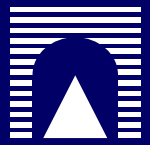

UNIVERSITÀ DEGLI STUDI ROMA TRE
Dipartimento di Informatica e Automazione

Netkit: Easy Emulation of Complex Networks on Inexpensive Hardware

Maurizio Pizzonia

Massimo Rimondini

4th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (TridentCom 2008)
Mar 18th, 2008

Netkit in a Nutshell

software integration project

based on user mode linux

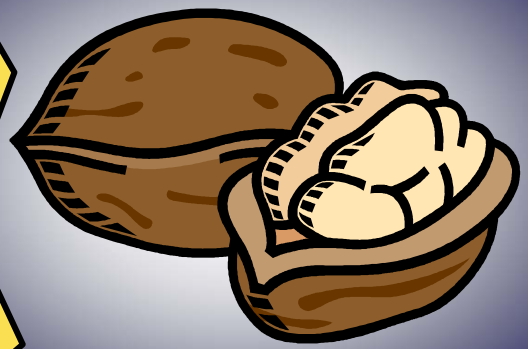

tools to set up virtual labs

ready to use labs & teaching material

Overview

- Architecture
- Supported technologies
- Virtual network setup
- Applications
- Related work
- Scalability

Architecture

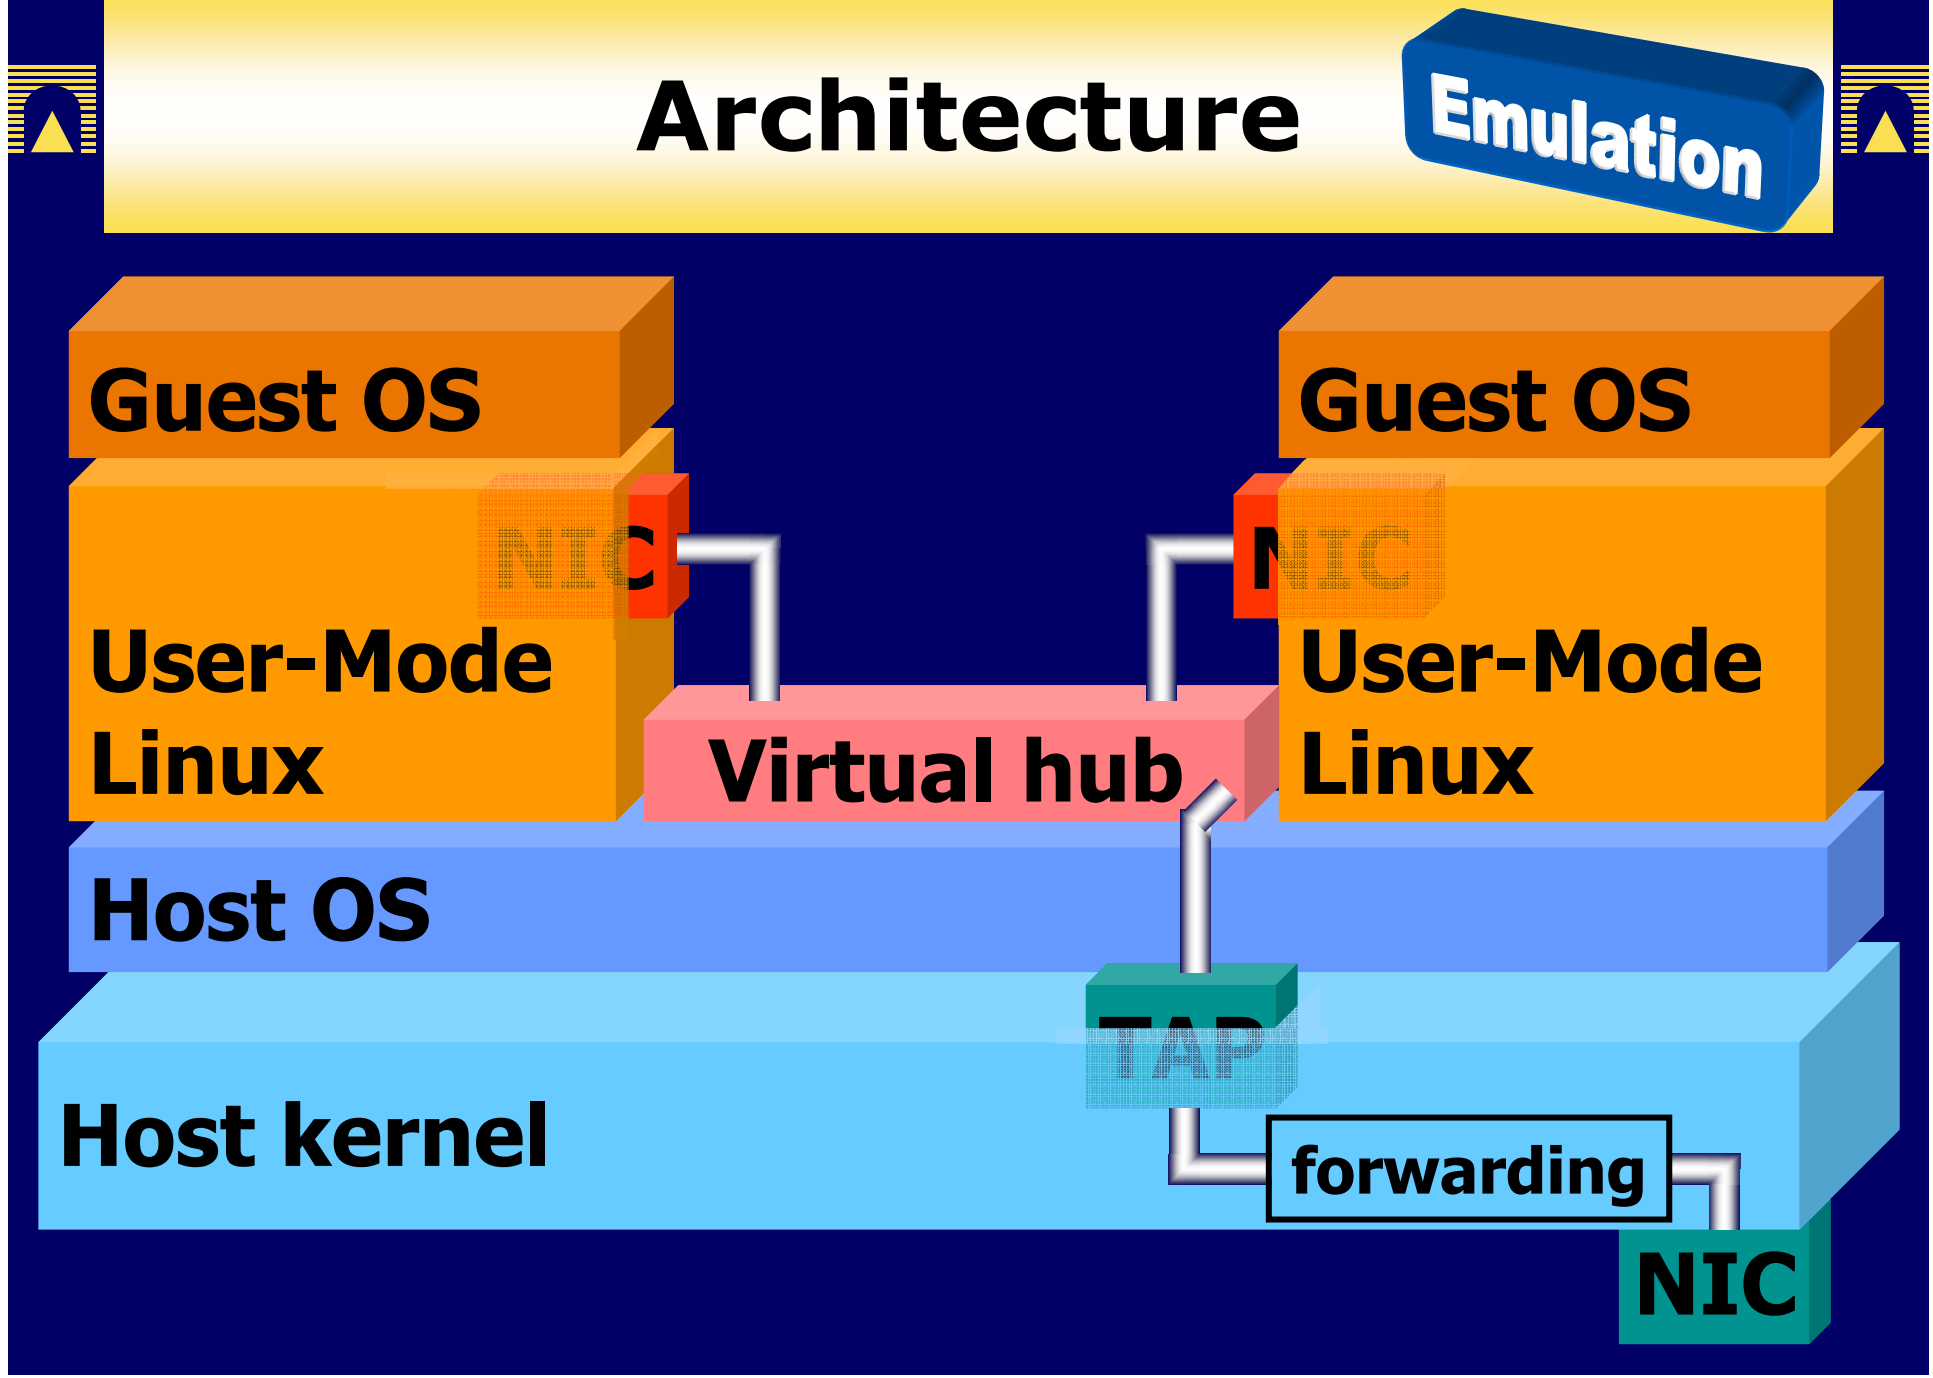

Architecture

Emulation

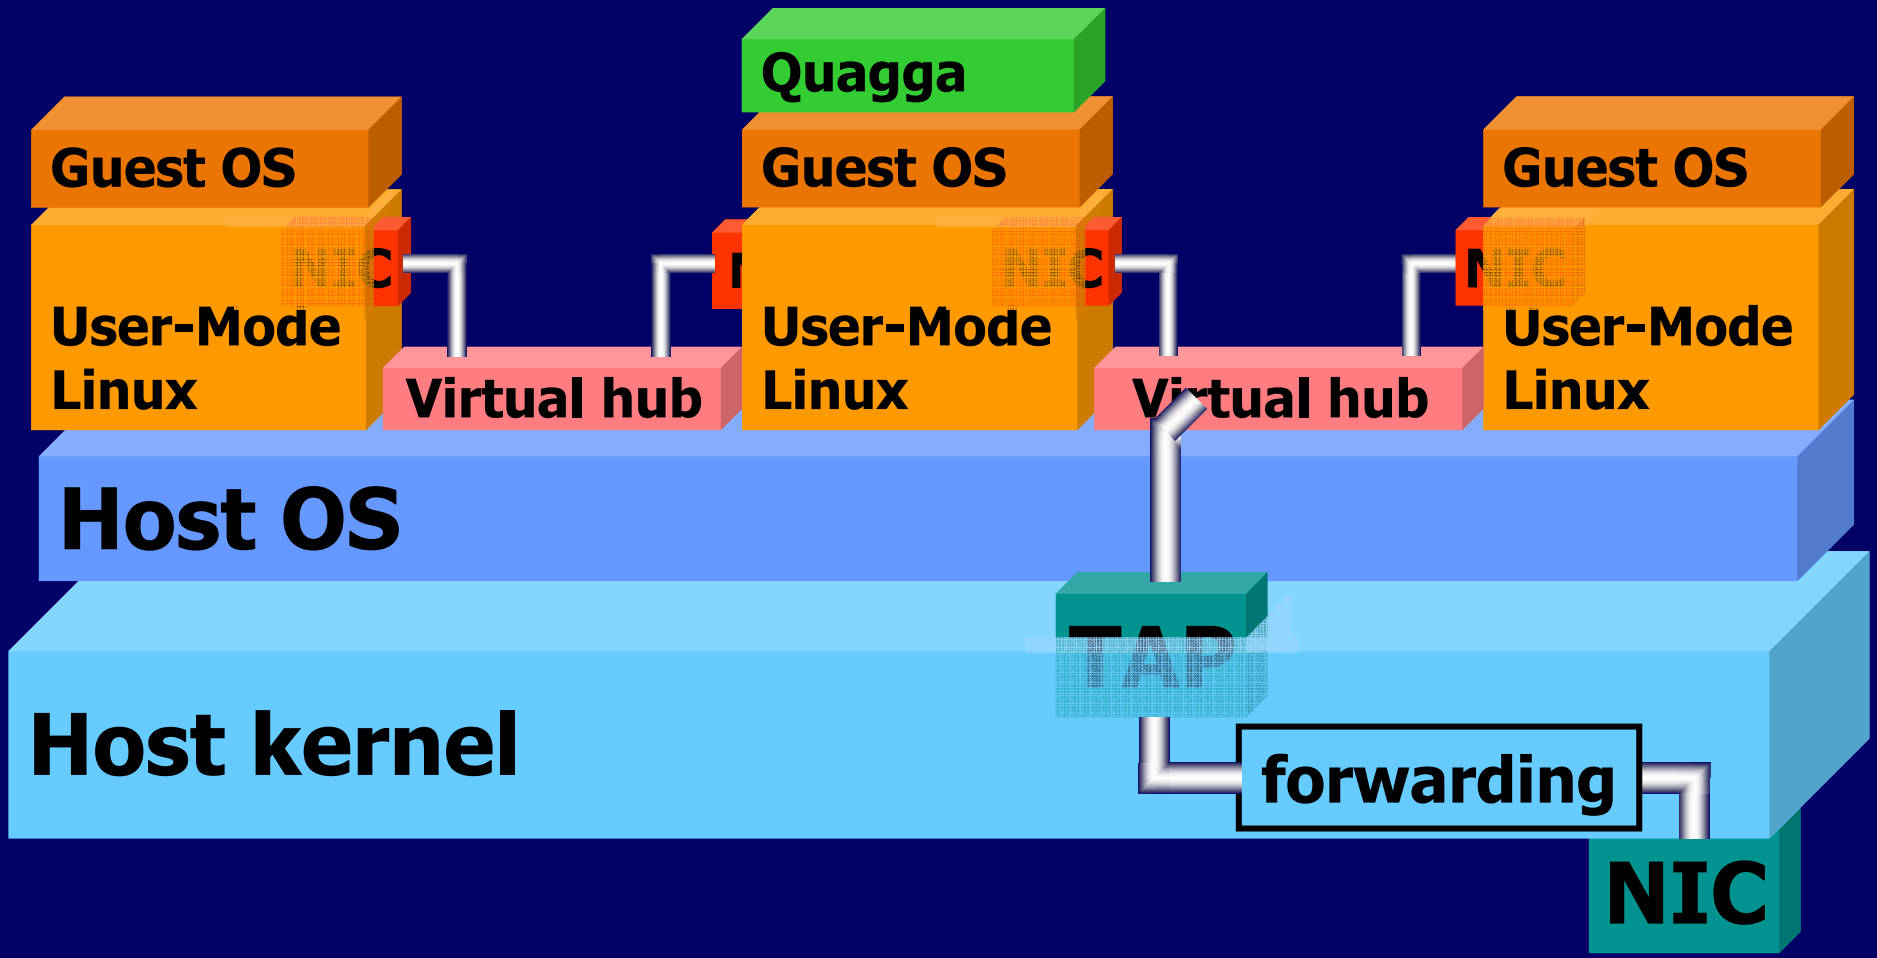

Architecture

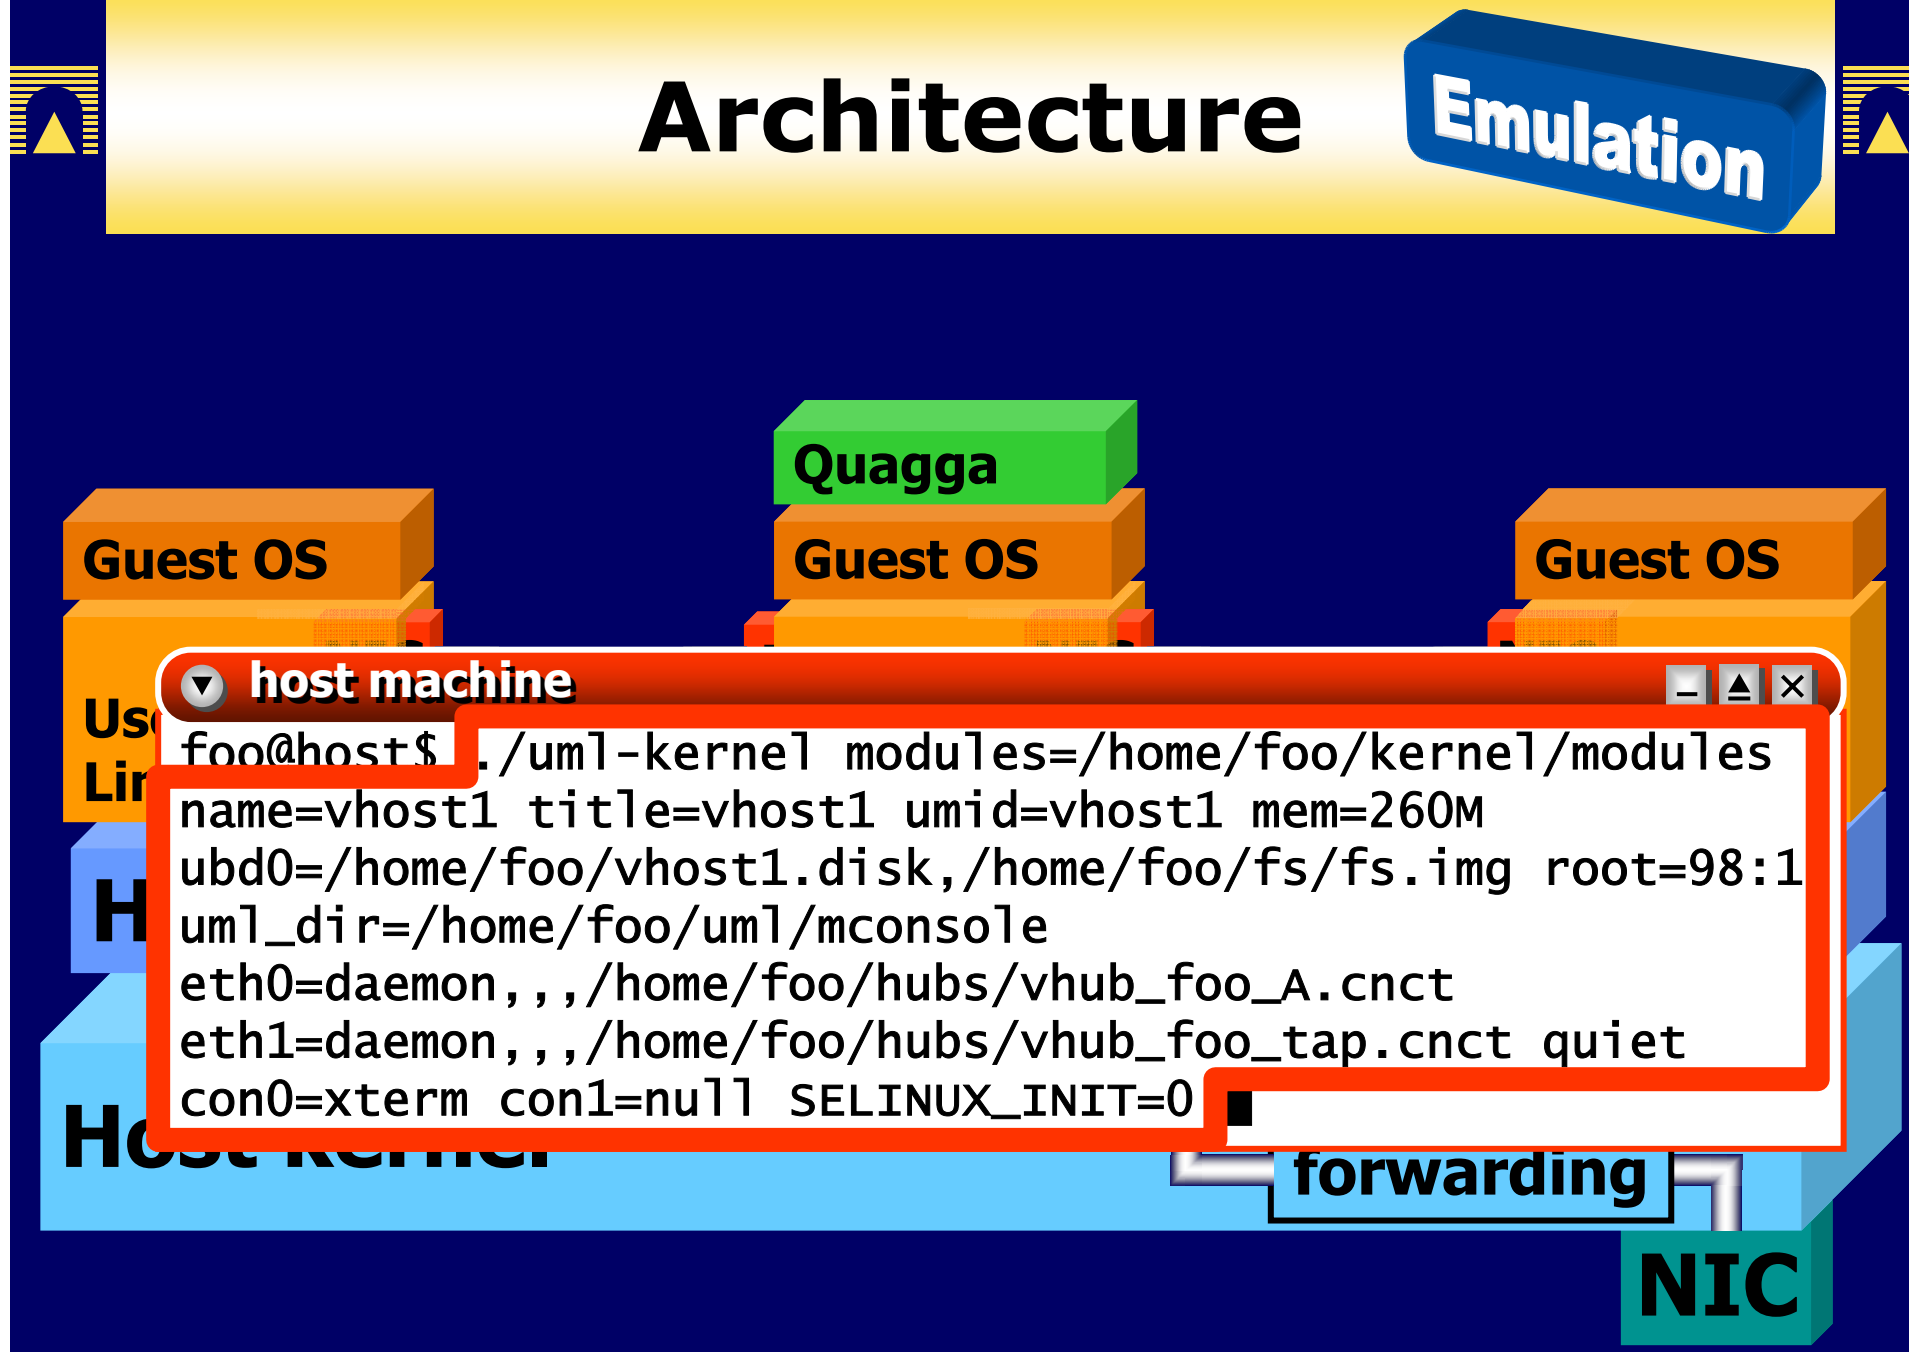

Architecture

