

The Onelab2 project and research on federations

Panayotis Antoniadis

(joint work with Serge Fdida and Timur Friedman)

Kassel, March, 2009

Laboratoire LIP6 – CNRS

UPMC Univ Paris 06



Outline

- **Network facilities**
 - Distributed, stand-alone, infrastructure providers
 - The Onelab2 project and PlanetLab Europe
- **The federation concept**
 - When and how to federate?
- **Current federation activities**
 - PLC/PLE/PLJ, G-lab, Federica, GENI, ...
- **Research agenda on federation policies**



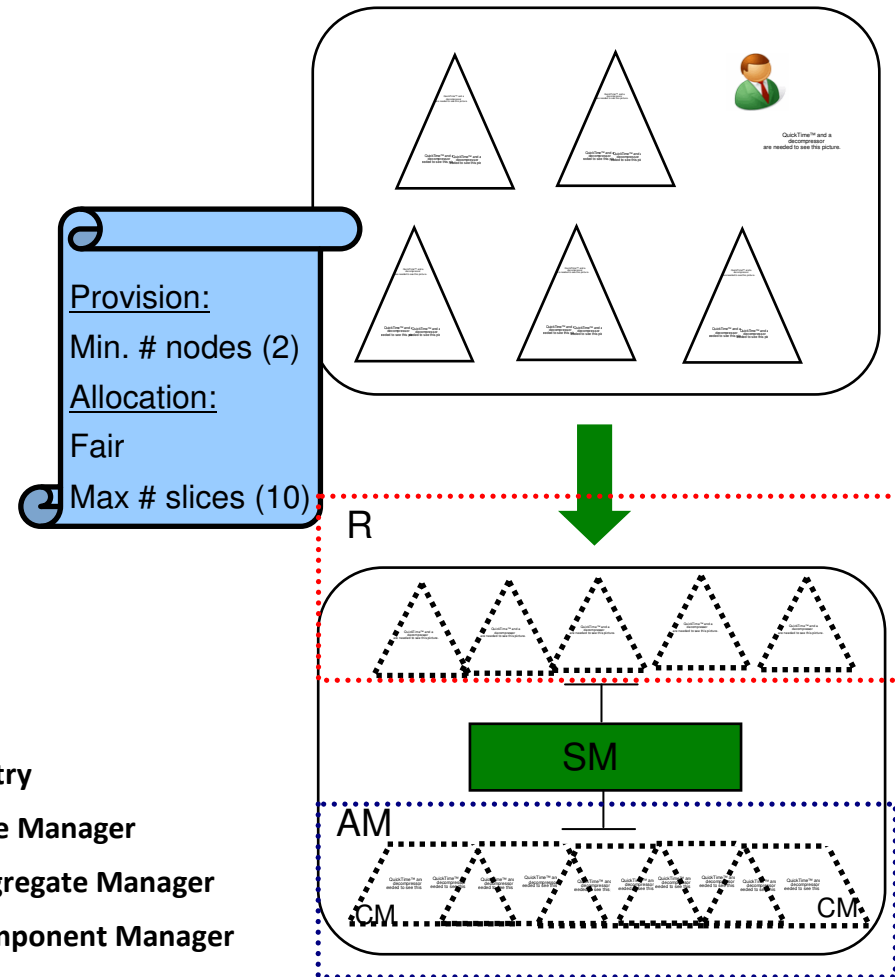
Distributed facilities

- PlanetLab

- Sites contribute *nodes*
- Researchers belonging to a site create *slices* that consume resources
- An Internet overlay
- Public good approach

- Slice-based Facility Architecture

- MyPLC
- Private vs. Public



R: Registry

SM: Slice Manager

AM: Aggregate Manager

CM: Component Manager

Network infrastructure providers

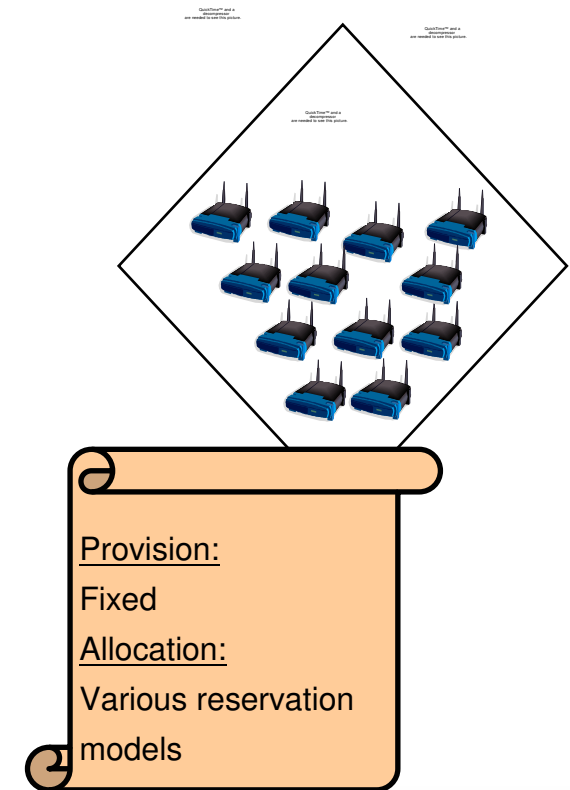
- “Geographic” initiatives
 - FEDERICA (Europe)
 - VINI (US)
 - SINET3 (Japan)
- Low-level network virtualization tools
- Main issue: resource allocation



