



Federating

FEDERICA

Mauro Campanella - GARR

GENI - FIRE Workshop
Madrid, December 8th, 2008




Agenda



Basic federation analysis
FEDERICA and virtualization
Federating FEDERICA

Facilities Usage

The user accesses two facilities and operates them **separately**.




Facility Requirement	W/out Clearinghouse	With Clearinghouse
Common control plane	No	No
Common data plane	No	No
Standards for resources and services	No	No
Common calendar	No	No

The complexity is totally on **user's** shoulders, easy for facilities.

Facilities Usage

The user accesses two facilities and operates them **separately** (through the clearinghouse).




Facility Requirement	W/out Clearinghouse	With Clearinghouse
Common control plane	No	No
Common data plane	No	No
Standards for resources and services	No	No
Common calendar	No	No

The complexity is on **clearinghouse** shoulders, easy for facilities. The data plane from the user to the Facilities may or may not traverse the clearinghouse

Facilities Usage

The user accesses two facilities and operates them **simultaneously** as a single facility




Facility Requirement	W/out Clearinghouse	With Clearinghouse
Common control plane	No	No
Common data plane	Yes	Yes
Standards for resources and services	Some	Some
Common calendar	No	No

The complexity is on **user's** shoulders and limited on **facilities**. Facilities must have at least a common data plane to exchange data

Facilities Usage

The user accesses two facilities and operates them **simultaneously** as a single facility




Facility Requirement	W/out Clearinghouse	With Clearinghouse
Common control plane	No	No
Common data plane	Yes	Yes
Standards for resources and services	Some	Some
Common calendar	No	No

The complexity is on **Clearinghouse'** shoulders and limited on **facilities**.
Facilities must have at least a common data plane

Facilities Usage

If facilities share both a control plane and a data plane the user may connect only to a single facility and organize a joint experiment. The control protocol must be complex and has to export facility's information description (like a routing protocol). In this case the clearinghouse may not be needed.

The availability of a clearinghouse simplifies the tasks in each facility, but creates a single point of failure and a scaling issue.



Requirement	Without Clearinghouse	With Clearinghouse
Common control plane	Yes	Yes
Common data plane	Yes	Yes
Standard representation	Yes	Yes
Common calendar	No but common format	No

Federation summary

There are two desirable characteristics:

1. Lower the usage **complexity** for the user
2. Allow the use **simultaneous use** of more than one facility by the same user

For the first requirement there are two techniques:

- a **clearinghouse**, that takes care of the protocol conversions and the different data and control planes, offering to the user a single interface and a “standard” representation of “services”
- A **standard control and data plane** plus a connection between facilities

For the second requirement the need is:

- At least a **data plane** between facilities

If a common control (inter- facility) plane is added a full peer to peer model is enabled lowering the need for the clearinghouse (its availability may in any case make the overall architecture simpler)



