Distributed Virtual Scenarios over Multi-host Linux Environments

David Fernández, Alejandro Cordero, Jorge Somavilla, Jorge Rodríguez
Departamento de Ingeniería de Sistemas Telemáticos
Universidad Politécnica de Madrid
Madrid, Spain
david@dit.upm.es

Aitor Corchero, Luis Tarrafeta
S21Sec
Pamplona, Spain
ltarrafeta@s21sec.com

Fermín Galán
Telefónica I+D
Madrid, Spain
fermin@tid.es

dit UPM

SVM PARIS 2011

Contents

- Motivation
- Previous work:
  - Virtual Networks User Mode Linux (VNUML)
  - Distributed deployment: EDIV
  - VNUML/EDIV Limitations
- Virtual Networks over LinuX (VNX)
  - Goals and architecture
  - Implementation
  - Validation and tests
- Conclusions and future work

Context

- Systems development life cycle:

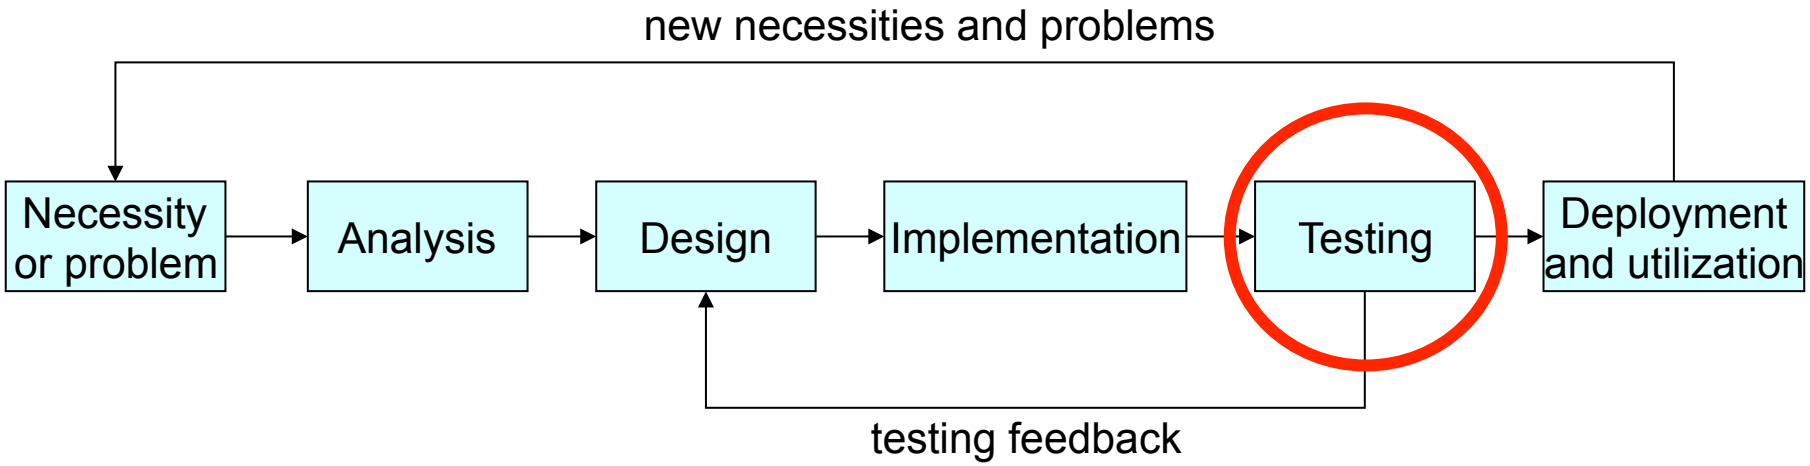

- Network and services testbeds (informal definition):
  - infrastructure platform used to experiment with networking systems and technologies under controlled conditions that often resemble those found in production networks

Testbed example: 6to4 network laboratory

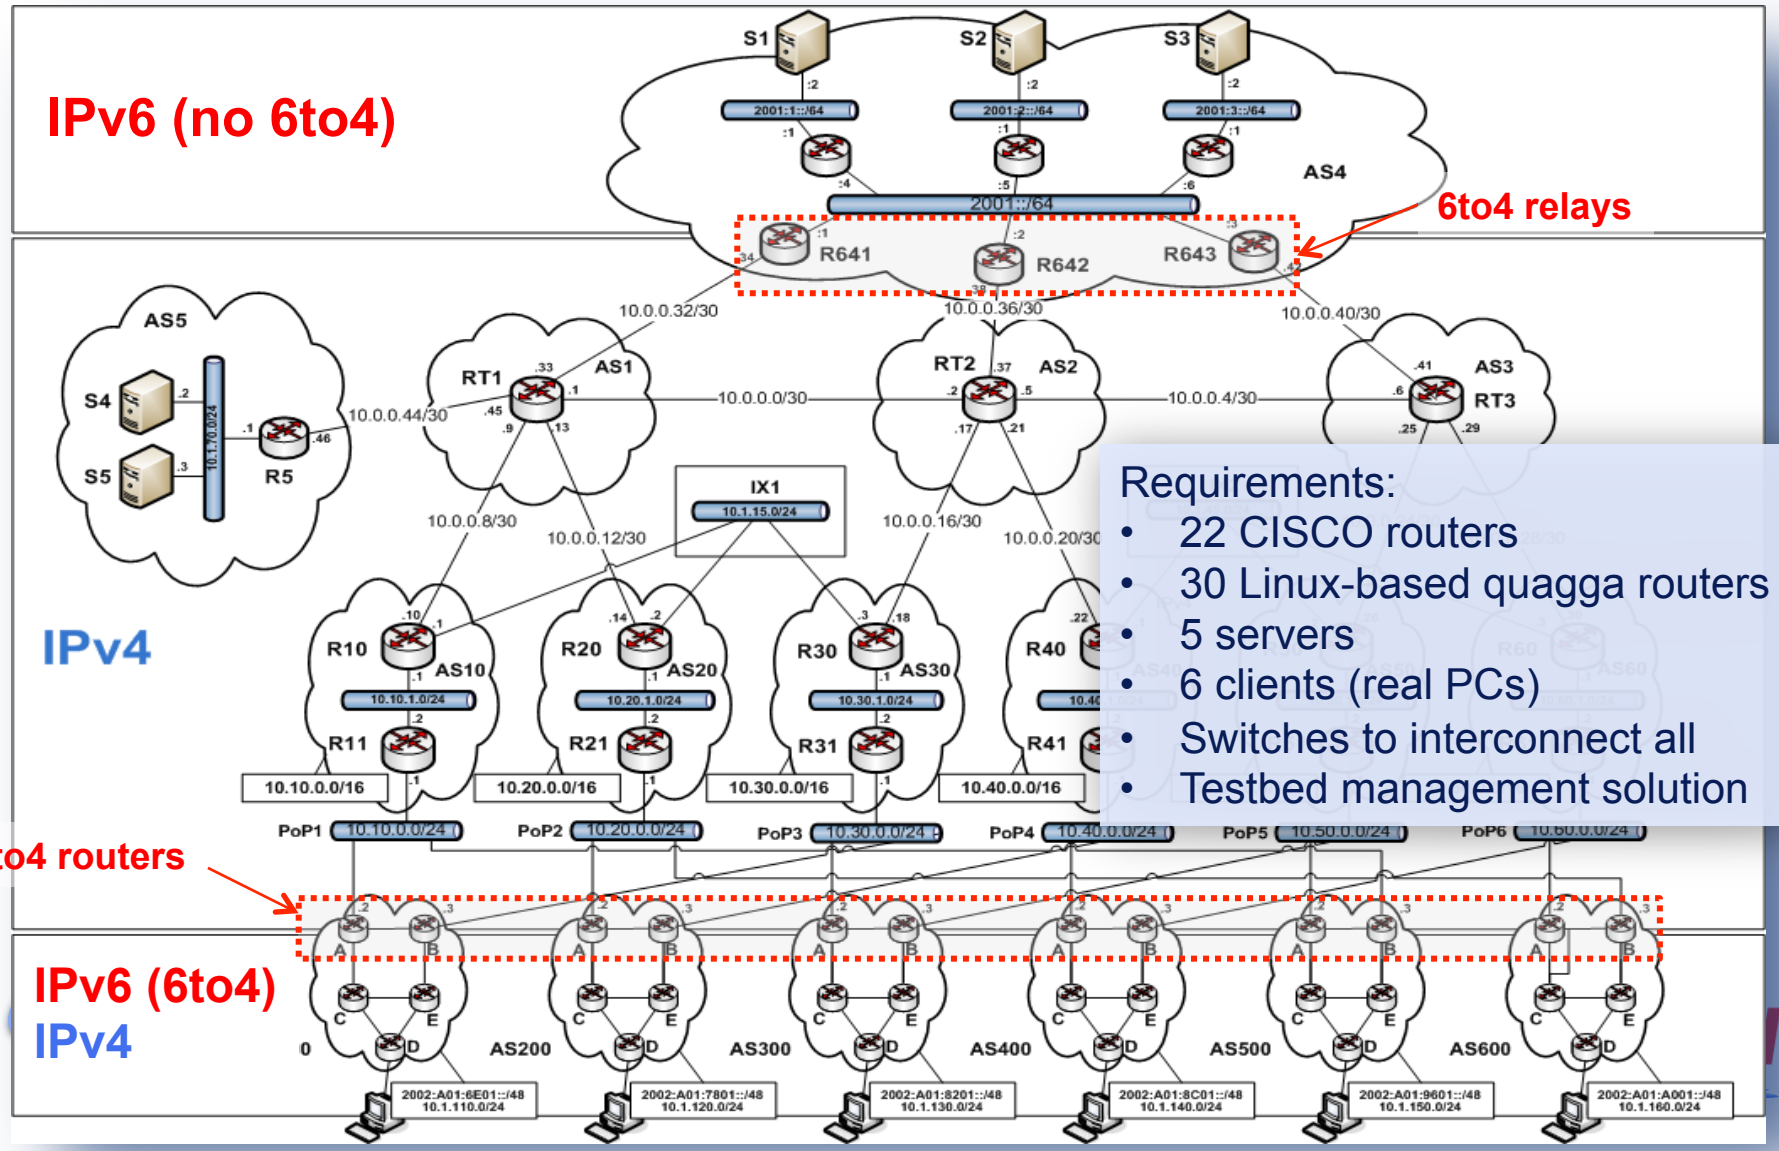

Requirements:

- 22 CISCO routers
- 30 Linux-based quagga routers
- 5 servers
- 6 clients (real PCs)
- Switches to interconnect all
- Testbed management solution

Virtualization in testbeds

- Virtualization techniques allow to execute multiple virtual machines over a physical host (Ex.: KVM, Xen, VMware, User Mode Linux, etc.)
- Combined with the use of virtual emulated networks allow the creation of **Virtual Network Scenarios (VNS)**
  - Following a user-defined topology
  - Possibly including connection to external equipment and networks

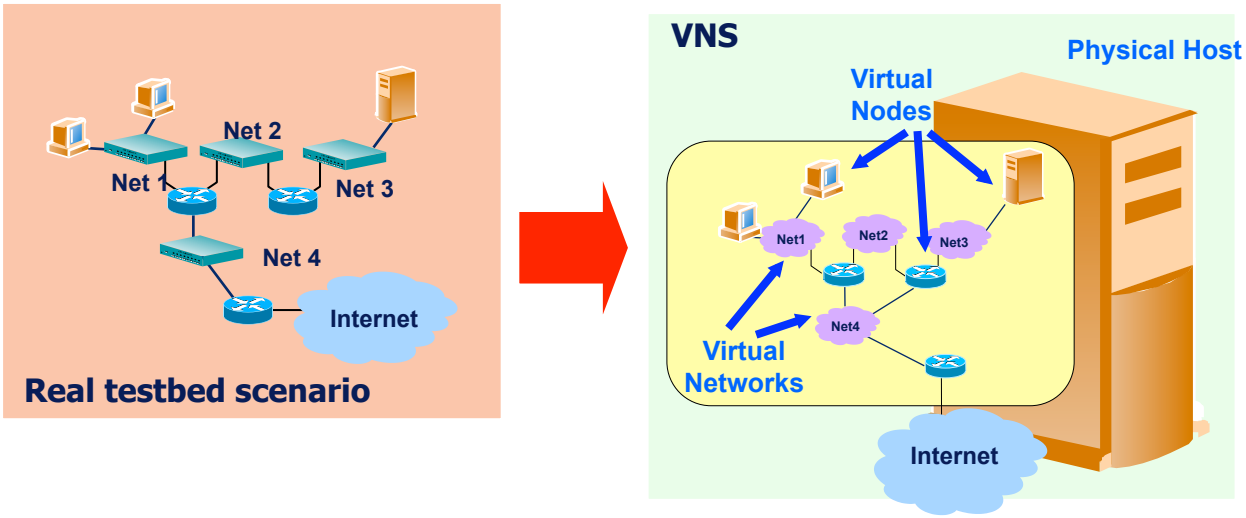

- Several tools available to manage VNS:
  - GNS3, Netkit, MNL, Marionnet, VNUML, etc.