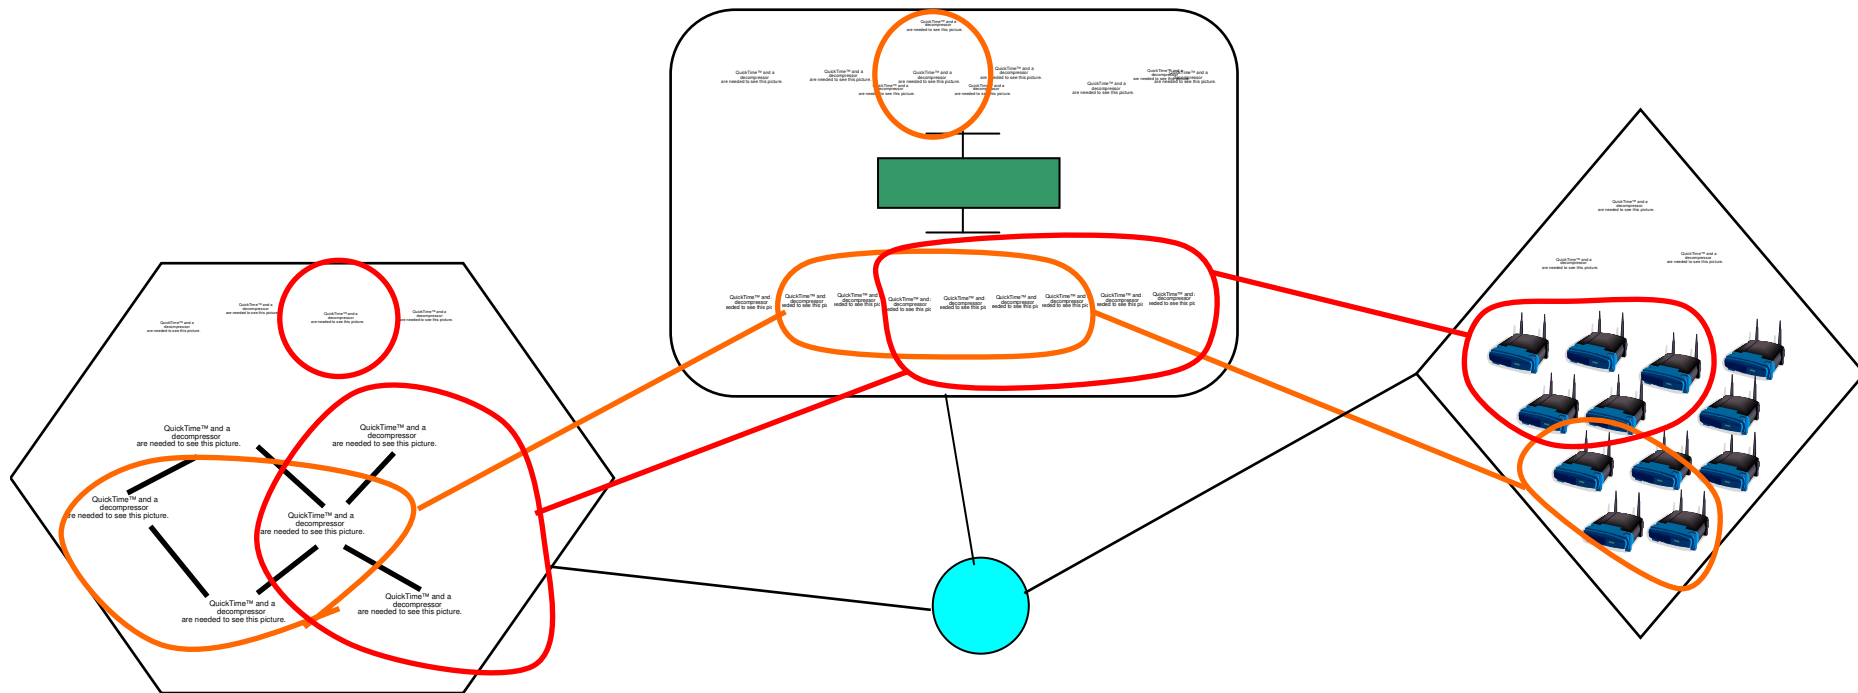
Provision:
Fixed
Allocation:
Access control
Constrains (time,
bandwidth)
Payments?

