An Introduction to #CloudComputing

jesse@jdonkervliet.com

@Alosup

Jesse Donkervliet
Distributed Systems Group

dr. ir. Alexandru Iosup
Distributed Systems Group

TU Delft

DELFT DATA SCIENCE

ICT Addresses Major Societal Challenges

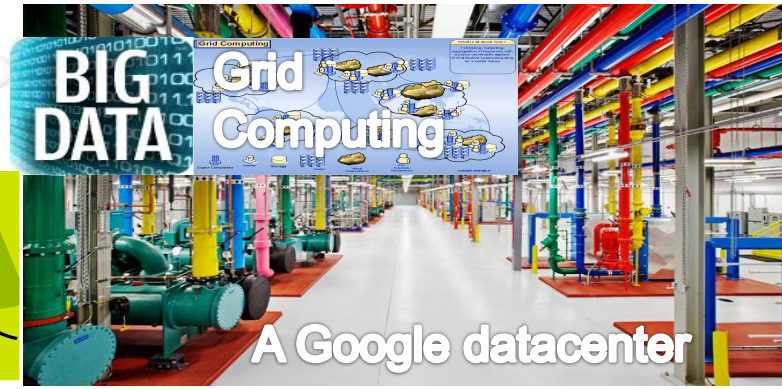

The quadruple helix: prosperous society & blooming economy & inventive academia & wise governance depend on ICT

- Enable data access & processing as a fundamental right in Europe
- Enable big science and engineering (2020: €100 bn., 1 mil. jobs)
- “To out-compute is to out-compete”, but with energy footprint <5%
- Keep Internet-services affordable yet high quality in Europe
- The Schiphol of computation: Netherlands as a world-wide ICT hub

Societal Challenges, A Concrete Example: Graph Processing for Everyone

LinkedIn

Oracle 1.2M followers,
132k employees

company/day:
40-60 posts, 500-700 comments

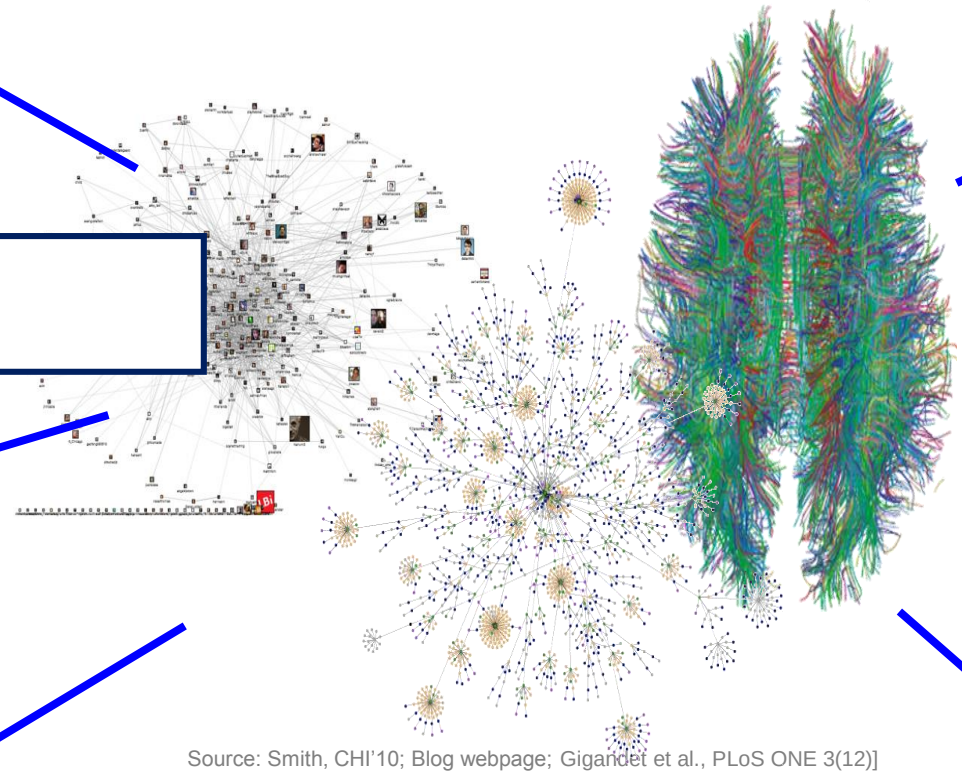

Source: Smith, CHI'10; Blog webpage; Gigandet et al., PLoS ONE 3(12)]

270M MAU
200+ avg followers

>54B edges

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

YAHOO!

friendster
(SME*)

XFIRE
(SME*)

* SMEs in EU/NL= 60% gross value added, little to no ICT expertise

3

LinkedIn: Job-Seeker & Job-Creator Match

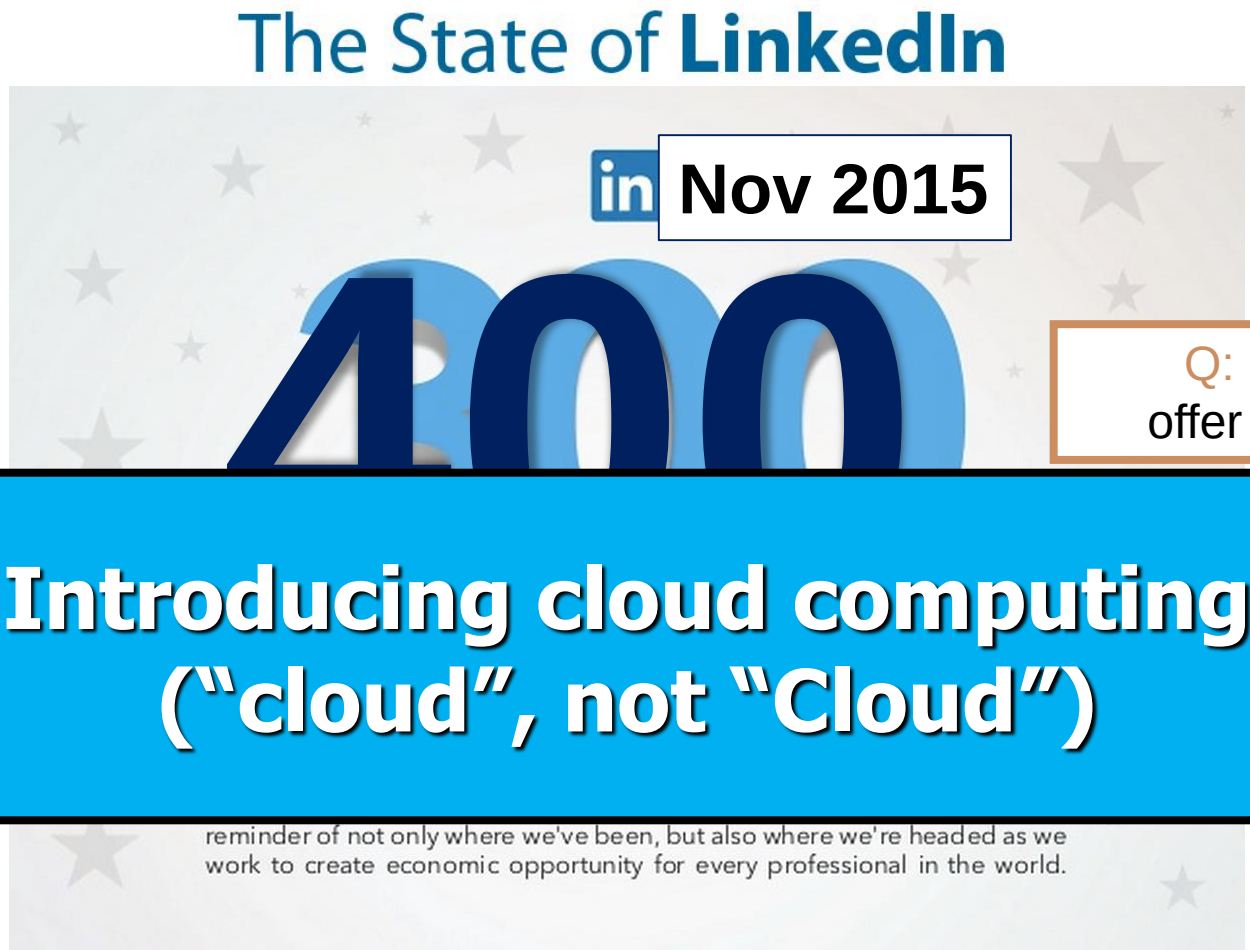

Q: How can LinkedIn offer this service to you?

Introducing cloud computing ("cloud", not "Cloud")

Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

What is Cloud Computing?
1. A Cloudy Buzzword

- 18 definitions in computer science (ECIS'10).
  NIST has one. Cal has one. We have one.
- "We have redefined cloud computing to include everything that we already do." Larry Ellison, Oracle, 2009

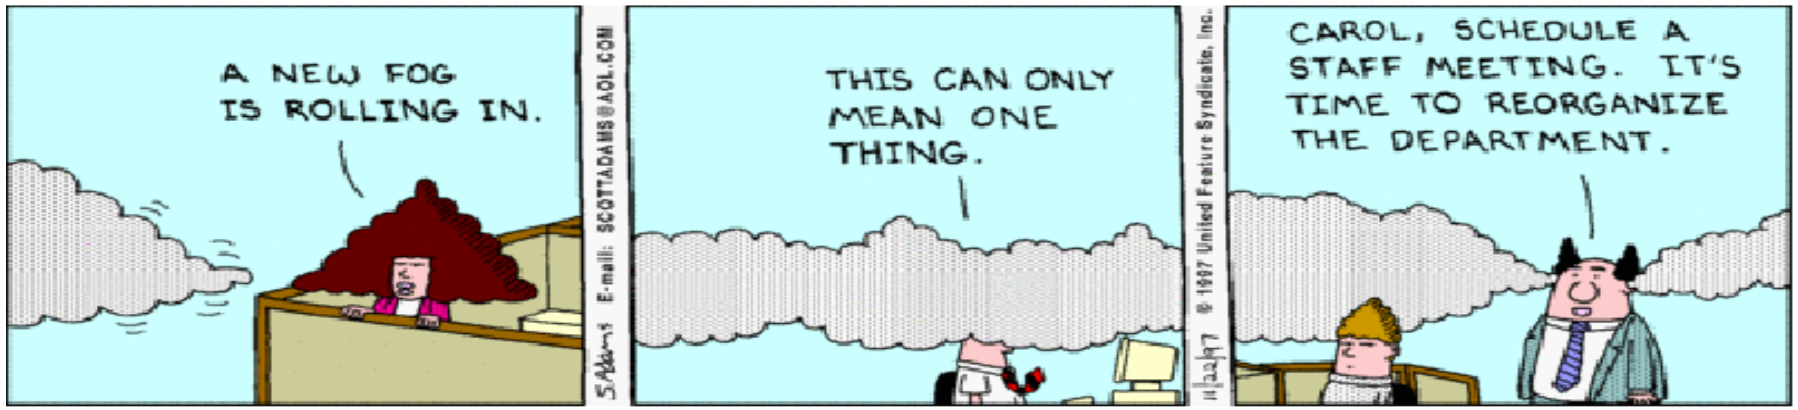

Source: http://dilbert.com/strips/comic/1997-11-22/

What is Cloud Computing?
2. A Descendant* of the Grid Idea

Source: http://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities [+ for] nontrivial QoS." I. Foster, 1998 + 1999

* Subset.

Cloud ~~Grid~~ Applications

~~Cloud~~/Grid MW Stack

- ~~Cloud~~ Grid Very High Level MW
- Cloud ~~Grid~~ High Level MW
- Cloud ~~Grid~~ Low Level MW

Virtualized HW + OS

MW = Middleware

TU Delft

Have you noticed "QoS"? What is that?

What is Cloud Computing?
3. A Useful IT Service

"Use only when you want! Pay only for what you use!"

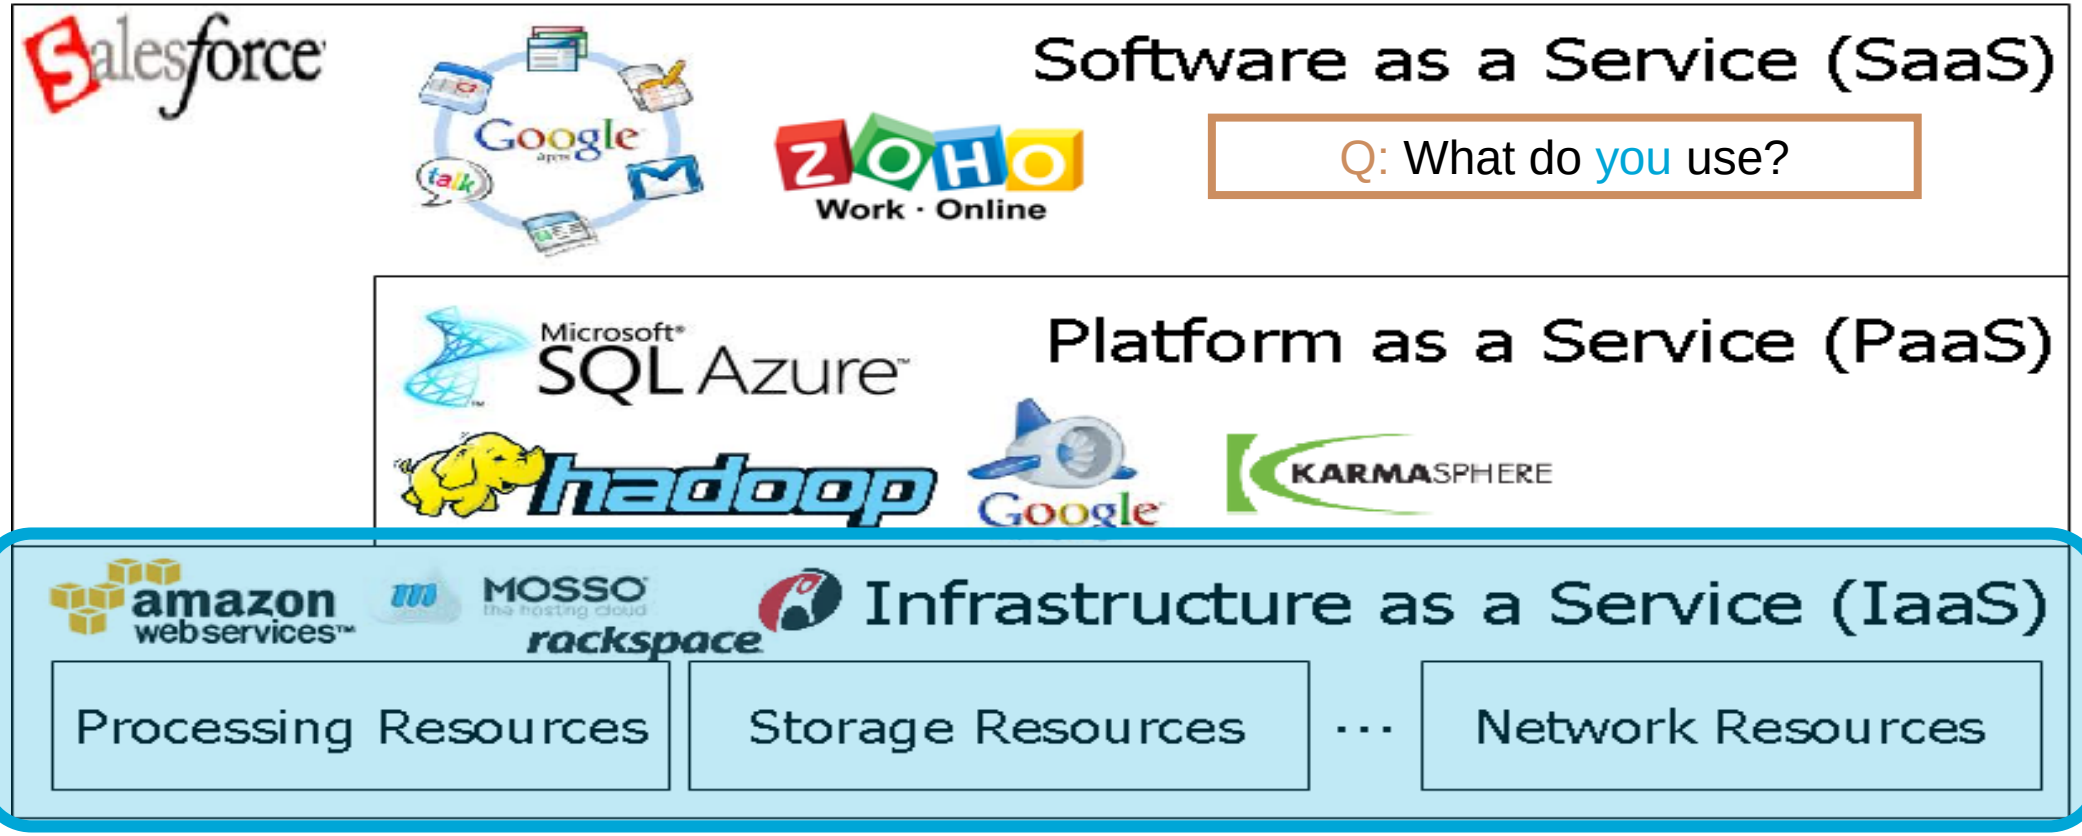

Main Characteristics of IaaS Clouds

1. On-Demand Pay-per-Use
2. Elasticity (cloud concept of Scalability)
3. Resource Pooling
4. Fully automated IT services
5. Quality of Service

Q: Sounds great, but …
How can we make
all this stuff happen?

Introducing datacenters & datacenter-based clouds = ICT service creation for everyone

Intro to Cloud Computing

Here or @home

Interactive

TU: 1 — Pitch on Datacenter-Based Cloud Computing

1 — The Golden Age of Datacenters

4 — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

7 — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

— Reality Check

This Is the Golden Age of Datacenters

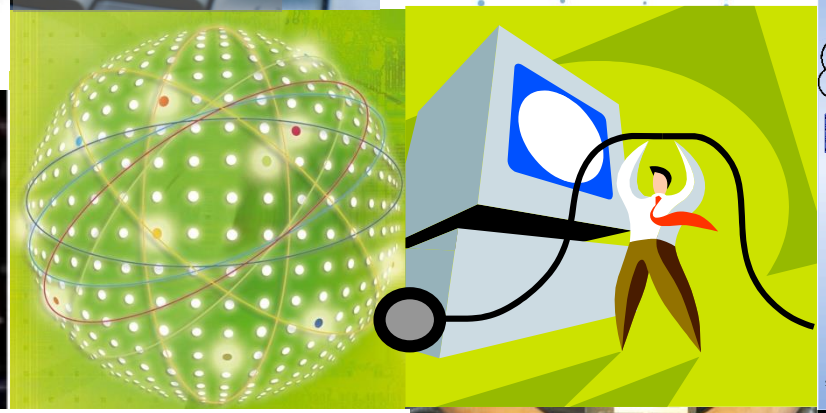

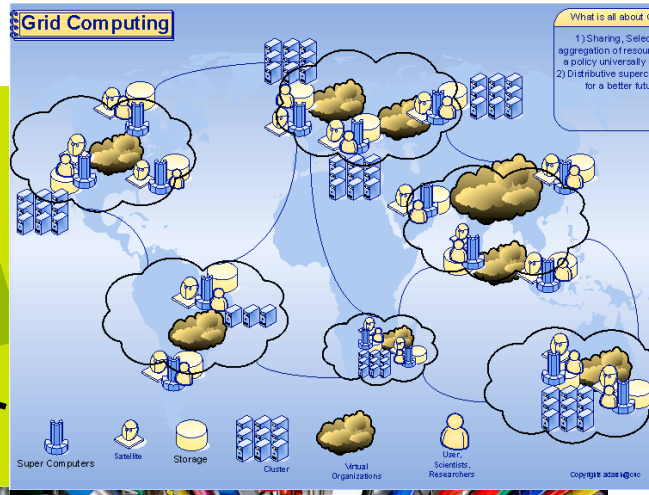

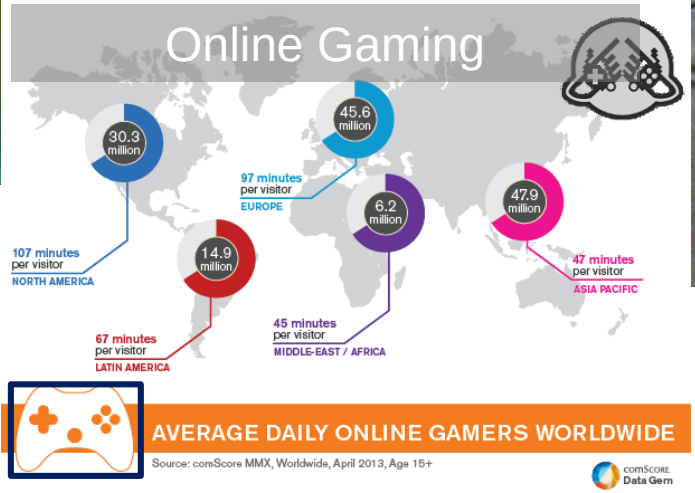

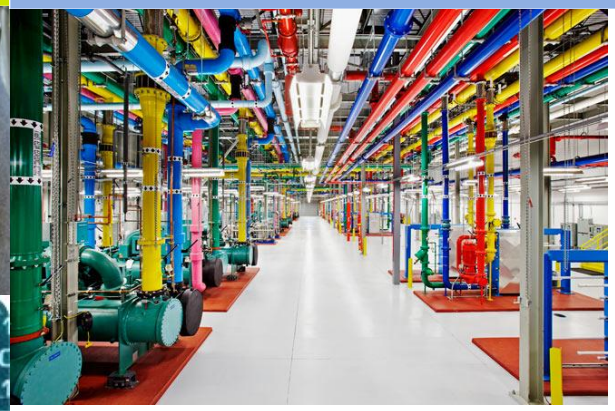

Factories Powering the Goods Economy and Better Living Standards

Creator

Factory (1980)

Consumer

TU Delft

Datacenters = Digital Factories Powering the Digital Economy and Better Living Standards

Creator

Datacenter (2010)

Consumer

We want more of this

"15% ICT market is simple DC services"
"Already 60+ bn.€/year"
"Consume more energy than the transport industry"

We want less of this (and other issues raised by DCs)

TU Delft

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

Joe Has an Idea ($$$)

