



Service level management in e-Infrastructures

the challenge of co-evolving, cross-
organisational user and provider groups

Matti Heikkurinen, Emergence Tech Ltd

Breaking down the title



- **Goal: service level management**
 - Organisational processes ensuring that Service Level Agreements (SLAs) are robust
 - In plain English: all parties know what to expect and do - also when things go wrong
- **First challenge: co-evolution**
 - Users and providers are re-organising their processes
 - New users and use cases emerge, services evolve
- **Second challenge: cross-organisational groups**
 - “What to expect” hard to know
 - Especially when things go wrong



- 1654: Regulation of Hackney-Coachmen in London
 - First policy-level SLM framework?
 - SLA and QoS issues related to taxis: availability, knowledge of the area, minimum size of the horse
- Later developments
 - ITIL: define key (organisational) processes, overview of best practices
 - MOF: templates and job aids
 - ISO/IEC 20000: minimum requirements, checklists
 - eTOM: process and information models
- All focus on organisational/procedural level
 - Independent of the technical aspects of SLA monitoring etc.

Challenge 1: Co-evolution

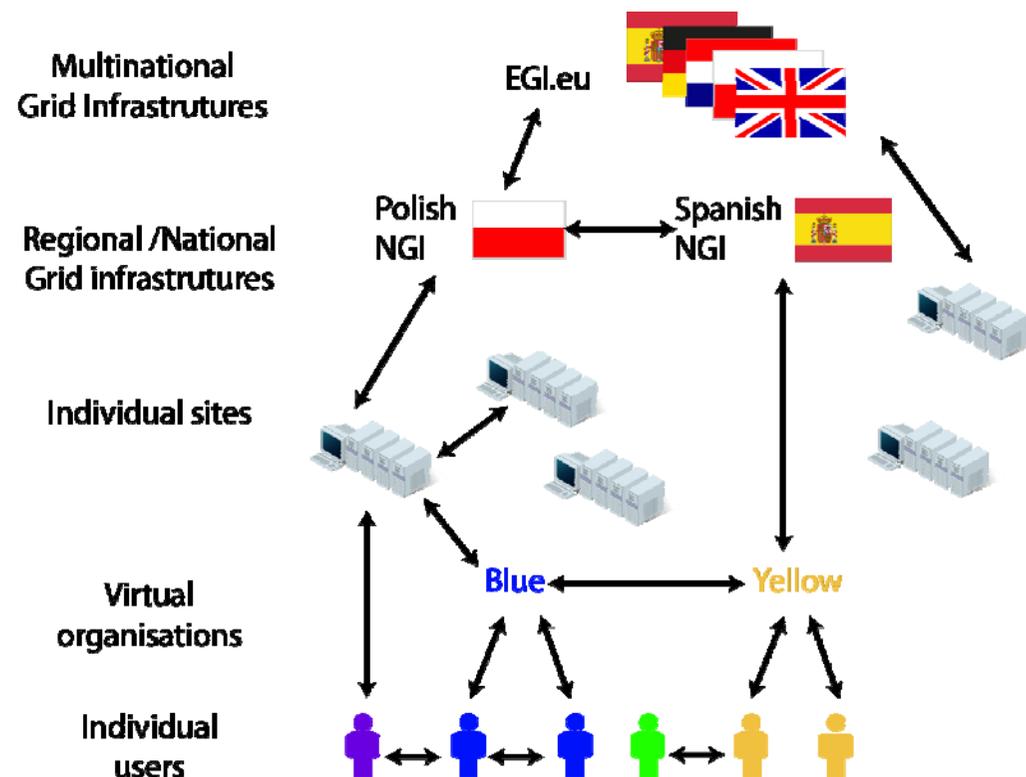


- Happens everywhere, multiple triggers
 - New HPC hardware -> code rewrite
 - New RI -> new computing challenge
 - New project -> new data models (often cross-disciplinary)
 - New e-Infrastructure service -> new opportunities, surprising use cases

Why is this a challenge?