VNUML Operation Workflow

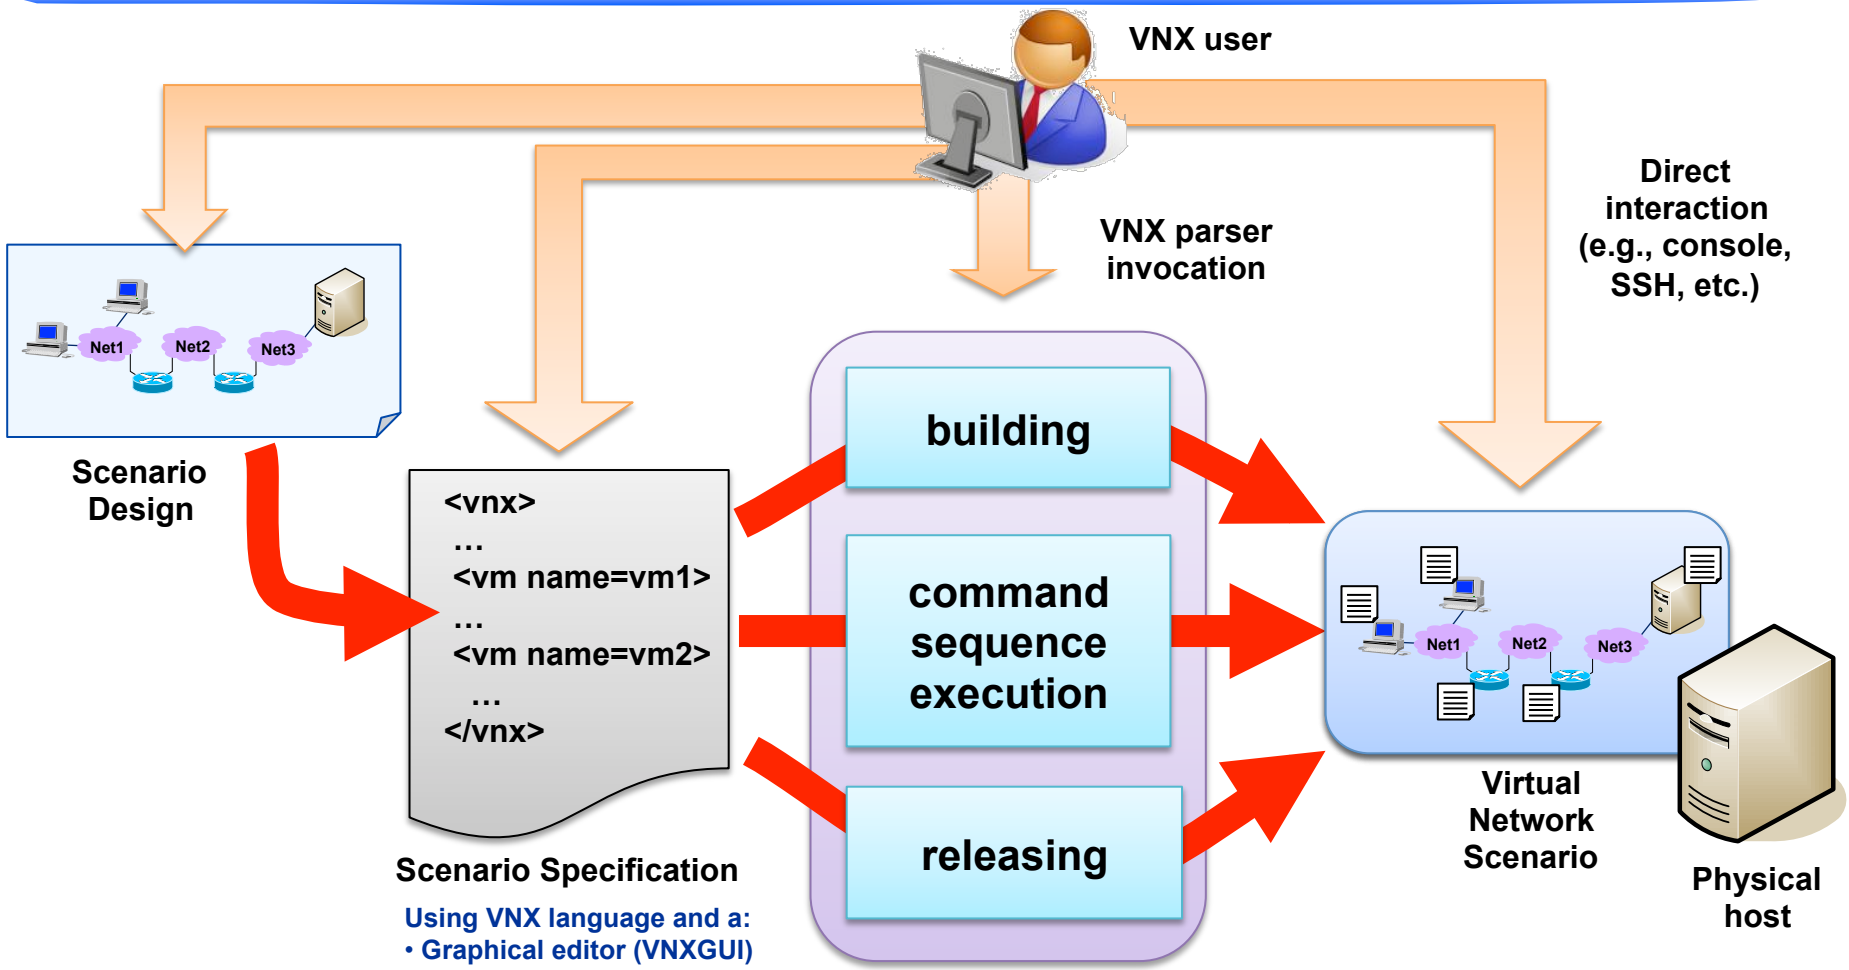

VNS Specification Language (I)

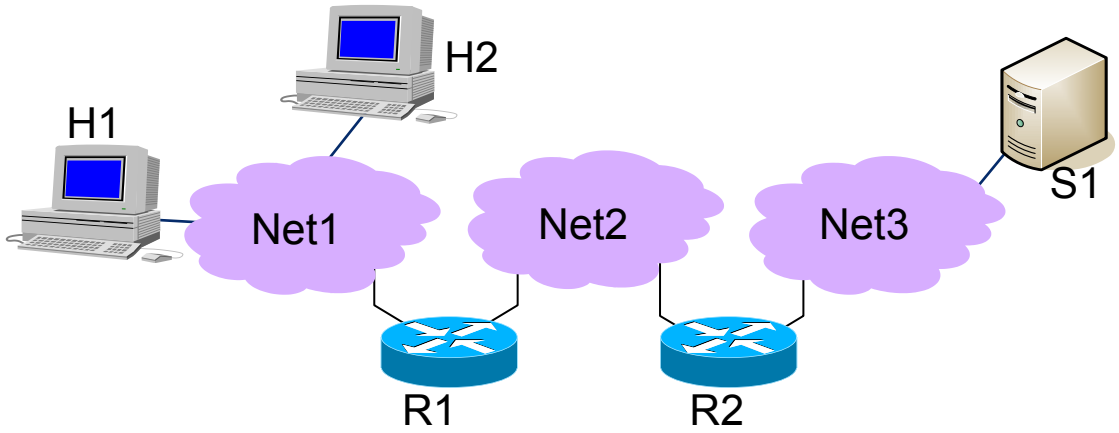


```
<?xml version="1.0" encoding="UTF-8"?>
<vnx>
    (global definitions: <global>)
    (virtual networks definitions: <net>)
    (virtual machine definitions: <vm>)
</vnx>
```

VNS Specification Language (II)

- Global definitions and networks

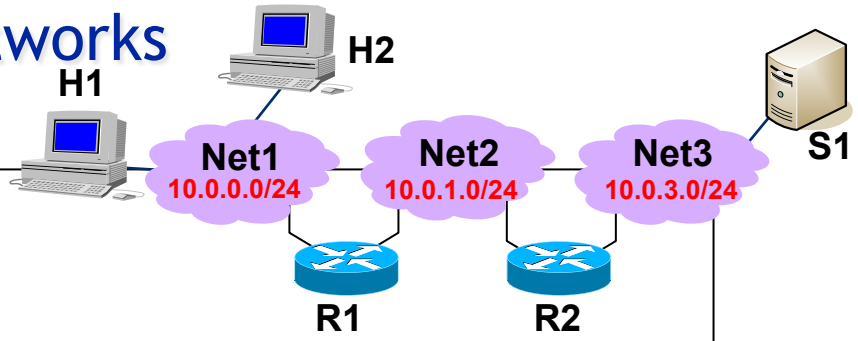


```
<vnx>
    <global>
            <simulation_name>cinco_vms</simulation_name>
            <vm_defaults>
                    …
            </vum_defaults>
    </global>

    <net name="Net1" mode="virtual_bridge" />
    <net name="Net2" mode="virtual_bridge" />
    <net name="Net3" mode="virtual_bridge" />
    …
```

VNS Specification Language (III)

- Virtual machines (end-systems)

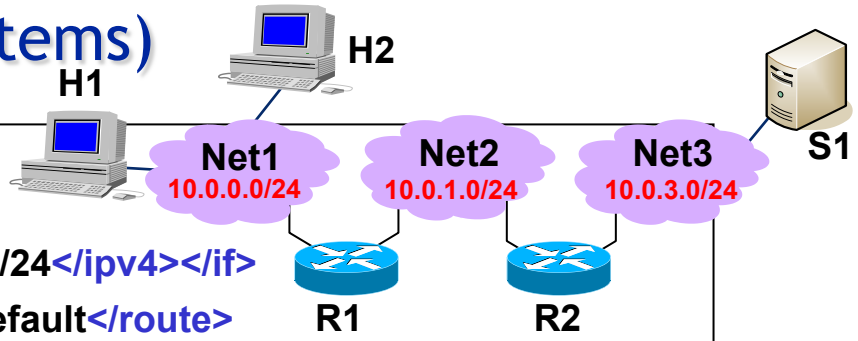


```
…
<vm name="H1">
    <if id="1" net="Net1"><ipv4>10.0.0.1/24</ipv4></if>
    <route type="ipv4" gw="10.0.0.3">default</route>
</vm>
<vm name="H2">
    <if id="1" net="Net2"><ipv4>10.0.0.2/24</ipv4></if>
    <route type="ipv4" gw="10.0.0.3">default</route>
</vm>
<vm name="S1">
    <if id="1" net="Net3"><ipv4>10.0.2.2/24</ipv4></if>
    <route type="ipv4" gw="10.0.2.1">default</route>
</vm>
…
```

dit UPM

SVM PARIS 2011

VNS Specification Language (IV)

- Virtual machines (routers)

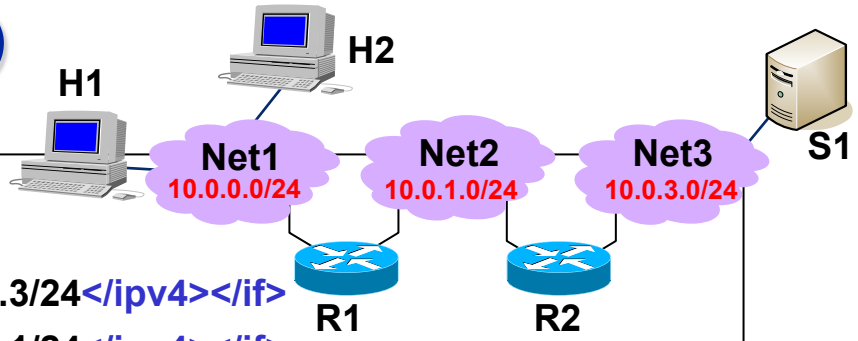


```
…
<vm name="R1">
    <if id="1" net="Net1"><ipv4>10.0.0.3/24</ipv4></if>
    <if id="2" net="Net2"><ipv4>10.0.1.1/24</ipv4></if>
    <route type="ipv4" gw="10.0.1.2">10.0.2.0/24</route>
    <forwarding type="ip" />
</vm>
<vm name="R2">
    <if id="1" net="Net2"><ipv4>10.0.1.2/24</ipv4></if>
    <if id="2" net="Net3"><ipv4>10.0.2.1/24</ipv4></if>
    <route type="ipv4" gw="10.0.1.1">default</route>
    <forwarding type="ip" />
</vm>
…
```

Scalability

- Complexity of virtual scenarios supported dependant on:
  - host resources (cpu, memory, disk, etc) available
  - VMs resources demanded
- Need to distribute virtual machines over multiple hosts to deploy bigger scenarios:
  - Clusters of virtualization servers
  - Need for mechanisms to interconnect virtual machines in different hosts
  - Requirements: transparency, efficiency, etc.

