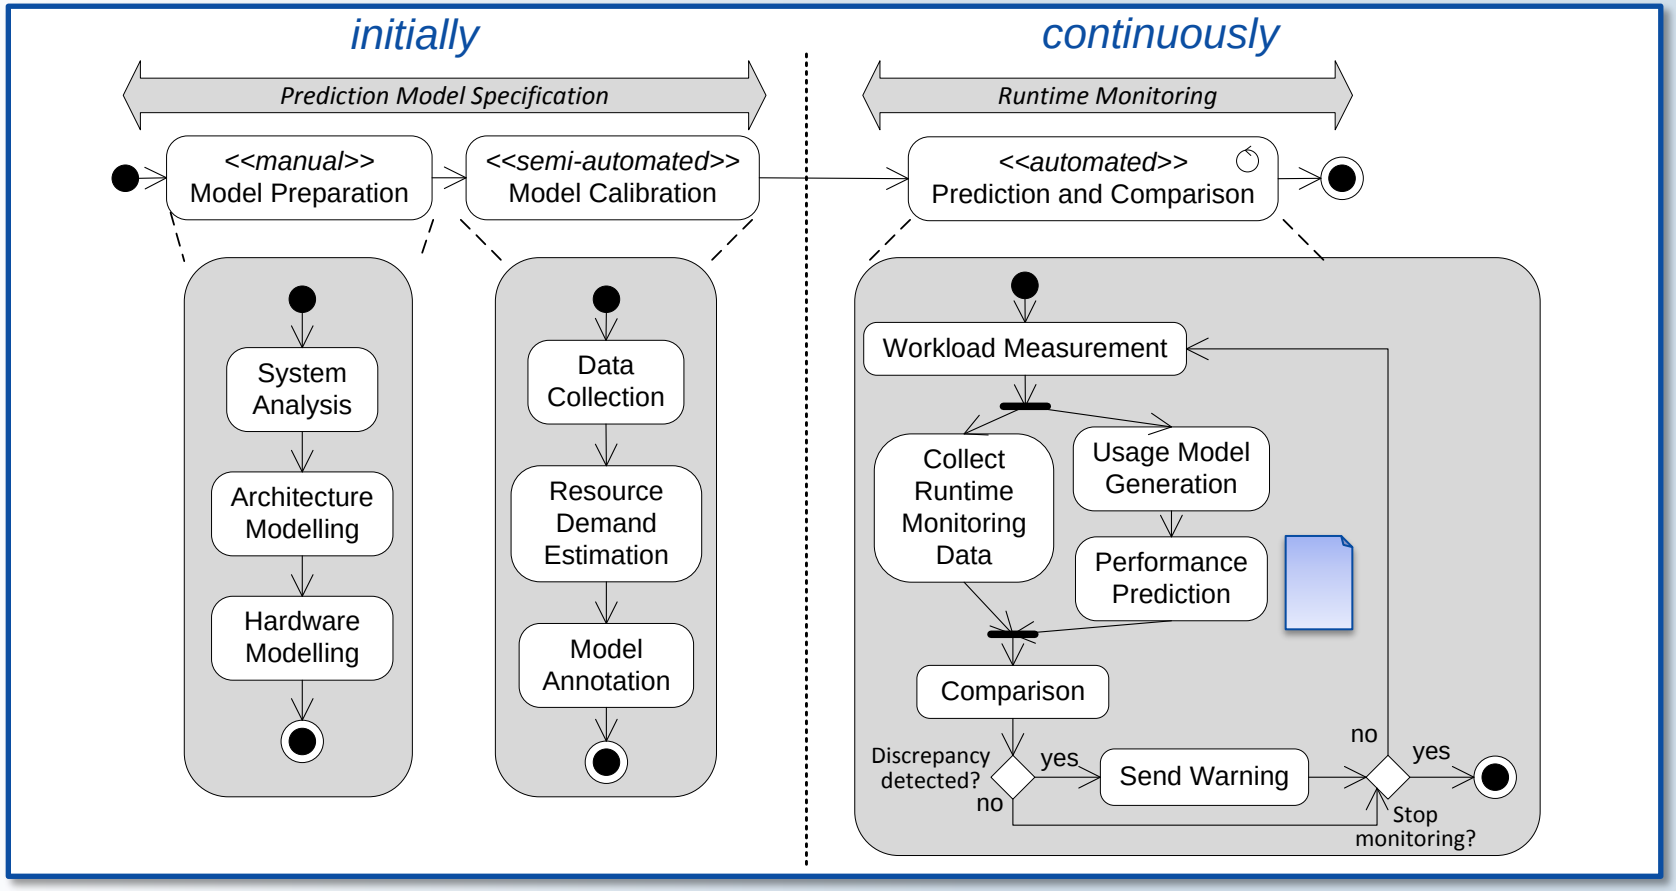

Foundations
Model-based Performance Predictions

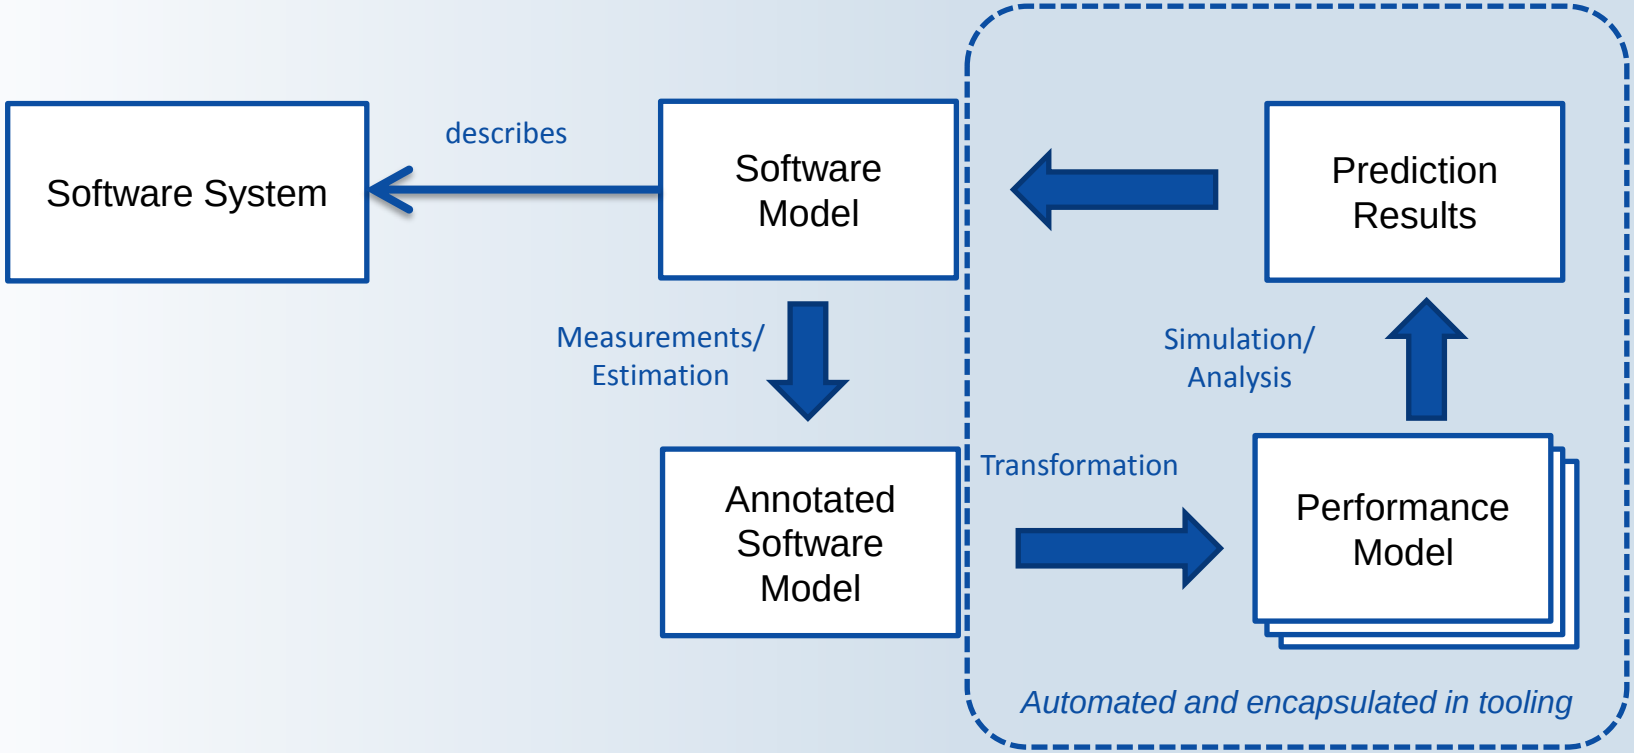

Foundations
Model-based Performance Prediction

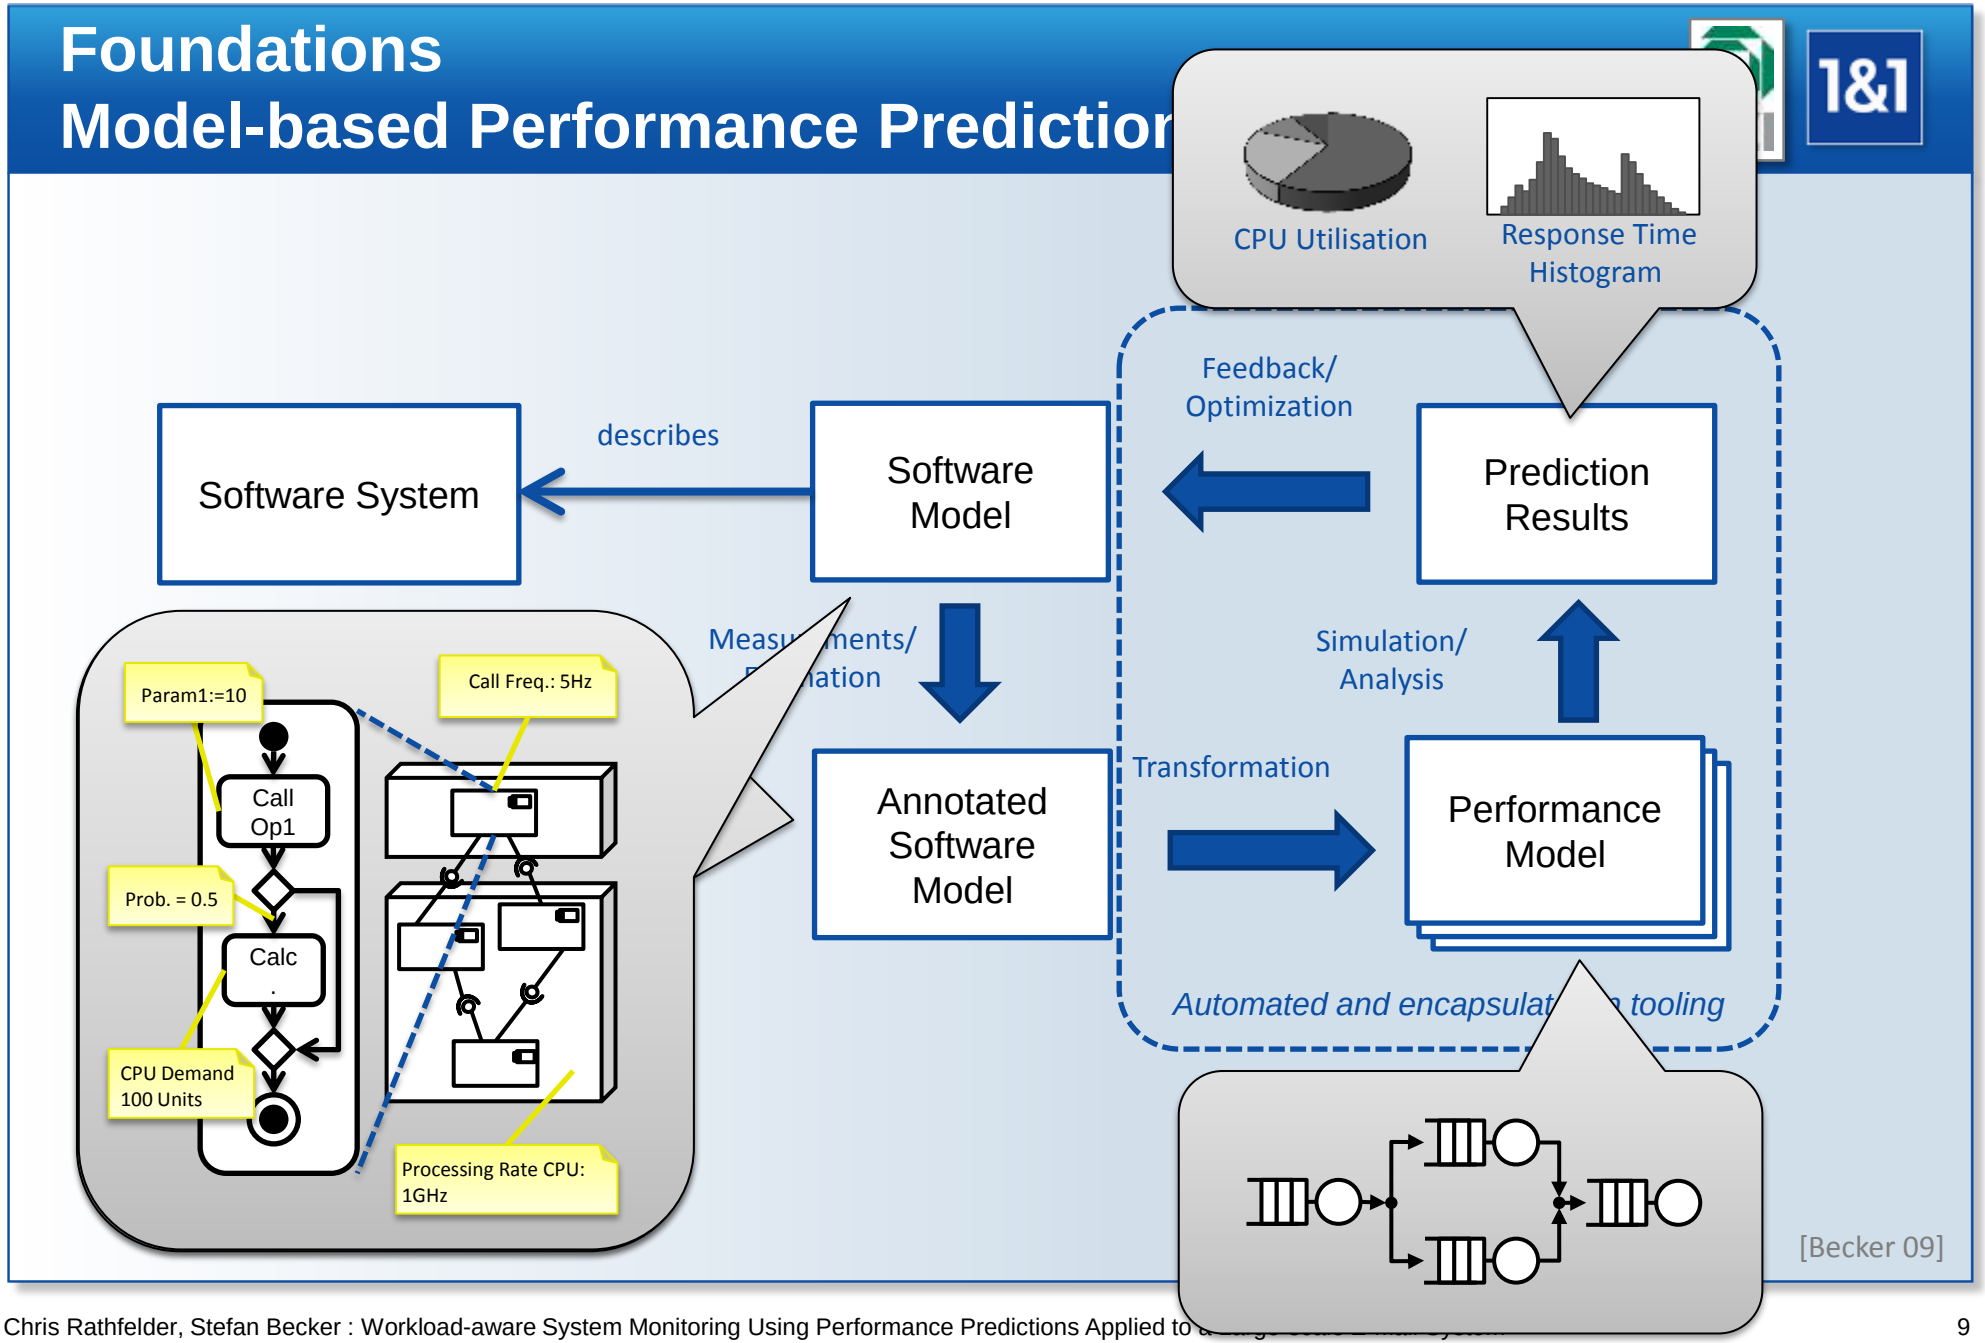

Foundations
Palladio - Tooling

- Palladio Component Model (PCM)
  - Domain-specific modelling language
  - Aligned with UML2 syntax

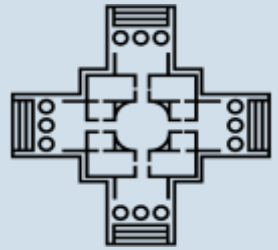

