



**The RESERVOIR Cloud Computing Infrastructure  
for On-demand Service Provisioning**

[www.reservoir-fp7.eu](http://www.reservoir-fp7.eu)



**Erik Elmroth**  
HPC2N & Dept. Computing Science  
Umeå University, Sweden  
elmroth@cs.umu.se

The research leading to these results has been partially funded by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°215605.



**Brief outline**

- A Game changing trend in IT use
- Revitalization of the datacenters
  - Infrastructure providers for service providers
- Compute Clouds
  - datacenter infrastructure providing virtual resources as utility
- The RESERVOIR approach – business-aware federated Clouds
  - Architecture & functionality
  - Scenarios illustrating the RESERVOIR capabilities
  - Research challenges

EchoGRID Conference 2008 Shenzhen



## A Game Changing Trend - Growth on Service Consumer Side

- Individuals use internet-based services such as YouTube, MySpace, Google documents, Adobe Photoshop Express, etc, for managing their private and professional everyday activities
 


- Companies use external services such as hosted Microsoft Exchange, external mail services, external customer relations management, accounting systems, or hosting of their complete IT environments
 




- Explosive growths in availability of Software-as-a-Service (SaaS) (and Infrastructure-as-a-Service (IaaS), Everything-as-a-Service (XaaS), ...)

EchoGRID Conference 2008 Shenzhen



## Crucial capacity characteristics – to be cost-efficiently met

Extremely rapid growths (from global scale)

- MySpace only needed 36 months to reach 100 million users (now per day: 300 000 new users, 65 billion page views)
- YouTube reached 20 million users within 16 months
- App Store: Over 1000 Iphone applications. Over 160 million sold

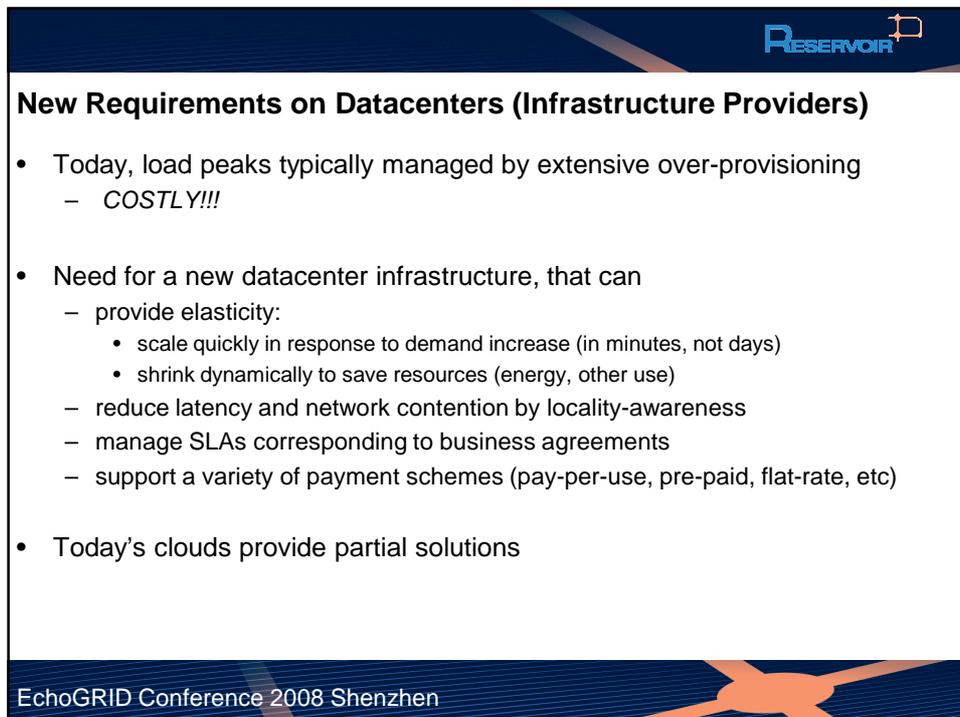
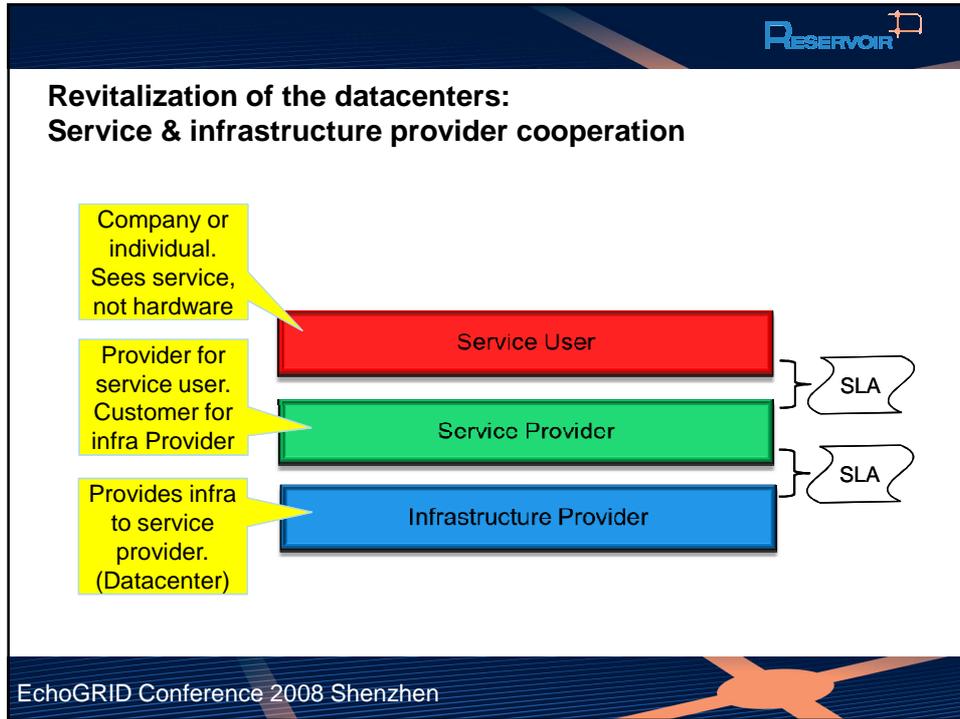
Regular/planned peaks

