PlanetLab and OneLab

Presentation for EuroView
1 August 2006

Timur Friedman

Université Pierre et Marie Curie
Laboratoire LIP6-CNRS

PlanetLab slides based on slides
provided courtesy of Larry Peterson

Outline

- PlanetLab
  - Architecture description
- OneLab
  - Extend PlanetLab
  - Deepen PlanetLab
  - Federate PlanetLab
- Interesting questions

PlanetLab

- An open platform for:
  - testing overlays,
  - deploying experimental services,
  - deploying commercial services,
  - developing the next generation of internet technologies.

- A set of virtual machines
  - *distributed virtualization*
  - each of 350+ network services runs in its own *slice*

PlanetLab nodes

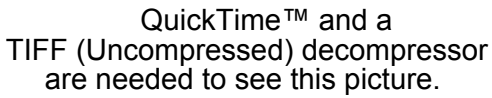

- 693 machines spanning 335 sites and 35 countries
  - nodes within a LAN-hop of > 2M users
- Administration at Princeton University
  - Larry Peterson, six full-time systems administrators

Slices

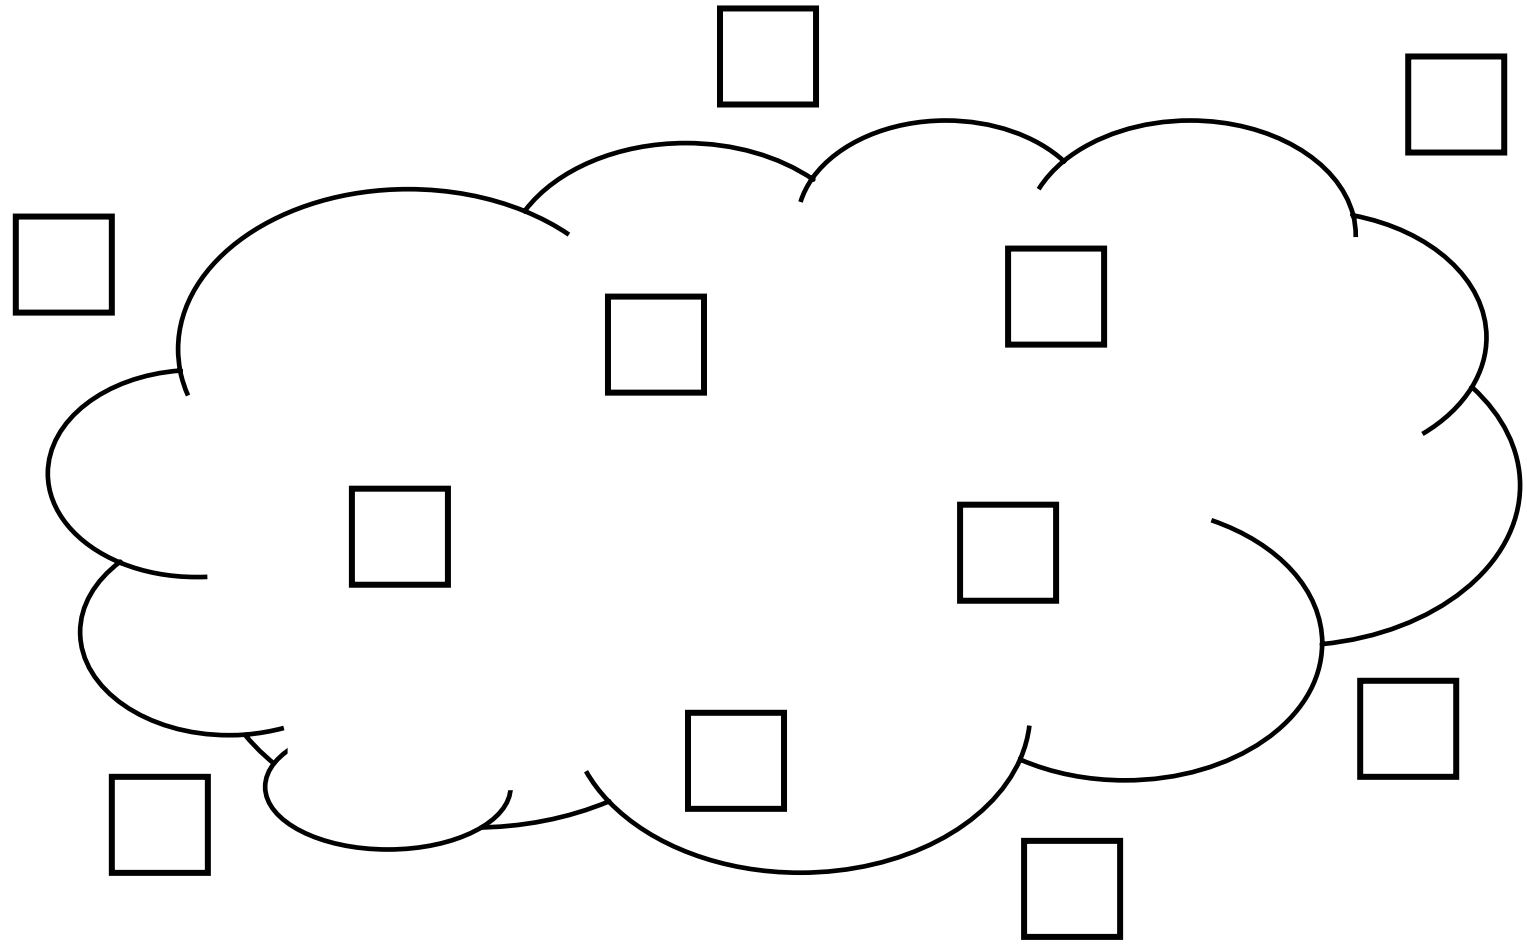

Slices

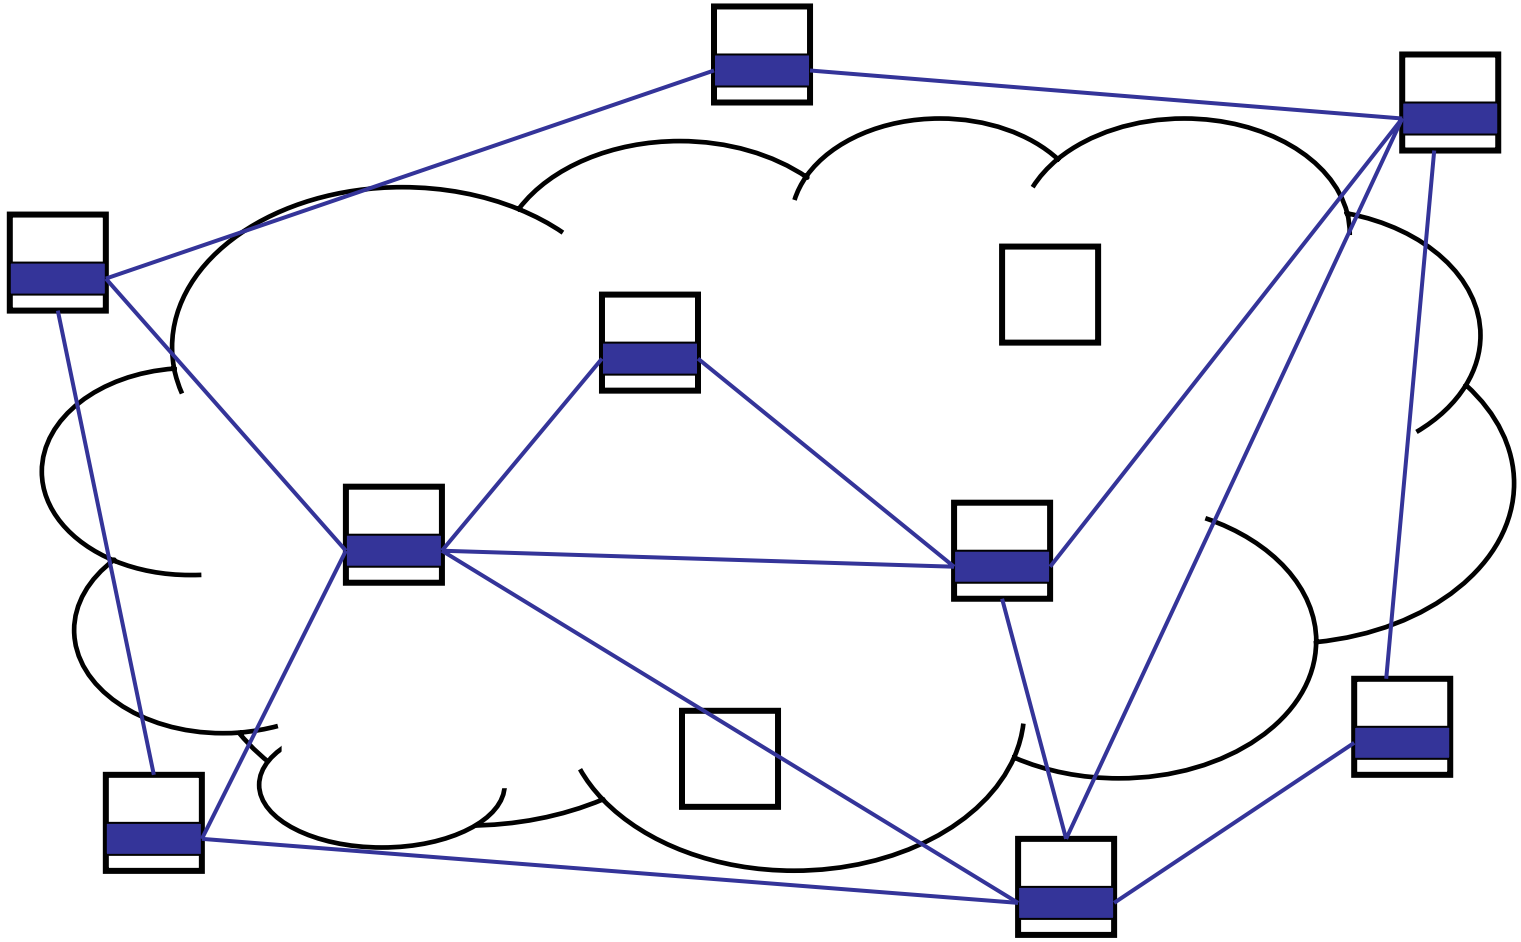

Slices

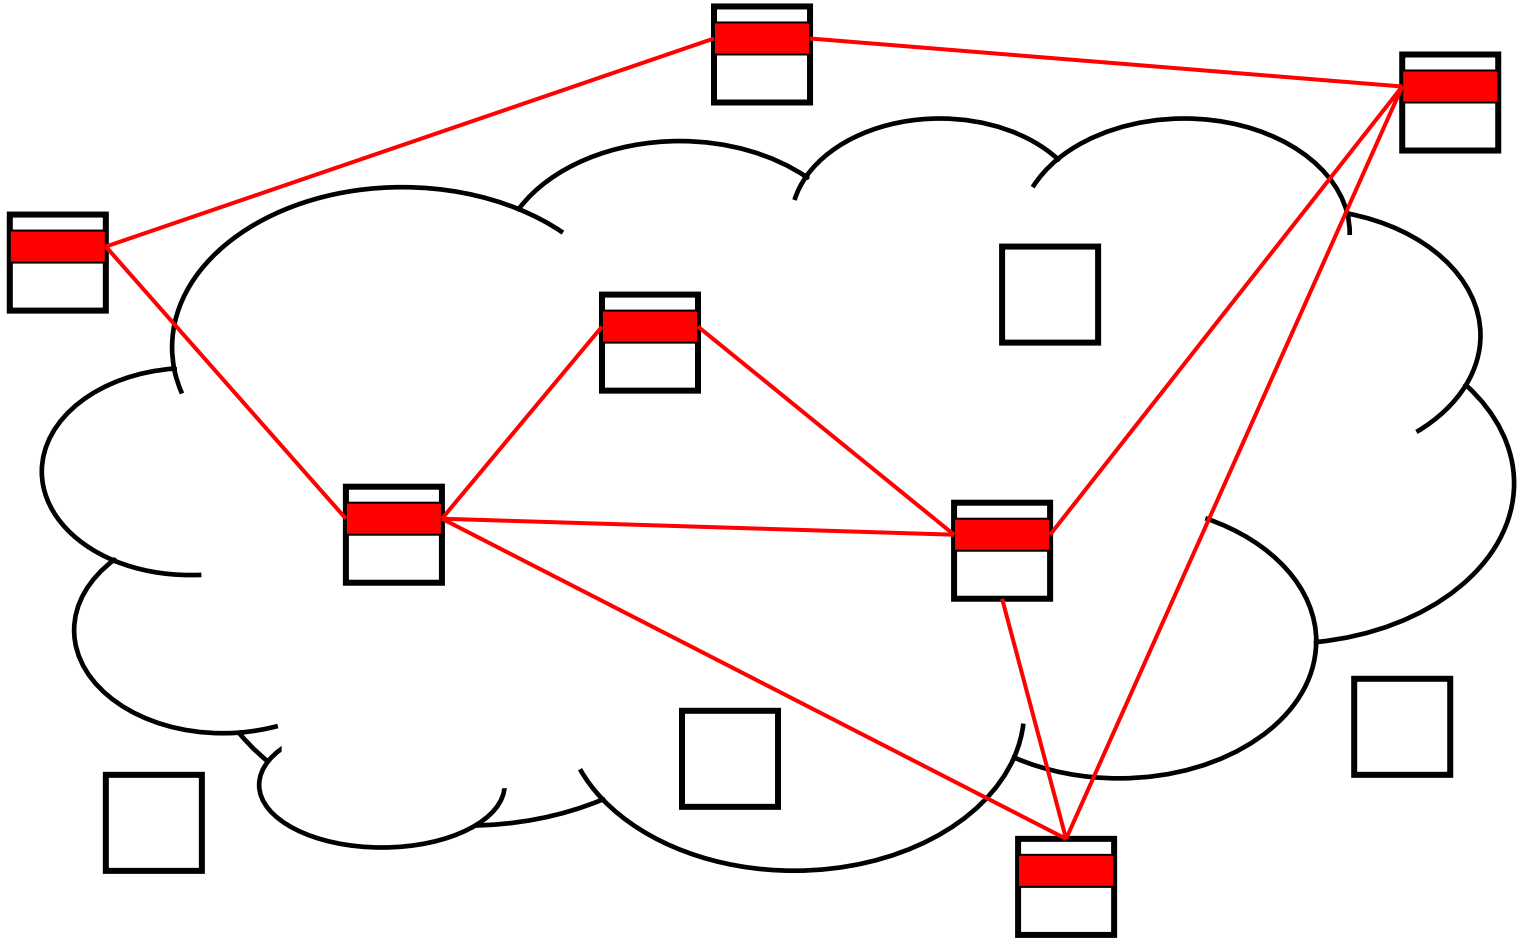

User Opt-in