- Transformations into predictions models and simulation
  - Performance
  - Reliability

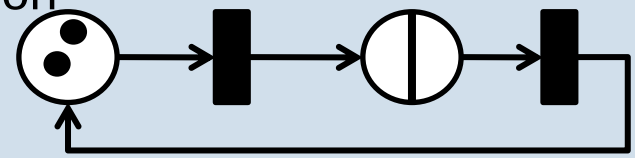

- Palladio Bench
  - Eclipse-based tool
  - Integrated modelling and prediction
  - Graphical editors

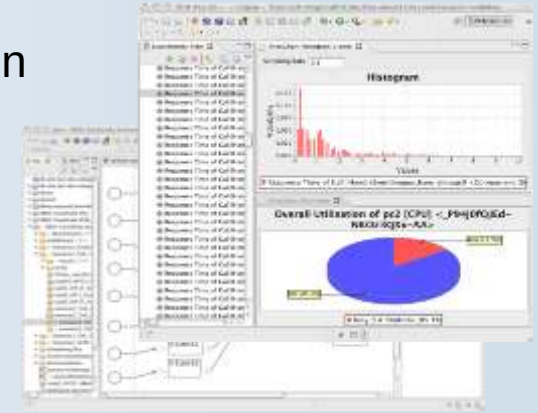

BUILT ON eclipse

- Tool maturity and availability
  - Open source
    - http://www.palladio-simulator.com
  - Development started in 2003
  - > 20 active contributor

