



# Guided Tour through a Cloud Datacenter

- The Umeå University approach to cloud resource management

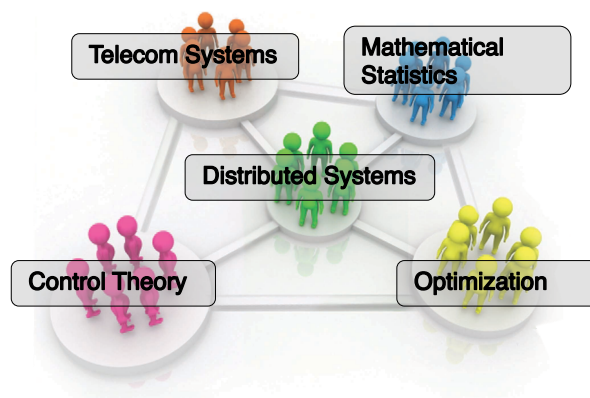
**Erik Elmroth**  
**Umeå University**  
**elmroth@cs.umu.se**

LCCC Focus Period / 4th Cloud Control Workshop  
Lund University,  
May 7-9, 2014

[www.cloudresearch.org](http://www.cloudresearch.org)



## Interdisciplinary collaborations



Images from <http://artsandhumanities.pressible.org>

Erik Elmroth, [elmroth@cs.umu.se](mailto:elmroth@cs.umu.se)

11

**Holistic Management**

**Admission Overbooking**

**Autoscaling**      **Placement**      **Knowledge Base**

**Guided Tour**

14

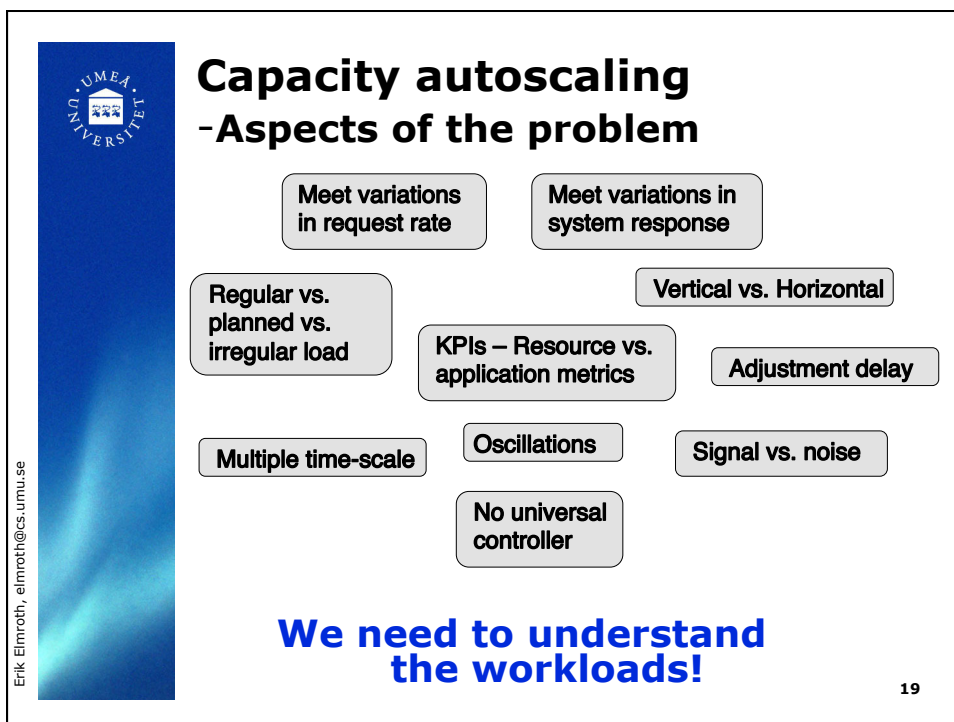
## Admission Control


- Determines load, revenue, and risks
- Risk theory
  - Utility versus Violations
  - Overbooking
  - Long term effects