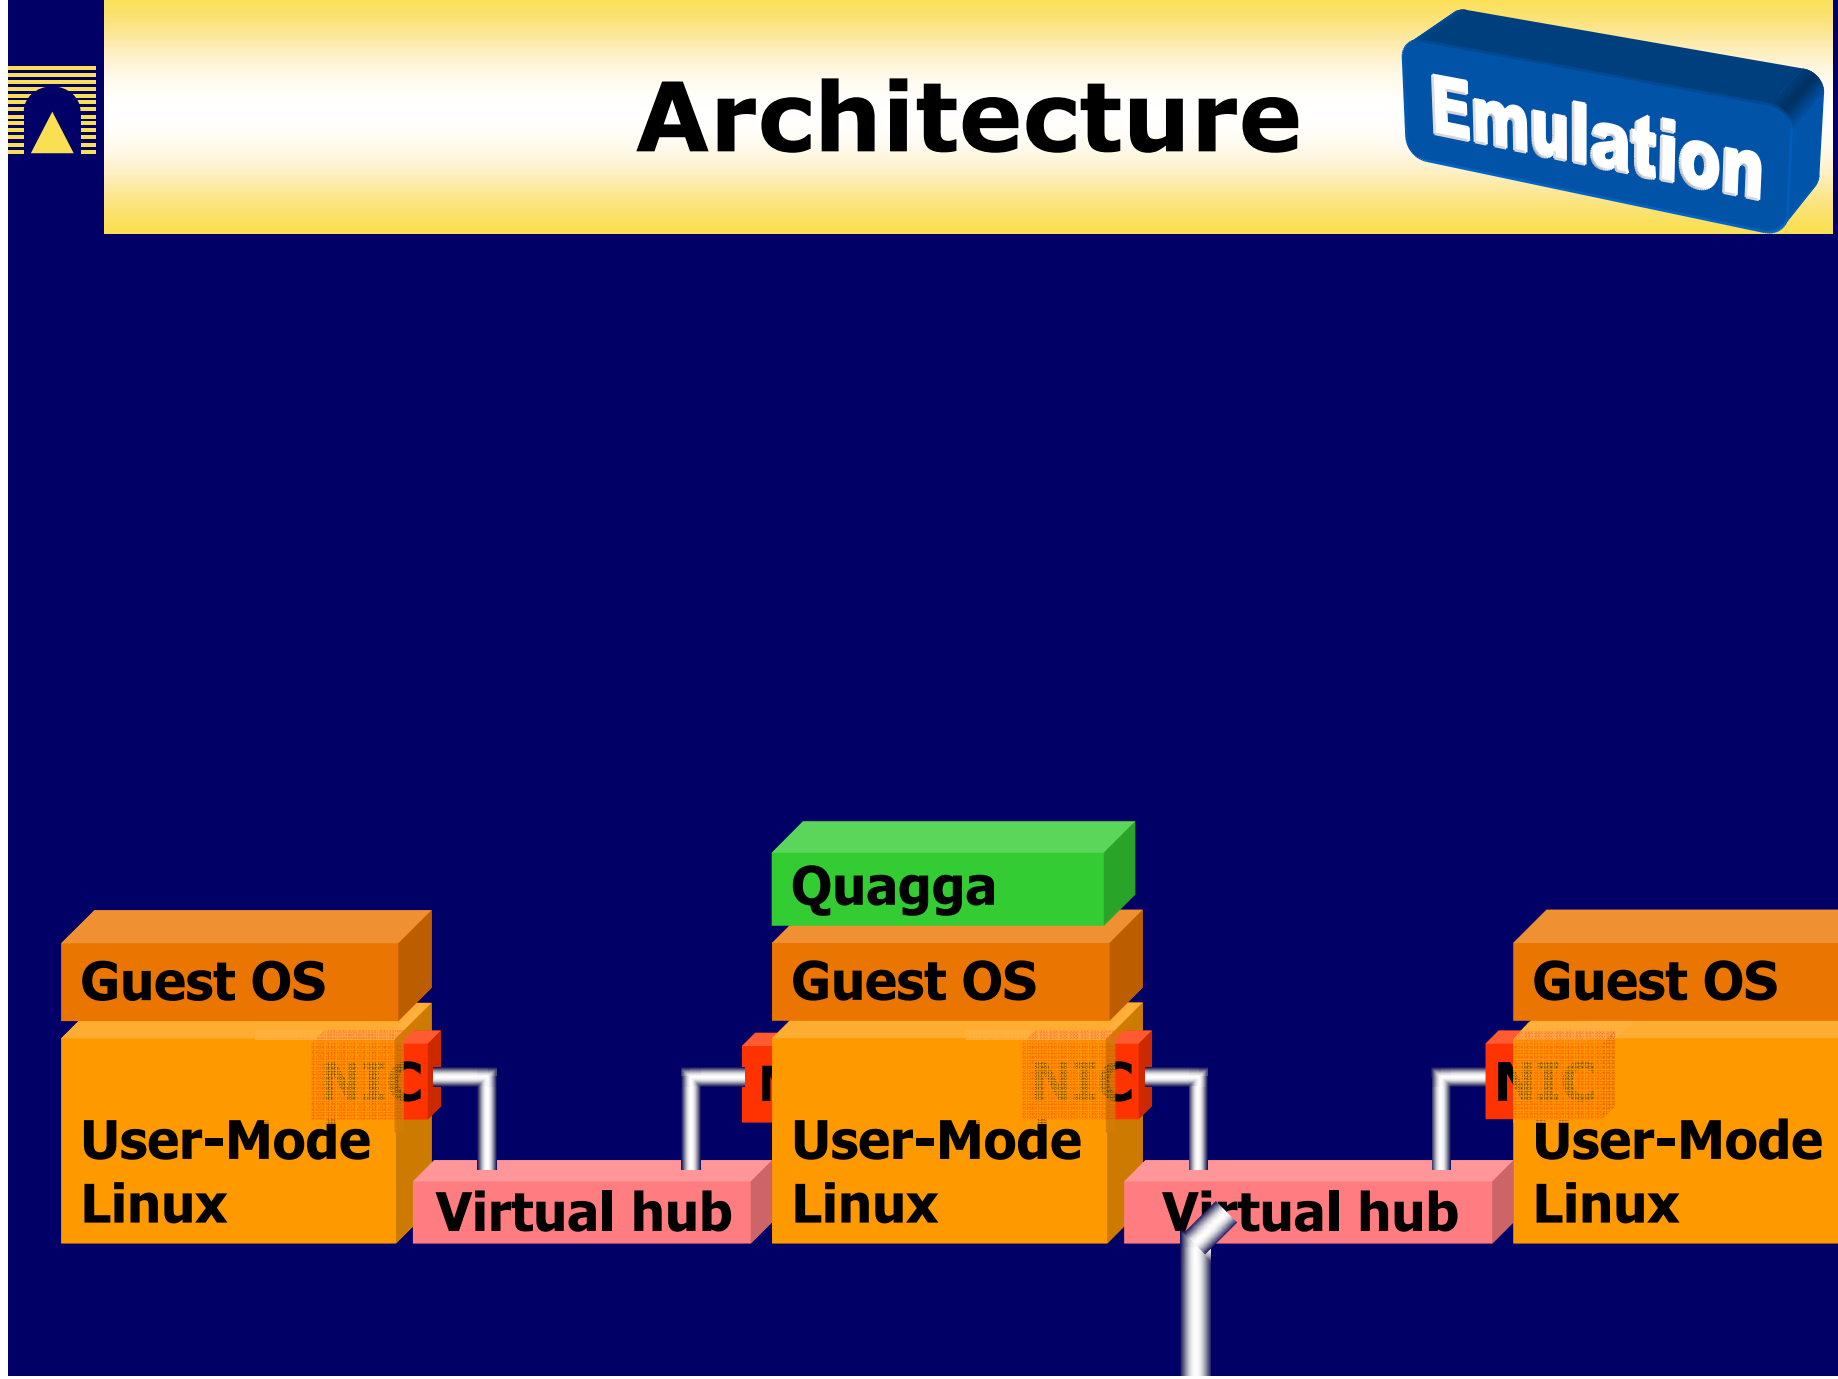

Architecture

- manage a single user mode linux instance (virtual machine)
- customizable hardware & net configuration

