Network Virtualization

A state of the art

Workshop on Grid'5000
and network virtualization

Paris - October 16th, 2008

Pascale Vicat-Blanc Primet
Fabienne Anhalt
INRIA RESO
Pascale.Primet@inria.fr
Fabienne.Anhalt@ens-lyon.fr

Definitions

- Virtualization is an abstraction/ a framework (ie tools) that
  - decouples the physical hardware from the service level
  - Combines (aggregate) and/or divides(split) resources
    - M-to-N mapping (M “real” resources, N “virtual” resources)
  - To present a *transparent* view of one or more environments
  - To deliver greater resource/infrastructure utilization and flexibility
  - Virtualization started as virtual memory (Kilburn, Atlas(1962)),
  - VM-concept for migration-and partitioning-purposes (~1980)
  - Virtual circuits exist in networks since X25, ATM, Frame Relay, MPLS, GMPLS…
  - Emerged into a key-technology to aggregate Internet servers in the new re-centralization scenario. (fujisu-siemens view point)

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Network virtualisation

- Packet-switched users use to require “Private Network for security & QoS”

- NV solves a security problem
  - L2 & L3 VPNs - Ipsec : tunneling

- NV solves a cost problem for providing Private Networks: (capex)
  - Splitting available physical bandwidth into independant channels (L2 VPN,VLANs)
  - resource sharing

- NV solve a routing problem for performance optimization
  - Combining channels into E2E path with « predefined routing » (ex MPLS, λ path)
  - Circuit switching paradigm emulation (LSP)

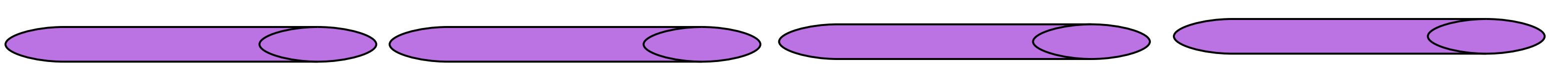

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

VLANs

- Segmentation of a LAN into several Virtual LANs
  - Packets are tagged so that the switches recognize their VLAN

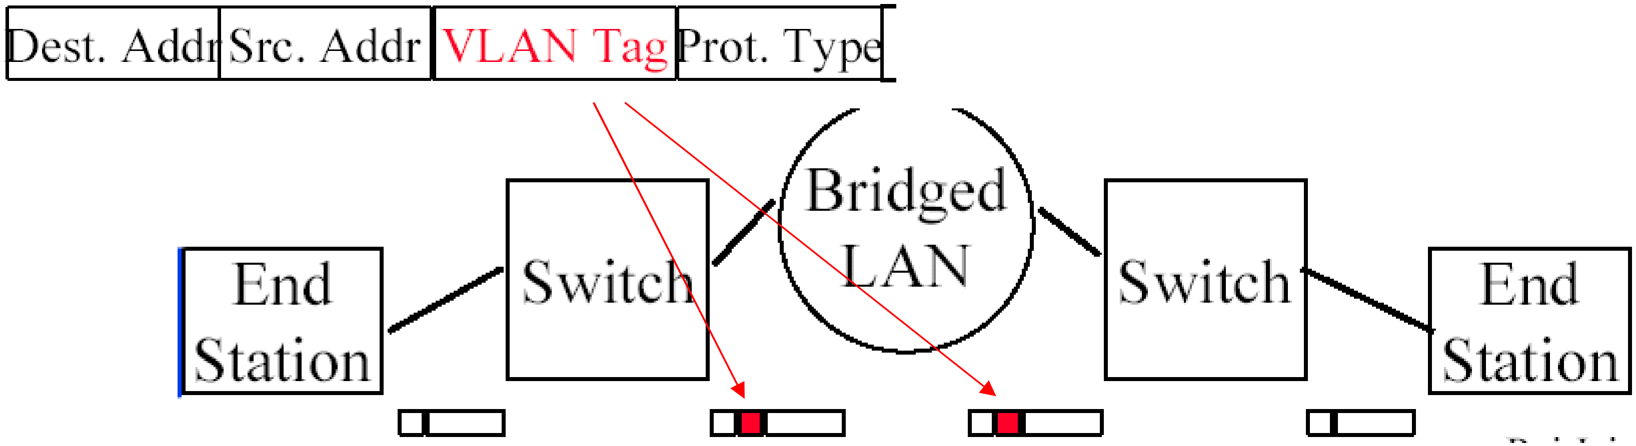

- Norme 802.1Q
  - Up to 4096 VLANs (#VLAN on 12 bits)
  - MAC@, port or higher level tagging

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

VLANs

- Benefits
  - Virtual Workgroups: Limitation of the broadcast domain
  - Performance: Less overhead in VLAN switches than in routers
  - Simple management:
    - Ease of adding users to the VLAN if MAC@ tagged
    - Ease of changing NIC if port tagged
  - Reduce cost: Less routers
  - Security: Restrict access by segmenting the network
- Issues
  - Add new components if port tagging
  - Big tag-correspondence tables and heavy administration if MAC@ tagging

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtual Private Networks

- Security
  - Communicate through the public Internet in a secure way
    - Tunnels (IP/IPsec, MPLS, L2TP)
- QoS
  - Isolation
  - Possibility of bandwidth provisioning per channel

Hose-Model

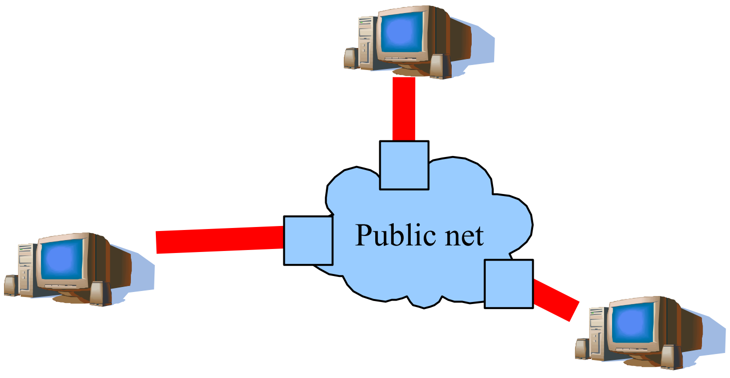

Full Mesh/Customer Pipe-Model

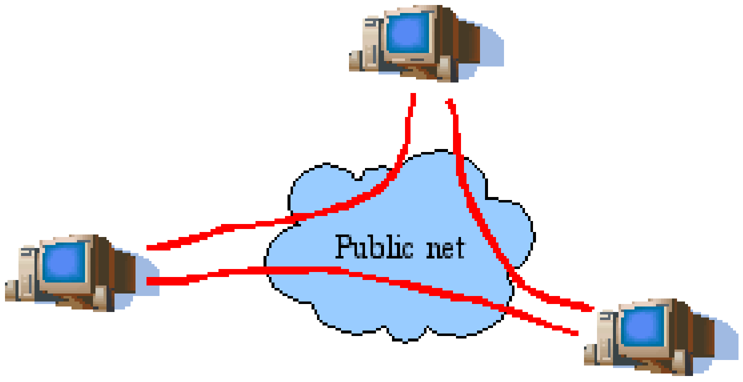

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

VPN with MPLS

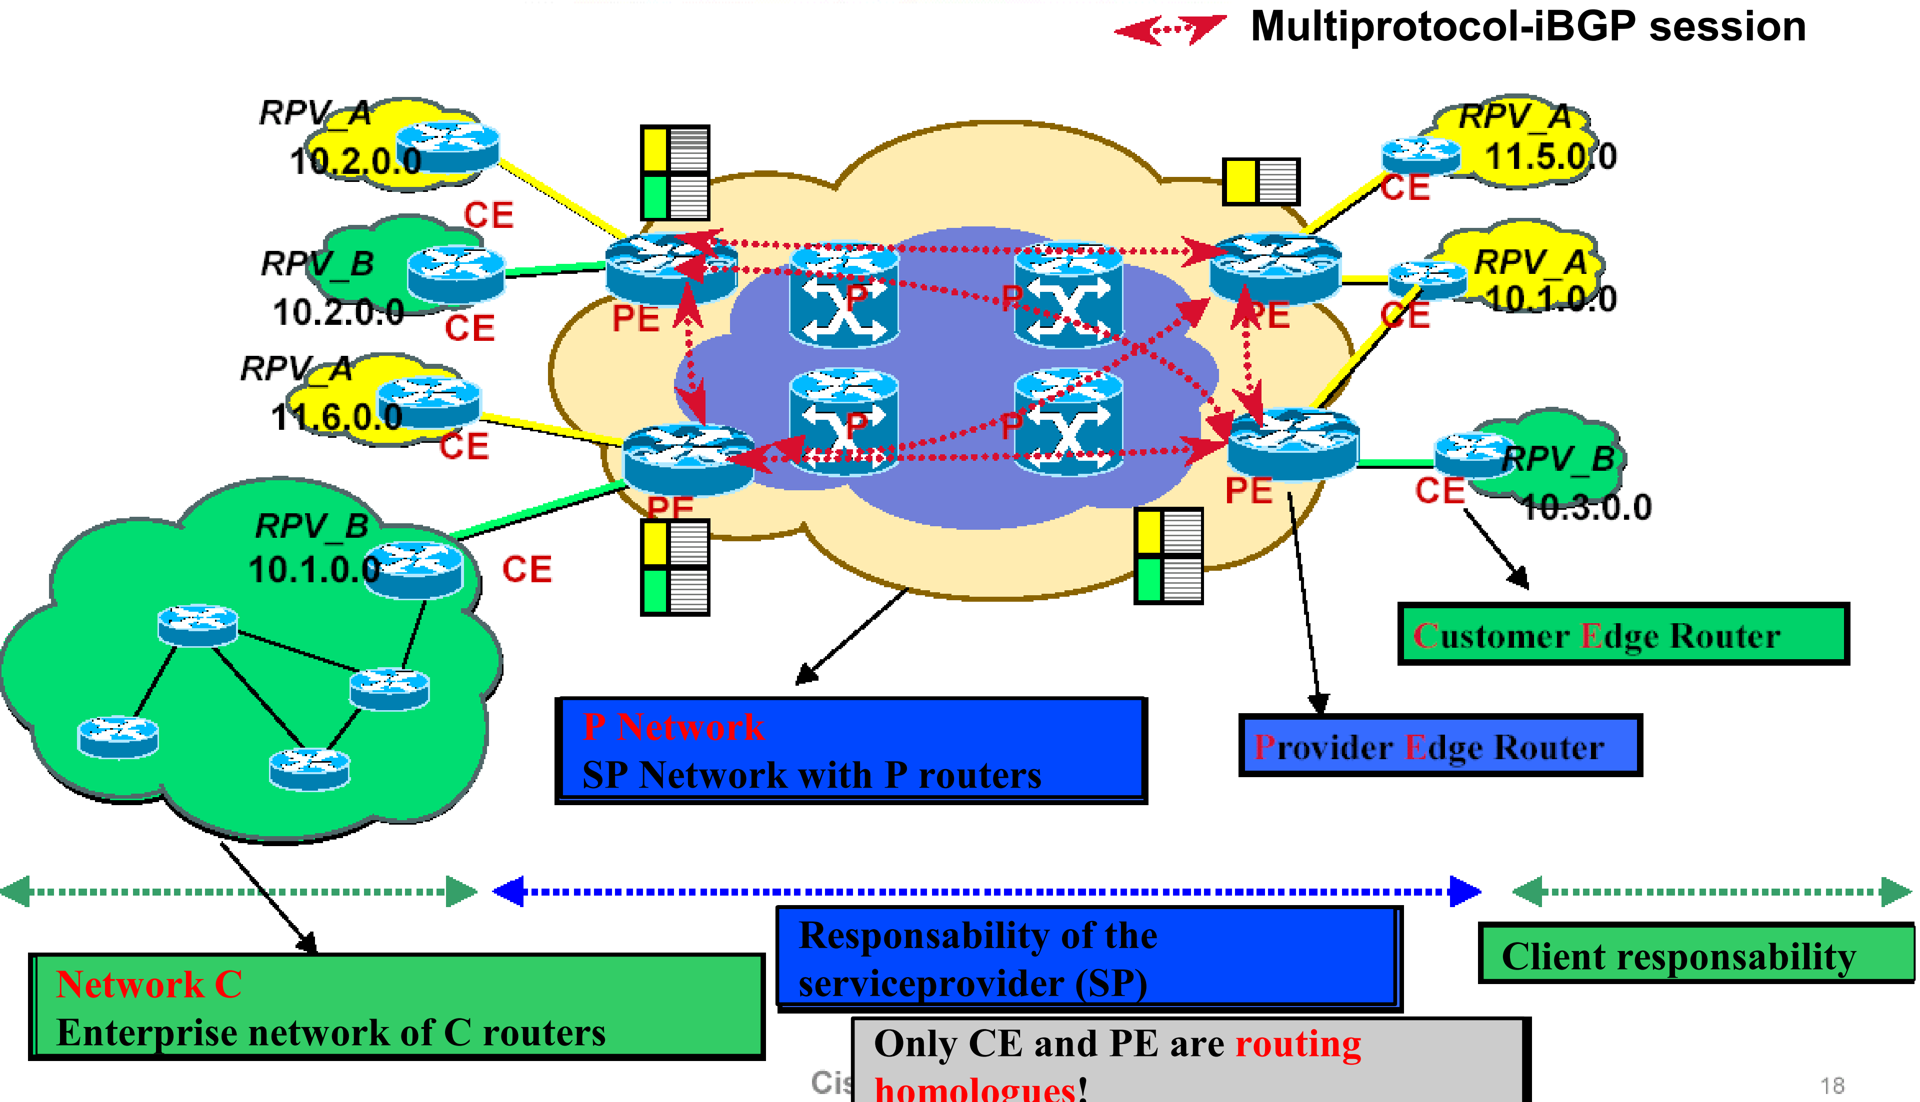

VPN with MPLS

- MPLS Label Switching Router Block
  - **Control plane**: IP + MPLS-IP Routing
  - **Data Plane**: MPLS forwarding (label switching)