Stand-alone testbeds

- **Testbeds managed by a single authority**
 - Geographically limited
 - With associated users or not
- **Different technologies**
 - Wireless, emulation, sensors, ...
 - E.g., ORBIT, Emulab, Hagggle/ANA, WISENET
- **Again, resource allocation is critical**
 - Scarce resources
 - In many cases virtualization is hard



The federation vision



1. Share user credentials and resource descriptions
2. Agree on slice management API and allocation policy
3. Allow experiments to run across facilities

Facility research challenges

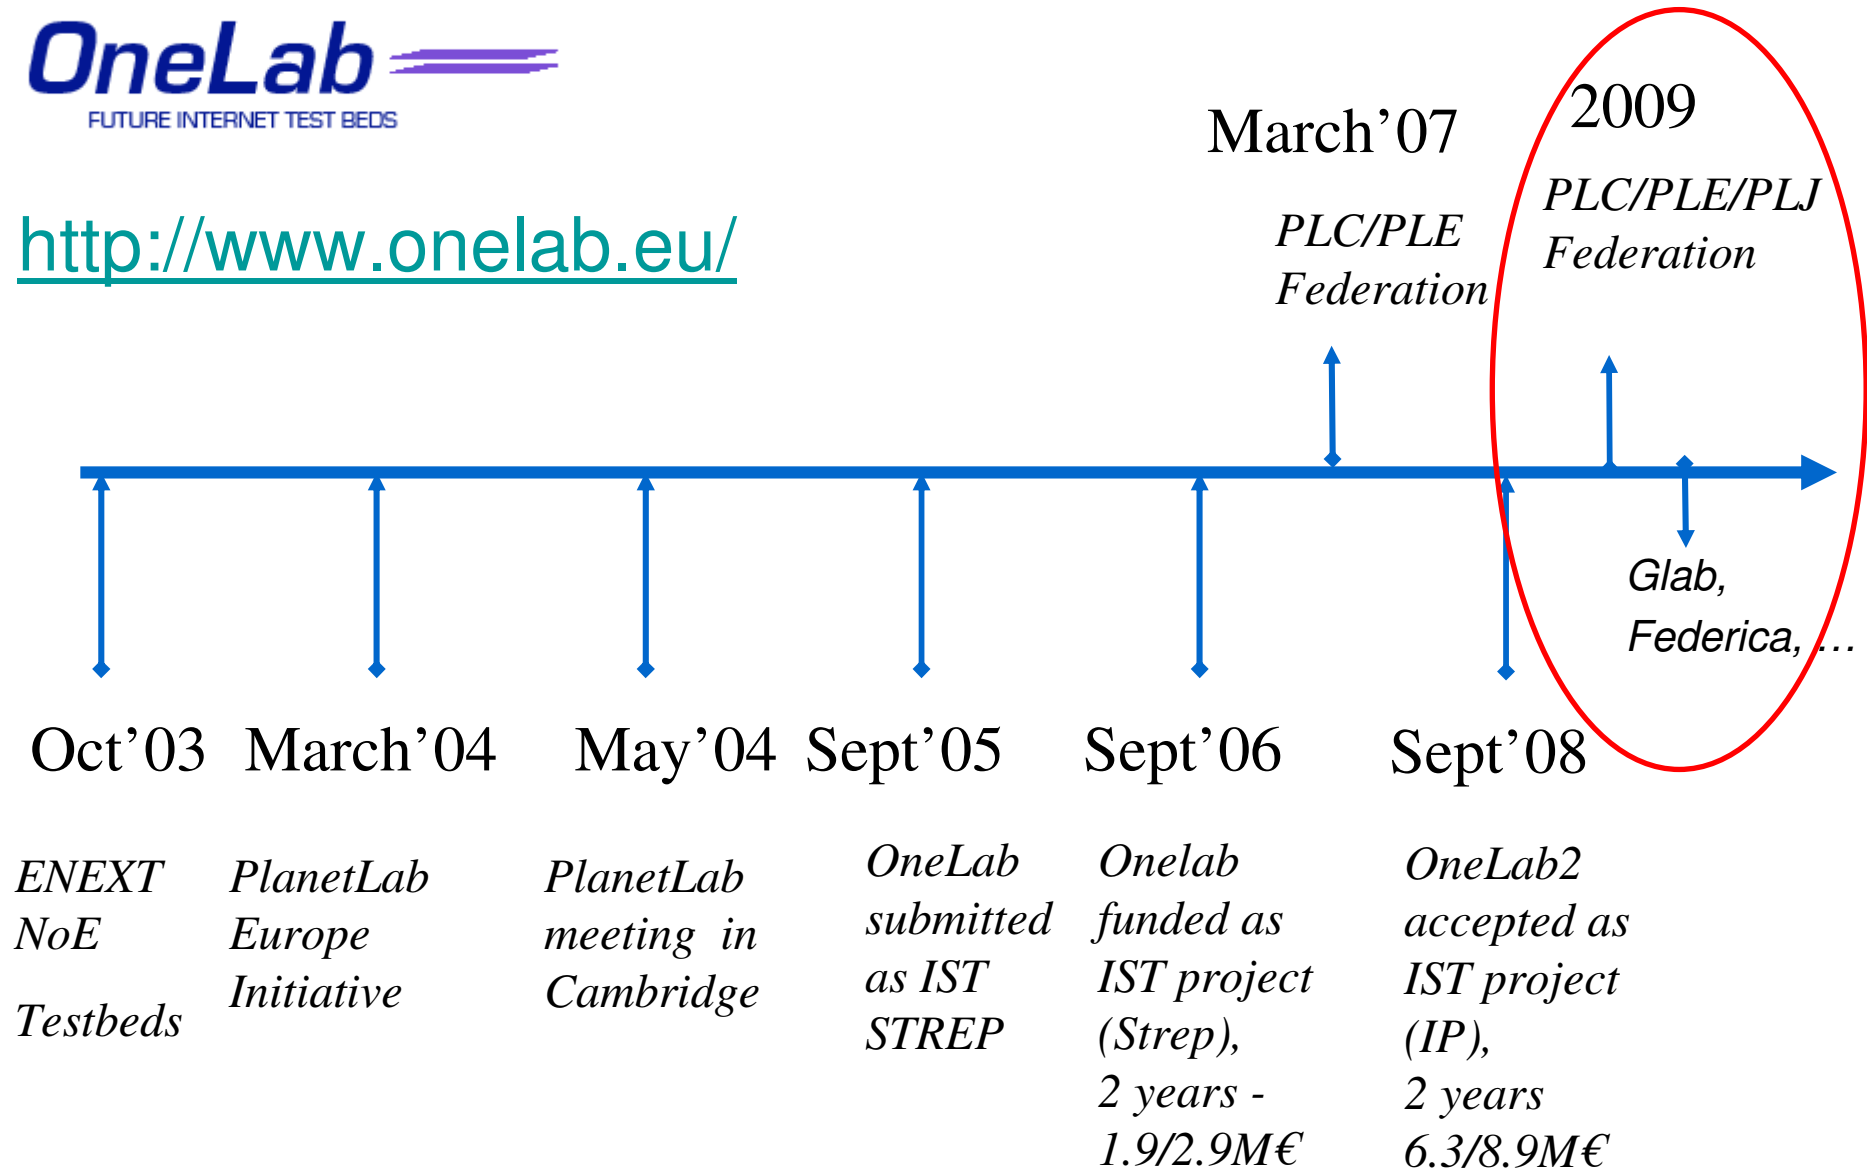
- **Virtualization:** Run concurrent experiments, support services
- **Monitoring:** Collect data and make them available
- **Legal:** Responsibilities and liabilities, IPR, ...
- **Benchmarking:** Assessment of the results produced, reproducibility
- **Security:** Robust and secure facility
- **Economics:** Of the facility, for the users, the operators
- **Federation**
 - Inter-operability framework
 - Data and resource representation
 - Control plane, resource management policies, incentives

