

# The Descartes Modeling Language and Tool Chain

**Samuel Kounev**

University of Würzburg

<http://se.informatik.uni-wuerzburg.de/>

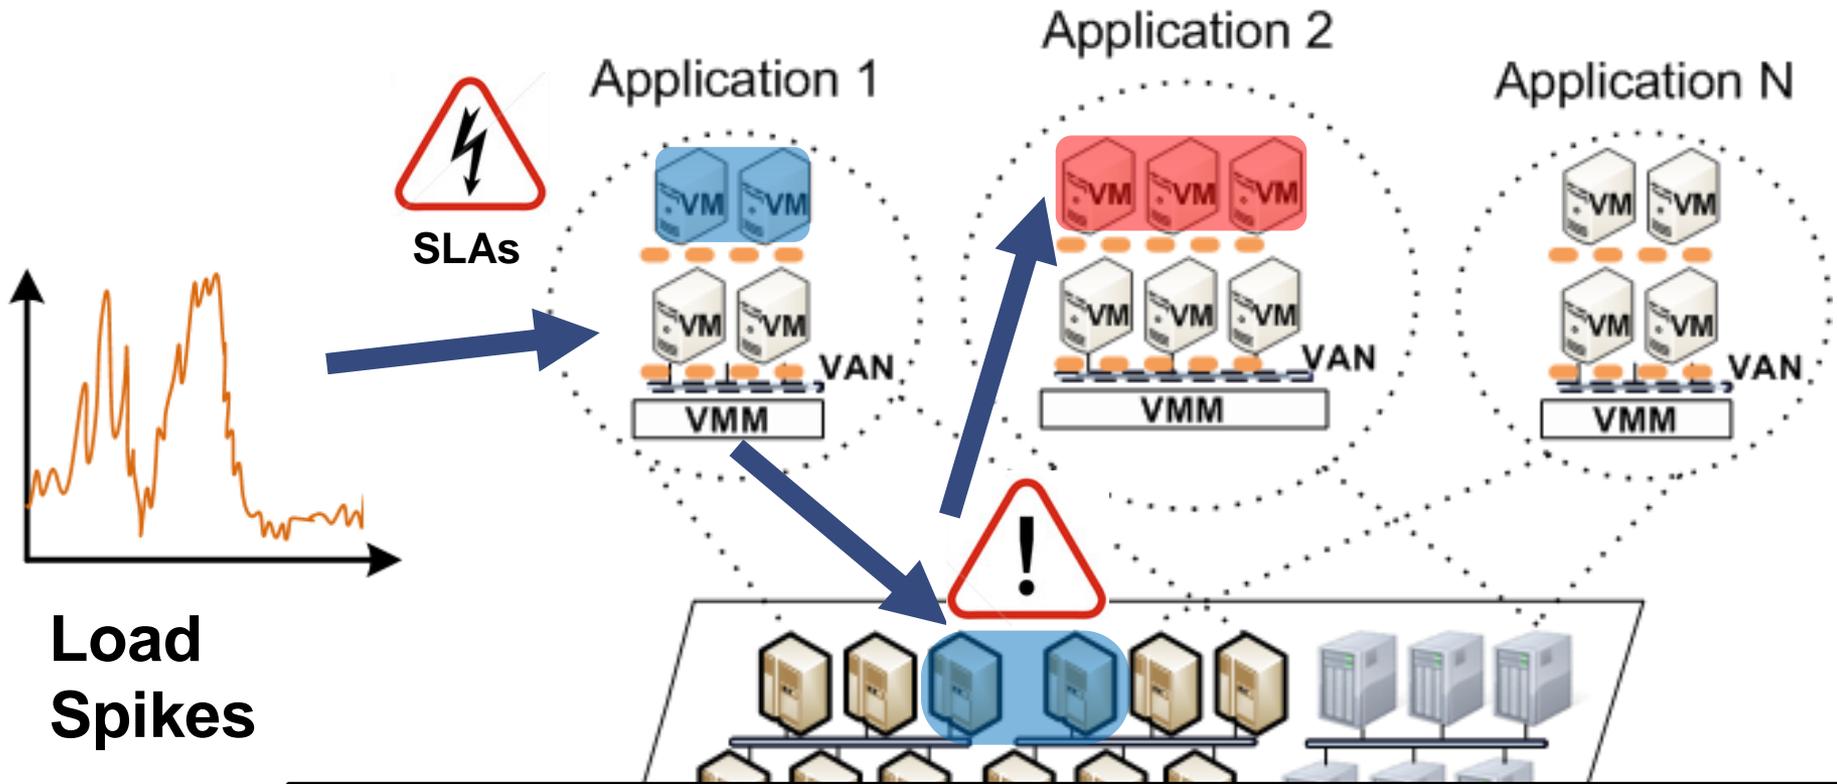
Jan 21, 2015

Dagstuhl Seminar 15041:

“Model-driven Algorithms and Architectures for Self-Aware Computing Systems”

- S. Kounev, F. Brosig, N. Huber, and X. Zhu. **Model-Based Approach to Designing Self-Aware IT Systems and Infrastructures**. Under review. IEEE Computer Special Issue on Self-Aware and Self-Expressive Computing Systems, 2015. *Available on request*.
- S. Kounev, F. Brosig, and N. Huber. **The Descartes Modeling Language**. Technical report, Department of Computer Science, University of Wuerzburg, October 2014. [ [http](#) | [http](#) | [.pdf](#) ]
- F. Brosig, N. Huber, and S. Kounev. **Architecture-Level Software Performance Abstractions for Online Performance Prediction**. *Elsevier Science of Computer Programming Journal (SciCo)*, Vol. 90, Part B:71-92, 2014, Elsevier. [ [DOI](#) | [http](#) | [.pdf](#) ]
- N. Huber, A. van Hoorn, A. Koziolok, F. Brosig, and S. Kounev. **Modeling Run-Time Adaptation at the System Architecture Level in Dynamic Service-Oriented Environments**. *Service Oriented Computing and Applications Journal (SOCA)*, 8(1):73-89, 2014, Springer-Verlag. [ [DOI](#) | [.pdf](#) ]
- F. Brosig, P. Meier, S. Becker, A. Koziolok, H. Koziolok, and S. Kounev. **Quantitative Evaluation of Model-Driven Performance Analysis and Simulation of Component-based Architectures**. *IEEE Transactions on Software Engineering (TSE)*, 2014, IEEE, Preprint. [ [DOI](#) | [.pdf](#) ]
- F. Gorsler, F. Brosig, and S. Kounev. **Performance Queries for Architecture-Level Performance Models**. In *5th ACM/SPEC International Conference on Performance Engineering (ICPE 2014)*, Dublin, Ireland, 2014. ACM, New York, NY, USA. 2014. [ [DOI](#) | [.pdf](#) ]
- N. Herbst, N. Huber, S. Kounev and E. Amrehn. **Self-Adaptive Workload Classification and Forecasting for Proactive Resource Provisioning**. *Concurrency and Computation - Practice and Experience, John Wiley and Sons, Ltd.*, 26(12):2053-2078, 2014. [ [DOI](#) | [http](#) | [.pdf](#) ]
- S. Spinner, G. Casale, F. Brosig, S. Kounev. **Evaluating Approaches to Resource Demand Estimation**. Elsevier Performance Evaluation Journal, 2015. Under publication. Available on request.
- N. Herbst, S. Kounev and R. Reussner. **Elasticity: What it is, and What it is Not**. In *10th Intl. Conference on Autonomic Computing (ICAC 2013)*, San Jose, CA, June 24-28, 2013. [ [slides](#) | [http](#) | [.pdf](#) ]
- J. von Kistowski, N. Herbst, and S. Kounev. **Modeling Variations in Load Intensity over Time**. In *3rd Intl. Workshop on Large-Scale Testing (LT 2014)*, Dublin, Ireland, March 22, 2014, pages 1-4. ACM, New York, NY, USA. March 2014. [ [DOI](#) | [slides](#) | [http](#) | [.pdf](#) ]
- A. Milenkoski, B. Payne, N. Antunes, M. Vieira and S. Kounev. **HInjector: Injecting Hypercall Attacks for Evaluating VMI-based Intrusion Detection Systems** (Poster Paper). In *2013 Annual Computer Security Applications Conf. (ACSAC 2013)*, New Orleans, Louisiana, USA, 2013. [ [.pdf](#) ]