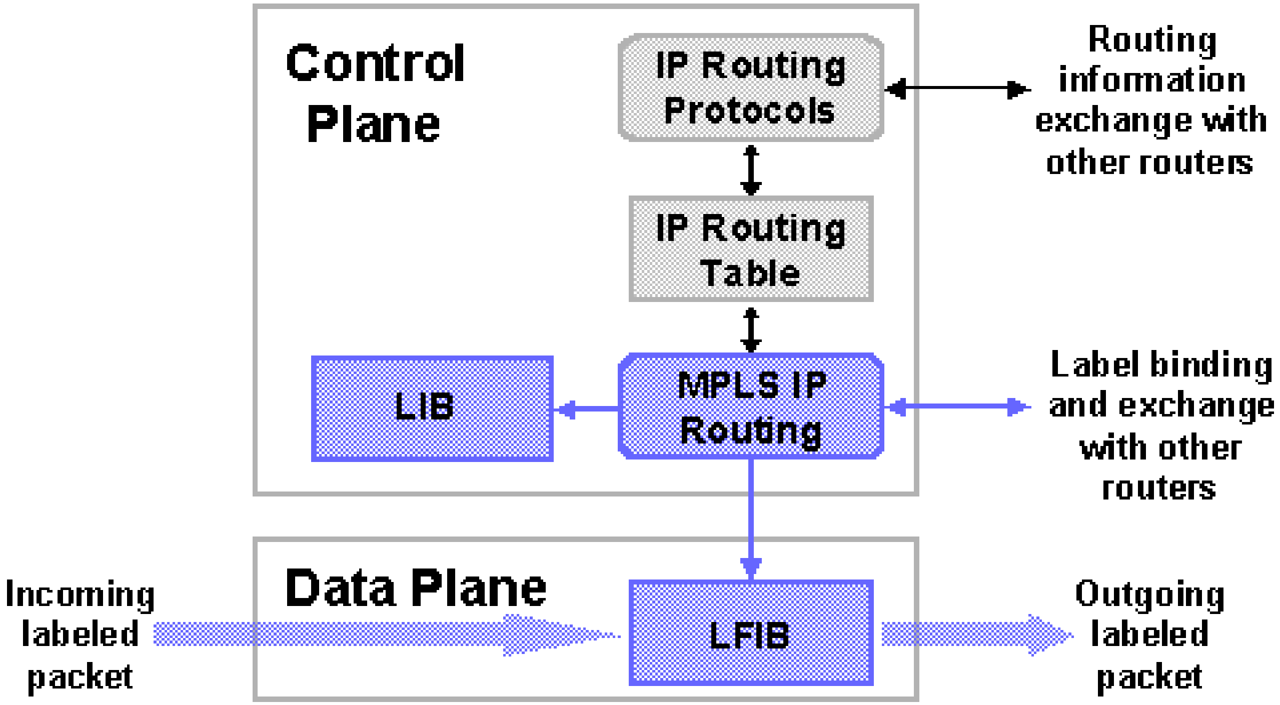

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

GMPLS

- Generalized MPLS
  - Separation of control and data planes
  - **Control plane**: IP
  - **Data plane**: multiple types of switching
    - TDM, Lambda, packet, and fiber, etc
  - Dynamic bandwidth provisioning
    - Automatic selection of the most efficient optical or wavelength connection

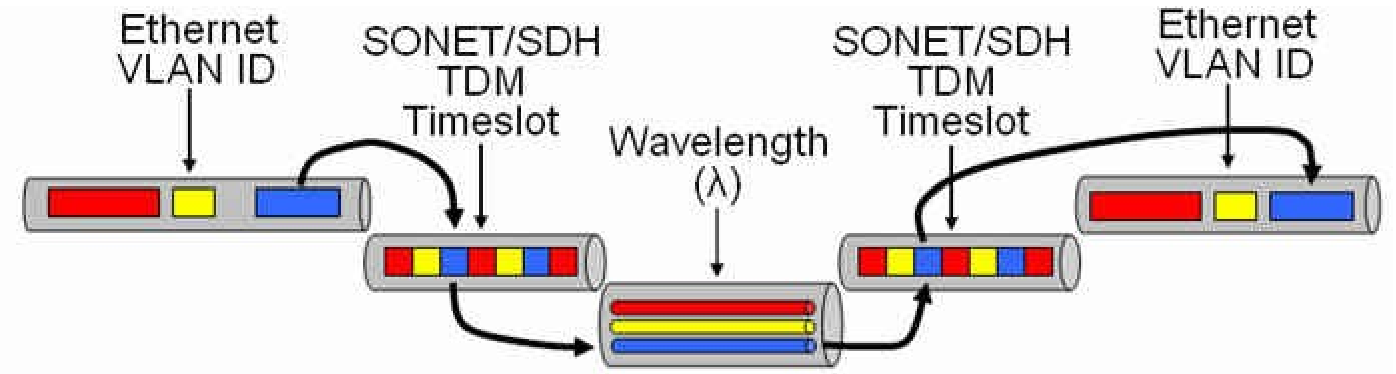

[Turnin Networks – Ralph Santitoro]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Overlay Networks [Jennifer Rexford - Princeton]

Motivation

- Problems with the underlying routing system
- Source routing, overlay networks

Argument

- The end-user has better end-to-end view

Overlay networks

- Pros: flexibility, limited overhead, & value-added
- Cons: data-path overhead, probes, & feedback

Putting More Power in End Hosts

Source routing (e.g., Nimrod)

- End host selects the end-to-end path
- Routers simply forward packets on the path
- Requires the routers to agree to participate

Overlay networks (e.g., RON)

- Conventional computers act as logical routers
- Real routers deliver packets to intermediate hosts
- No need for cooperation from the real routers

Hybrid schemes

- Source routing at the AS level
- Source routing in the overlay network

Overlay Network

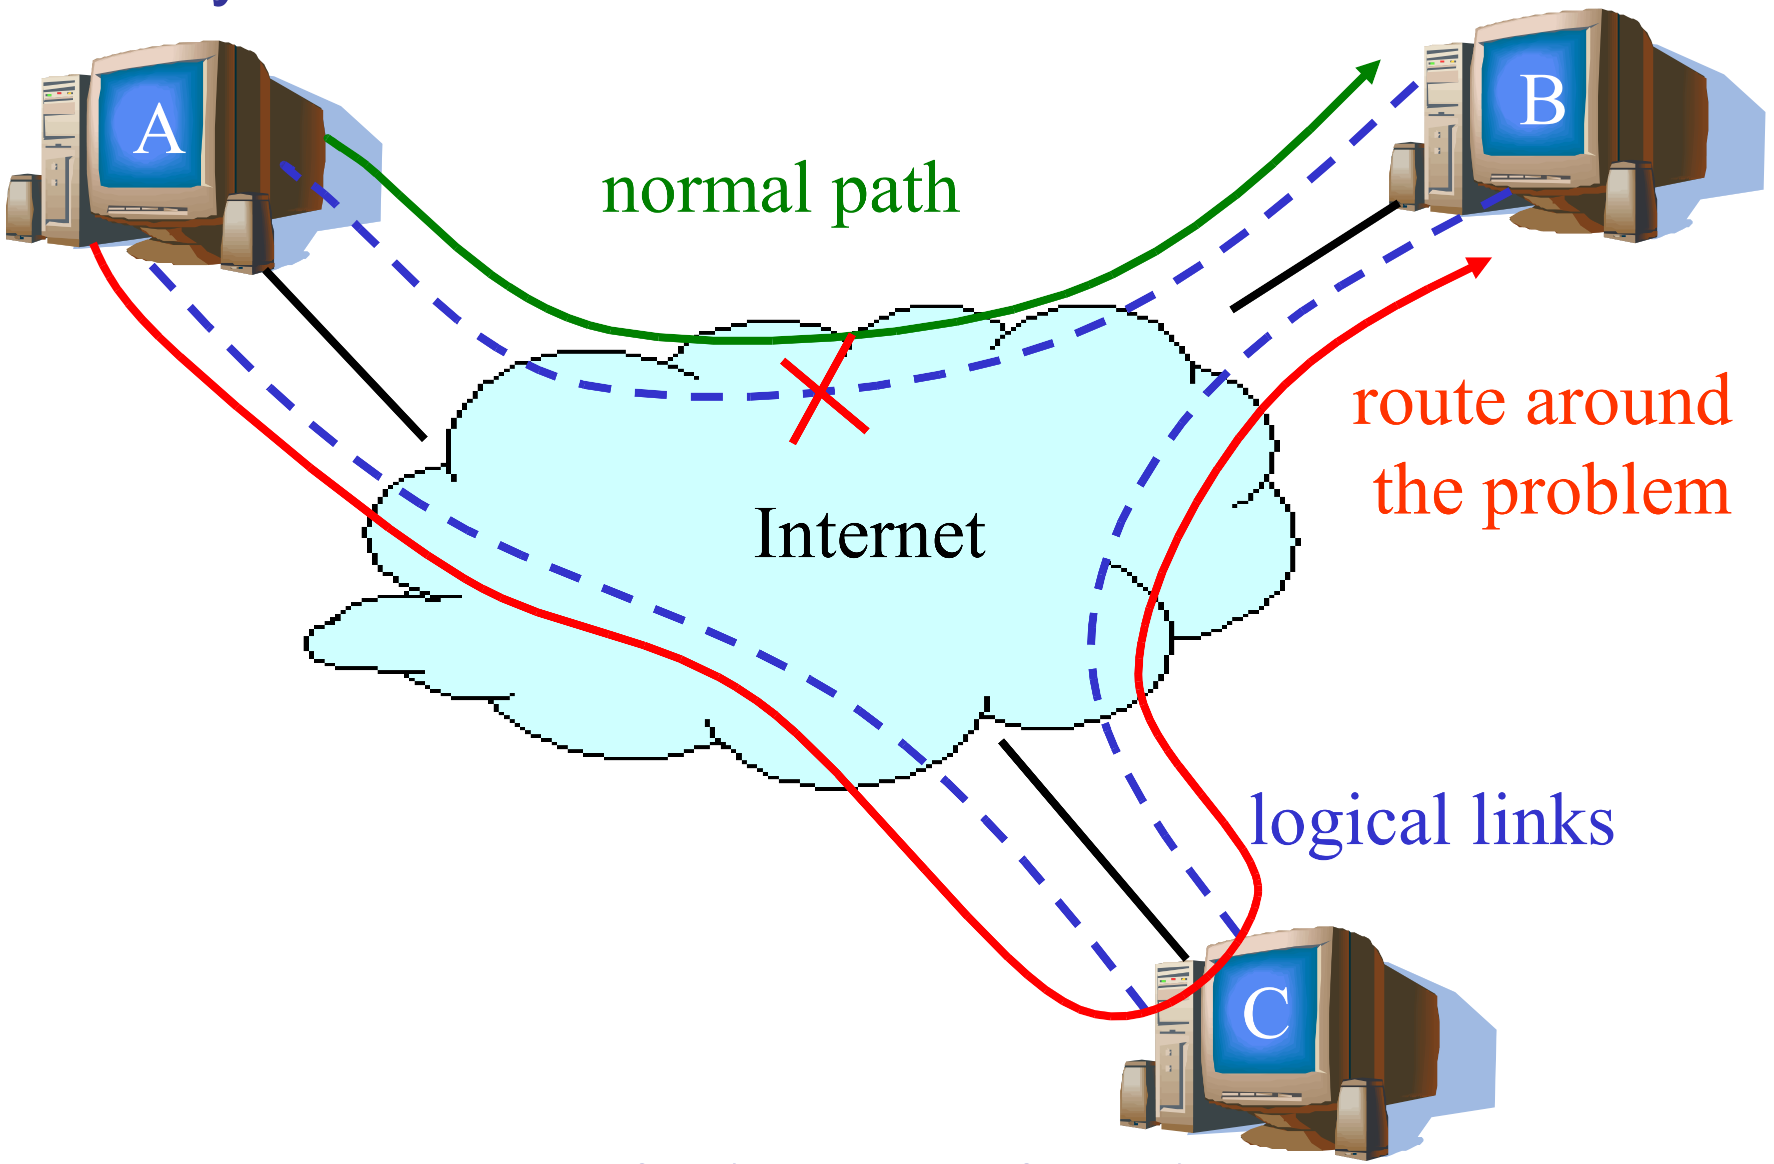

Overlay Network : QoS

- Bandwidth-aware routing (BARON)
  - Monitor the network with overlay nodes to optimize bandwidth
- OverQos: Bandwidth/loss-rate control

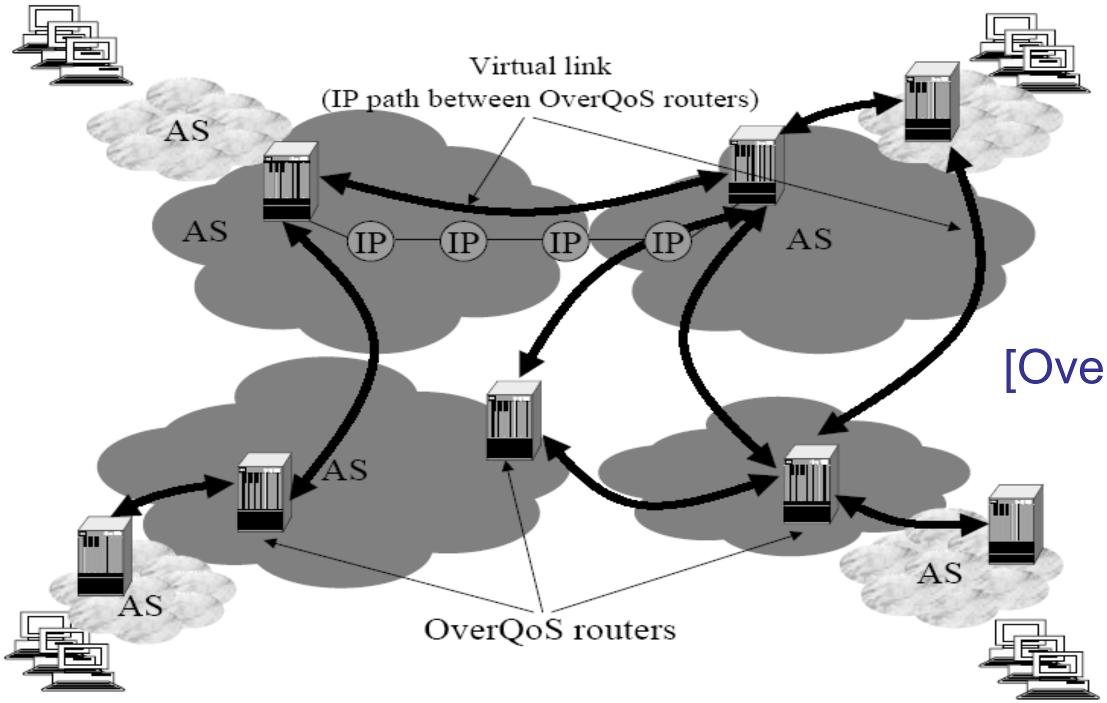

[OverQoS – L. Subramanina et al.]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Overlay Networks : Infrastructures on demand

- Create on-demand Network infrastructures with overlay nodes (end-hosts)
  - e. g. the X-Bone