Foundations
Palladio Approach

- Parameterized Models

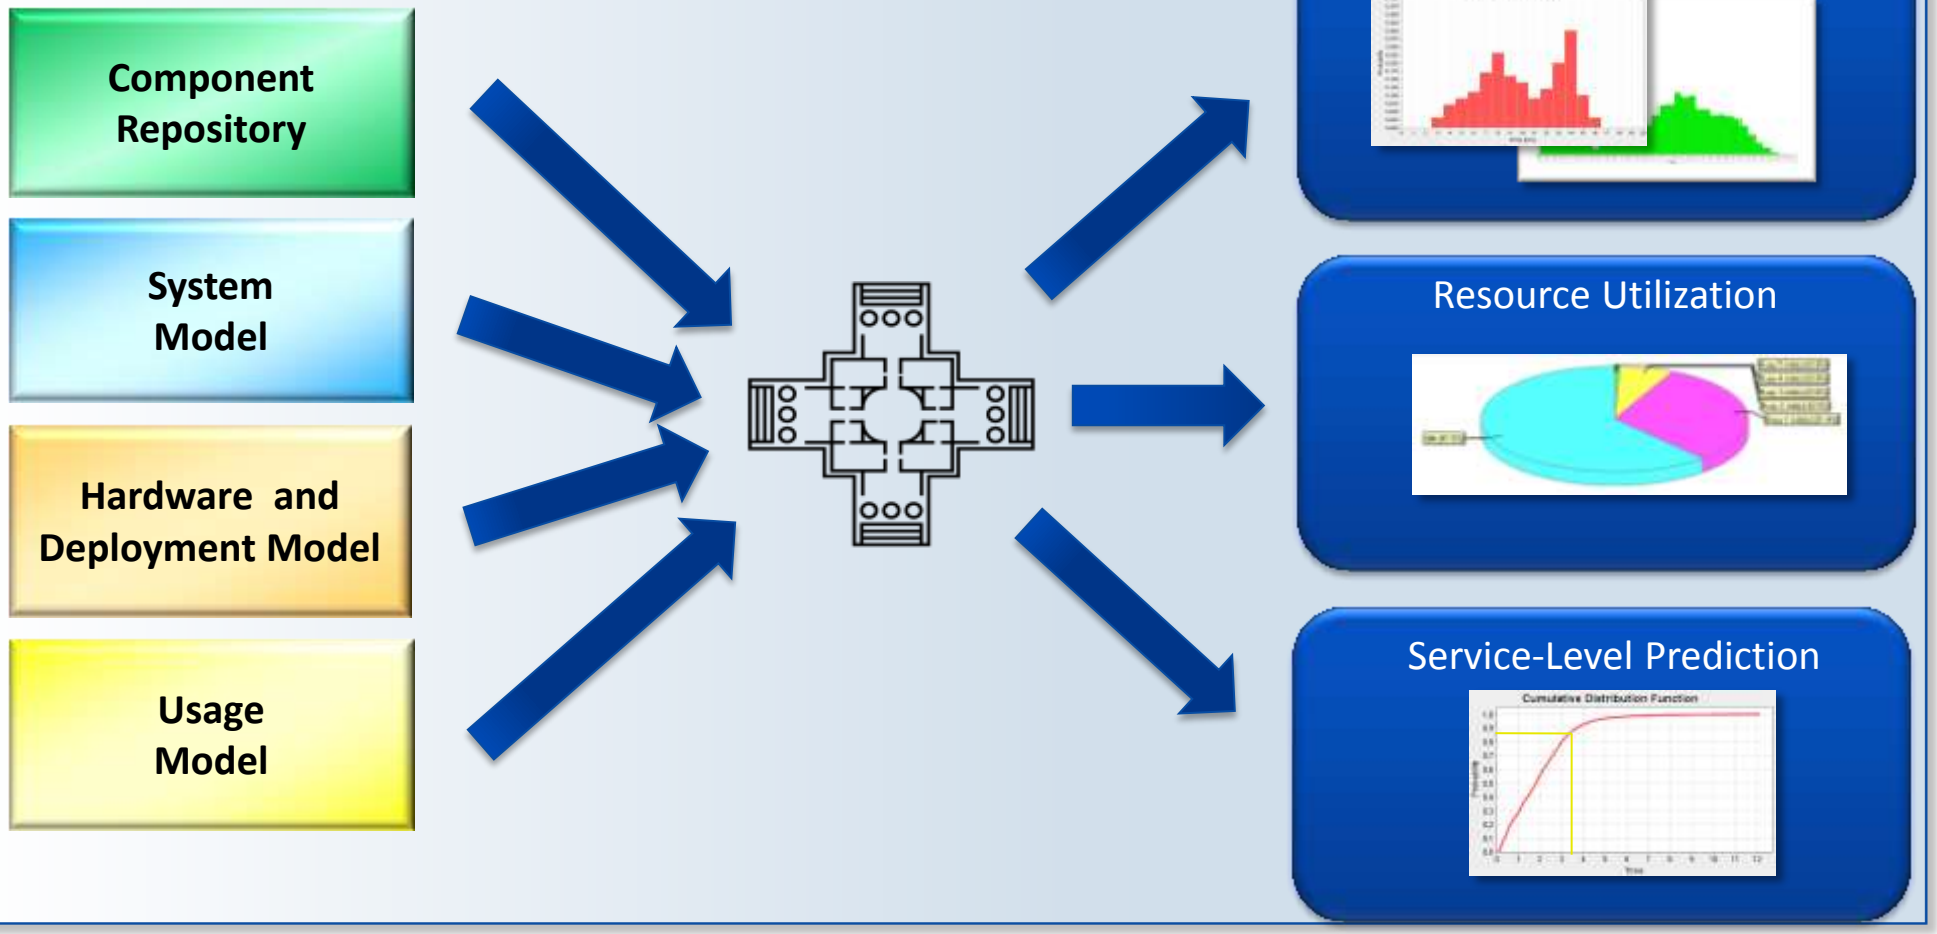

Agenda

1. Quick Facts about the 1&1 Mail System
2. Motivation and Approach
3. Foundation
    1. Model-based Performance Predictions
    2. Palladio Approach
4. Case Study
5. Validation and Applicability of our Approach
6. Lessons Learned
7. Conclusion and Future Work