Agenda



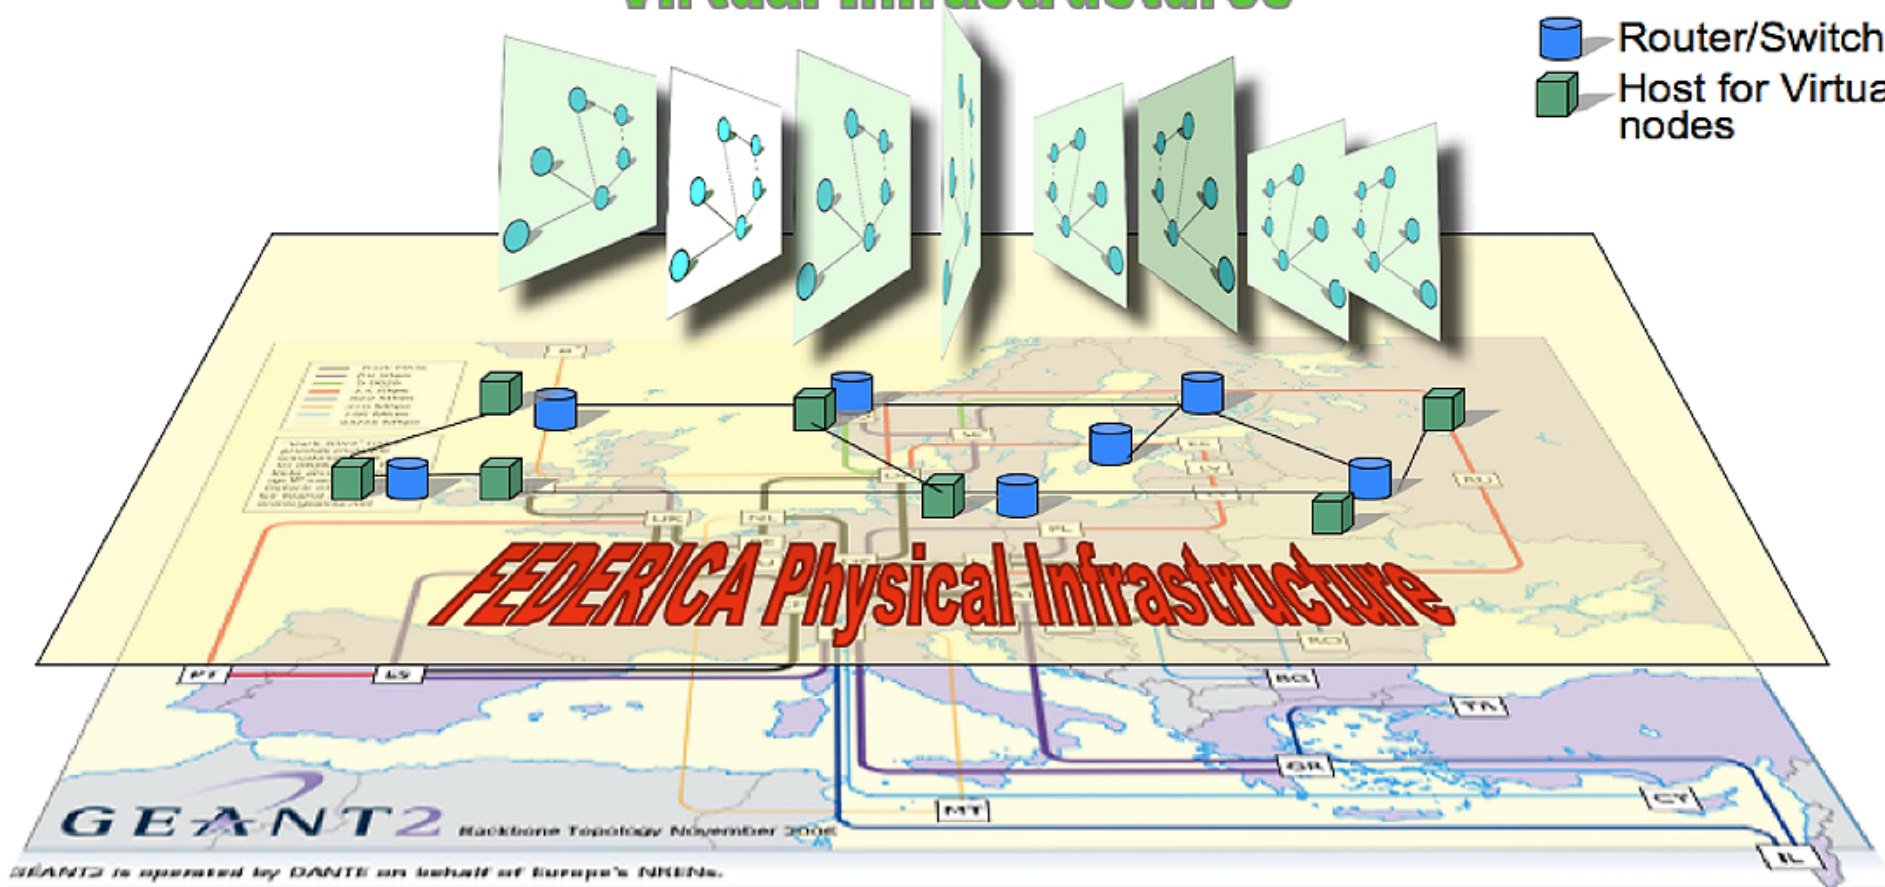
Basic federation analysis

FEDERICA and virtualization

Federating FEDERICA

Virtual Infrastructures

- Router/Switch
- Host for Virtual nodes

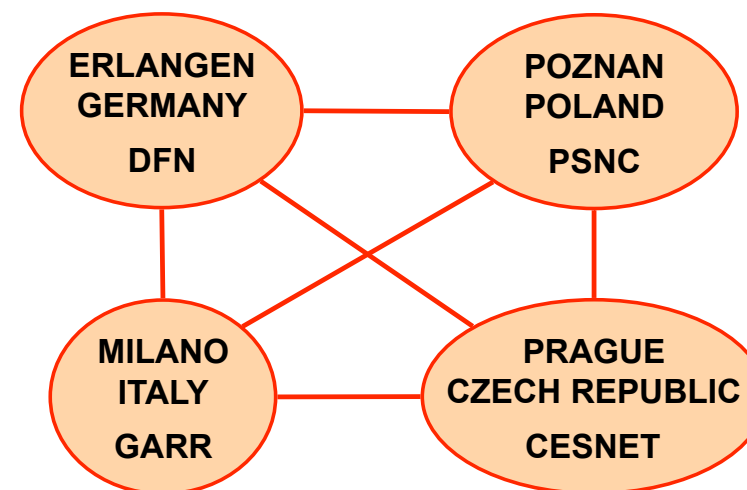


GEANT2 and NRENs Infrastructure

FEDERICA Key Principles


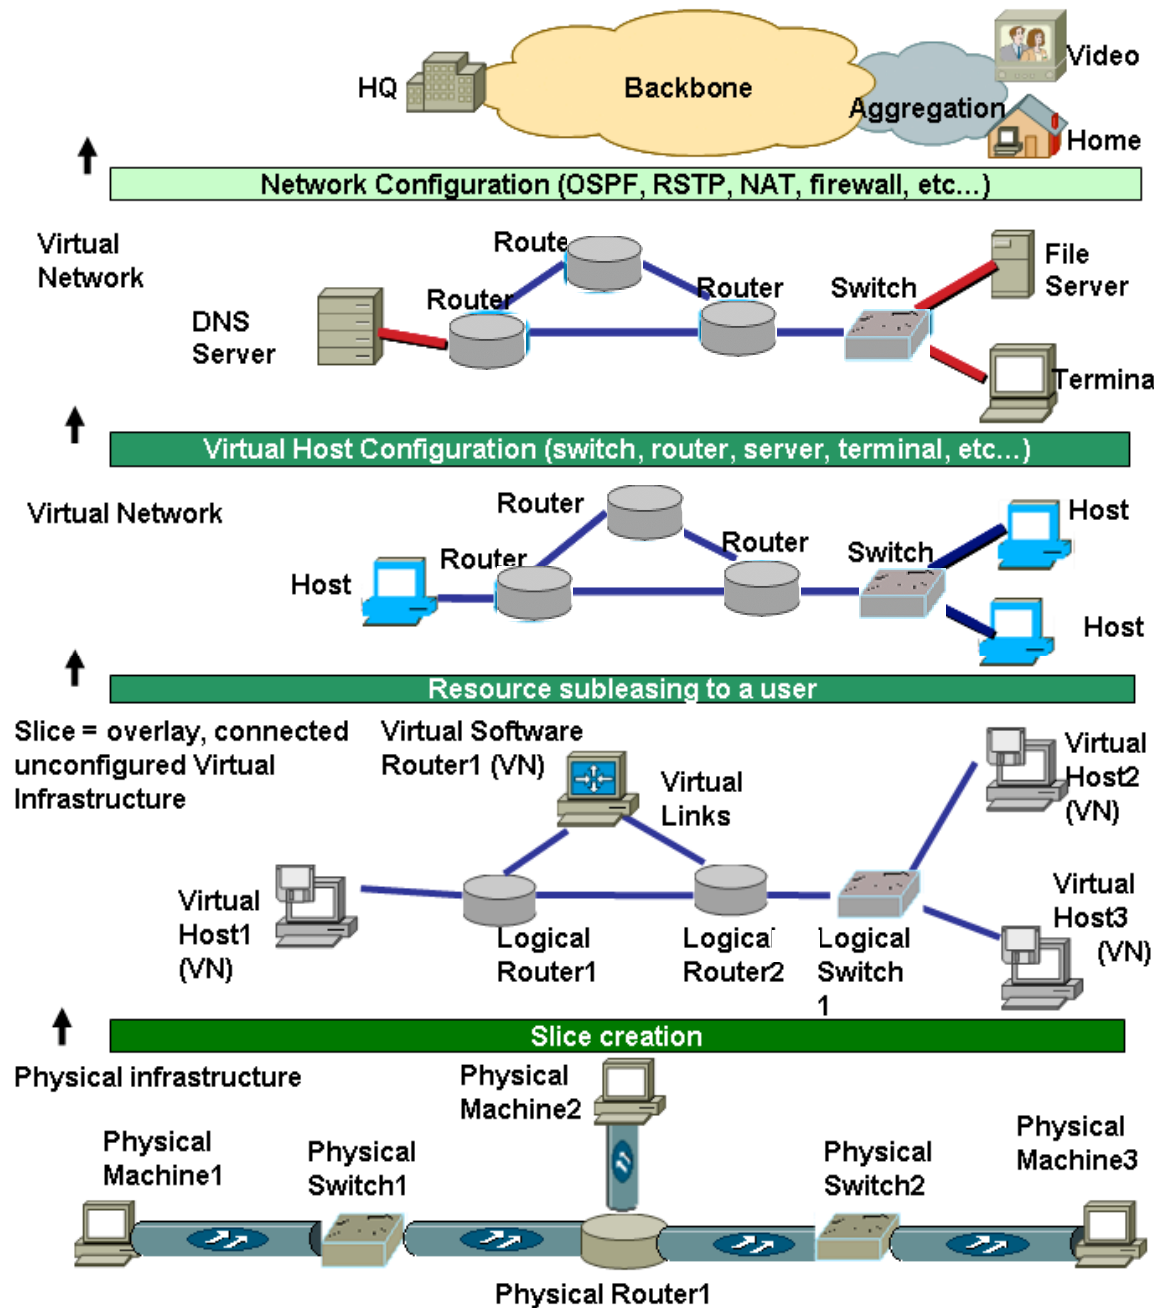
1. Create an **agnostic** and **neutral** (transparent) infrastructure using **virtualization** technologies
2. Create “**slices**” which are a set of (virtual) network and computing resources according to user’s request and are “disruptible”. Slices may communicate with General Internet.
3. Provide to the user **complete control and configuration capabilities** within a slice, down to the lowest possible network layer (in particular allow testing of any application and protocol)
4. Strive/engineer for **reproducibility** of experiments, i.e. given the same initial conditions, the results of an experiment are the same
5. Open **to interconnect / federate** with other e-Infrastructures
6. Access granted through a **User Policy Board** for a suggested **maximum time of three months**

Phase 2 Core Infrastructure



Now operational.

