NWO STW COMMIT/

Massivizing Computer Systems = Making Computer Systems Scalable, Reliable, High-Performance, etc., Yet Form an Efficient Ecosystem

TUDelft

VU VRIJE UNIVERSITEIT AMSTERDAM

@AIosup

dr. ir. Alexandru Iosup

This Is the Golden Age of Computer Systems

Education for Everyone (Online)

Business Services

Big Science

Grid Computing

Online Gaming

Datacenters

Daily Life

AVERAGE DAILY ONLINE GAMERS WORLDWIDE

Source: comScore MMX, Worldwide, April 2013, Age 15+

This Is the Golden Age of Computer Systems

Do you recognize this App?

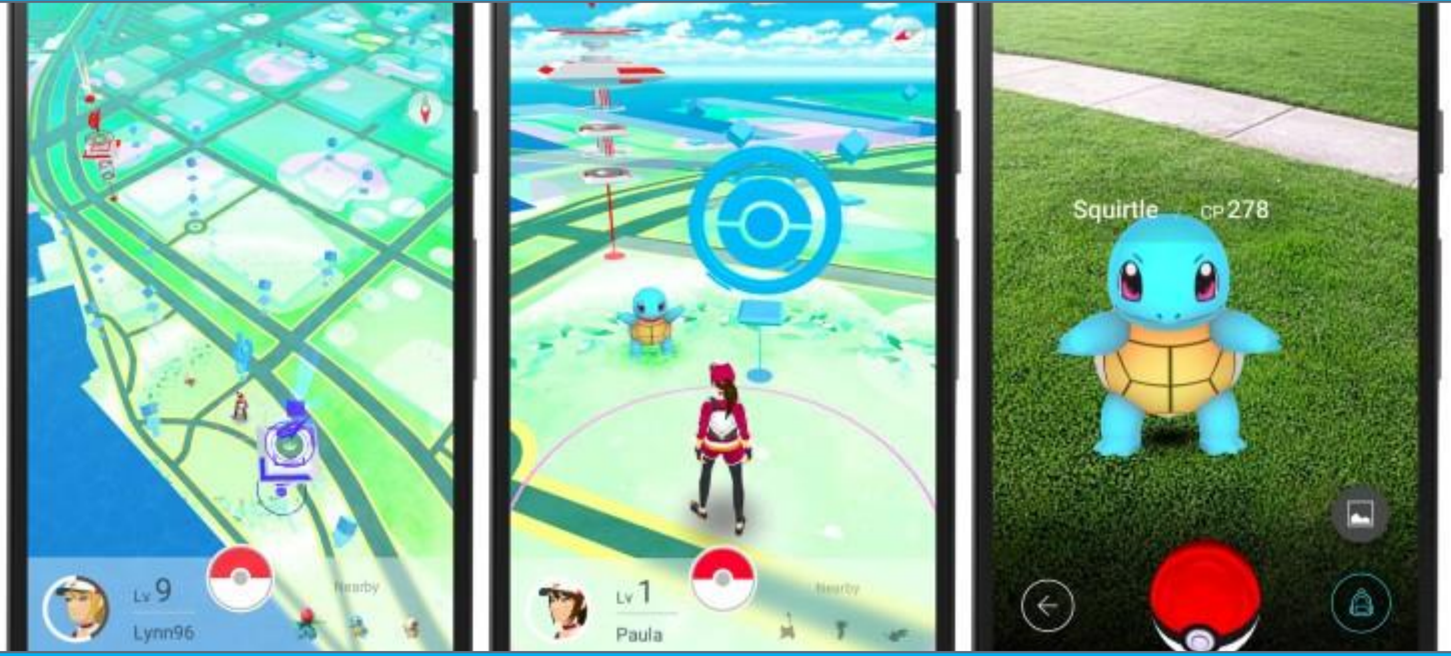

Here is how it operates…

Current Technology: Scheduler? Datacenter? Etc.