L. Tomás and J. Tordsson, Cloudy with a Chance of Load Spikes: Admission Control with Fuzzy Risk Assessments, Proc of 6th IEEE/ACM International Conference on Utility and Cloud Computing, 2013

15

Erik Elmroth, elmroth@cs.umu.se

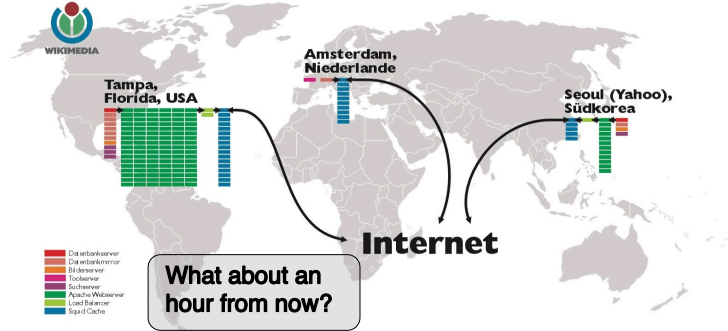




# Workload Analysis


What will your workload look like six years from now?

Wikipedia Pages:  
11.3 million (2008)  
29.6 million (2013)



A. Ali-Eldin, A. Rezaie, A. Mehta, S. Razroev, S. Sjöstedt-de Luna, O. Seleznev, J. Tordsson, and E. Elmroth, How will your workload look like in 6 years? Analyzing Wikimedia's workload. *Proceedings of the 2014 IEEE International Conference on Cloud Engineering (IC2E 2014)*, pp. 349-354, 2014. 20