- banks see a substantial peak over a few days around the 20th every month
- on-line tax filing - exceptional load a few days per year
- meet a rapid increase of use, e.g., after a marketing campaign

Unexpected peaks

- News-related video streaming
- Stock trading peaks at financial crises

EchoGRID Conference 2008 Shenzhen





## Compute Clouds

- Virtual “cloud” of IT resources (within a data center)
- Services run on virtual resources, unaware of the physical resources
- Infrastructure – compute, storage, and network
- Utility model – provision on demand, charge back on use

**Compute Cloud – Amazon EC2 & S3, pay-per-use**

- \$0.10 per instance-hour consumed (or part of an hour consumed)
- \$0.15 per GB-Month of Amazon S3 storage used
- \$0.10 per GB of data transferred outside of Amazon

EchoGRID Conference 2008 Shenzhen




## The RESERVOIR Vision

The Next Generation Infrastructure for Service Delivery

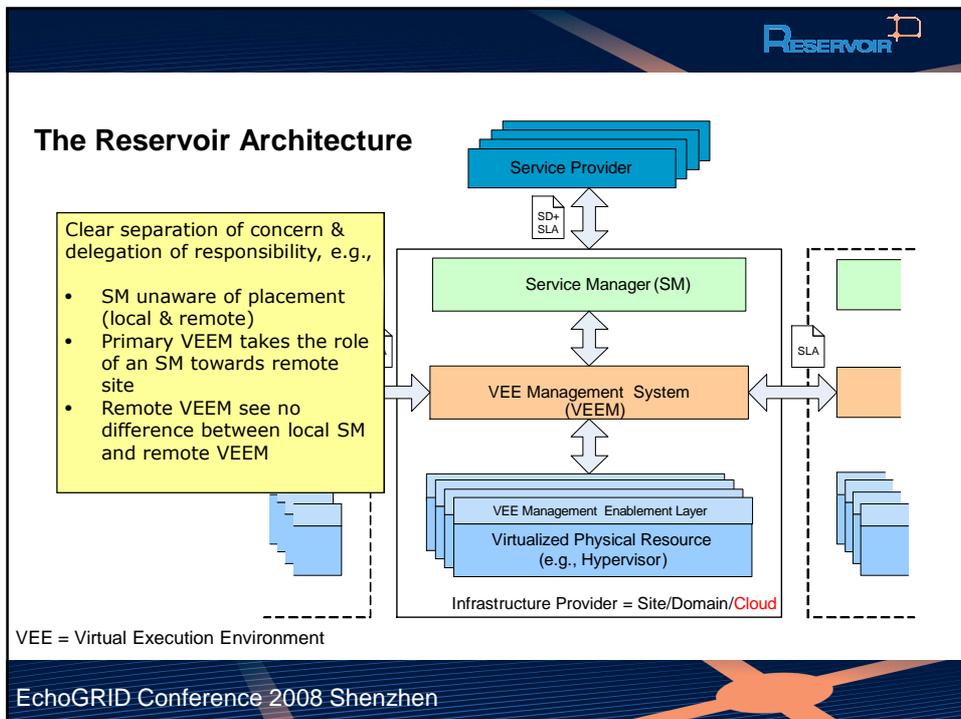
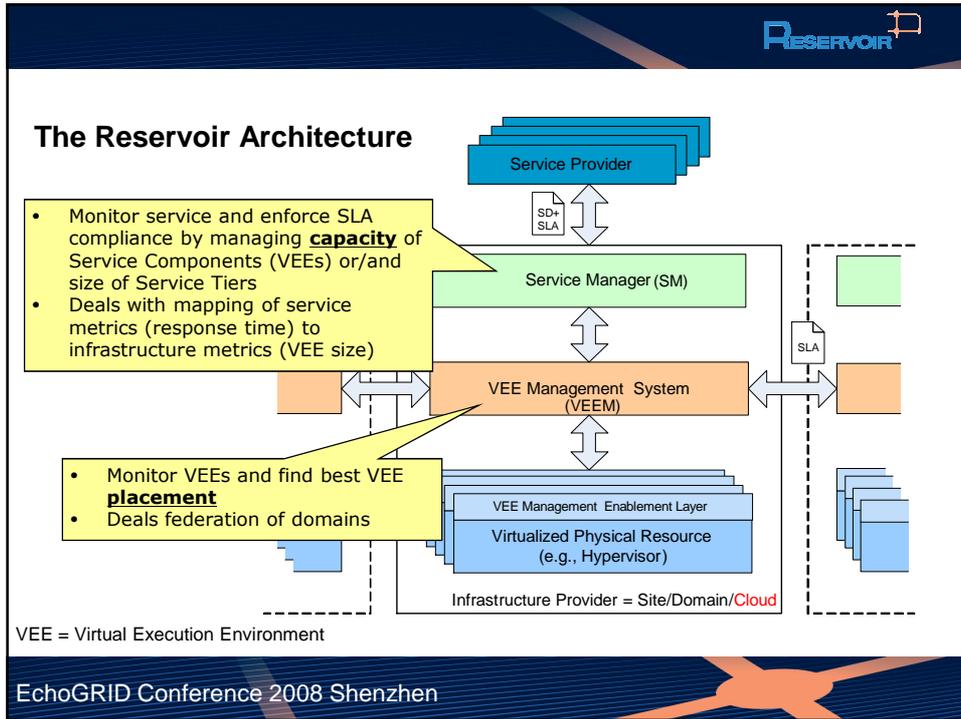
- Federation of clouds
- Leverage locality – enable migration
- Service definition, SLA management, accounting and billing

Analogies exists in areas outside services:

- Electrical power delivery: capacity can be shifted to guarantee supply and lower costs
- Roaming cellular communications: Talk wherever you are

<b>Grid-aware Virtualization</b> <small>Service-oriented capacity provisioning across sites</small>	+	<b>Virtualization- aware Grid</b> <small>Optimal placement of VMs on a federated cloud</small>	+	<b>Business &amp; Service Management</b> <small>Policy-based management of service-level agreements</small>	= SOI
--	---	---	---	--	-------