EDIV: Distributed Architecture

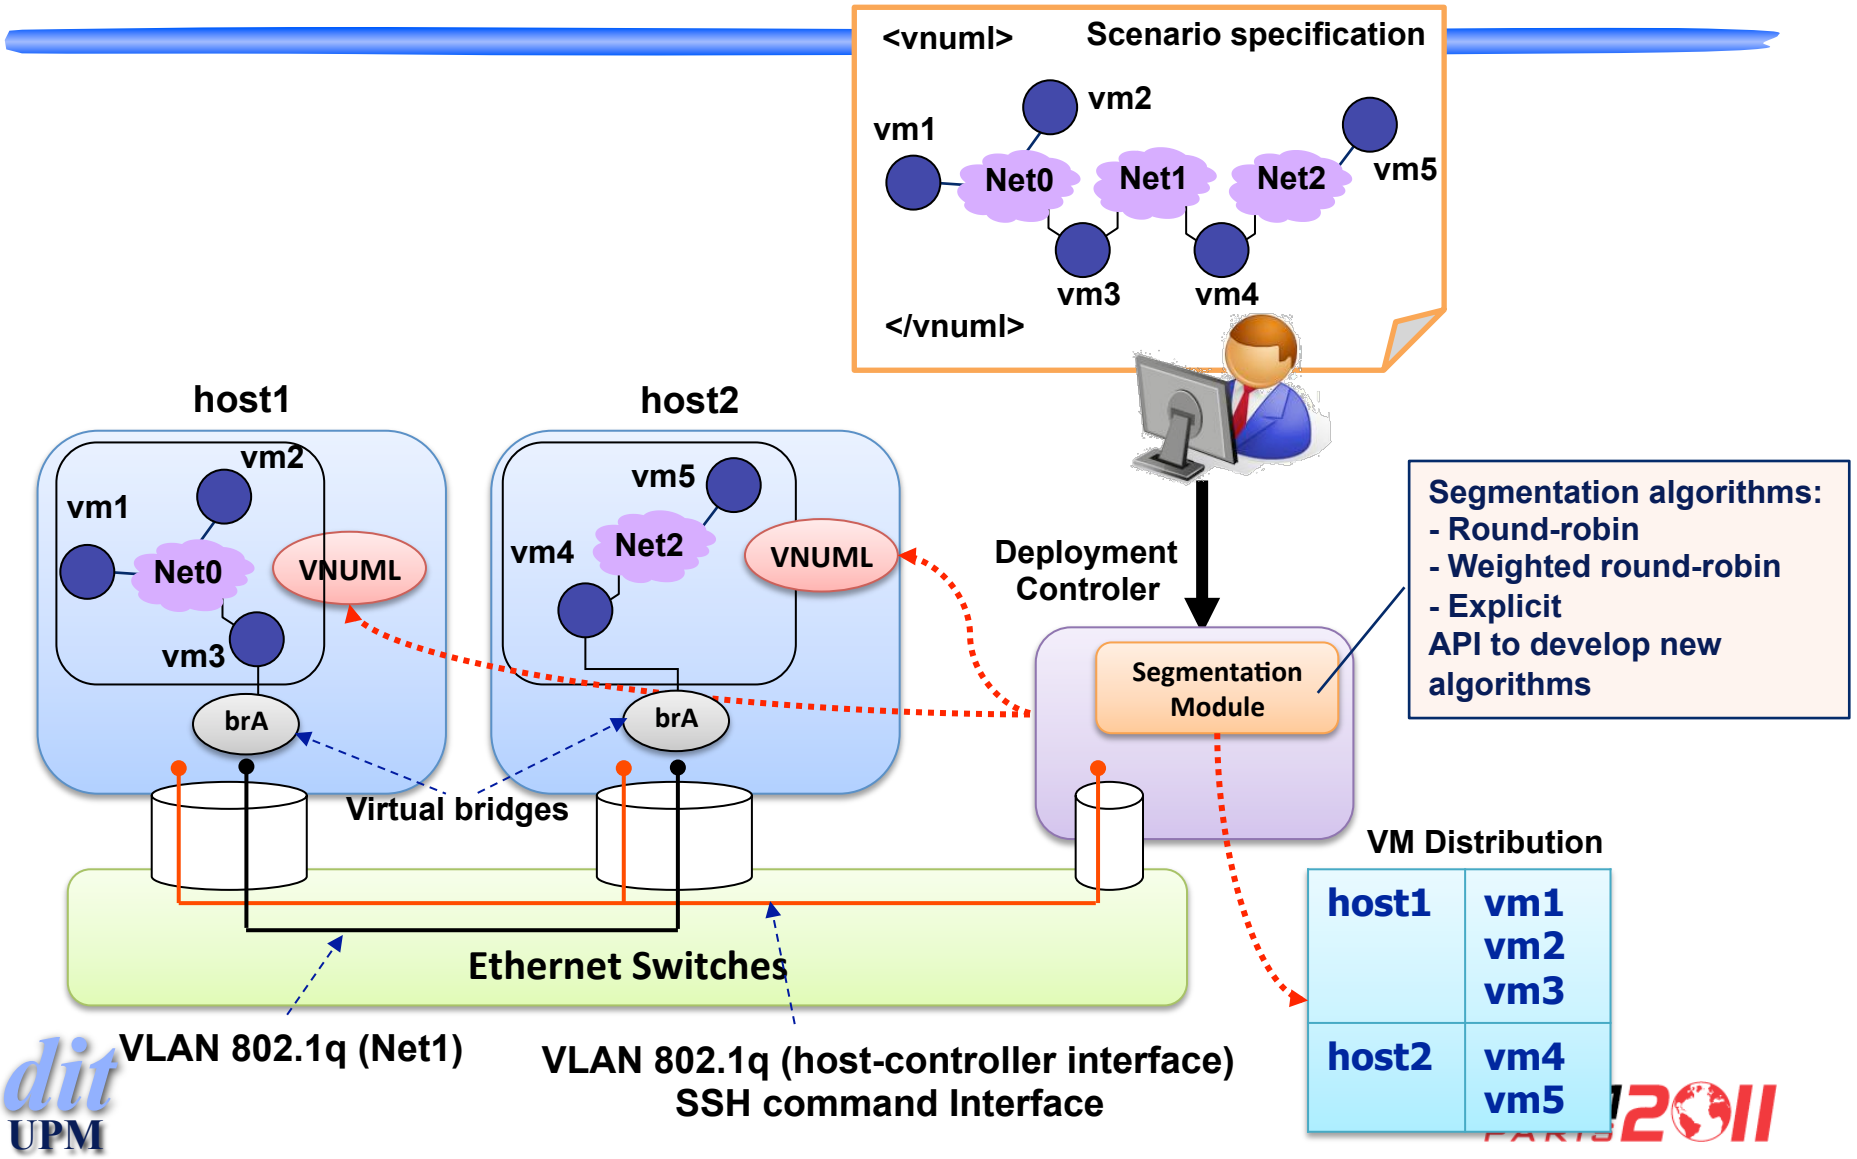

VNUML/EDIV Limitations

- Limitations of VNUML and EDIV tools:
  - Only Linux virtual machines (User Mode Linux limitation)
  - Performance problems
  - Inability to manage virtual machines individually
  - Autoconfiguration and command execution limited
  - Distributed version (EDIV) limitations: manual network configuration for disperse clusters, lack of monitoring tools, etc
- All these limitations led us redesign and rewrite VNUM create:

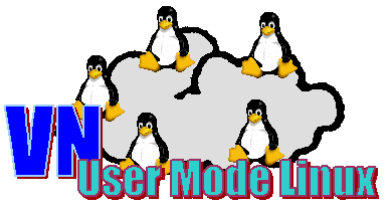

http://www.dit.upm.es/vnuml

dit UPM

VNX Objective

- Deployment of large VNS over distributed clusters made of virtualization hosts and physical equipment