# Workload Decomposition



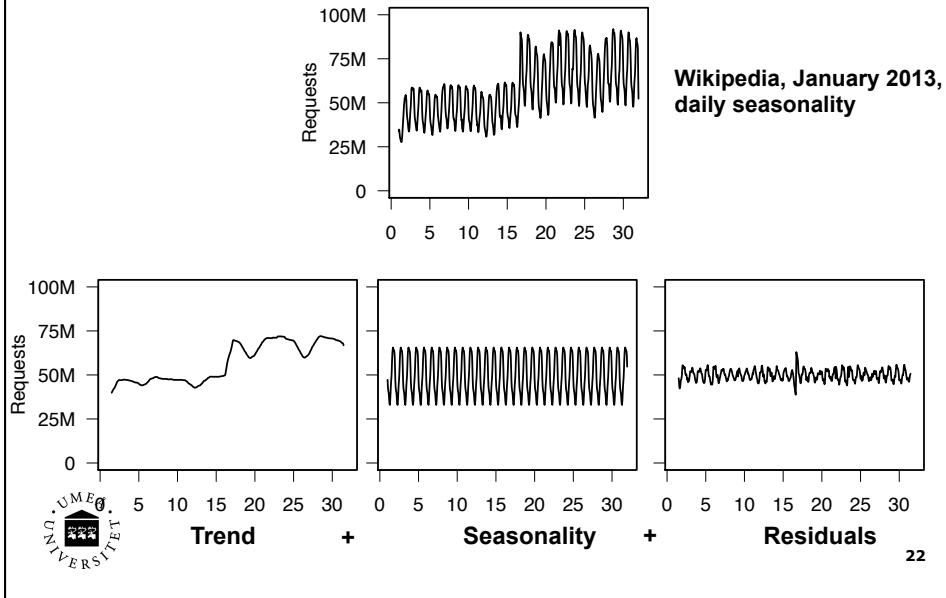
Wikipedia, January 2013



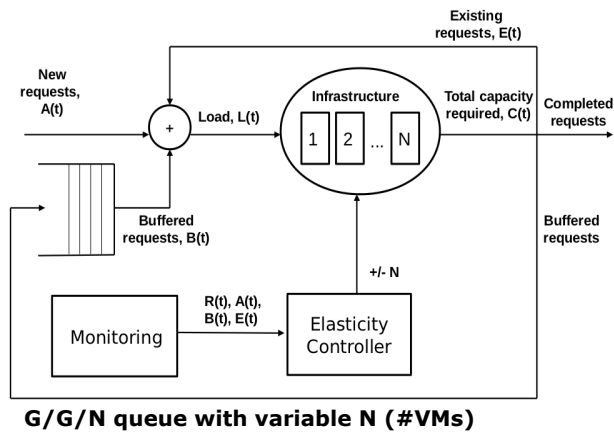
21



# Workload Decomposition




# Capacity Autoscaling Sample control theoretic model



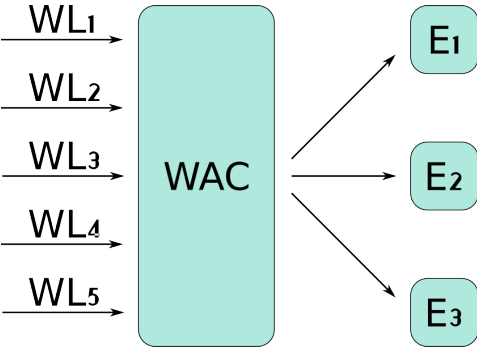
A. Ali-Eldin, M. Kihl, J. Tordsson, and E. Elmroth. Efficient Provisioning of Bursty Scientific Workloads on the Cloud Using Adaptive Elasticity Control, In Proceedings of the 3rd Workshop on Scientific Cloud Computing (ScienceCloud 2012), ACM New York, pp. 31-40, 2012.

A. Ali-Eldin, J. Tordsson, and E. Elmroth. An Adaptive Hybrid Elasticity Controller for Cloud Infrastructures, The 13th IEEE/IFIP Network Operations and Management Symposium (NOMS 2012), IEEE, pp. 204-212, 2012.


28




# Several Autoscaling Methods + Auto-selection



A. Ali-Eldin, J. Tordsson, M. Kihl, and E. Elmroth. Workload Classification for Efficient Cloud Infrastructure Elasticity Control, *submitted*, 2013.



# Scaling and Repacking



M. Sedaghat, F. Hernandez-Rodriguez, and E. Elmroth. A virtual machine re-packing approach to the horizontal vs. vertical elasticity trade-off for cloud auto scaling, *Proc. ACM Cloud and Autonomic Computing Conference (CAC)*, ACM, 2013




## VM placement

- Map VMs to resources
  - After admission
  - After scaling
  - To reconsolidate
- Across datacenters
  - e.g., linear programming problem
- Within datacenter
  - Load mixing
  - Multi-dimensional multi-knapsack problem

Erik Elmroth, elmroth@cs.umu.se

36



## Inter Cloud VM Placement

### Modeling (Cost Goals)

Minimize **Total cost**

$$TTP = \sum_{i=1}^l \sum_{j=1}^m \sum_{k=1}^n P_{ijk}$$

Subject to

**Capacity constraints**

$$\sum_{j=1}^m \sum_{k=1}^n x_{ijk} \leq C_{ijk} \quad \forall i \in [1..n]$$

$\forall i \in [1..n]$ :

$$\sum_{j=1}^m \sum_{k=1}^n x_{ijk} = 1 \quad (2)$$


$\forall k \in [1..m]$ :

**Load balance constraints**

$$\sum_{i=1}^l \sum_{j=1}^m x_{ijk} \leq C_{mjk} \quad (3)$$


W. Li, J. Tordsson, E. Elmroth. Modelling for Dynamic Cloud Scheduling via Migration of Virtual Machines, 2011 Third IEEE International Conference on Cloud Computing Technology and Science (Cloudcom 2011), IEEE Computer Society, pp. 163-171, 2011.