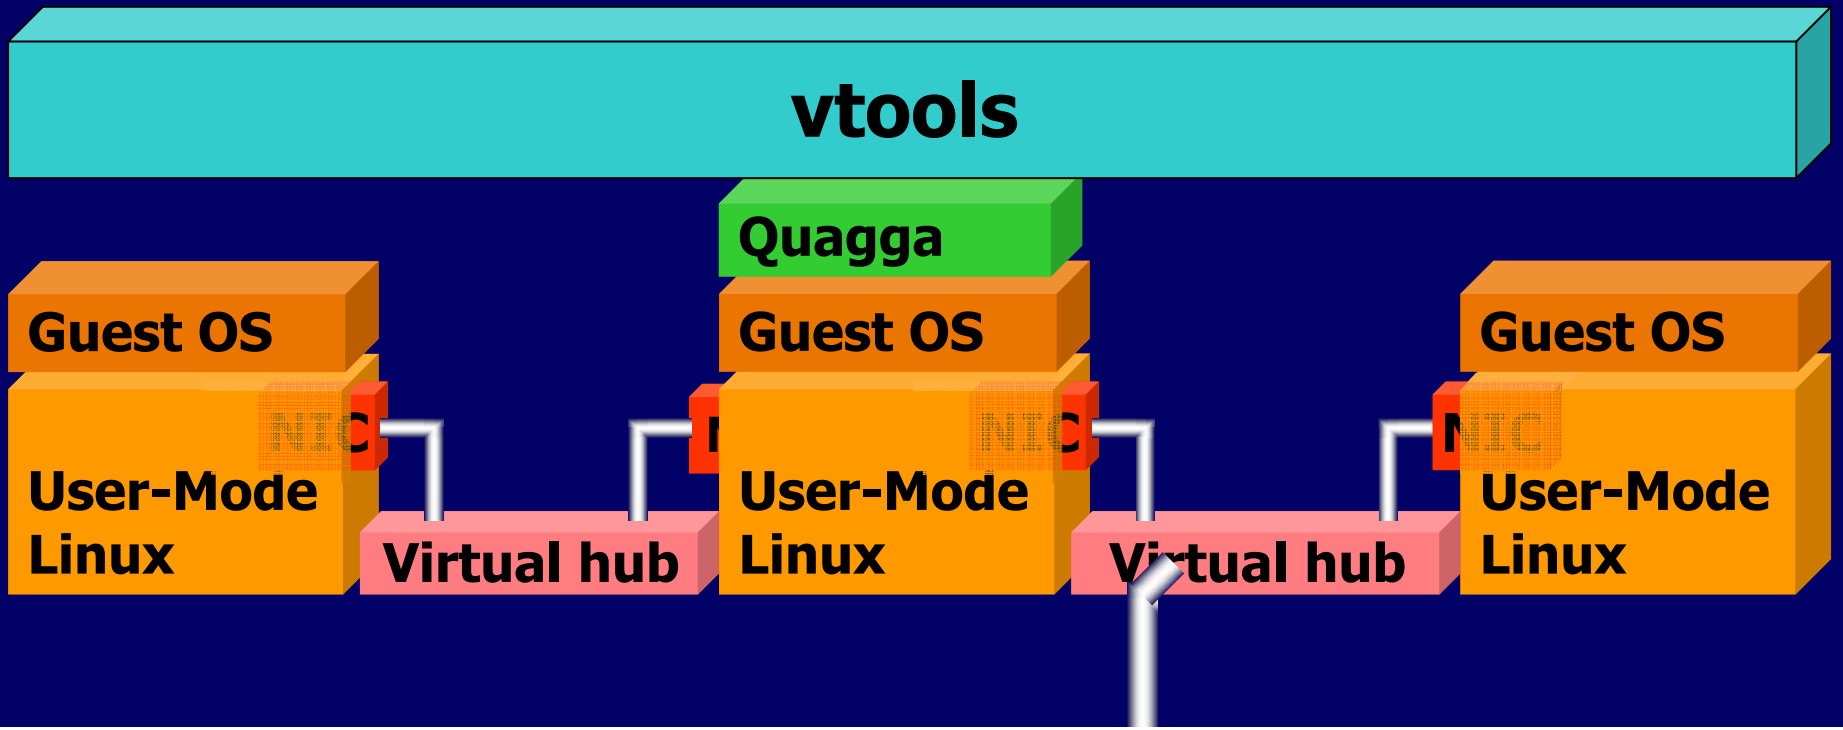

Architecture

- manage a single user mode linux instance (virtual machine)
- customizable hardware & net configuration