The 1&1 Mail System – Top Level View

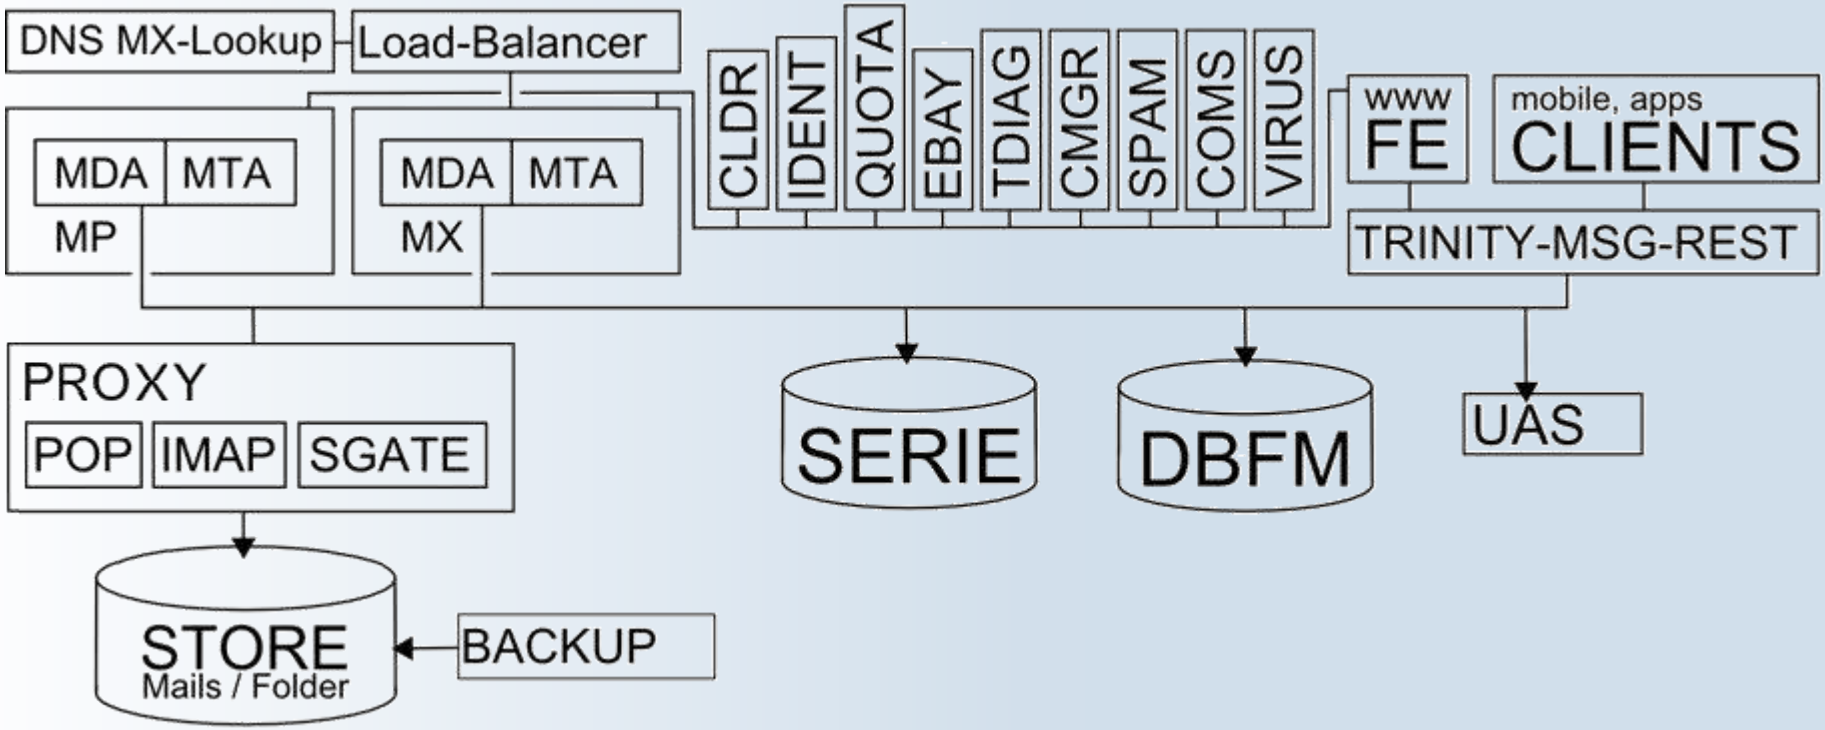

The 1&1 Mail System – Top Level View

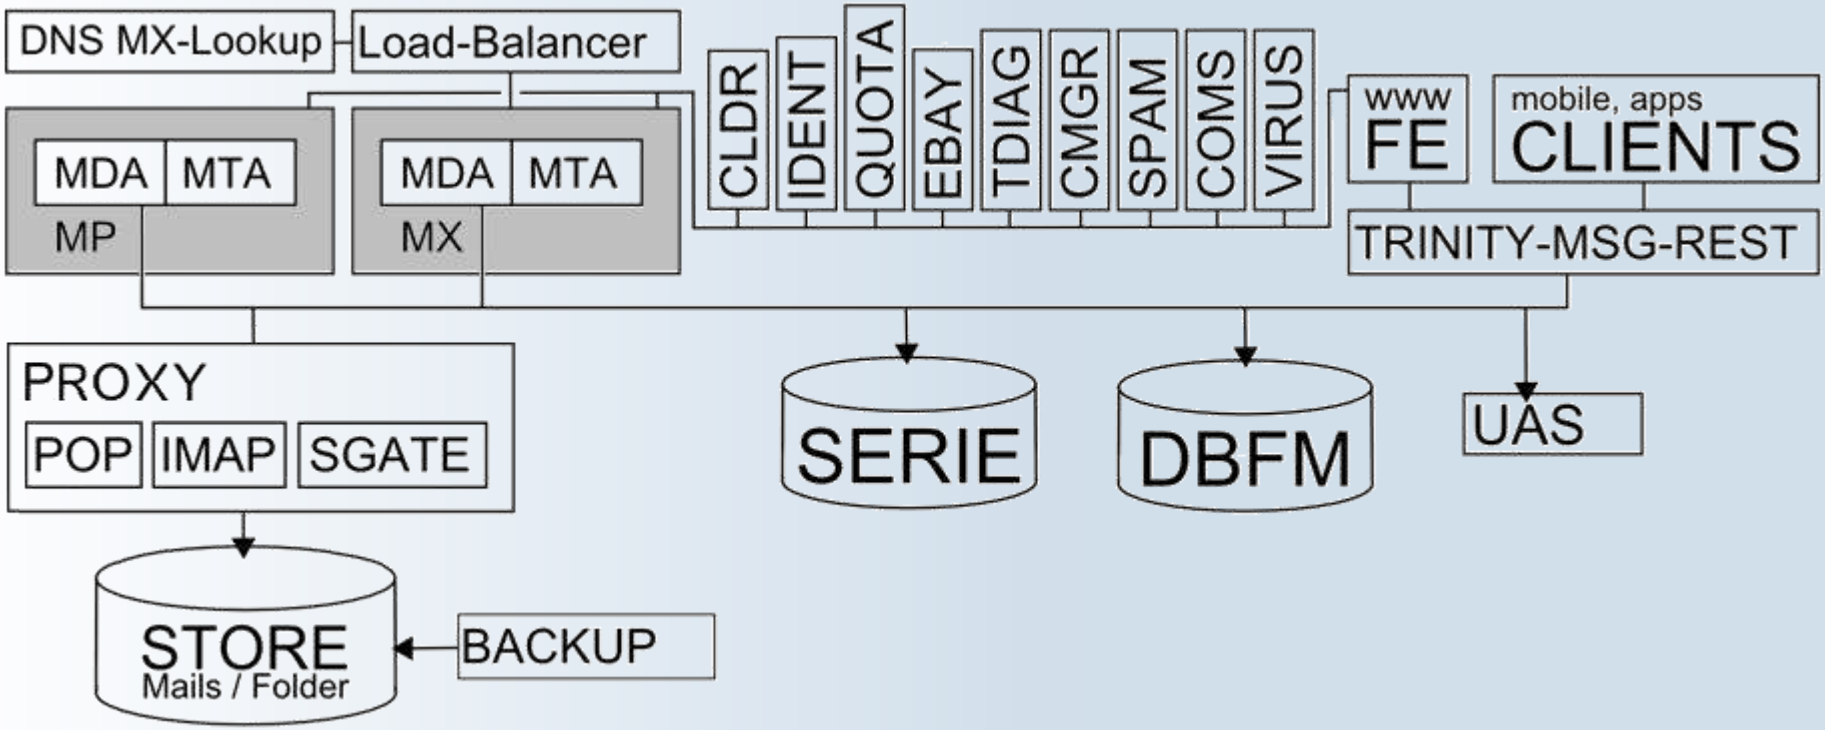

The 1&1 Mail System – Top Level View

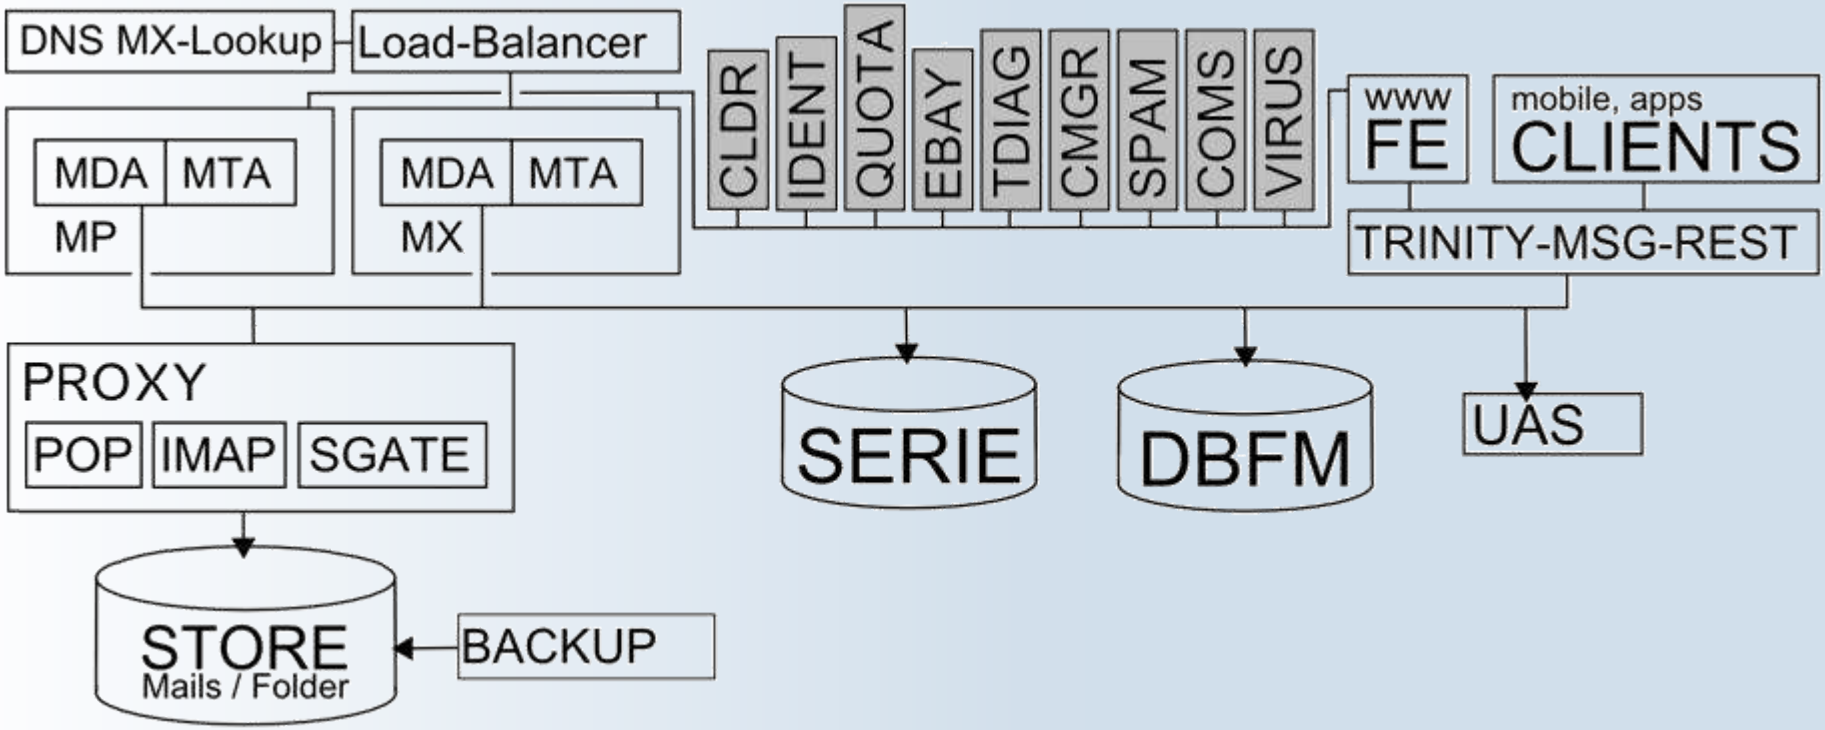

The 1&1 Mail System – Top Level View