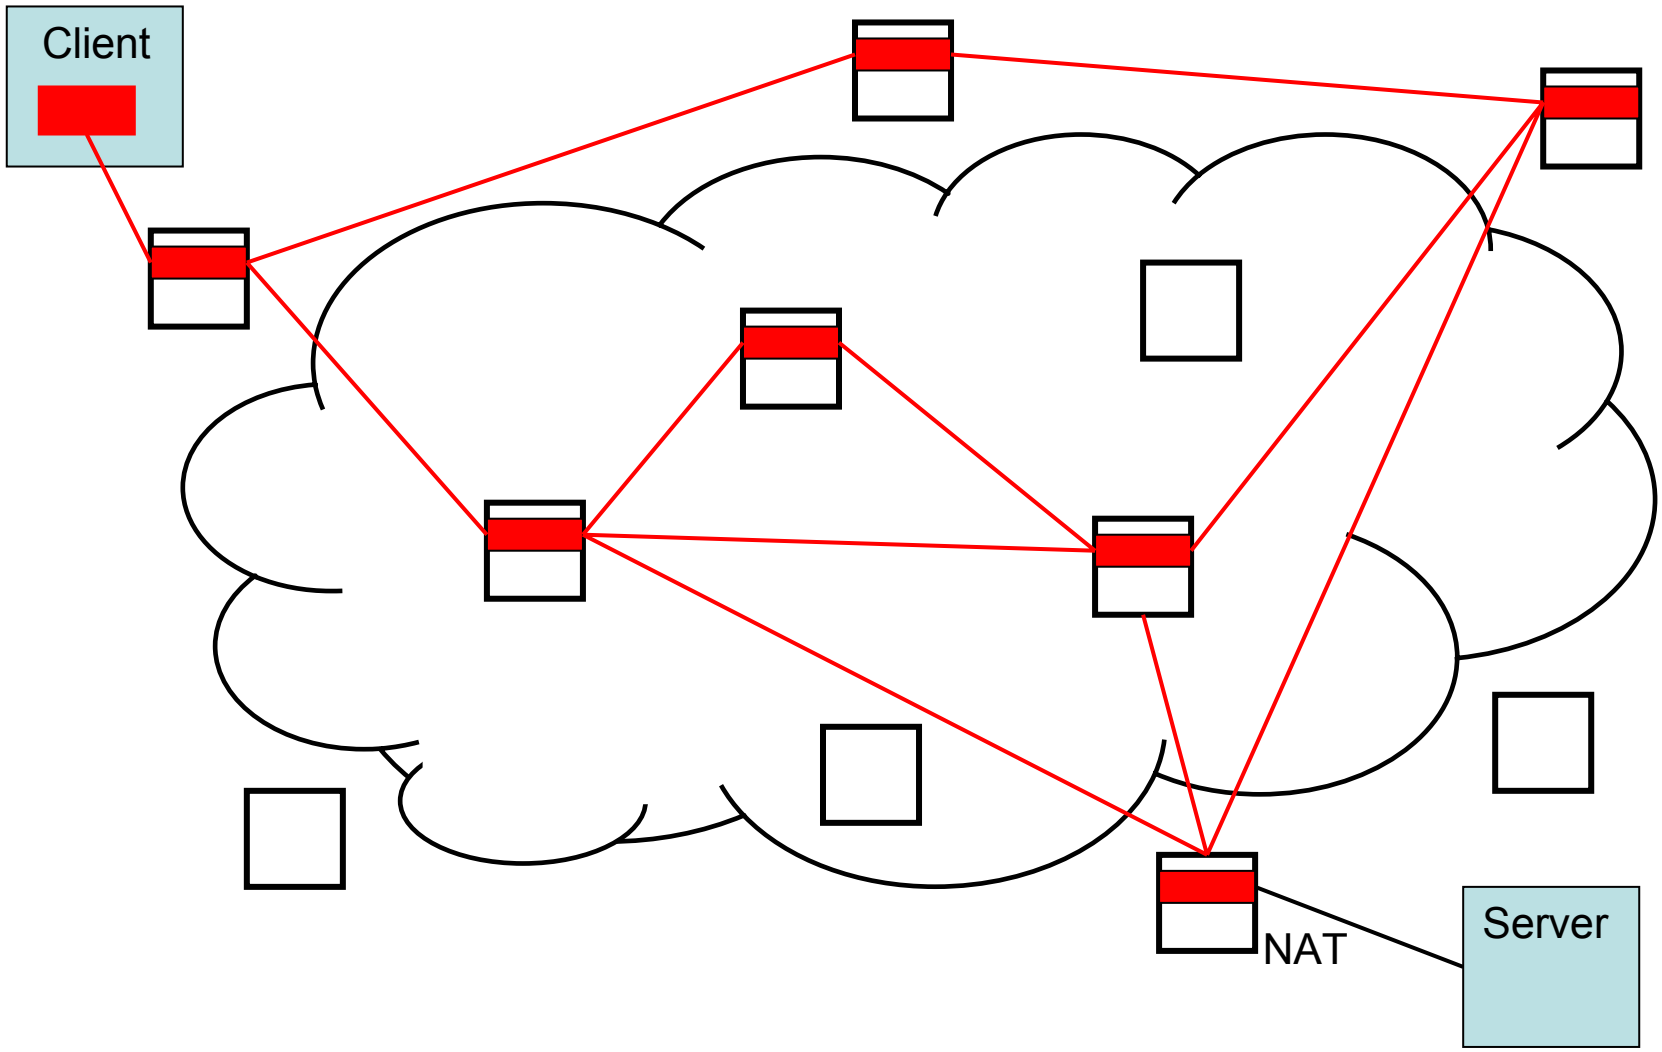

Per-Node View

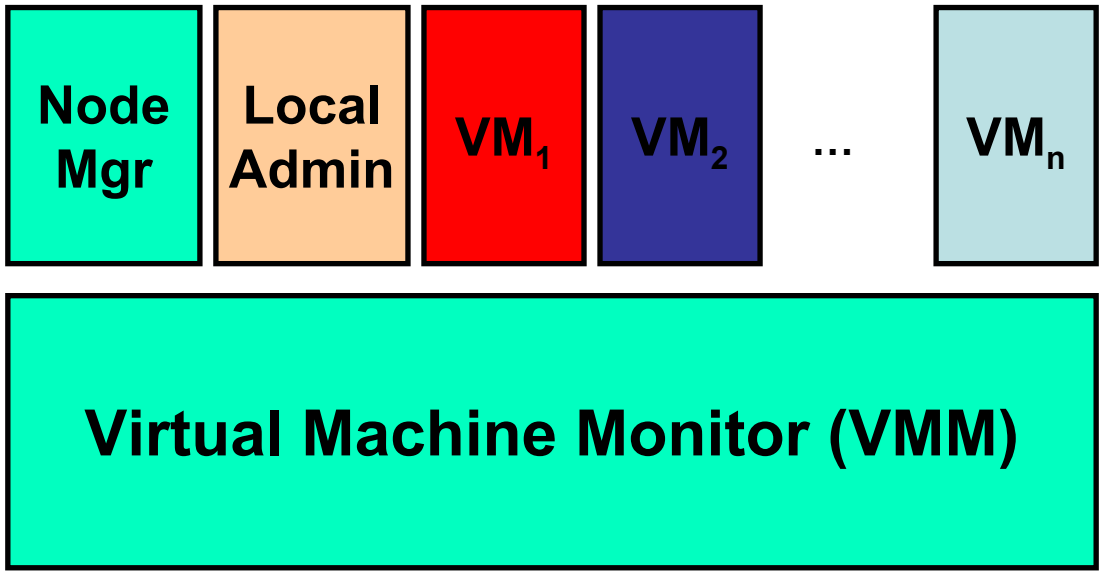

VMM: Currently Linux with vserver extensions
Could eventually be Xen

Architecture (1)

- Node Operating System
  - isolate slices
  - audit behavior
- PlanetLab Central (PLC)
  - remotely manage nodes
  - bootstrap service to instantiate and control slices
- Third-party Infrastructure Services
  - monitor slice/node health
  - discover available resources
  - create and configure a slice
  - resource allocation

Architecture (2)