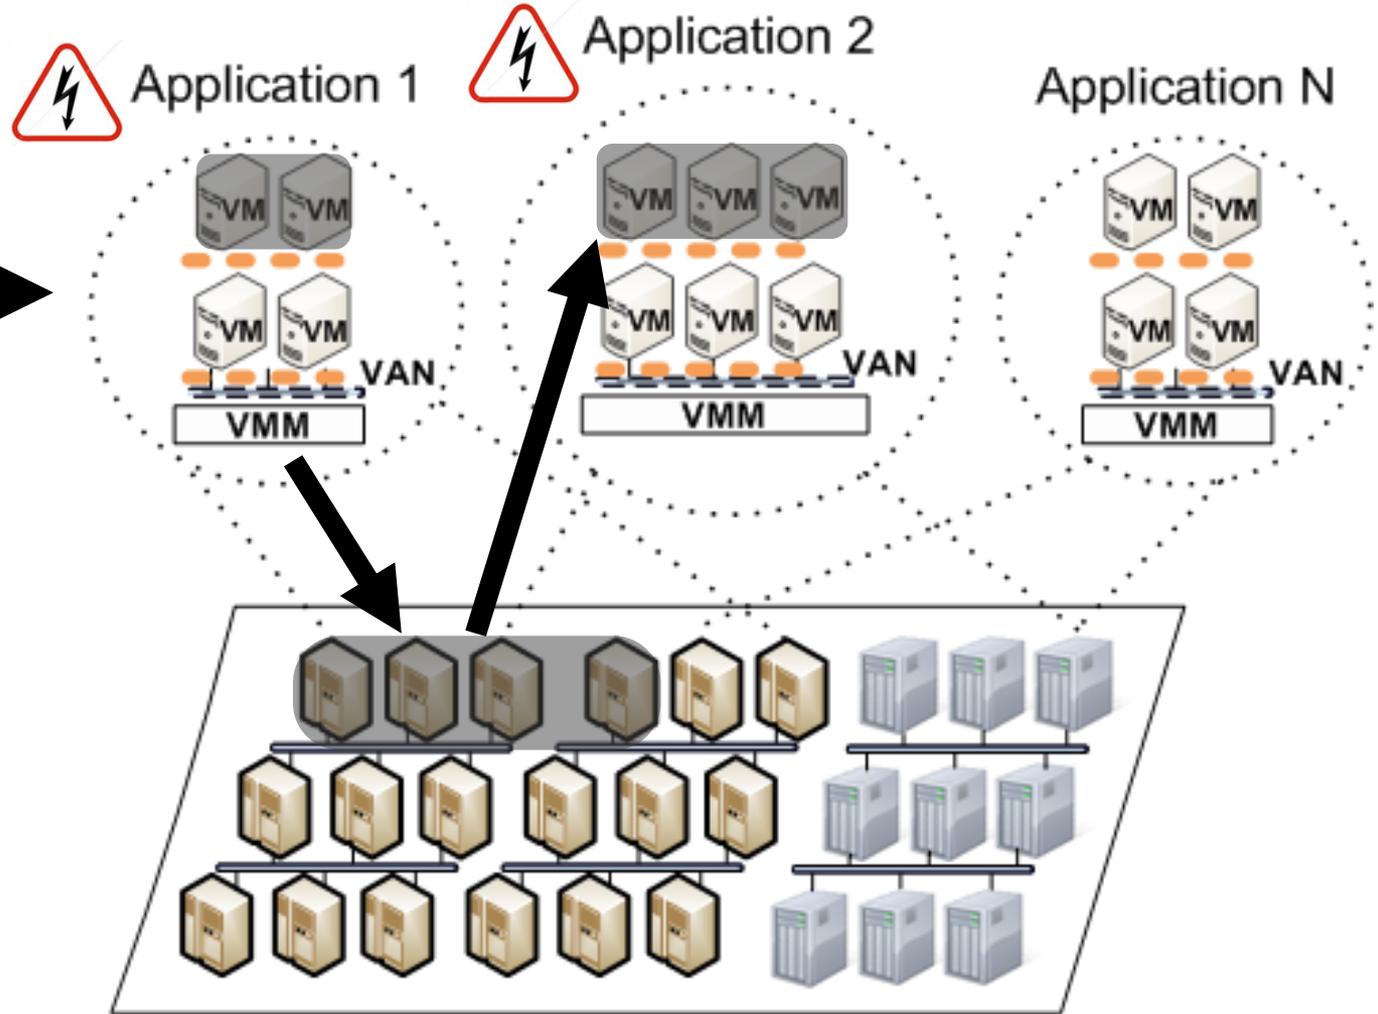




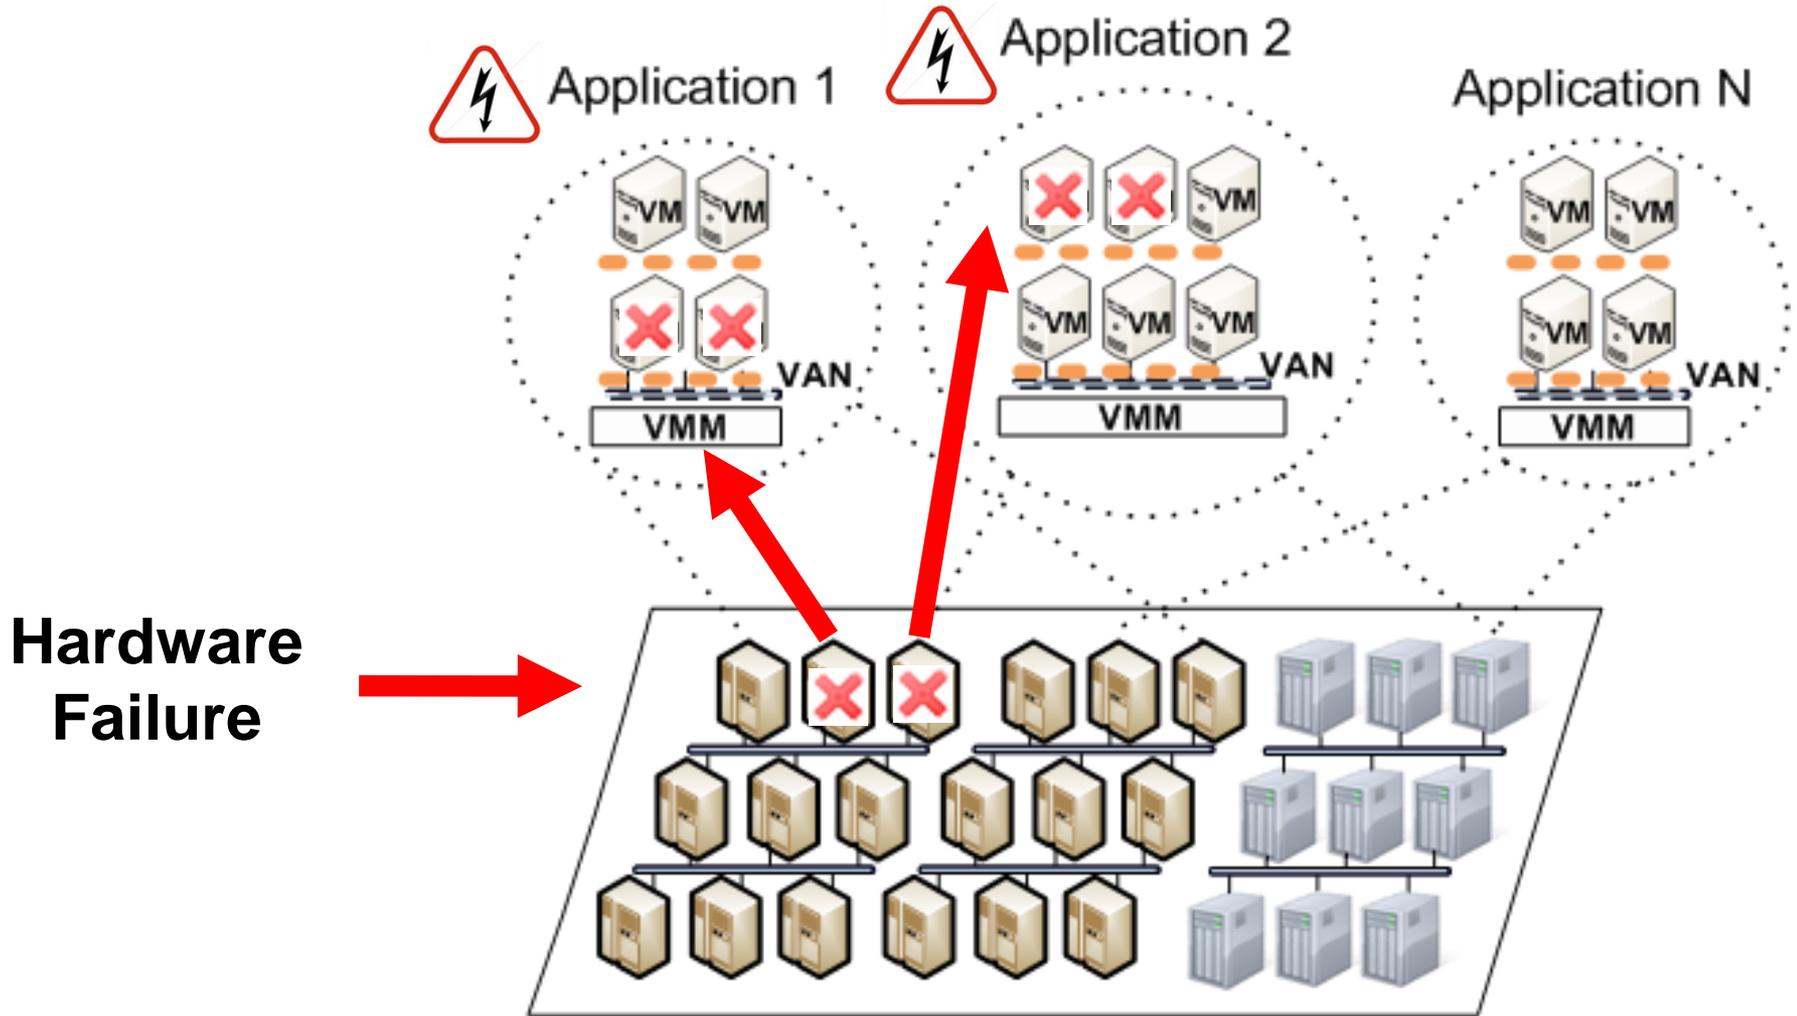
## Challenges

- When exactly should a reconfiguration be triggered?
- Which particular resources should be scaled?
- How quickly and at what granularity?

  
**Security  
Attack**



# Motivation



**Hardware Failure**

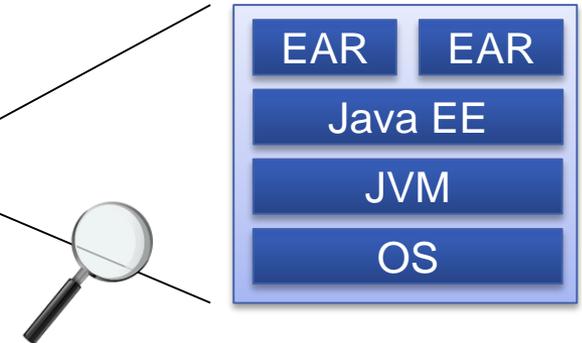
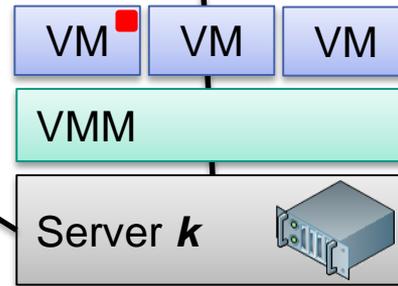
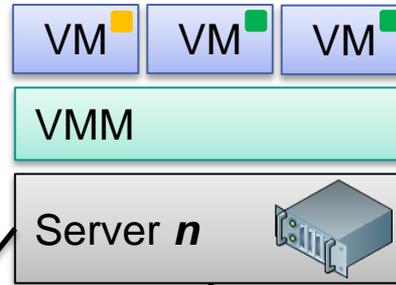
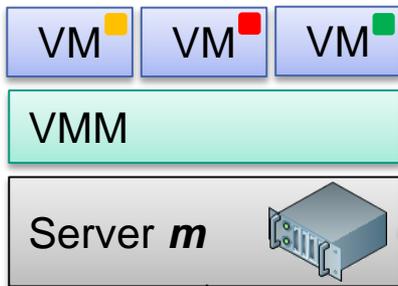


# Semantic Gap Problem

## Applications ■ ■ ■

- Multiple tiers
- Multiple resource types

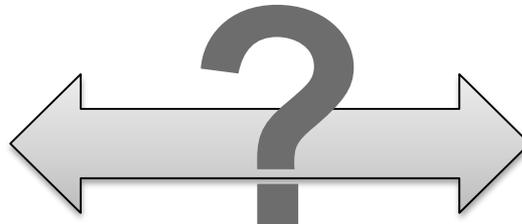
Resource Allocation

## Complex Software Stacks

- Multiple layers
- Heterogeneous

High-level Application Goals (e.g., SLOs)



Configuration of System Components, Layers & Tiers

# Semantic Gap Problem

## ■ Performance

- # requests that can be processed per sec > 1000
- Response time of service  $x < 20$  ms
- Server utilization > 60% on average
- ...

## ■ Availability / Reliability

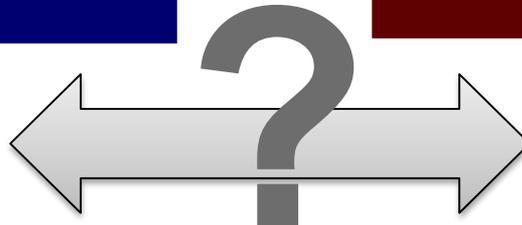
- Time to recover after a server failure < 1 min
- ...

## ■ Security

- Intrusion attempts are detected on time and prevented
- ...

- On which server to deploy software component  $y$ ?
- How many vCPUs to allocate to VM  $n$ ?
- How much memory to allocate to VM  $n$ ?
- When exactly should a reconfiguration be triggered?
- Which particular resources to scale / replicate / migrate?
- How quickly and at what granularity?