37



## Intra Datacenter Placement

- Workload mixing (time & space)
- Multi-dimensional, multi-knapsack
- Application Specific
- Heterogeneous hardware

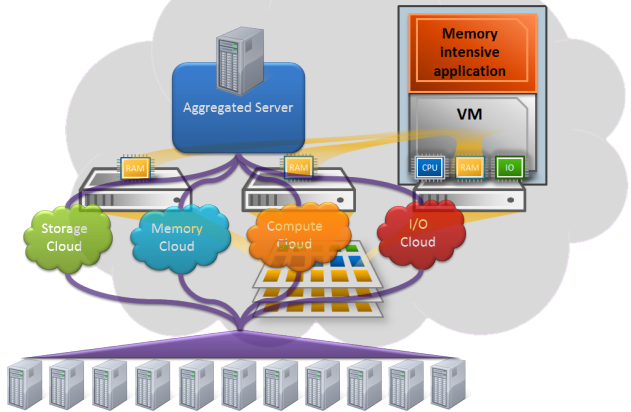


W. Li, J. Tordsson, and E. Elmroth. Virtual Machine Placement for Predictable and Time-Constrained Peak Loads, GECON 2011, Springer LNCS 7150, pp. 120-134, 2012.

L. Tomás and J. Tordsson, Cloudy with a Chance of Load Spikes: Admission Control with Fuzzy Risk Assessments, Proc of 6th IEEE/ACM International Conference on Utility and Cloud Computing, 2013

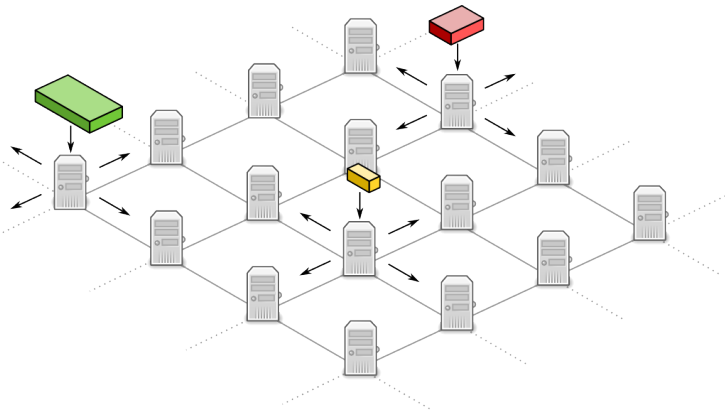
# Relaxed box model virtualization

For enhanced workload mixing (space)



P. Svärd, J. Tordsson, B. Hudzia, E. Elmroth. Hecatonchire: Enabling Multi-Host VMs by Resource Aggregation and Pooling. *Submitted, 2014* 44

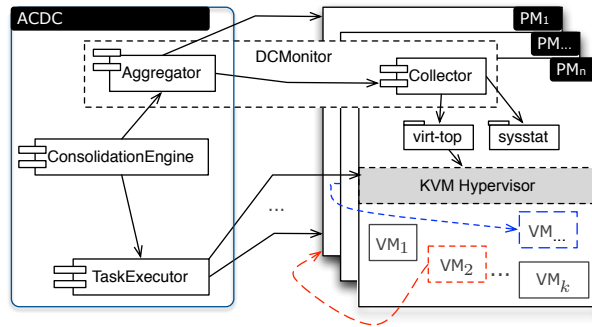
# Decentralized Placement



M. Sedaghat, F. Hernández, and E. Elmroth, Peer to peer resource management for cloud data centers, *Submitted, 2013* 48



## Datacenter Reconsolidation




- **Concerns**


- Optimal solution most likely infeasible
  - Gradual improvement
- Heuristic approach

## Live VM migration (without service interruption)

### Pre-copy migration




### Post-copy migration




	Pre	Post	Hybrid
Continuous service		(✓)	✓
Resource usage		✓	✓
Robustness	✓		
Predictability		✓	✓
Transparency	✓	✓	✓

P. Svård, J. Tordsson, E. Elmroth, B. Hudzia. The Noble Art of Live VM Migration - Principles and Performance, *submitted*, 2013.