- e-Infrastructure today can provide acceptable level of service
- Often based on informal, “personal” links
- However: model does not scale
 - Number of links becomes unmanageable
 - More use -> oversubscription and unaccounted costs become (politically) impossible

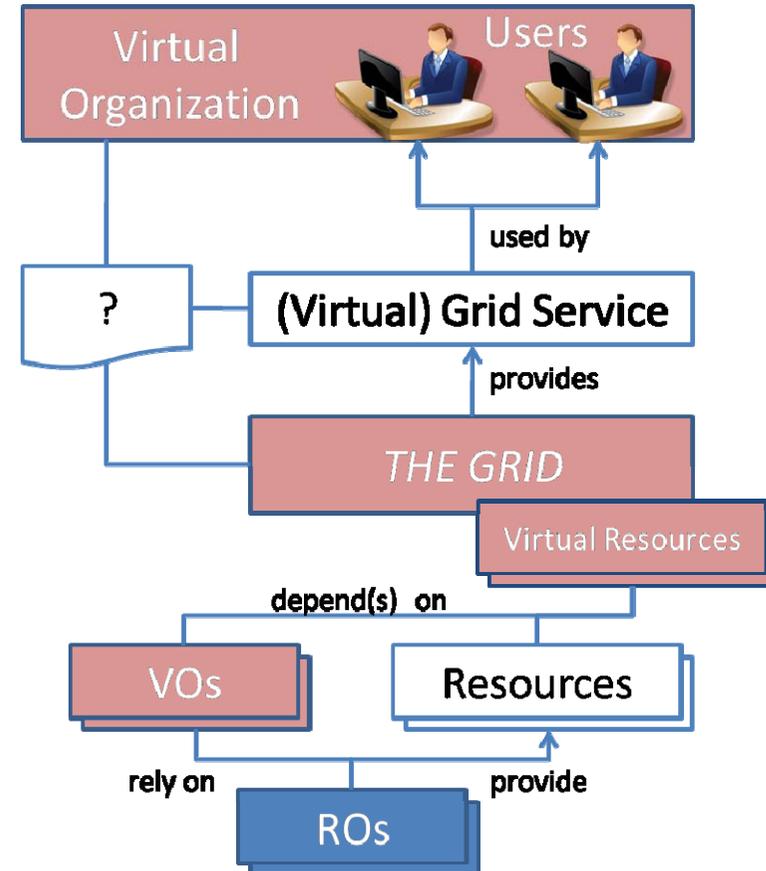
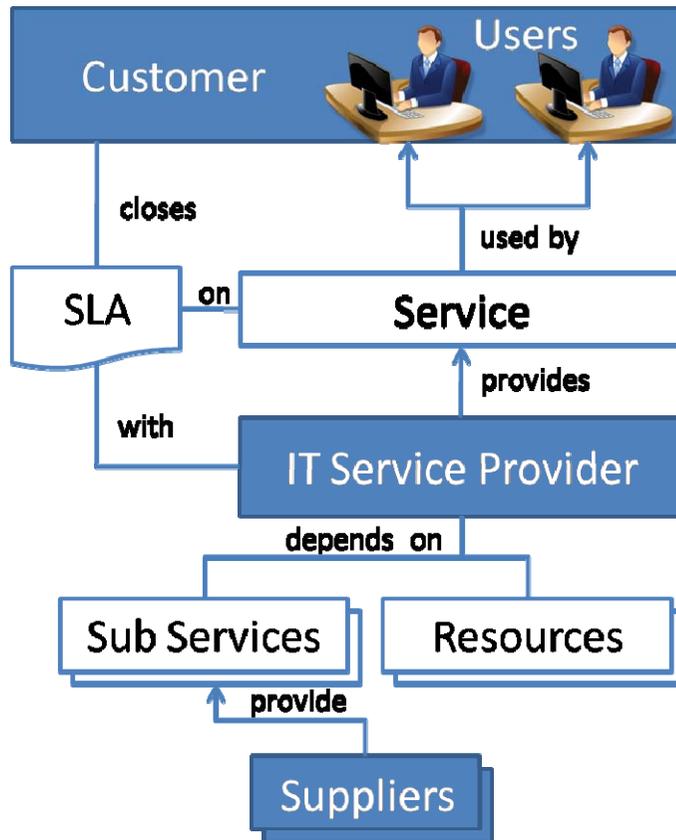


Challenge 2: cross-organisational



- It is unavoidable
 - Modus operandi of scientific discovery
 - Capacity/capability limits of single provider
 - Emerging: awareness of vendor lock-in
- It is messy
 - Breaks underlying assumptions of traditional AAI and ITSM
- It (almost) guarantees misunderstandings
 - No *lingua franca* to describe needs and services
 - Same data format, different semantics
 - Starting from metric vs. imperial (Mars Climate Orbiter)

Cross-organisational issues (provider side)



What about clouds?