High-level Application Goals (e.g., SLOs)



Configuration of System Components, Layers & Tiers



**<http://descartes.tools>**

# Current Tools

- **DML** – Descartes Modeling Language ([homepage](#), [publications](#))
- **DML Bench** ([homepage](#), [publications](#))
- **DQL** – Declarative performance query language ([homepage](#), [publications](#))
- **LibReDE** - Library for resource demand estimation ([homepage](#), [publications](#))
- **LIMBO** – Load intensity modeling tool ([homepage](#), [publications](#))
- **WCF** – Workload classification & forecasting tool ([homepage](#), [publications](#))
- **BUNGEE** – Elasticity benchmarking framework ([homepage](#), [publications](#))
- **hInjector** – Security benchmarking tool ([homepage](#), [publications](#))
- **Further relevant research**
  - [http://descartes-research.net/research/research\\_areas/](http://descartes-research.net/research/research_areas/)
  - **Self Aware Computing** ([publications](#))



- « Fakultät für Mathematik und Informatik
- « Institut für Informatik
- « Lehrstuhl für Informatik II
- News
- People ▶
- Research ▶
- Publications ▶
- Projects ▶
- Tools ▼**
- DML Bench ▶
- DNI
- LIMBO ▶
- WCF
- LibReDE ▶
- SPA
- DQL
- BUNGEE
- hInjector

## Descartes Tools

Below you see a list of the tools we develop. Please click on the tool name to get more information:

### Descartes Modeling Language:

[DML Specification](#)

Implementation in EMF (Eclipse Modeling Framework)

[DML Bench](#)

[DNI - Descartes Network Infrastructures Modeling](#)

### Workload Characterization & Model Extraction:

[LIMBO Load Intensity Modeling Tool](#)

[WCF \(Workload Classification and Forecasting Tool\)](#)

[LibReDE \(Library for Resource Demand Estimation\)](#)

[SPA \(Storage Performance Analyzer\)](#)

### Declarative Performance Engineering:

[DQL \(Descartes Query Language\)](#)

### Benchmarking:

[BUNGEE Cloud Elasticity Benchmark](#)

[hInjector Hypercall Attack Injector](#)

### Stochastic Modeling:

[QPME \(Queueing Petri net Modeling Environment\)](#)



### Important Links

[SPEC Research Group](#)



[Relate FP7 ITN](#)



[Descartes Modeling Language \(DML\)](#)



[Queueing Petrinet Modeling Environment \(QPME\)](#)



[Interval Standard Working Group P1788](#)

### Upcoming Events

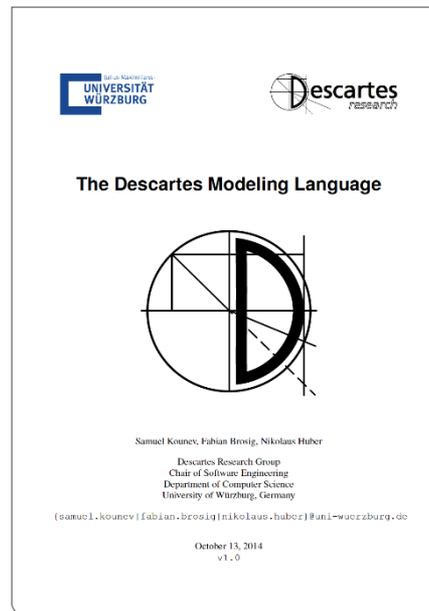
[Int. Conference on Performance Engineering \(ICPE\)](#)

[Dagstuhl Seminar on Self-Aware Computing](#)

[Int. Conference on Autonomic](#)

# Descartes Modeling Language

- **Problem: How to model performance and resource management related aspects of IT systems and infrastructures at the architecture-level?**
  - Prediction of the impact of dynamic changes at run-time
  - Current version focused on performance including capacity, responsiveness and resource efficiency aspects