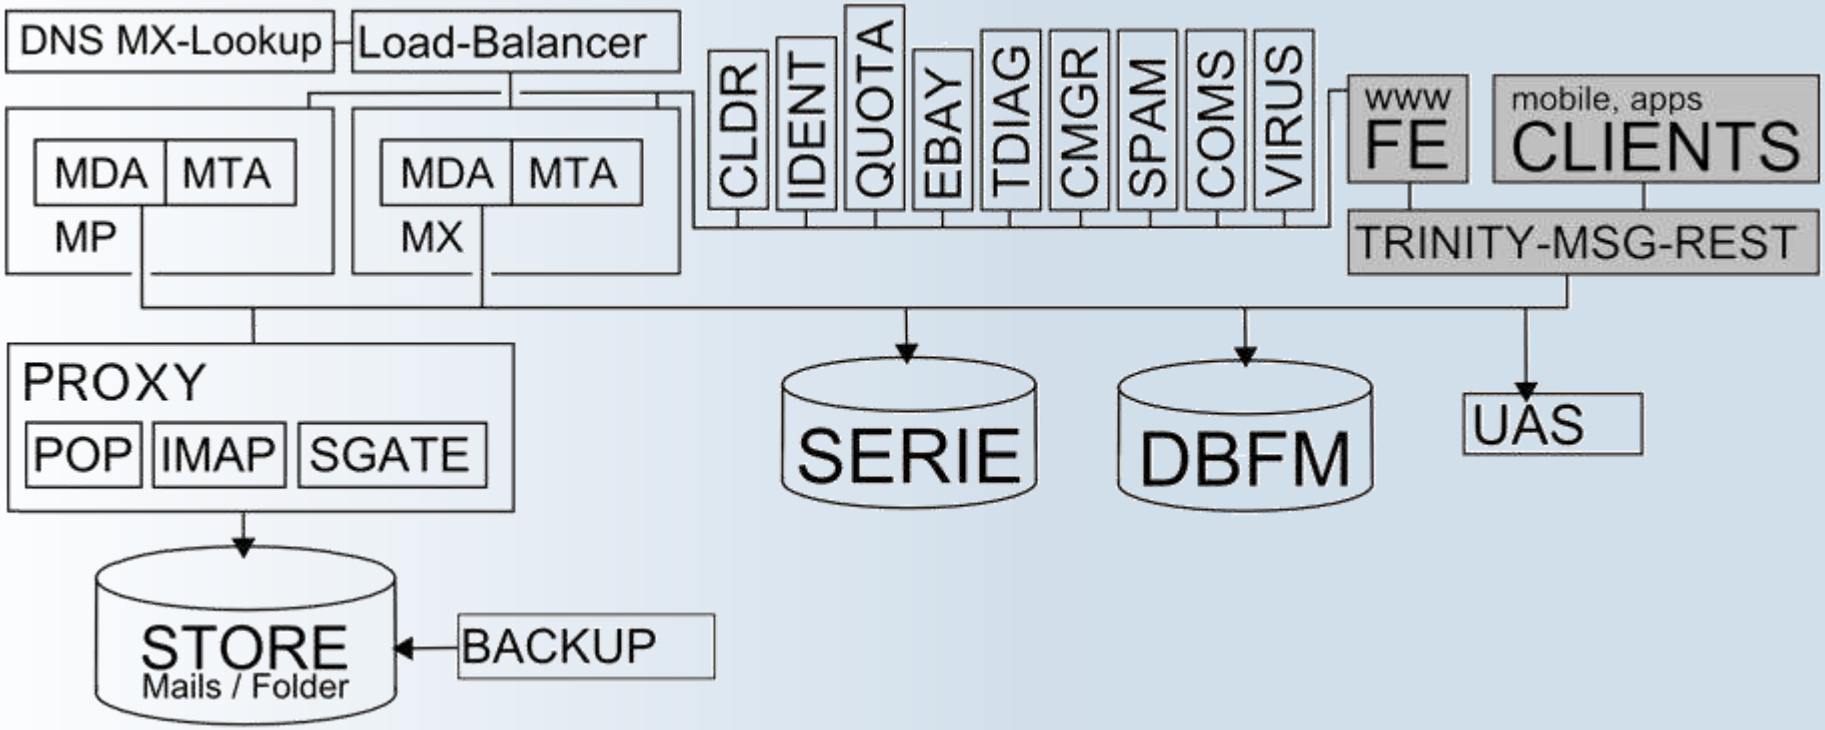

The 1&1 Mail System – Top Level View

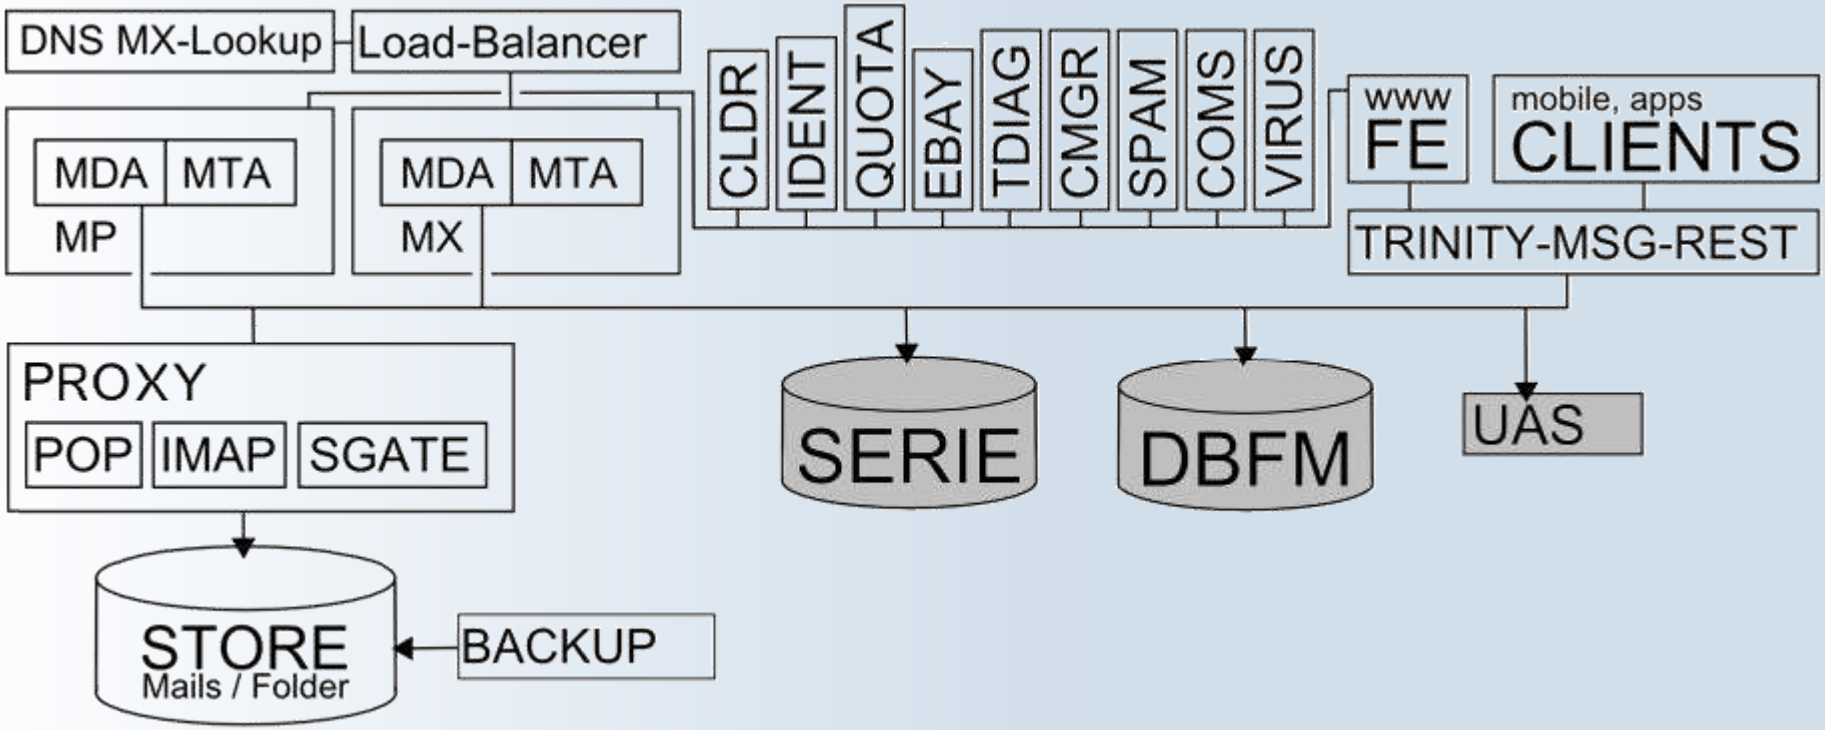

The 1&1 Mail System – Top Level View

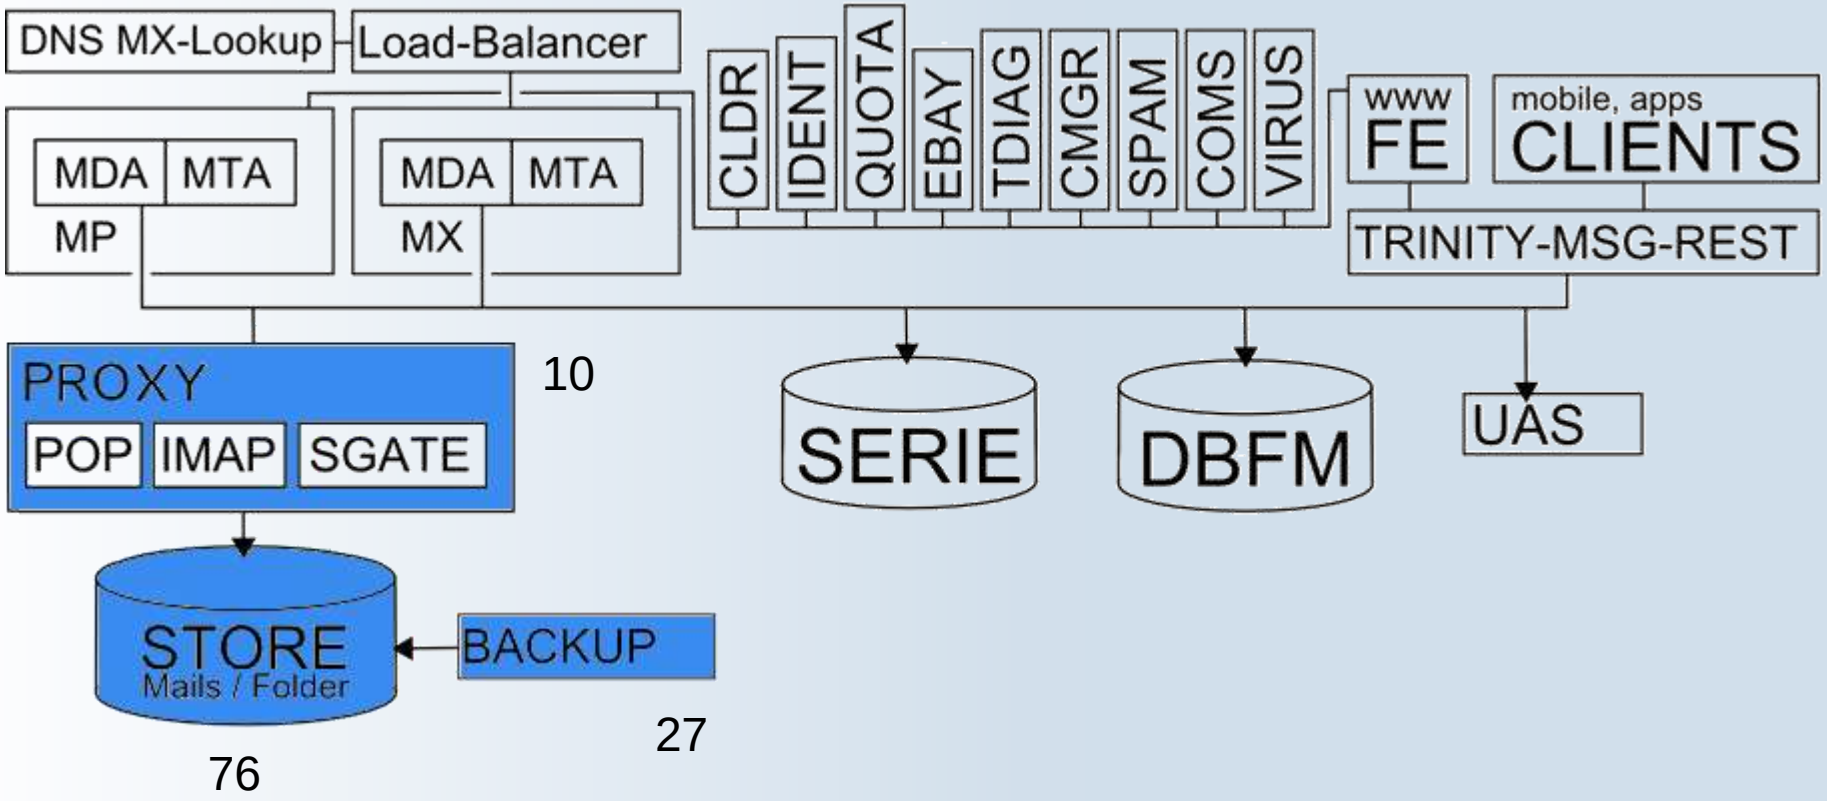

Case Study

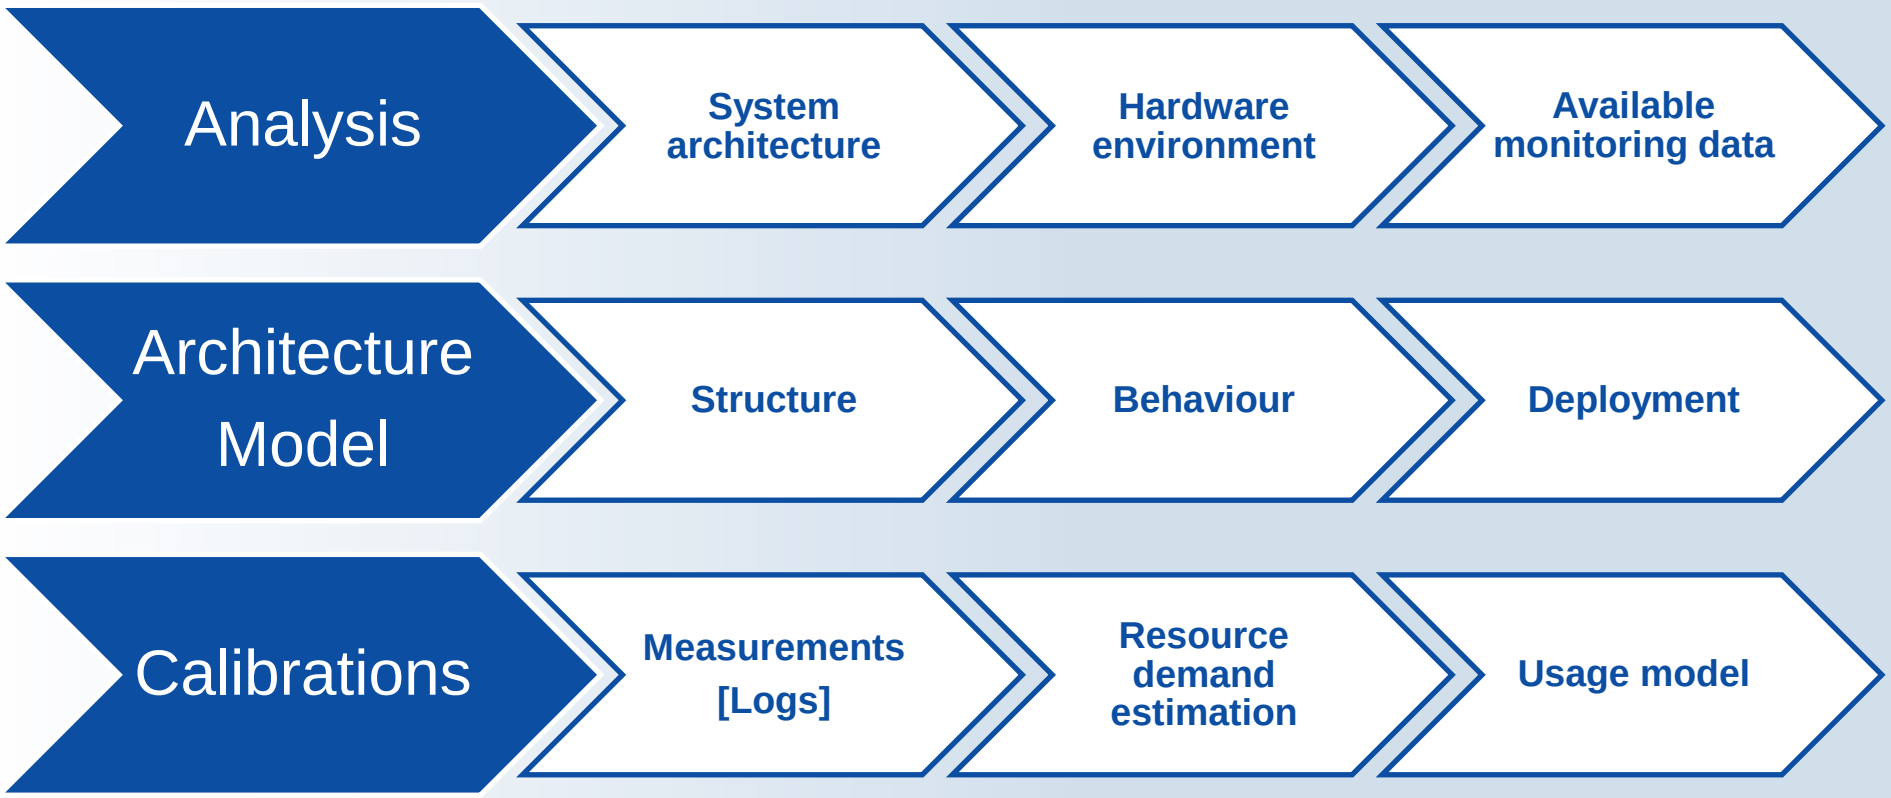

Case Study

Challenges

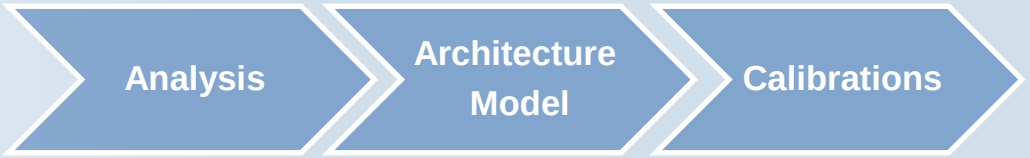

- Many servers
  - 113 even in the selected subsystem
- Live system
  - High availability must be guaranteed
  - No experiments
  - Expensive generation of load
  - QA System can not be used because of virtual hosts, mocks etc.
- Distributed knowledge
- Outdated documentation