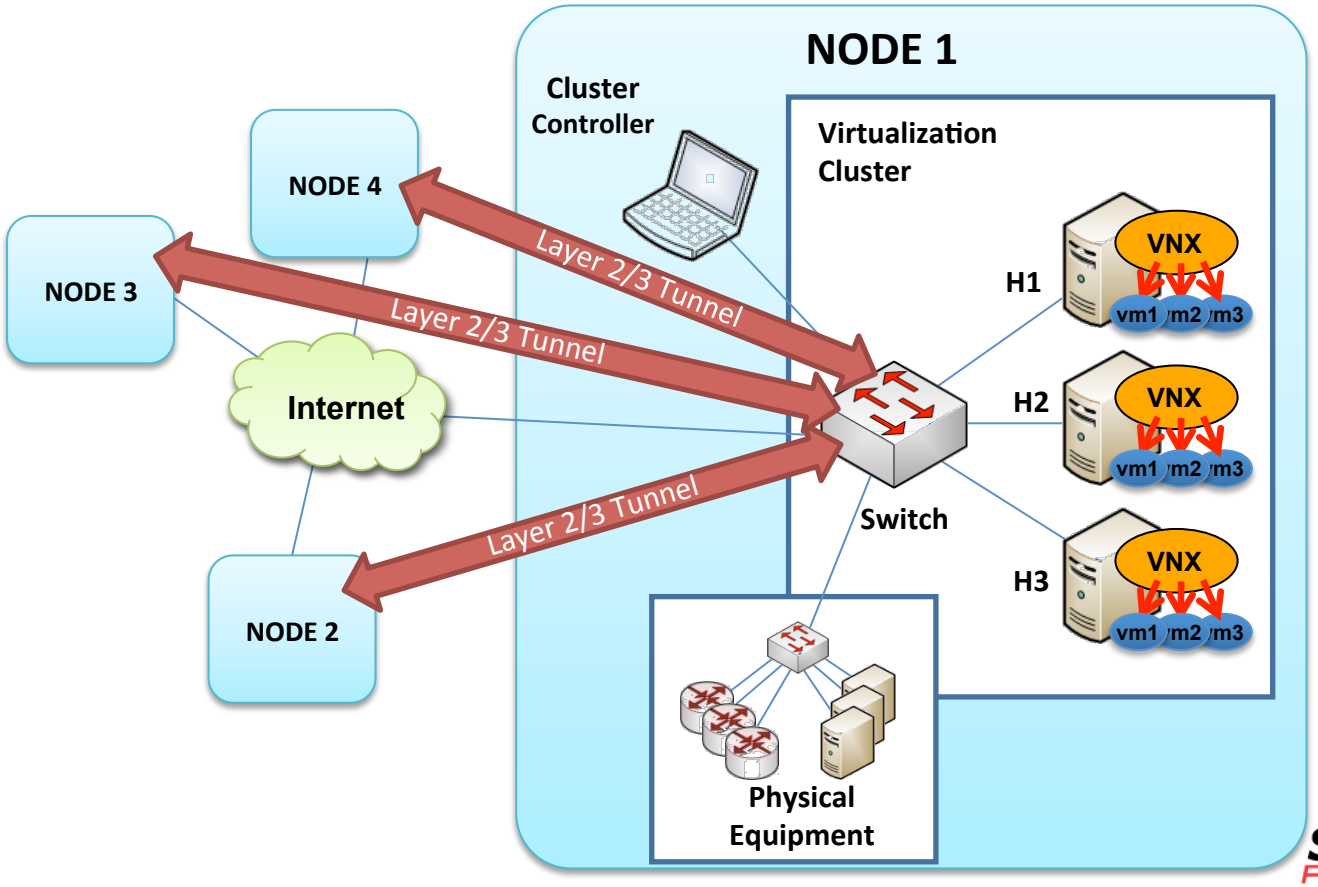

VNX Internal Architecture

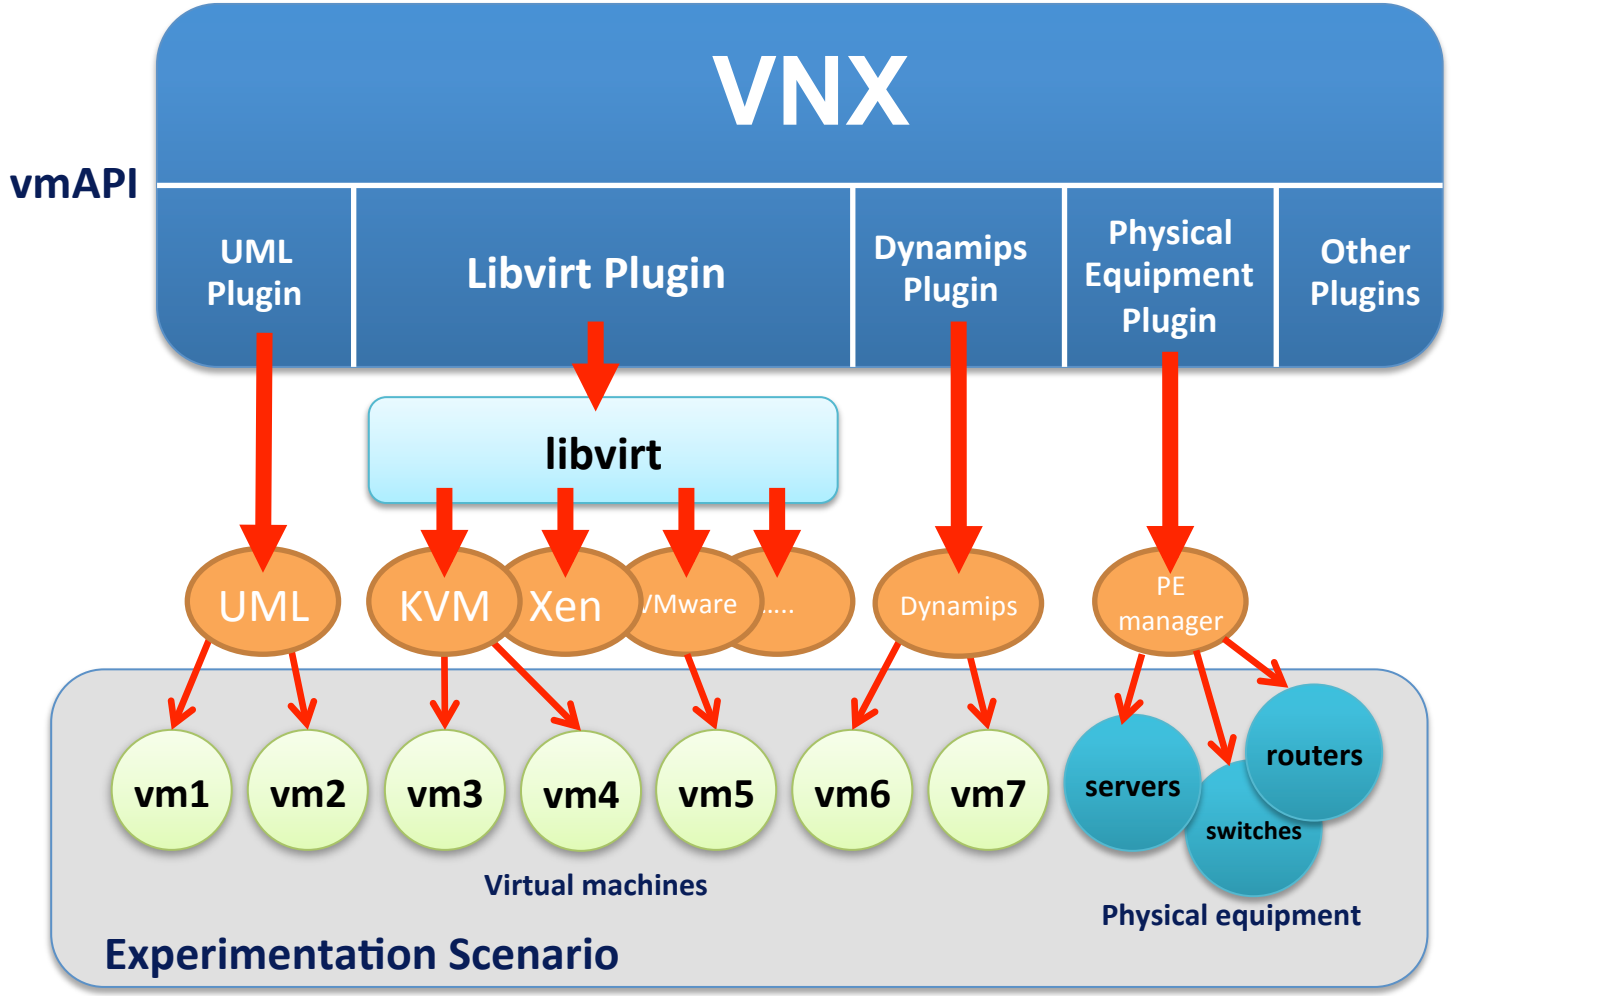

Example scenario screenshot: 6to4

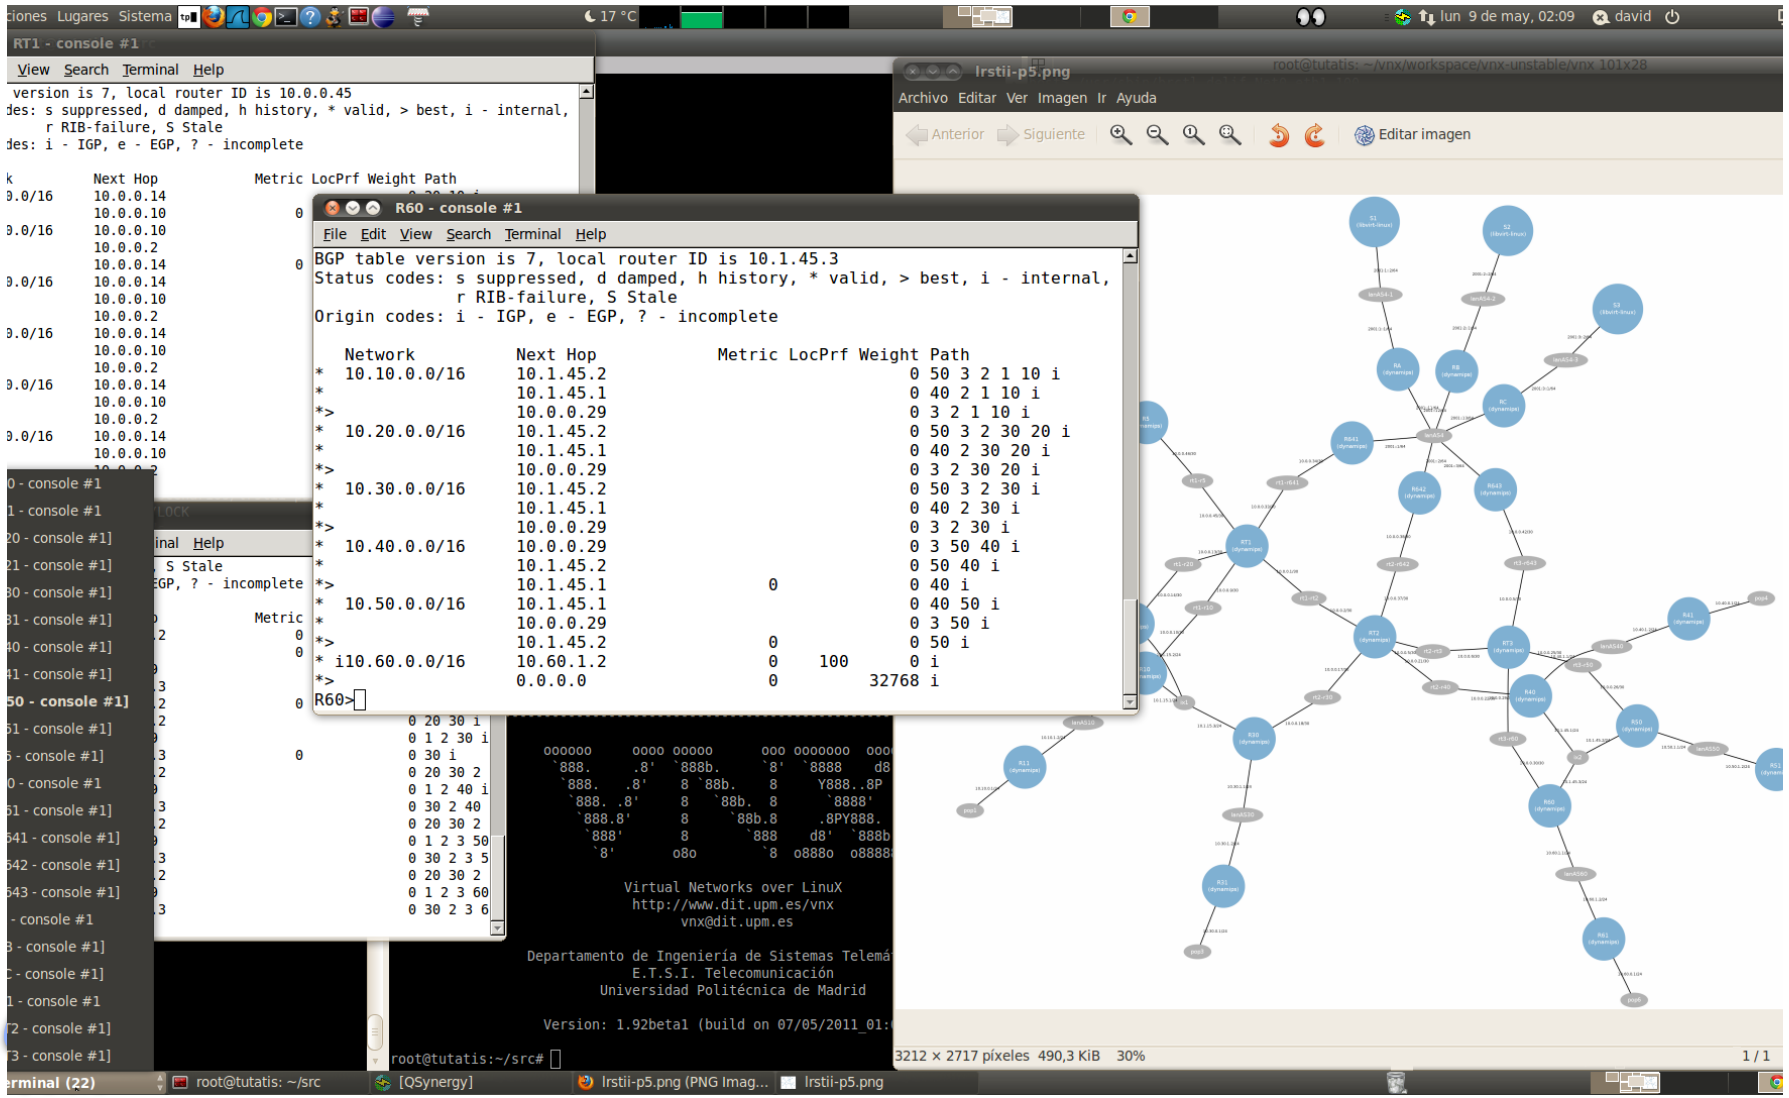

Implementation Details

- **VNX** is an open source tool (GPL) for the management of VNS based on Virtual Networks User Mode Linux (VNUML)
- First beta version available at http://www.dit.upm.es/vnx with:
  - **libvirt** support, tested with Linux (Ubuntu, Fedora, CentOS), FreeBSD and Windows (XP and 7).
  - **Dynamips and Olive** router emulation support
  - Individual management of virtual machines
  - General OVF-Environment-like autoconfiguration and command execution mechanism for Windows, Linux and FreeBSD
  - Plug-in architecture to allow extensions to VNX
  - Improved distributed deployment support (EDIV)
  - Library of root filesystems available
- VNX written in Perl (around 25000 lines of code); Windows autoconf daemon in C++.
  - ~40% of VNUML code reused with minor modifications

dit UPM

SVM PARIS 2011

Autoconfiguration and Command Execution

- Based on OVF Environment approach:
  - A dynamically created CDROM is offered to virtual machines with:
    - Initial configuration values
    - Commands to execute and files to copy
- Virtual machines run an Autoconfiguration and Command Execution Daemon (ACED) that:
  - Waits for CDROMs and,
  - Read XML files and process them
- ACED include auto-update functionality

```
<?xml version="1.0" encoding="UTF-8"?>
<create_conf>
  <vm name=”vm4” >
    <filesystem type="cow”>rootfs_ubuntu</filesystem>
    <mem>128M</mem>
    <if id="1" net="Net1" mac=",fe:fd:00:00:04:01">
      <ipv4 mask="255.255.255.0">10.0.1.2</ipv4>
    </if>
    <if id="2" net="Net2" mac=",fe:fd:00:00:04:02">
      <ipv4 mask="255.255.255.0">10.0.2.1</ipv4>
    </if>
    <route type="ipv4" gw="10.0.1.1">default</route>
    <forwarding type="ip"/>
  </vm>
</create_conf>
```

```
<command>
  <id>ubuntu-fYH3pA</id>
  <filetree seq="start-www" root="/var/www/"
       user="www-data" group="www-data"
       perms="644">conf/txtfile</filetree>
  <exec seq="start-www" type="verbatim"
       ostype="system">service apache2 start</exec>
</command>
```