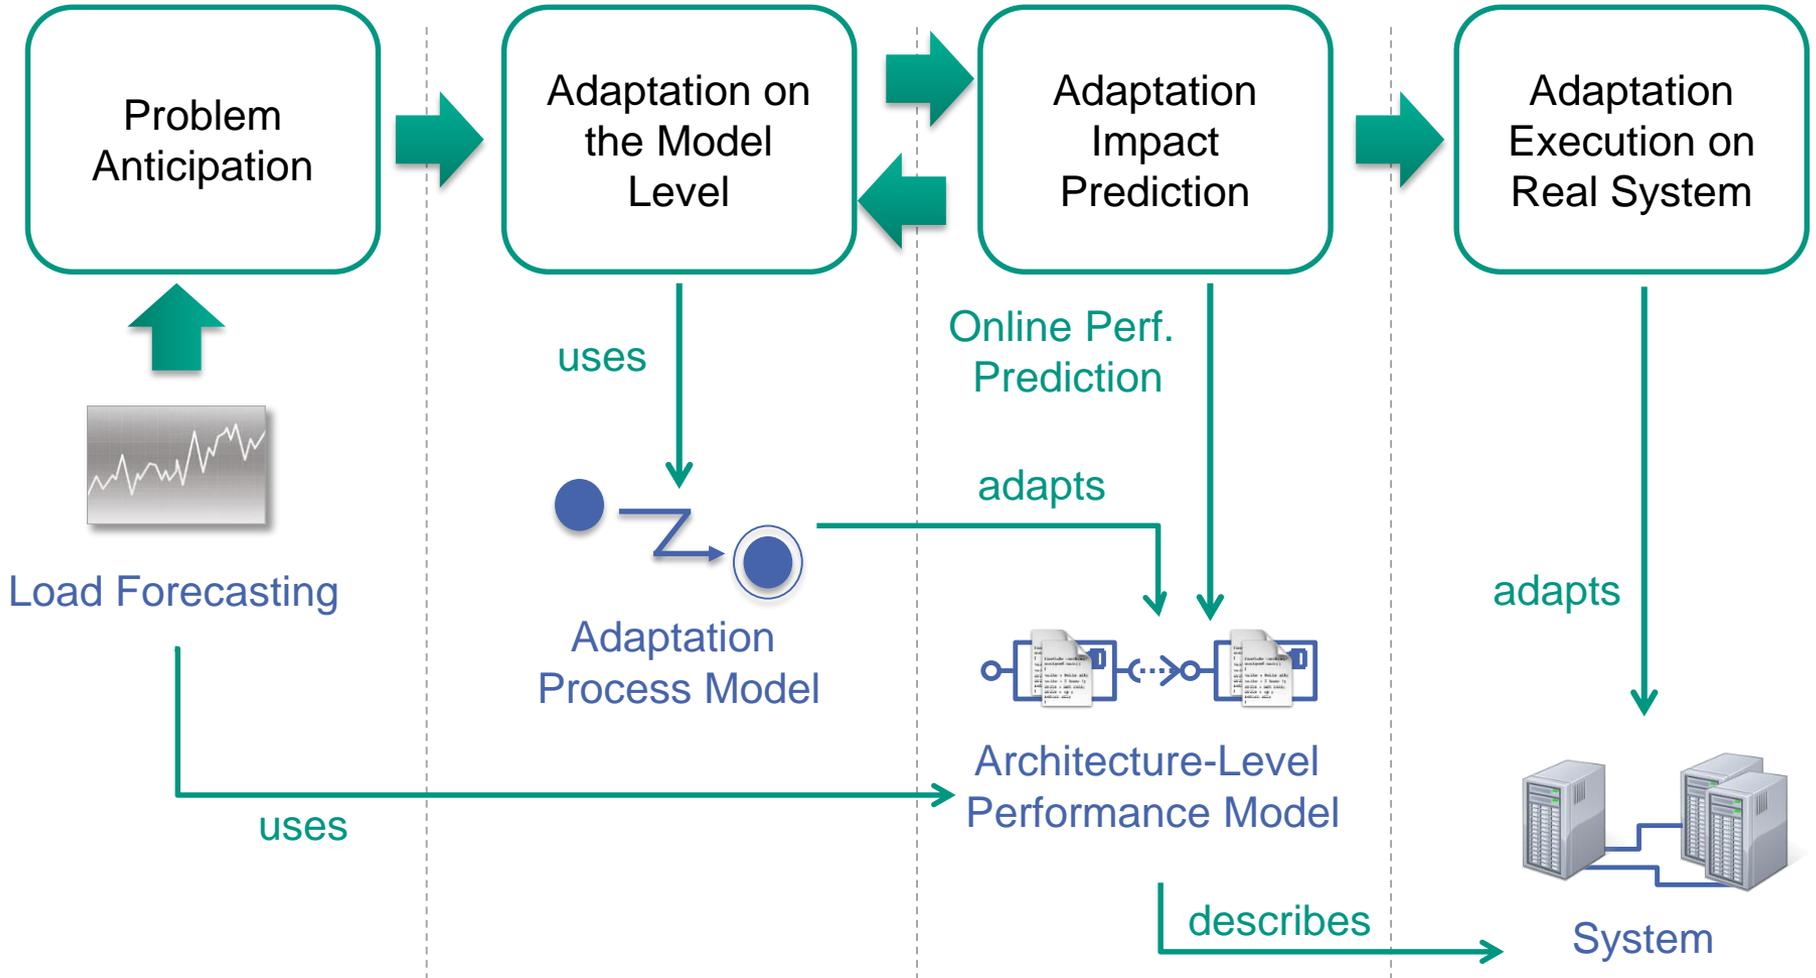


<http://descartes.tools/dml>





# Model-Based System Adaptation



- **Editors**
  - Textual and graphical editors for DML models
- **Solvers**
  - Solvers for conducting performance prediction
- **S/T/A Adaptation Framework**
  - Execution of adaptation process on the model level



**DML Bench**

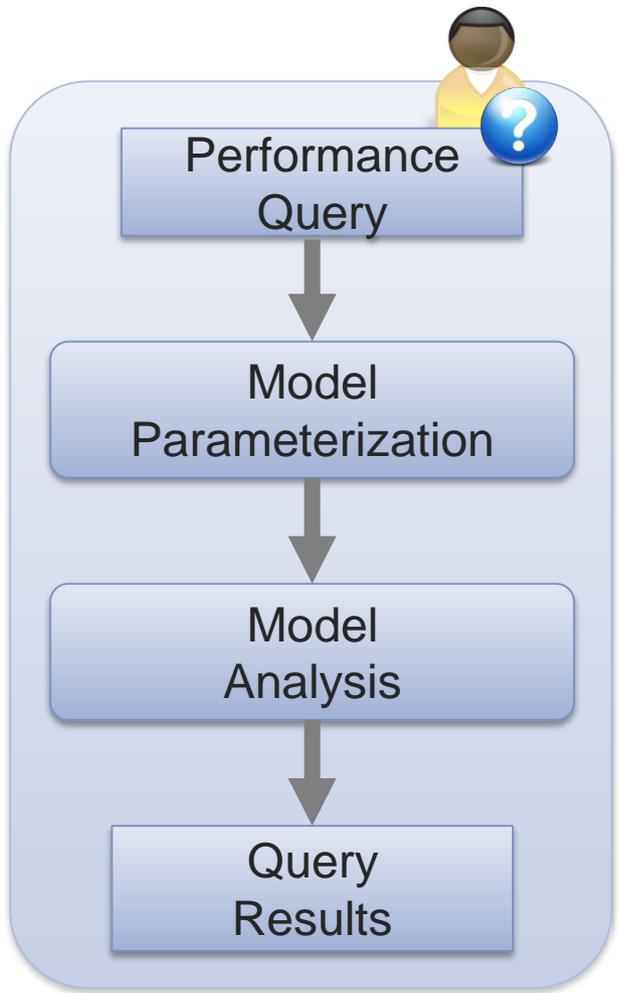
[http://descartes.tools/dml\\_bench](http://descartes.tools/dml_bench)



# DQL (Descartes Query Language)

Example of a performance query specified with DQL

```
SELECT s.avgResponseTime,  
app.utilization,  
dbs.utilization  
CONSTRAINED AS FAST  
FOR RESOURCE  
'ApplicationServer' AS app,  
RESOURCE 'DBServer' AS dbs,  
SERVICE 'processOrder' AS s;
```



<http://descartes.tools/dql>

- **Problem: How to estimate the total service time of a given type of request/job at a given resource?**
- **Library for Resource Demand Estimation**
  - Ready-to-use implementations of estimation approaches
  - Comparison of the accuracy of different approaches
  - Selection of a suitable approach for a given scenario



**<http://descartes.tools/librede>**

S. Spinner, G. Casale, X. Zhu, and S. Kounev. **LibReDE: A Library for Resource Demand Estimation (Demo Paper)**. In *5th ACM/SPEC Intl. Conf. on Performance Engineering (ICPE 2014)*, Dublin, Ireland, March 22-26, 2014, pages 227-228. ACM Press, New York, NY, USA. March 2014. [ [http](http://descartes.tools/librede) | [.pdf](#) ]

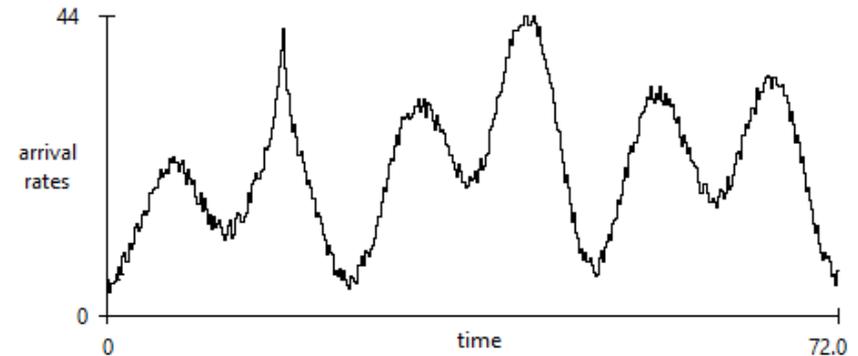


# Estimation Approaches

Technique	Variant	References
Approximation with response times		Urgaonkar et al. [13] Nou et al. [14] Brosig et al. [15]
Service Demand Law		Lazowksa [4] Brosig et al. [15]
Linear regression	Least squares	Bard and Shatzoff [16] Rolia et al. [17, 18] Pacifici et al. [19] Kraft et al. [20, 21]
	Least absolute differences	Zhang et al. [22, 23, 24]
	Least trimmed squares	Casale et al. [25, 26]
Kalman filter		Zheng et al. [27, 28] Kumar et al. [29] Wang et al. [30, 31]
Optimization	Non-linear constrained optimization	Zhang et al. [32] Menascé [33]
	Quadratic programming	Liu et al. [34, 35, 36] Kumar et al. [37]
Machine learning	Clusterwise linear regression	Cremonesi et al. [38]
	Independent component analysis	Sharma et al. [39]
	Support vector machine	Kalbasi et al. [40]
	Pattern matching	Cremonesi et al. [41, 42]
Maximum likelihood estimation		Kraft et al. [20] Perez et al. [21]
Gibbs sampling		Sutton and Jordan [43] Wang et al. [44]
Demand Estimation with Confidence (DEC)		Kalbasi et al. [45, 46]

S. Spinner, G. Casale, F. Brosig, S. Kounev. **Evaluating Approaches to Resource Demand Estimation**. Elsevier Performance Evaluation Journal. Under publication. Available on request.