Case Study

Store System

Analysis → Architecture Model → Calibrations

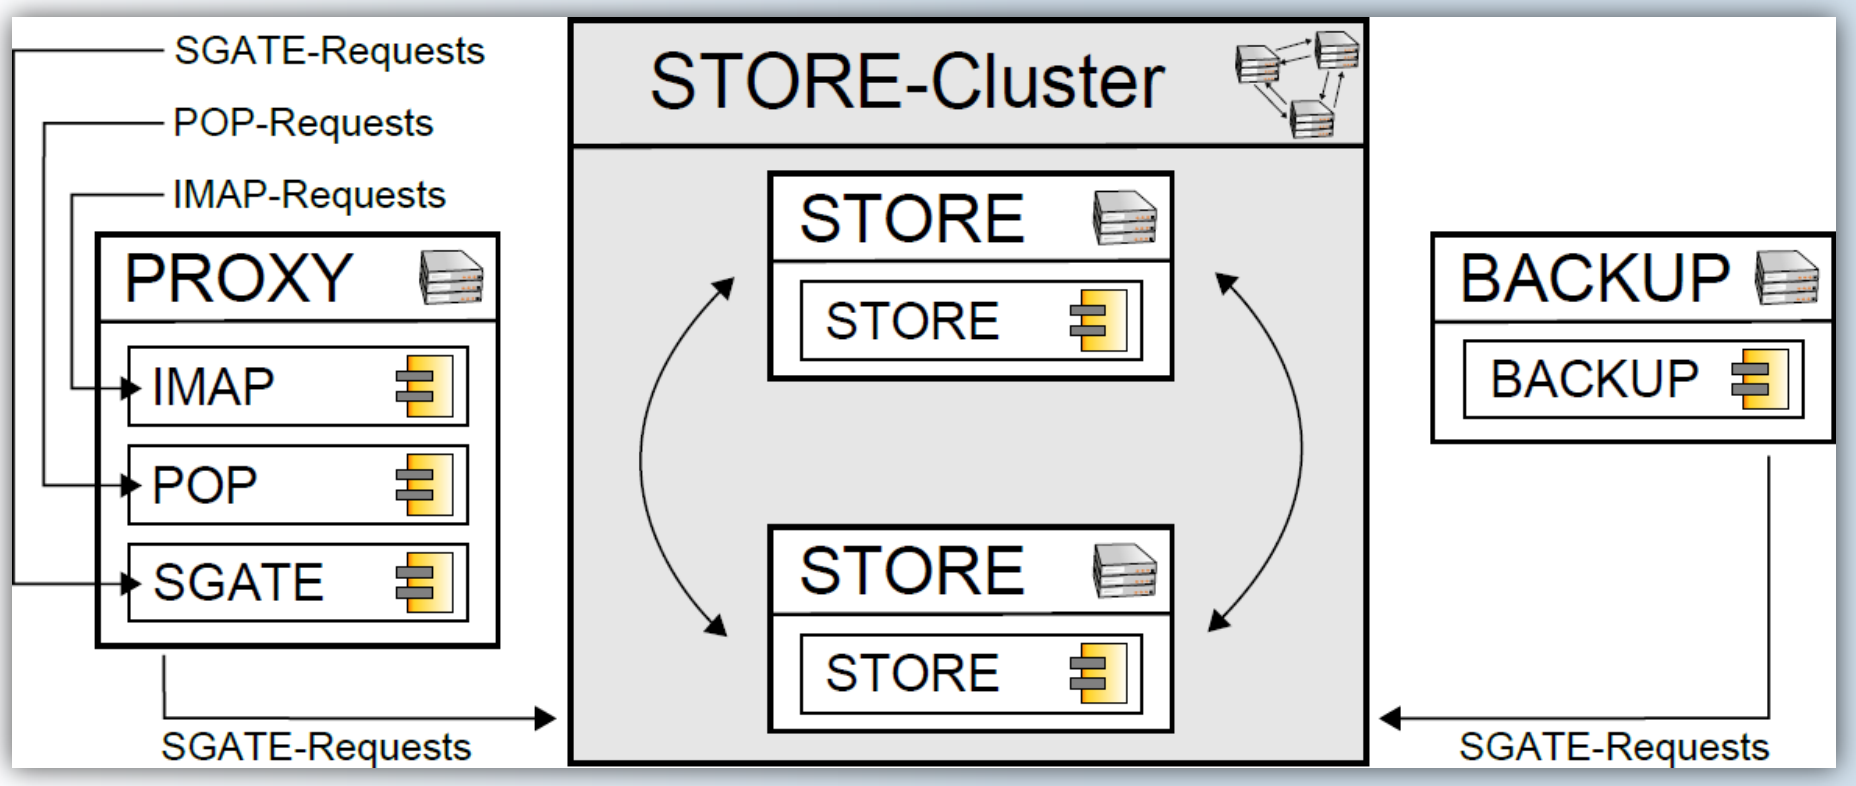

Case Study

Repository

Analysis → Architecture Model → Calibrations

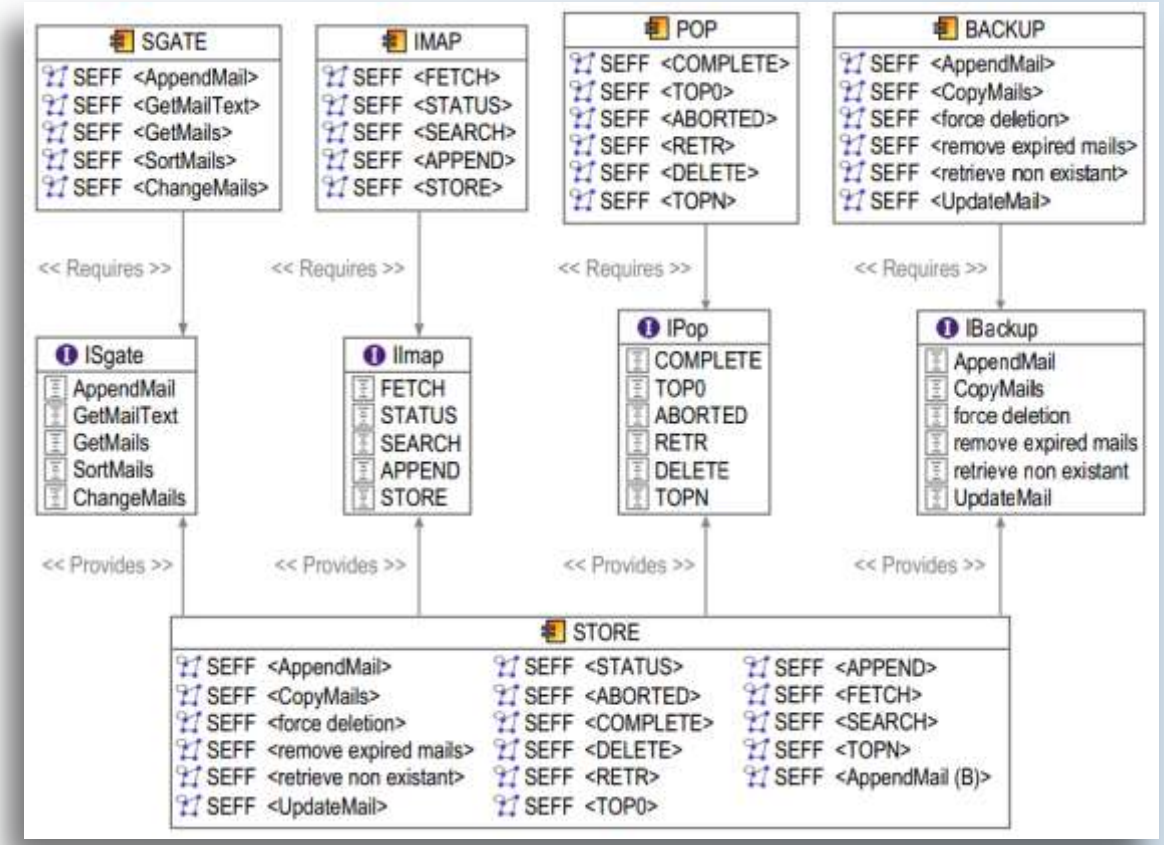

Case Study

System Model

Analysis | Architecture Model | Calibrations

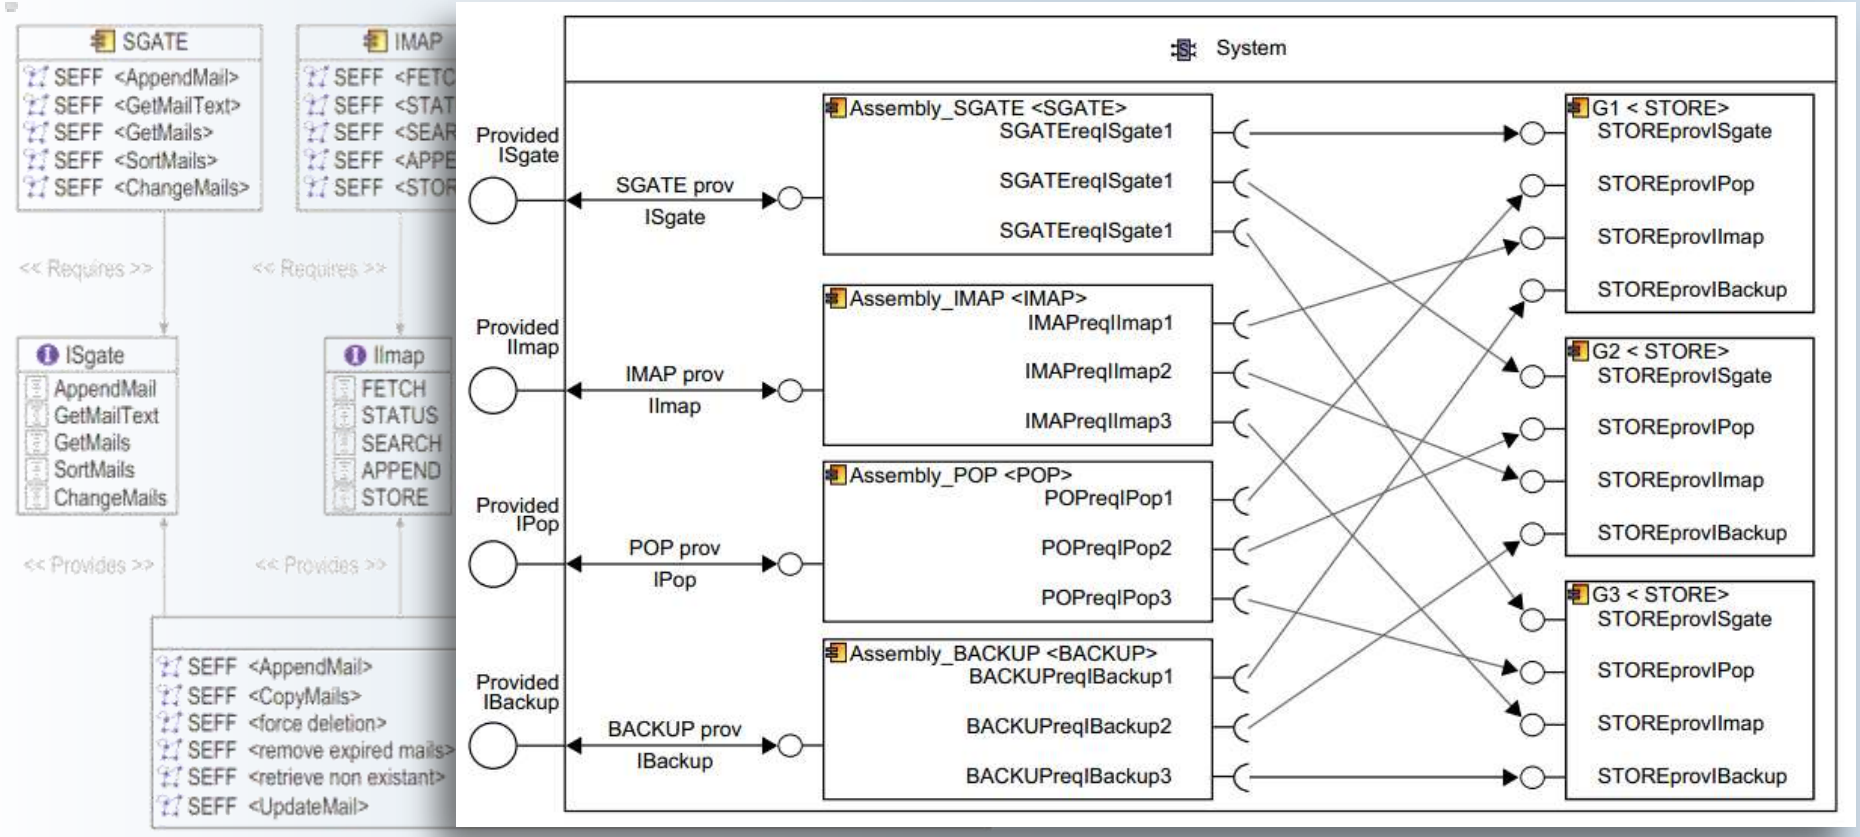

Case Study

Resource Demand Estimation

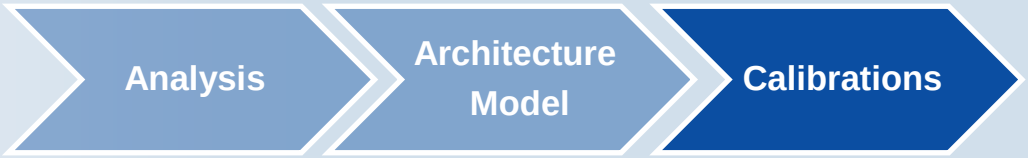

Model function for incoming traffic

Bytes received	Function	Relation
1	AppendMail	Byte
1	GetMailText	Byte
1	RETR	Byte
530966	APPEND	Request
143349	GetMails (+)	Request
83067	TOPN	Request
62479	TOP0	Request
23911	SortMails (+)	Request
3089	FETCH	Request
1298	STORE	Request