Solution #1

Buy then Maintain

- Big up-front commitment
- Load variability: NOT supported

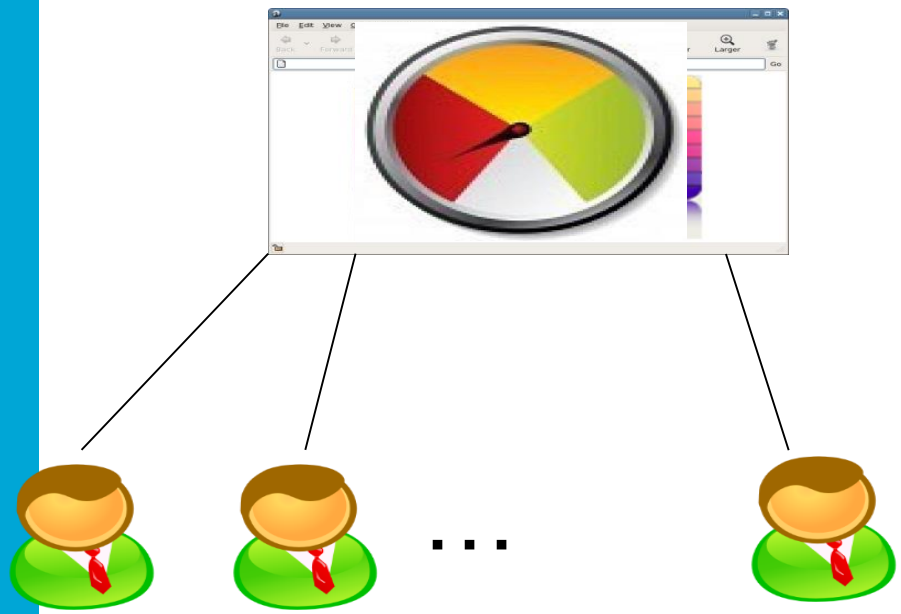

Solution #2

Deploy on IaaS Cloud

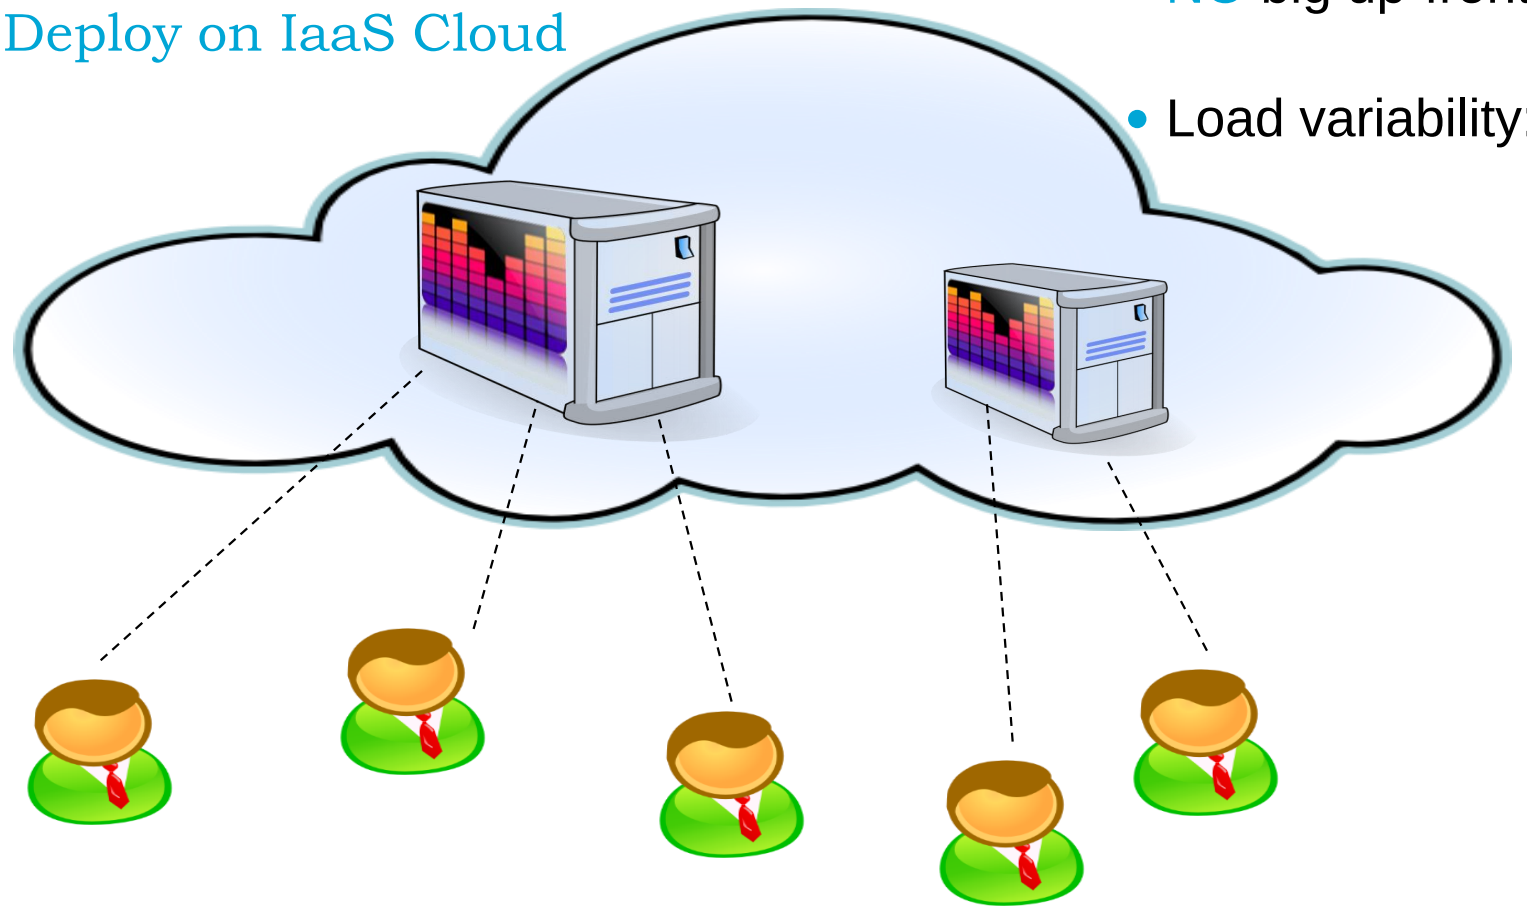

- NO big up-front commitment
- Load variability: supported

Inside a Cloud Datacenter: Infrastructure as a Service

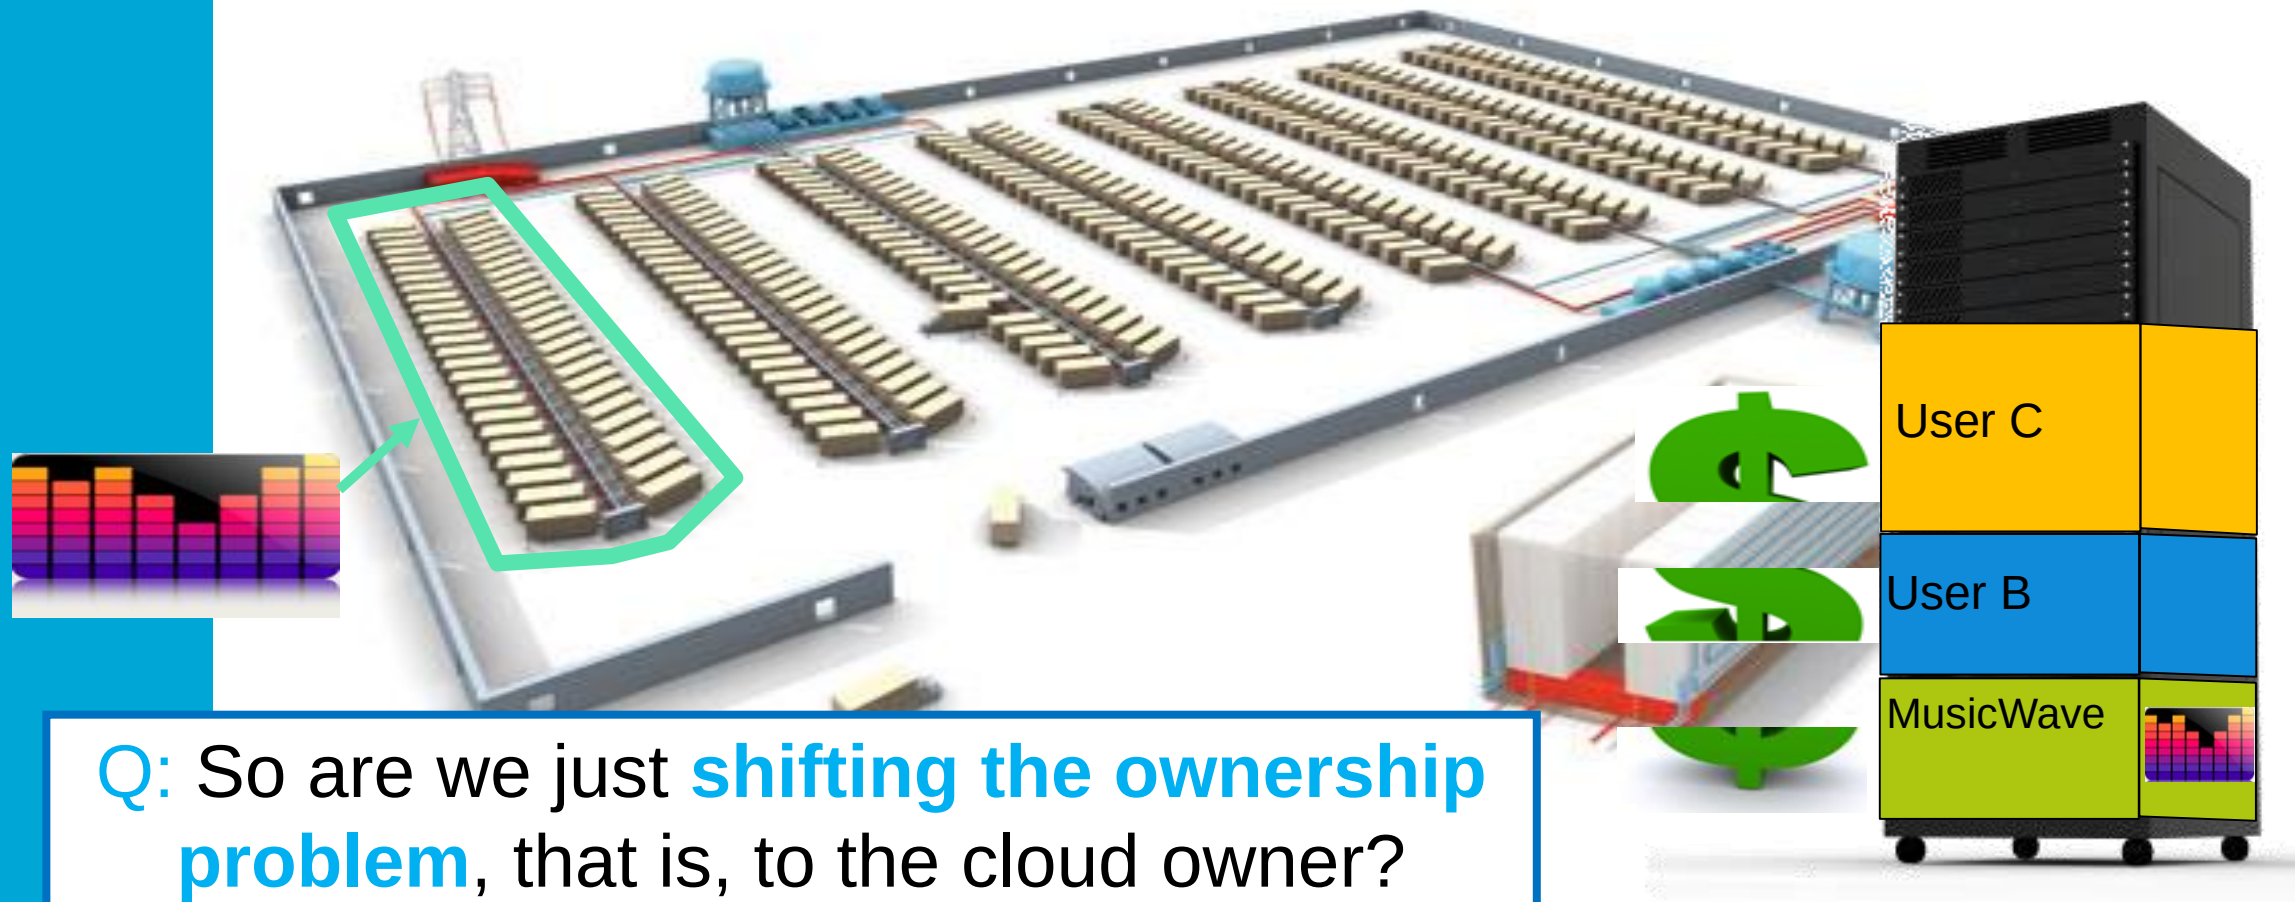

Q: So are we just **shifting the ownership problem**, that is, to the cloud owner?

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

TU Delft

The Cloud Owner Perspective

- Build the datacenter
- Operate the datacenter

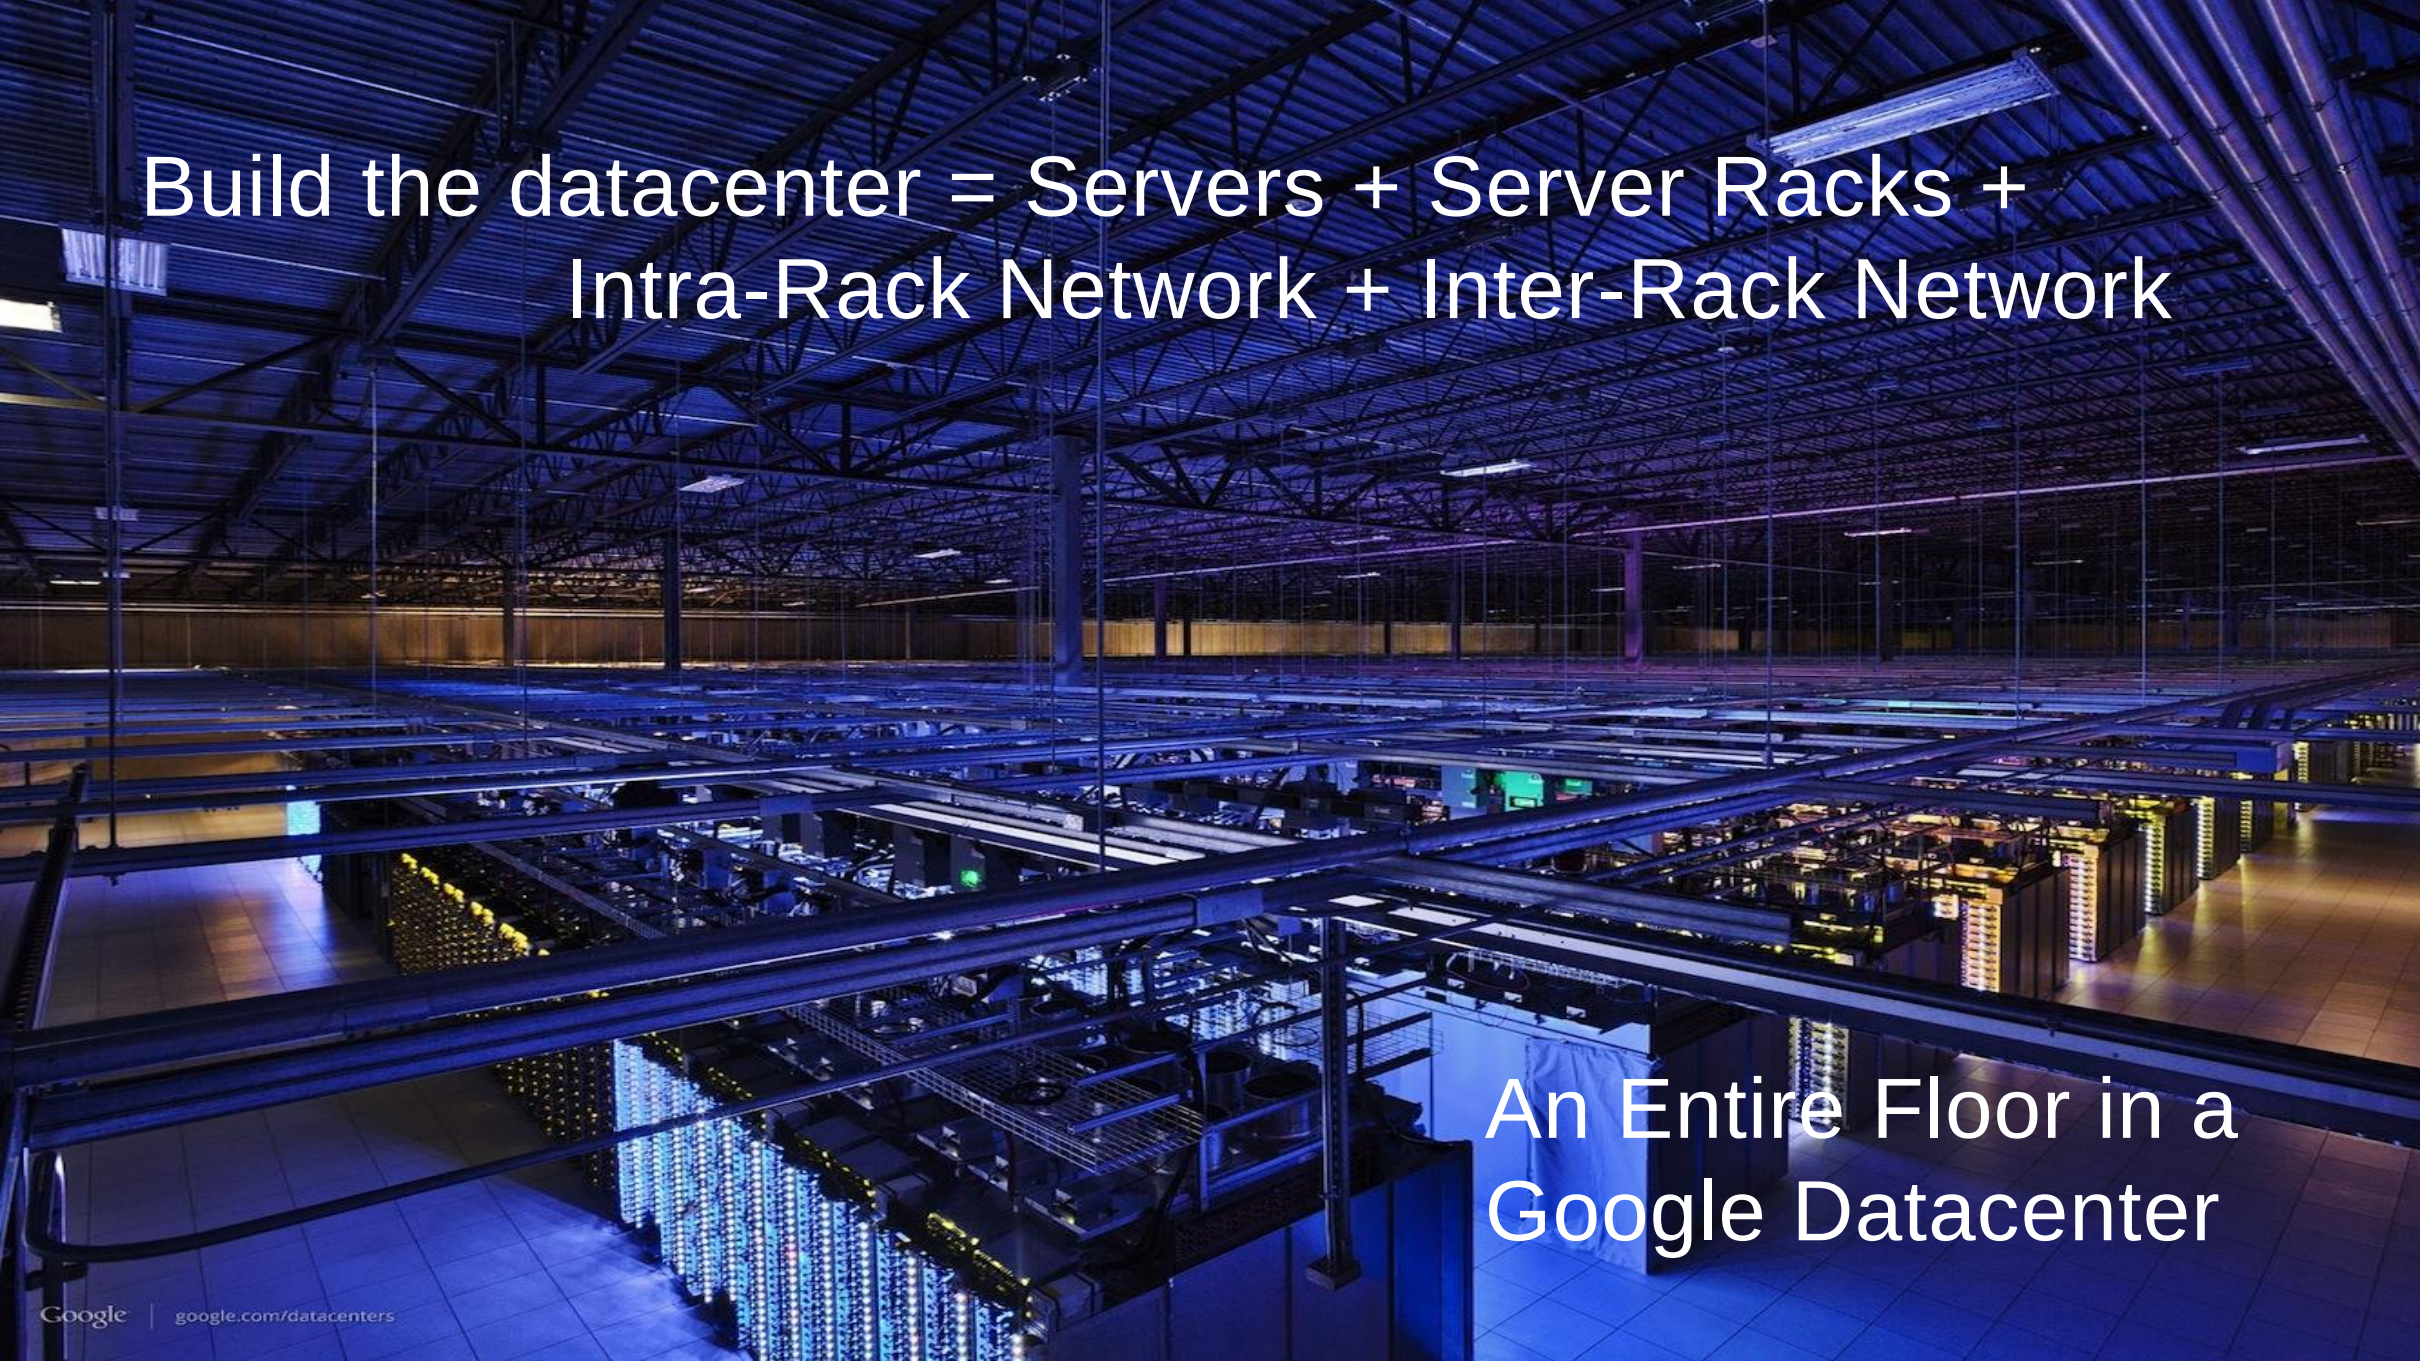
Build the datacenter = Servers + Server Racks +
Intra-Rack Network + Inter-Rack Network
An Entire Floor in a
Google Datacenter
Google | google.com/datacenters

Datacenter = commodity high-performance ICT

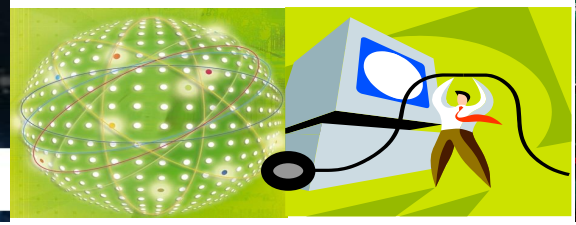

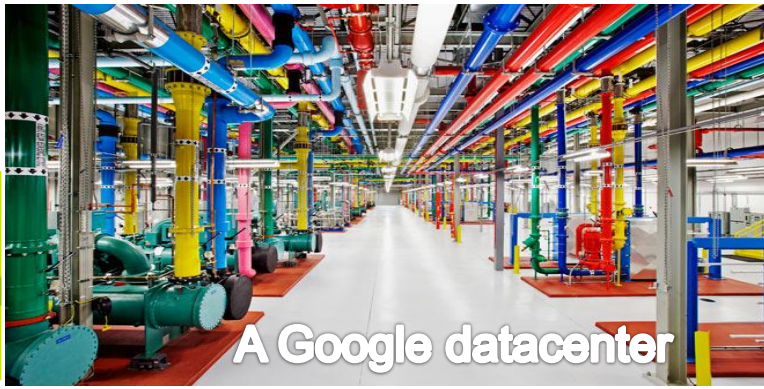

Large-scale infrastructure

- High-tech automated software to manage
- Inter-connected computer clusters
- High-end computation, storage, network
- Large memory capacity

"my other computer is a datacenter"

The Pizza-Box Stack

- The 1U server

The Pizza-Box Stack

- The 1U server

The Pizza-Box Stack

- The 1U server

The Pizza-Box Stack

- The 1U server
- The 19” server rack (42U is now standard)
  - Half-racks also common

Q: What is a **half-rack**, and why is it useful?

The Data Center Network

- Network bandwidth per rack
  - 1 x 48-port GigE switch = 40 UP-, 8 DOWN-links

Q: What is are the characteristics of **research**, **enterprise**, and **consumer** hardware (and other services)?

- Network bandwidth per socket
  - (consumer) 100 Mbps for 1 GigE rack switch
  - (enterprise) 1 Gbps for 10 GigE rack switch
  - (enterprise) 10 GBps for ncHT3 (supercomputing class)
  - (research) 1 Tbps!~100 GBps optical (not yet production-ready)

Source: Dennis Abts (Google, Inc.) and John Kim (KAIST), High Performance Data Center Networks, 2011.