— 1 Gbps Ethernet




Virtual network creation steps

Pictorial of creation of a Slice

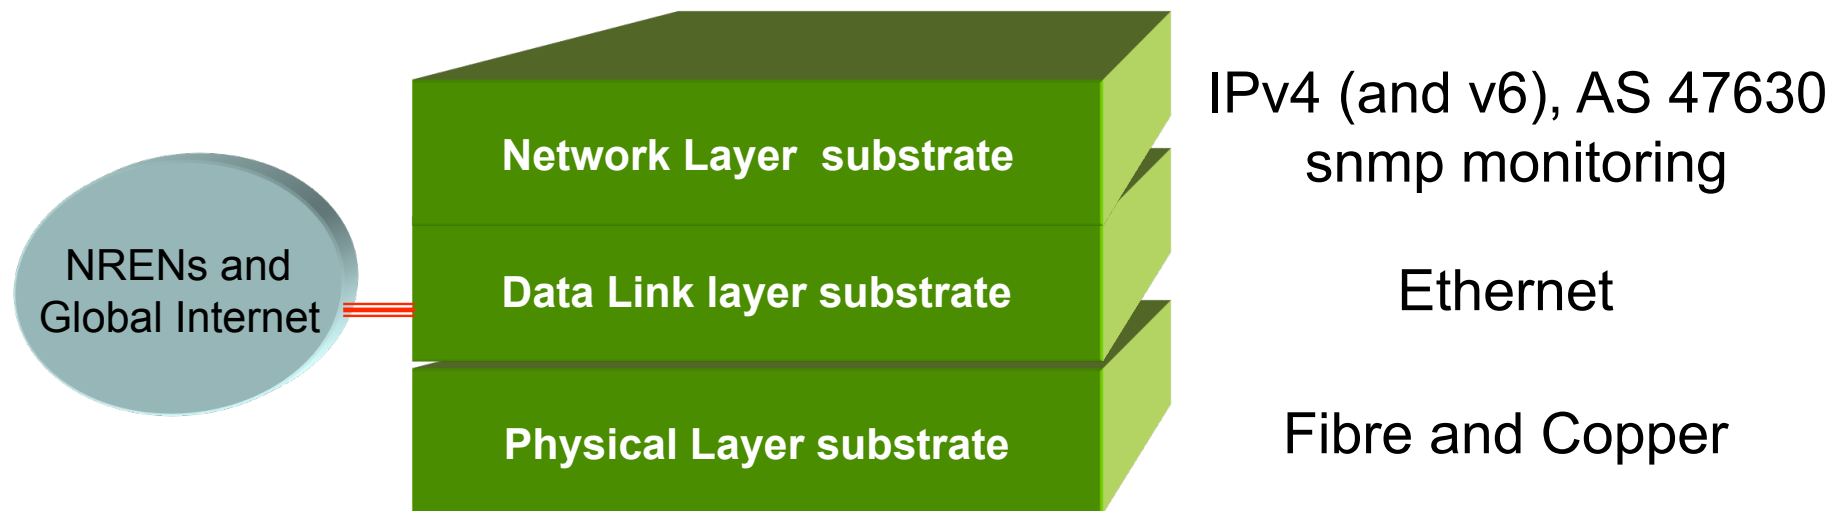
The user requests an Infrastructure made of L2 circuits, un-configured virtual nodes, to test a new BGP version.

1. Create user credentials and authentication, create entity "Slice"
2. Create Virtual Gateway (in red) to bridge the user from outside into the slice
3. Create resources and connect them as specified by the user