Research on federation policies

- **Difficult to avoid**
 - Demand may exceed total capacity (at least this is the goal :-))
 - Scarce resources (e.g., specialized testbeds)
 - Conflicting local policies
- **Challenging aspects**
 - Heterogeneous resources, users, testbeds
 - Different objectives and constraints
 - Political issues as well
 - Complexity
- **Interesting theoretical and technical aspects**

Onelab2

- **Follow on of Onelab**
 - IP project, 2 years, 6.3/8.9M Euro
- **Responsible for Planetlab Europe**
 - Already federated with PLC
 - http://www.planet-lab.eu/join_us
- **+ various stand-alone testbeds**
 - Wireless testbeds (WiMax, WiFi, Multi-link, Multi-radio)
 - SAC testbeds (ANA/Haggle)
- **Specific focus on measurements tools**
- **Currently 46 (35 operational) sites, 77 (65 available) nodes**



Federation

A federation is a union comprising a number of partially self-governing regions united by a central ("federal") government under a common set of objectives.



Why to federate?

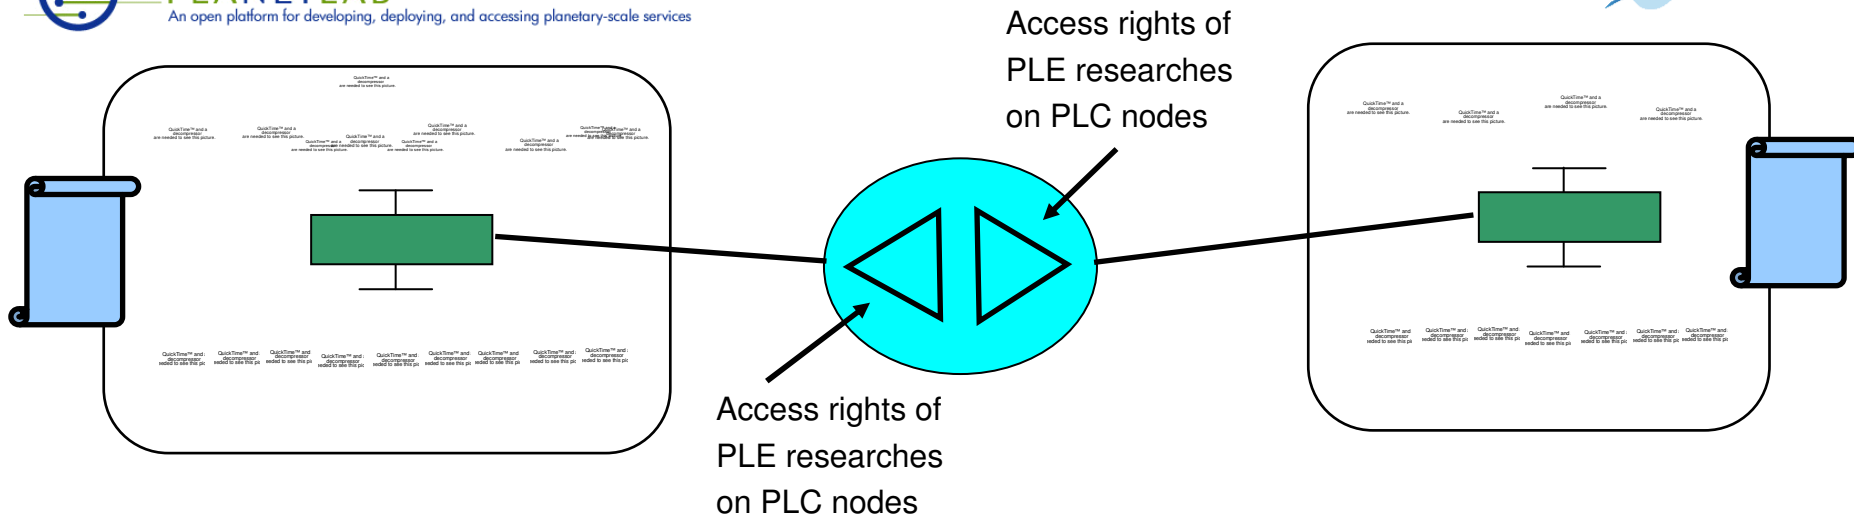
- **Many benefits**
 - Diversity, realism (geography, technology)
 - Scale (number of nodes, resources)
 - Multiplexing (more efficient resource usage)
 - Creation of a global research community
- **But also challenges**
 - Complexity
 - Legal and trust issues
 - Policies



When and how to federate?

- **Different objectives:** Scale, realism, reproducibility, heterogeneity
- **Different constraints:** Security, privacy, allocation policies
- **When (not)? (which types of facilities)**
 - Commercial vs. open testbeds
 - Reproducibility vs. realism
 - Heterogeneity/scale vs. capacity per node/simplicity
- **How? (under which policies)**
 - Fairness vs. efficiency
 - Sophisticated incentive mechanisms vs. accounting
 - Independence vs. simplicity

PLE-PLC federation policy



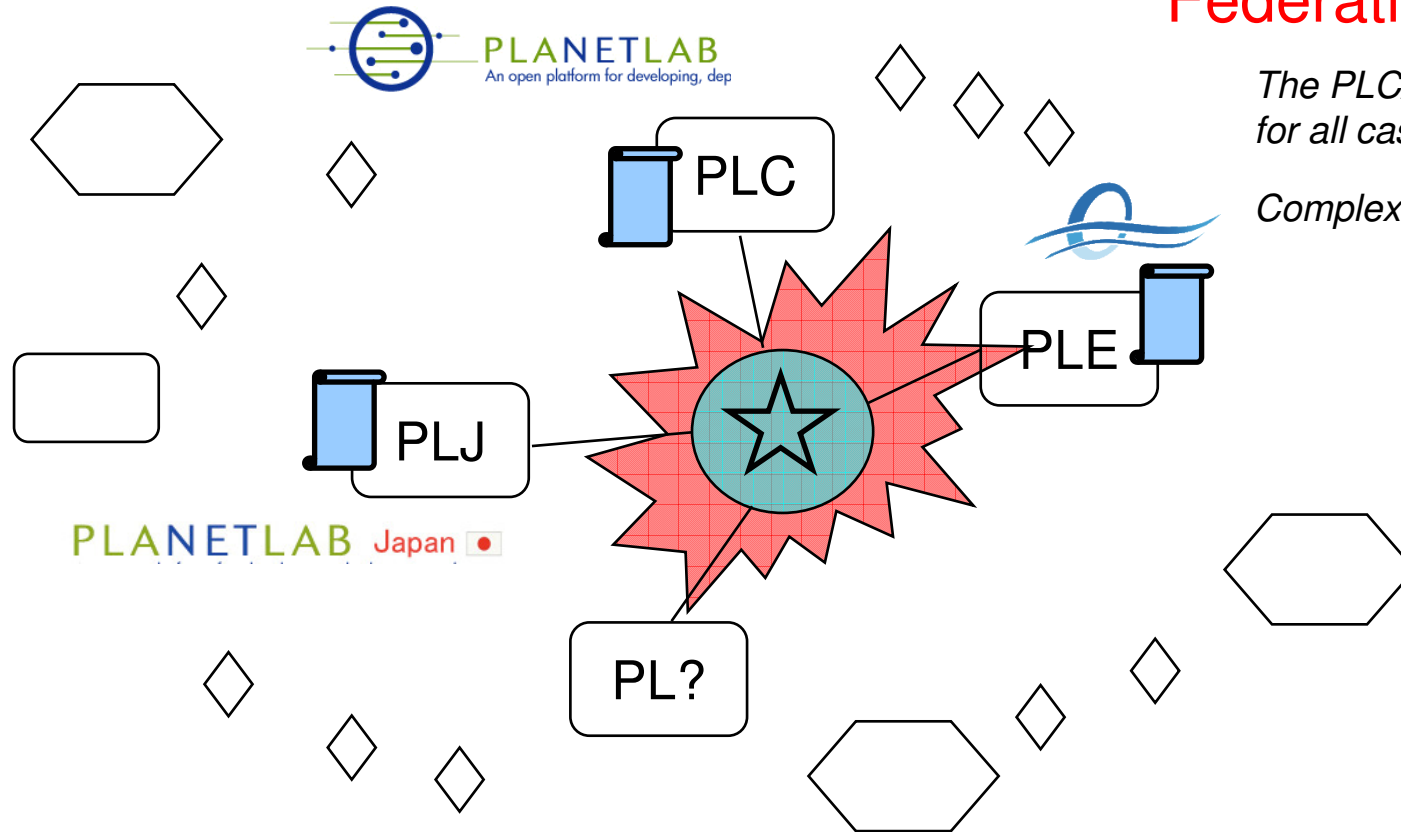
Current federation policy: “peering”