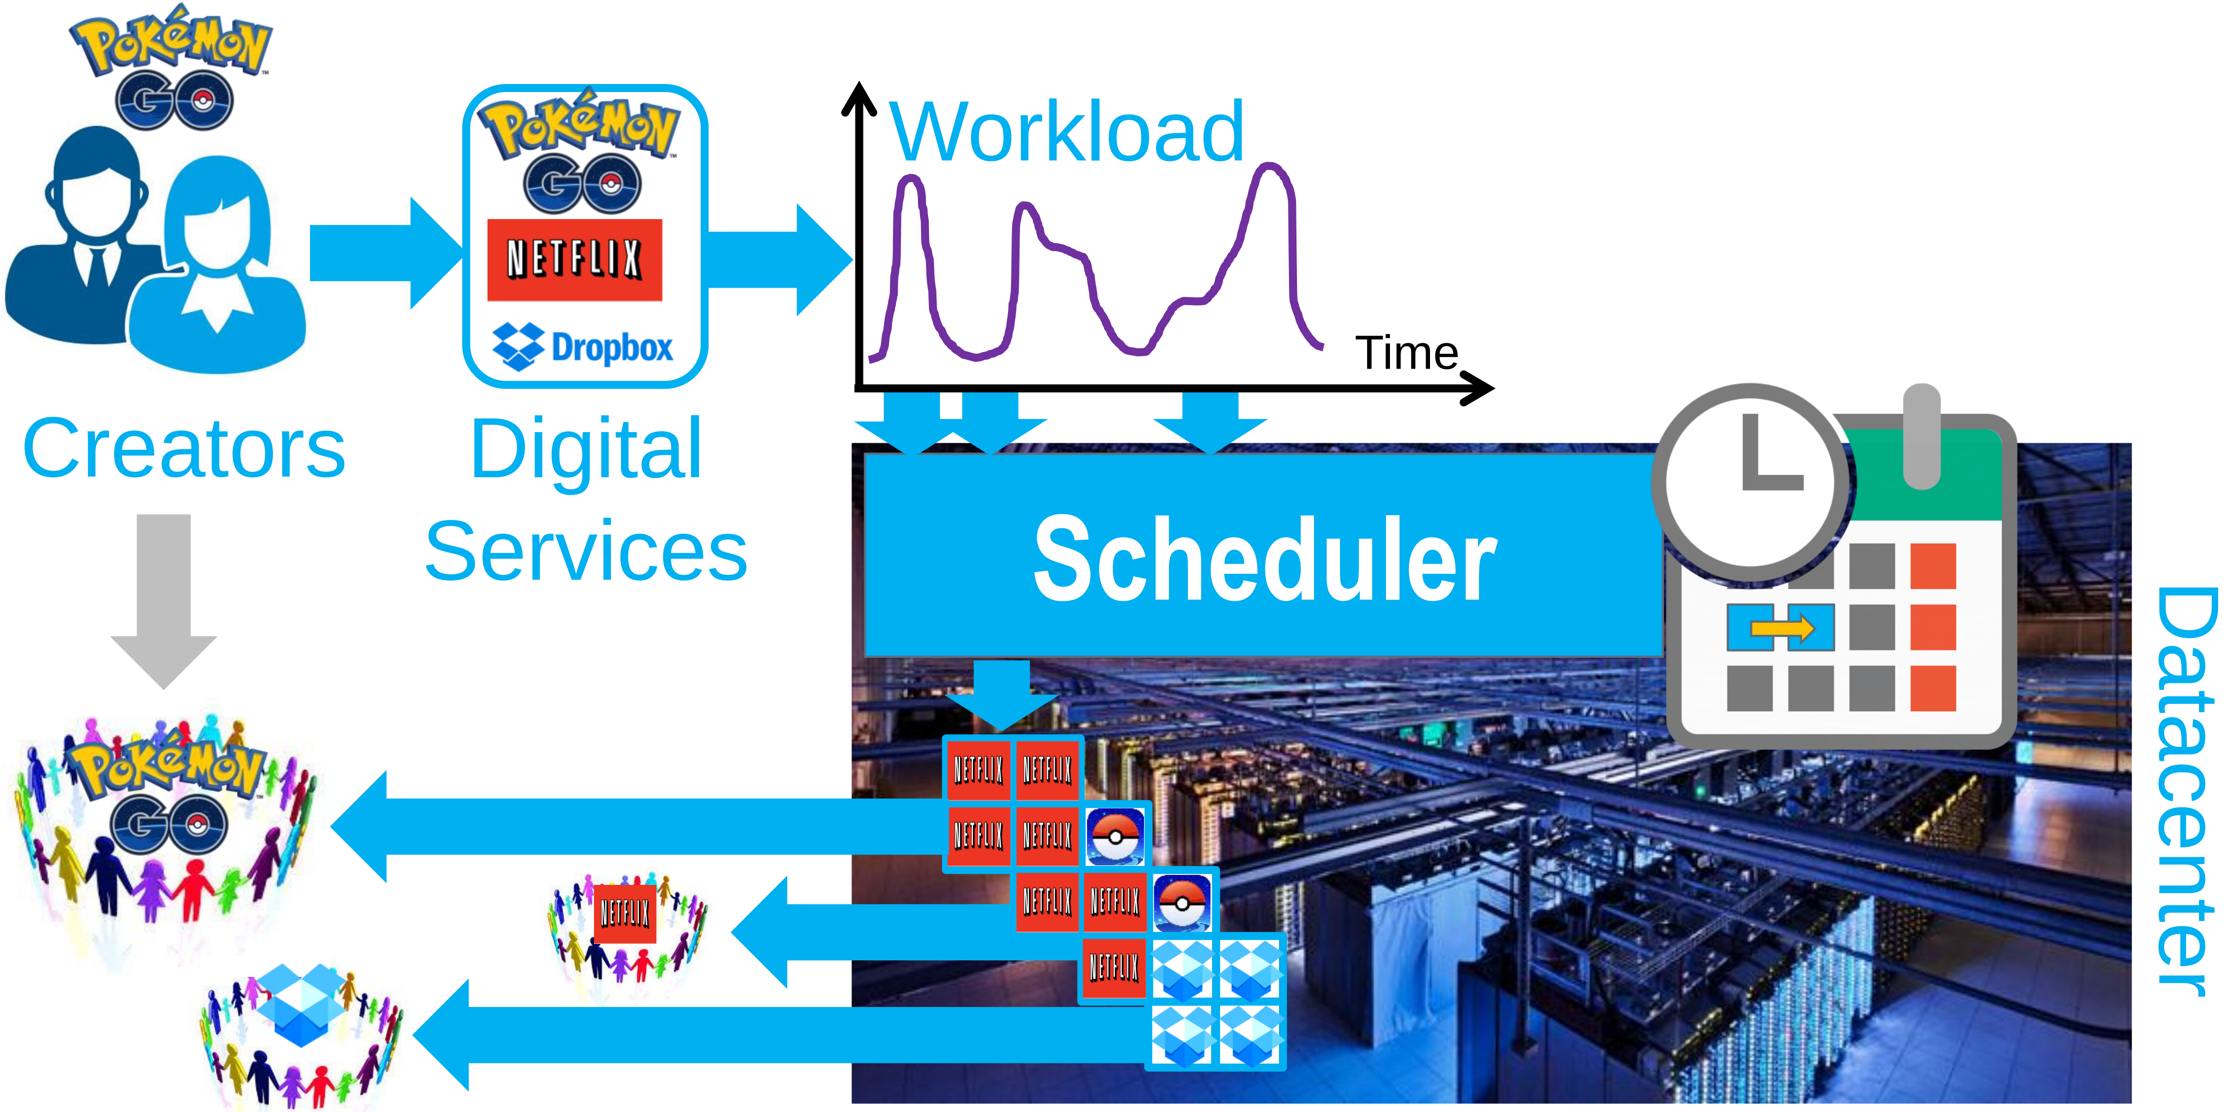

Creators

Digital Services

Workload

Time

Scheduler

Datacenter

Performance, Dependability, Efficiency

The Golden Age of Computer Systems
... Yet We Are in a Crisis

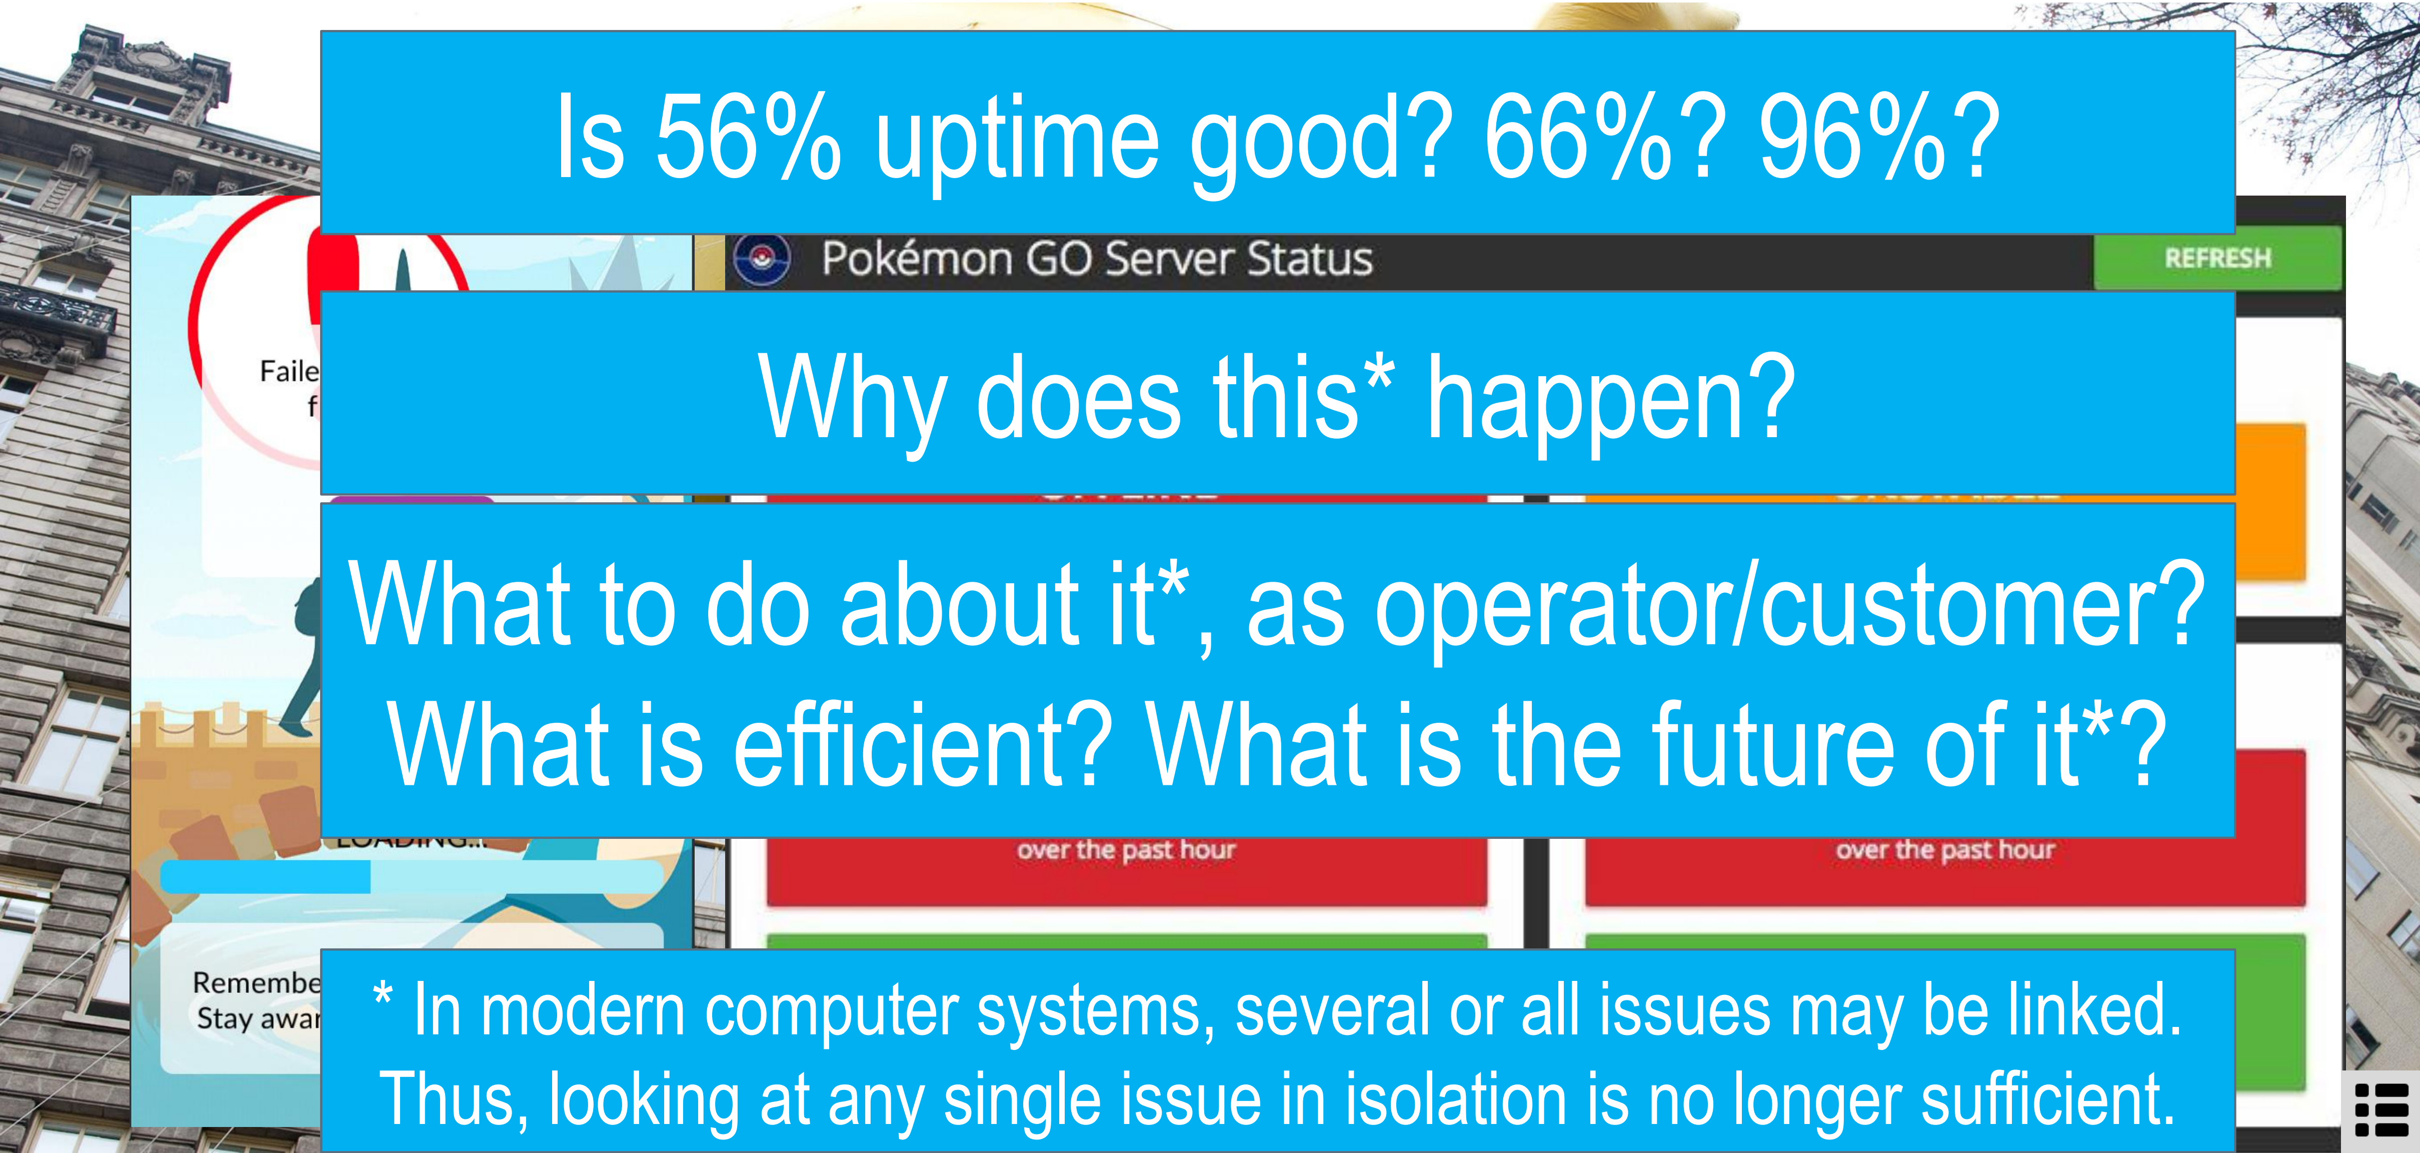

Is 56% uptime good? 66%? 96%?

Why does this* happen?

What to do about it*, as operator/customer? What is efficient? What is the future of it*?

* In modern computer systems, several or all issues may be linked. Thus, looking at any single issue in isolation is no longer sufficient.