Autoconfiguration and Command Execution (II)

- Important effort dedicated to virtual machine autoconfiguration and command execution
- OVF-Environment like CDROM based autoconfiguration is a general approach, however:
  - It is slow
  - Interferes with CDROM OS autoexecute mechanisms
  - Differences among operating systems and releases make the development of ACED costly
    - Tune and test for every OS and release
  - Unidirectional: no feedback from the virtual machine
- Alternative mechanisms implemented based on a shared filesystems and a serial line:
  - Files copied to shared filesystem
  - Simple signalling protocol over the serial line
  - Implemented for Olive routers; being extended to other VMs

dit UPM

SVM PARIS 2011

EDIV Validation Scenario: MPLS VPN

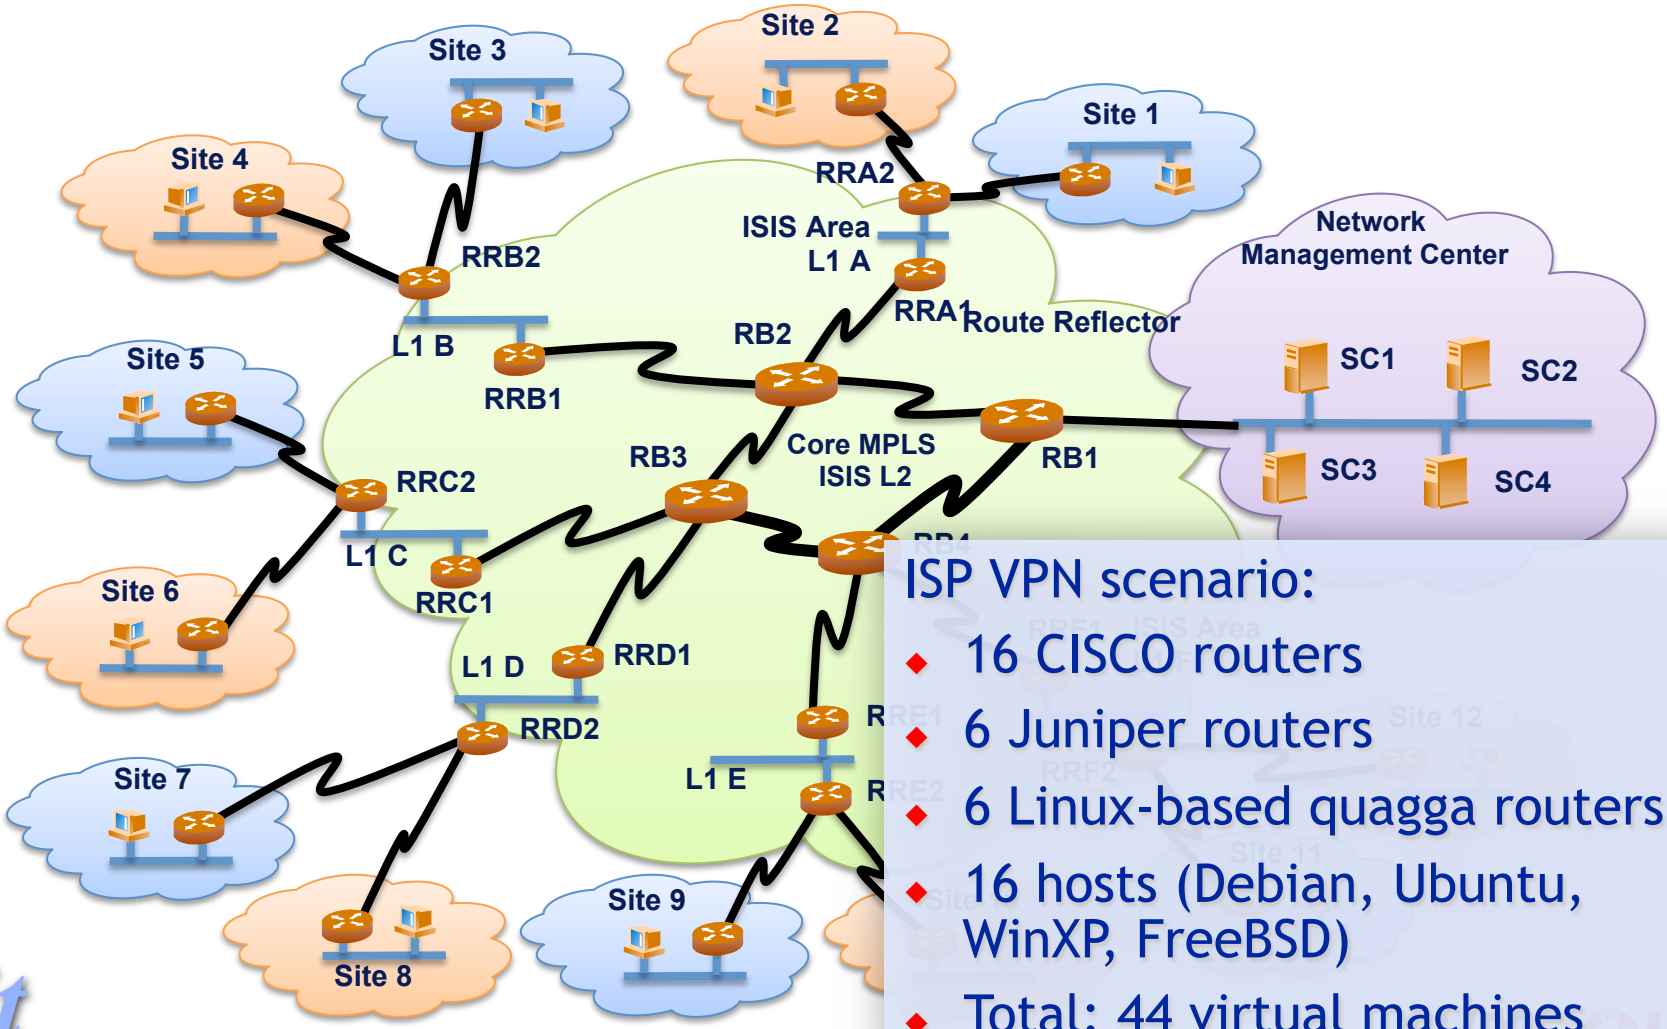

ISP VPN scenario:

- 16 CISCO routers
- 6 Juniper routers
- 6 Linux-based quagga routers
- 16 hosts (Debian, Ubuntu, WinXP, FreeBSD)
- Total: 44 virtual machines

dit UPM

EDIV Validation Scenario: MPLS VPN

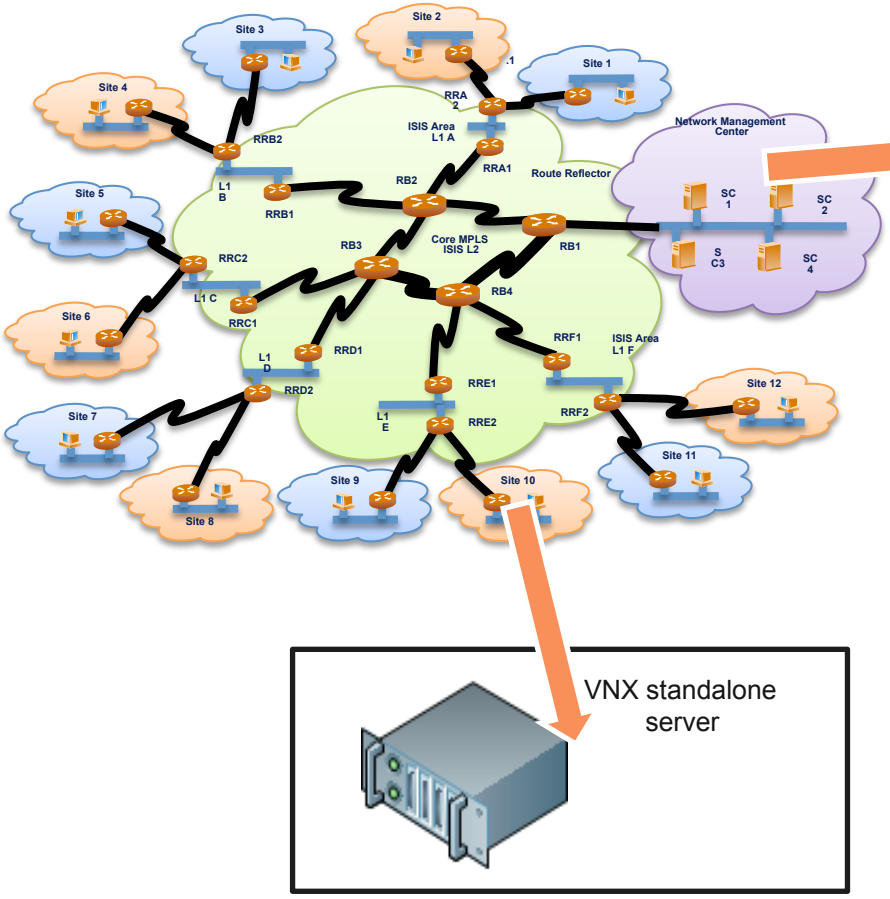

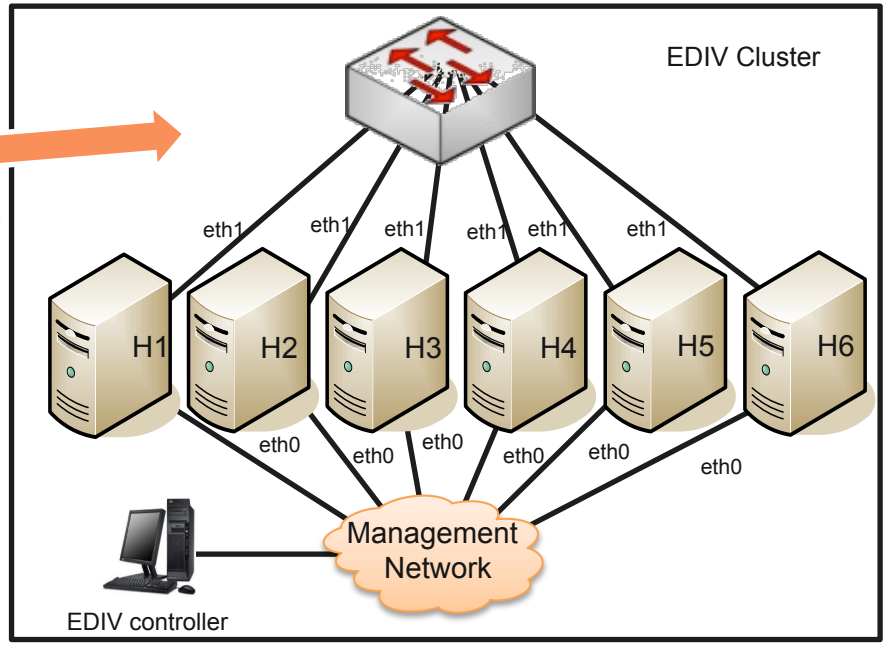

6 x Dell Optiplex 745:
- Core 2 E6400 2.13 GHz
- 3GB of RAM

Sun Fire X4150:
- dual Xeon E5440 (2.83GHz)
- 8 GB of RAM

Conclusions

- VNX helps the management of testbed scenarios, saving on equipment investment and management resources
  - Reusability of testbeds
  - Facilitates sharing testbed infrastructures
- First version of VNX/EDIV successfully used in university computer network laboratories
- Distributed version needs improvements to cope with VMs heterogeneity

Future work

- Finish VNX GUI
- Complete and improve distributed cluster support:
  - Improve cluster interconnection mechanisms (OpenvSwitch,TRILL?)
  - Improve management of server heterogeneity
- Support dynamic scenarios: adding/releasing VMs and networks, VM mobility (libvirt+Sheepdog?)
- Improve network emulation capabilities
- Integrate and test new virtualization platforms (i.e. VMware)
- New types of virtual machines (i.e. Android)
- Support the integration of physical equipment into testbeds (plug-in)
- Testbeds over the Cloud
- Full OVF Environment support
- New applications:
  - Security: dynamic creation of (honeynets)

dit UPM

SVM PARIS 2011

Thanks for your attention!