Workload

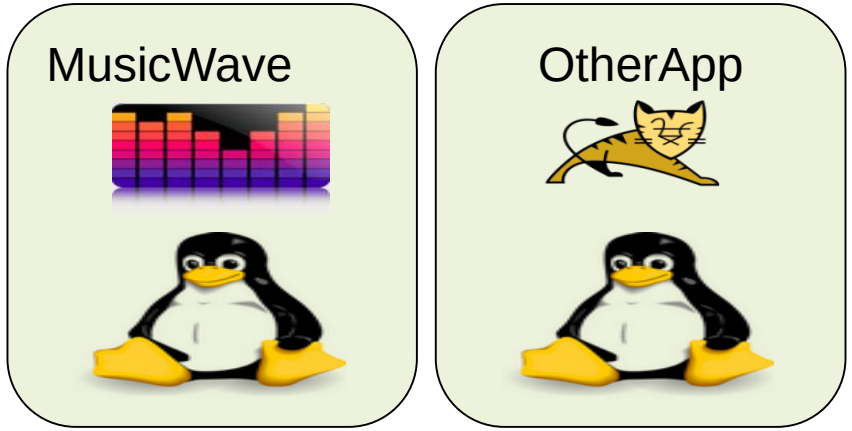
MusicWave
OtherApp

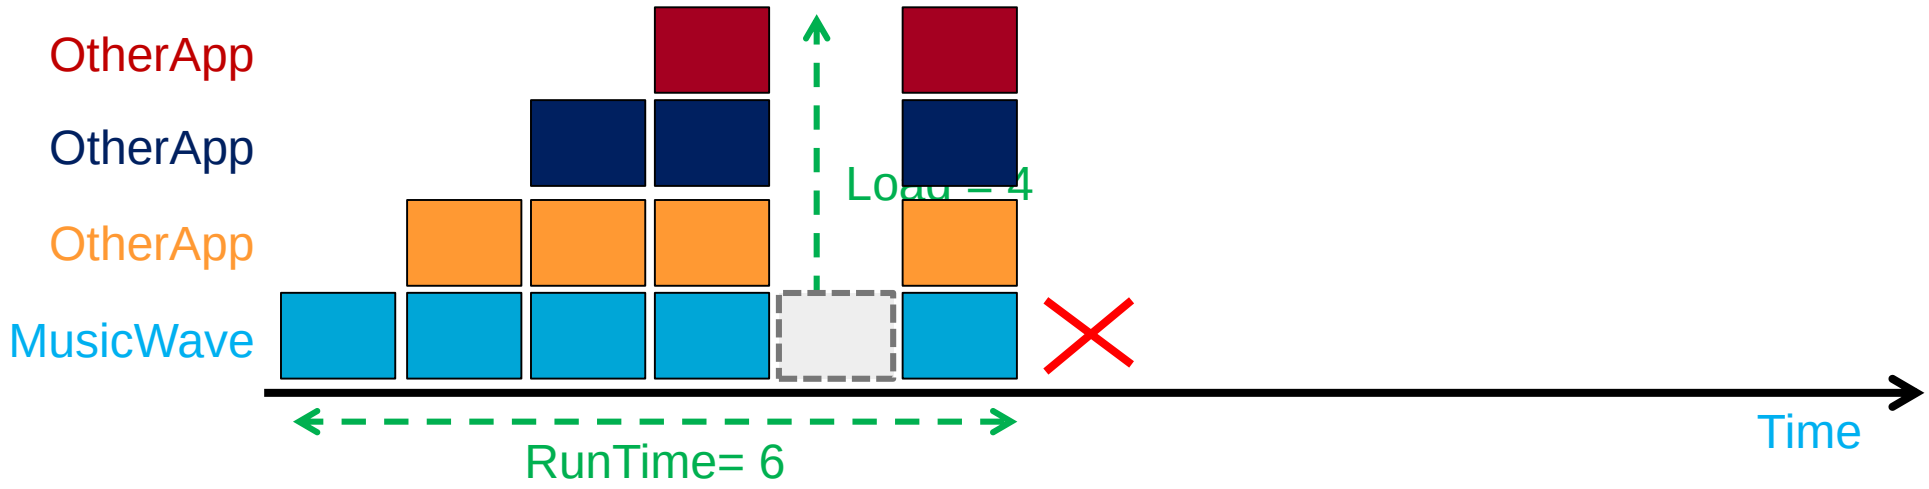
OtherApp
OtherApp
OtherApp
MusicWave
Load = 4
RunTime= 6
Time

Current Technology: What Is A Scheduler?

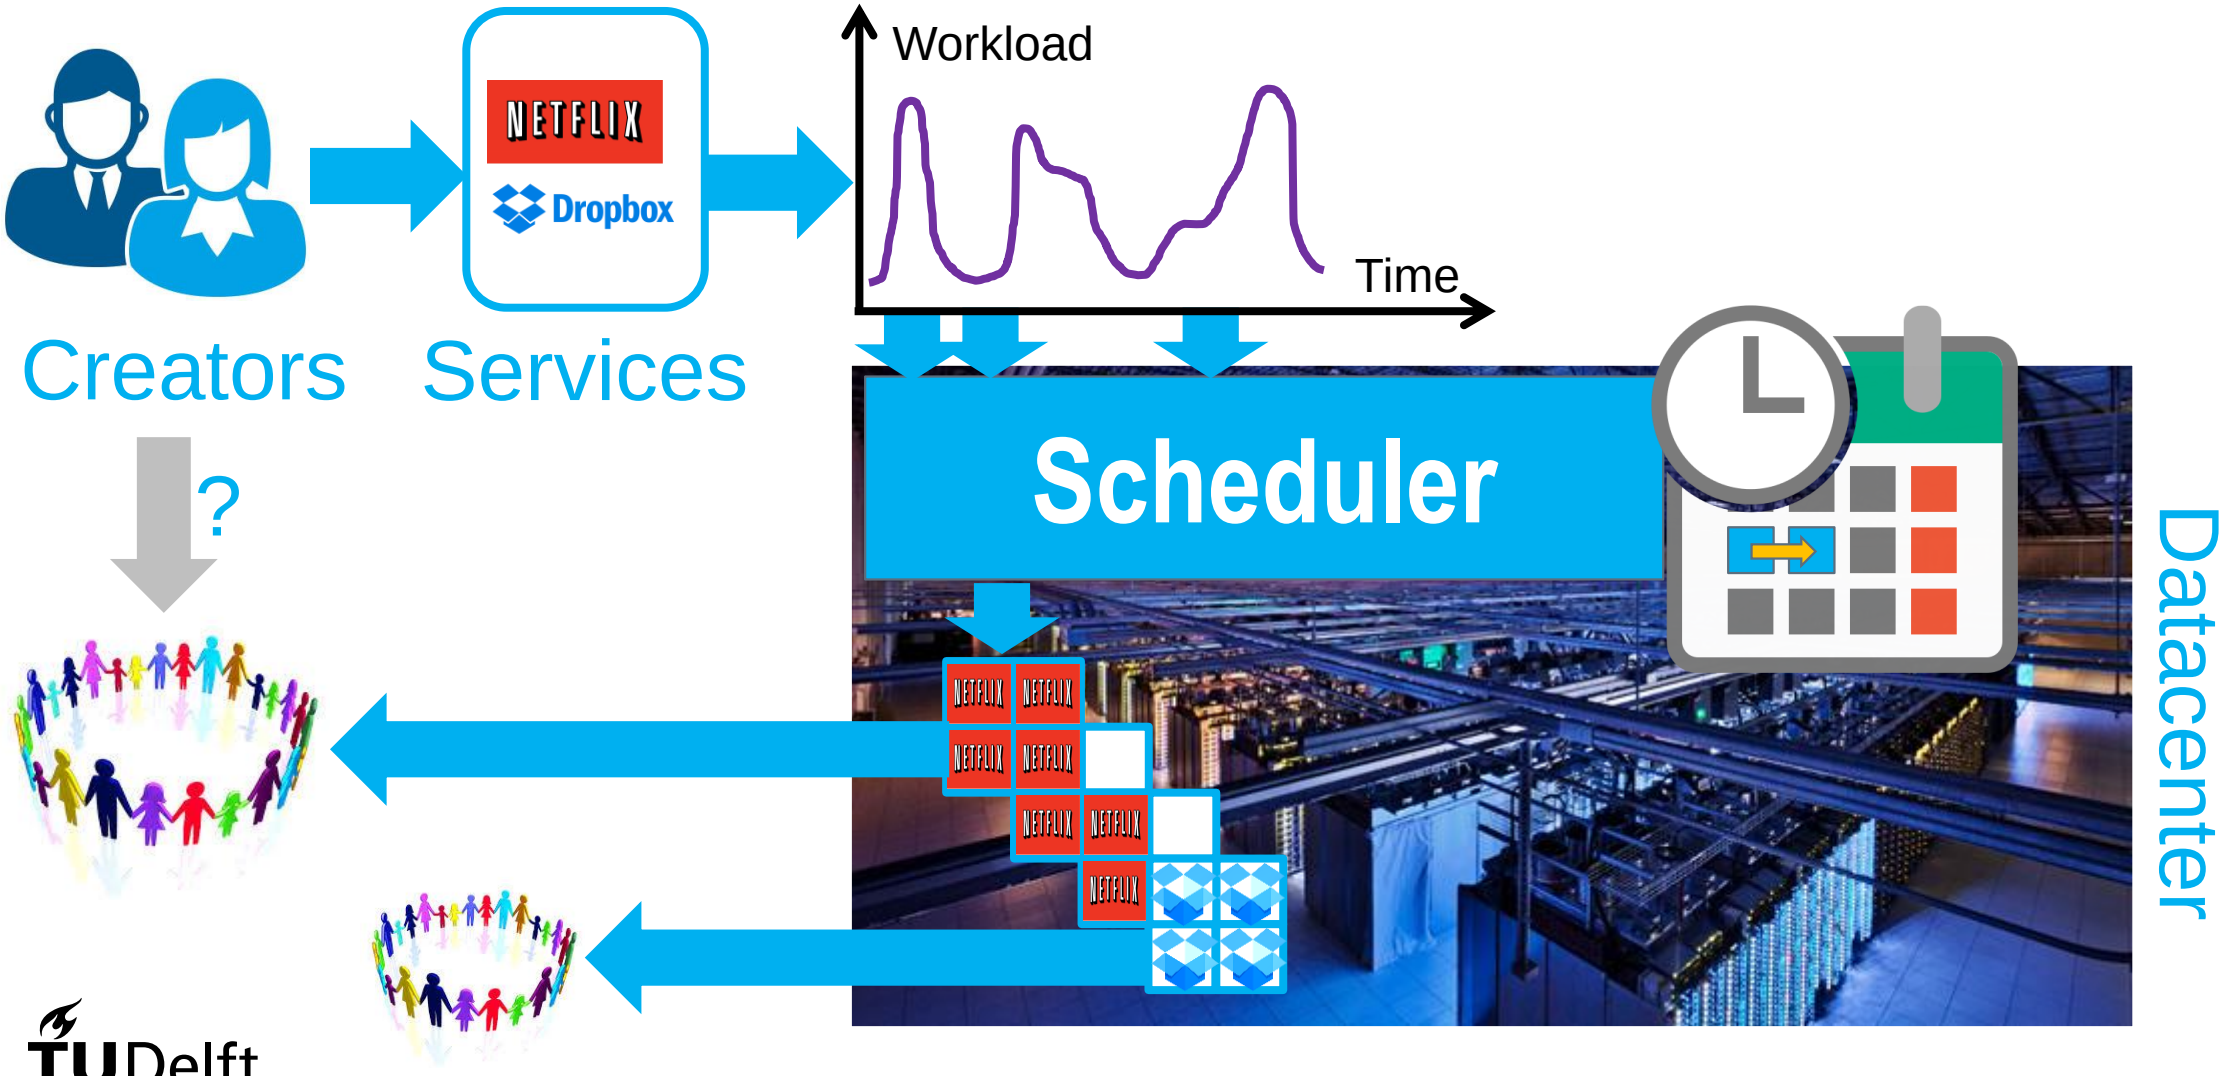

Resource Sharing Models

Grids
Space-Sharing

IaaS Clouds
Time-Sharing

Q: Which one is better?

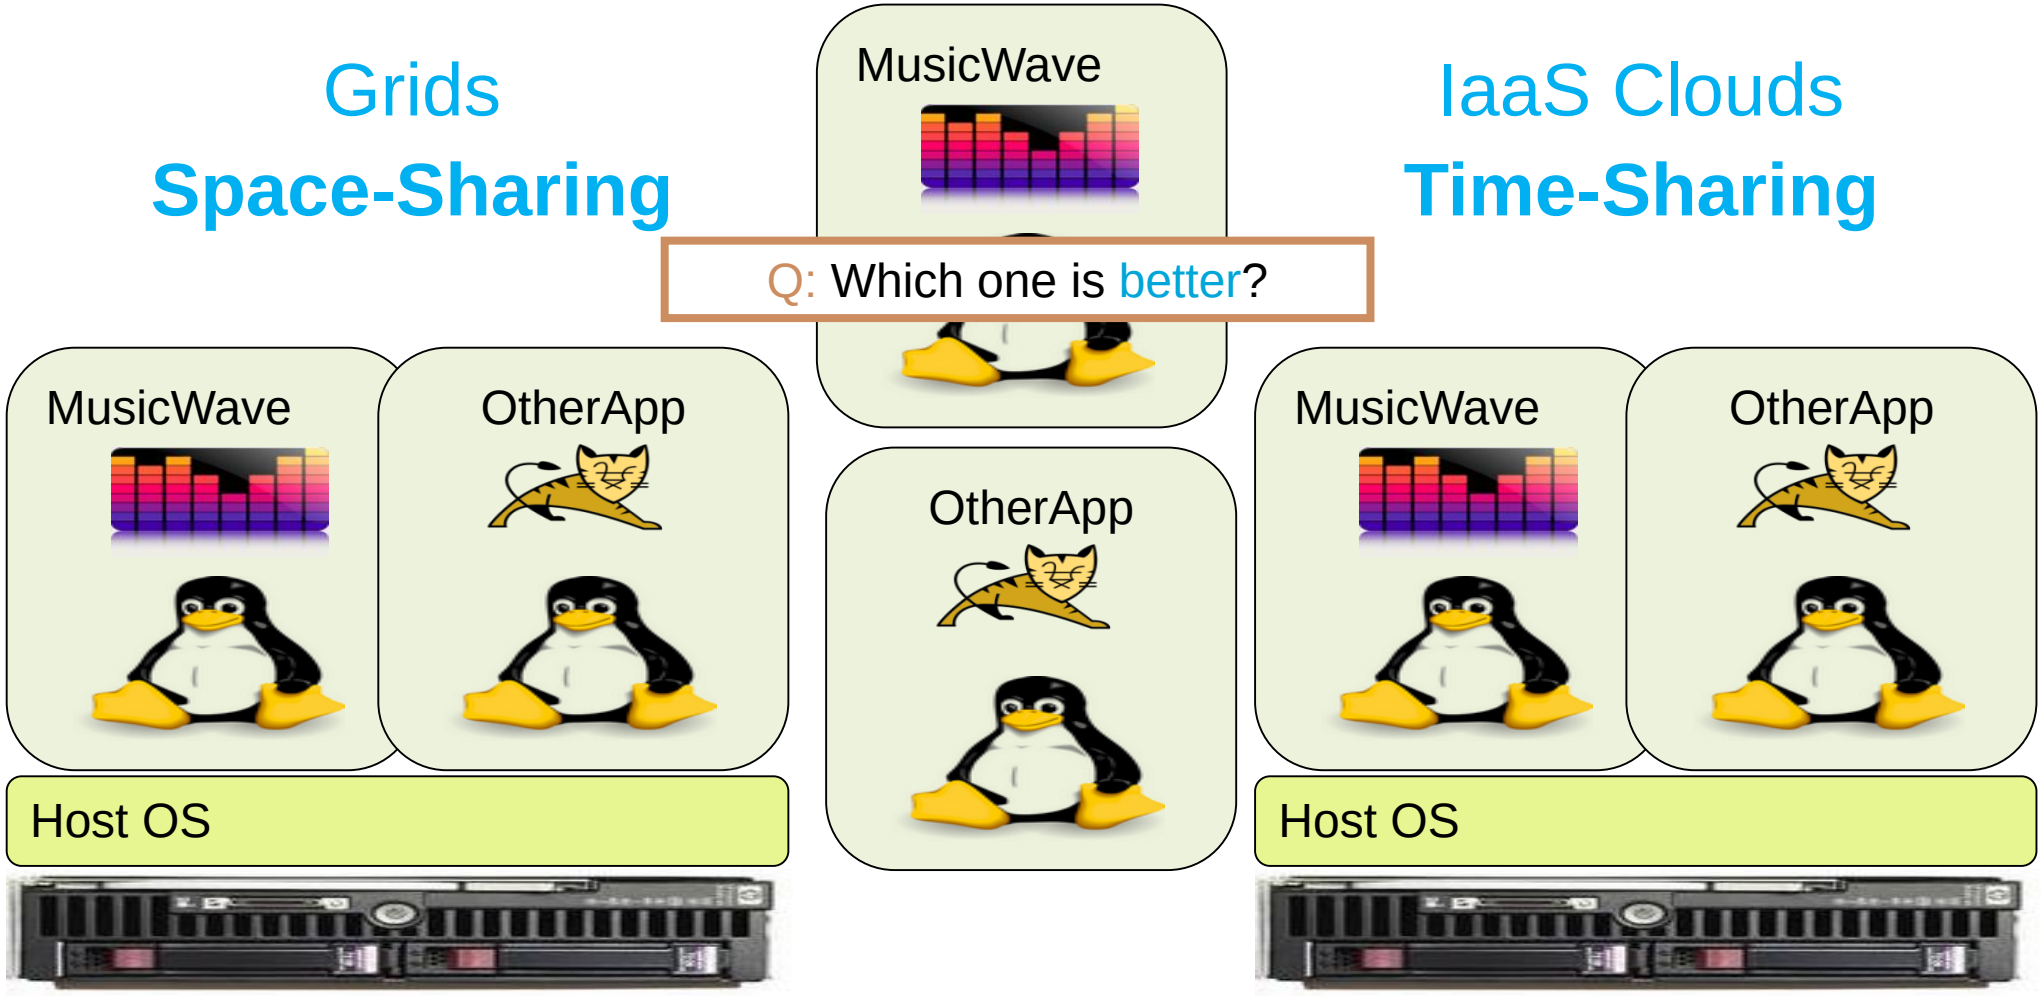

Virtualization

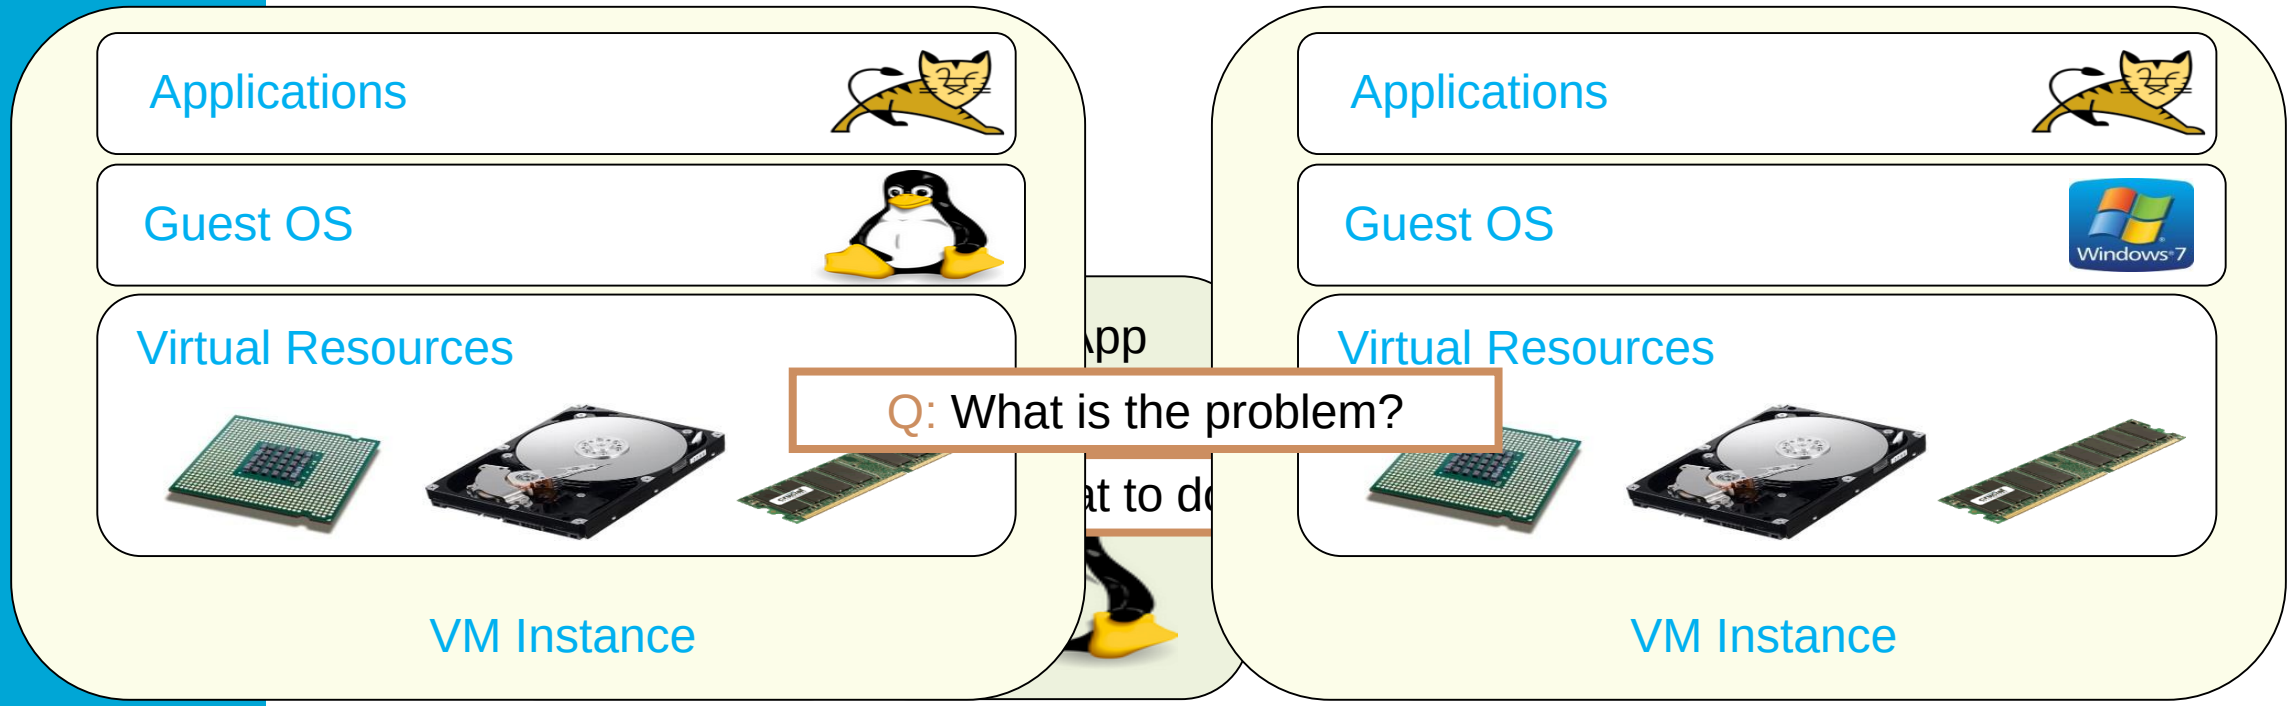

Virtualization

Host OS

(Source: A. Iosup, Course IN4392, TU Delft, 2014)