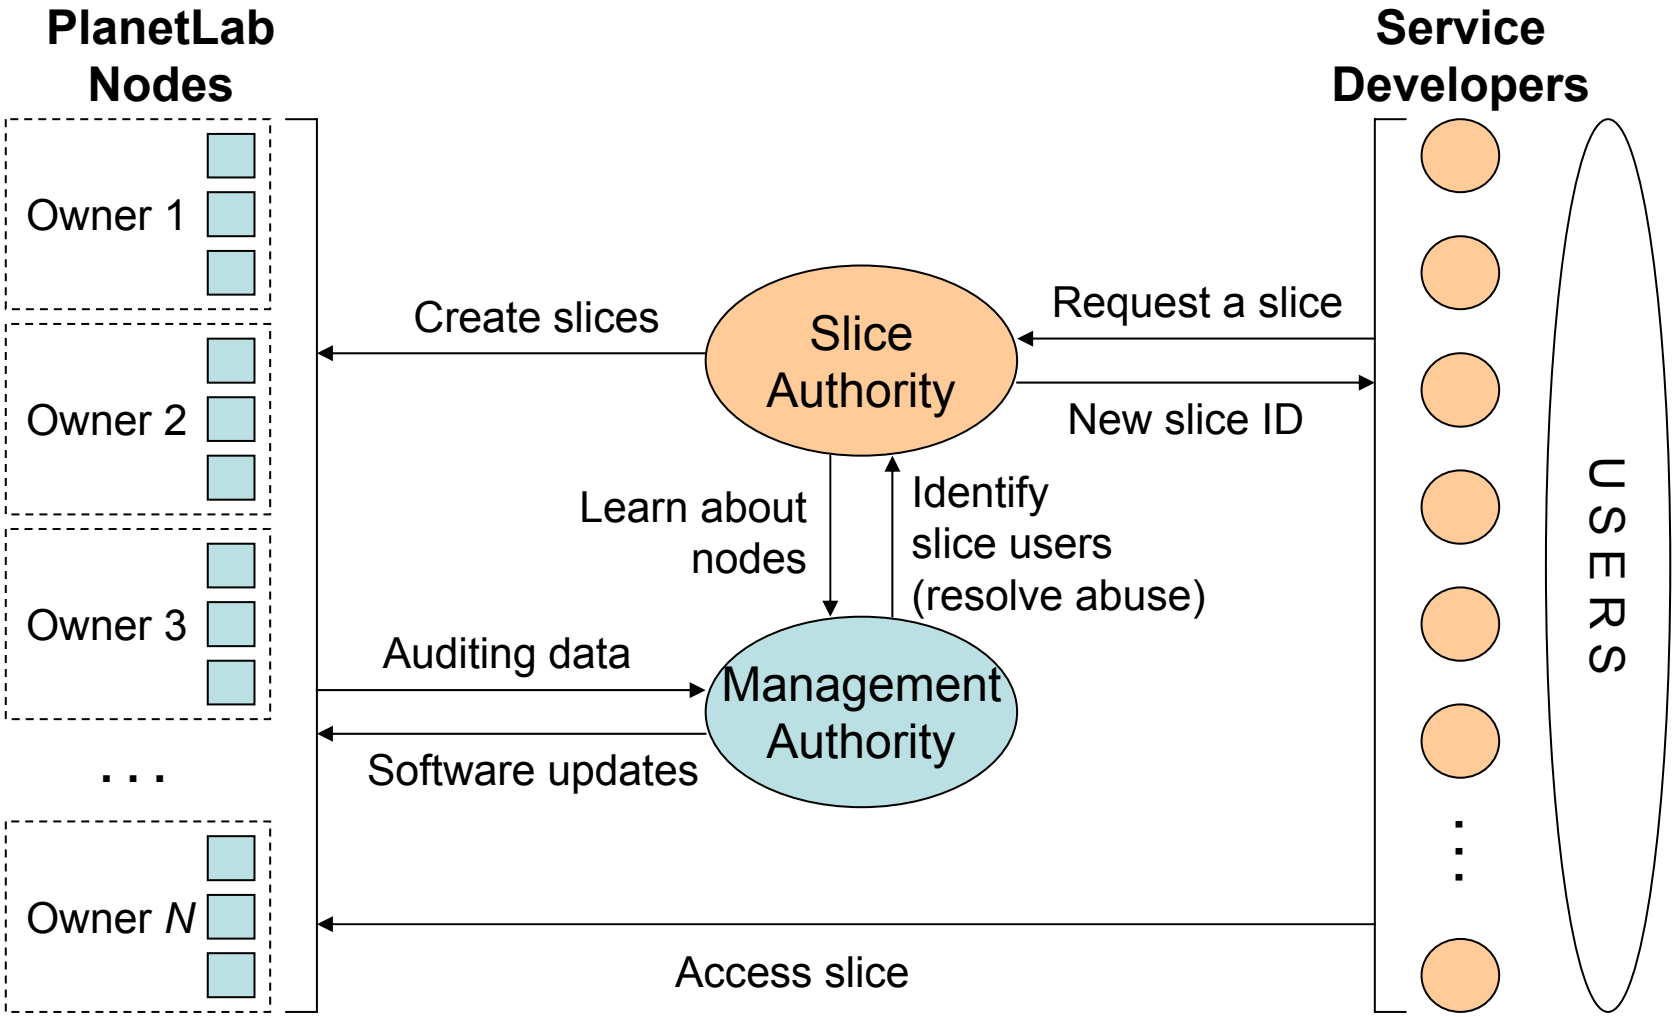

Architecture (3)

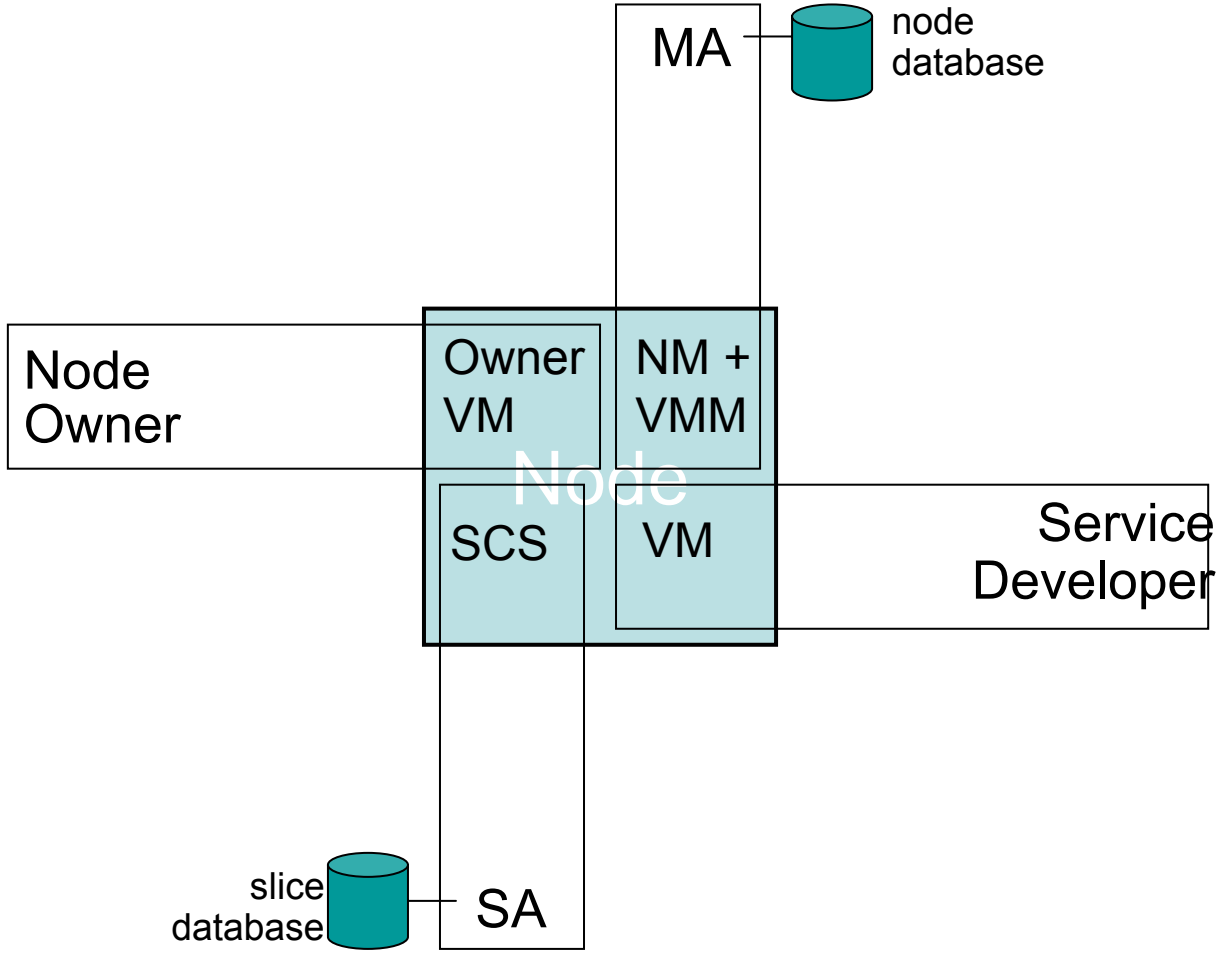

Long-Running Services

- Content Distribution
  - CoDeeN: Princeton (serving > 1 TB of data per day)
  - Coral CDN: NYU
  - Cobweb: Cornell
- Internet Measurement
  - ScriptRoute: Washington, Maryland
- Anomaly Detection & Fault Diagnosis
  - PIER: Berkeley, Intel
  - PlanetSeer: Princeton
- DHT
  - Bamboo (OpenDHT): Berkeley, Intel
  - Chord (DHash): MIT

Services (cont)

- Routing
  - i3: Berkeley
  - Virtual ISP: Princeton
- DNS
  - CoDNS: Princeton
  - CoDoNs: Cornell
- Storage & Large File Transfer
  - LOCI: Tennessee
  - CoBlitz: Princeton
  - Shark: NYU
- Multicast
  - End System Multicast: CMU
  - Tmesh: Michigan

Usage Stats

- Slices: 350 - 425
- AS peers: 6000
- Users: 1028
- Bytes-per-day: 2 - 4 TB
  - Coral CDN represents about half of this
- IP-flows-per-day: 190M
- Unique IP-addrs-per-day: 1M

- Experiments on PlanetLab figure in many papers at major networking conferences

Outline

- PlanetLab
  - Architecture description
- OneLab
  - Extend PlanetLab
  - Deepen PlanetLab
  - Federate PlanetLab
- Interesting questions

OneLab

- STREP
  - A STREP (Specific Targeted Research Project) funded by the European Commission
  - start: September 2006
  - duration: 2 years

- Consortium
  - Project leader: U. P. & M. Curie LIP6-CNRS lab
  - Technical direction: INRIA Sophia-Antipolis
  - Partners:
    - Intel Research Cambridge
    - Universidad Carlos III de Madrid
    - Université Catholique de Louvain
    - Università di Napoli
    - France Telecom (Lannion)
    - Università di Pisa
    - Alcatel Italia
    - Telekomunikacja Polska

OneLab Goals

- Extend
  - Extend PlanetLab into new environments, beyond the traditional wired internet.
- Deepen
  - Deepen PlanetLab's monitoring capabilities.
- Federate
  - Provide a European administration for PlanetLab nodes in Europe.