- Problem: How to capture the load intensity variations (e.g. requests per sec) in a compact mathematical model?
- **Load Intensity Modeling Tool**
  - Automated model extraction from recorded traces
  - Creation and composition of custom models
  - Emulation of job arrivals for load generation

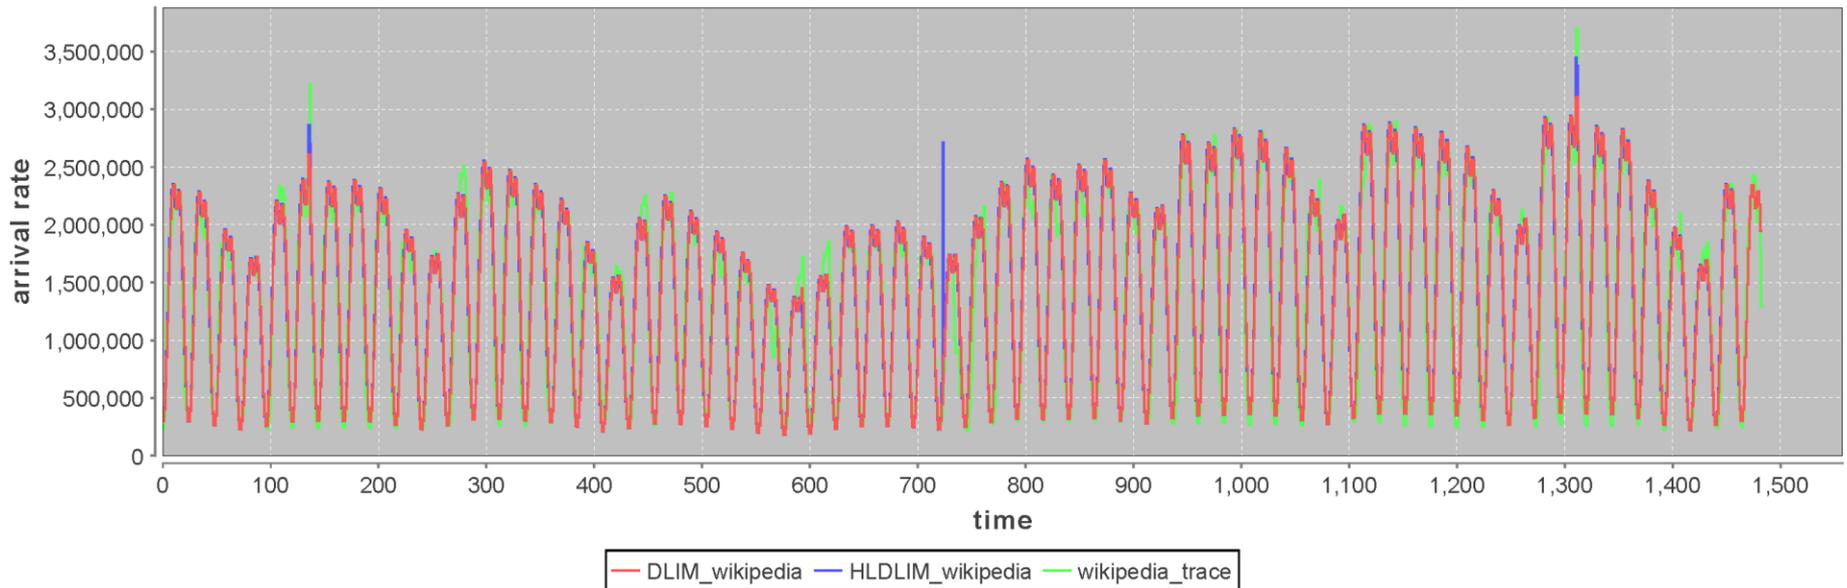


<http://descartes.tools/limbo>

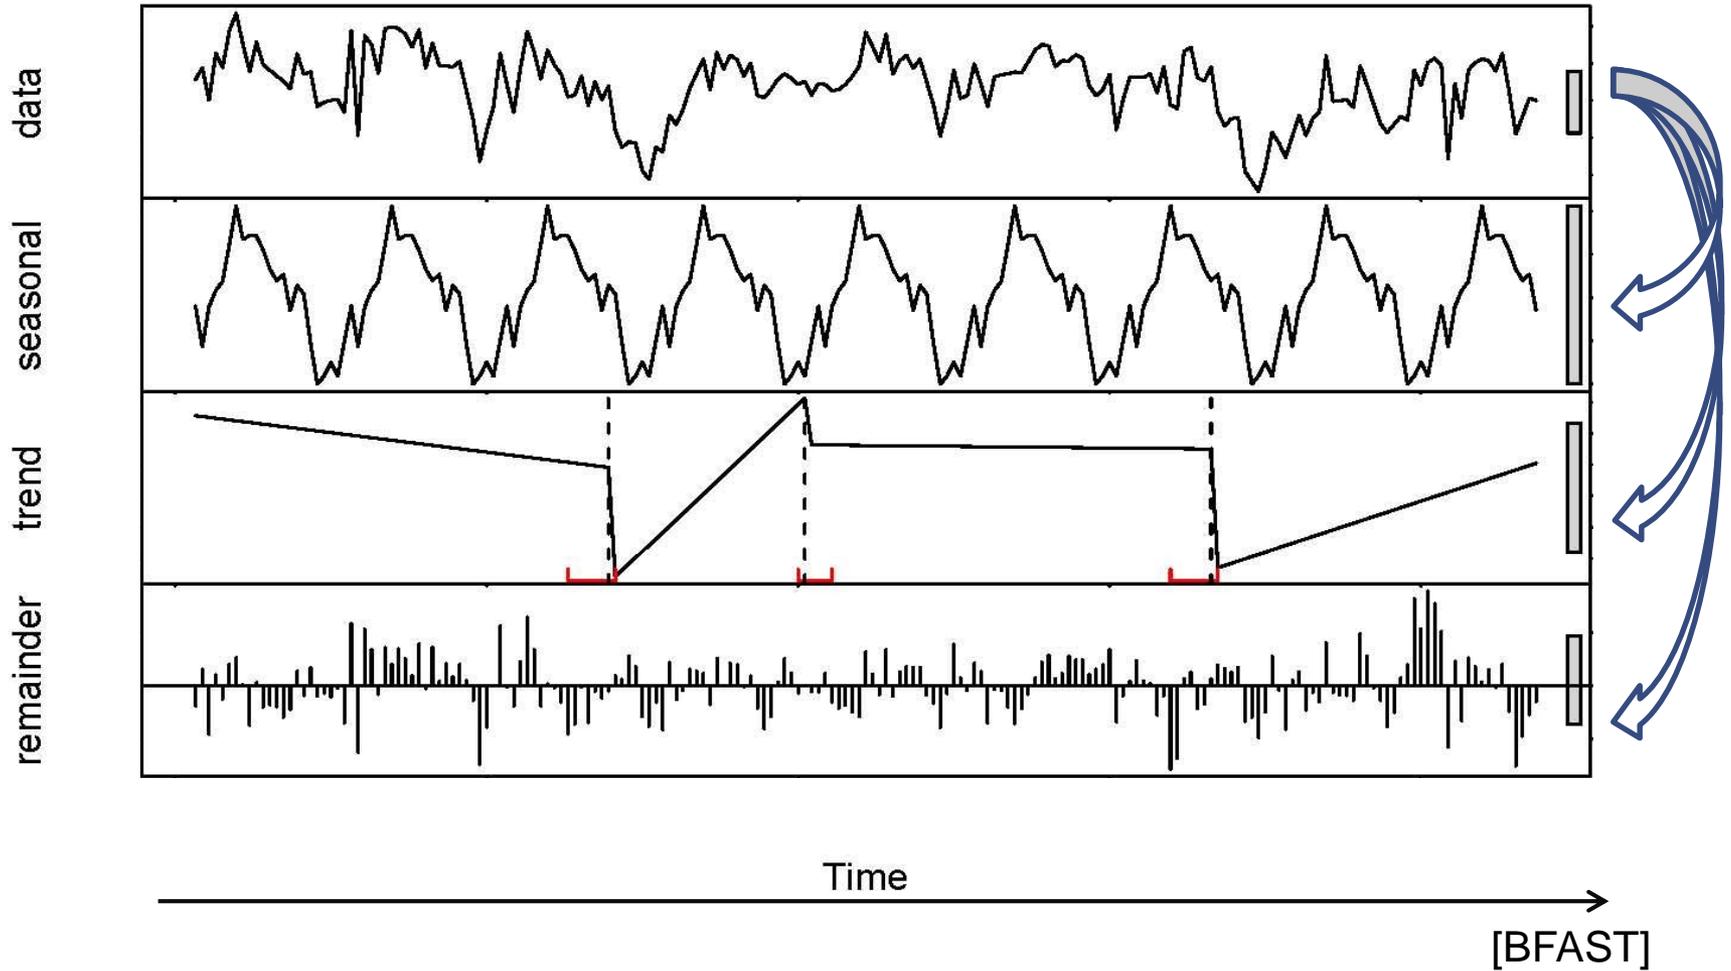


# Example: Wikipedia Workload

**DLIM\_wikipedia Arrival Rates**



# Time Series Analysis





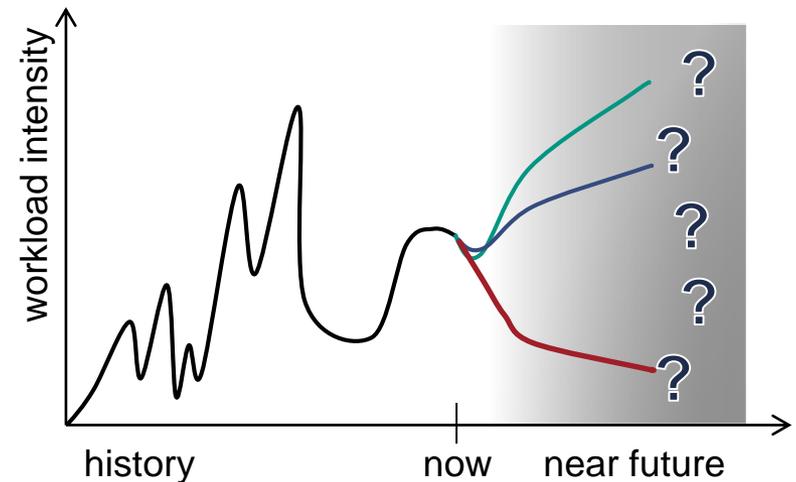
<http://descartes.tools/limbo>

J. von Kistowski, N. Herbst, and S. Kounev. **LIMBO: A Tool For Modeling Variable Load Intensities (Demo Paper)**. In *5th ACM/SPEC International Conference on Performance Engineering (ICPE 2014)*, Dublin, Ireland, March 22-26, 2014, ICPE '14, pages 225-226. ACM, New York, NY, USA. March 2014. [ [DOI](#) | [slides](#) | [http](#) | [.pdf](#) ]

J. von Kistowski, N. Herbst, and S. Kounev. **Modeling Variations in Load Intensity over Time**. In *3rd Intl. Workshop on Large-Scale Testing (LT 2014)*, Dublin, Ireland, March 22, 2014, pages 1-4. ACM, New York, NY, USA. March 2014. [ [DOI](#) | [slides](#) | [http](#) | [.pdf](#) ]



- Problem: How to forecast the load intensity (requests per sec) in future time horizons?
- **Workload Classification & Forecasting Tool**
  - Use of multiple alternative forecasting methods in parallel
  - Selection of method based on its accuracy in the past



<http://descartes.tools/wcf>

## Basic Methods

(initial)

Naïve, Moving Averages, Random Walk