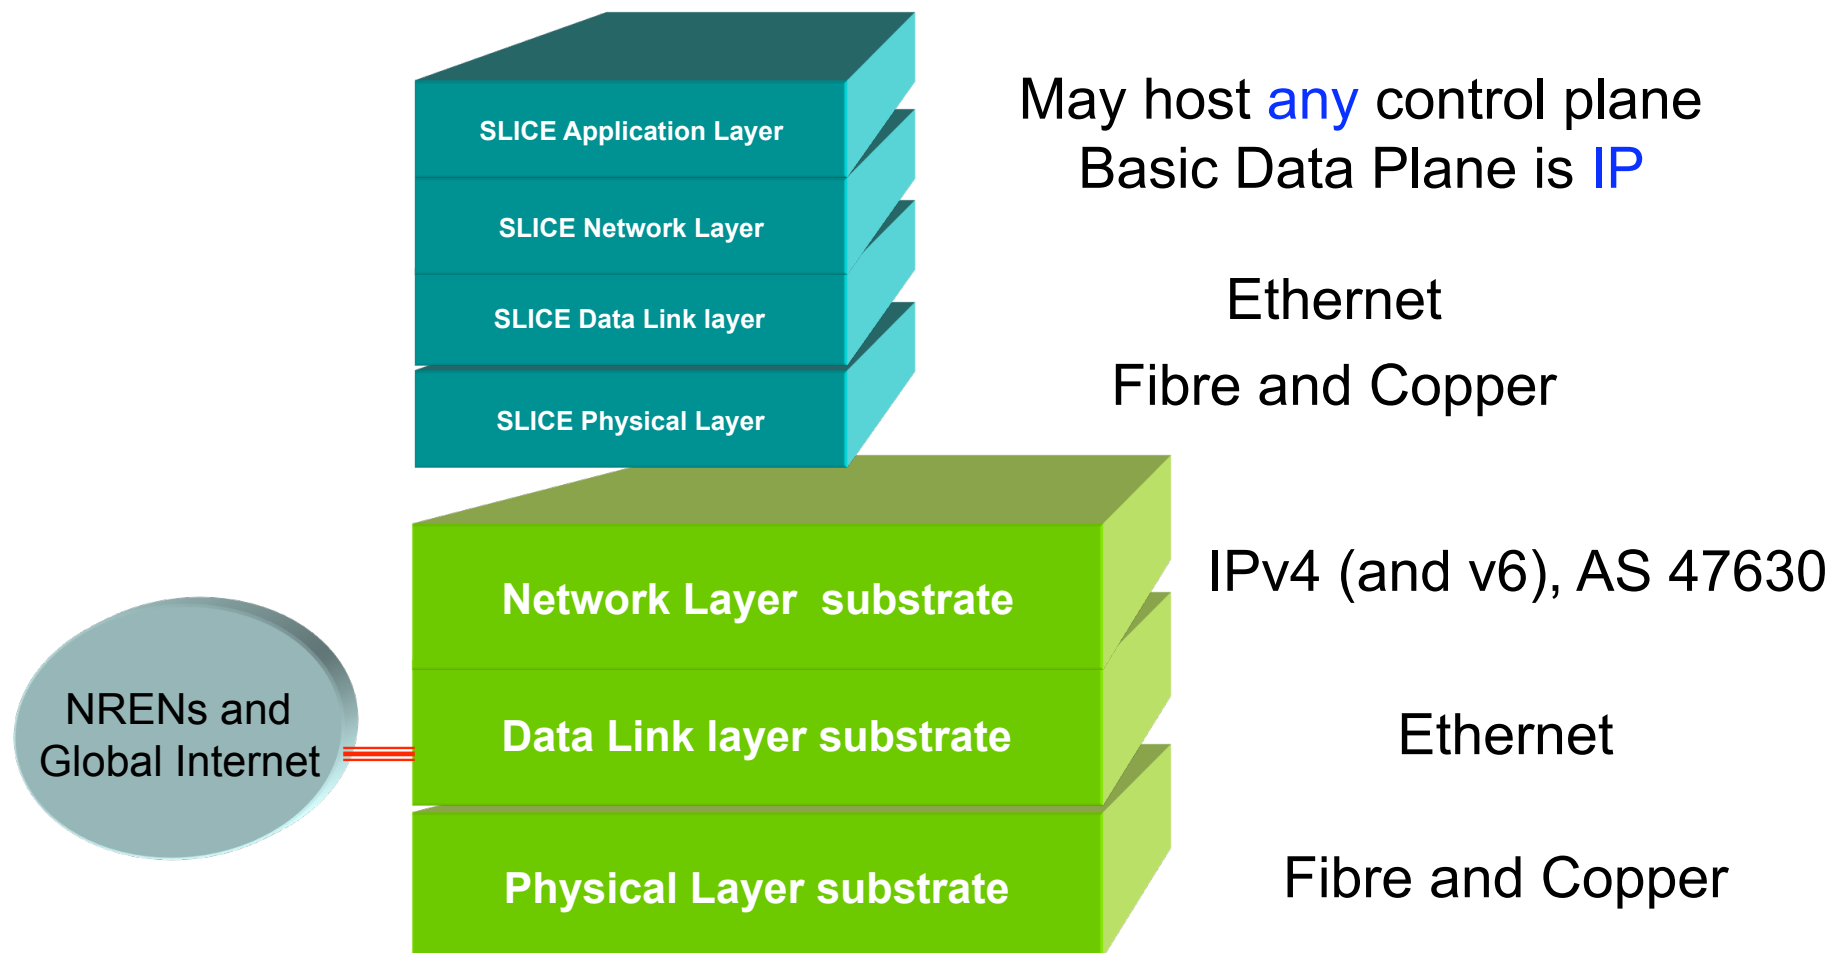


Access protocol : initially paper due to need for scheduling, security and technical agreements (no first come, first serve policy). Next step may be based on SOA (need standard representation of resources)