This Is the Golden Age of Computer Systems

… Yet We Are in a Crisis

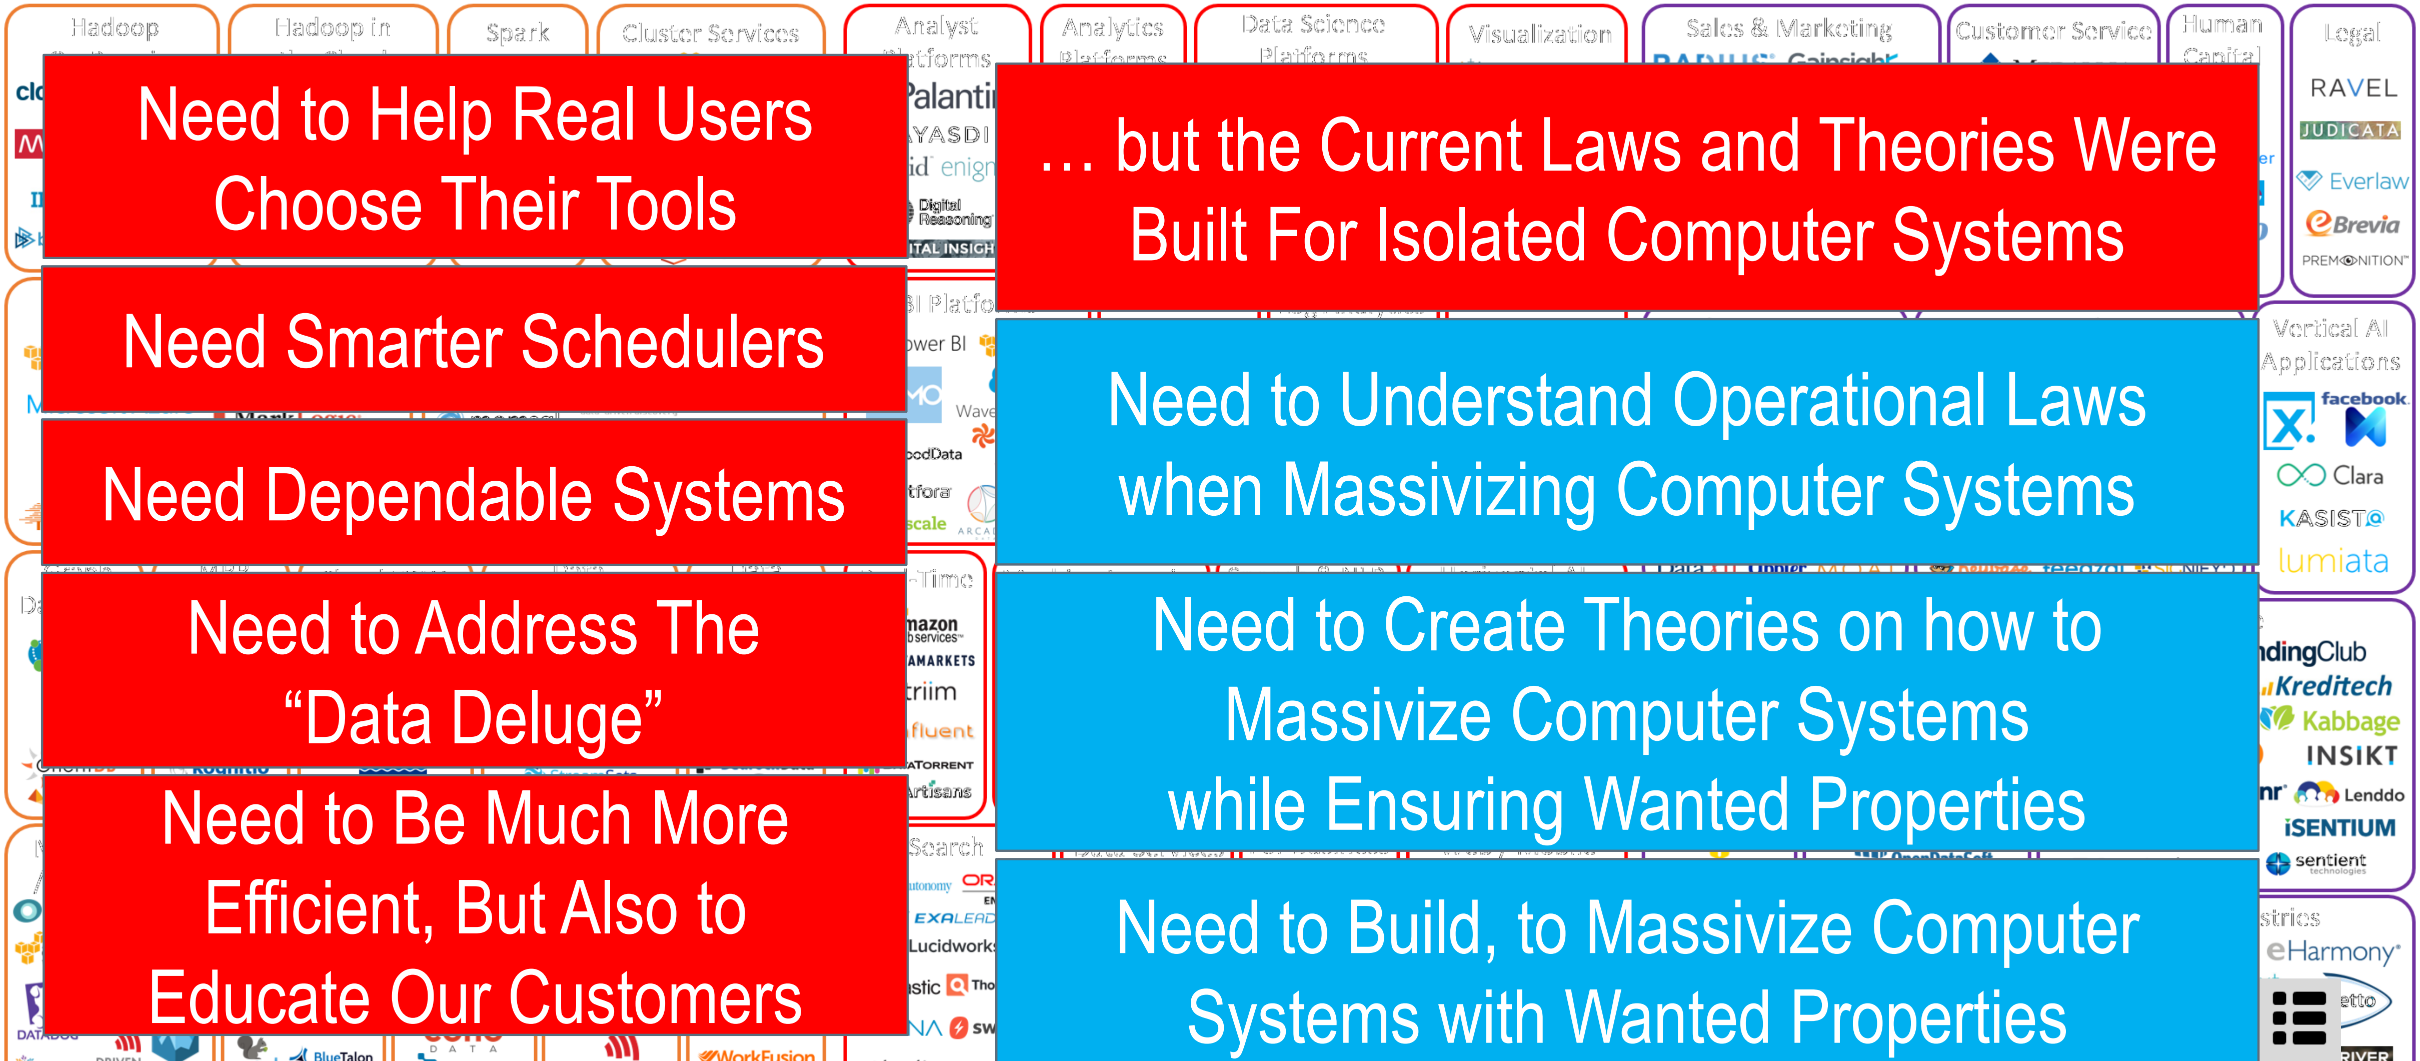

Need to Help Real Users Choose Their Tools

Need Smarter Schedulers

Need Dependable Systems

Need to Address The “Data Deluge”

Need to Be Much More Efficient, But Also to Educate Our Customers

… but the Current Laws and Theories Were Built For Isolated Computer Systems

Need to Understand Operational Laws when Massivizing Computer Systems

Need to Create Theories on how to Massivize Computer Systems while Ensuring Wanted Properties