## Trend Interpolation

(fast)

Simple Exponential Smoothing (SES)

[Hynd08]

Cubic Smoothing Splines

[Hynd02]

Croston's method for intermittent time series

[Shen05]

Autoregressive Moving Averages (ARMA11)

[Box08]

## Estimation and Modelling of Seasonal Pattern

(complex)

Extended Exponential Smoothing (ETS)

[Hynd08, Hyn08]

ARIMA framework with automatic model selection

[Box08, Hynd08]

tBATS for complex seasonal patterns

[Live11]



Metrics and benchmarks for quantitative evaluation of

1. Resource elasticity
2. Performance isolation
3. Intrusion detection (and prevention)
4. ...

S. Kounev. **Quantitative Evaluation of Service Dependability in Shared Execution Environments** (Keynote Talk). In 11th Intl. Conf. on Quantitative Evaluation of SysTems (QEST 2014), Florence, Italy, September 8-12, 2014. [ [slides](#) | [extended abstract](#) ]



Def: The degree to which a system is able to **adapt** to **workload changes** by **provisioning and deprovisioning** resources in an **autonomic manner**, such that at each point in time the **available resources match** the **current demand** as closely as possible.

*N. Herbst, S. Kounev and R. Reussner*

***Elasticity: What it is, and What it is Not.***

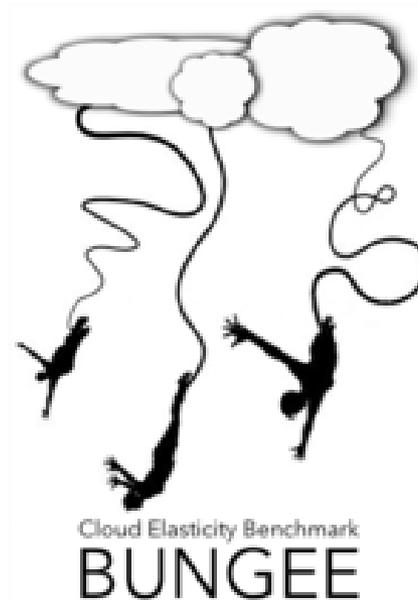
*in Proceedings of the 10th International Conference on Autonomic Computing (ICAC 2013), San Jose, CA, June 24-28, 2013.*

[ [slides](#) | [http](#) | [.pdf](#) ]

[http://en.wikipedia.org/wiki/Elasticity\\_\(cloud\\_computing\)](http://en.wikipedia.org/wiki/Elasticity_(cloud_computing))