- **Co-evolution**
 - Users and providers are still re-organising their processes
 - More of the responsibility for supporting them is shifted to users and applications (by default)
 - Economies of scale on CPU-hour level?
 - Fragmentation of knowledge on higher levels?
- **Cross-organisational groups continue to exist**
 - “What to expect” still hard to know
 - Especially when things go wrong
- **Service level management**
 - All parties know what to expect - on # CPU hour per € level
 - Application level SLM at least as hard as with Grids?

- gSLM mapping and modeling exercise



- Scaling up the e-Infrastructures requires formal SLM approach
 - Reliability, predictability, cost-effectiveness, managing customer relationships,...
- Approach needs to incorporate best practices that have made e-Infrastructure successful
 - Formalise and refine cross-organisational SLM, do not re-invent it!
- gSLM wants to act as a facilitator
 - “SLM people wo speak Grid” and vice versa

Concrete project activities

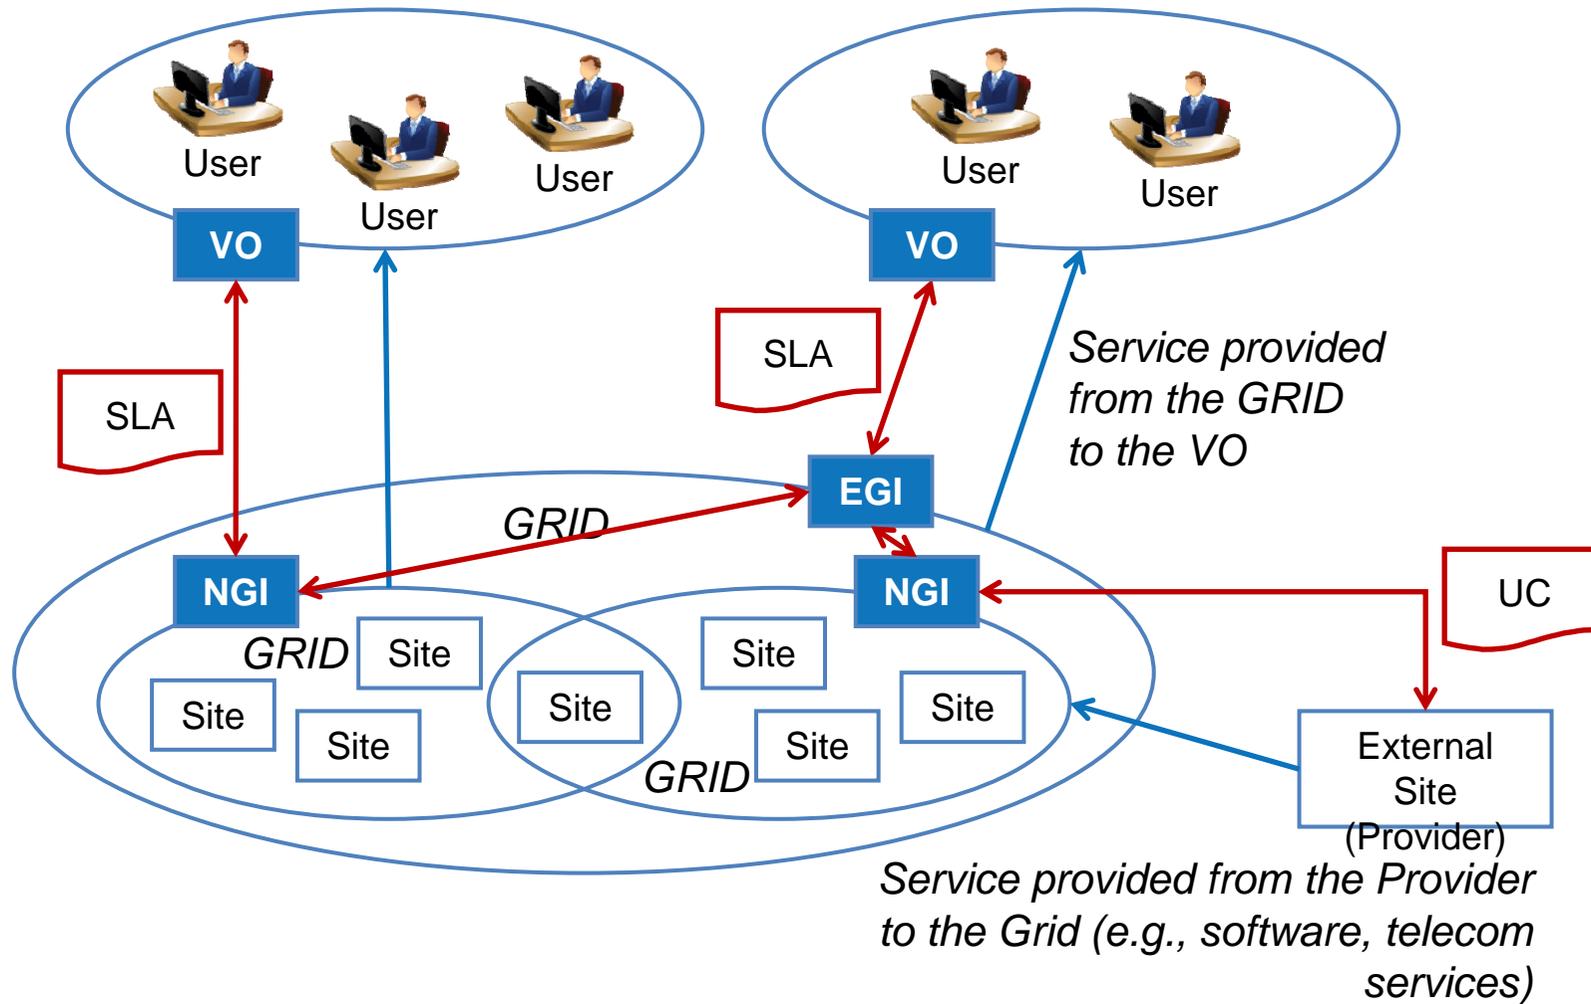


- Propose common model
 - Communication tool (*lingua franca* beta)
 - Best practices, e-Infrastructure SLM use cases,...
 - Distilled from analysis, workshops, consultations
 - Focus on Grids, but aim at broader relevance
- Documents aiming at making
 - e-Infrastructures understandable to SLM professionals
 - SLM relevant to e-Infrastructures
- Community building
 - Workshops
 - Online collaboration