The Cloud User Perspective

- Lease the resources
- Use the resources

Where to Lease Resources From?
IaaS Cloud Deployment Models

Private
On-premises

Public
Off-premises

Hybrid
Two or more clouds

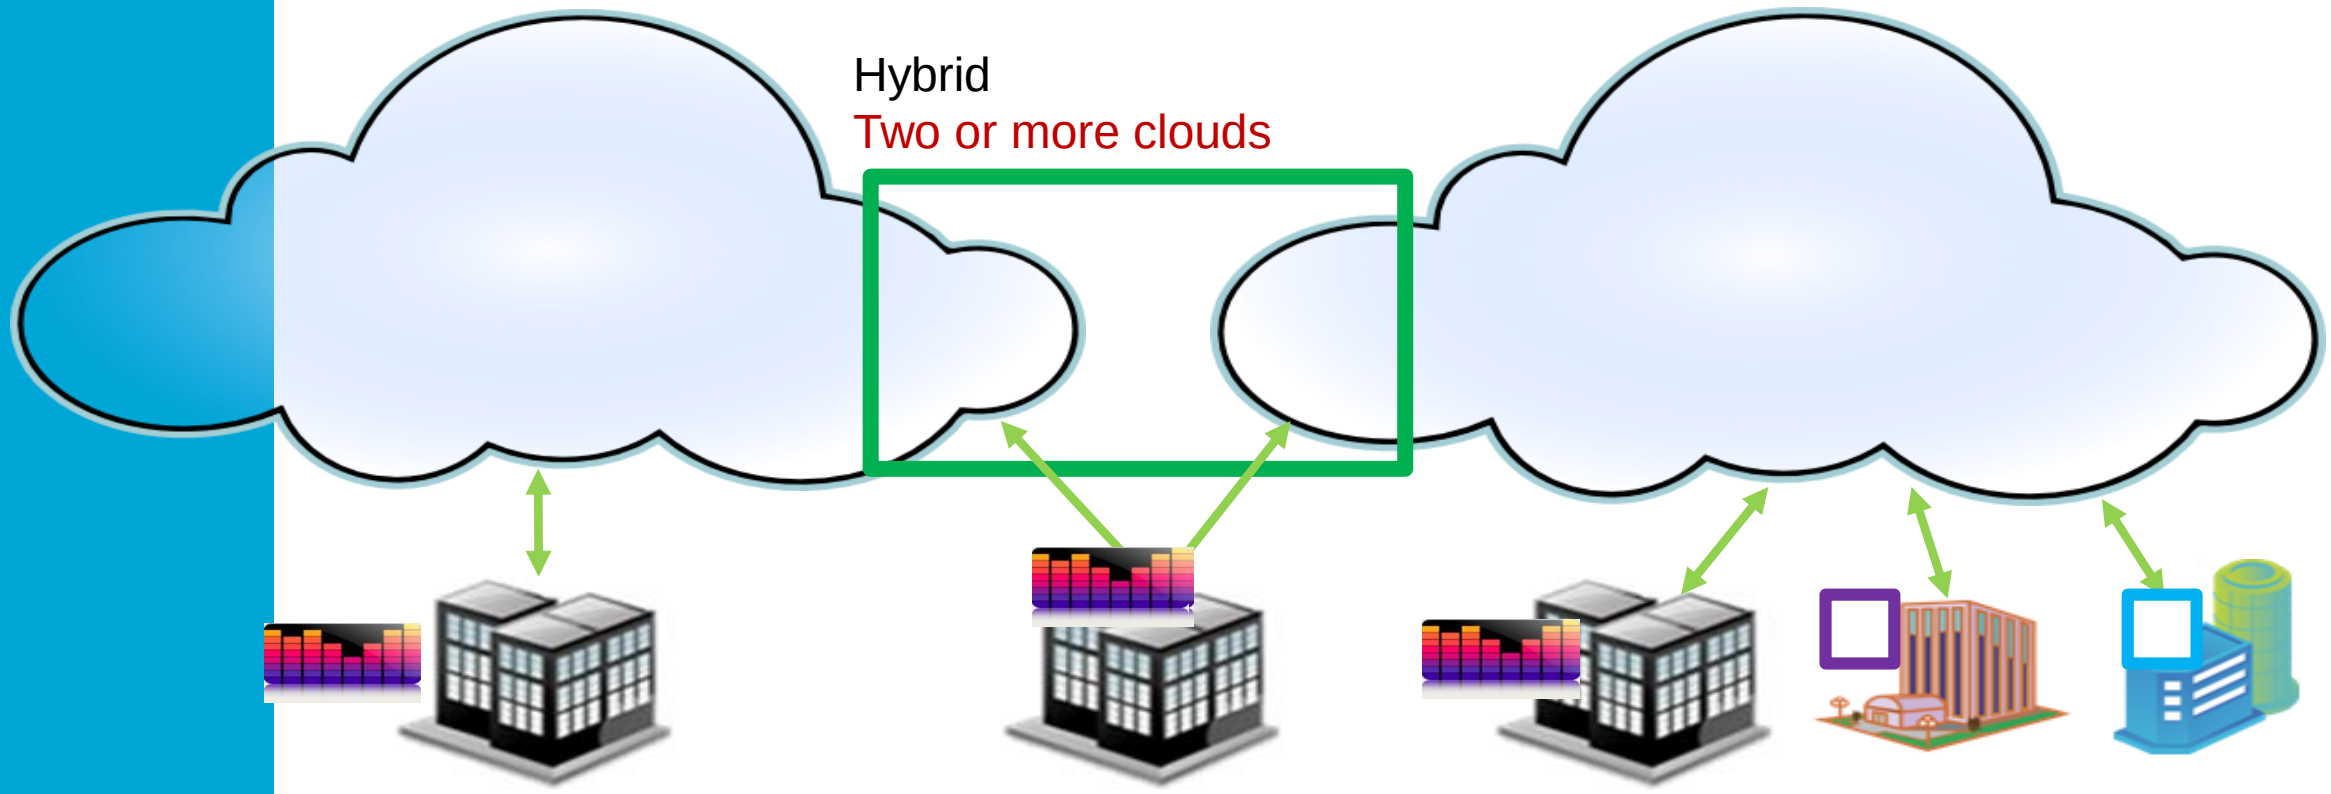

(Source: A. Antoniou, MSc Defense, TU Delft, 2012. Original idea: Mell and Grance, NIST Spec.Pub. 800-145, Sep 2011.)

Use Case:
Amazon Elastic Compute Cloud (EC2)

- Prominent IaaS provider

- Datacenters all over the
- Many VM instance types

Instance	Capacity	US$/hour
m1.small		0.10
m1.large	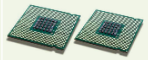	0.38
c1.xlarge	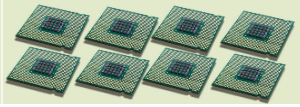	0.76

- Per-hour **charging**
- **Auto-scaling with simple policies**

What Does It Mean To Lease A Resource?
The Virtual Machine Lifecycle

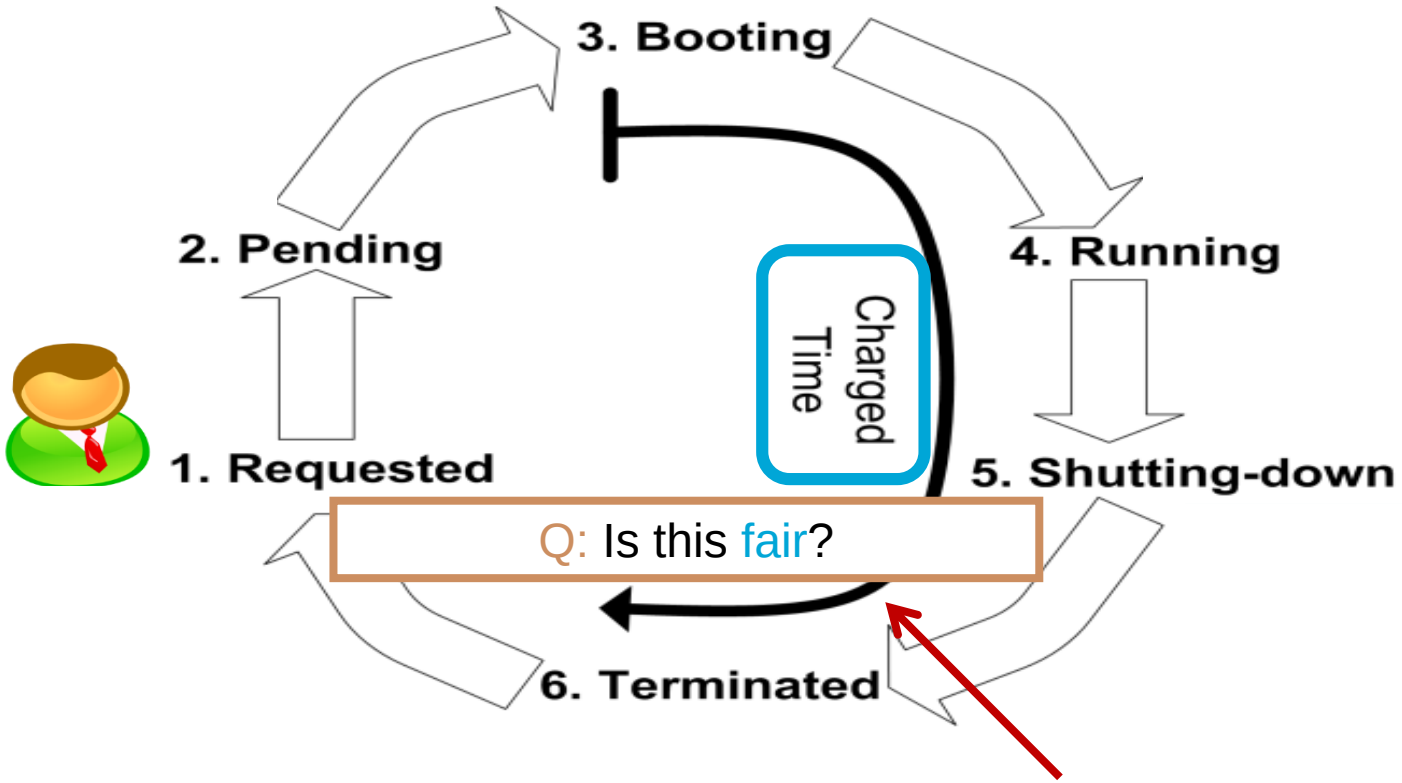

At this scale, computer science is also about ethics, law, society, etc.

(Source: A. Antoniou, MSc Defense, TU Delft, 2012.)

Provisioning and Allocation of Resources

Q: How do service creators get resources?

Provisioning

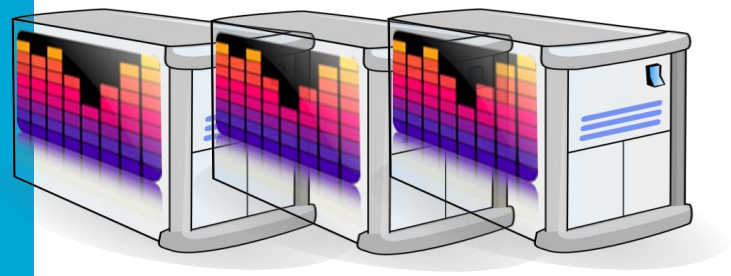

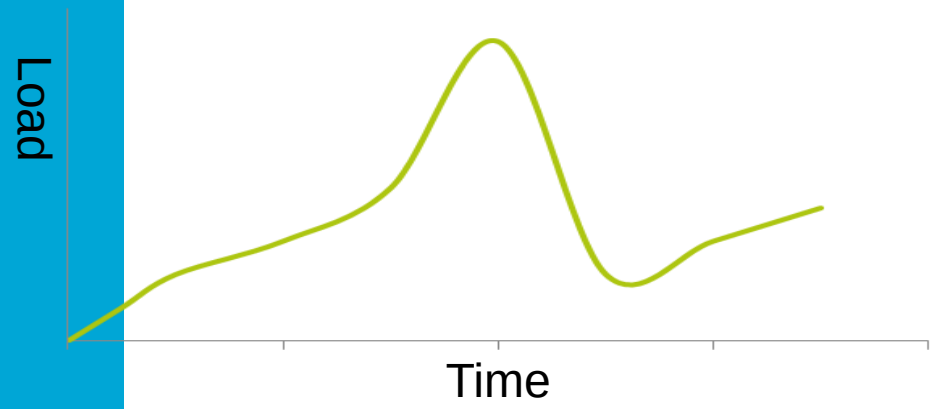

Allocation

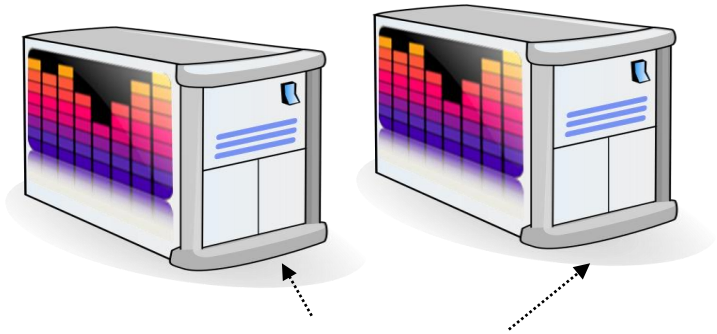

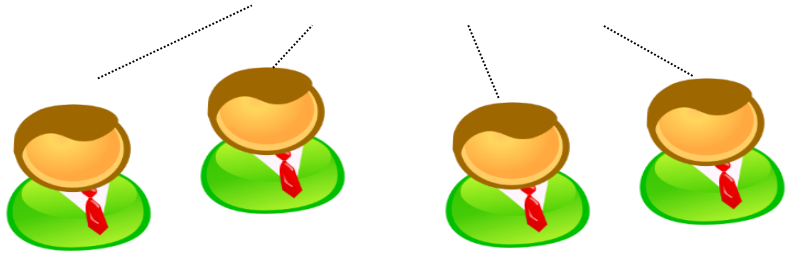

Provisioning and Allocation *Policies*

Q: How many policies exist?

Q: How to select a policy?

Provisioning

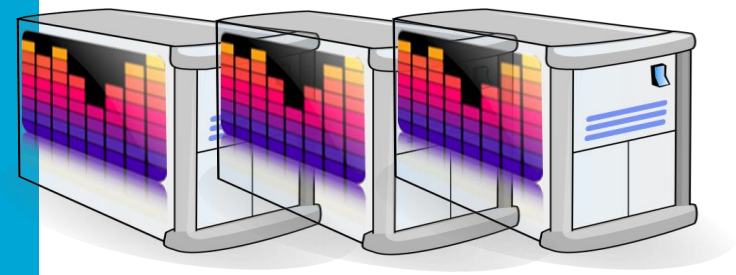

When?

From where?

How many?

Which type?

etc.

Allocation

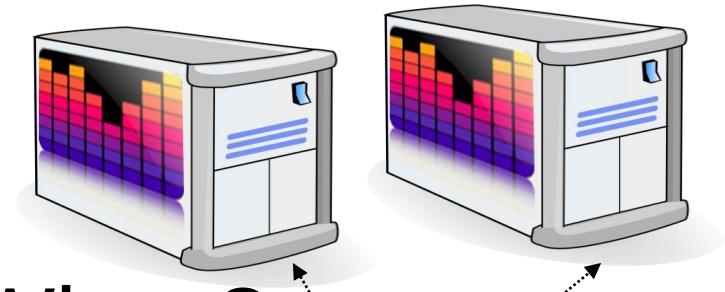

When?

Where?

etc.

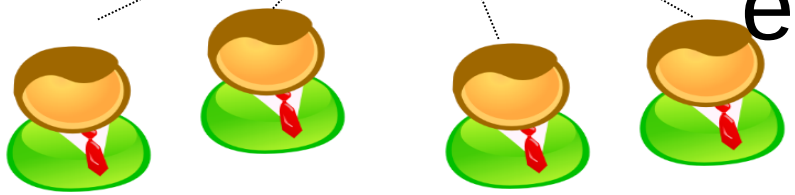

(Source: A. Antoniou, MSc Defense, TU Delft, 2012.)

Auto-Scalers = automatically provision resources, on-demand

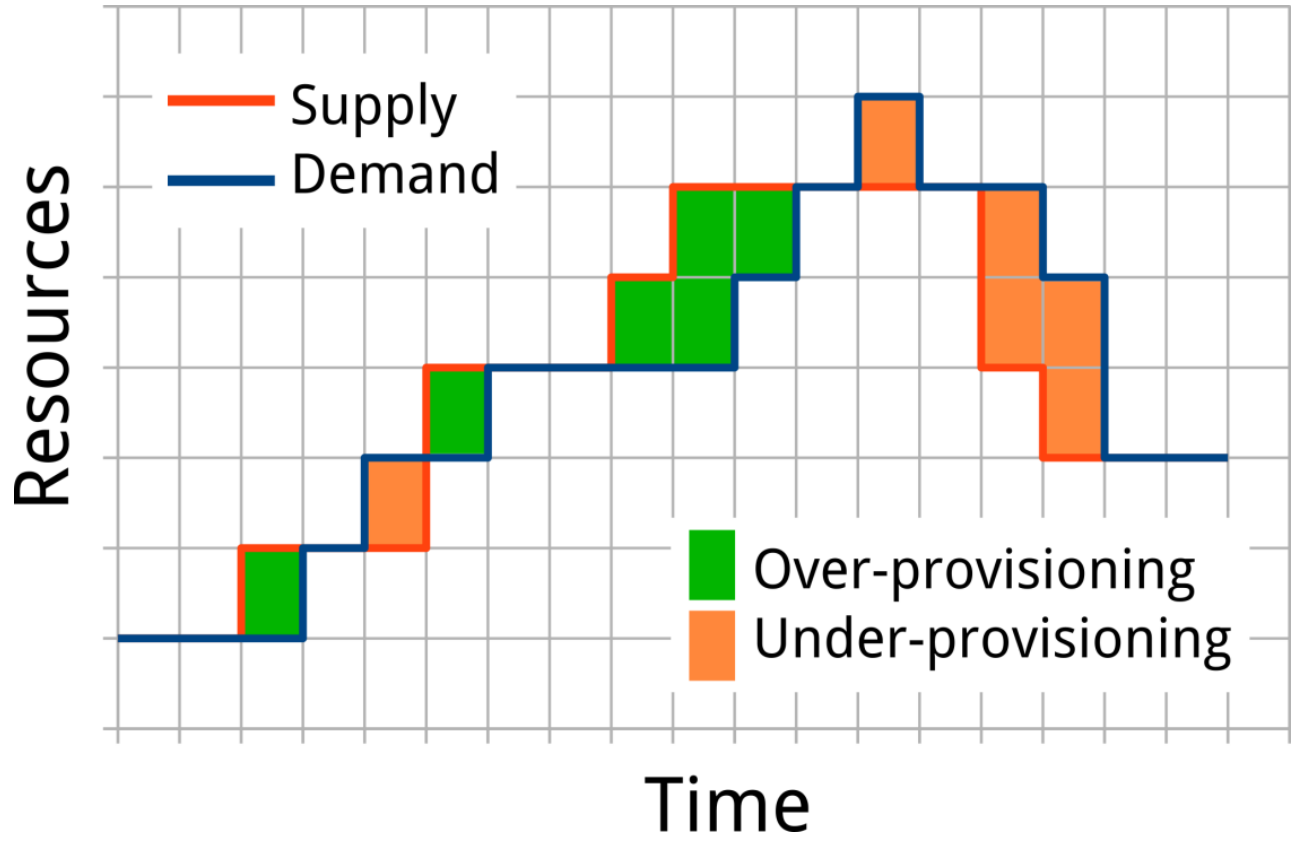

(Source: A. Ilyushkin, N. Herbst, A. Ali El Din, SPEC RG Cloud Group, 2016.)

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

Scientific Challenges to Get This Done

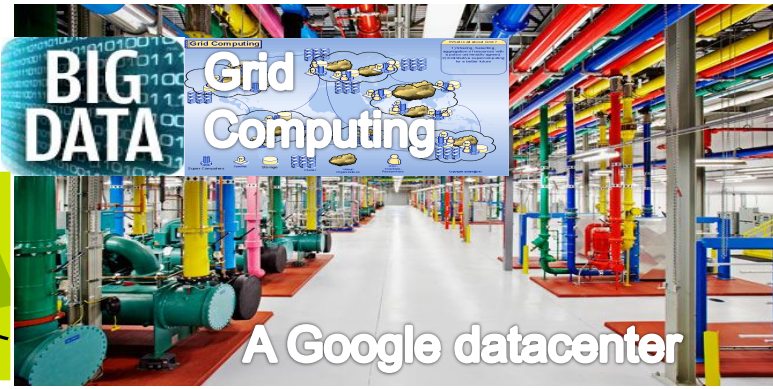

How to massivize?

- Super-scalable, super-flexible, yet efficient ICT infrastructure
- End-to-end automation of large-scale, simple and complex processes
- Dynamic, compute- and data-intensive workloads
- Evolving, heterogeneous hardware and software
- Strict performance, cost, energy, reliability, and fairness requirements
- … all these, without needing much expertise from customers

↖ There's a lot we don't know how to do yet… you can help!

The Scheduling Challenge

Cloud operator:

Which resources to lease?
Where to place? Penalty v reward?

Need scheduling policies for both the cloud user and the cloud operator

Cloud customer:

Which resources to lease?
When? How many? When stop?
Utility functions?

The Ecosystem Navigation Challenge

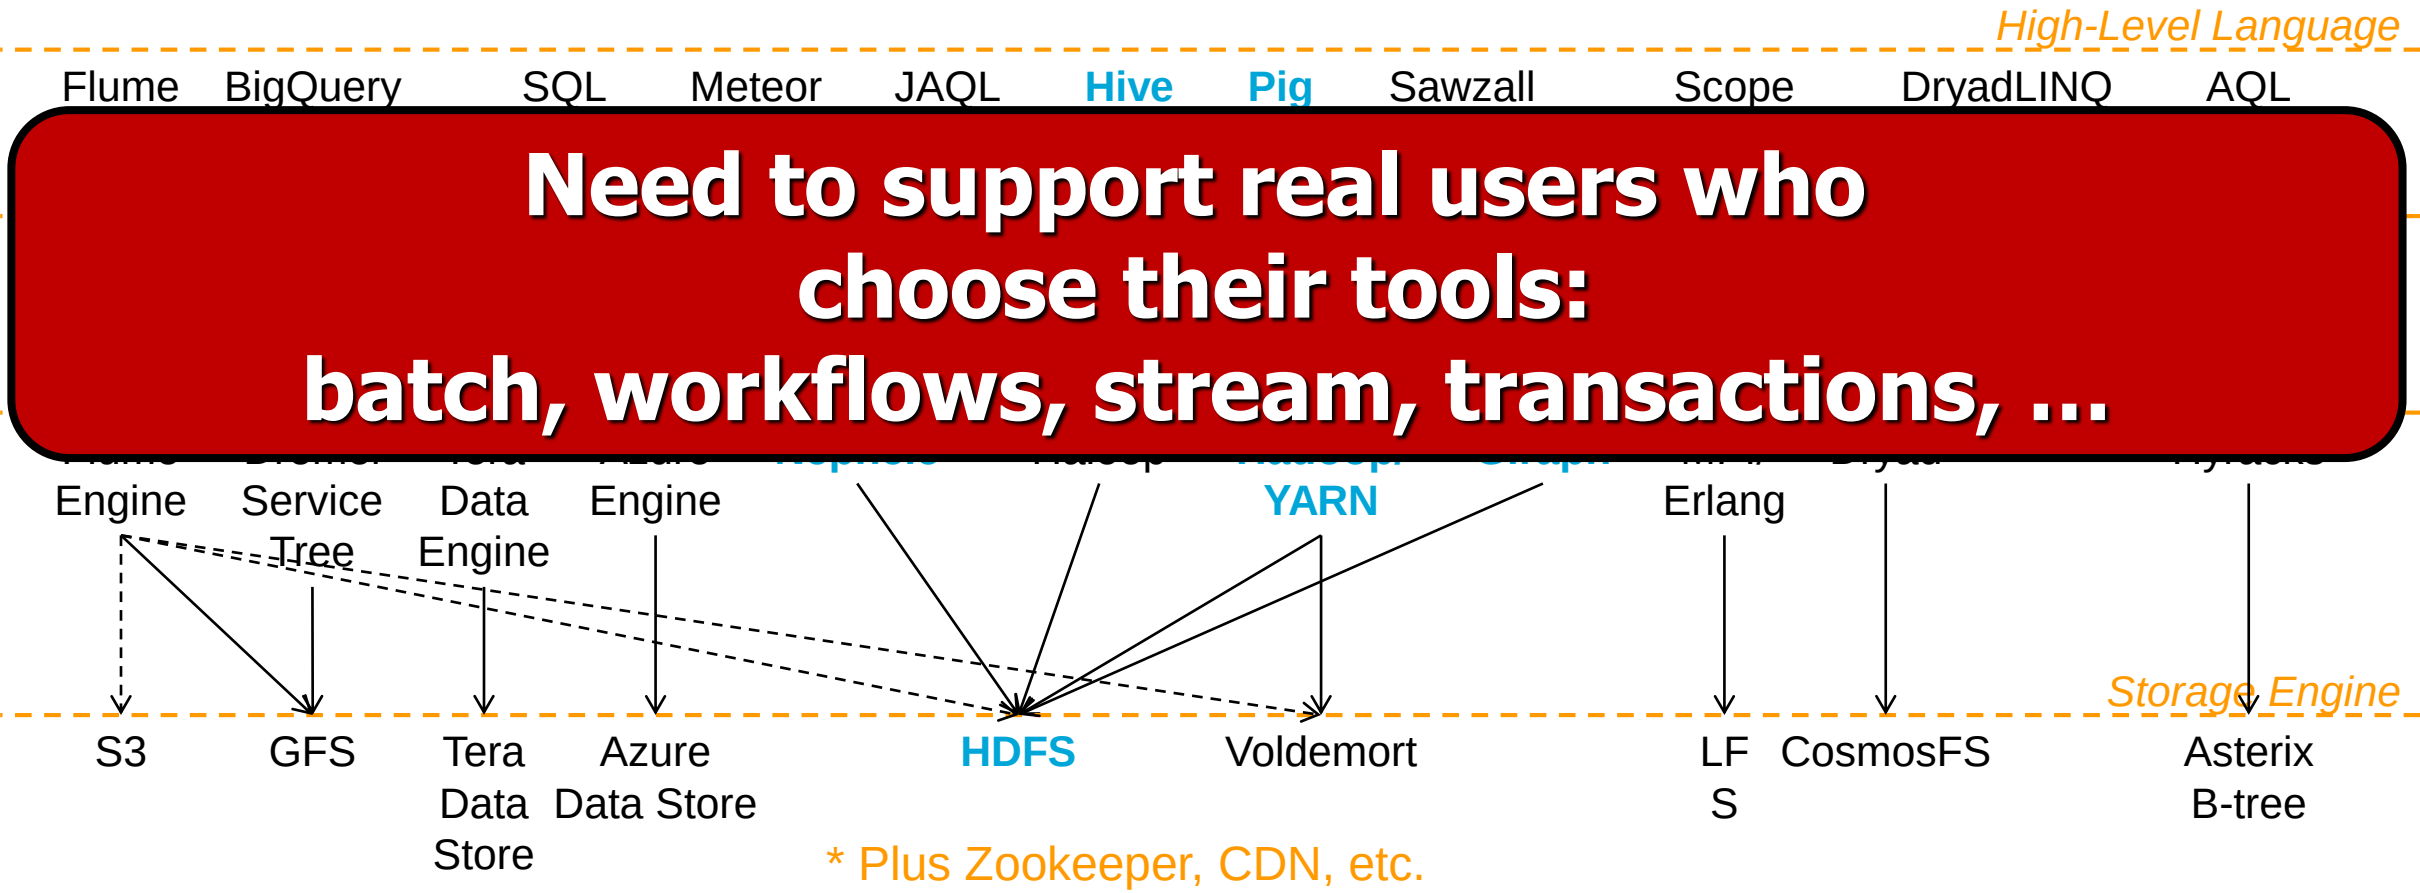

The “Big Cake” Challenge In the Datacenter

Online Social Networks

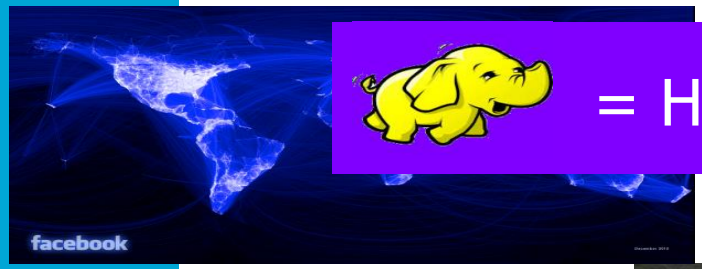

= Hadoop / MapReduce framework

Financial Analysts

Need multi-tenant, self-aware schedulers and resource managers

Universe Explorers

Multiple frameworks = Isolation, especially performance

Big Data Enthusiast

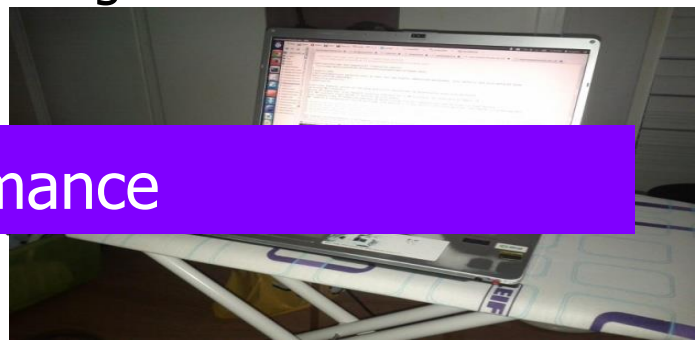

(Source: B. Ghit et al., SIGMETRICS 2014)

Jevons Effect: More Efficient, Less Capable?