Need to Build, to Massivize Computer Systems with Wanted Properties

This Is the Golden Age of Computer Systems

… Yet We Are in a Crisis

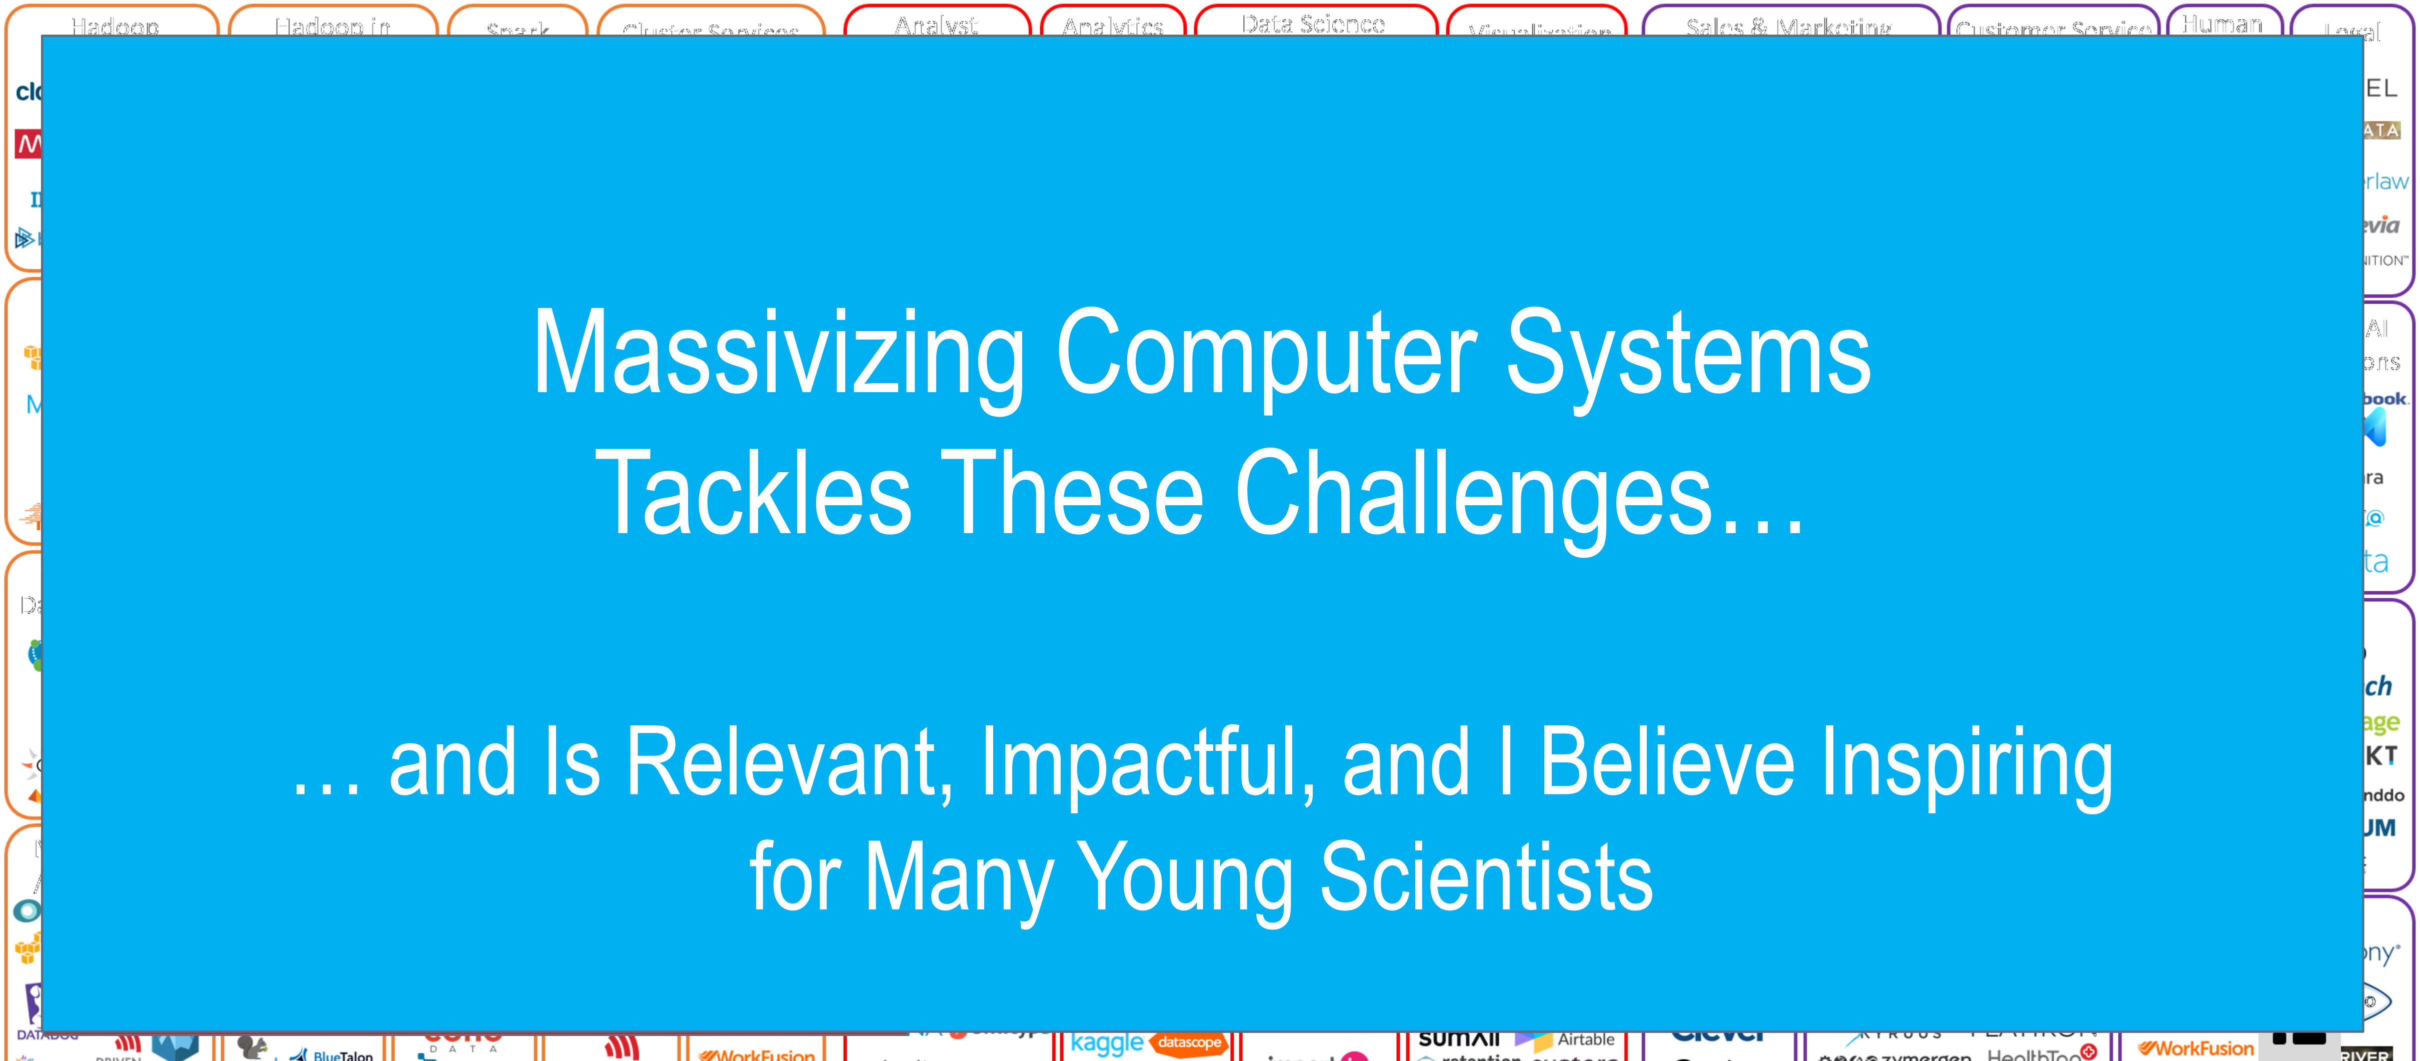

Massivizing Computer Systems Tackles These Challenges…

… and Is Relevant, Impactful, and I Believe Inspiring for Many Young Scientists

My Story From Now On…
Massivizing Computer Systems
(Problem Found)

In Pasteur’s Quadrant+:

- Fundamental research
- Inspired by real use
- Experimental in nature

+ Please ask for an example

Fundamental Research in Massivizing Comp. Sys.

Scheduling

Bags-Of-Tasks

Workflows !

Portfolio !

Dependability

Failure Analysis*

Space-/Time-Correlation

Availability-On-Demand

New World+

Workload Modeling

Business-Critical

Online Gaming

Ecosystem Navigator+

Performance Variability

Grid*, Cloud, Big Data

Benchmarking

Longitudinal Studies

Scalability/Elasticity+

Delegated Matchmaking*

POGGI*, AoS

BTWorld*

Auto-Scalers

Heterogeneous Systems

Socially Aware+

Collaborative Downloads*

Groups in Online Gaming

Toxicity Detection*

Interaction Graphs !

Software Artifacts !

Graphalytics, etc.

Data Artifacts !

A Distributed Systems Memex*

Fundamental Problems/Research Lines

My Contribution So Far

! VIDI-funded

+ Please ask for a definition

* Award-winning (level of ambition)

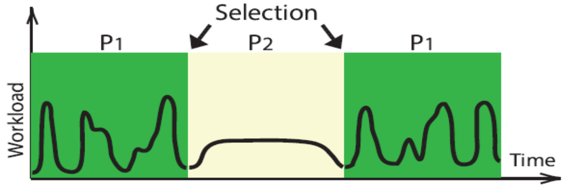

An Example:
Portfolio Scheduling for Datacenters

(what's in a name)

Portfolio Scheduling, In A Nutshell

- Datacenters cannot work without one or even several schedulers
- Instead of ephemeral, risky schedulers, I propose to

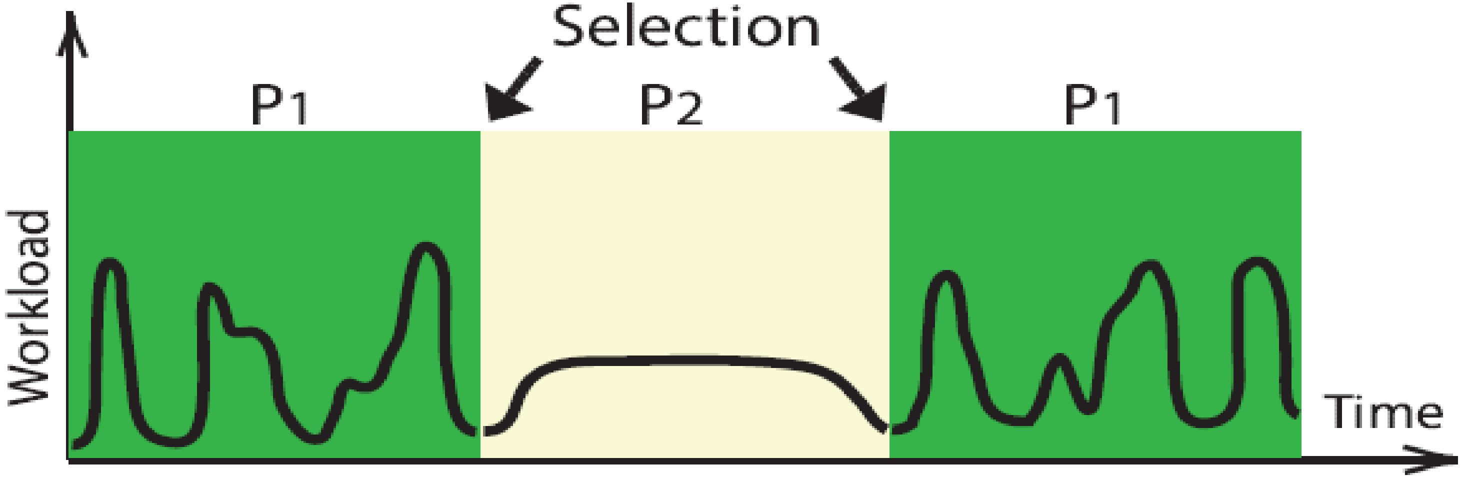

1. Create a set of schedulers
   - Resource provisioning and allocation policies for datacenters
2. Select active scheduler online, apply for the next period, analyze results

(Repeat)