Control plane is not fully automated and it is a set of tools and manual configuration (due to the combined network and system resources)



Slice Data and Control Plane





Agenda



Basic federation analysis

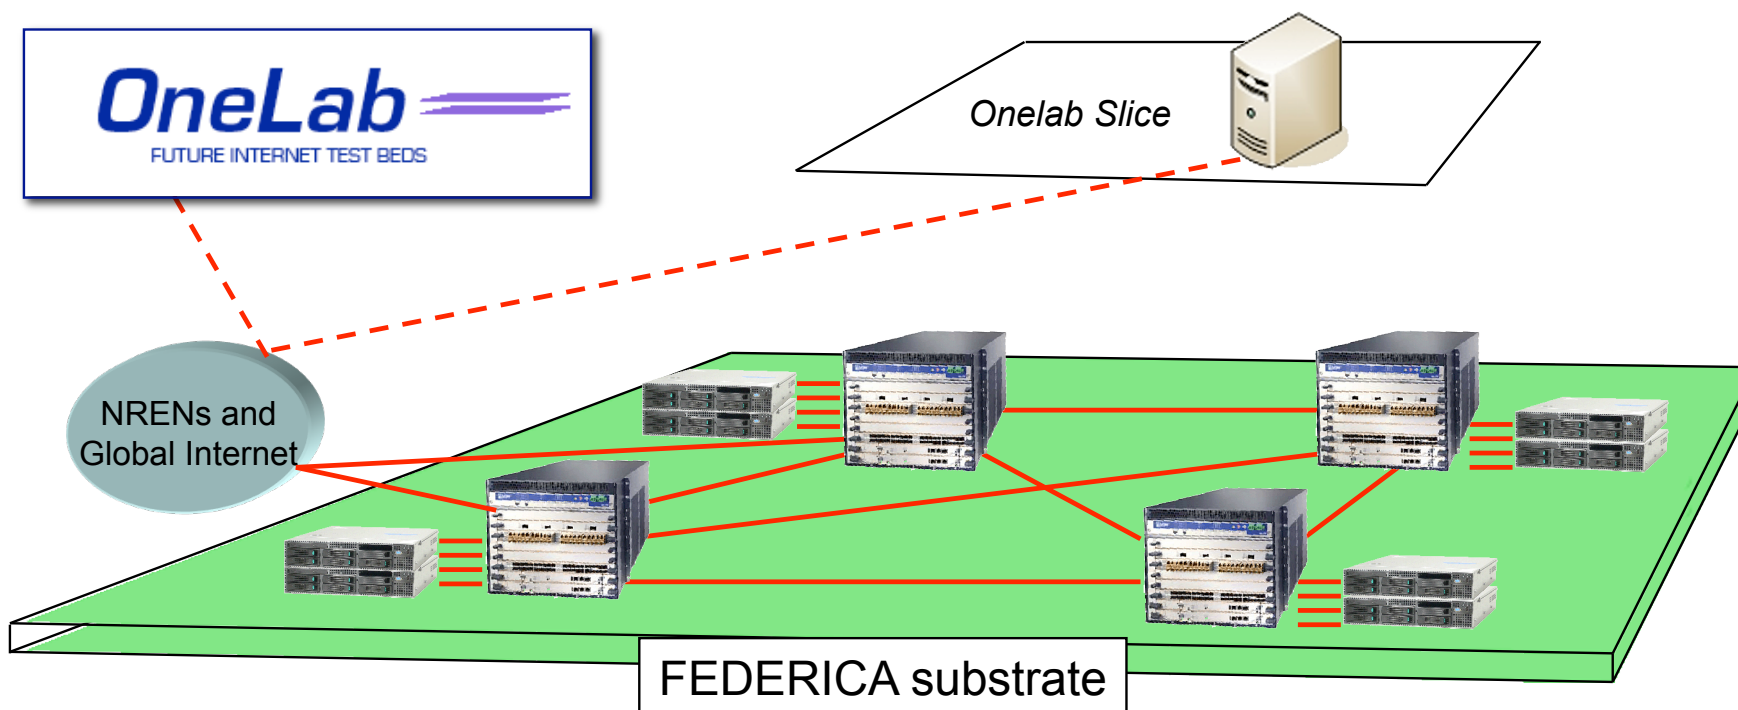
FEDERICA and virtualization

Federating FEDERICA

Federating FEDERICA

- Data plane is IP based (packet switched Ethernet)
- Physical connectivity can be accepted, currently with wired Ethernet
- Access is regulated by humans for first access, automated protocols (control) can be used later (trust and AAI needed)
- Not yet resources representation schemas available (needed to describe the available services)
- Inter-facility control plane not available (SOA?)
- Intra-facility control plane is very complex, due to scheduling and slice mapping to physical topology tasks, still manual
- A slice may host whatever internal control plane

A Panlab node can be hosted in a slice. That specific node has full control of its network interface and circuits up to the egress from FEDERICA into General Internet