Case Study

Service Effect Specification

Analysis → Architecture Model → Calibrations

APPEND MAIL request on a PROYS server

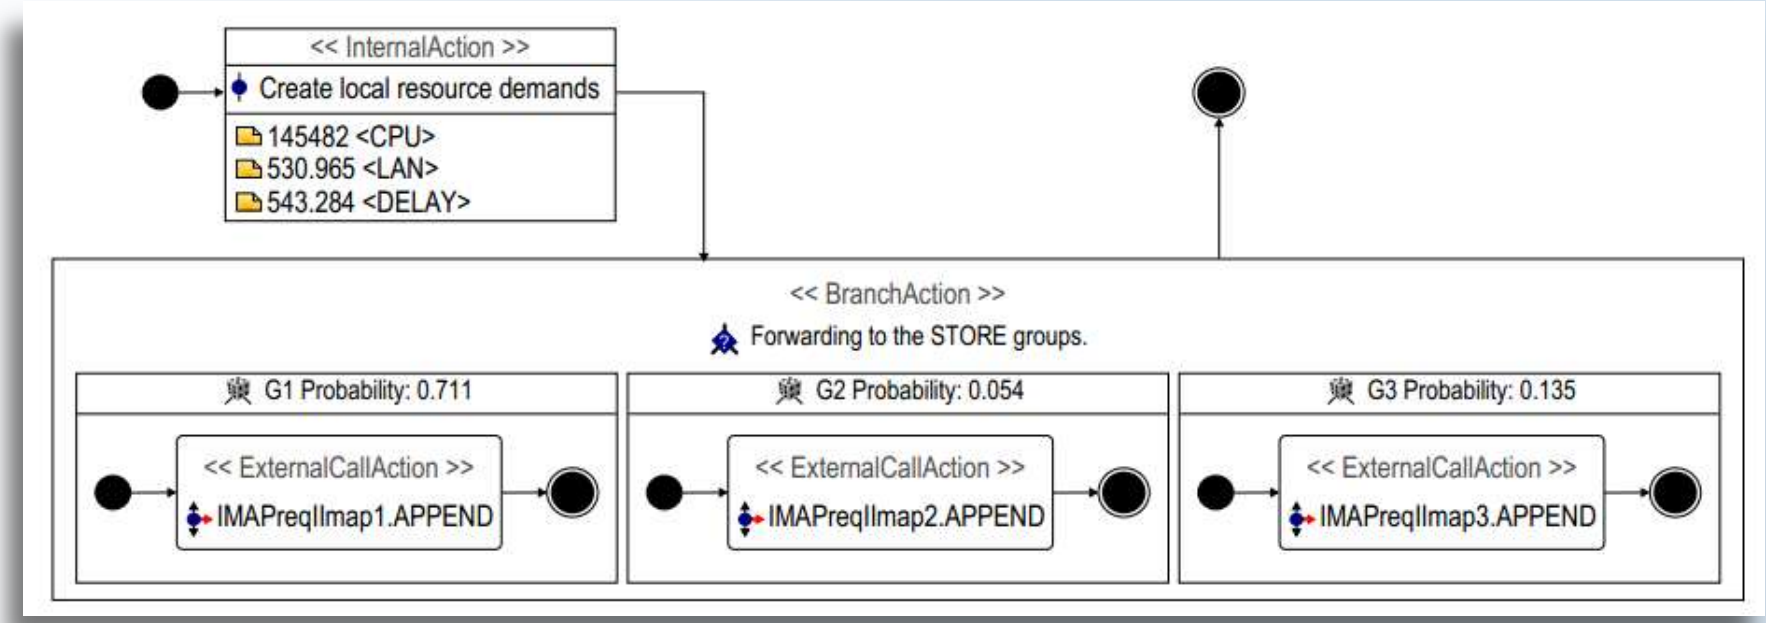

Validation

Predicition Accuracy

- Comparison between real and predicted utilisations per resource
- 1 reference day
  - 48 Measurements
  - 30 minutes each
  - Over all mean error
    - 8,6 %

	CPU	NET in	NET out	DISK
Whole Model				
Average	12,94	5,99	5,92	9,64
Decile	36,7	13,46	12,37	20,1
PROXY-Server				
Average	5,9	2,93	5,92	
Decile	10,1	6,4	11,09	
STORE-Server				
Average	15,29	7,01	6,01	9,64
Decile	42,1	16,19	13,31	20,1

Model errors in percent

Applicability of our Approach

Simulated System Failure

- Real and predicted resource utilisation over the day
- Planned update between 2:00 p.m. and 6:00 p.m.
- Situation for the model to raise an alarm

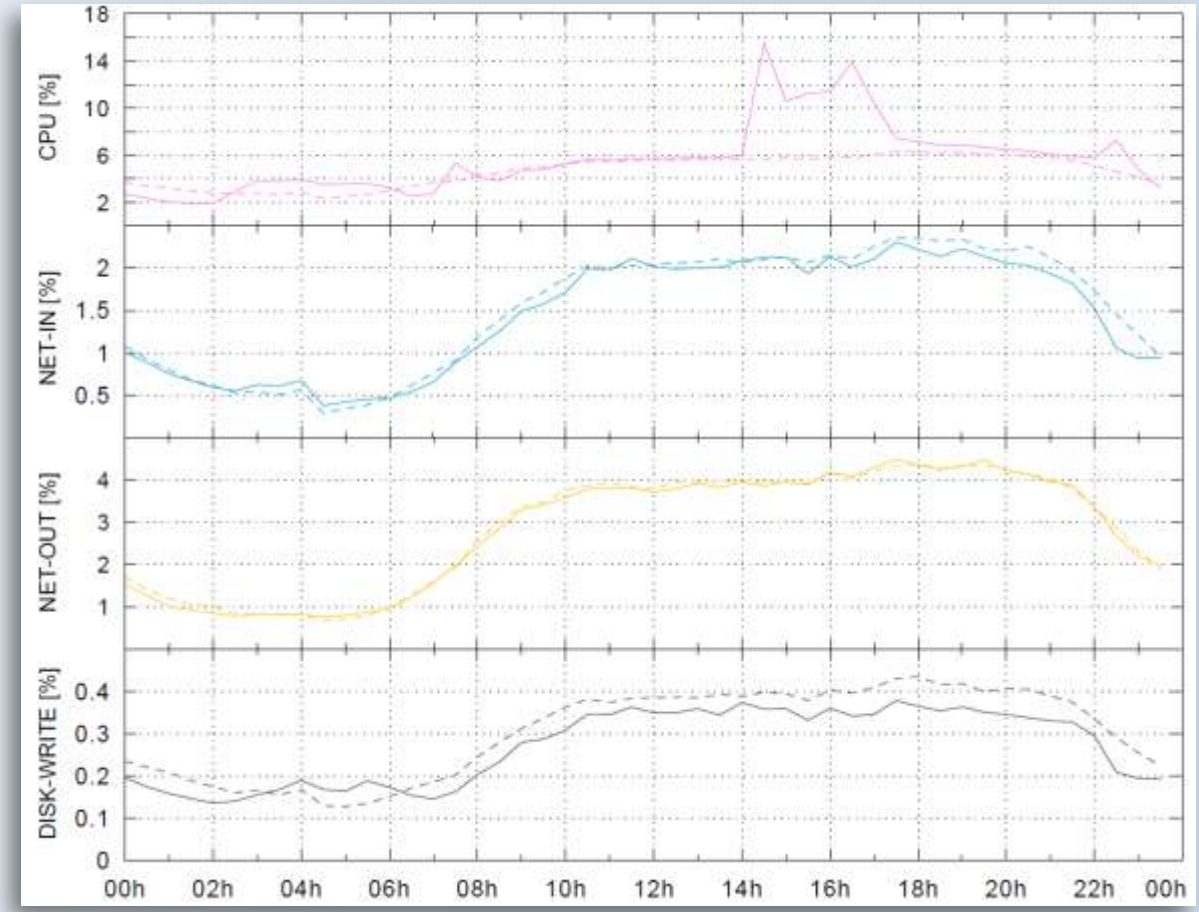

Lessons Learned

- Distributed knowledge and outdated documentation increase efforts
  - Performance models can serve as documentation
- Very large systems can not be meassured out by performing experiments
  - High availability => too dangerous
  - Creating the needed load needs many resources
  - Existing logfiles often contain useful data
- Redundant parts can be depicted together in a model
  - Combining servers with identical hardware software and load
- Detection of missconfigurations
  - Possible already in early steps of the modelling process
  - I.E. we found a server with disabled hyperthreading.

Conclusions and Outlook

Conclusions

- Enhancing a system monitoring process with performance prediction results
- Applicability of PCM in an industrial environment shown (mean model error of 8,6 %)

Future Work

- Improved support for replicated software components
- Automated model extraction and calibration
- Runtime management based on performance predictions

http://www.palladio-simulator.com