Portfolio Scheduling

Portfolio Scheduling for Computer Systems

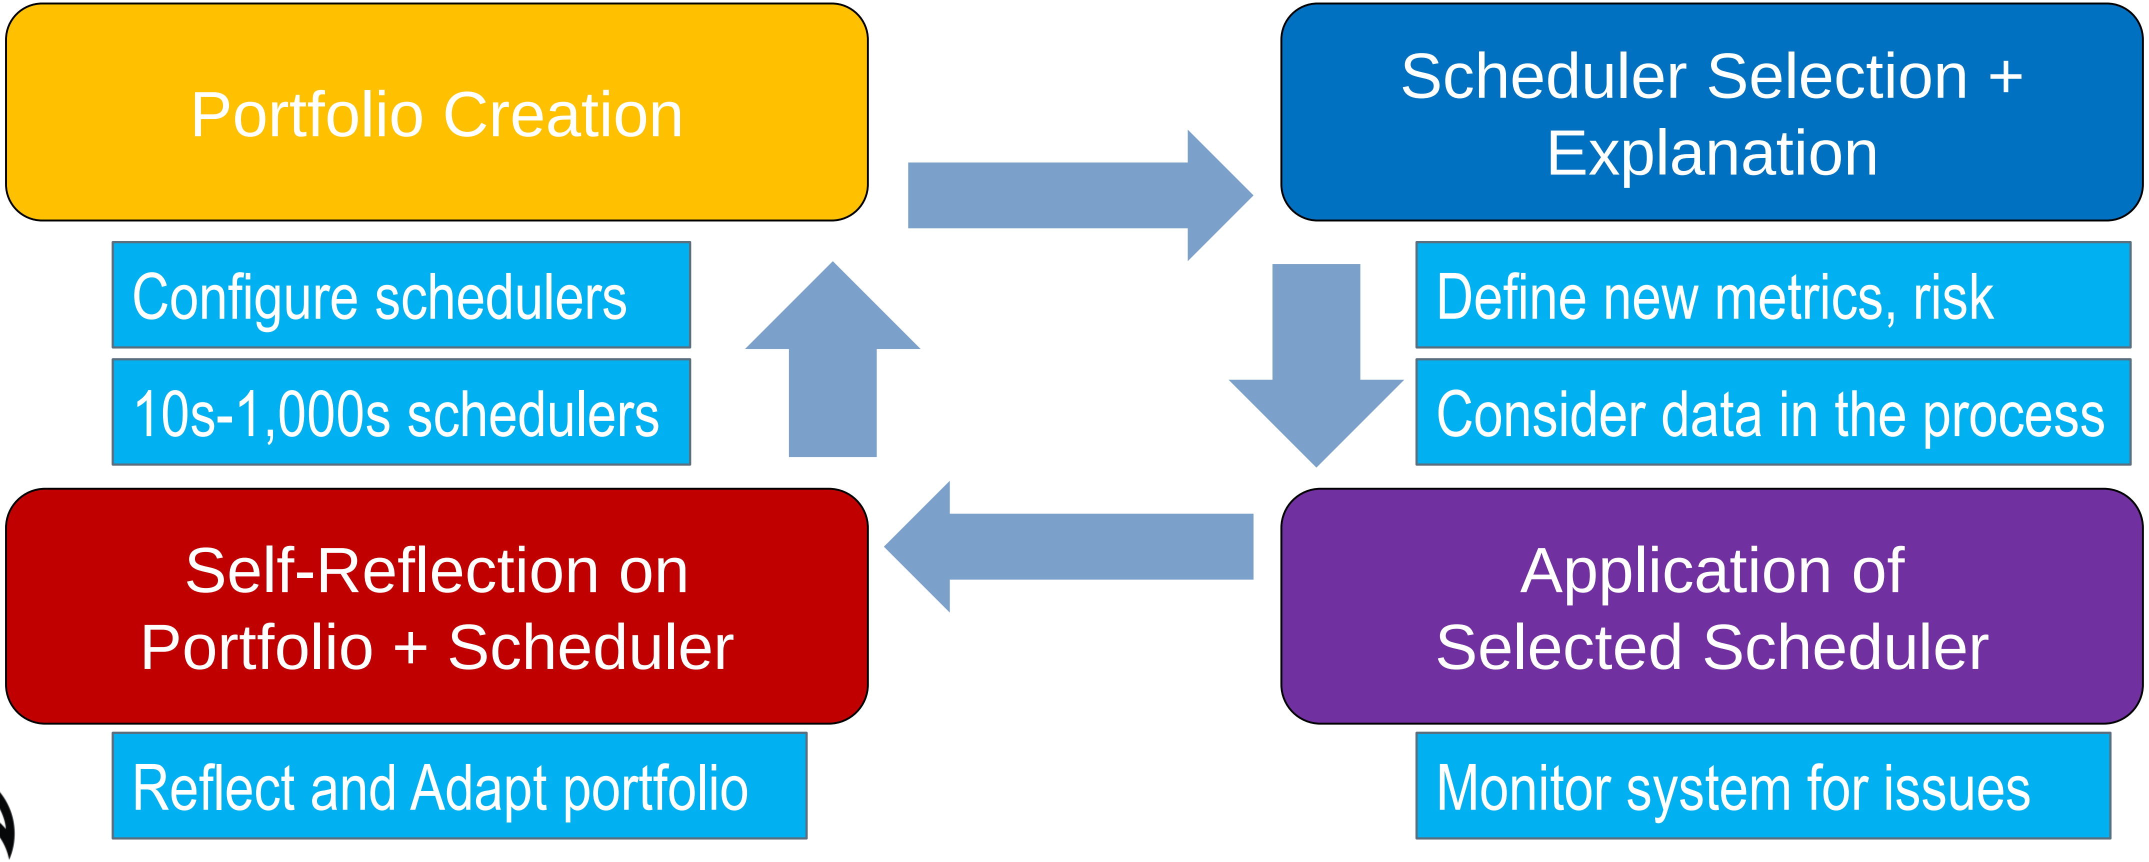

Portfolio Creation

Configure schedulers

10s-1,000s schedulers

Scheduler Selection + Explanation

Define new metrics, risk

Consider data in the process

Application of Selected Scheduler

Monitor system for issues

Self-Reflection on Portfolio + Scheduler

Reflect and Adapt portfolio

Experimental Research Methodology