## Aequus – Prioritization support

- Offers prioritization between competing potential utilizers
- Based on target – usage relation
- Priority applied hierarchically
- Decentralized system




P-O. Östberg and E. Elmroth. Decentralized Prioritization-Based Management Systems for Distributed Computing, *The 9th IEEE International Conference on e-Science (eScience 2013)*, pp. 228-237, 2013.

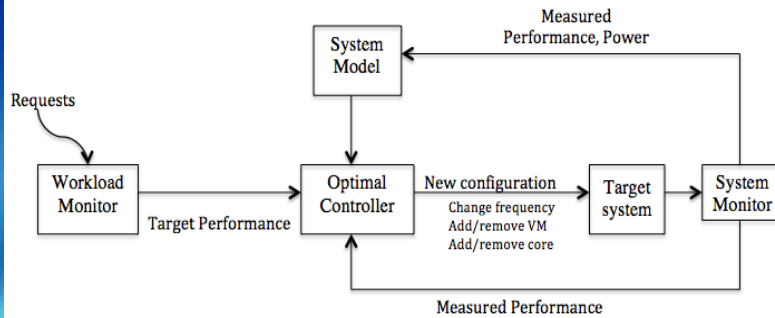
P-O Östberg, D. Espling, and E. Elmroth. Decentralized Scalable Fairshare Scheduling. *Future Generation Computer Systems*, Vol. 29, No. 1, pp. 130-143, 2013.

G. Rodrigo, P-O. Östberg and E. Elmroth. Priority Operators for Fairshare Scheduling. In *Proc. 18th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP 2014)*, Accepted, 2014.

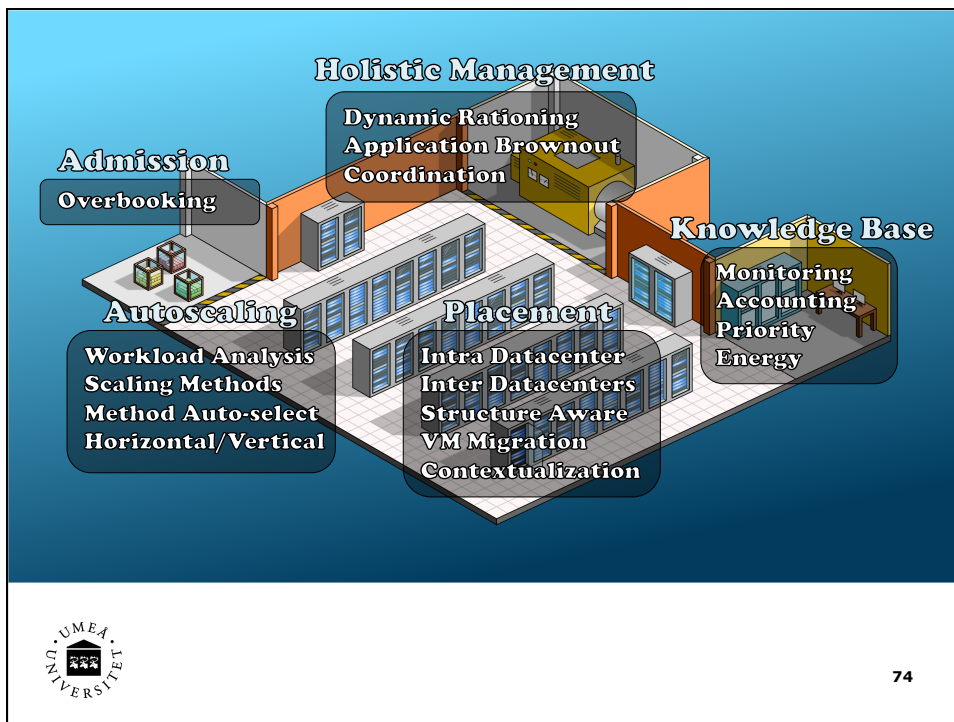





# Energy-efficient management



S. K. Tesfatsion, E. Wadbro, J. Tordsson, A Combined Frequency Scaling and Application Elasticity Approach for Energy-Efficient Clouds, *submitted*, 2013.