Users from both facilities have the same access rights over the whole infrastructure

Both facilities apply the same local policy



Future Plans

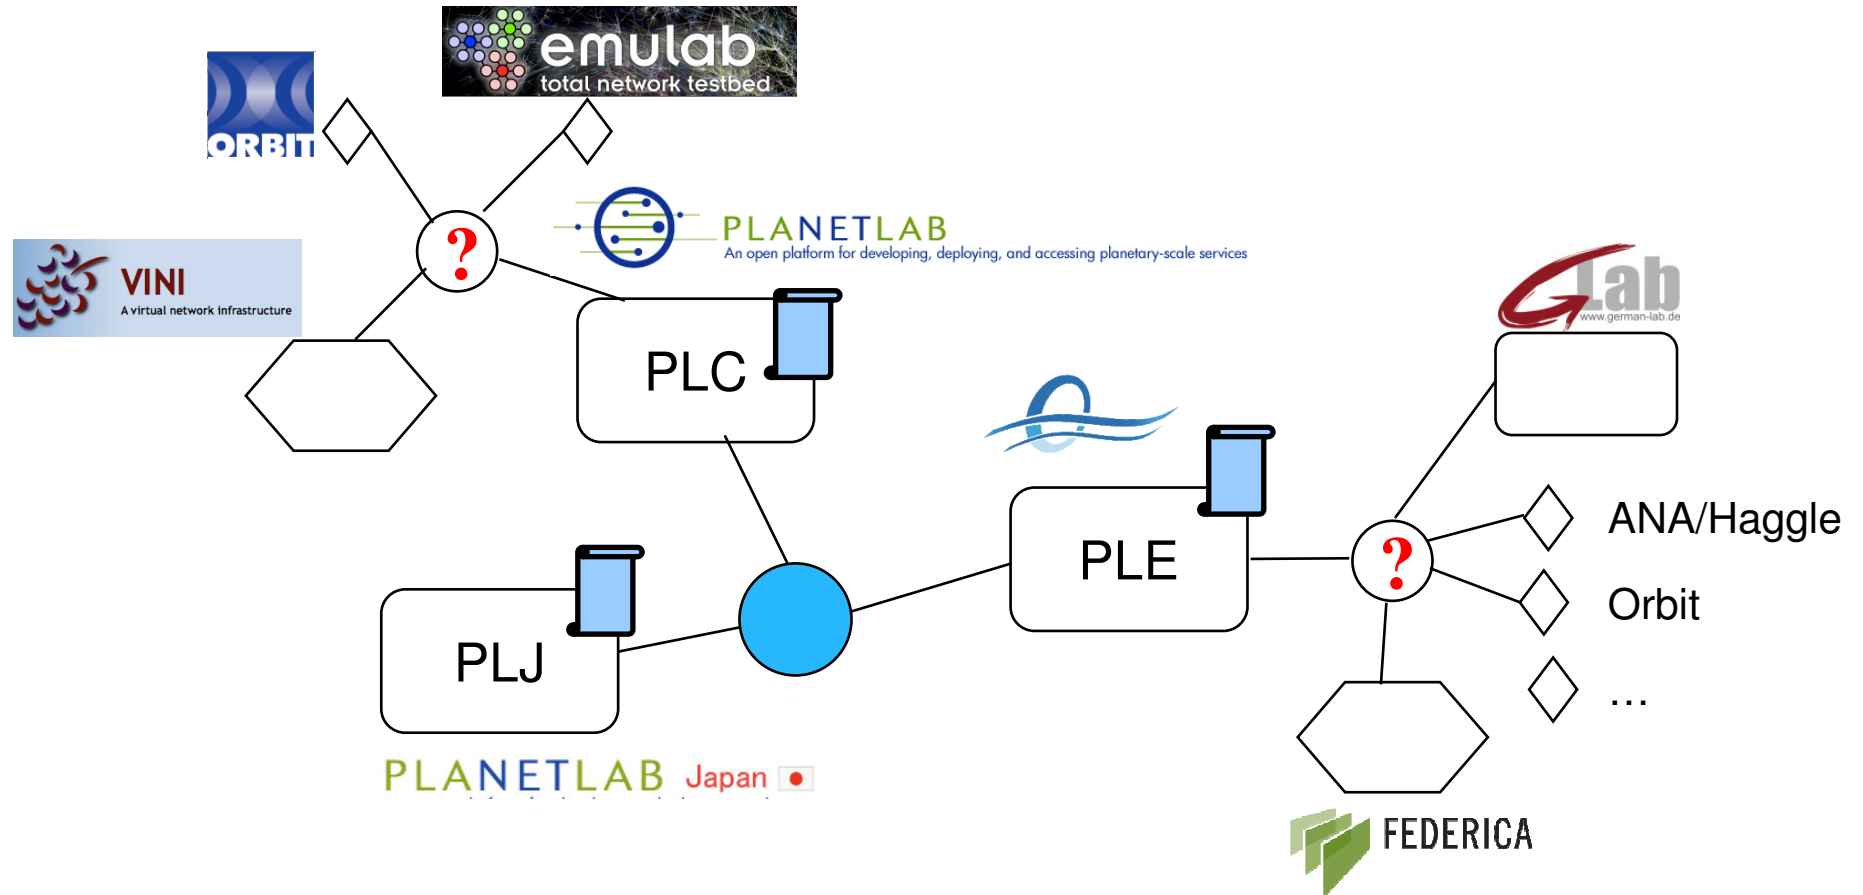


Federation policies?

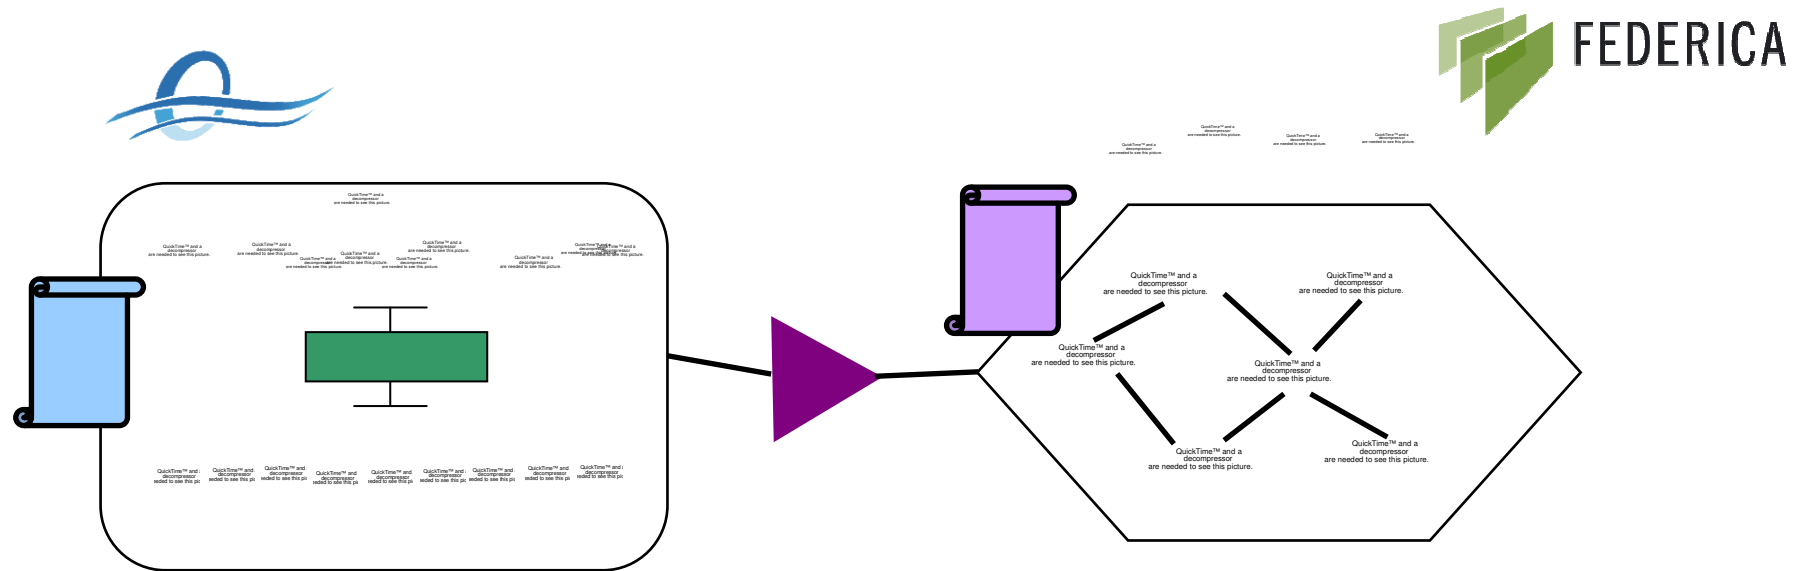
The PLC/PLE/PLJ model not appro for all cases

Complexity/scalability issues

Need for hierarchical federation




Other types of federation policies (PLE-Federica)



PLE users a subset of Federica users

A consumer-provider relationship?