PlanetLab Today

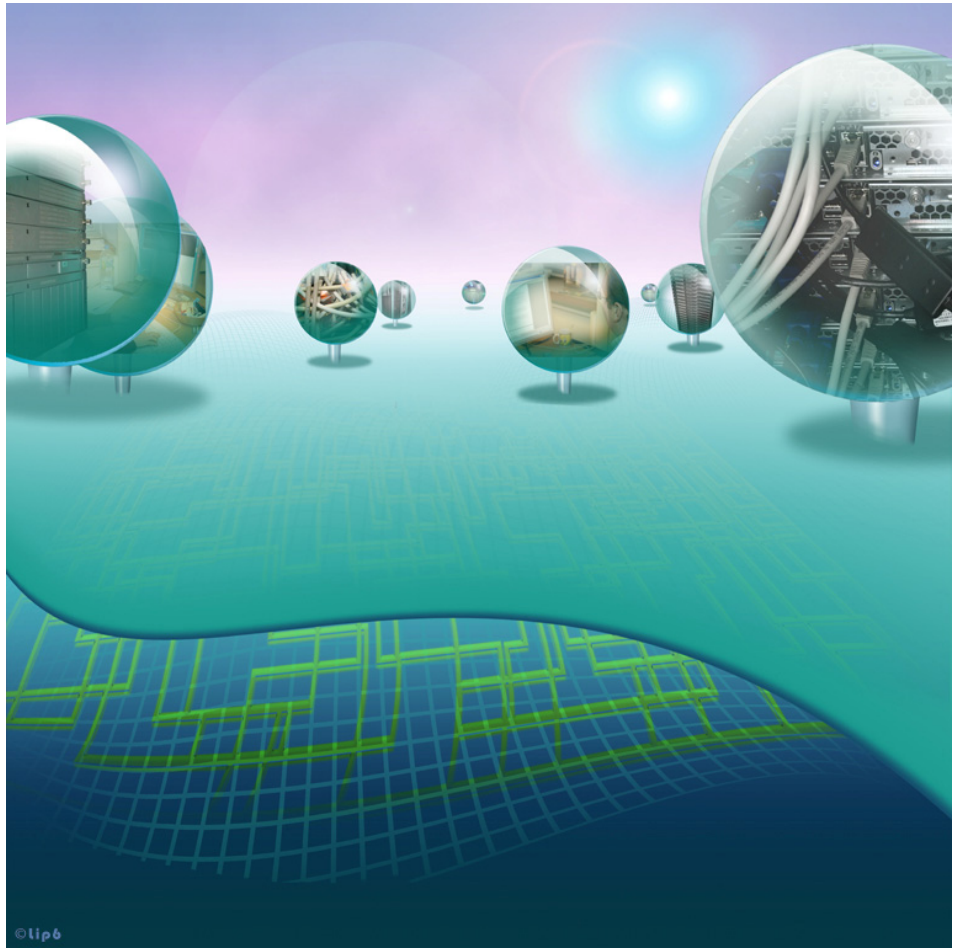

A set of end-hosts

A limited view of the underlying network

Built on the wired internet

OneLab Vision for PlanetLab

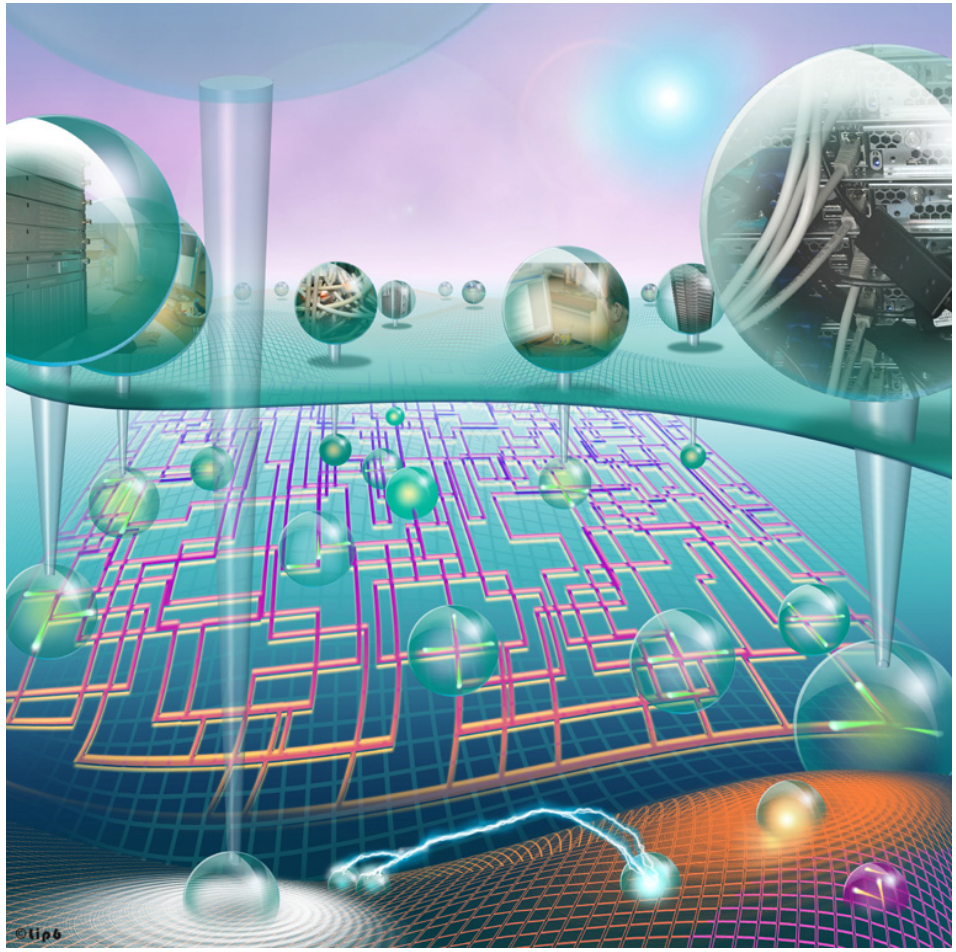

Reveal the underlying network

Extend into new wired and wireless environments

Goal: Extend

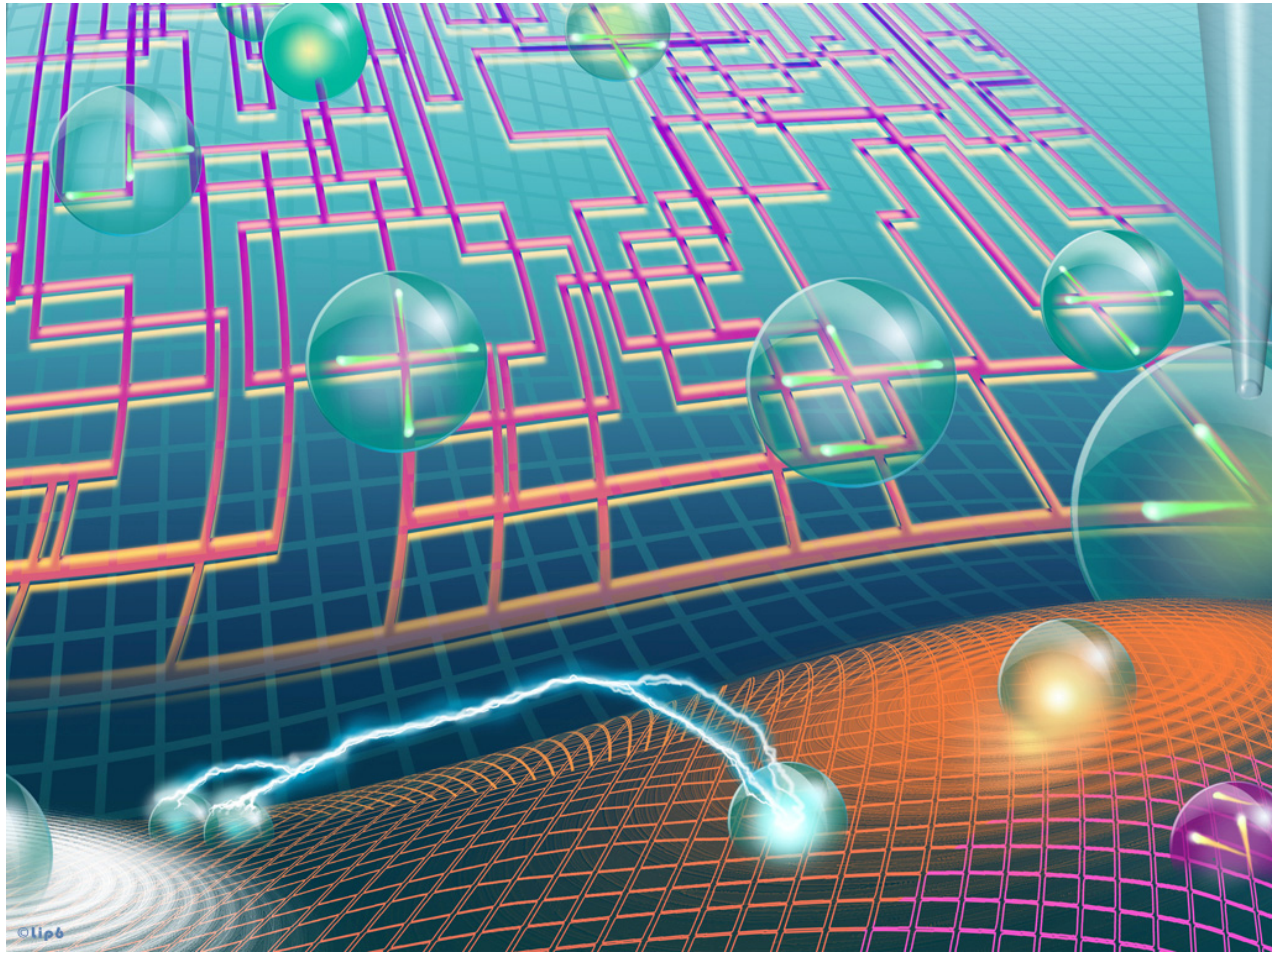

Why Extend PlanetLab?

- Problem: PlanetLab nodes are connected to the traditional wired internet.
  - They are mostly connected to high-performance networks such as Abilene, DANTE, NRENs.
  - These are not representative of the internet as a whole.
  - PlanetLab does not provide access to emerging environments.

OneLab's New Environments

- WiMAX (Université Catholique de Louvain)
  - Install two nodes connected via a commercial WiMAX provider
  - Nodes on trucks (constrained mobility)
- UMTS (Università di Napoli, Alcatel Italia)
  - Nodes on a UMTS micro-cell run by Alcatel Italia
- Wireless ad hoc networks (France Telecom at Lannion)
  - Nodes in a Wi-Fi mesh network (like ORBIT)
- Emulated (Università di Pisa)
  - For emerging wireless technologies
  - Based on *dummynet*
- Multihomed (Universidad Carlos III de Madrid)