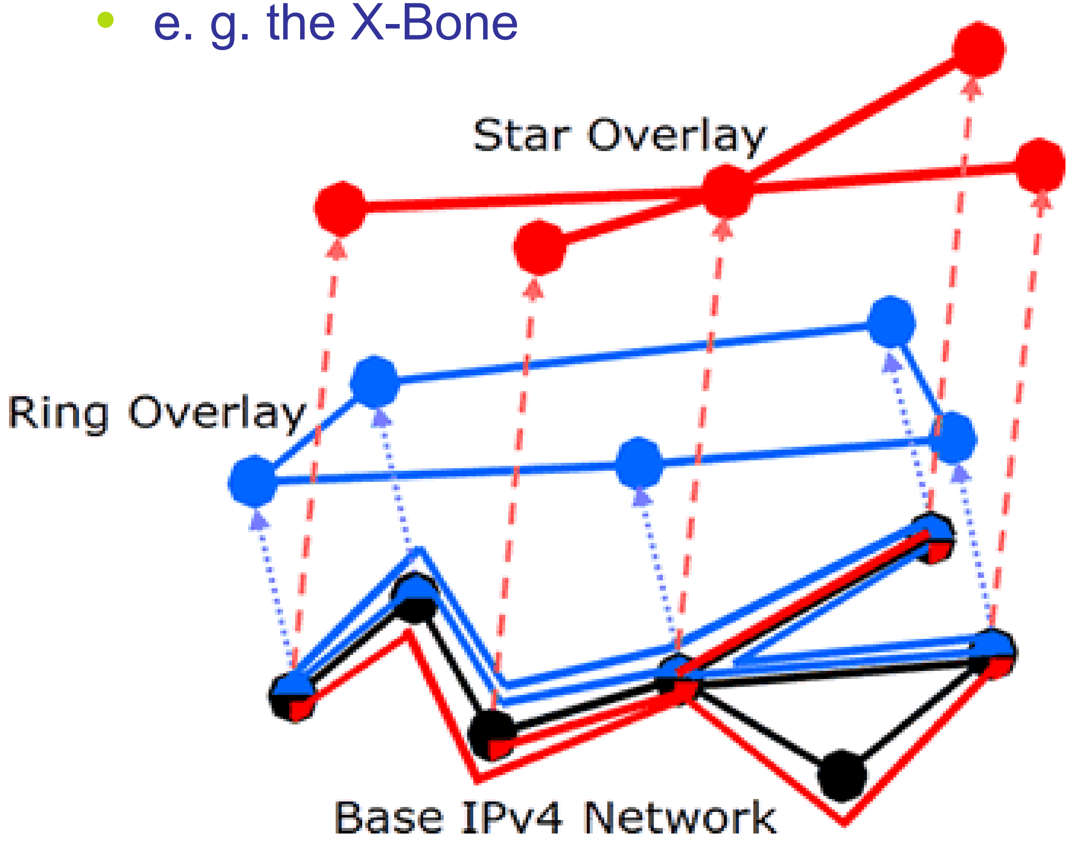

[http://www.isi.edu/xbone/]

Goals :

- To test new protocols in « real »
- Partition capacity
- Present simplified topologies

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Network virtualization benefits

- Network virtualization solves a security problem
  - L2 & L3 VPNs - Ipsec tunnels

- Link Virtualization solves a cost problem :
  - For offering VPNs to many users (capex, opex)
    - Splitting available physical bandwidth into independant channels (L2 VPN)
    - L1 resource sharing

- Path Virtualization accelerates the forwarding process
  - Combining channels into Path with « predefined routing » (ex MPLS, λ path)
  - Circuit-switching paradigm emulation (LSP)

- Overlays networks used to optimize IP routing problems (reliability)
  - Virtualisation of the data path and of the routing function
  - Key point: end user has a better end to end view
  - Negative interactions
    - With other overlays: the price of anarchy
    - With the underlay: influence on traffic engineering

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtualization of the underlay

- Classical approaches: **link & path virtualisation**
  - VPN, VLAN
- Overlay approach : end user path control

- New approach : **Network Devices Virtualization**
  - Logical routers, virtual routers
  - Benefits
    - Isolation/security
    - Reconfigurability/ customization
    - Mobility / Fault tolerance
    - Increases « overlay » performance
  - Overheads:
    - Processing
    - Local resource sharing

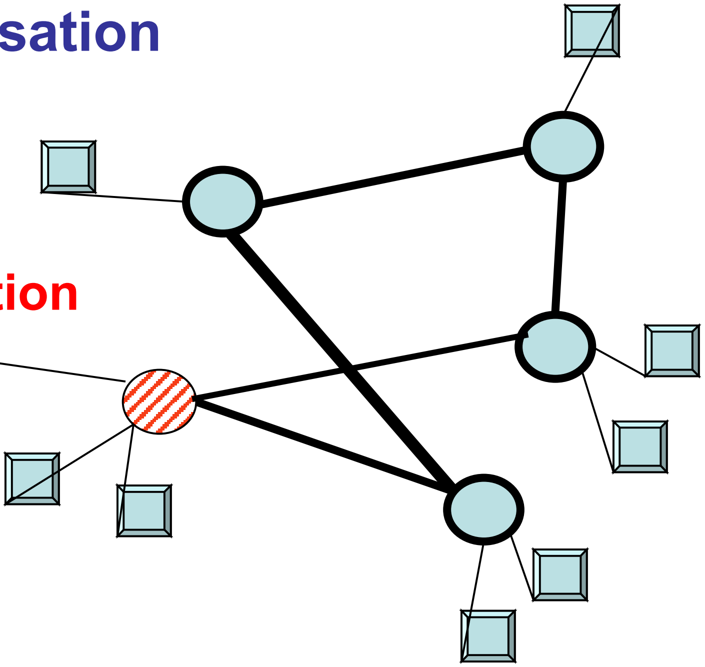

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtualization of the underlay

- New approach : **Virtualization of the network core devices**

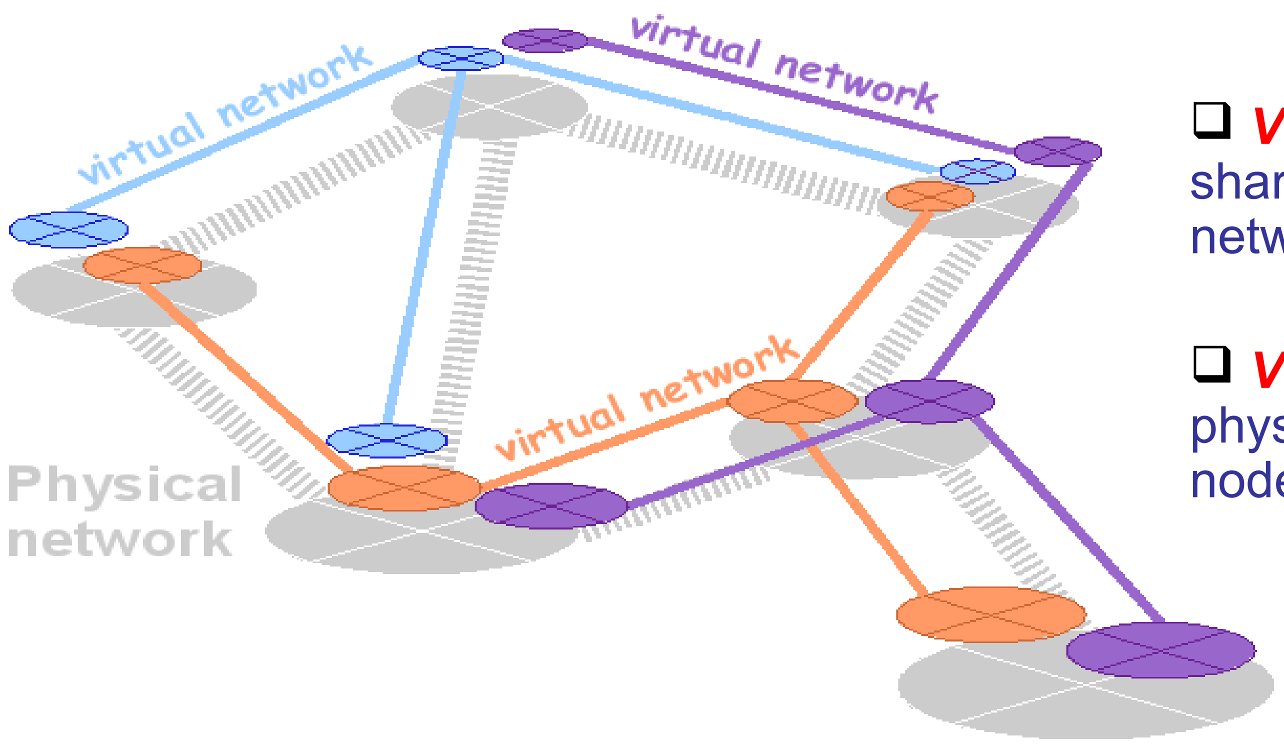

- ***Virtual networks*** share the physical network
- ***Virtual routers*** share physical interconnection nodes

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

APN and virtualized infrastructures

- Emerging approach : **link/path + device virtualization**
  - Articulated private network (APN)
    - through web services
  - User-defined autonomous systems
    - UCLP, Manticore, Feredica
  - Virtual infrastructures
    - Cabo, HIPerNet

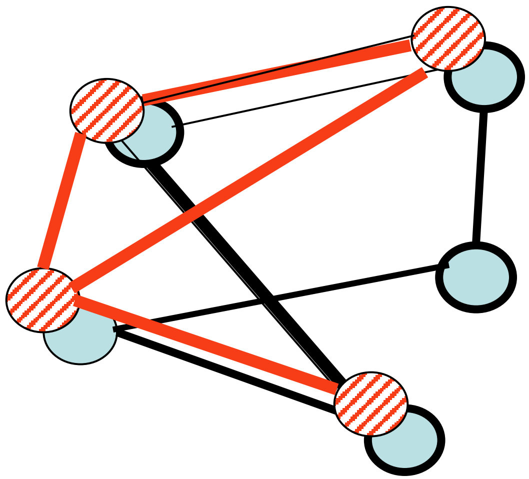

•**Key points:**

- **All network elements seen as resources (links & devices)**
- **Full separation of the infrastructure level and the service level**
- **End to end (or SP?) control and ownership of resources**

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtualization in networks

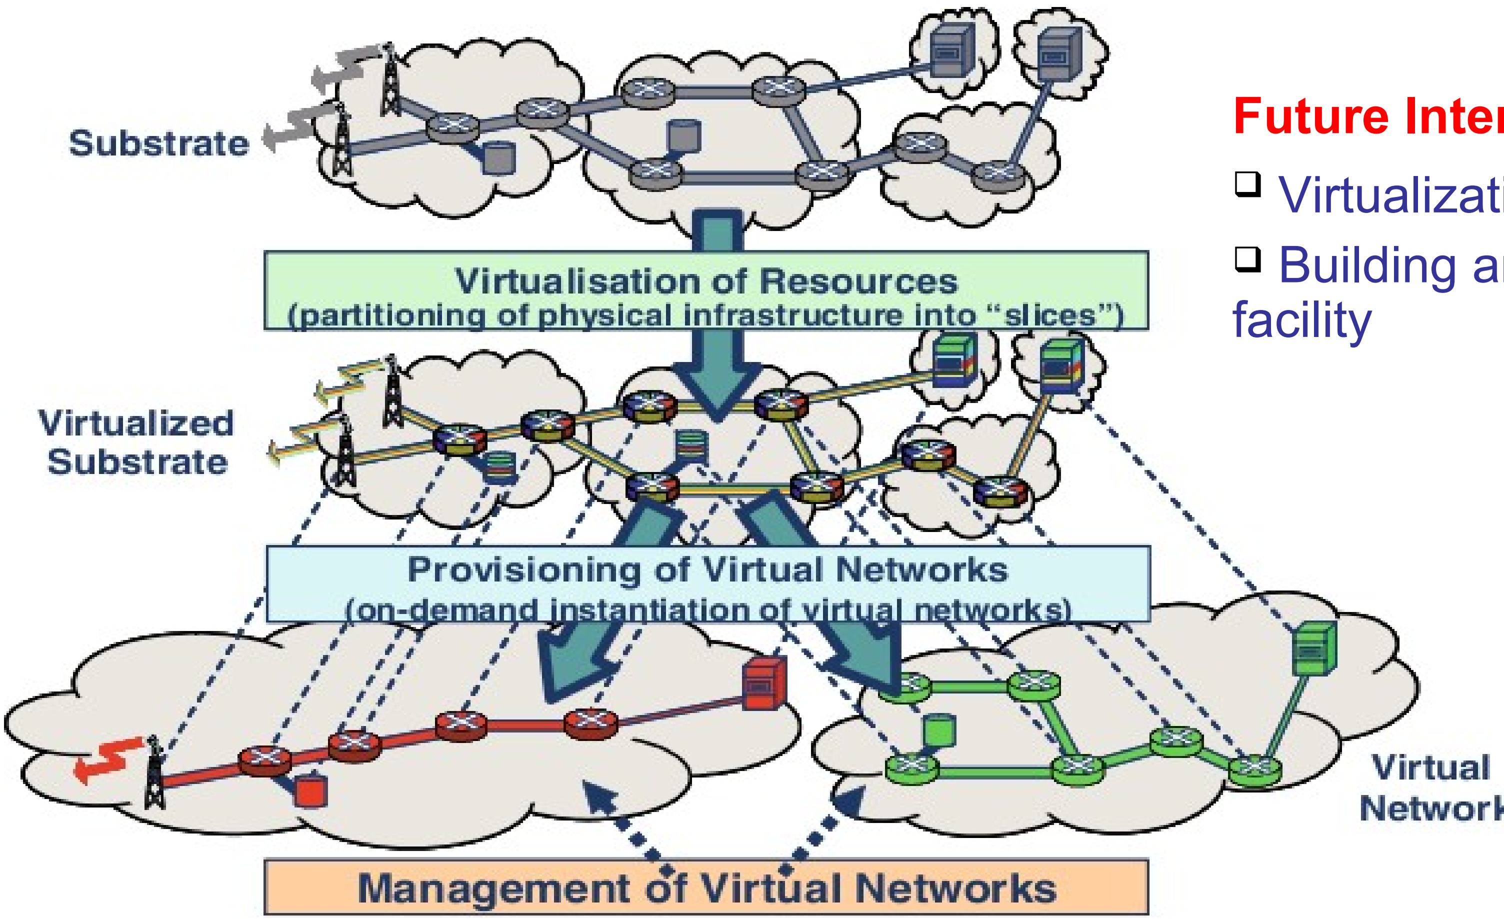

Future Internet

- Virtualization Framework
- Building an experimental facility

[Network Virtualization: A Viable Path Towards the Future Internet – N. Niebert et al.]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Benefits of « net » virtualization

1. **Security**: provide a confined network where non-trusted « protocols » can be run;
2. **Isolation**: limit network resource access and usage, through isolation techniques, or expand it transparently for the applications
3. **Appliance**: adapt the « network » to the application instead of porting the application to the « network» environment;
4. **Customization** : use dedicated or optimized mechanisms (schedulers, buffer management, network protocol) for each application;
5. **Ease management**: manage as a whole applications and processes running within a virtual network.
6. **Cost reduction**: resource consolidation, load balacing, dynamic resizing, resource provisioning, power saving
7. **Ease development process**: Testing, experimenting