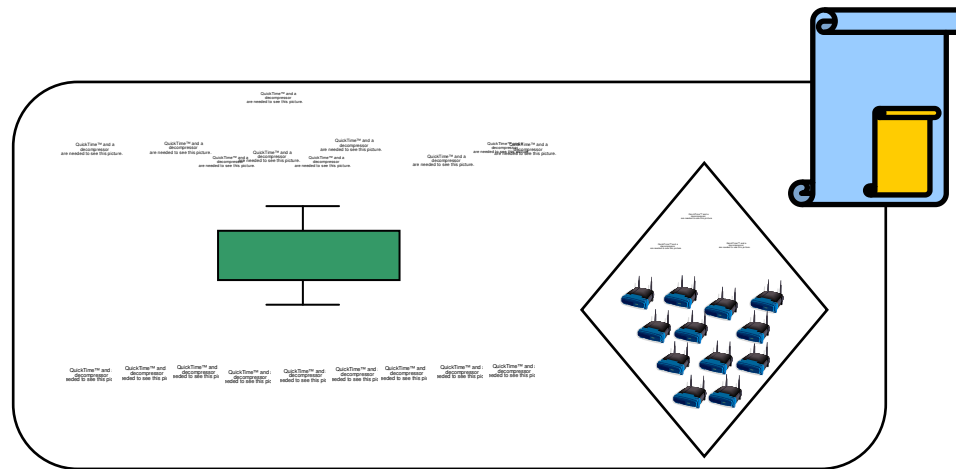
PLE-Private PlanetLabs (Glab, EverLab)



We may have to constrain access of PLE users to the private PlanetLab.


Perhaps at the level that Glab users consume PLE resources?

PLE-Stand-alone testbeds



Stand-alone testbeds as “special” sites?

And it gets more complicated ...



Local vs. global policies, 1st vs. 2nd federation level

Competition. "back" edges

The economic point of view

- **Incentives are required for facilities and users**
 - Resource contribution
 - Reliability
 - Diversity
 - Early entry
 - Responsible use
 - Consumption
 - Security

Possible punishments/rewards

- **Service differentiation**
 - Number of nodes, coverage
 - Number of slices, duration
 - Resource sharing (priority in scheduling, quotas)
 - Privileged access to scarce resources, reservations
- **Relate allocation with contribution**
 - Both for local and federation policies
- **Or maybe use markets?**
 - Regulation vs. “self-selection”
 - Fairness vs. efficiency

Important aspects

- **Resource availability vs. resource consumption**
 - Public good vs. market model
 - Independence?
- **Relation of demand with capacity**
 - Are there congestion effects?
- **Timescales of fairness**
 - Long-term vs. short-term experiments
 - But some experiments offer services!

First step: resource characterization

- **How can we estimate the total capacity/value of a facility?**
- **Three main value generators**
 - Diversity (number of nodes, position, technology)
 - Capacity (CPU, memory, bandwidth)
 - Time (duration, reliability)
- **Can we define a “typical” node?**
 - Different for different experiment profiles?
- **We need feedback from the users!**

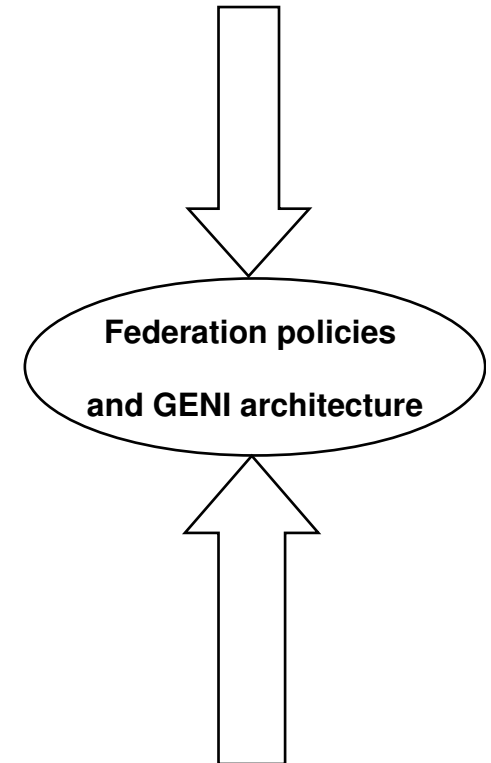
Our research agenda

- **Top-down**

- The notion of federation at large
- Economic and business models
- Lessons learned from ISP interconnection and federalism
- Evaluation of different policies based on different assumptions

- **Bottom-up**

- Simple resource allocation/provision rules
 - Encourage contribution, avoid overconsumption, and protect newcomers
- A friendly slice management framework
 - Encourage participation, community building and feedback
- A powerful meta-monitoring tool
 - User preferences and identify different experiment profiles
 - Strategic behaviour



Check on our progress :-)

- <http://www.planet-lab.eu/>
- **And join us!**