My Main Scientific Instrument: DAS-5

My (& Your) Prototypes

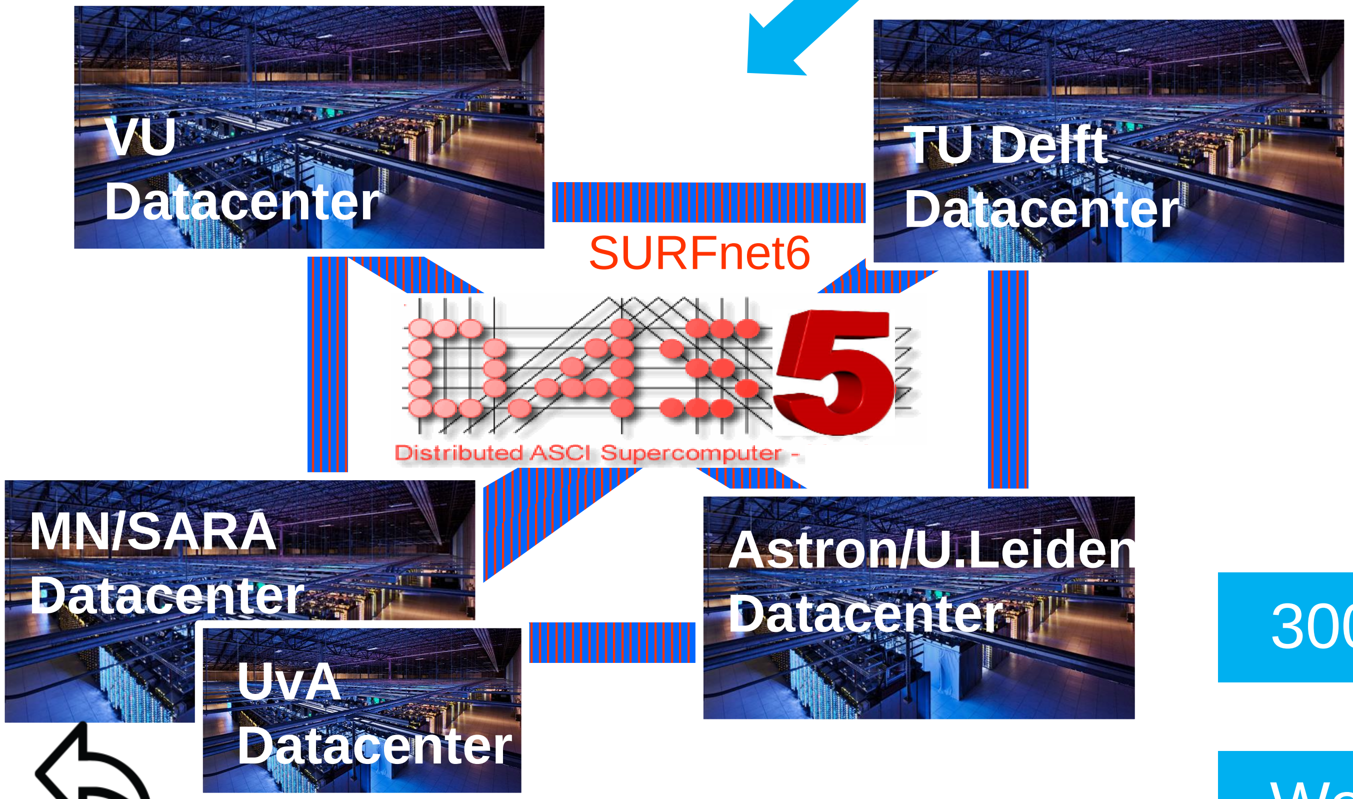

300+ scientists as users

We won IEEE Scale Challenge 2014

Portfolio Scheduling in Practice: Massive Datacenters

Policy selected, fraction of decisions

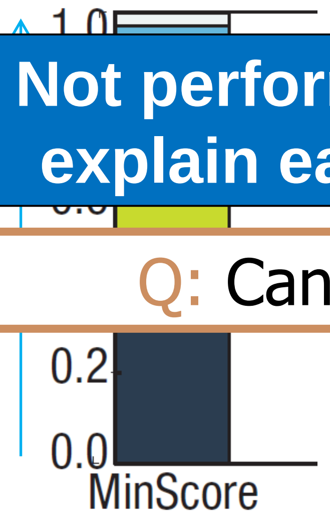

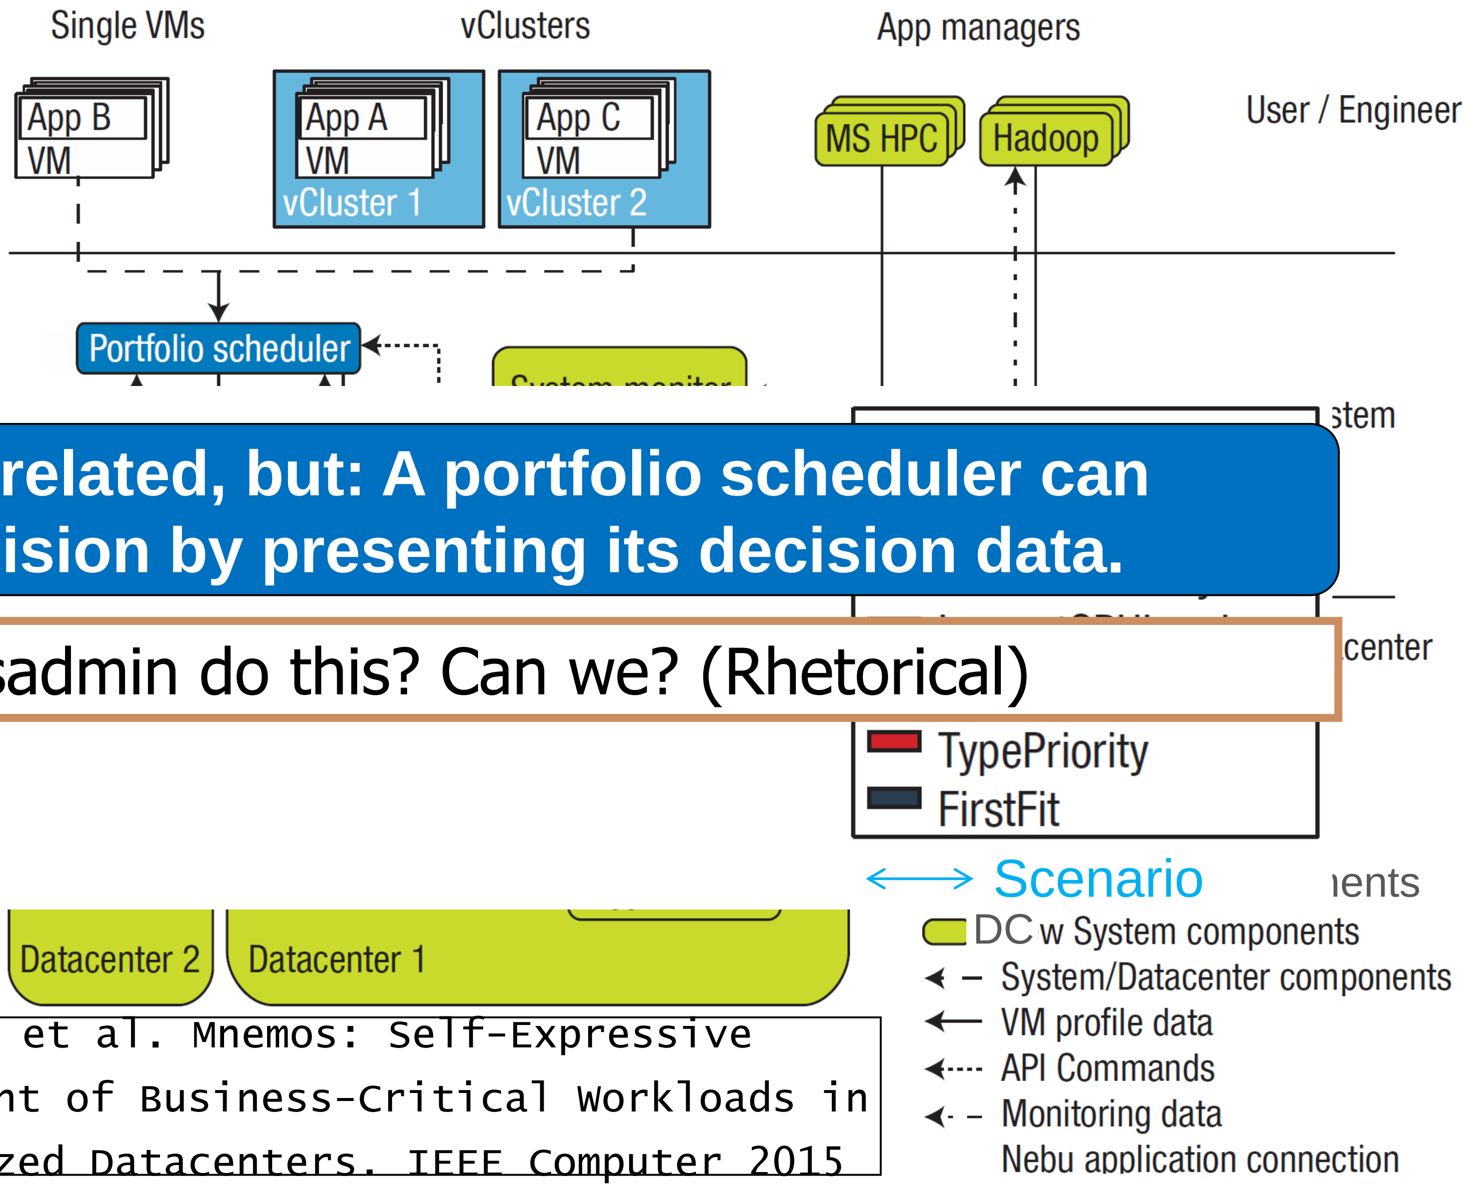

Scenario

Not performance-related, but: A portfolio scheduler can explain each decision by presenting its decision data.