Over 500 YouTube videos have at least 100,000,000 viewers each.

Need to be more efficient in how we use our resources, (also educate others to not abuse "infinite" capacity)

PSY Gangnam consumed ~500GWh

= more than entire countries* in a year (*41 countries),
= over 50MW of 24/7/365 diesel, 135M liters of oil,
= 100,000 cars running for a year, ...

Source: Ian Bitterlin and Jon Summers, UoL, UK, Jul 2013.
Note: Psy has >3 billion views (Nov 2015).

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

The Scheduling Challenge

Cloud operator:

Which resources to lease?
Where to place? Penalty v reward?

Need scheduling policies for both the cloud user and the cloud operator

Cloud customer:

Which resources to lease?
When? How many? When stop?
Utility functions?

Portfolio Scheduling, In A Nutshell

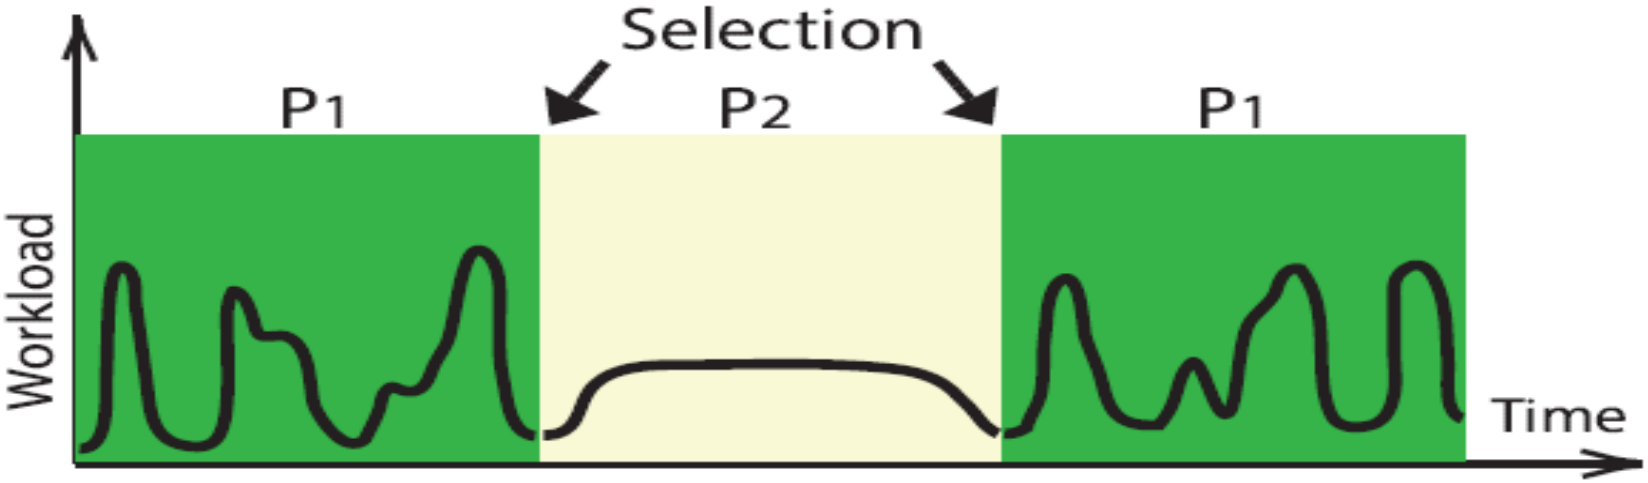

- Create a set of scheduling policies
  - Resource provisioning and allocation policies for datacenters
- Online selection of the active policy, at important moments

Portfolio Scheduling: Process

Which policies to include?

Creation

Which policy to activate?

Selection

Application

Which resources? What to log?

Reflection

Which changes to the portfolio?

Portfolio Scheduling in Practice

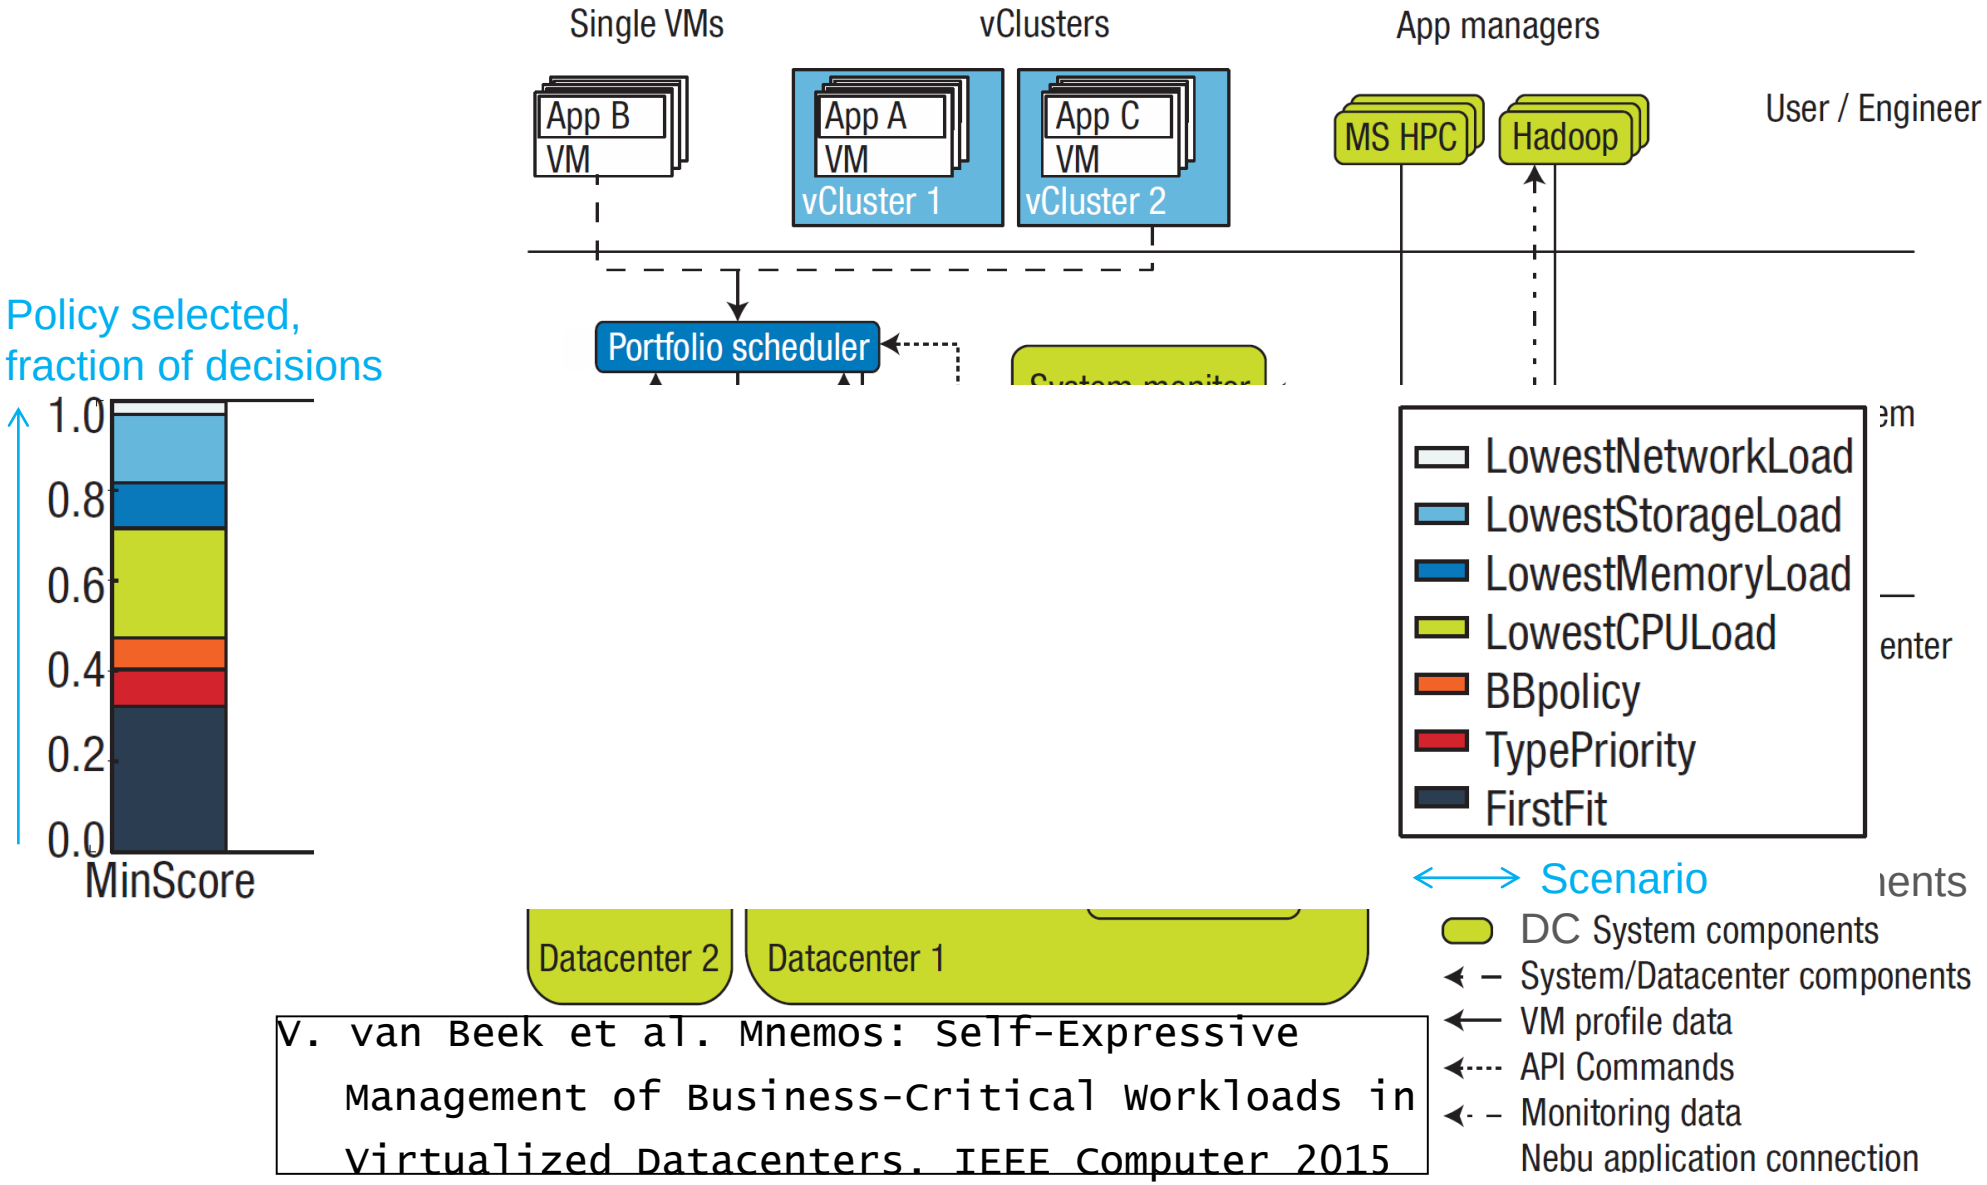

V. van Beek et al. Mnemos: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters. IEEE Computer 2015

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

The Ecosystem Navigation Challenge

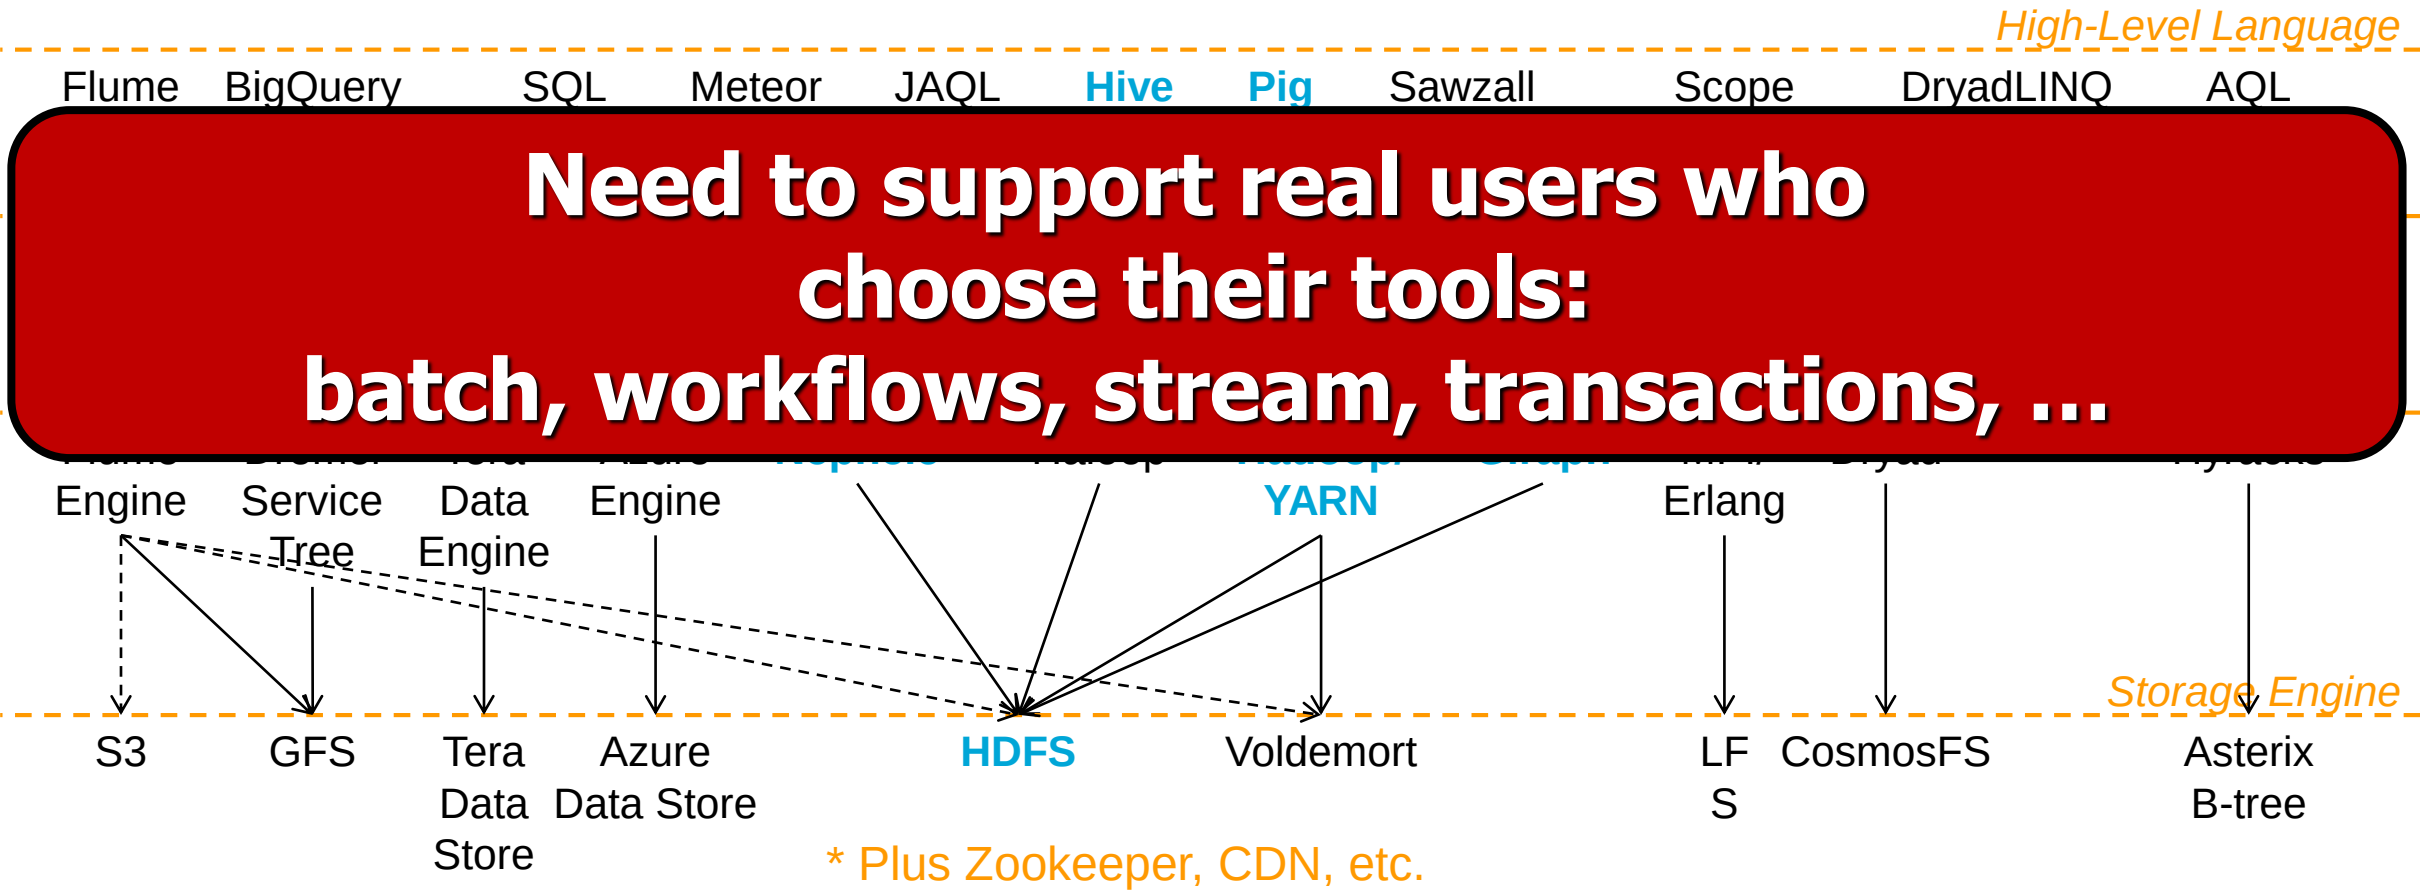

Use Case:
Provisioning Policies, Compared

Startup

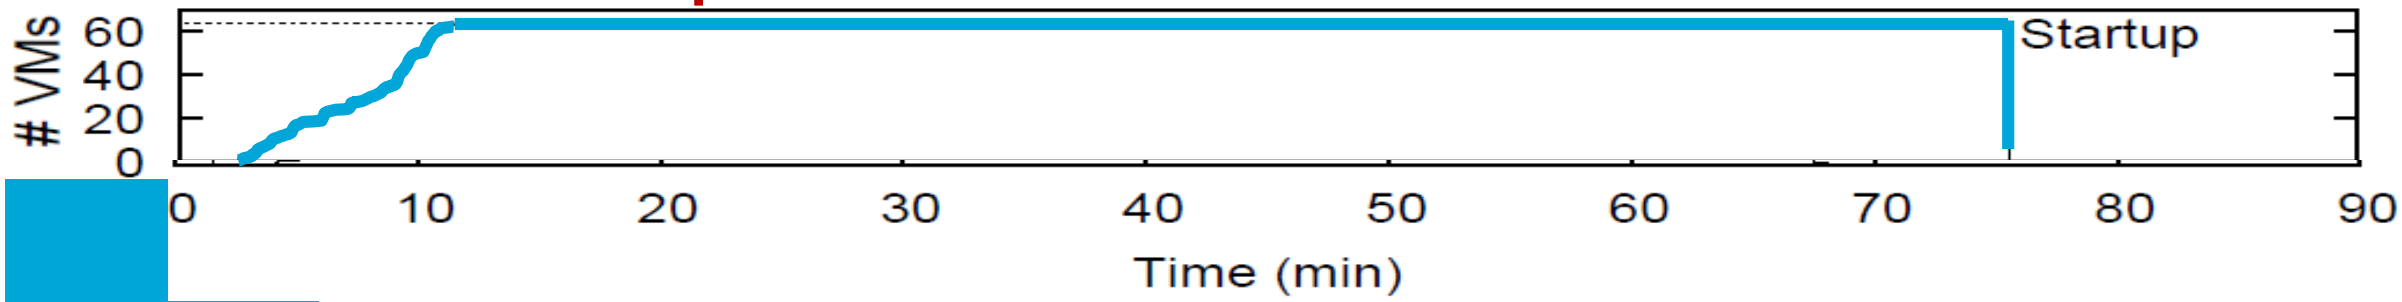

OnDemand

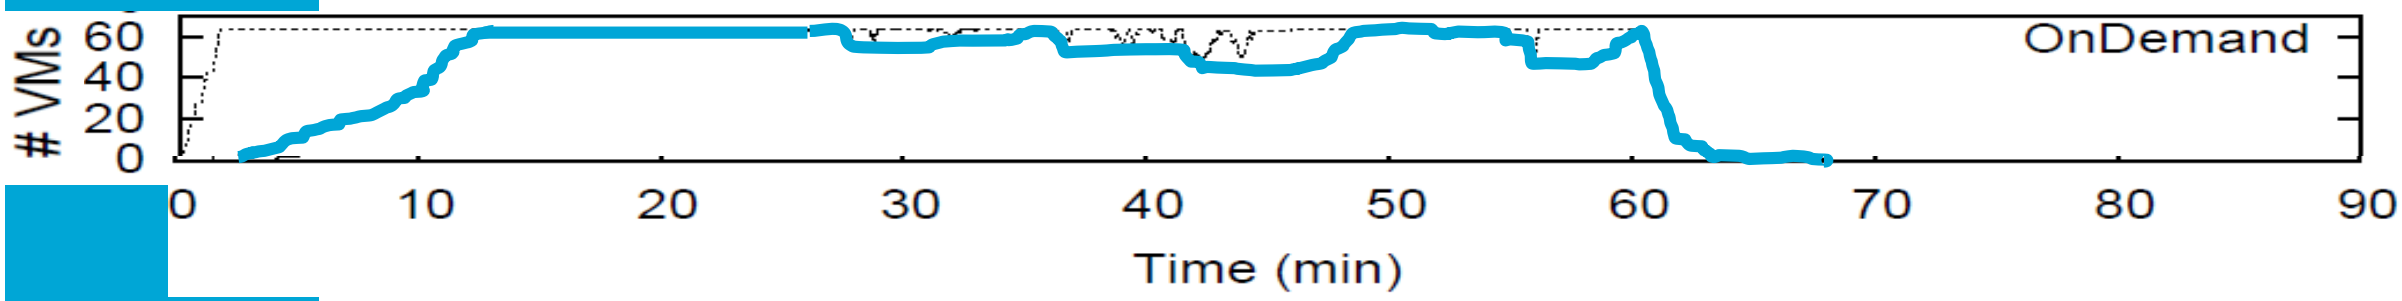

Requested Instances ---------- Accessible Instances ——

Villegas, Antoniou, Sadjadi, Iosup. An Analysis of Provisioning and Allocation Policies for Infrastructure-as-a-Service Clouds, CCGRID'12