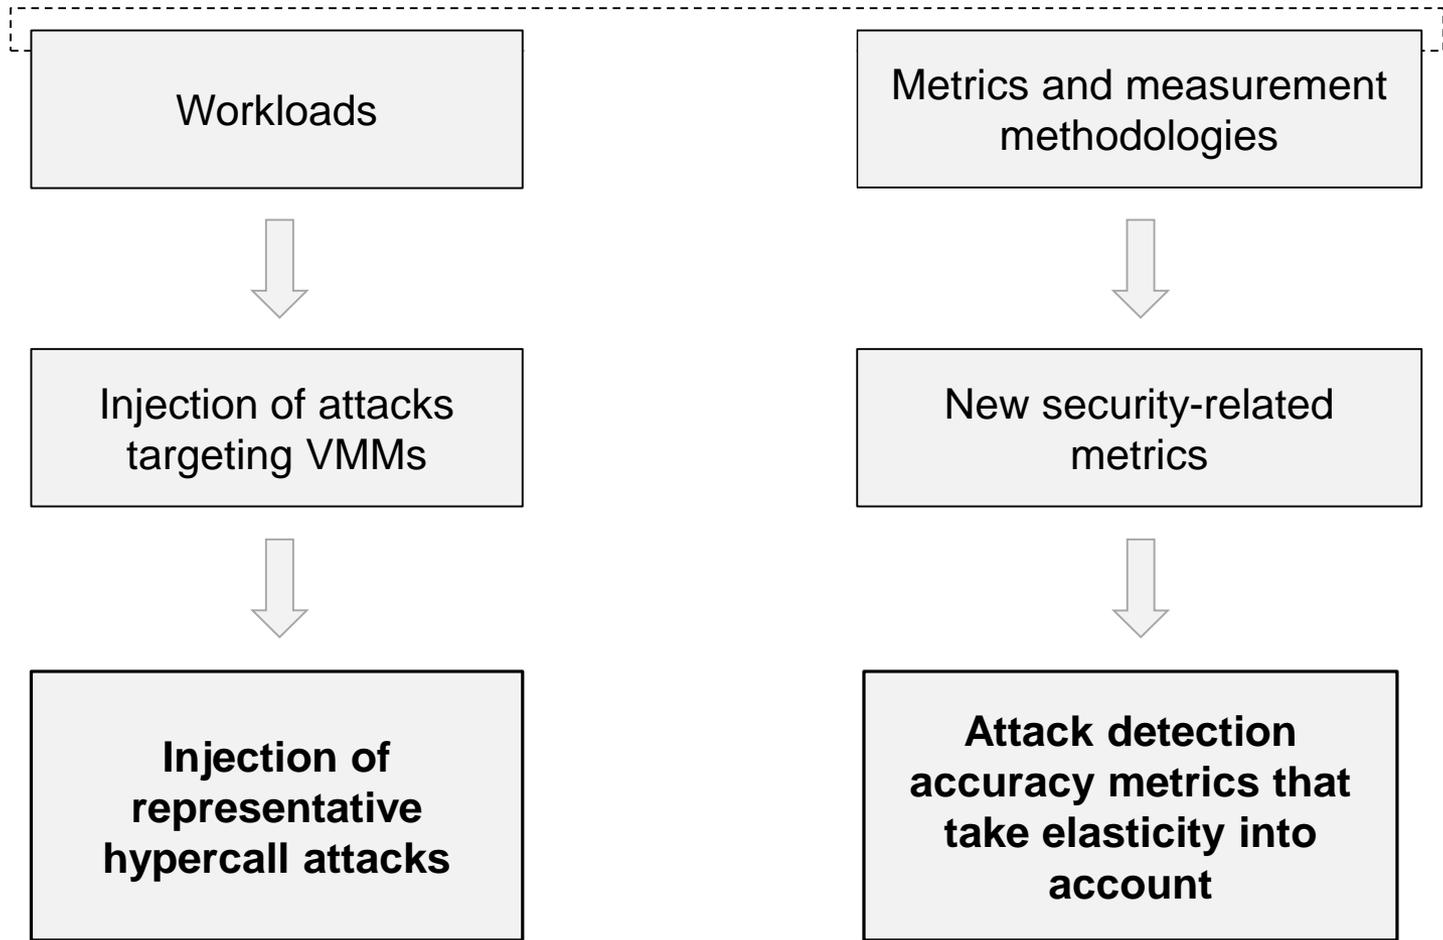


- Problem: How to measure the level of elasticity provided by an IaaS platform?
- Framework for benchmarking elasticity
  - Current focus: IaaS cloud platforms



<http://descartes.tools/bungee>

## IDS evaluation in virtualized environments



## Collaboration with

Marco Vieira and Nuno Antunes, University of Coimbra, Portugal  
Bryan D. Payne, Department of Security Research, Nebula Inc.  
Alberto Avritzer, Siemens Corporate Research, USA

## Main references

A. Milenkoski, B. Payne, N. Antunes, M. Vieira and S. Kounev. **An Analysis of Hypercall Handler Vulnerabilities**. In *Proc. of 25th IEEE Intl. Symp. on Software Reliability Engineering (ISSRE 2014) - Research Track*, Naples, Italy, November 2014. IEEE.

A. Milenkoski, B. Payne, N. Antunes, M. Vieira and S. Kounev. **hInjector: Injecting Hypercall Attacks for Evaluating VMI-based Intrusion Detection Systems** (Poster Paper). In *2013 Annual Computer Security Applications Conf. (ACSAC 2013)*, New Orleans, Louisiana, USA, 2013. [ [.pdf](#) ]

## Further references

A. Milenkoski, S. Kounev, A. Avritzer, N. Antunes and M. Vieira. **On Benchmarking Intrusion Detection Systems in Virtualized Environments**. Technical Report SPEC-RG-2013-002 v.1.0, SPEC Research Group - IDS Benchmarking Working Group, Standard Performance Evaluation Corporation (SPEC), June 2013. [ [.pdf](#) ]

A. Milenkoski, M. Vieira, B. Payne, N. Antunes and S. Kounev. **Technical Information on Vulnerabilities of Hypercall Handlers**. Technical Report SPEC-RG-2014-001 v.1.0, SPEC Research Group - IDS Benchmarking Working Group, Standard Performance Evaluation Corporation (SPEC), August 2014. [ [.pdf](#) ]



# Links for Further Information

- **DML** – Descartes Modeling Language ([homepage](#), [publications](#))
- **DML Bench** ([homepage](#), [publications](#))
- **DQL** – Declarative query language ([homepage](#), [publications](#))
- **DNI** – Descartes network infrastructure modeling ([homepage](#), [publications](#))
- **LibReDE** - Library for resource demand estimation ([homepage](#), [publications](#))
- **LIMBO** – Load intensity modeling tool ([homepage](#), [publications](#))
- **WCF** – Workload classification & forecasting tool ([homepage](#), [publications](#))
- **BUNGEE** – Elasticity benchmarking framework ([homepage](#), [publications](#))
- **hInjector** – Security benchmarking tool ([homepage](#), [publications](#))
- **Further relevant research**
  - [http://descartes-research.net/research/research\\_areas/](http://descartes-research.net/research/research_areas/)
  - **Self Aware Computing** ([publications](#))



- **Open-Systems-Group (OSG)**
  - Processor and computer architectures
  - Virtualization platforms
  - Java (JVM, Java EE)
  - Message-based systems
  - Storage systems (SFS)
  - Web-, email- and file server
  - SIP server (VoIP)
  - Cloud computing
- **High-Performance-Group (HPG)**
  - Symmetric multiprocessor systems
  - Workstation clusters
  - Parallel and distributed systems
  - Vector (parallel) supercomputers
- **“Graphics and Workstation Performance Group” (GWPG)**
  - CAD/CAM, visualization
  - OpenGL

<http://www.spec.org>



# SPEC Research Group (RG)

- Founded in March 2011: <http://research.spec.org>
  - Transfer of knowledge btw. academia and industry
- Activities
  - Methods and techniques for experimental system analysis
  - Standard metrics and measurement methodologies
  - Benchmarking and certification
  - Evaluation of academic research results
- Member organizations (Feb 2014)



# Questions?



<http://www.descartes-research.net>

<http://descartes.tools>

# Thank You!

[skounev@acm.org](mailto:skounev@acm.org)

<http://se.informatik.uni-wuerzburg.de>