EchoGRID Conference 2008 Shenzhen



**RESERVOIR**

### Service Applications on Reservoir

- One multi-VEE application on:
  - One VEE host
  - Multiple VEE hosts
  - Multiple sites

Reservoir Site 1    Reservoir Site 2

EchoGRID Conference 2008 Shenzhen

**RESERVOIR**

### SOI: Grid Computing

Grid node or Service Site (datacenter)

Physical Resources

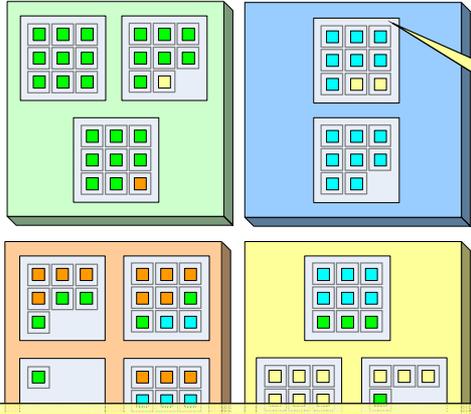
Service Tasks

Color of service task illustrates owning organization

EchoGRID Conference 2008 Shenzhen



### SOI: Grid Computing + Virtualization



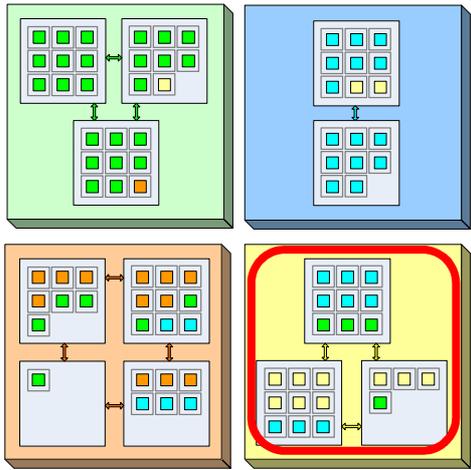
Virtual Execution Environment (VEE)

Improved isolation, Relax dependencies, Well defined billing units

EchoGRID Conference 2008 Shenzhen



### SOI: Grid Computing + Virtualization + BSM

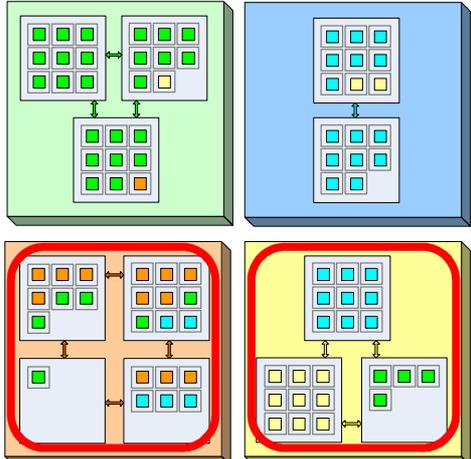


Policy 1:  
If possible keep VEEs from the same organization in the same physical box

EchoGRID Conference 2008 Shenzhen



### SOI: Grid Computing + Virtualization + BSM



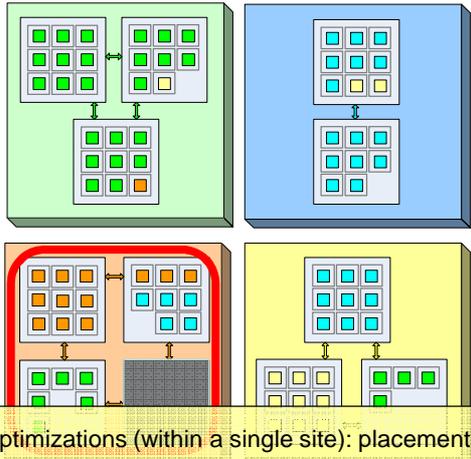
Policy 1:  
If possible keep  
VEEs from the  
same organization  
in the same  
physical box

Policy 2:  
Turn off  
underutilized  
physical boxes

EchoGRID Conference 2008 Shenzhen



### SOI: Grid Computing + Virtualization + BSM



Policy 1:  
If possible keep  
VEEs from the  
same organization  
in the same  
physical box