Use Case: Provisioning Policies, Compared

Environments (values for 2012 study, 2016 study larger)

System	Hardware	VIM	Hypervisor	Max VMs
DAS4/Delft	20 Dual quad-core 2.4 GHz 24 GB RAM 2x1 TB storage	OpenNebula	KVM	64
FIU	7 Pentium 4 3.0 GHz 5 GB RAM 340 GB Storage	OpenNebula	Xen	7
Amazon EC2	unkown/various	-	Xen	20

Villegas, Antoniou, Sadjadi, Iosup. An Analysis of Provisioning and Allocation Policies for Infrastructure-as-a-Service Clouds, CCGRID'12

Use Case:
Two Provisioning Policies, Compared

Workloads

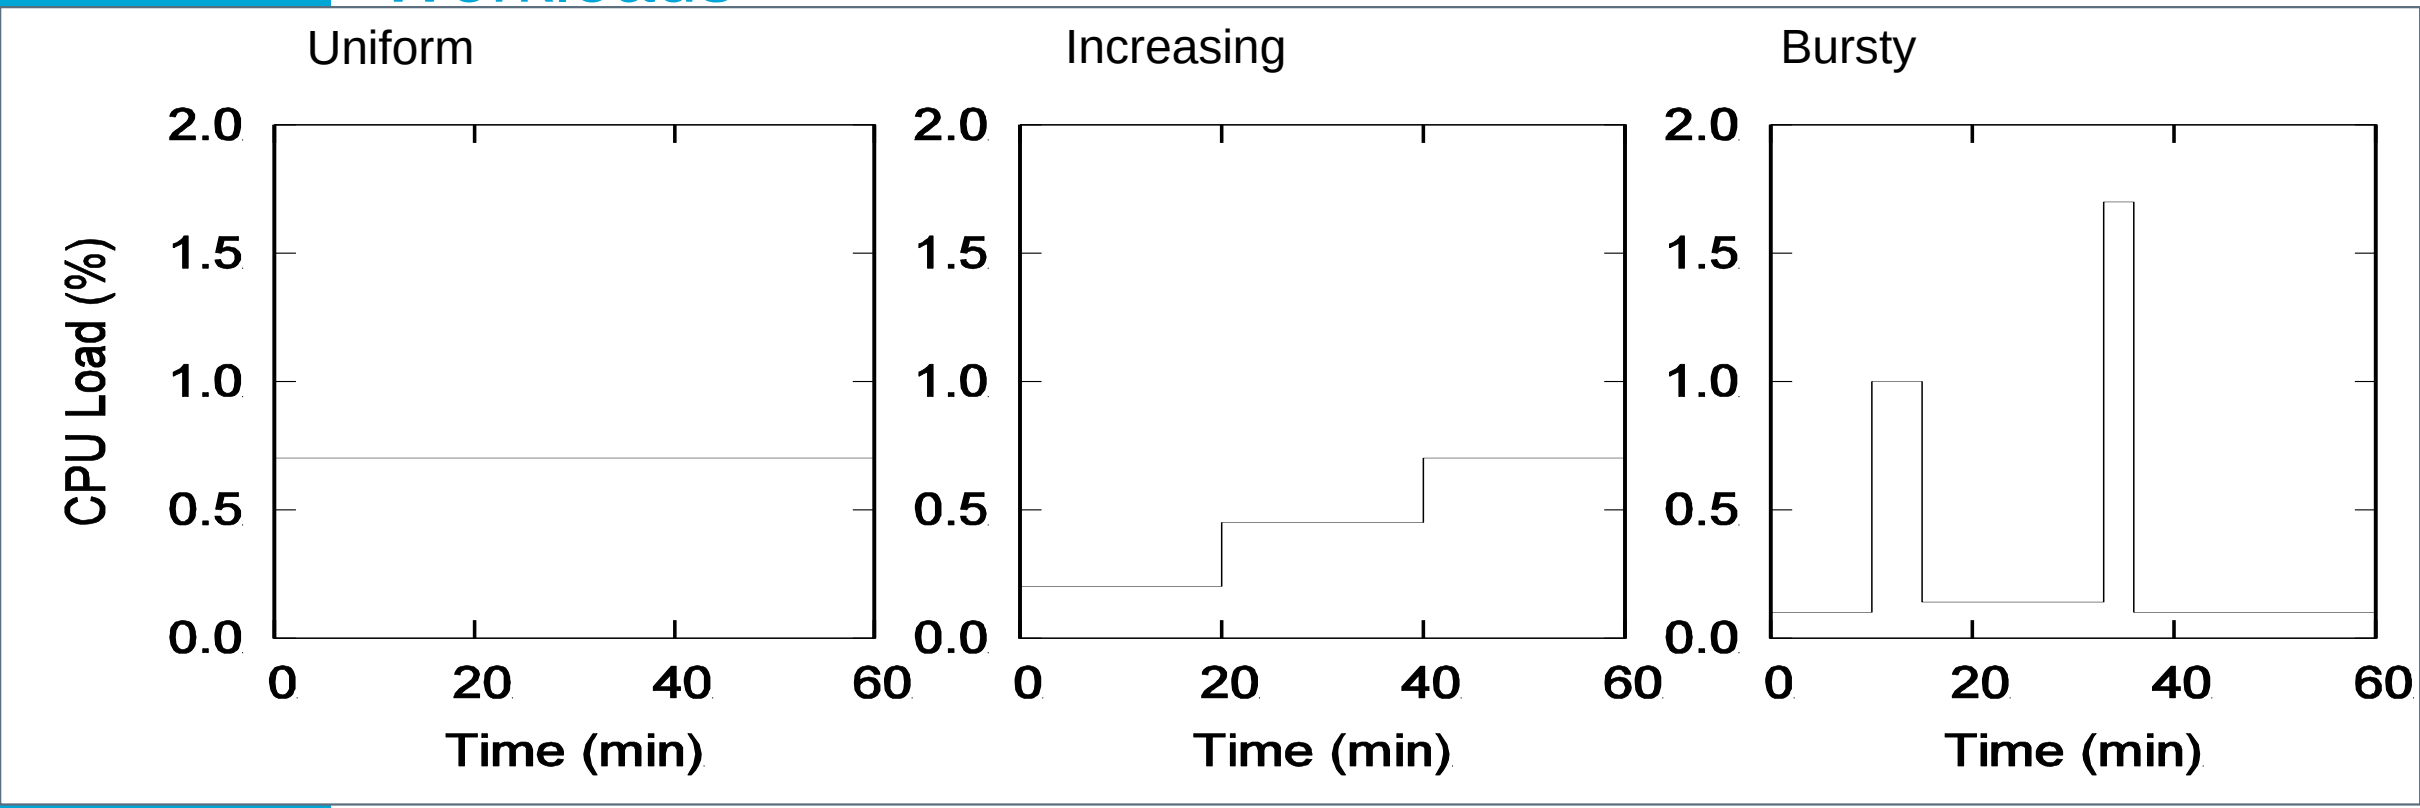

Villegas, Antoniou, Sadjadi, Iosup. An Analysis of Provisioning and Allocation Policies for Infrastructure-as-a-Service Clouds, CCGRID'12

Use Case:
Two Provisioning Policies, Compared

Metrics for comparison

- Job Slowdown (***JSD***): Ratio of actual runtime in the cloud and the runtime in a dedicated non-virtualized environment

Q: Charged cost vs Total RunTime?

- Charged Cost (C_c)

$$C_c(W) = \sum_{i \in leased\ VMs} \lceil t_{stop}(i) - t_{start}(i) \rceil$$

- Utility (***U***)

$$U(W) = \frac{SU_1(W)}{C_c(W)}$$

Use Case:
Many Provisioning Policies, Compared
Job Slowdown (JSD)

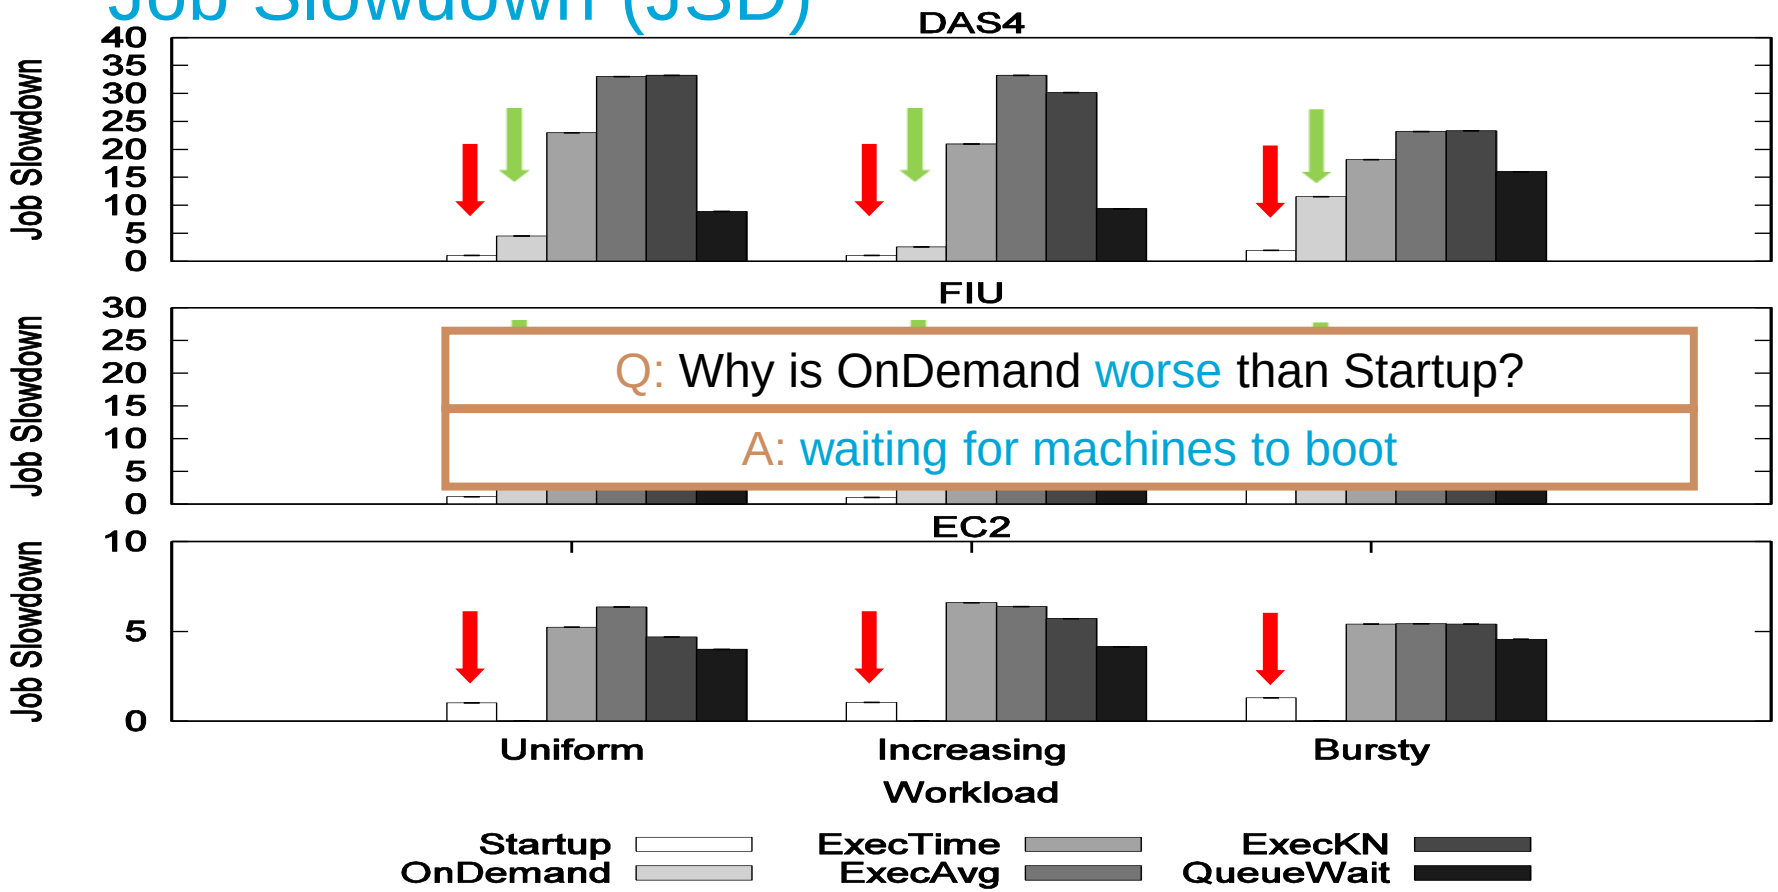

Use Case:
Many Provisioning Policies, Compared

Charged Cost (C_c)

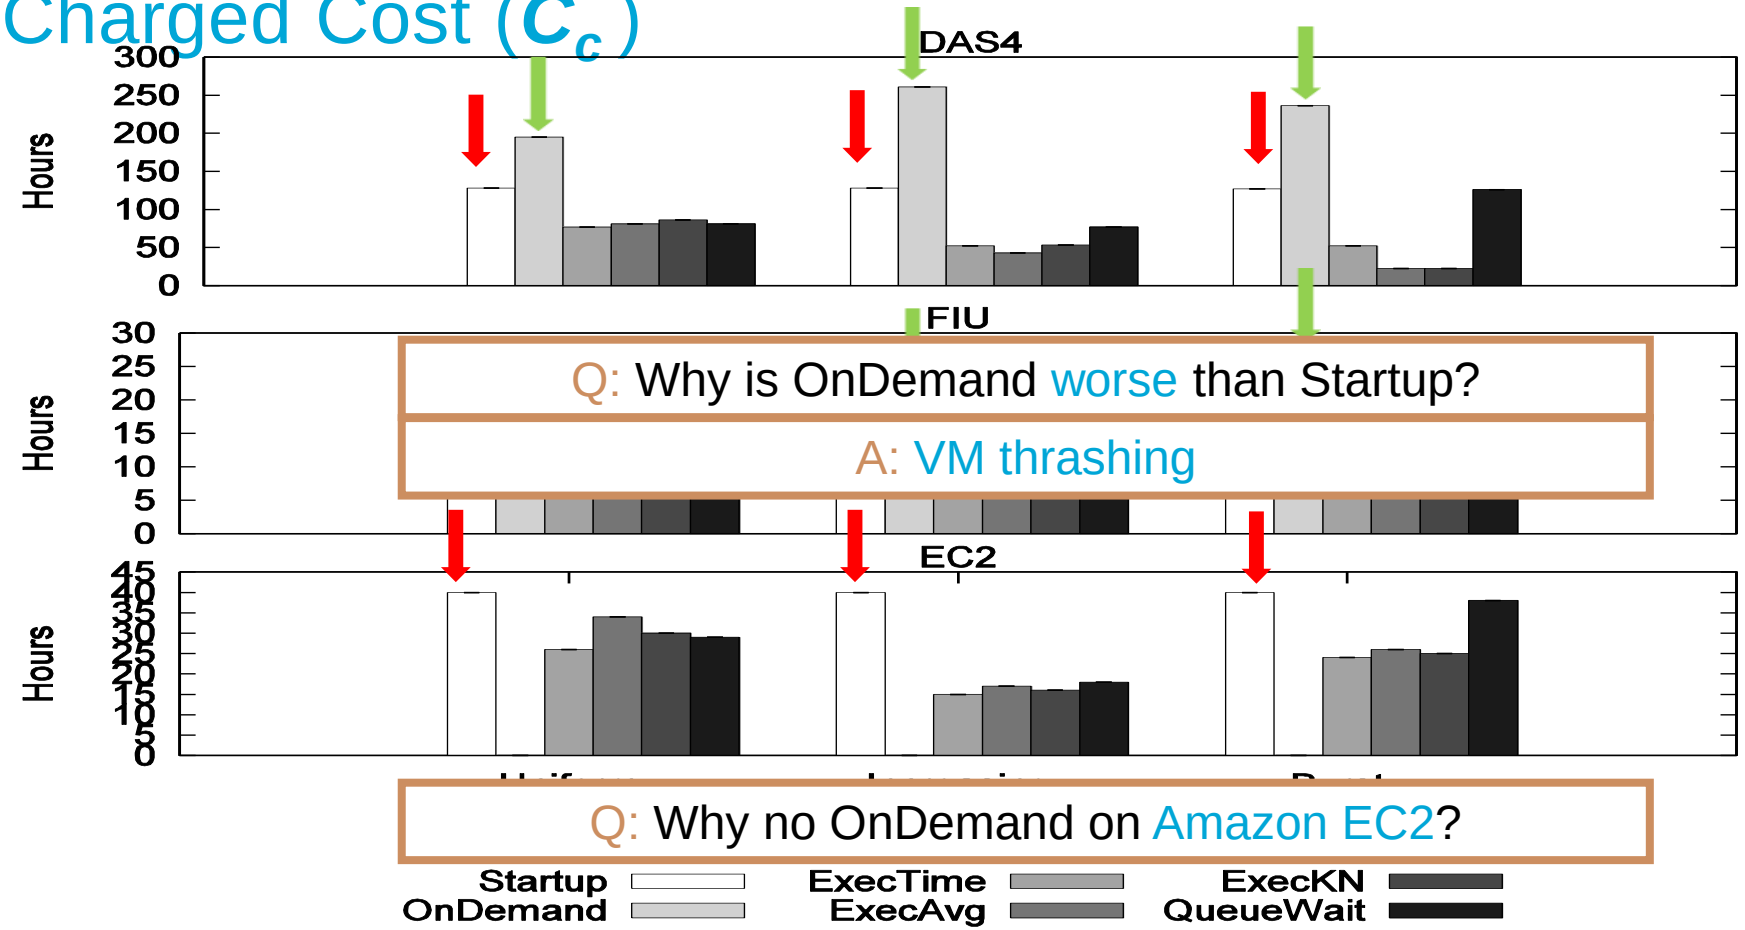

Jul 13, 2017

With SPEC RG Cloud Group: Currently Conducting Largest Auto-Scale Study

Alexey Ilyushkin

Challenges for auto-scalers across many application domains

- **Scientific workflows** have varying resource requirements
- **Big data apps** use HDFS, which is difficult to scale
- A **complex web-server** consists of many components

Graph-Processing Platforms

- Platform: the combined hardware, software, and programming system that is being used to complete a graph processing task

Graphalytics, in a nutshell

- An LDBC benchmark*
- Advanced benchmarking harness
- Diverse real and synthetic datasets
- Many classes of algorithms
- Granula for manual choke-point analysis
- Modern software engineering practices
- Supports many platforms
- Enables comparison of community-driven and industrial systems

http://graphalytics.ewi.tudelft.nl
https://github.com/tudelft-atlarge/graphalytics/

Graphalytics Capabilities: An Example

Graphalytics enables deep comparison of many systems at once, through diverse experiments and metrics