# Dynamic Resource Rationing

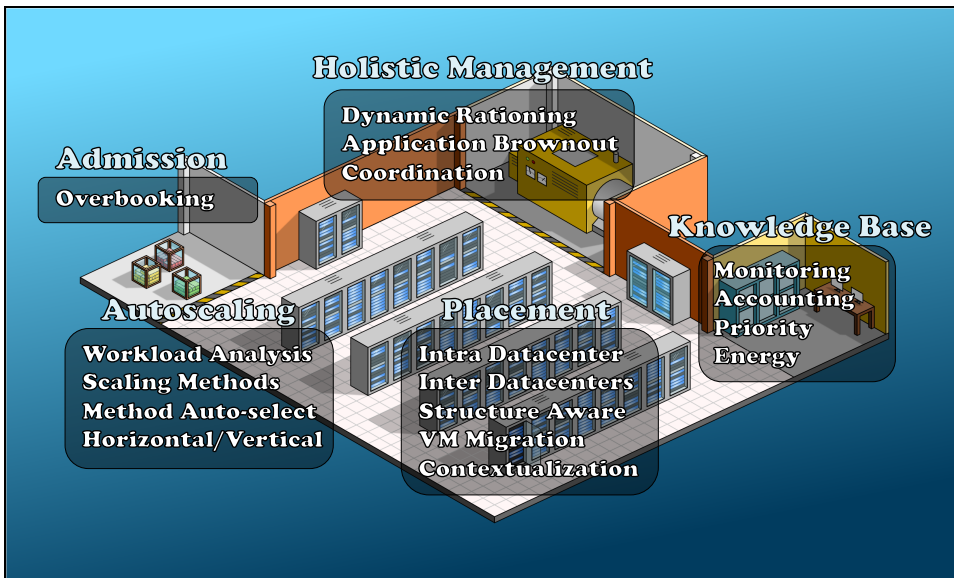

Where to cut when resources are insufficient?

Two approaches

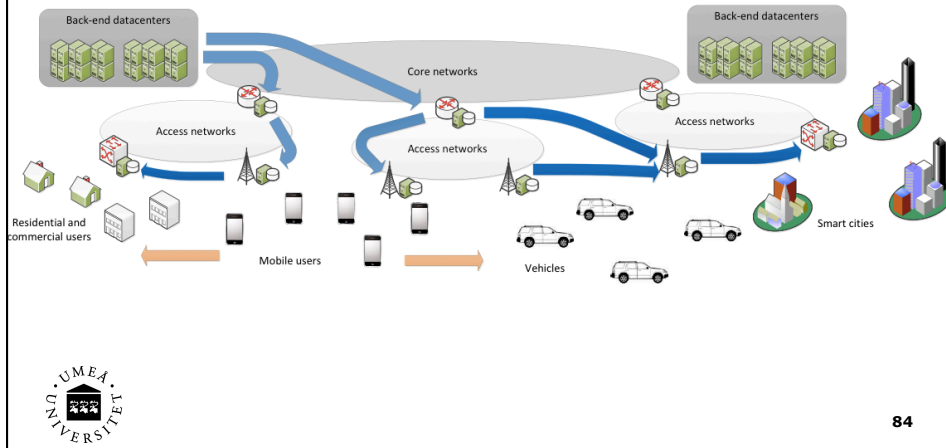
1. Strict QoS-level adherence
2. Overall cost-benefit with QoS-level weights