Virtual Networks over linuX

http://www.dit.upm.es/vnx

Additional Slides

Testbed Example: Firewalls

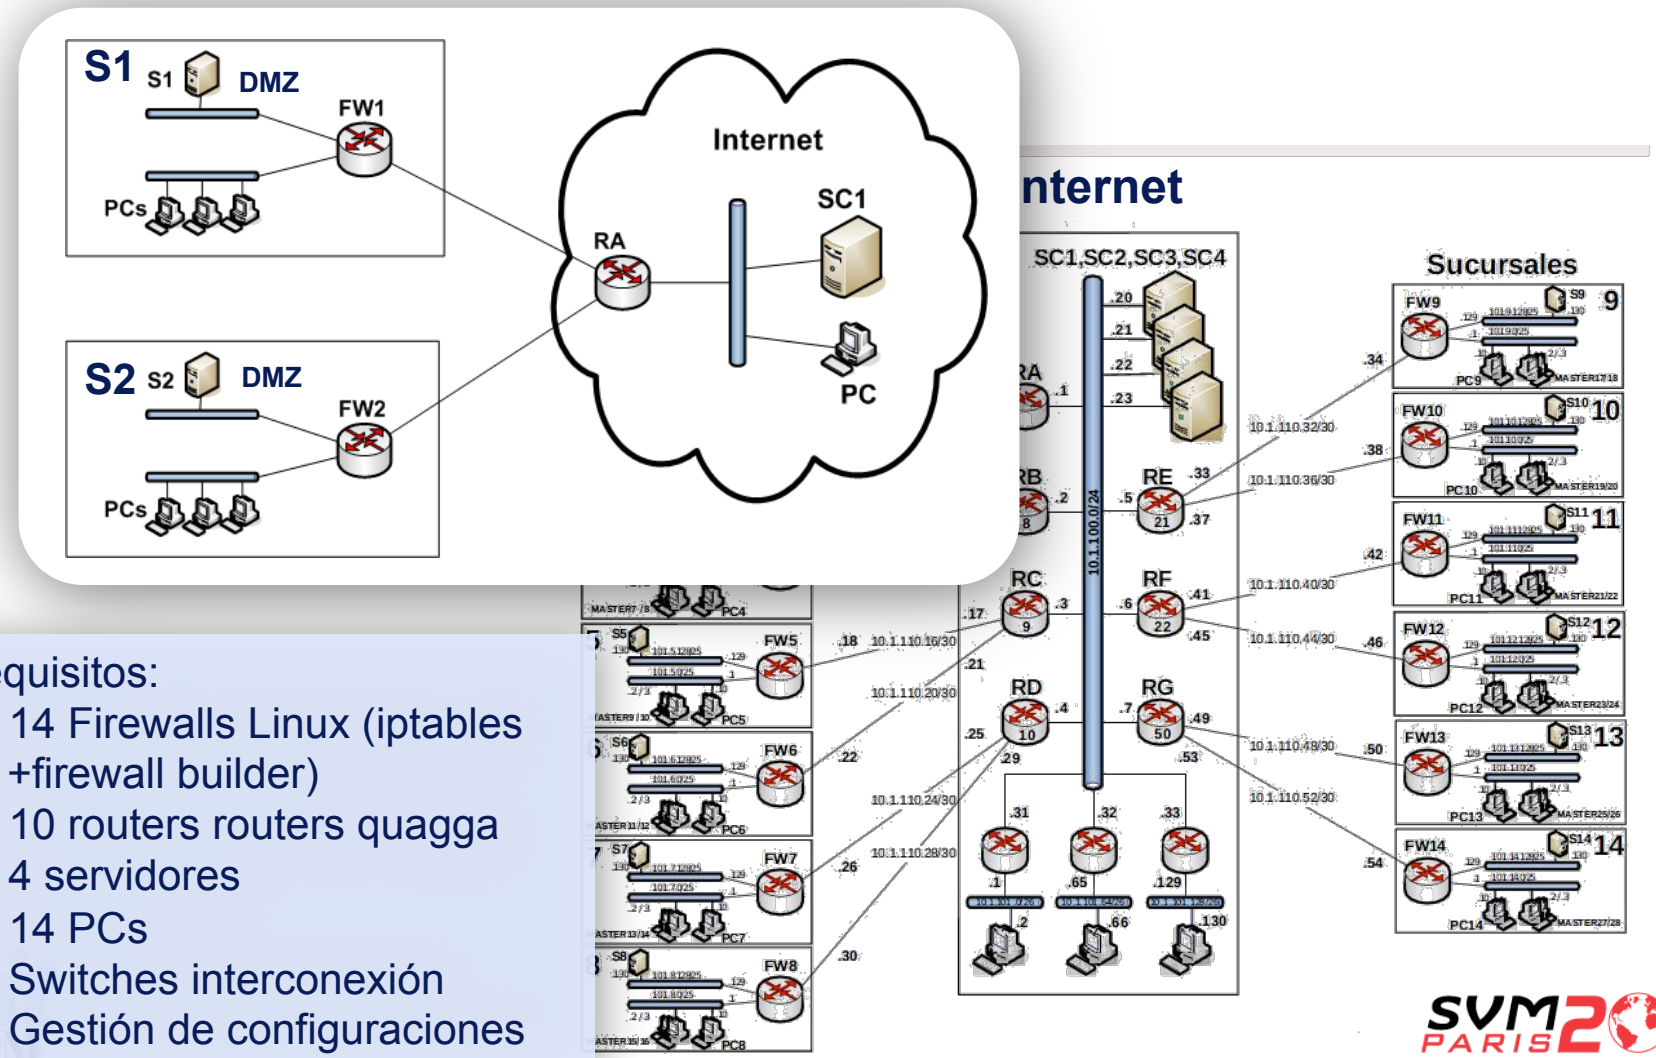

Requisitos:

- 14 Firewalls Linux (iptables +firewall builder)
- 10 routers routers quagga
- 4 servidores
- 14 PCs
- Switches interconexión
- Gestión de configuraciones

VNX Internal API

Primitive	Description
defineVM	Defines a new virtual machine
undefineVM	Undefines an existent virtual machine
startVM	Starts a virtual machine
shutdownVM	Shutdowns a virtual machine in an ordered way.
destroyVM	Kills (switches off) a virtual machine
saveVM	Hibernates a virtual machine (saves state to disk)
restoreVM	Restores a virtual machine previously hibernated
suspendVM	Suspends a virtual machine (saves state to memory)
resumeVM	Resumes a previously suspended virtual machine
rebootVM	Reboots a virtual machine (=shutdown+define+start)
resetVM	Resets a virtual machine (=destroy+define+start)
executeCMD	Executes a command inside the virtual machine

Use Examples

- Starting a VNS:
  - vnx –f escenario.xml --create
- Accessing consoles:
  - vnx –f escenario.xml --console –M vm1
- Executing commands:
  - vnx –f escenario.xml --execute start
- Restarting a VM:
  - vnx –f escenario.xml --reboot –M vm1
- Stopping/releasing the VNS:
  - vnx –f escenario.xml --shutdown
  - vnx –f escenario.xml --destroy