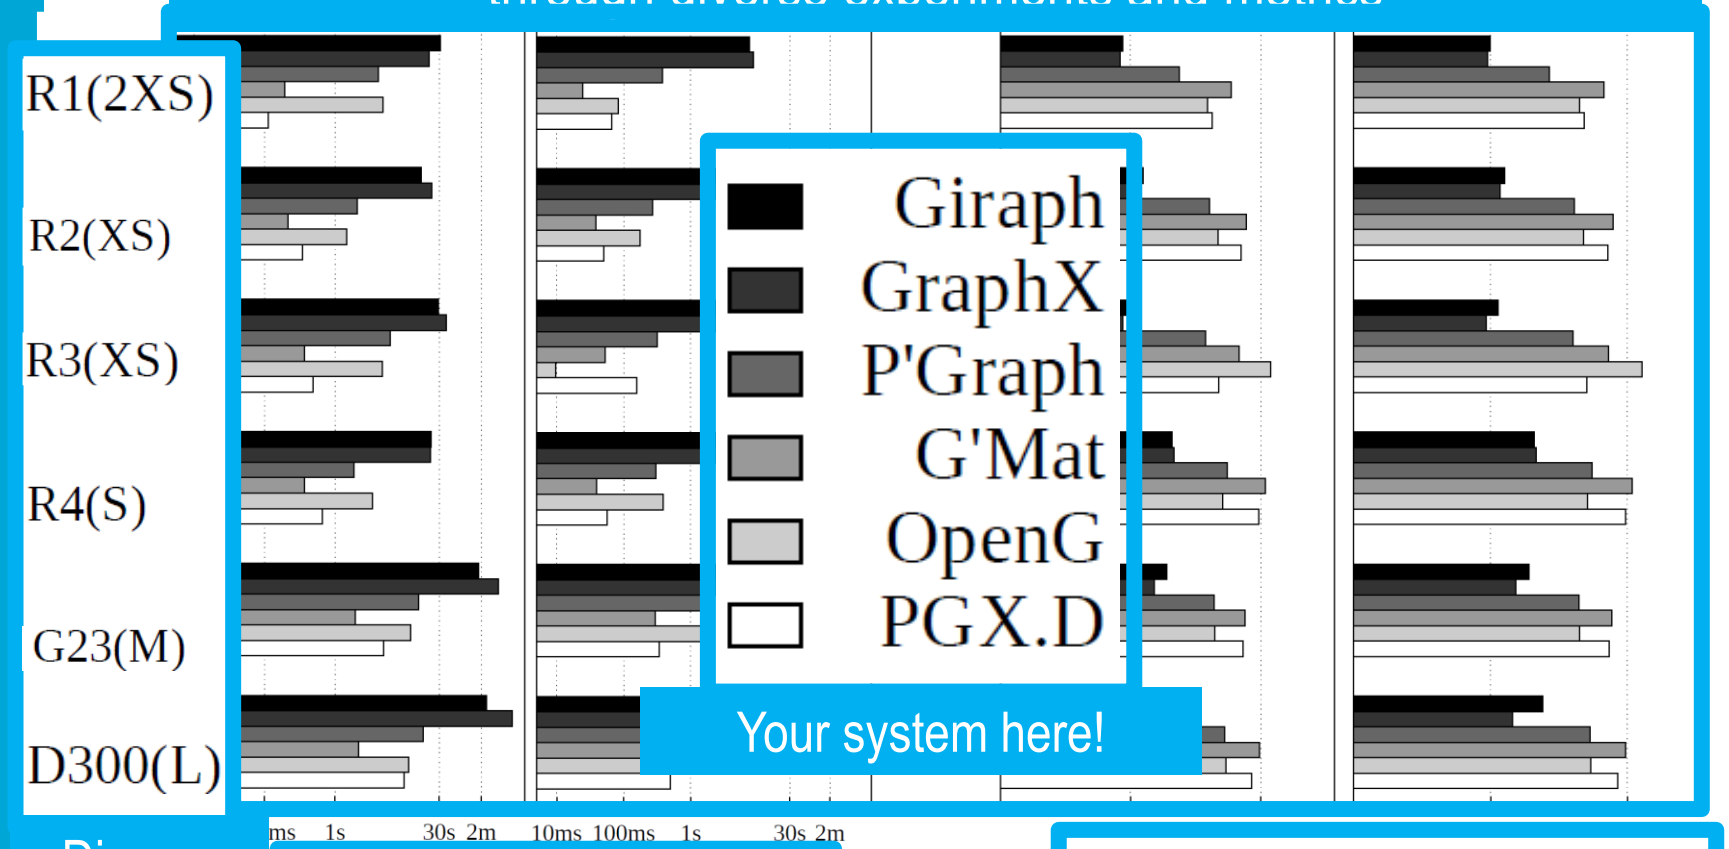

Processing time (s) + Edges[+Vertices]/s

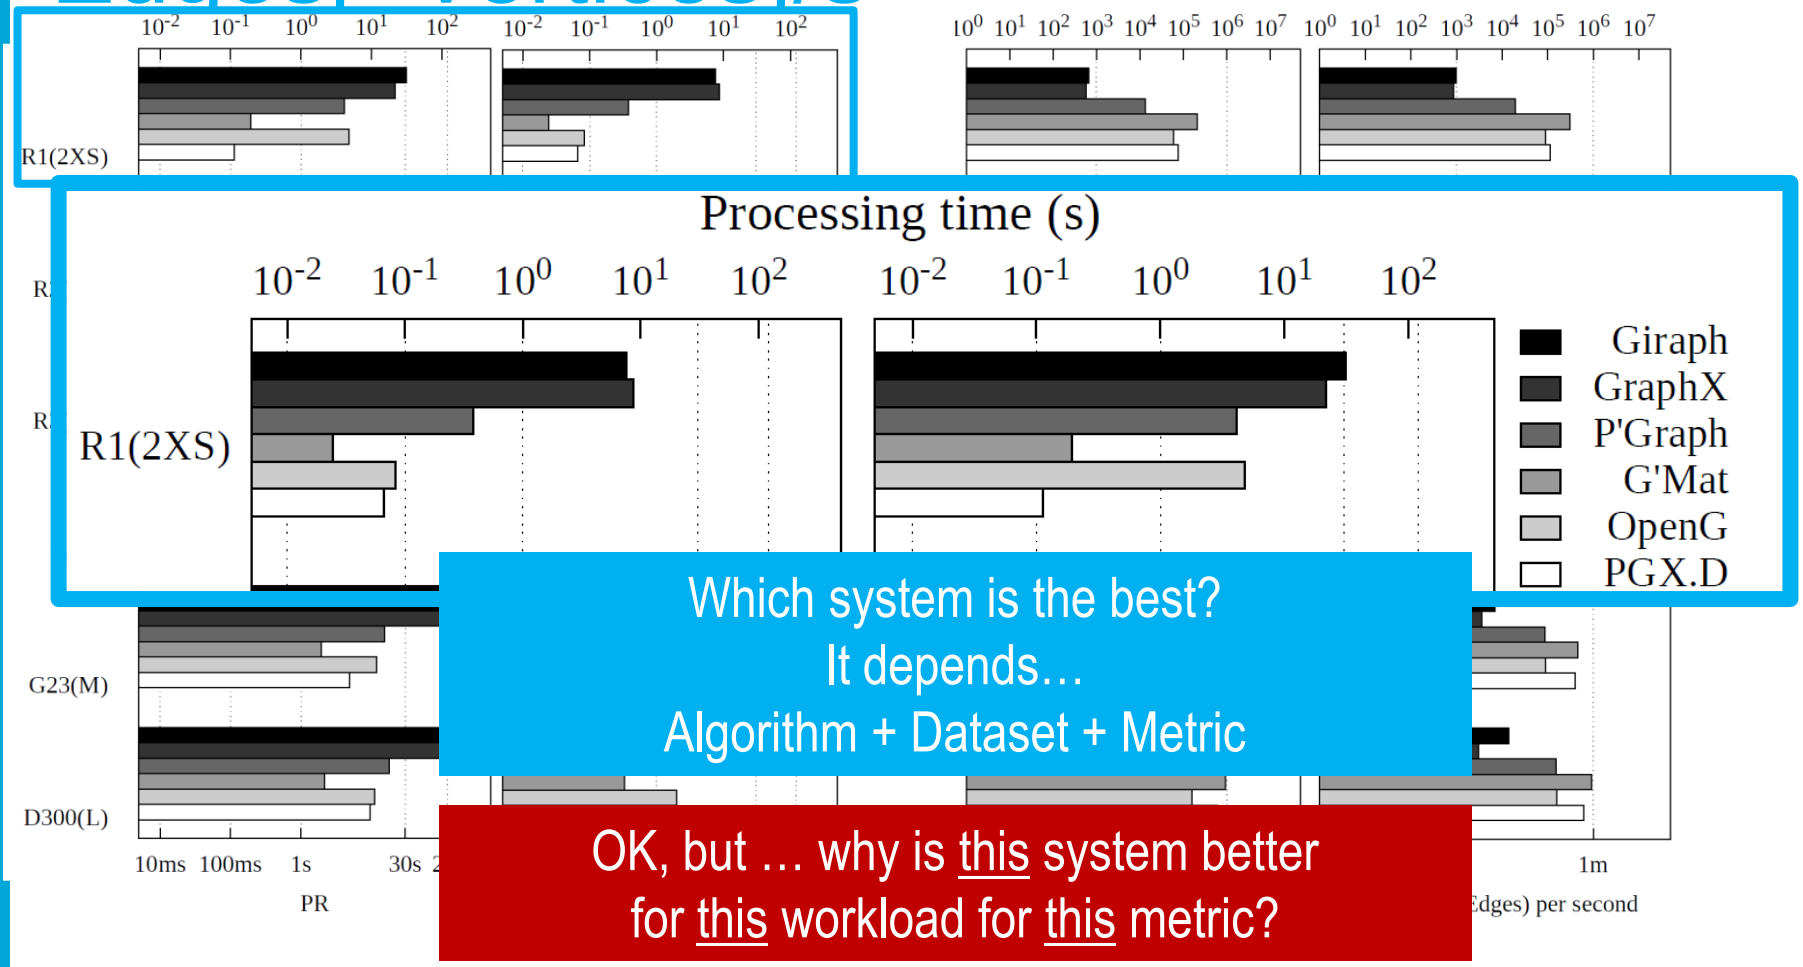

Which system is the best?
It depends…
Algorithm + Dataset + Metric

OK, but … why is this system better for this workload for this metric?

66

Granula Visualizer

Portable choke-point analysis for everyone!

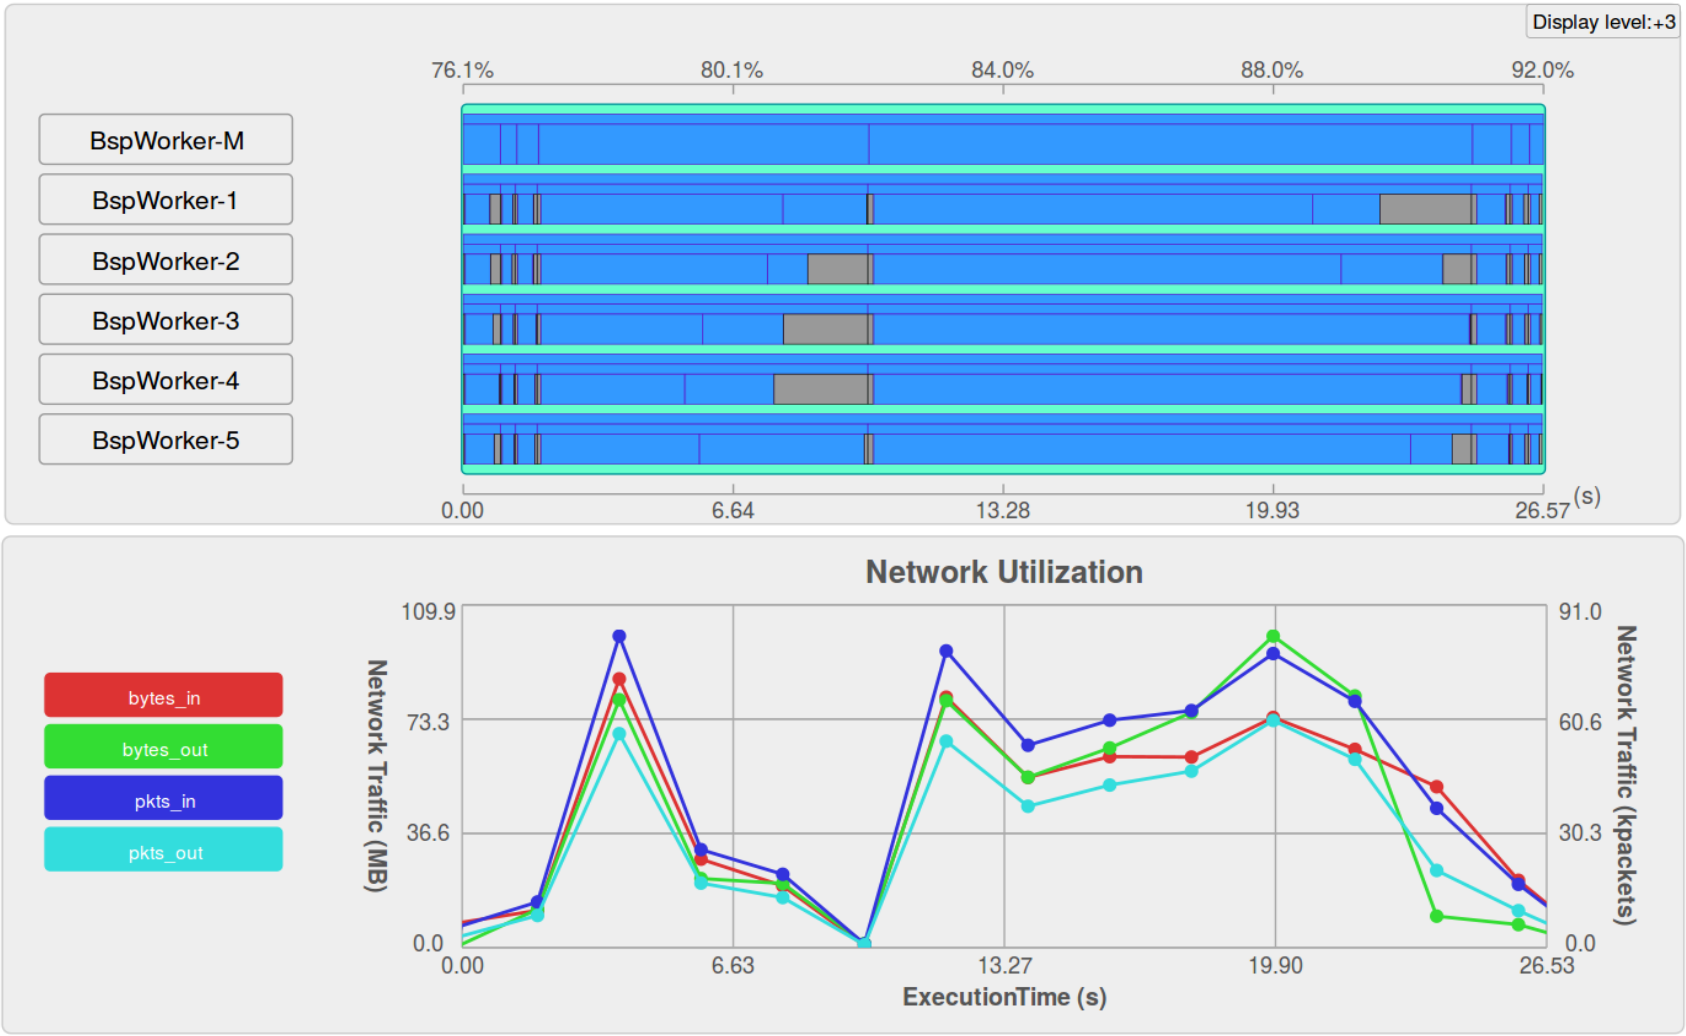

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

The “Big Cake” Challenge In the Datacenter

Online Social Networks

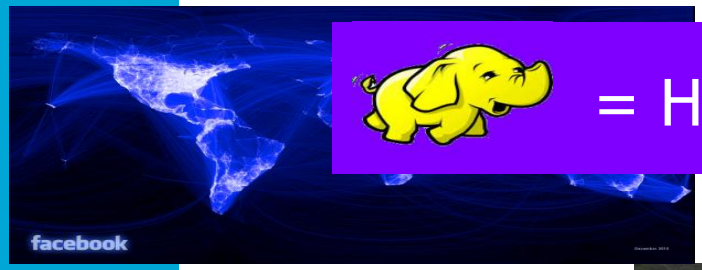

Financial Analysts

= Hadoop / MapReduce framework

Need multi-tenant, self-aware schedulers and resource managers

Universe Explorers

Big Data Enthusiast

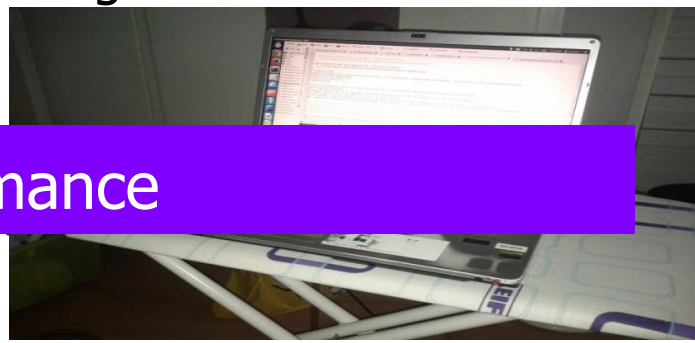

Multiple frameworks = Isolation, especially performance

Dynamic Big Data Processing

Fawkes = Elastic MapReduce

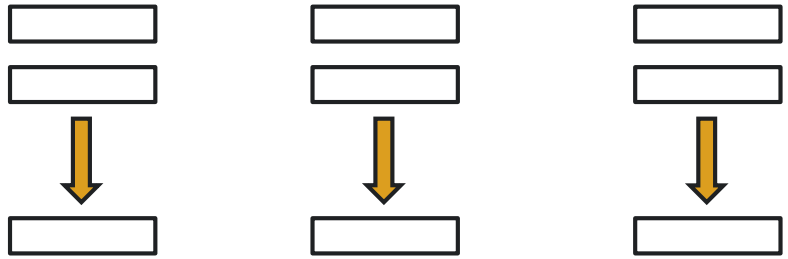

Job submissions

Frameworks

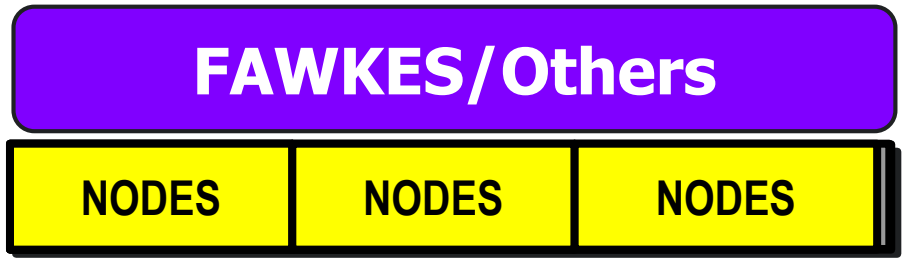

Resource manager

Infrastructure

Elasticity for MapReduce Frameworrks

Core nodes

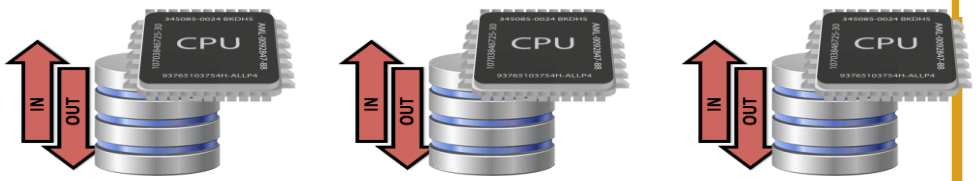

INPUT/OUTPUT DATA

- Classical deployment
- Uniform data distribution
- **No removal**

Transient nodes (TR)

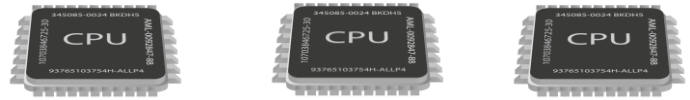

NO DATA

- No local storage
- R/W from/to core nodes
- **Instant removal**

Trans-core nodes (TC)

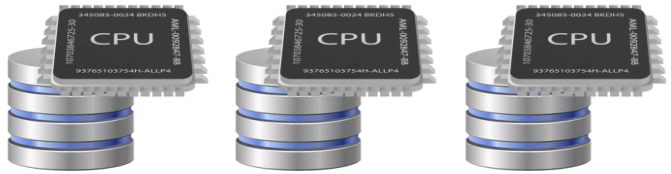

OUTPUT DATA

- Local storage, no input
- Only R from core nodes
- **Delayed removal**

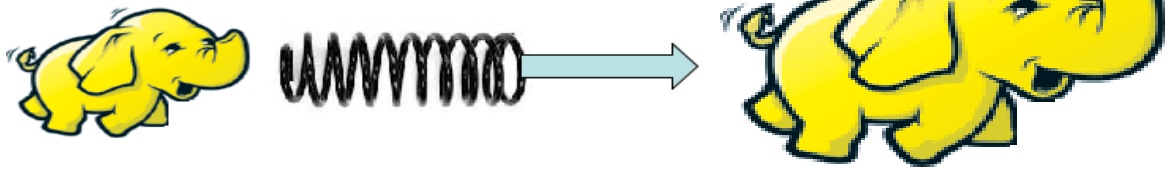

TU Delft

Fawkes in a Nutshell [1/2]

Because workloads may be time-varying:

- Poor resource utilization
- Imbalanced service levels

1. Fair framework size:

$$s_i = \frac{w_i}{w_1 + w_2 + w_3}, \quad i = 1,2,3$$

Fawkes in a Nutshell [2/2]

2. Updates dynamic weights when:
   - New frameworks arrive
   - Framework states change

3. Shrinks and grows frameworks to:
   - Allocate new frameworks
   - Give fair shares to existing frameworks
   - Eliminate unused frameworks

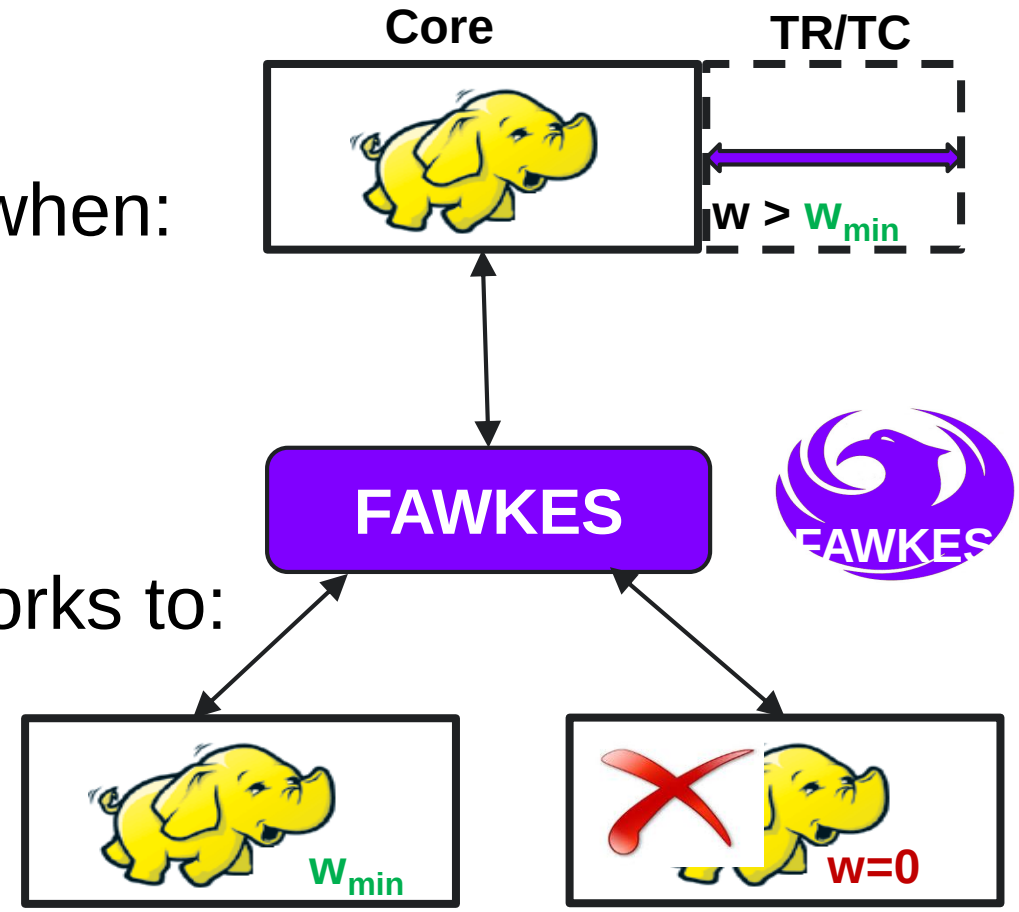

MapReduce Applications Tested with Fawkes

Application	Type	Input	Output
Wordcount (WC)	CPU	200 GB	5.5 MB
Sort (ST)	Disk	200 GB	200 GB
PageRank (PR)	CPU	50 GB	1.5 MB
K-Means (KM)	Both	70 GB	72 GB
TrackerOverTime (TT)	CPU	100 GB	3.9 MB
ActiveHashes (AH)	Both	100 GB	90 KB
BTWorld (BT)	Both	100 GB	73 GB

Synthetic benchmarks:
- HiBench suite
- Single applications
- Random datasets

Real-world applications:
- BTWorld workflow
- 14 Pig queries
- BitTorrent monitoring data

Performance of dynamic MapReduce

10 core + 10xTR
10 core + 10xTC
vs.
20 core nodes (baseline)

TRansient - good for compute-intensive workloads.

TCore - needed for disk-intensive workloads.

Dynamic MapReduce:
< 25% overhead

Fawkes also reduces imbalance

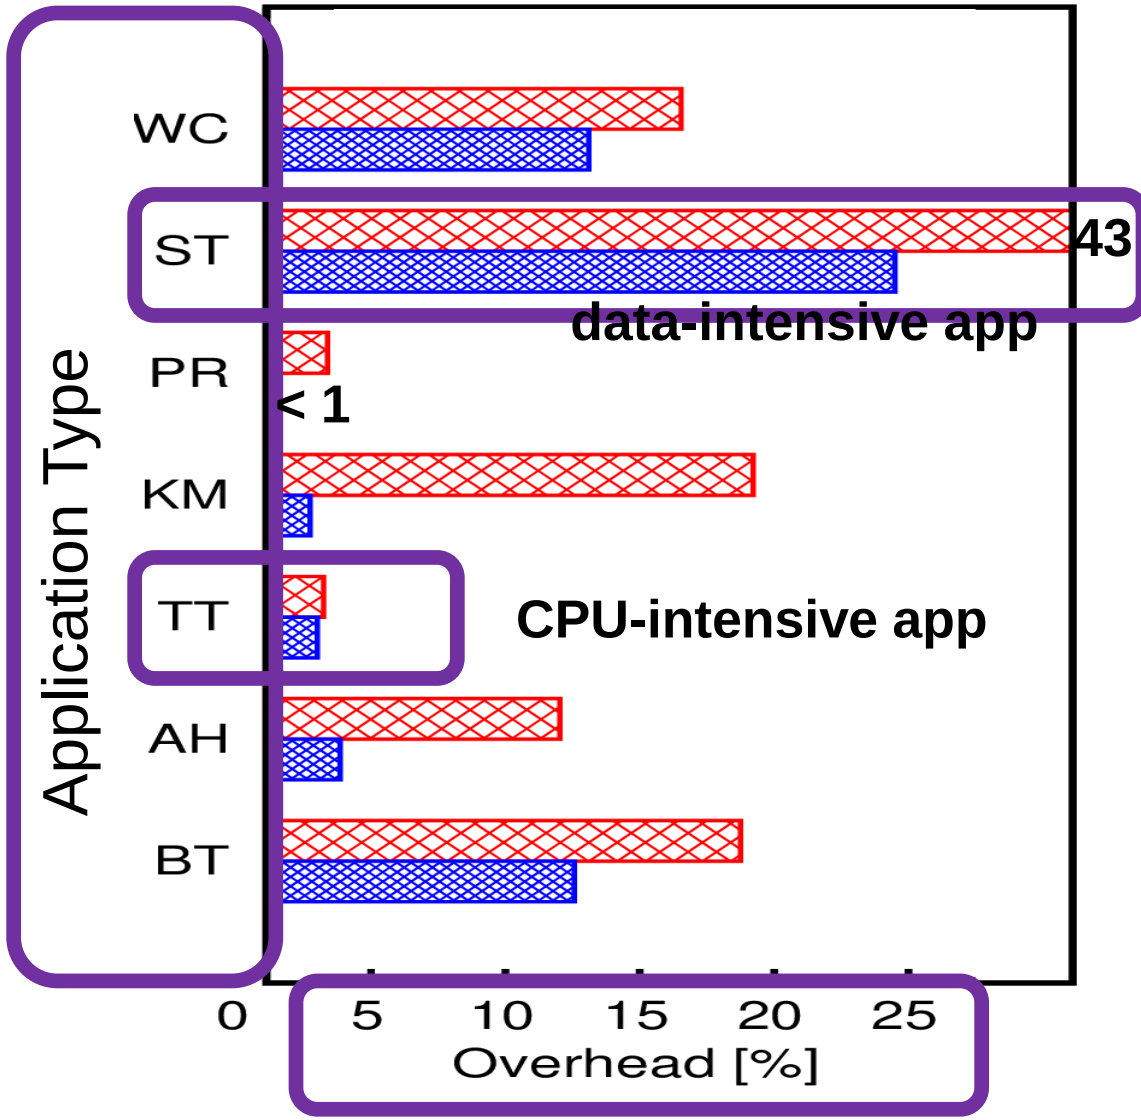

Intro to Cloud Computing

Here or @home

Interactive

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

Jevons Effect: More Efficient, Less Capable

Over 500 YouTube videos have at least 100,000,000 viewers each.