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Outline

- Virtualization concept
- Virtualization in networks
  - Why and how
  - History of virtualization in network
  - Virtualizing the underlay
- Examples
  - PlanetLab/VINI/Cabo
  - GENI/Federica/Manticore
  - Carriocas/HIPerNet
- Emerging building block
  - Virtual core network resources
- Virtual networks on Grid'5000
- Conclusion

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

PlanetLab
[Larry Peterson - Princeton University]

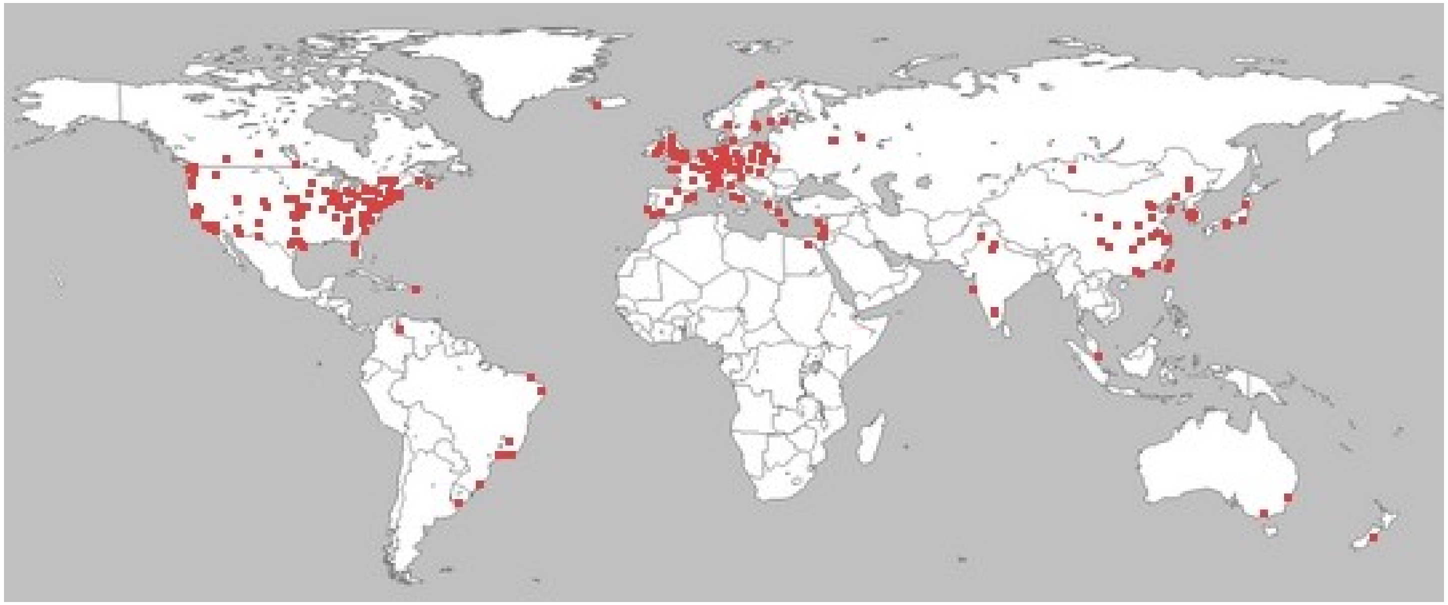

- 800+ machines spanning 400 sites and 40 countries
- Supports *distributed virtualization*