dit UPM

SVM PARIS 2011

Example: tutorial_root1_all

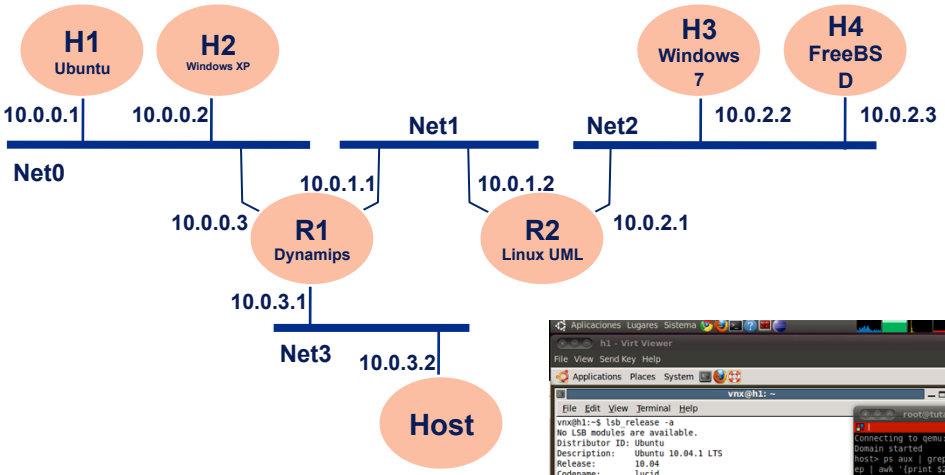

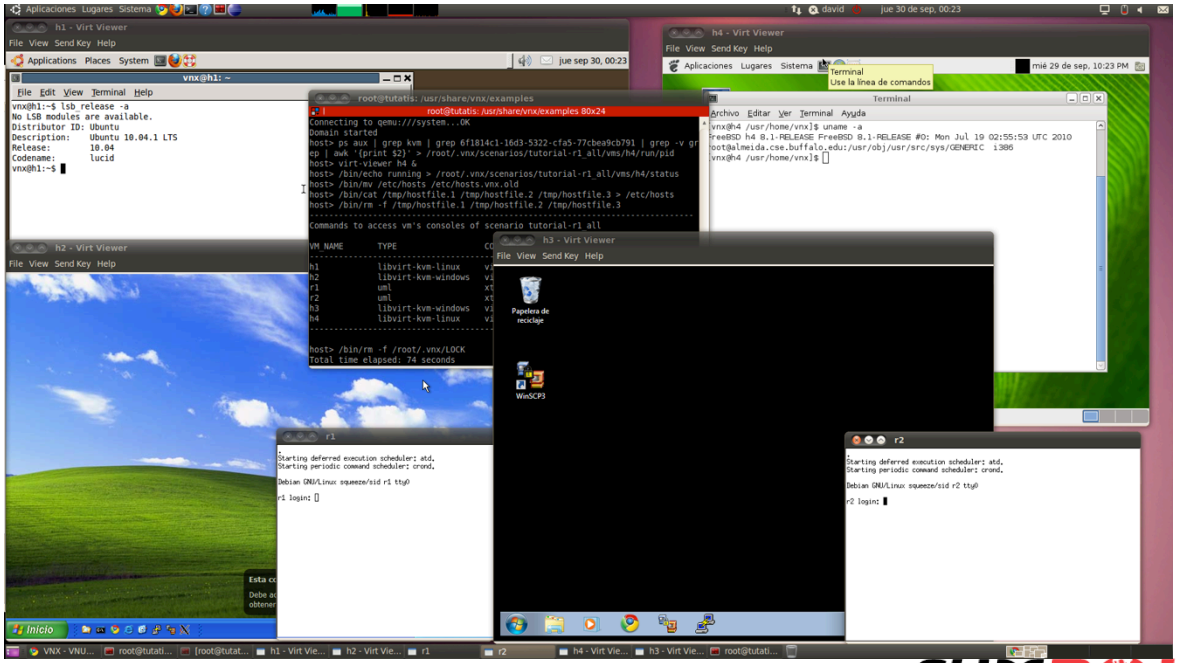

Example: tutorial_root1_all

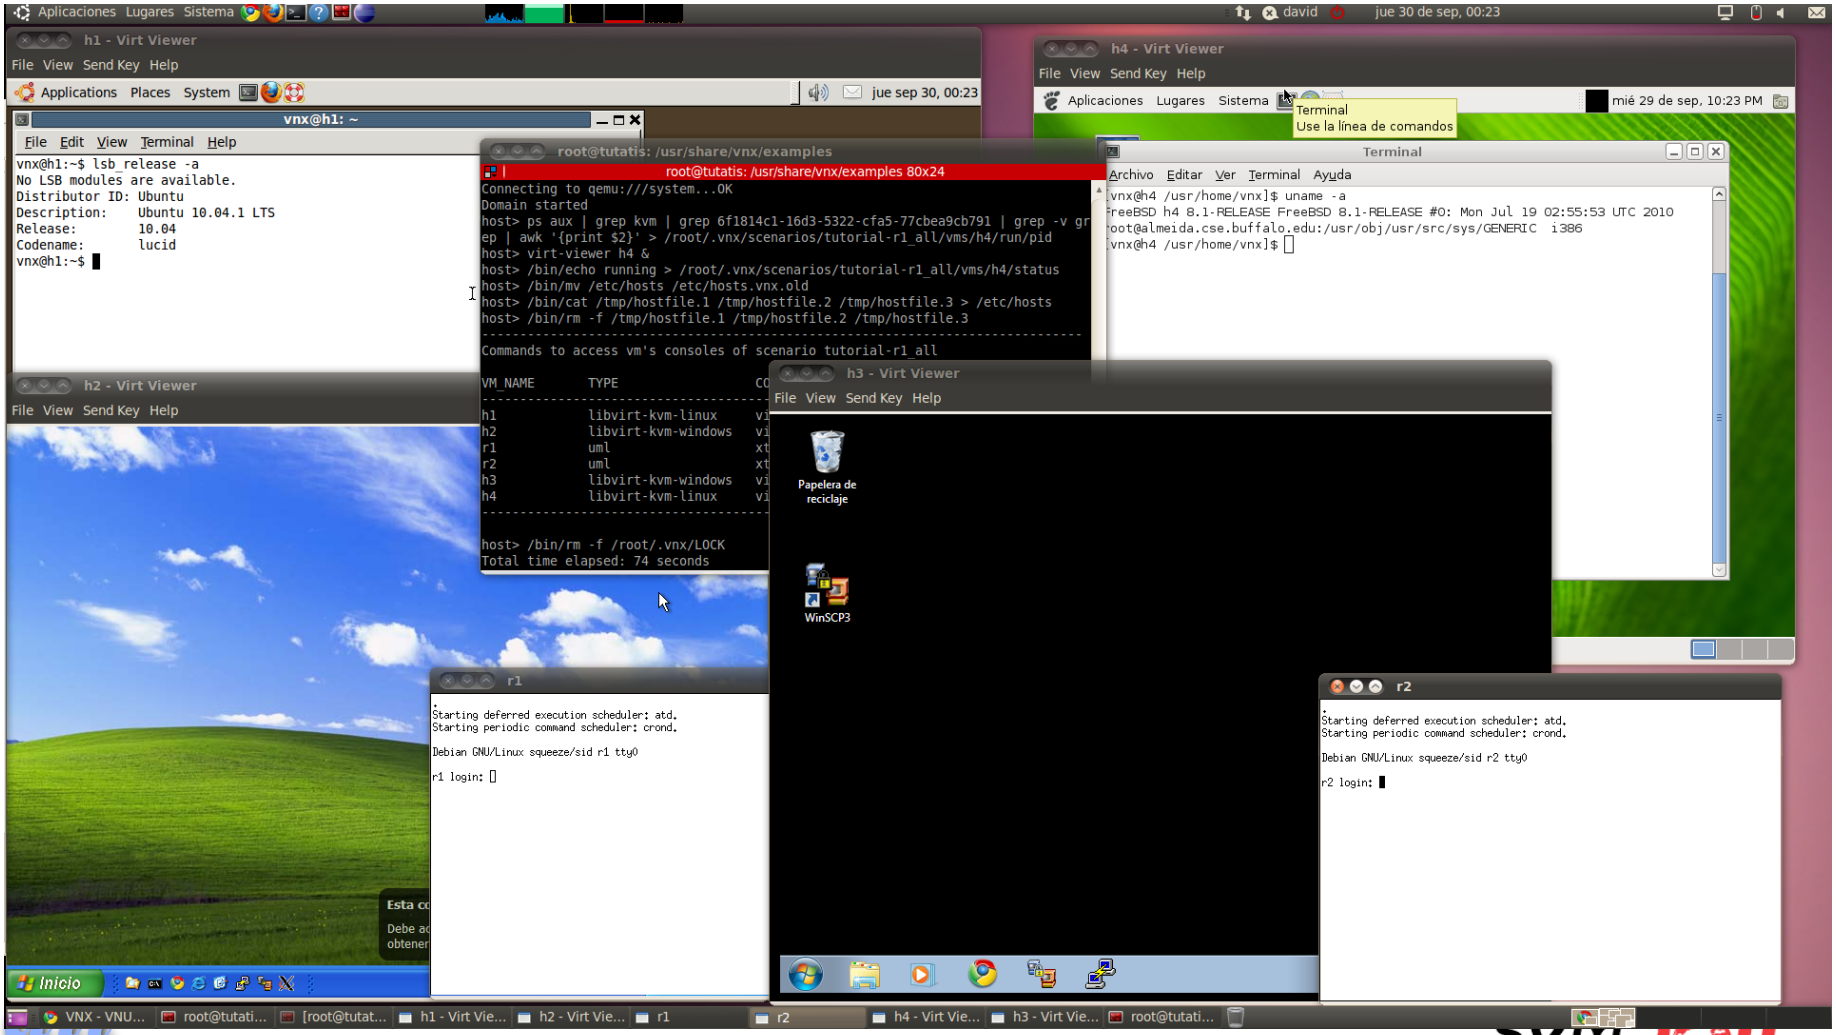

Segur@: Dynamic Deployment of Honeynets

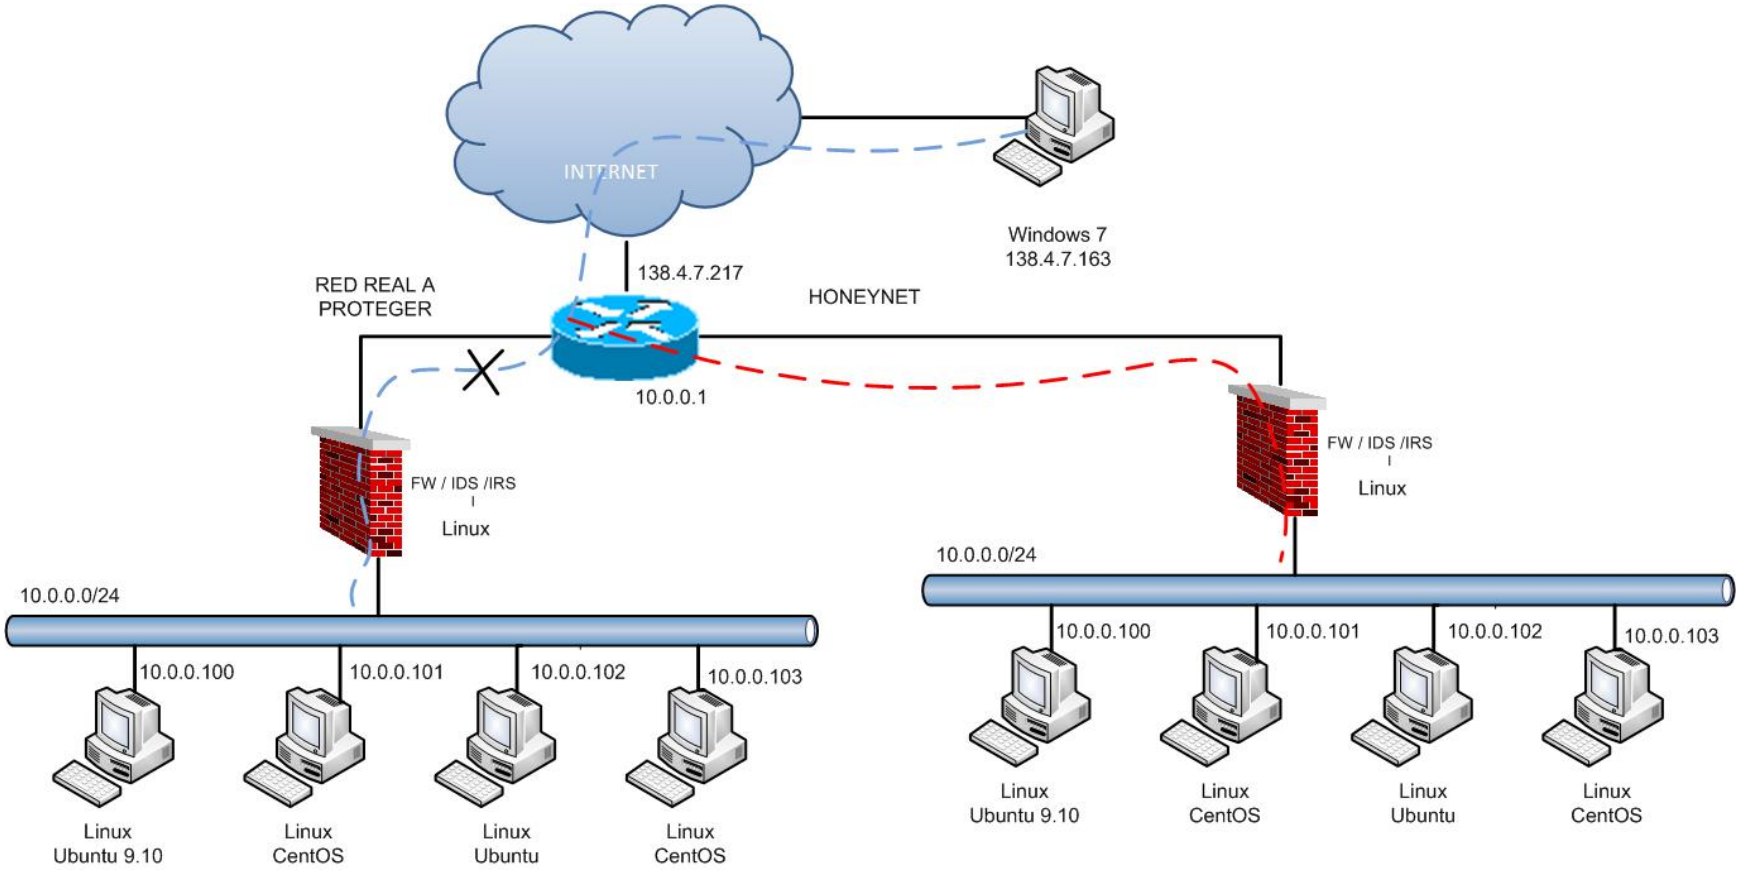

Proyecto Euro6IX

- VNUML se desarrollo inicialmente en el contexto del proyecto Euro6IX
  - Diseño de un nuevo modelo de punto de intercambio (IX) para IPv6 con asignación de direcciones basadas en IX
- Validación basada en VNUML:
  - Escenarios de hasta 20 vms
  - Quagga, BGP, RPSLng, etc

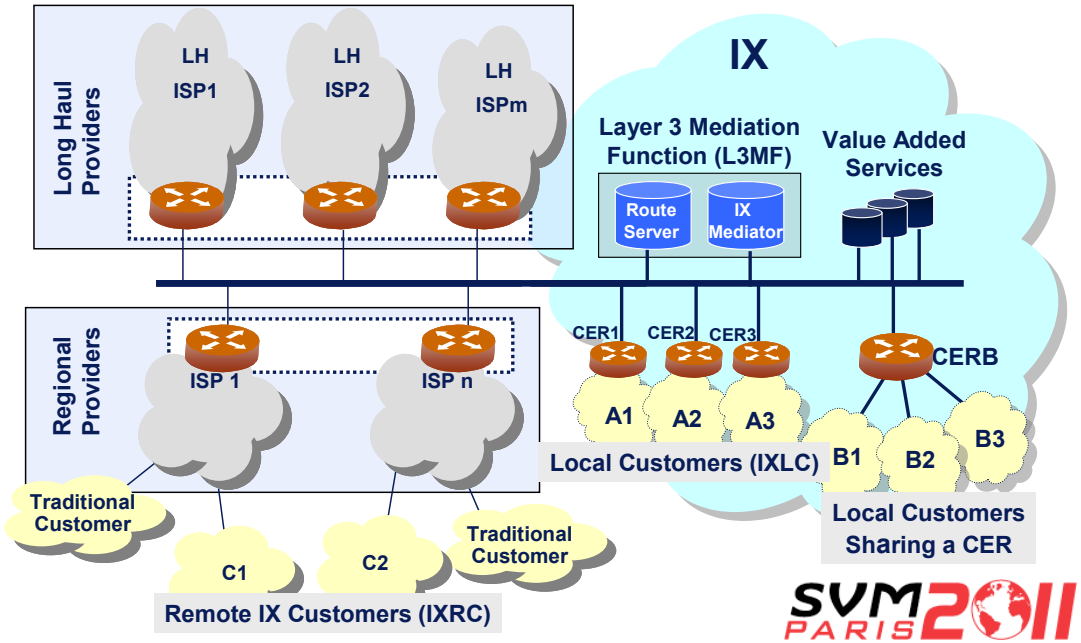

3GPP System Architecture Evolution (SAE)