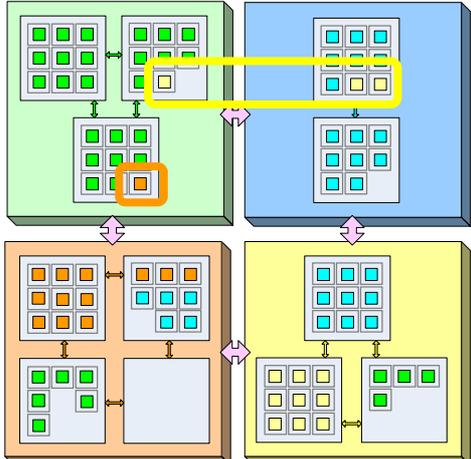
Policy 2:  
Turn off  
underutilized  
physical boxes

Local optimizations (within a single site): placement, power, etc.

EchoGRID Conference 2008 Shenzhen



### SOI: Grid Computing + Virtualization + BSM – Boundaries

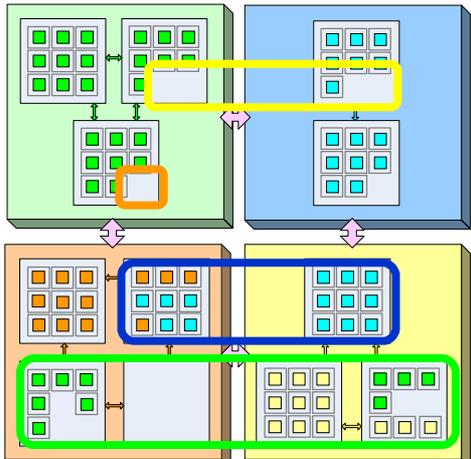


Policy 3:  
If possible keep  
VEEs in "owning"  
organization

EchoGRID Conference 2008 Shenzhen



### SOI: Grid Computing + Virtualization + BSM – Boundaries



Policy 3:  
If possible keep  
VEEs in "owning"  
organization

Policy 4:  
If possible keep  
VEEs in least  
number of  
external  
organizations

EchoGRID Conference 2008 Shenzhen

### SOI: Grid Computing + Virtualization + BSM – Boundaries

Policy 3:  
If possible keep VEEs in “owning” organization

Policy 4:  
If possible keep VEEs in least number of external organizations

EchoGRID Conference 2008 Shenzhen

### SOI: Grid Computing + Virtualization + BSM – Boundaries

Policy 5:  
“Follow” the service customer

Migration across sites → **Global optimizations**: placement, cost, bandwidth, etc.

EchoGRID Conference 2008 Shenzhen

**RESERVOIR**

### Virtualize the Network

Create virtual networks connecting VEEs regardless of physical server location

EchoGRID Conference 2008 Shenzhen

**RESERVOIR**

### Virtualize the Network and the Storage

Enable secure access to relevant data regardless of storage location

EchoGRID Conference 2008 Shenzhen





### Sample Research Challenges

- Overall architectural issues
  - Separation of concern, delegation of responsibility
  - Minimize inter-component dependencies while preserving degrees of freedom for optimizing functionality
  - Information models
  - Policy infrastructure

### Sample Service Level Challenges

- Translate business concept requirements to infrastructure requirements
  - E.g., response time to CPU utilization
  - Define a Service Definition Language to characterize all information and context required to enable lifecycle management of services across RESERVOIR sites
  - Must be able to handle rollback on deployment failures
- SLA definitions
- Support multiple levels of QoS