each of 600+ network services running in their own *slice*

http://coblitz.org/www.princeton.edu/

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Slices

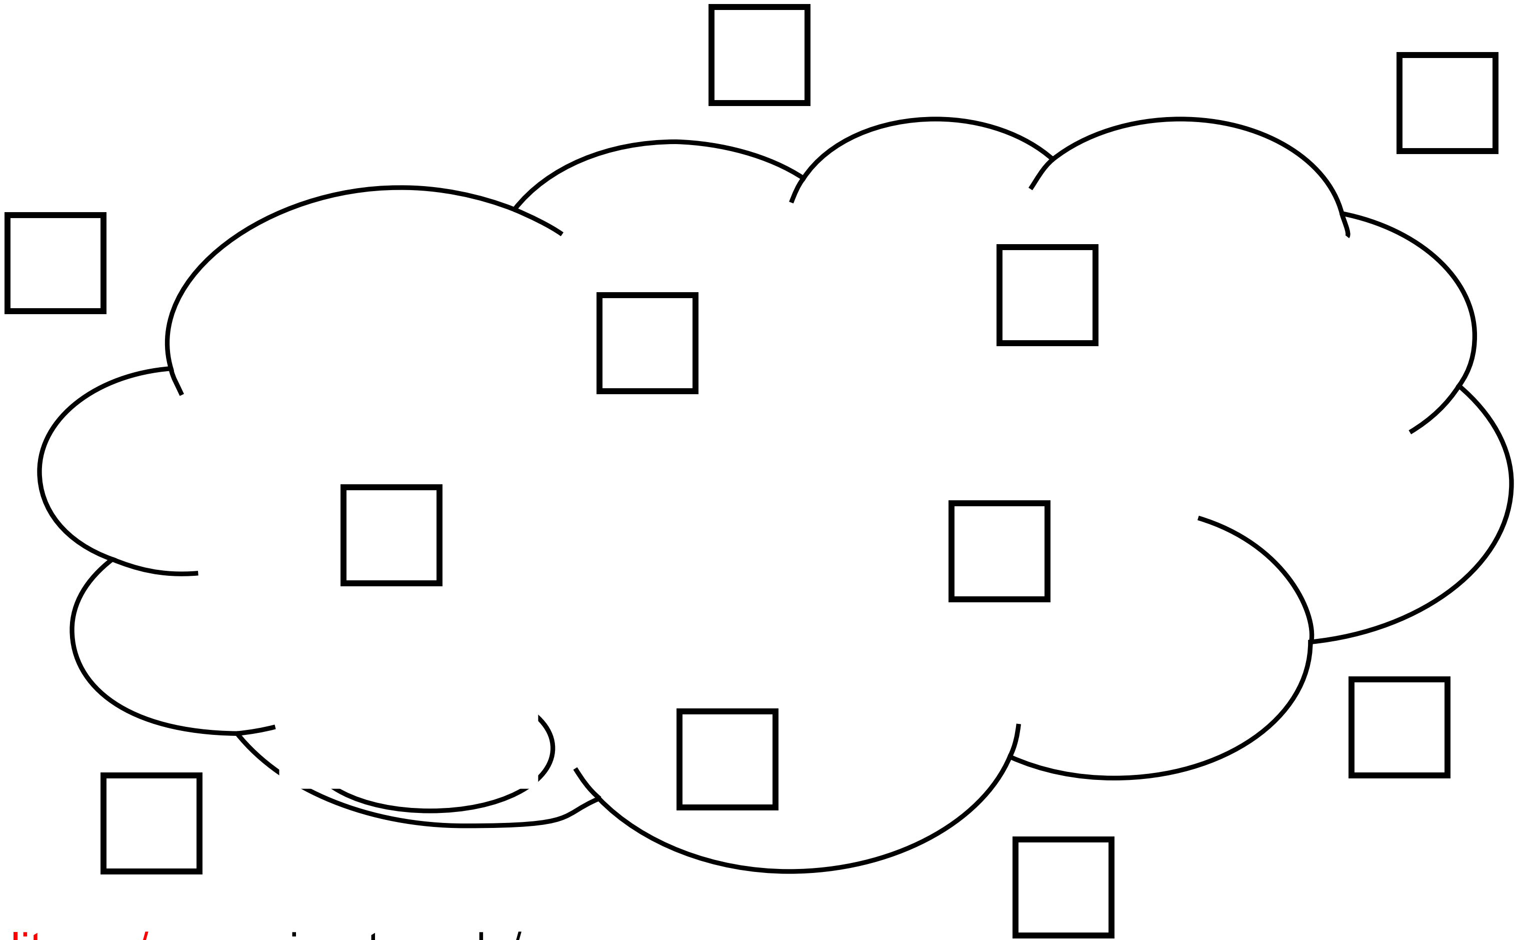

http://coblitz.org/www.princeton.edu/

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Slices

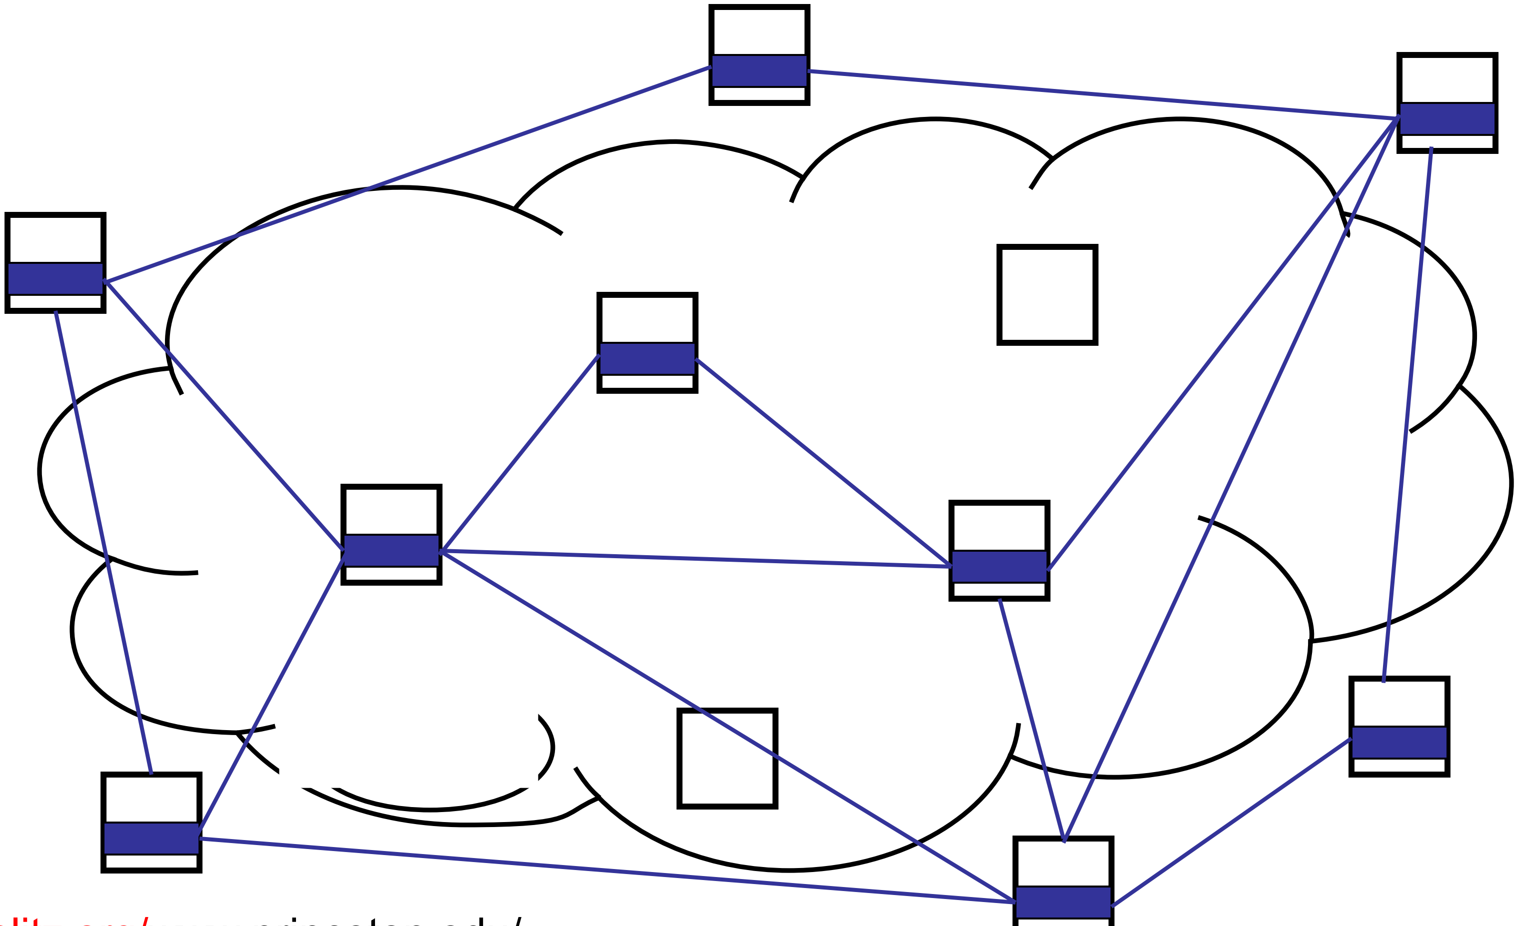

http://coblitz.org/www.princeton.edu/

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Slices

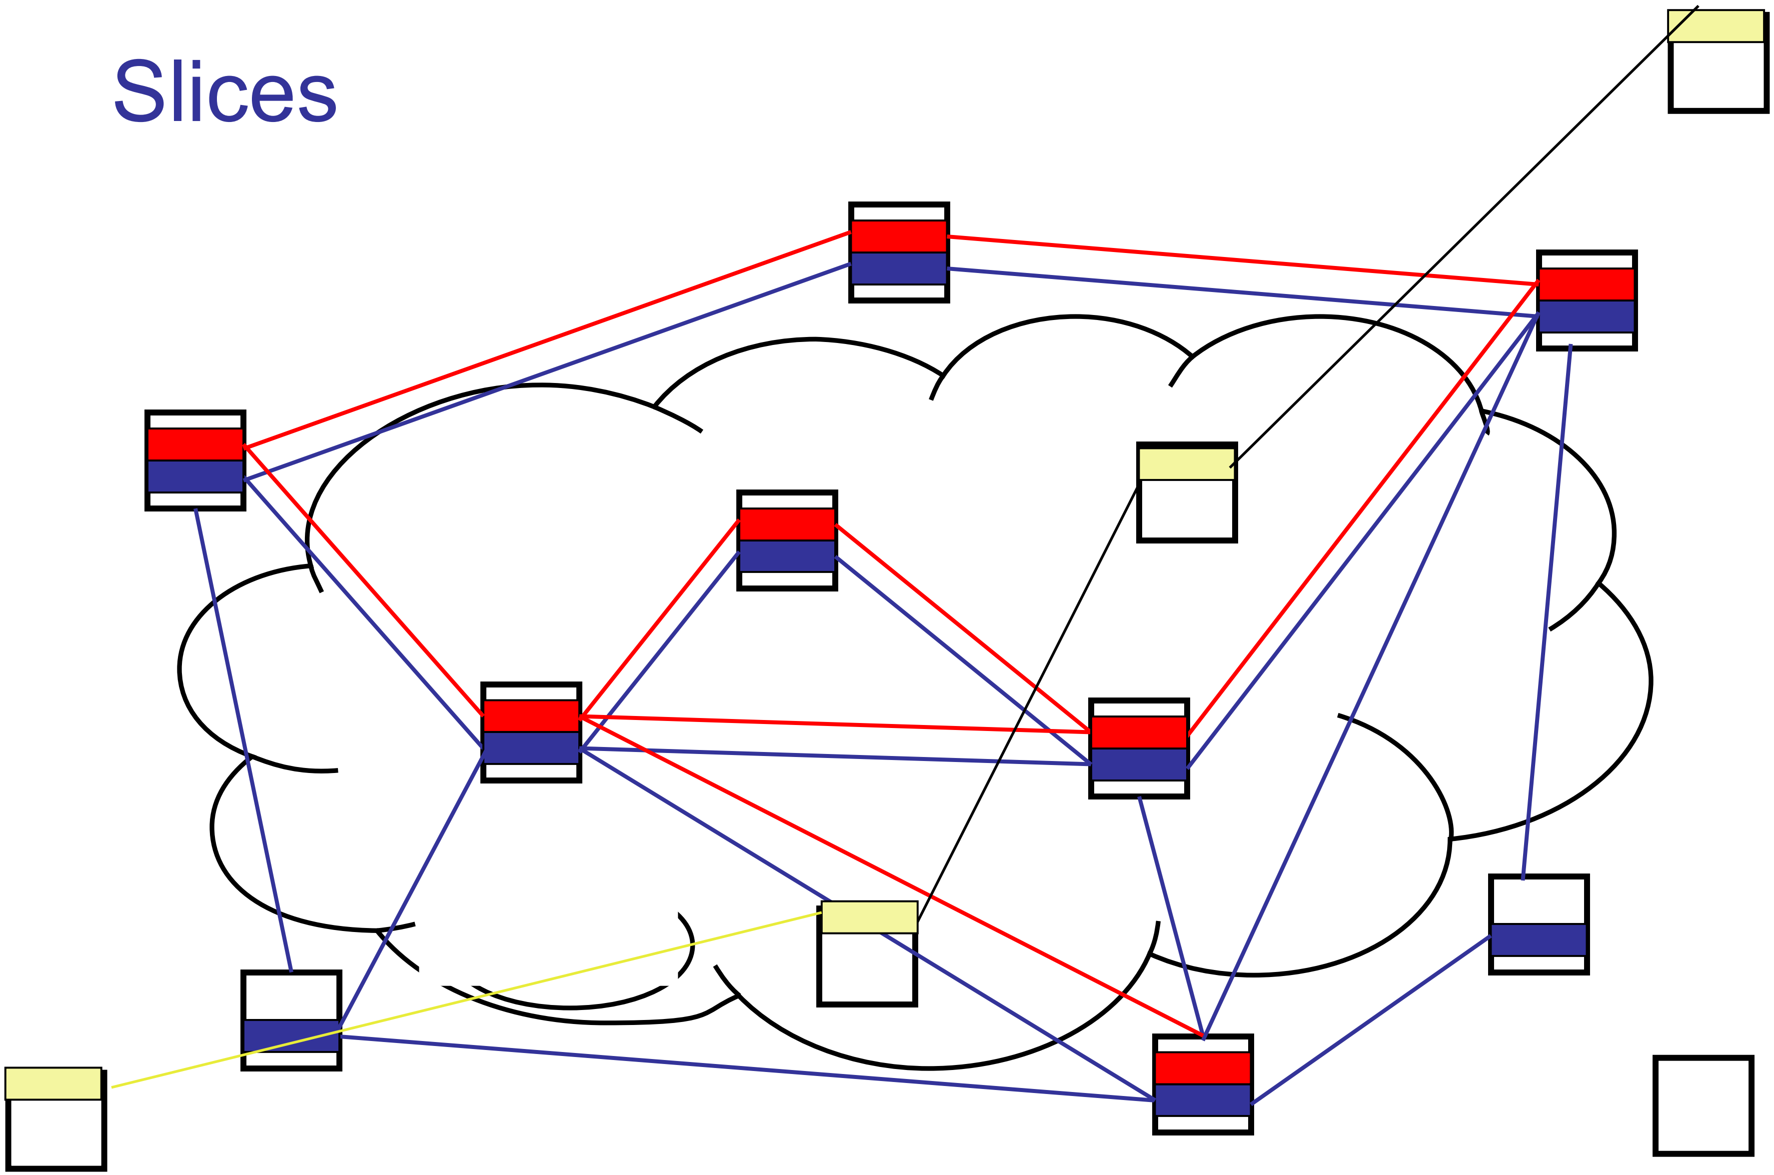

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

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

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

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

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

VINI: A Virtual Network Infrastructure

[Andy Bavier slides - Princeton University]

Balance reproducibility and reality

- Dedicated resources for large-scale networking experiments like IIAS
- **Expose L2 circuits to slices**
- OS support for high speed packet forwarding

Federate with public PlanetLab

- Moving away from PlanetLab's *best effort* model
- Experiment with new policies, kernels
- Approach federation problem from both sides

Software toolkit for network experimentation

- Observe routing protocols "in the wild"
- Evaluate new protocols and proposed changes
- Carry real traffic on behalf of real clients
- Hands-on teaching aid

Leverage research, open source:

- Click modular software router
- XORP routing protocol suite
- OpenVPN, User-Mode Linux

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

VINI: IIAS

[Andy Bavier slides - Princeton University]

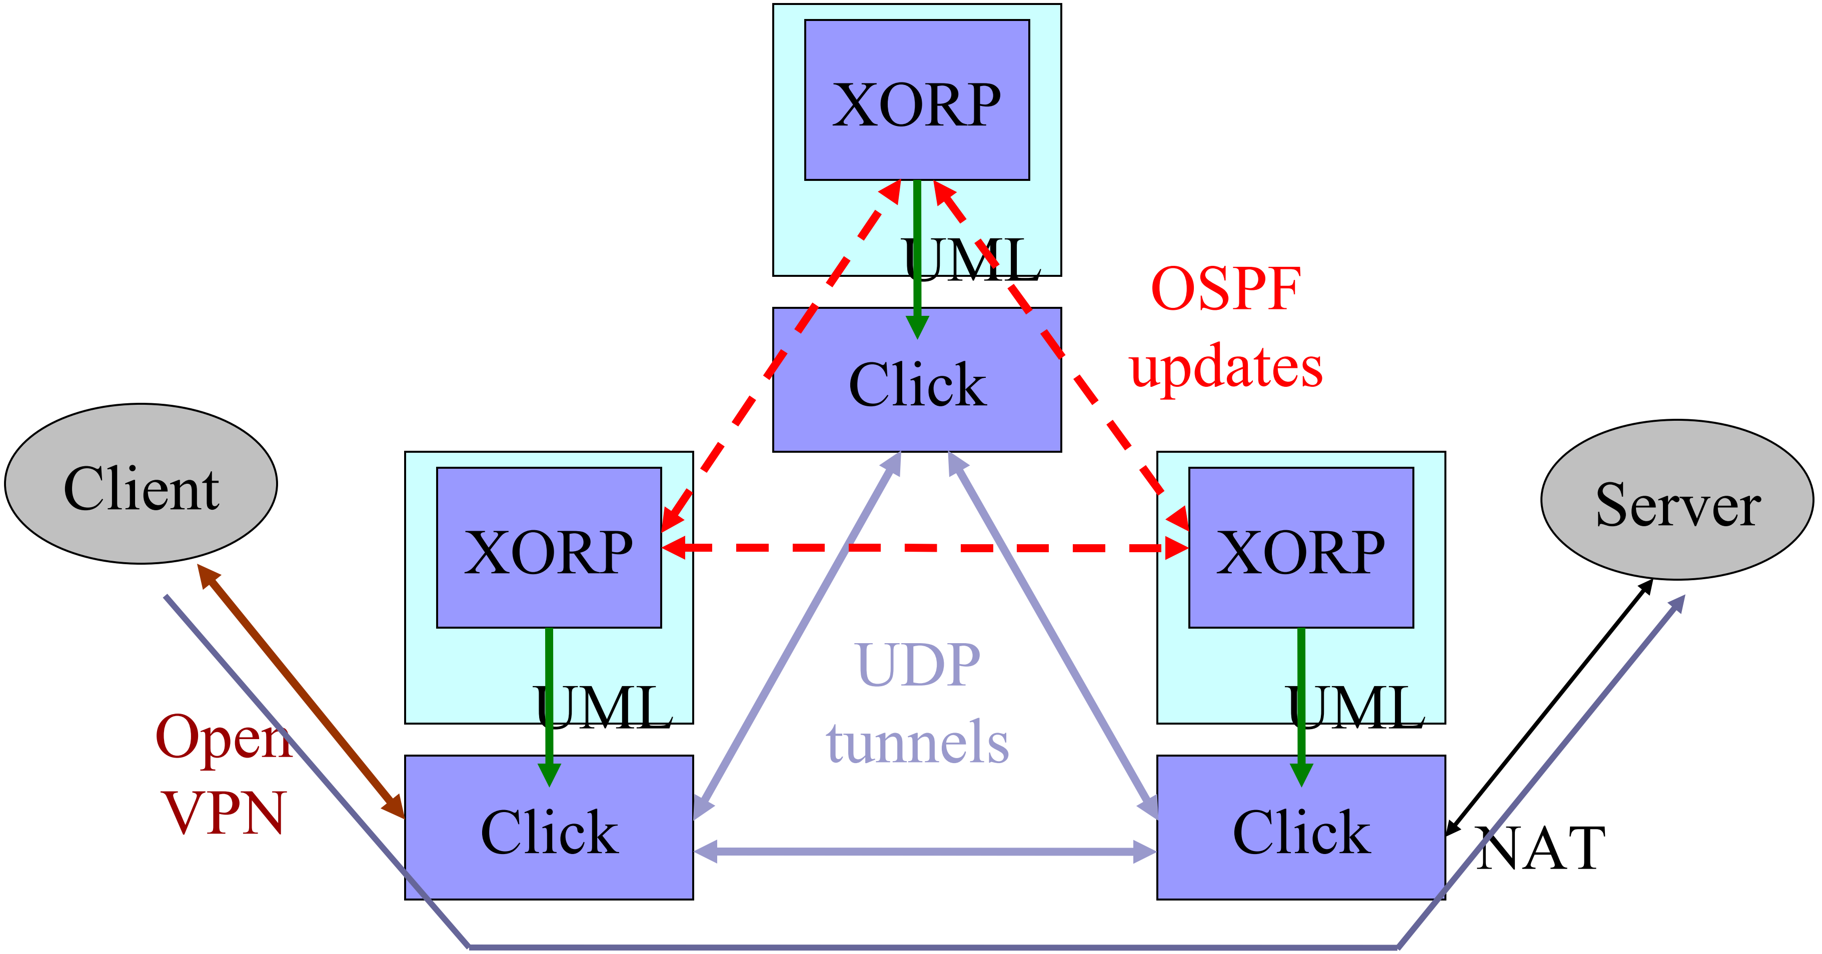

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

GENI (Global Environment for Network Innovations)

- For realistic network experiments in end-to-end virtualized slices

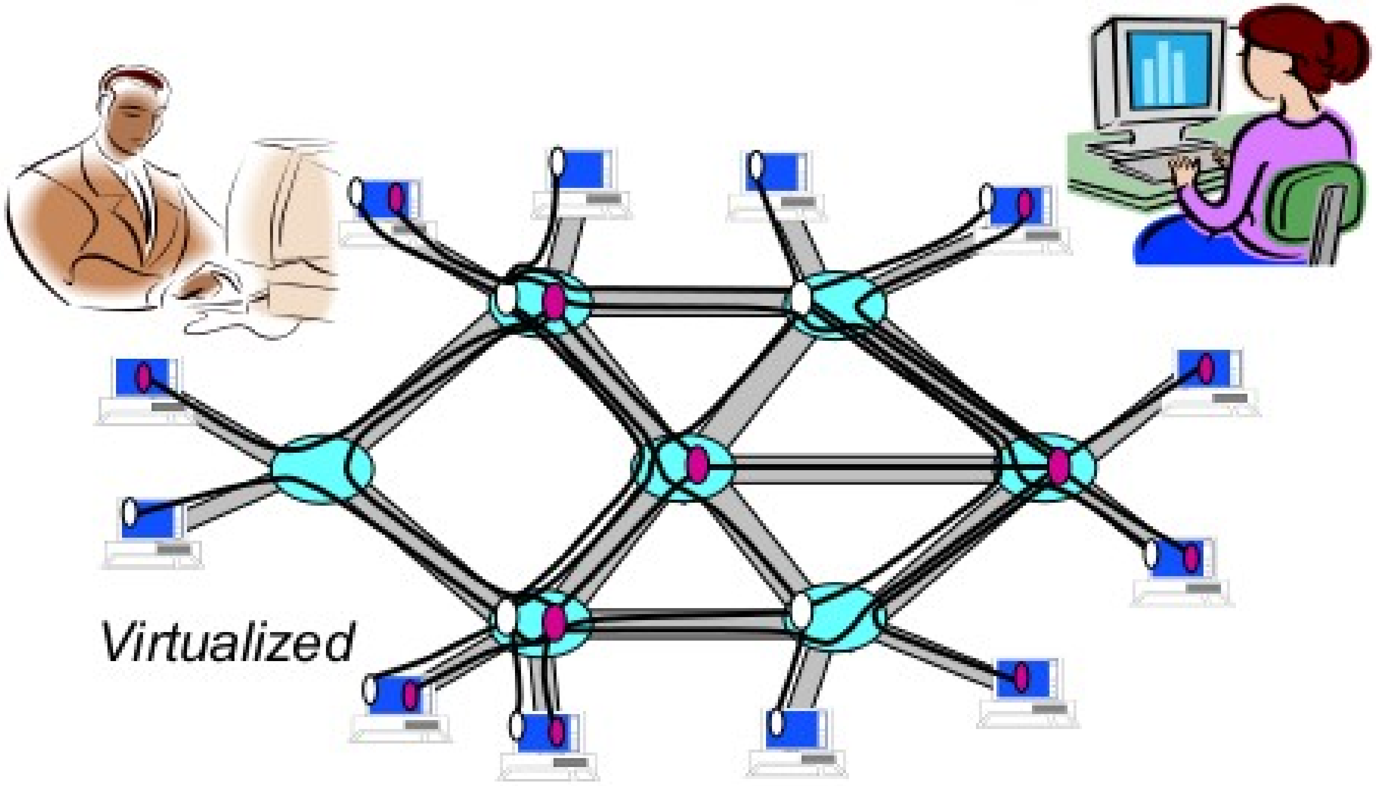

[www.geni.net]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Example: MANTICORE-II

(Victor Reijs, HEAnet)