Q: Can our sysadmin do this? Can we? (Rhetorical)

V. van Beek et al. Mnemos: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters. IEEE Computer 2015

End of Example:
Portfolio Scheduling for Datacenters

(what’s in a name)

Take-Home Message: Massivizing Computer Systems

- The Golden Age of Computer Systems
- My Research is about Massivizing Computer Systems
  - Research approach: Pasteur’s Quadrant
  - Fundamental research lines
- An Example: Portfolio Scheduling
- **General Questions** ← we are here now

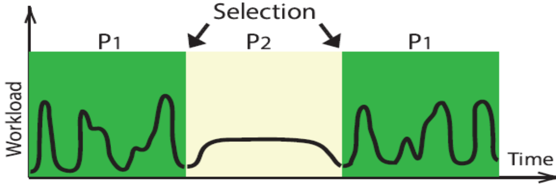

Education for Everyone (Online)
Busines s Service s
Cloud Computing
Grid Computing
What is all about Grid ?
1) Sharing, Selecting, aggregation of resources with a policy universally agreed.
2) Distributive supercomputing for a better future
Grid

Many Thanks to Many!
(My @large Team at TU Delft)

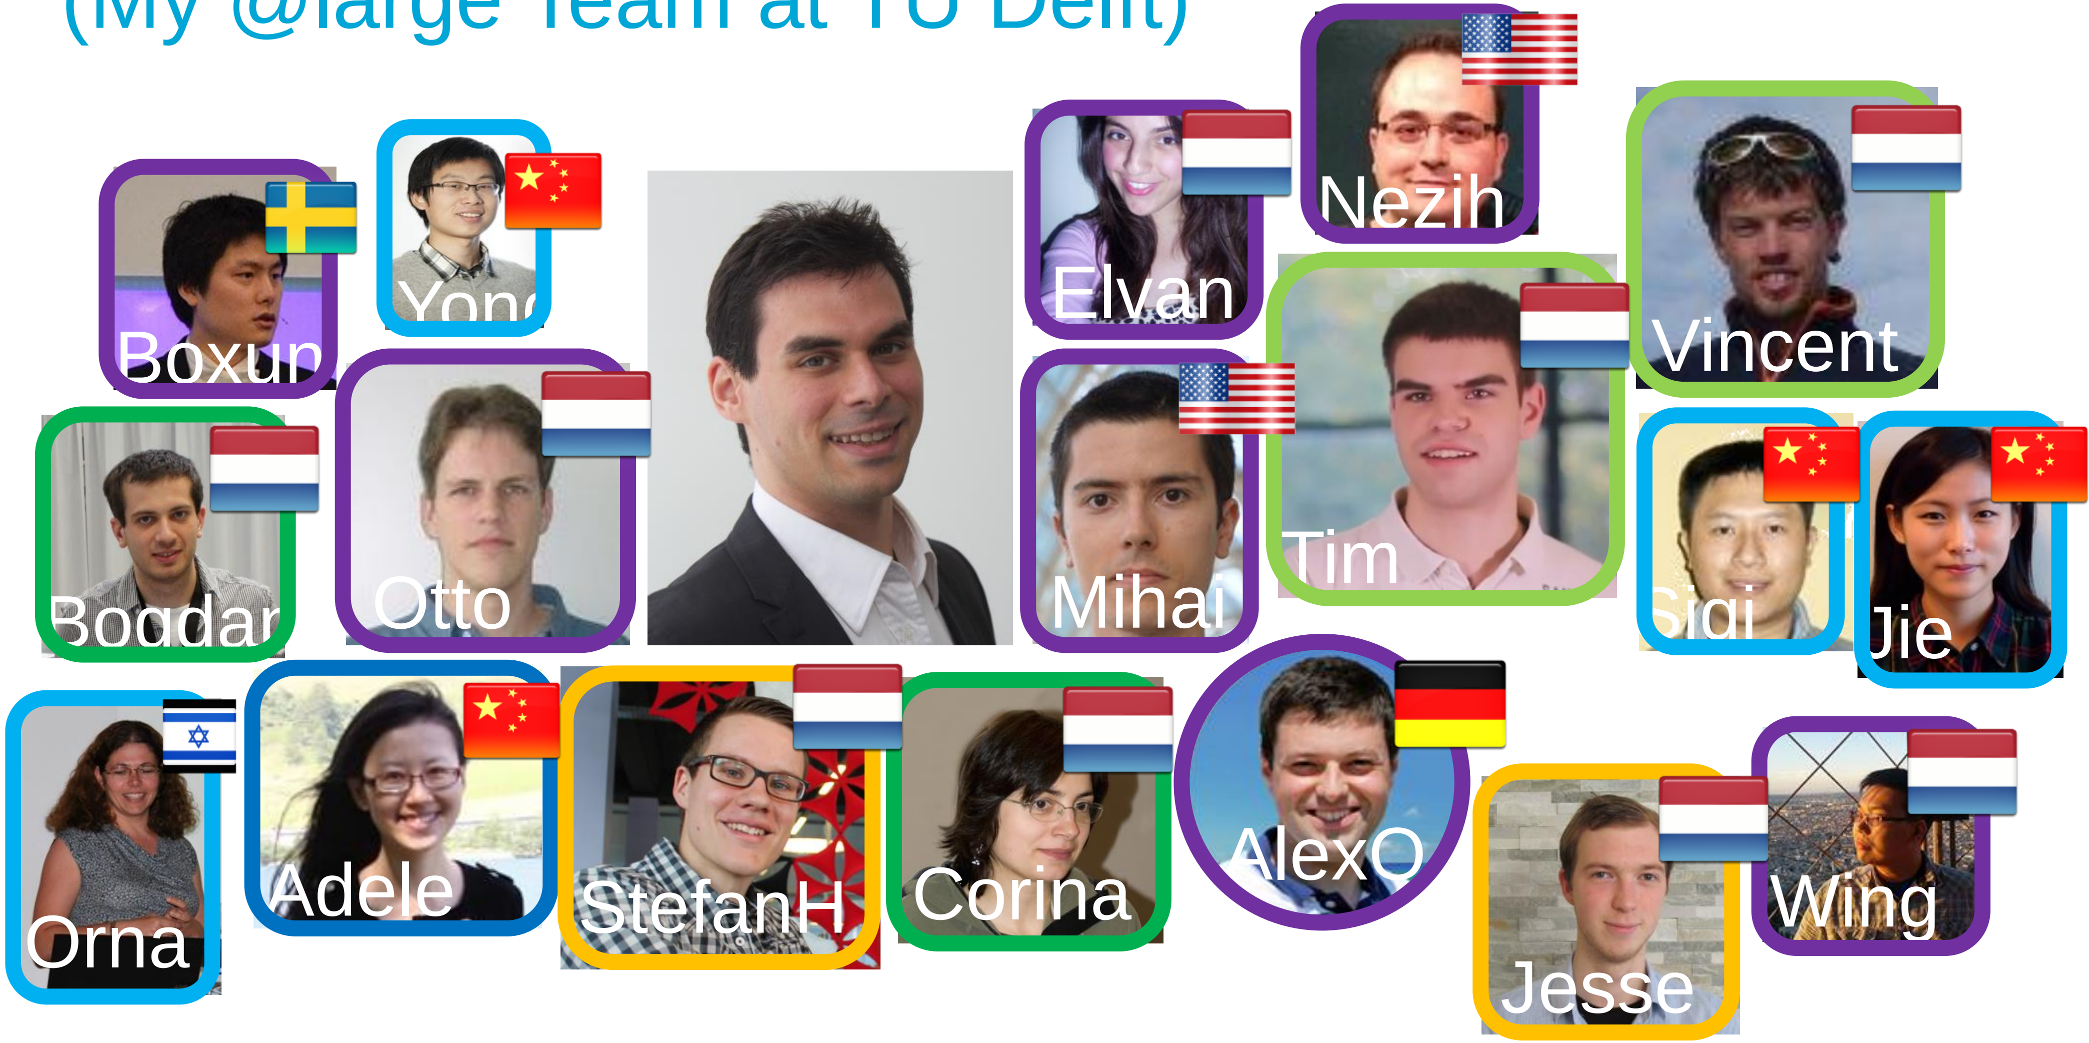

🟩 Ph.D. student 🟩 Post-doc 🟦 Assistant Prof. 🟦 AssociateProf. 🟪 Scientist 🟨 Teacher

Consider Reading the Following:

1. Iosup et al. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328 (2016)
2. Guo et al.: Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems. CCGrid 2016: 203-212
3. van Beek et al.: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters. IEEE Computer 48(7): 46-54 (2015)
4. Jia et al.: Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games. TKDD 10(2): 11 (2015)
5. Ghit et al.: V for Vicissitude: The Challenge of Scaling Complex Big Data Workflows. CCGRID 2014: 927-932
6. Guo et al.: How Well Do Graph-Processing Platforms Perform? An Empirical Performance Evaluation and Analysis. IPDPS 2014: 395-404
7. Javadi et al.: The Failure Trace Archive: Enabling the comparison of failure measurements and models of distributed systems. J. Parallel Distrib. Comput. 73(8): 1208-1223 (2013)
8. Iosup and Epema: Grid Computing Workloads. IEEE Internet Computing 15(2): 19-26 (2011)
9. Iosup et al.: On the Performance Variability of Production Cloud Services. CCGRID 2011: 104-113
10. Iosup et al.: Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. IEEE Trans. Parallel Distrib. Syst. 22(6): 931-945 (2011)