- Constrained optimization
- Substantial dependency on KPI-type (e.g. latency vs. throughput)
- System feedback on KPI and dimmer effect

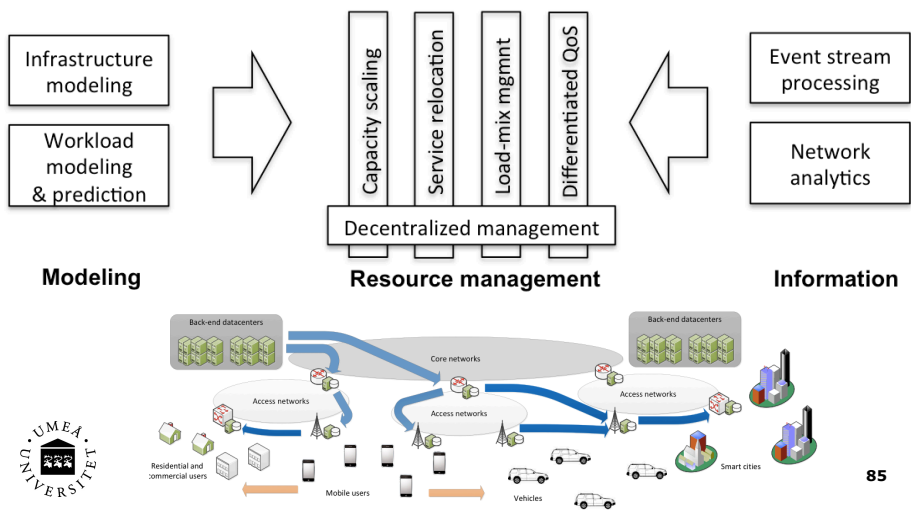
## System Architecture




# Managing the infinite (or telco) cloud



# Managing the infinite (or telco) cloud







## Large-scale Collaborations



EU FP7 IP. Introduced federated clouds. EU's first major cloud project. (Completed.)



EU FP7 IP. Optimized cloud services over complete lifecycle. Non-functional aspects. (Completed.)



EU FP7 IP. Pioneering federated storage clouds. Raised level of abstraction. Media- and telecom applications. (Completed.)



Governments strategic efforts. Methods and software for eScience applications.



Umeå initiative for innovation and industry benefits within simulation, visualisation, computation and infrastructure.



Swedish Research Council framework project. A control theoretic approach to cloud management.



EU FP7 STREP. Context-Aware Cloud Topology Optimisation and Simulation.




EU FP7 STREP. Business continuity through fault management extending on virtual consolidation.



Vinnova VINNVÄXT. Innovation project, focused on resource efficiency and green ICT for datacenters.

**Key partners:** IBM Haifa Research Labs, SAP, ATOS Orgin, Univ. Compl. de Madrid, Leeds Univ, Barcelona SC, Telefonica, British Telecom, Uppsala Univ., Luleå Univ. Tech., KTH, Lund Univ., etc

86



### Senior researchers



Erik Elmroth, Professor  
Francisco Hernandez, Assistant Professor  
Johan Tordsson, Assistant Professor  
P-O Östberg, Researcher

### Post docs



Cristian Klein, PhD  
Luis Tomás, PhD  
Daniel Esping, PhD

### PhD students



Ahmed Ali-Eidin, PhLic  
Jakub Krzywda  
Ewnetu Bayuh Lakew, PhLic  
Wubin Li, PhLic  
Mina Sedaghat, PhLic  
Petter Svärd, PhLic  
Kosten Selome Tesfatsion  
Gonzalo Rodrigo

### PhD students



Amardeep Mehta  
Olumuyiwa Ibidunmoye

### Others



Peter Gardfjäll, Sys. Dev, PhLic  
Lars Larsson, System developer  
Lennart Edblom, Senior lecturer  
Tomas Forsman, Systems expert  
Niclas Lockner, Research assistant

87

[www.cloudresearch.org](http://www.cloudresearch.org)