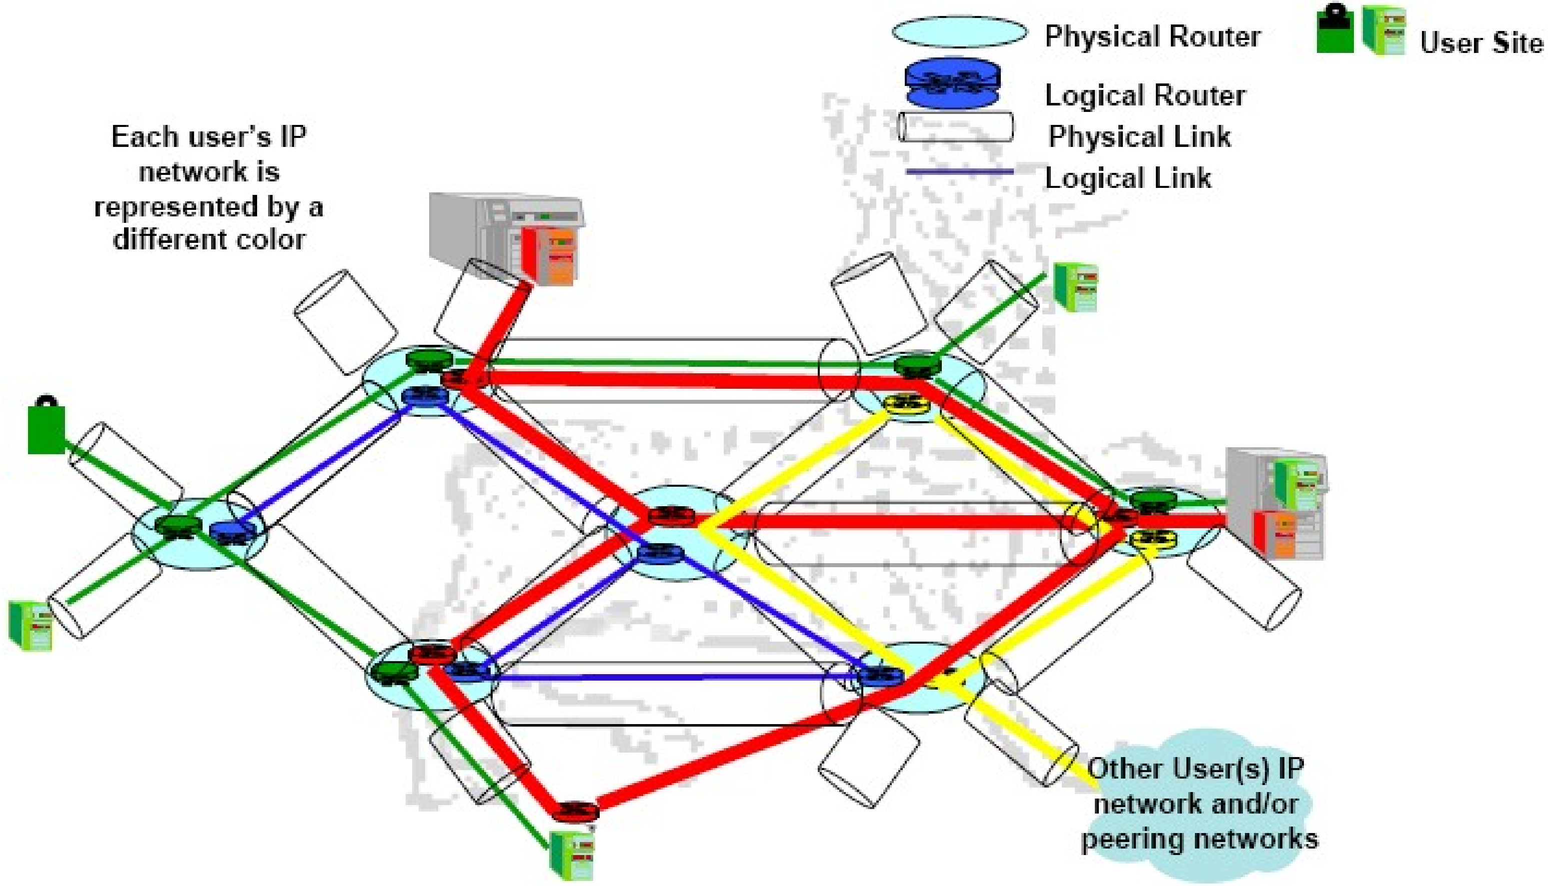

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

FEDERICA

(Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures)

- Provide Virtual network infrastructures over production networks for testing

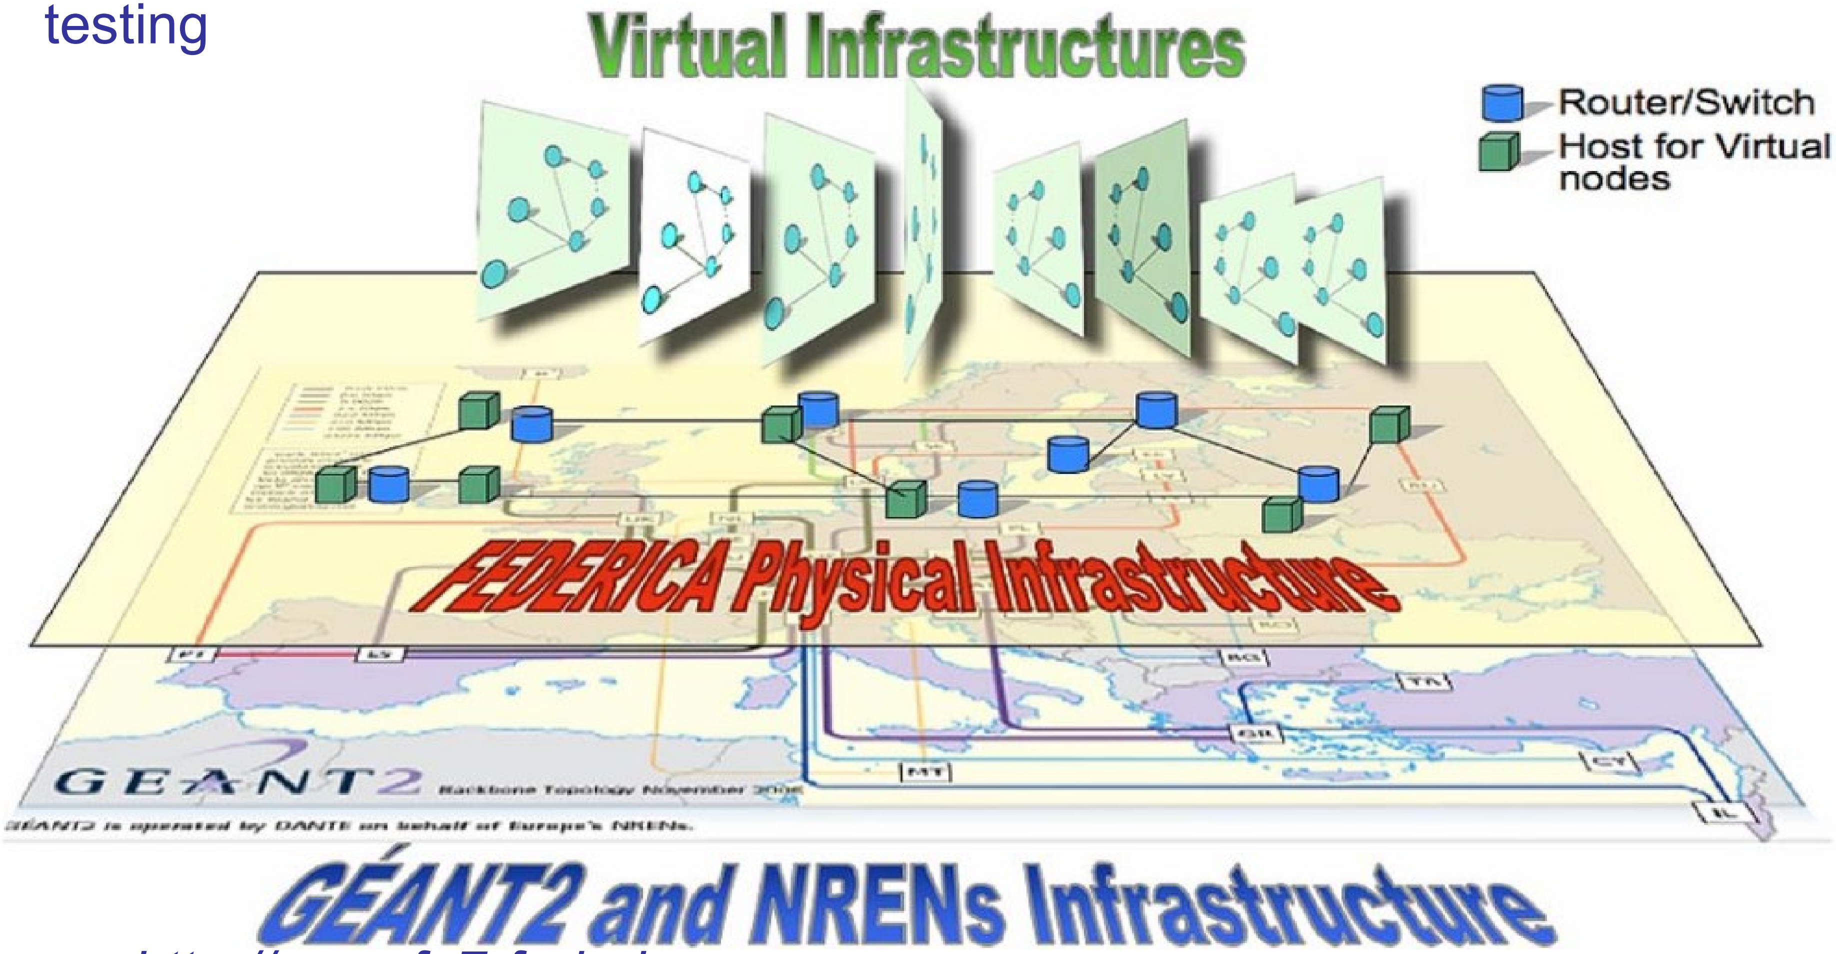

[http://www.fp7-federica.eut]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

CARRIOCAS virtualization layer [INRIA,ALU..]

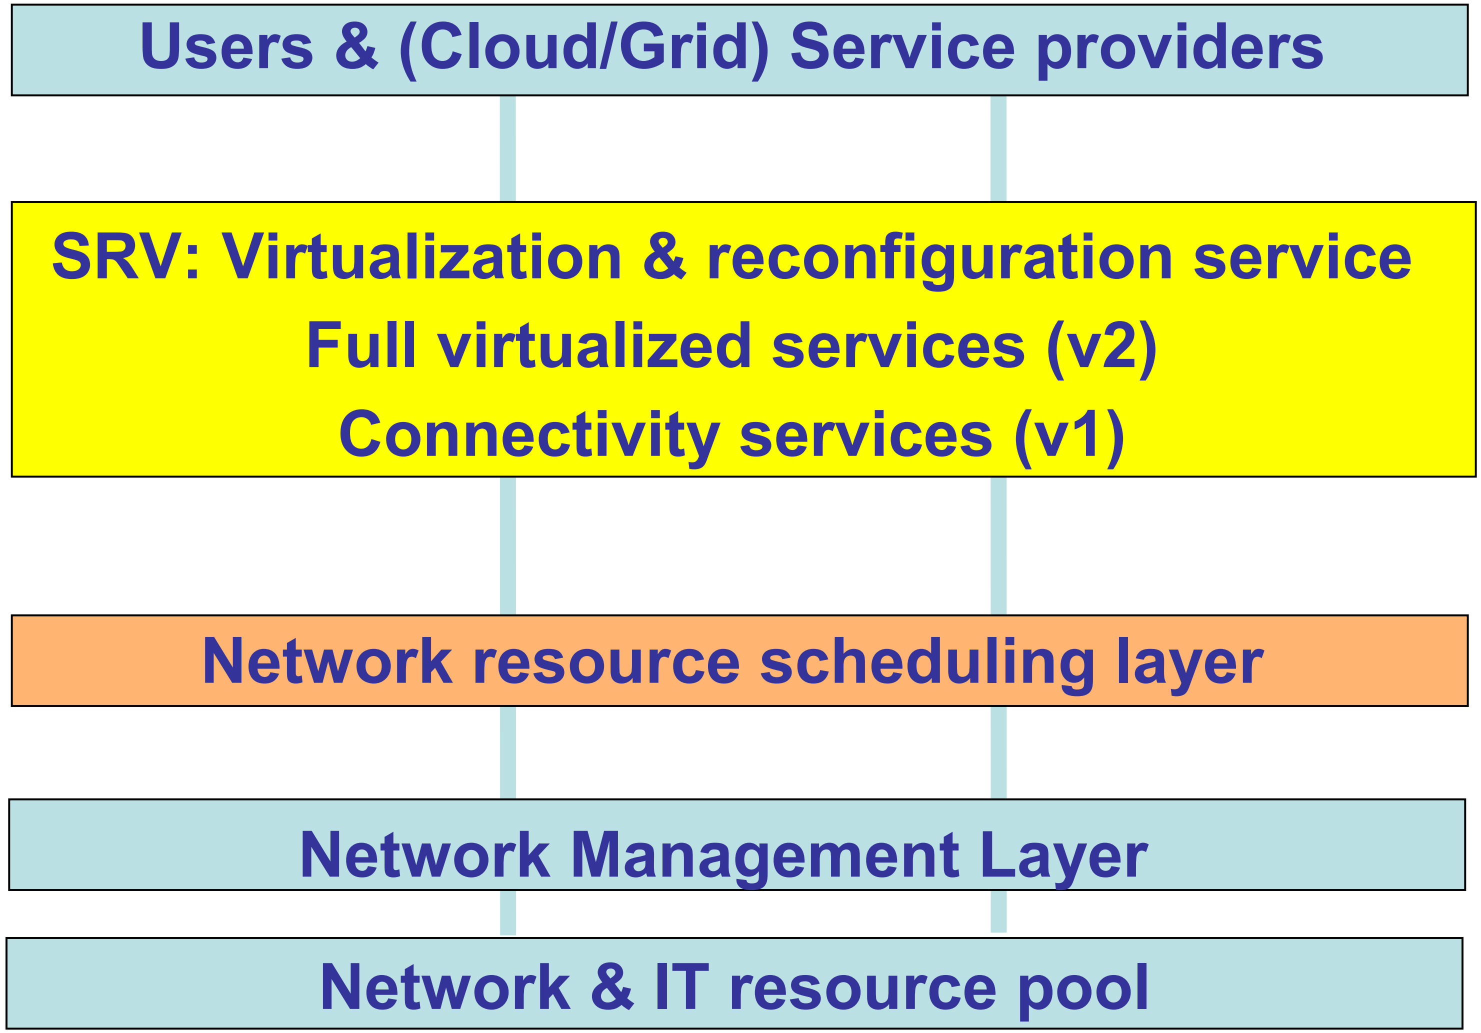

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

33

CARRIOCAS Pilot Network and test bed (FT)