EchoGRID Conference 2008 Shenzhen



### Sample Management Challenges

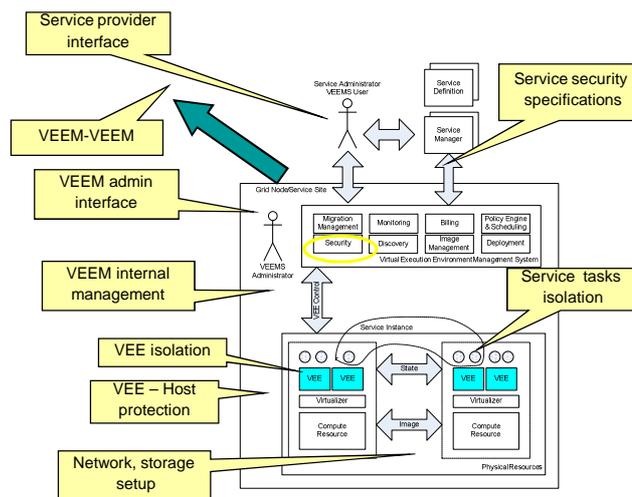
- Support policy based management across administrative domains (clouds)
  - Automatically hire additional “power” from another clouds (based on pre-existing framework agreements about guaranteed capacity or best effort dynamic requests)
- Create an inter-site protocol to allow for federation of RESERVOIR sites
- Protect Service Level Agreements
  - Detect violations (SLA monitoring)
  - Predict and avoid SLA violations
  - Elasticity rules (definition, execution)
  - Provide for dynamic relocation of resources
  - Provide accountability
- Bill for services used, even across RESERVOIR sites
  - Different billing and accounting systems may be used

EchoGRID Conference 2008 Shenzhen

### Sample Infrastructure-level Challenges

- Provide for relocation of resources without boundaries
  - Live migration across subnet boundaries
  - Migration to a different physical host without shared storage
- Provide standardized interfaces for lifecycle management to Virtualized Execution Environment
- Analyze end-to-end performance in a virtualized environment to understand bottlenecks
- Be able to handle surges in 3-5 orders of magnitude in service requests
- Provide cost-efficient VEE placement
  - "Optimal" local placement decisions
  - "Optimal" remote placement decisions
  - Local vs. remote placement
  - Note differences compared to job scheduling – services typically have "infinite" lifetime and varying resource utilization

### Security hot spots in RESERVOIR





## The RESERVOIR Project – EU FP7 Integrated Project (IP)

- 3-year project, started February 2008 (17 M Euro budget)

Partner	Role	Comment
1. IBM, Israel	Technology	Virtualization/SOC Infrastructure. Project coordinator.
2. Telefonica I+D, Spain	Technology	Service Technology, Billing Infrastructure
3. UC London, UK	Technology	Virtualization Technology
4. Umeå Univ., Sweden	Technology	Grid technology, Resource Management, Accounting
5. SAP AG, Germany	Use-Cases	Use-Cases, Contribution to Requirement, Standards
6. Thales, France	Technology	Security, Virtualization Infrastructure, Hosting
7. Sun Microsystems	Use-Cases+Tech	Contribution to Standards, Java Services, Monitoring
8. DATAMAT	Technology	Service Management Technologies
9. UC Madrid, Spain	Technology	Grid, Dynamic Allocation Technology
10. CETIC, Belgium	Technology	Security
11. Univ Logano, Switzerl.	Technology	Monitoring and SLA Management
12. Univ. Messina, Italy	Technology	Grid Experience, Testbed Development
13. Open Grid Forum, UK	Technology	Standardization

[www.reservoir-fp7.eu](http://www.reservoir-fp7.eu)

EchoGRID Conference 2008 Shenzhen



## The Grid computing group at Umeå University

Erik Elmroth (Associate Professor)\*  
 Bo Kågström (Professor)  
 Francisco Hernandez (PhD, Postdoc)\*  
 Daniel Henriksson (PhD Student)\*  
 Lars Larsson (Software Developer)\*  
 Johan Tordsson (PhD Student)\*  
 P-O Östberg (PhD Student)  
 Raphaela Bieber (Systems Developer)  
 Mats Nylén (PhD, SweGrid project lead)  
 Roger Oscarsson (EGEE)  
 Åke Sandgren (SweGrid)  
 Mattias Wadenstein (NDGF)

**Other major topics:**

- SweGrid Accounting System (SGAS)
- Grid-wide fairshare scheduling (FSGrid)
- Grid Job Management Framework (GJMF)
- Job Submission Service (JSS)
- Grid Workflow Execution Engine (GWEE)
- Resource brokering
- Co-allocation
- Advance reservation (WS-Agreement)
- Interoperability, portability, SOA
- Some applications, portals, etc

\* = Involved in RESERVOIR

**GIRD**  
 Grid Infrastructure Research & Development  
[www.gird.se](http://www.gird.se)

EchoGRID Conference 2008 Shenzhen