- Scalability

- IPv6 unconfigured, to be enabled according to users' requests

- Ethernet framing (large MTUs) as data link

- Packet switching and statistical multiplexing assumed by default

- Less powerful switches outside the core

- Initial manual provisioning

- Not all technologies available (e.g. wireless, nomadic nodes)

- Larger virtual slices can be obtained reducing the number of concurrent users, user's equip. may be added

- Equipment is ready for IPv6

- Not considered a limiting factor, can be overcome later using WDM equipment

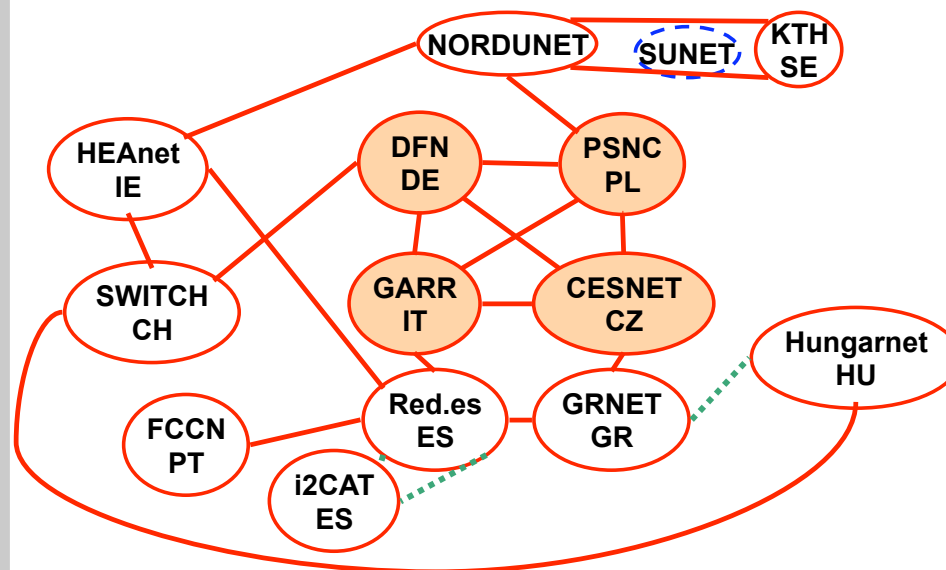
- Hardware QoS is available on two Juniper MX480 switches

- Rely more on software emulation

- Slower initial provisioning, compatible with decision process and overall management

- Equipment hosting, federation with other testbed

Infrastructure growth



Each new PoP will be equipped with a smaller switch/router (Juniper EX family) and one (or two) V-Nodes

— 1 Gbps Ethernet

*Thank you
for
your attention*