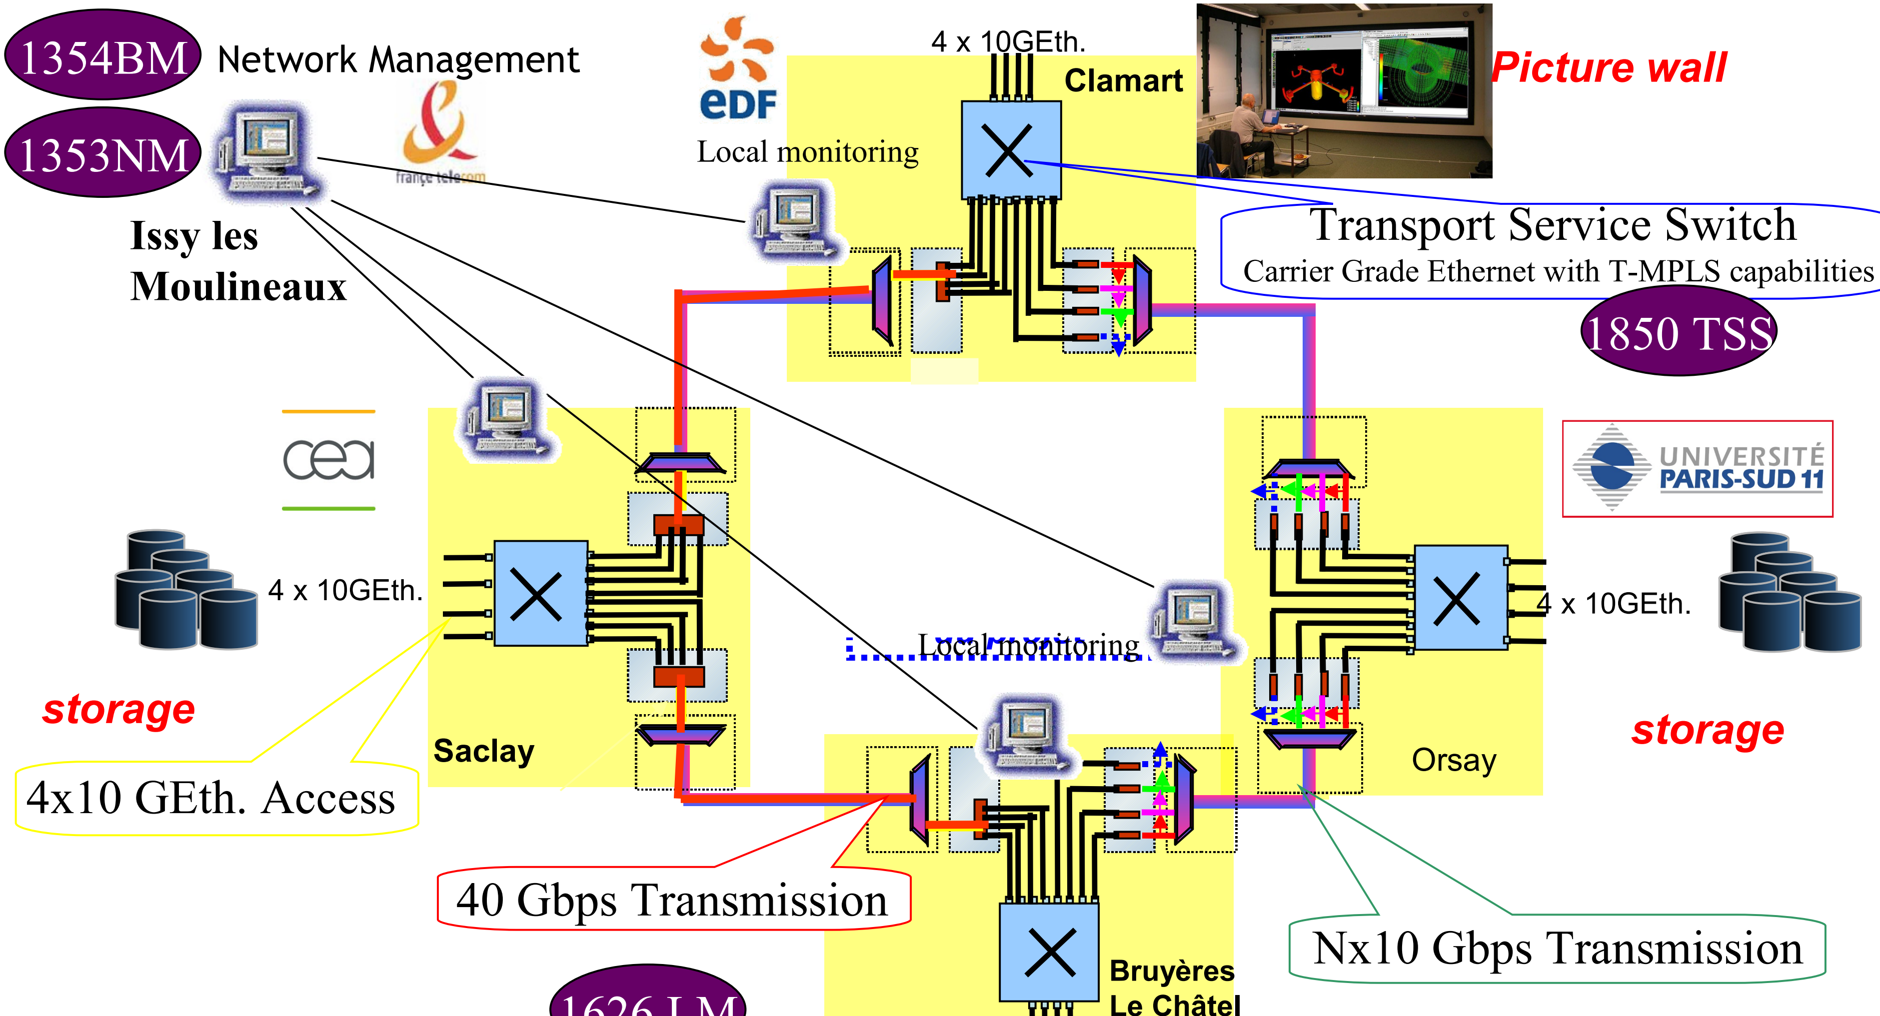

Deployment phase 1 (10Gb/s): June 2007

Deployment phase 2 (40Gb/s): June 2008

Cabo: “Concurrent Architectures are Better then One”

[How to Lease the Internet in Your Spare Time – N. Feamster et al.]

A new **architecture** for the Internet

- **An architecture which supports virtual networks**
- **For more flexibility and core network control**

Infrastructures Providers ≠ Service Providers

- **Service providers share the physical infrastructure**
- **A Service provider should have end-to-end path control over all its network ressources**
- **New network protocols could be deployed at a time**
  - **Incremental deployment**

Today

Competing ISPs with different goals must coordinate

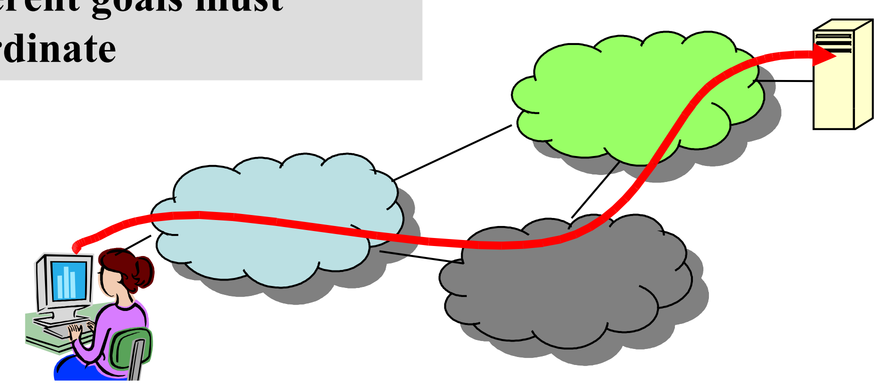

Cabo

Single service provider controls end-to-end path

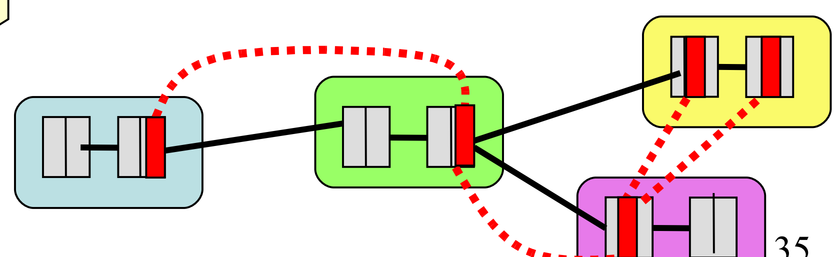

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

HIPerNet: Virtual private execution infrastructures (VPXI)

[RESO - INRIA]

- Coexistence of virtual networks

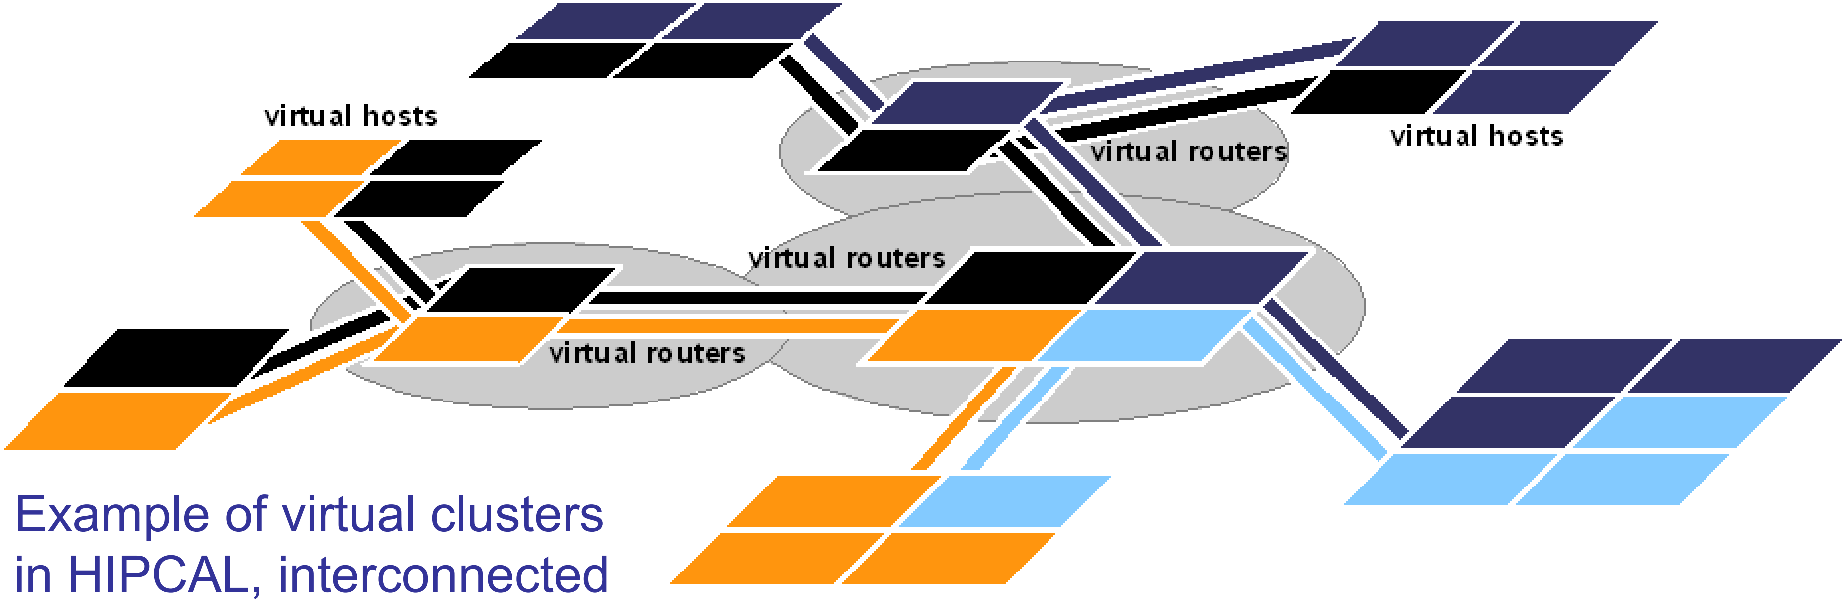

Example of virtual clusters in HIPCAL, interconnected by virtual routers

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Outline

- Virtualization concept
- Virtualization in networks
  - Why and how
  - History of virtualization in network
  - Virtualizing the underlay
- Examples
  - PlanetLab/VINI/Cabo
  - GENI/Federica/Manticore
  - Carriocas/HIPerNet
- Emerging building block
  - Server virtualization
  - Virtual core network resources
- Virtual networks on Grid'5000
- Conclusion

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Server virtualization: Virtualisation tech

KVM
VMware
} Full Virtualization

XEN (3.1, 3.2)
Denali
} Paravirtualization

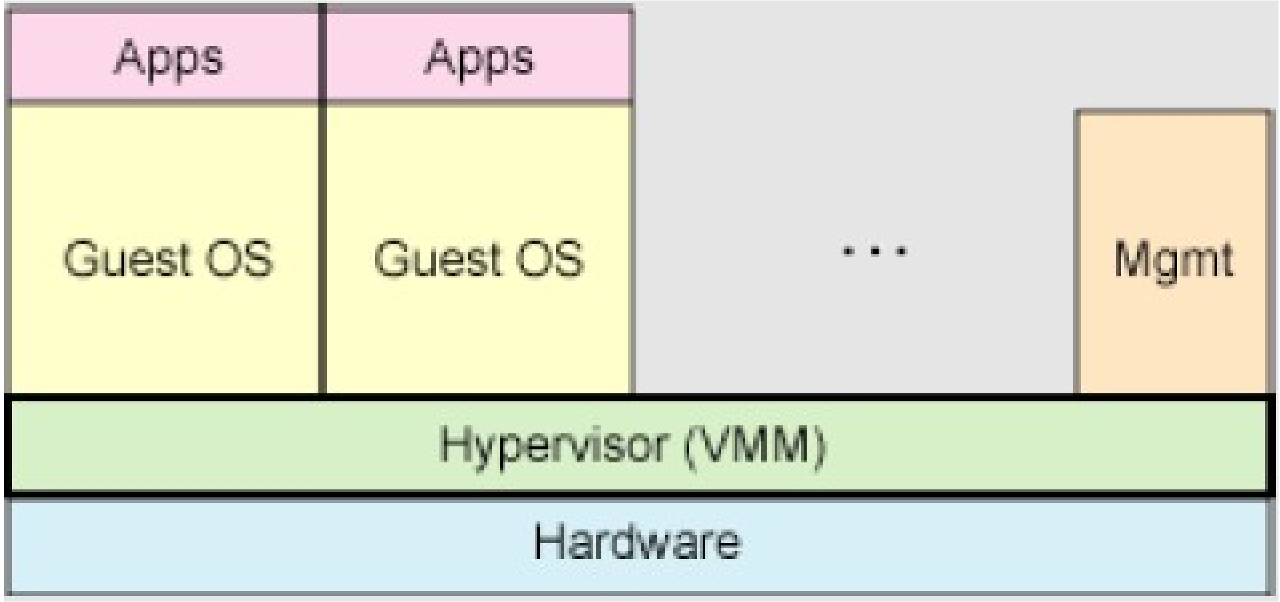

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Server virtualization: Virtualisation tech