- Plataforma de pruebas de escenarios 3GPP SAE multi-dominio basados en movi-lidad e IMS
  - Open source IMS
  - Movilidad IPv6 intra e interdominio
- Disponible en la sección de ejemplos:
  - http://www.dit.upm.es/vnuml

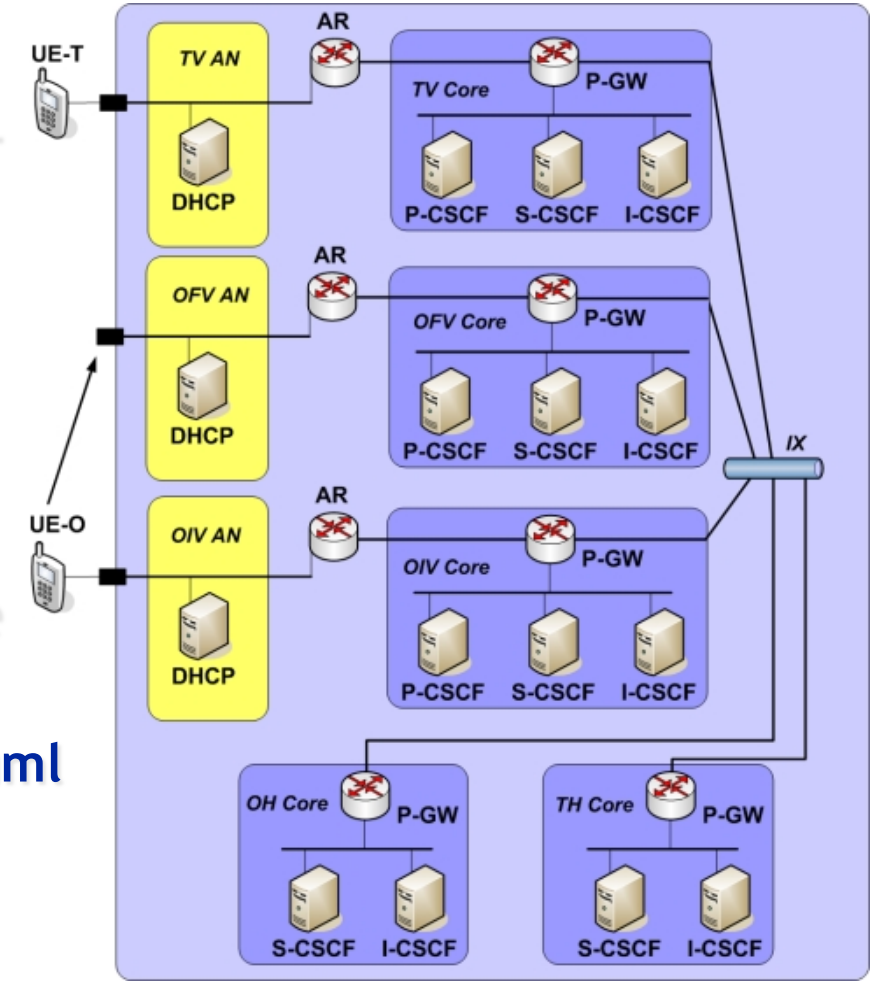

Movilidad interdominio

http://www.dit.upm.es/vnumlwiki/index.php/3gpp-emulator

Segur@: Escenario Red Corporativa

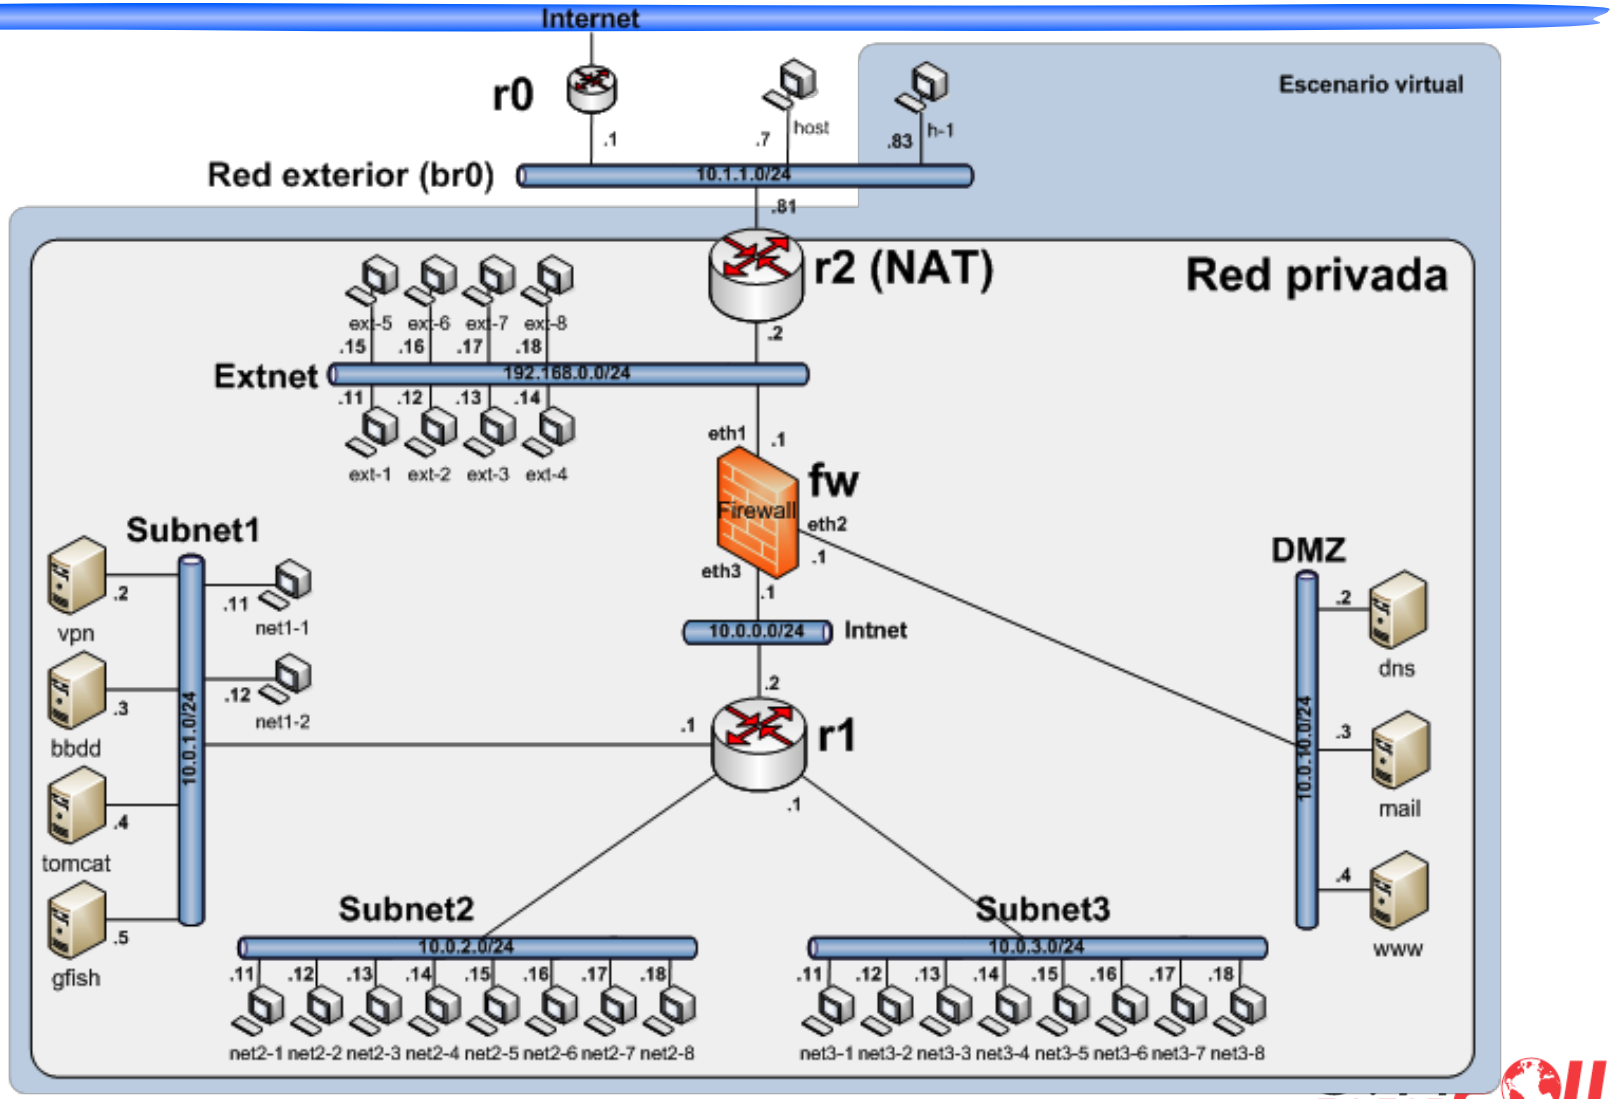

EDIV Demo in TridentCom 2009

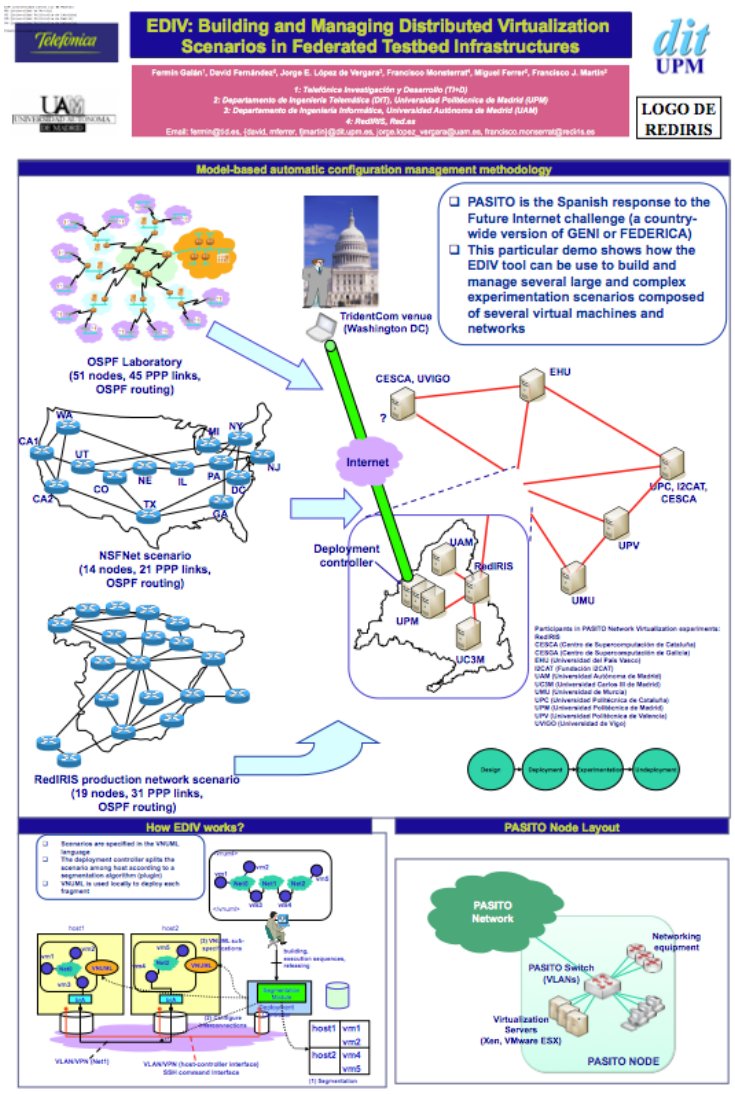

This work is being supported by the Business Oriented Infrastructure (BOI) research initiative within the IT Systems unit at Telefónica I+D and the Spanish Ministry of Industry, Tourism in the framework of the PASITO project.

Escenario de la Demostración: Red Corporativa