Goal: Deepen

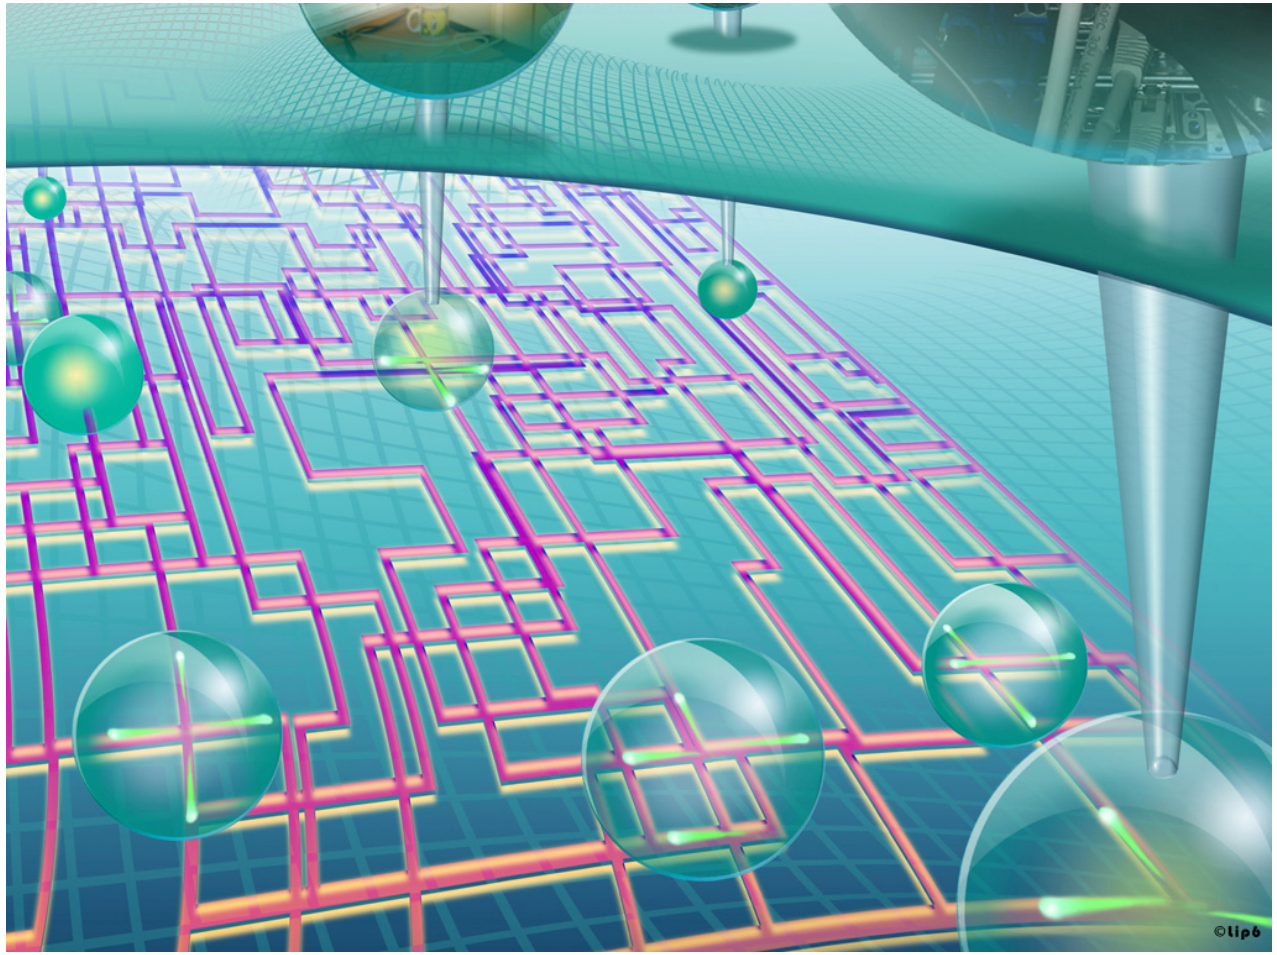

Expose the underlying network

Why Deepen PlanetLab?

- Problem: PlanetLab provides limited facilities to make applications aware of the underlying network
  - PlanetLab consists of end-hosts
  - Routing between nodes is controlled by the internet
    (This will change with GENI)
  - Applications must currently make their own measurements

OneLab Monitoring Components

- Passive monitoring (Intel Research Cambridge)
  - Track packets at the routers
  - Use *CoMo* boxes placed within DANTE
- Topology monitoring (U. P. & M. Curie)
  - Provide a view of the route structure
  - Build on *Scriptroute* to provide scalable active measurements

Goal: Federate

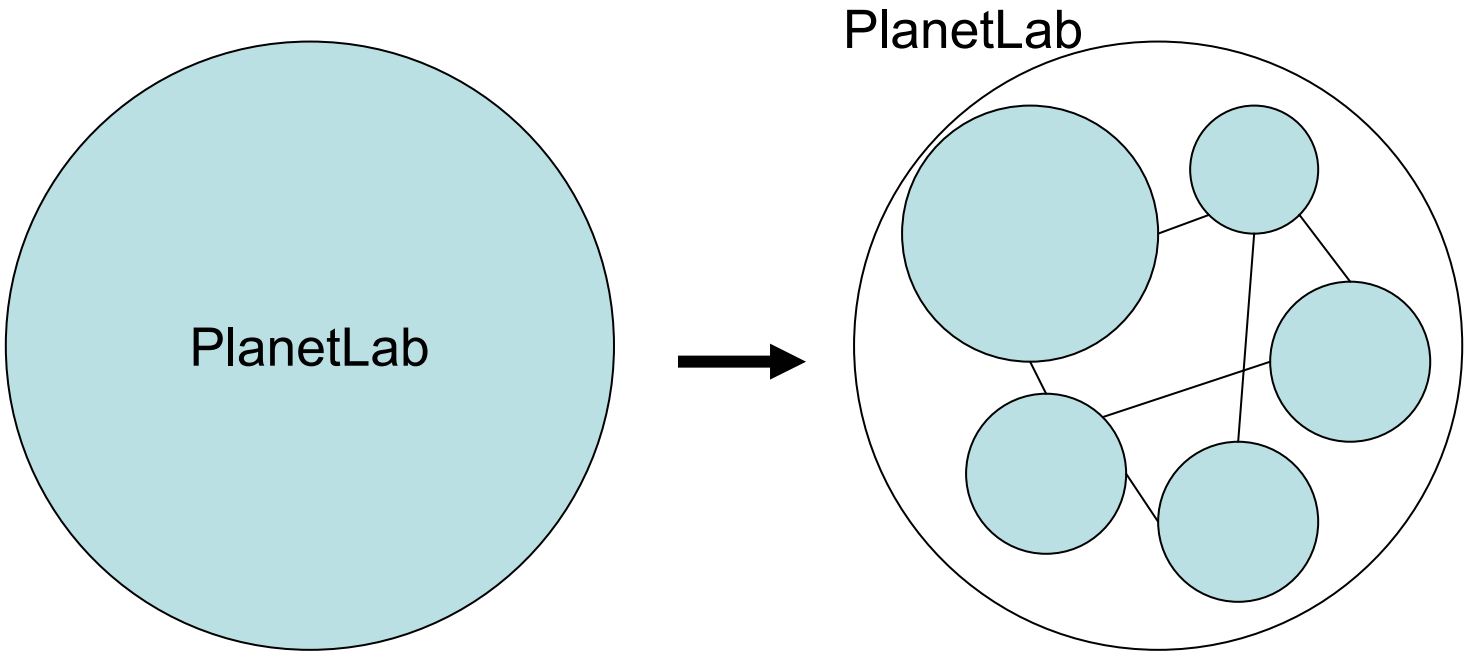

Why Federate PlanetLab?

- Problem: Changes to PlanetLab must come through the administration at Princeton.
  - PlanetLab in the US is necessarily less responsive to European research priorities.
    - What if we want to study a particular wireless technology, and this requires changes to the source code?
    - What if we wish to change the cost structure for small and medium size enterprises (currently $25,000/yr.)?

OneLab and Federation

- OneLab will create a PlanetLab Europe.
  - It will federate with PlanetLab in the US, Japan, and elsewhere.
    - Eventually federate with “Private PlanetLabs” as well
  - The federated structure will allow:
    - PlanetLab Europe to set policy in accordance with European research priorities,
    - PlanetLab Europe to customize the platform, so long as a common interface is preserved.

Outline

- PlanetLab
  - Architecture description
- OneLab
  - Extend PlanetLab
  - Deepen PlanetLab
  - Federate PlanetLab
- Interesting questions

How to virtualize wireless?

- We understand how to virtualize an end-host
  - Time-share a CPU
- We understand how to virtualize a wired link
  - MPLS for, e.g., VPNs
  - GENI will be doing this
- Less certain how to virtualize a wireless link
  - There is interference
  - One must make power and channel choices
  - What if the node is poorly connected?

Secure, scalable measurements?

- Security
  - Packet traces are highly sensitive
  - How can an application track only its own packets?
  - How can network providers set their own access policies for a common infrastructure?

- Scalability
  - Can active measurements be mutualized among applications?
  - How can information be shared among nodes to reduce measurement redundancy?

How to federate?

- A Private PlanetLab might have a rare resource
  - e.g., a node behind a wireless link
  - What are the right incentives to…
    - encourage the Private PlanetLab to share the resource
    - discourage over-subscription by other users?

- A Private PlanetLabs might wish to customize the source code
  - What are the right abstractions (APIs) that will allow both…
    - inter-operability
    - flexibility?