Container-base virtualization

- e. g. VServer

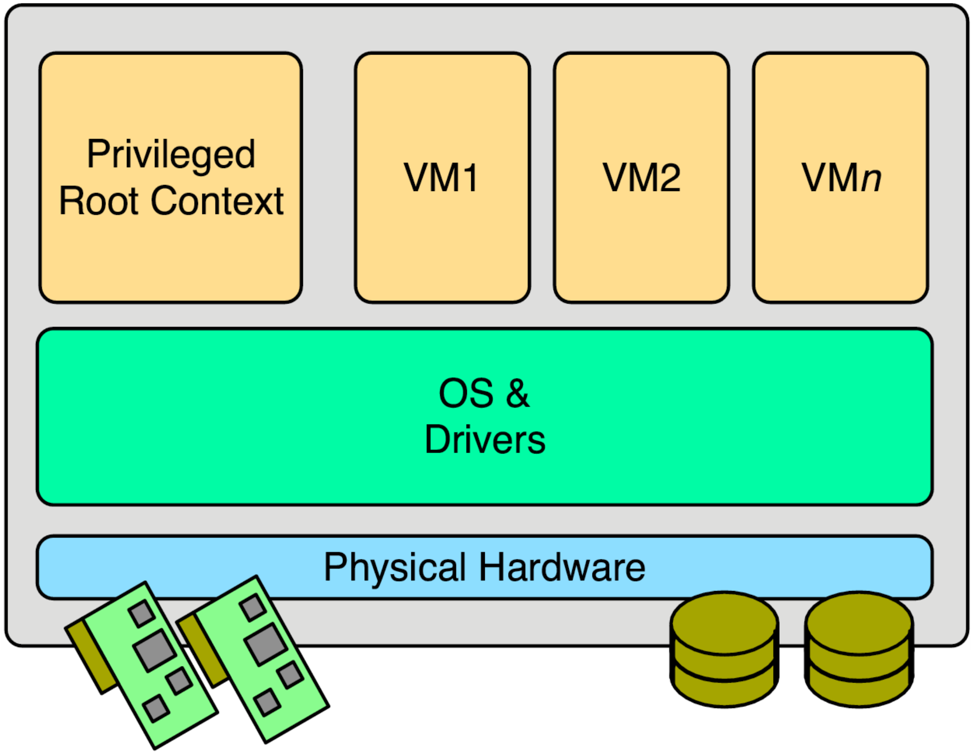

- Single Kernel
- Less isolation
- Good performance
- Live Migration
- Live System Update

[Container-based Operating System Virtualization Slides – S. Soltesz et al.]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtual core network resources

- OpenFlow switch [N. McKeown et al.]
  - Ethernet switch with flow-table
  - Goal: run experimental protocols in real networks
- Logical Routers (Cisco/Juniper)
  - Run several logical routers in parallel
  - Application Specific Routing

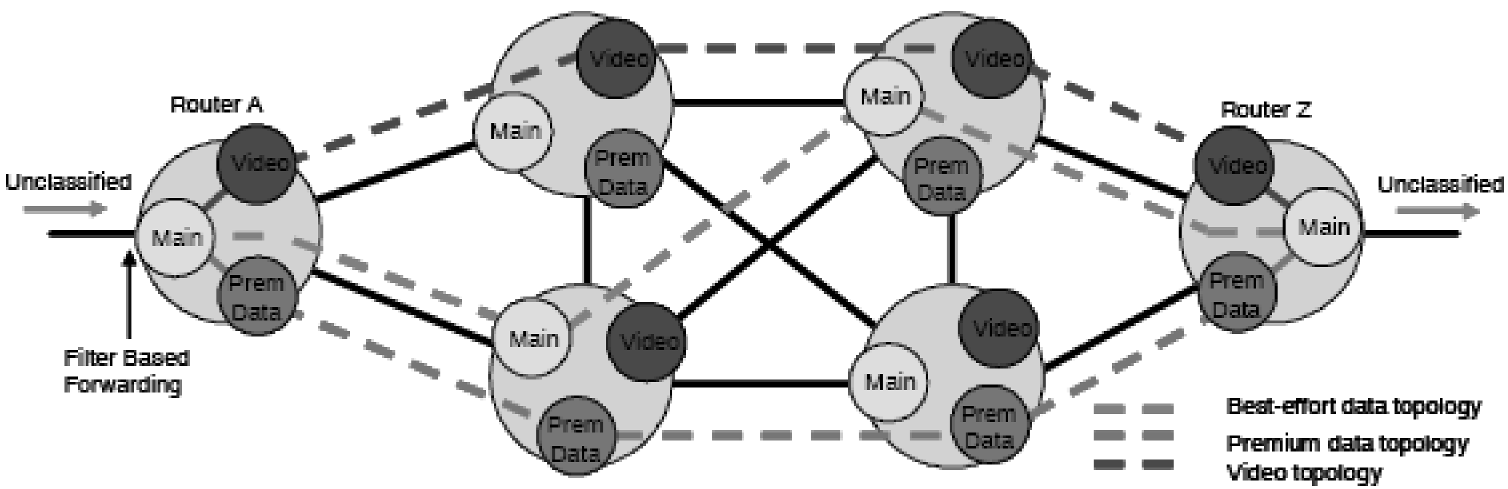

[Intelligent logical routing service - Juniper]

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Router virtualization

Advantages:

- Isolation/security
- Reconfigurability
- Mobility
- Network Customization (multi-path routing, multicast, failure recovery, TE, protocols…)(as in overlays)

Potential drawbacks

– processing overhead

– Resource sharing

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Software router

Classical software router

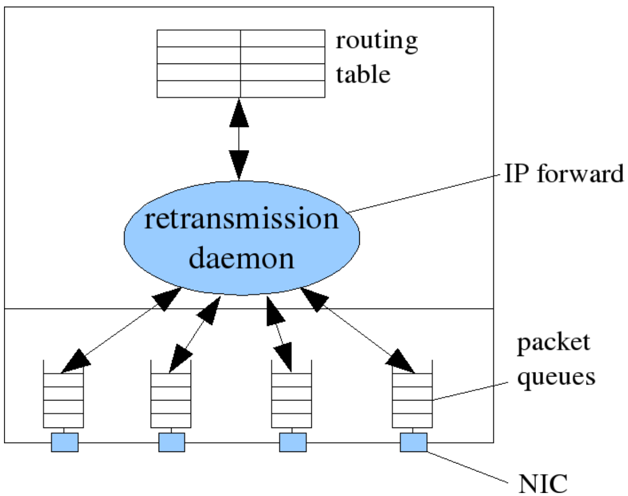

Virtual software router

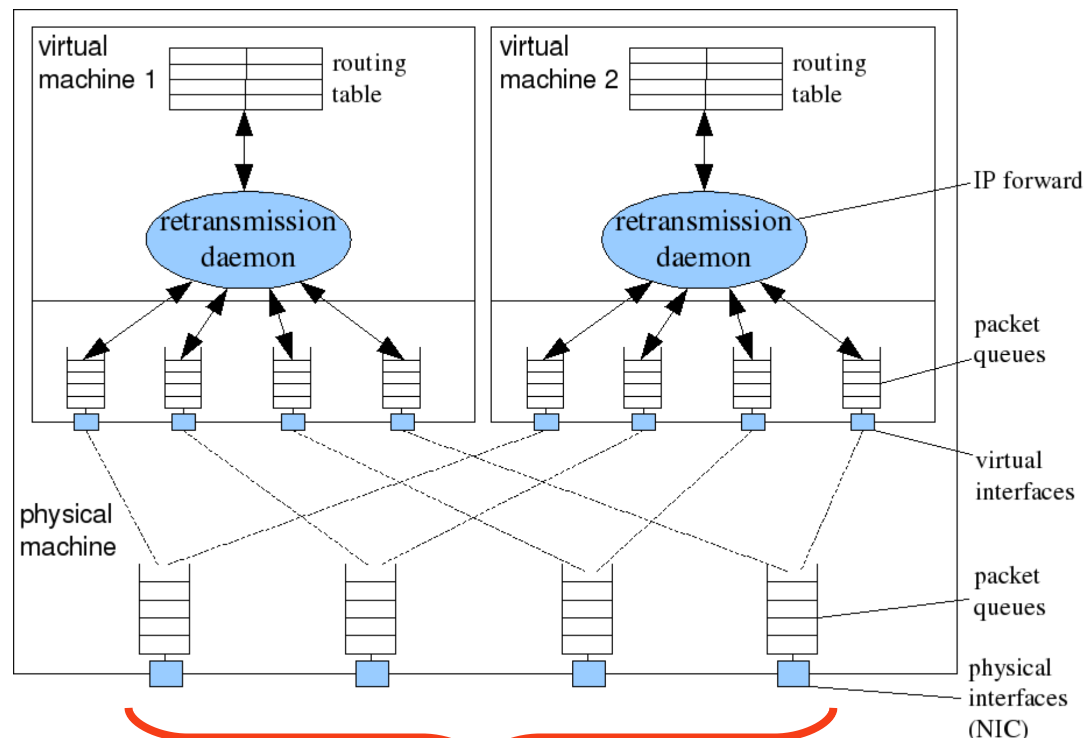

Additional processing

Resource sharing

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Network Virtualization with XEN

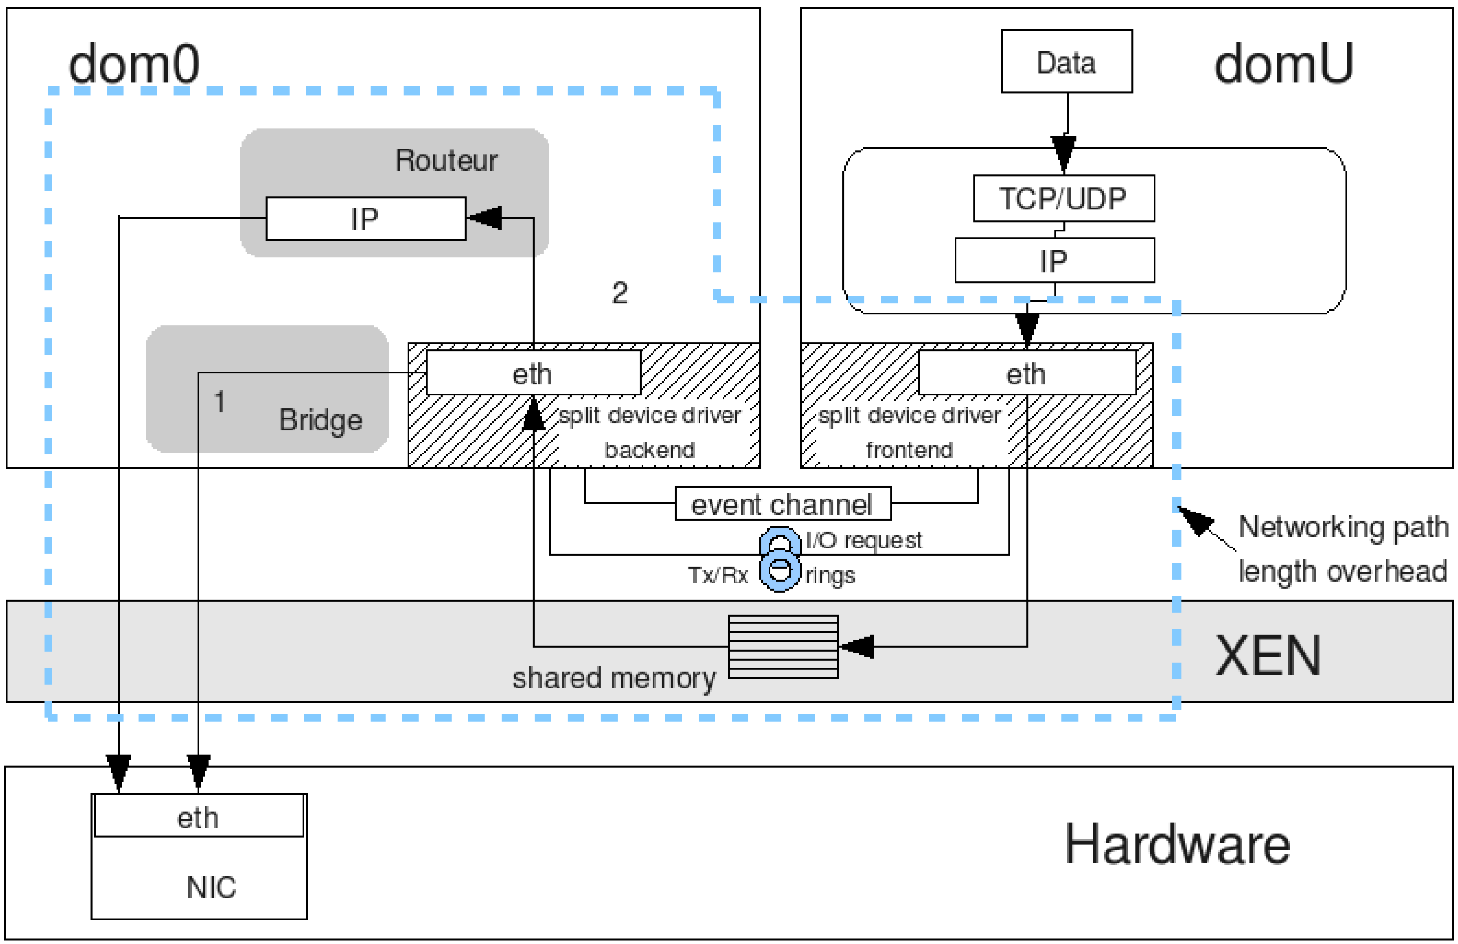

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtual software router performance

Virtualization of the

- Data plane (Packet forwarding)
- Control plane (Routing)

Performance issues in forwarding and resource sharing

- Forwarding in dom0 for Xen
- Hardware virtualization
- Adapt the physical substrate to carry virtual routers
  - e. g. Super PlanetLab Platform (SPP) providing special physical nodes with NP

With performant physical resources and adapted virtualization techniques, router virtualization becomes a promising approach

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Virtual networks on Grid'5000

HIPerNET on Grid'5000

- Resource « Virtual Router » available
- Bandwidth control of execution environments

Provide virtual network infrastructures on Grid'5000

- For isolated experience execution
- Virtual network creation on demand
  - Specify a graph of virtual resources with VXDL
  - Share the physical infrastructure with other users
- Resource control (Computing resources and bandwidth/delay)

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr

Conclusion

Resource Virtualization is a « very hot topic»

Lot of benefits envisioned

Key point: decouples the physical hardware from the service level

Very attractive for IT market as well as for Grid & network research.

« Virtual infrastructures » may be an important paradigm shift for the Internet & Grids

Issues to be further explored:

- performance, security, policies, stability, scalability…
- management: M physical resources, N virtual resources, Mapping

Pascale.Primet@inria.fr - Fabienne.Anhalt@ens-lyon.fr