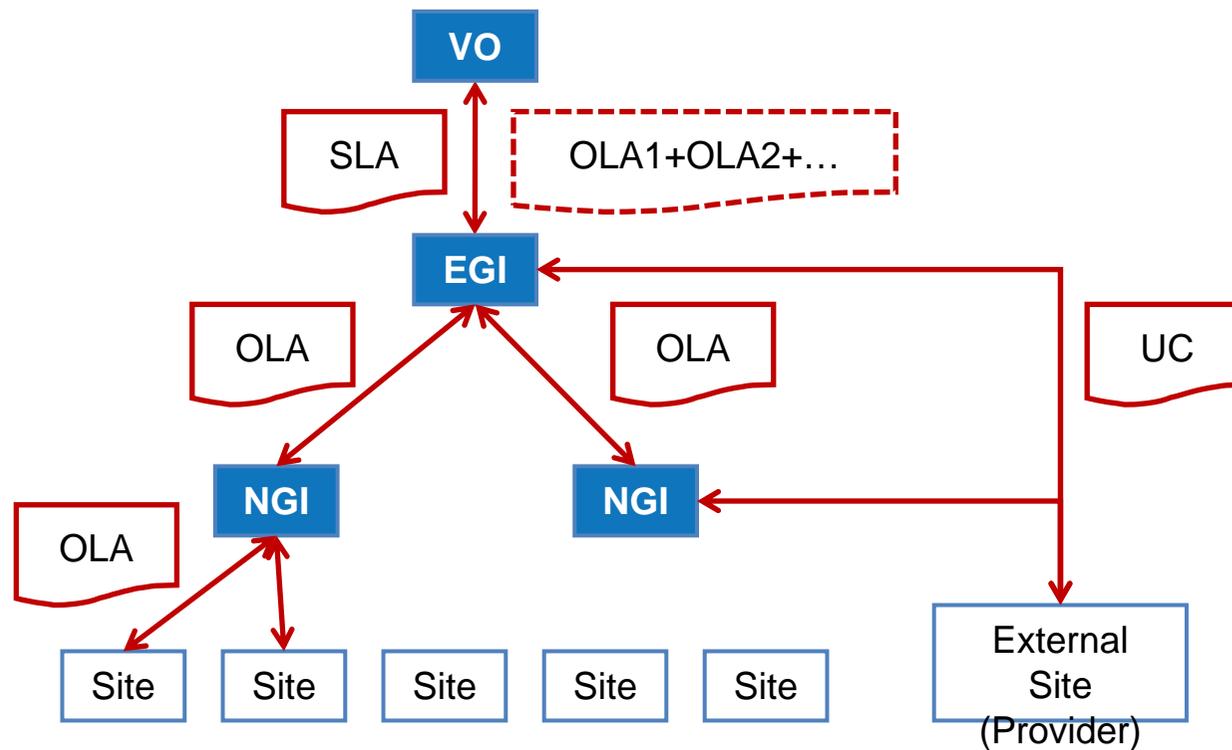


- e-Infrastructure
 - Federated resource
 - Network, HPC or Grid - or a resource on top of them (e.g. data)
- User community
 - Identifiable group of collaborating individuals (usually represented by a VO)
- Service
 - This is where things get interesting
 - E.g. is Cloud
 - A core in someone else's datacenter?
 - Turnkey solution ensuring the user has the computing capacity and access to it?

Conceptual model, idealised Grid



Idealised Grid relationships



Basic axioms



- The purpose of any Grid is to provide resources as a service to the subscribed VOs
 - Typical examples include storage and computational capacities
- The VO is the "customer" of a Grid Initiative
 - A VO may register itself with a Grid Initiative (GI)
 - Becomes a customer upon approval
- GI represents resource owners (sites)
 - Different forms of mandate

Relationship Assumptions



- A VO may register with one (or more) NGI(s) and/or EGI
- If a VO is registered with EGI, it is "known" by the NGIs as a customer of the EGI(!)
- The members of a VO are the users of the Grid.
- SLAs are agreed between VOs and GIs.
- OLAs are agreed between
 - EGI and NGIs
 - NGIs and Sites
- Underpinning Contracts are closed between ...
 - EGI and External Providers
 - NGIs and External Providers
 - Sites and External Providers

Intent assumptions



- Underpinning Contracts are established with the aim of making a Grid infrastructure operational
 - E.g. deliver cores or bandwidth to site or as a service
- SLAs are established to make utility and warranty commitments to VOs
 - Utility: service fits its purpose (technical, functional)
 - Warranty: attributes/commitments to making service fit for use?
- OLAs are established with the aim of supporting the delivery of services through the Grid to one or more VOs.
 - The OLA framework within a GI: supports the fulfillment of SLAs between the GI and its VOs
 - OLAs between VO and GI: general and/or preparatory basis for establishing new services/SLAs



- SLAs and OLAs are subject to changes for different reasons.
 - Degree of formalism
- The GI is a Single Point of Contact for a VO, representing the Grid as a whole.
- The added value of a GI may range from a simple aggregation (GI as "mediator") to full integration (GI as "service provider") of the underlying resources.

Next steps



- BDIM workshop, May 27th
 - Dublin, Ireland
- Public consultation of gSLM documents
 - Mailing lists being setup
- Live documents planned
 - Public wiki coming soon

Where do we need help?



- Broaden the pool of expertise
 - Tacit knowledge about running services
 - Experiences from user communities
 - Especially ones migrating to new e-Infrastructures
 - Cross-organisational SLA/SLM -related R&D
- Access to experiences from service provisioning
 - Pilots with new approaches of particular interest, e.g. data-related initiatives
- Feedback
 - Please keep us honest!

How to get involved?



- Mailing lists are being set up
 - Sign up at <http://gslm.eu/contact>
- Expert groups for targeted consultation
 - Volunteer (yourself and others) at <http://gslm.eu/contact>
- Join in the events
 - BDIM workshop, May 27th
 - <http://bit.ly/fD40Sb>
 - MDGS workshop at EuroPar conference (end of August)
 - <http://europar2011.bordeaux.inria.fr/>
 - Next gSLM events will be published at <http://gslm.eu/>

- Thank you!
- Questions, comments?
- More information: <http://gslm.eu/>
- The gSLM project is co-funded by the European Commission through the Seventh Framework Programme under contract number 261547