Need to be more efficient in how we use our resources, (also educate others to not abuse "infinite" capacity)

PSY Gangnam consumed ~500GWh

= more than entire countries* in a year (*41 countries),
= over 50MW of 24/7/365 diesel, 135M liters of oil,
= 100,000 cars running for a year, ...

Source: Ian Bitterlin and Jon Summers, UoL, UK, Jul 2013.
Note: Psy has >3 billion views (Nov 2015).

Existing Graph-Processing Systems: *Either* Distributed *or* Heterogeneous

- Distributed CPU-based systems cannot use additional computational power of accelerators

Oracle Labs PGX

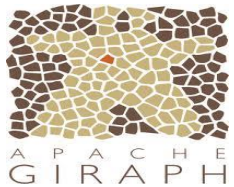

GraphLab

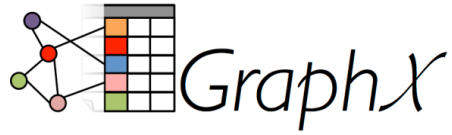

- GPU-enabled systems are (mostly) single-machine systems, cannot handle large-scale graphs

medusa-gpu

Medusa: Simplified Graph Processing on GPUs

mapgraph Beta

Massively Parallel Graph processing on GPUs

TOTEM

Gunrock

High-performance Graph Primitives on GPU

VertexAPI2

Y. Guo, A. L. Varbanescu, D. Epema, and A. Iosup, "Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems," Submission to *CCGrid, 2016.*

78

Our approach: 3 Families of Distributed *and* Heterogeneous (CPU+GPU) Graph-Processing Systems

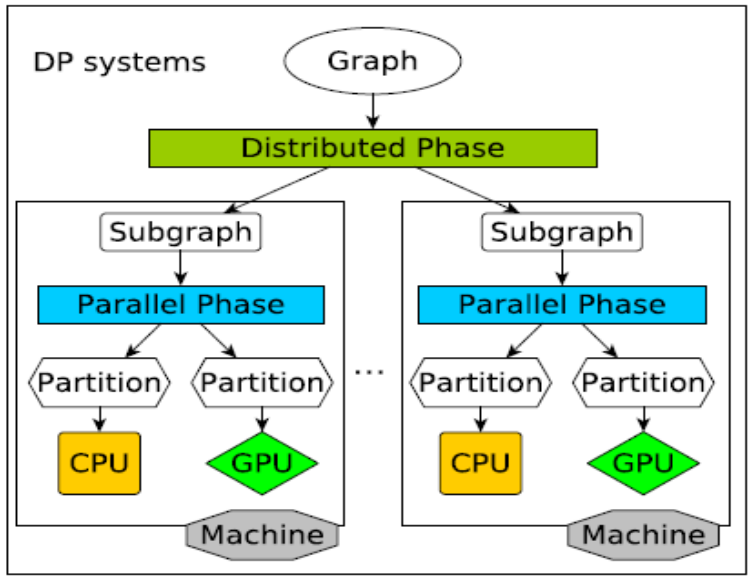

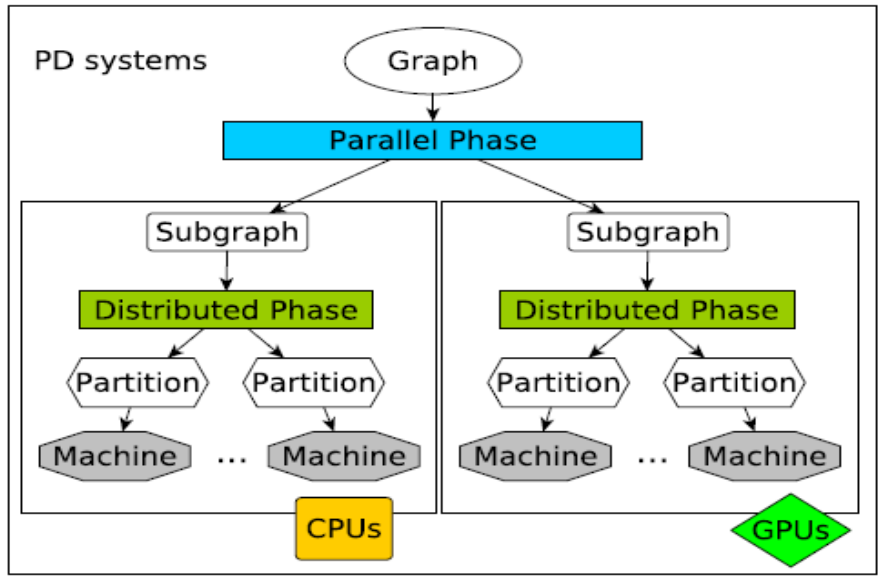

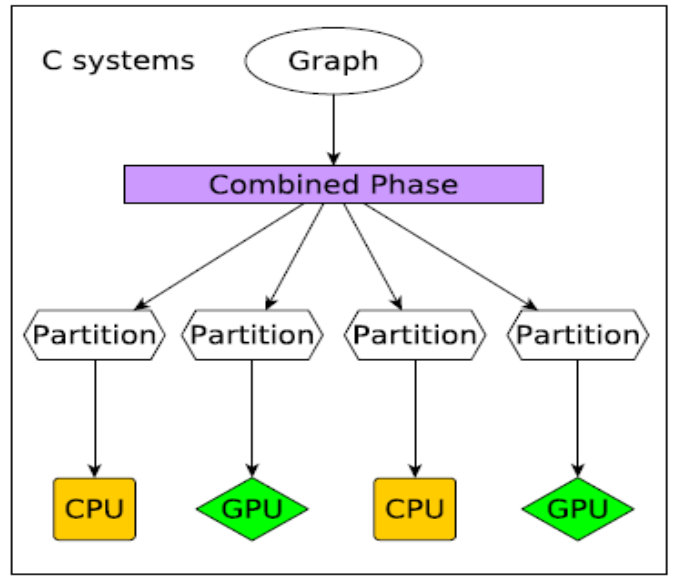

Distributed-then-Parallel (DP) Systems

Parallel-then-Distributed (PD) Systems

(Combined Par.-and-Distributed (C) Systems

Y. Guo, A. L. Varbanescu, D. Epema, and A. Iosup, "Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems," Submission to *CCGrid, 2016.*

79

Promising Results for Distributed *and* Heterogeneous Graph-Processing Systems

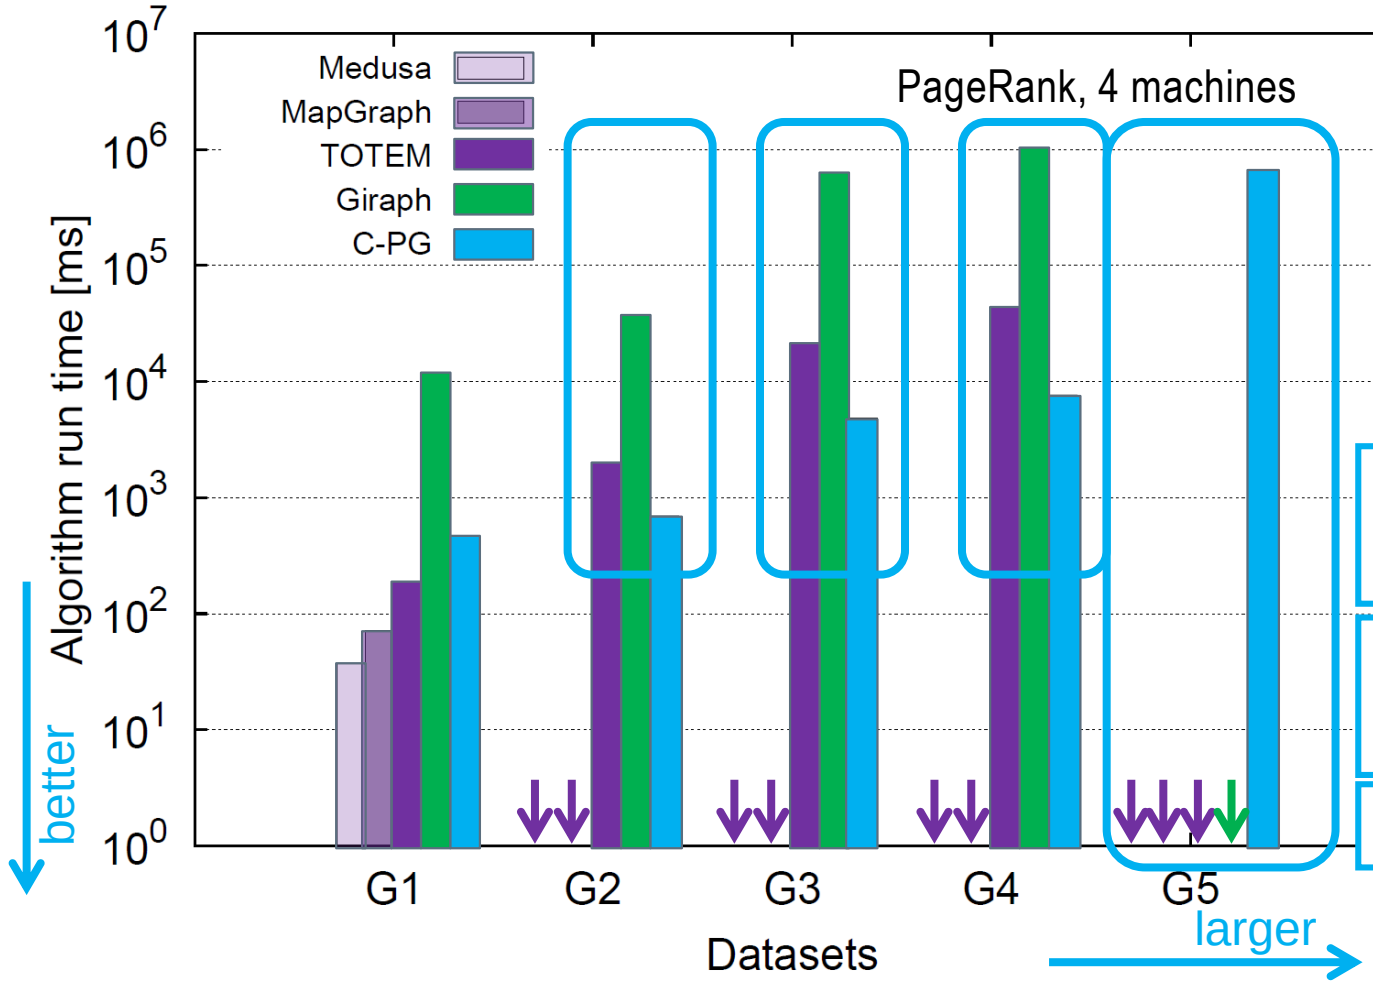

Distributed *and* Heterogeneous

Distributed only

Heterogeneous only

Missing bar = → = system failure

C-PG system has good performance iff. large data

TOTEM has good performance iff. in-memory

C-PG system scales

Y. Guo, A. L. Varbanescu, D. Epema, and A. Iosup, "Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems," Submission to *CCGrid, 2016.*

Intro to Cloud Computing

Here or @home

Interactive Masterclass

5’ — Pitch on Datacenter-Based Cloud Computing

5’ — The Golden Age of Datacenters

20’ — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35’ — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10’ — Reality Check

THE REAL IAAS CLOUD

http://www.flickr.com/photos/dimitrisotiropoulos/4204766418/

VS

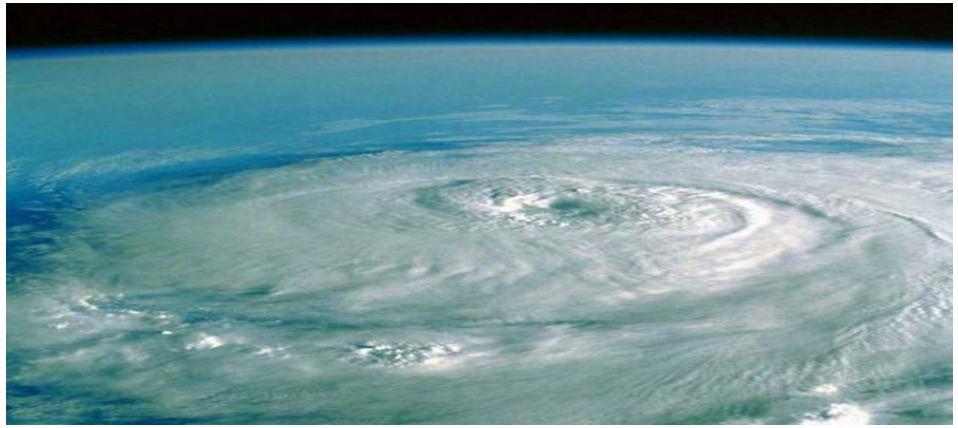

Tropical Cyclone Nargis (NASA, ISSS, 04/29/08)

- "The path to abundance"
- On-demand capacity
- Cheap for short-term tasks
- Great for web apps (EIP, web crawl, DB ops, I/O)

- "The killer cyclone"
- Not so great performance for scientific applications (compute- or data-intensive)

Animoto: Video App on Amazon EC2

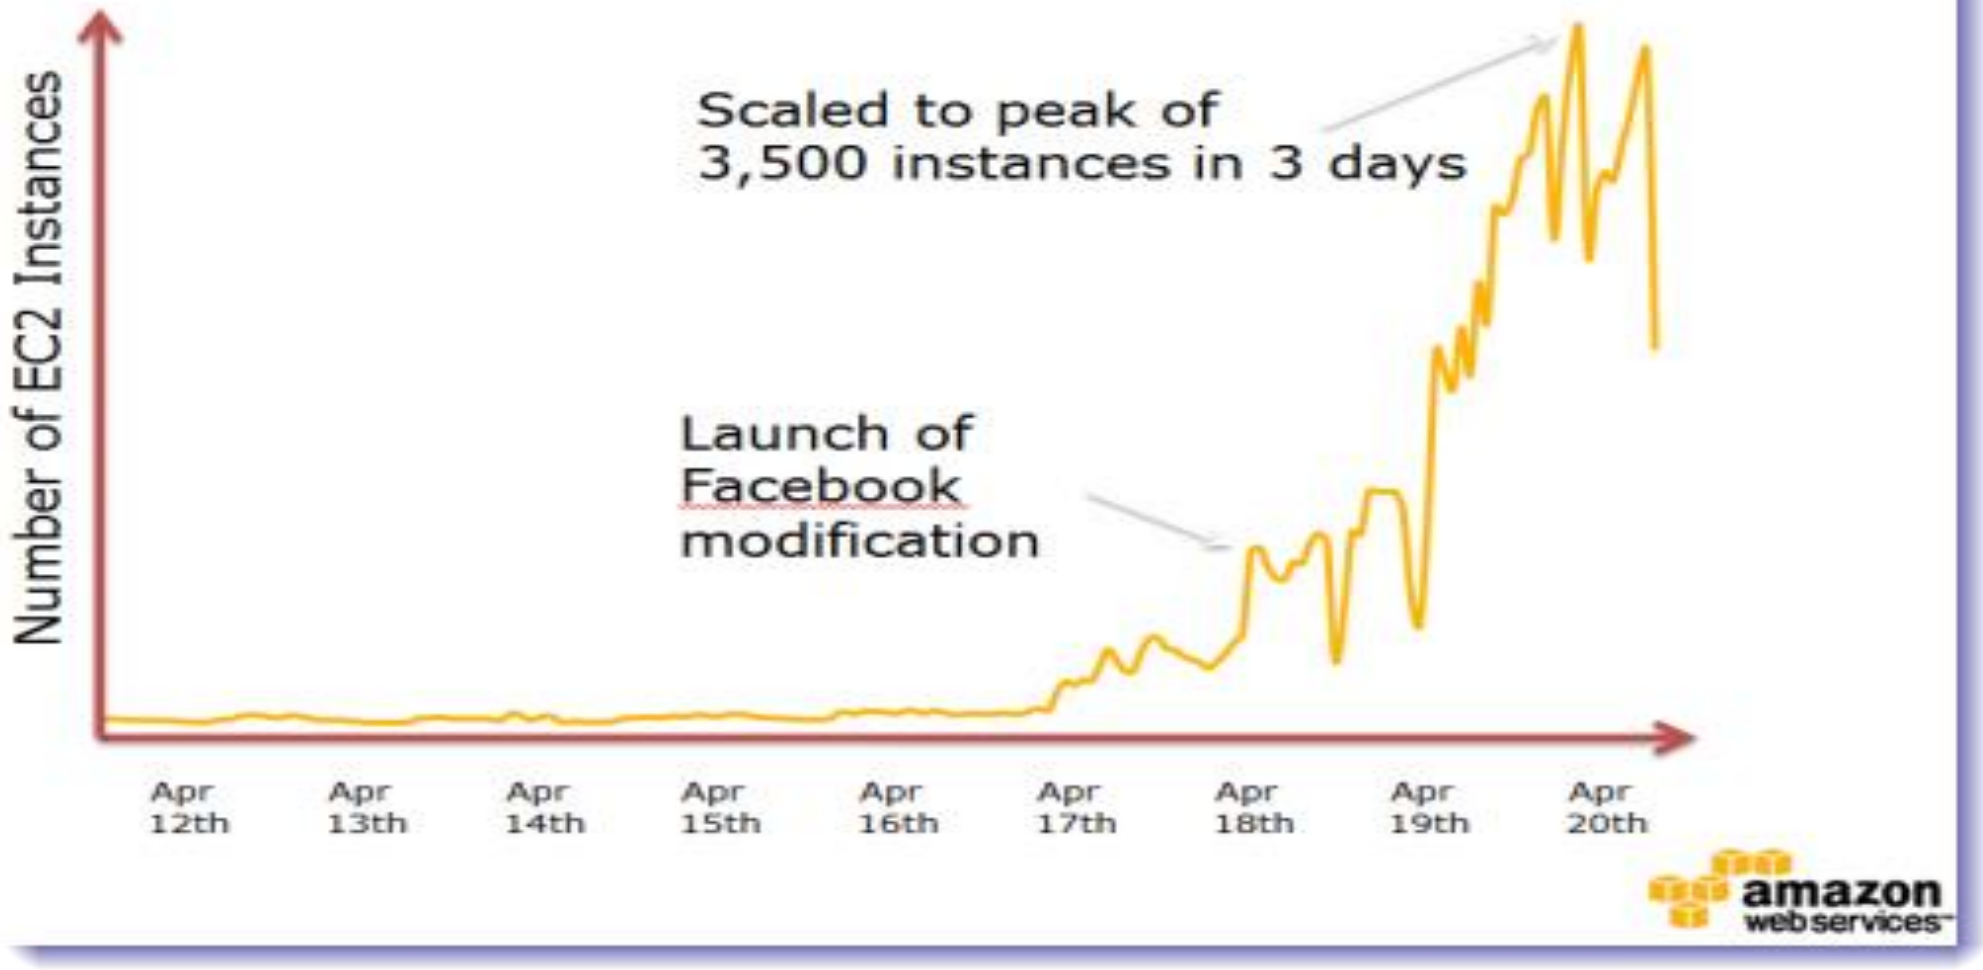

(Source: http://www.cca08.org/files/slides/w_vogel.pdf)

84

Zynga zCloud: Hybrid Self-Hosted/EC2

- After Zynga had large scale

- More efficient self-hosted servers
  - Run at high utilization

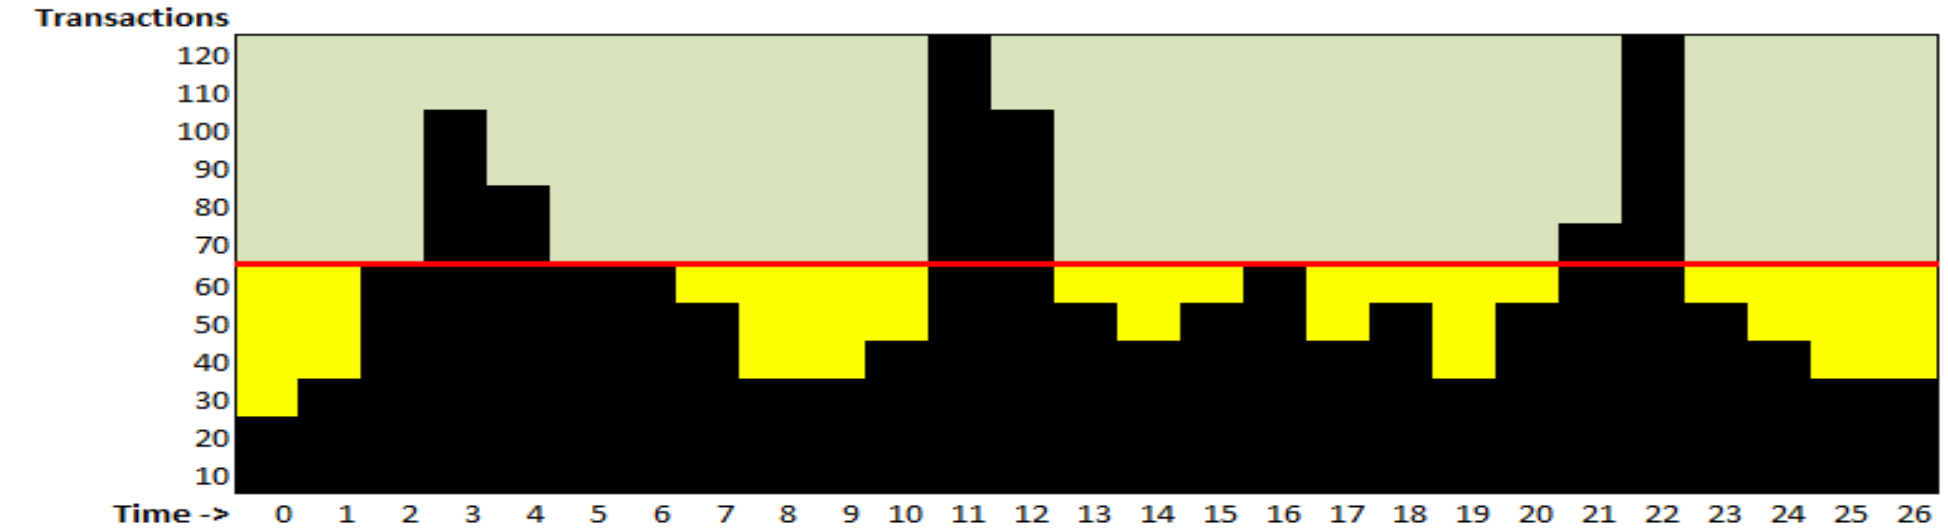

(Sources: http://seekingalpha.com/article/609141-how-amazon-s-aws-can-attract-ugly-economics and http://www.undertheradarblog.com/blog/3-reasons-zynga-is-moving-to-a-private-cloud/)

85

Other Cloud Customers

reddit

- 218 virtual CPUs
- 9TB/2TB block/S3 storage
- 6.5TB/2TB I/O per month

Dropbox

(Source: http://markbuhagiar.com/technical/businessinthecloud/)

86

Our Industry Collaborators

AZAVISTA



Take-Home Message

The Golden Age of Datacenters

Cloud computing = lease vs self-own

- On-Demand, Pay-per-Use, Elastic, Pooled, Automated, QoS
- Owner vs User perspective


Important New Challenges

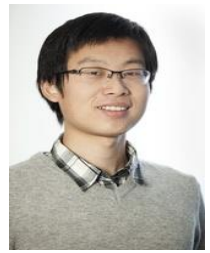
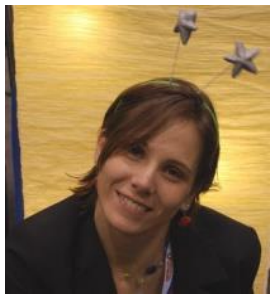

1. The scheduling challenge
2. The ecosystem navigation challenge
3. The big cake & 4. Jevons (Efficiency) challenges

Recommended Reading

Elastic Big Data and Computing

- B. Ghit, N. Yigitbasi (Intel Research Labs, Portland), A. Iosup, and D. Epema. Balanced Resource Allocations Across Multiple Dynamic MapReduce Clusters. SIGMETRICS 2014
- L. Fei, B. Ghit, A. Iosup, D. H. J. Epema: KOALA-C: A task allocator for integrated multicluster and multicloud environments. CLUSTER 2014: 57-65
- K. Deng, J. Song, K. Ren, A. Iosup: Exploring portfolio scheduling for long-term execution of scientific workloads in IaaS clouds. SC 2013: 55

Time-Based Analytics

- B. Ghit, M. Capota, T. Hegeman, J. Hidders, D. Epema, and A. Iosup. V for Vicissitude: The Challenge of Scaling Complex Big Data Workflows. Winners IEEE Scale Challenge 2014

- http://www.pds.ewi.tudelft.nl/epema
- http://www.st.ewi.tudelft.nl/~iosup/research.html
- http://www.st.ewi.tudelft.nl/~iosup/research_cloud.html
- http://www.pds.ewi.tudelft.nl/

Disclaimer: images used in this presentation obtained via Google Images.

- Images used in this lecture courtesy to many anonymous contributors to Google Images, and to Google Image Search.
- Many thanks!