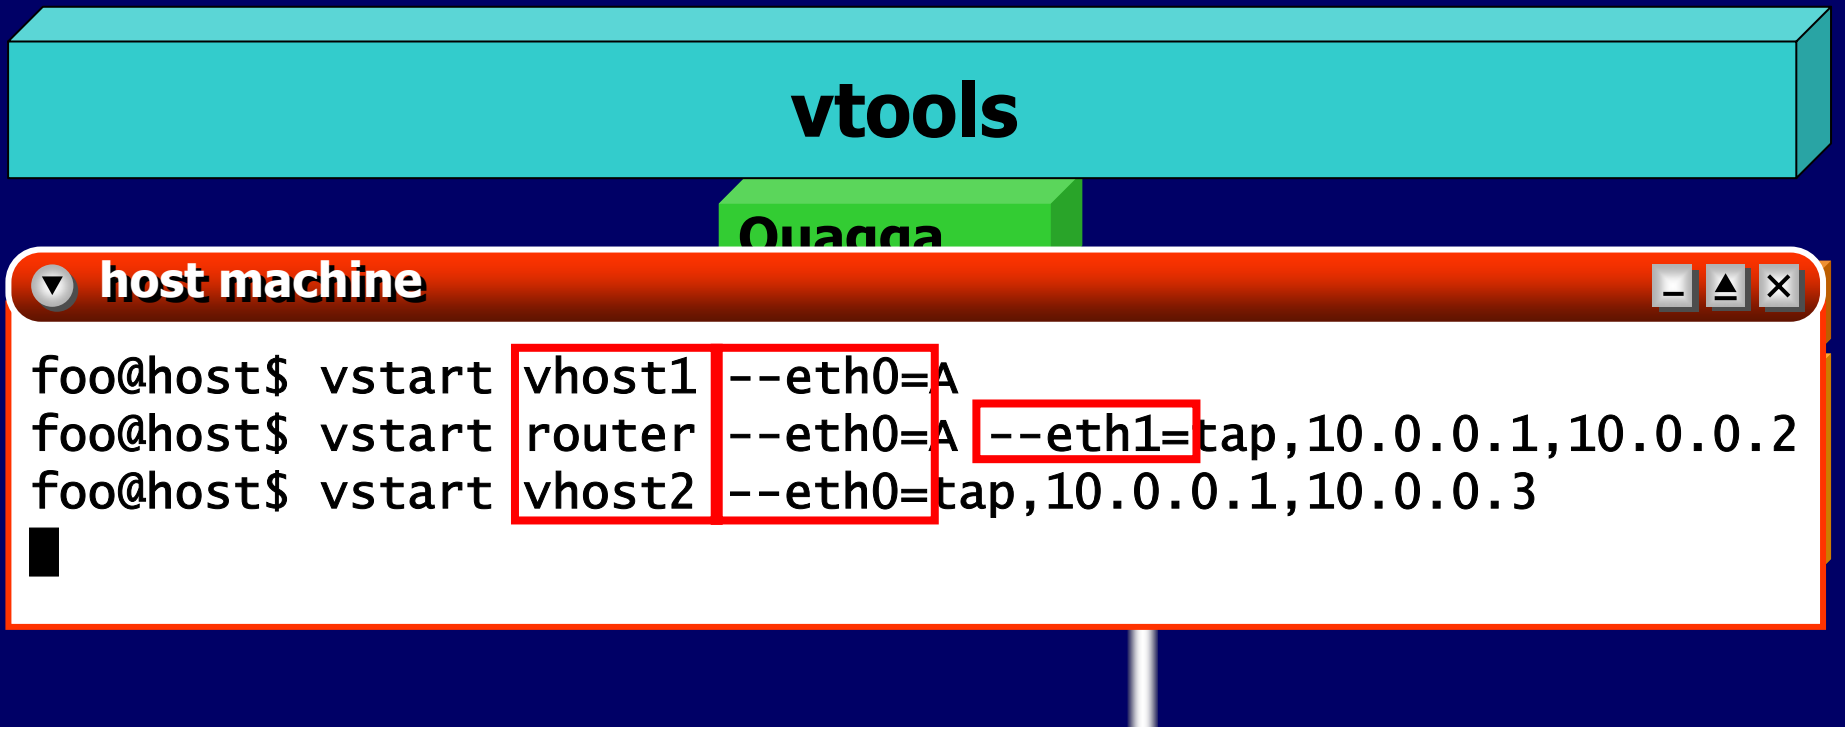

Architecture

- configure multiple virtual machines
- set up a virtual lab

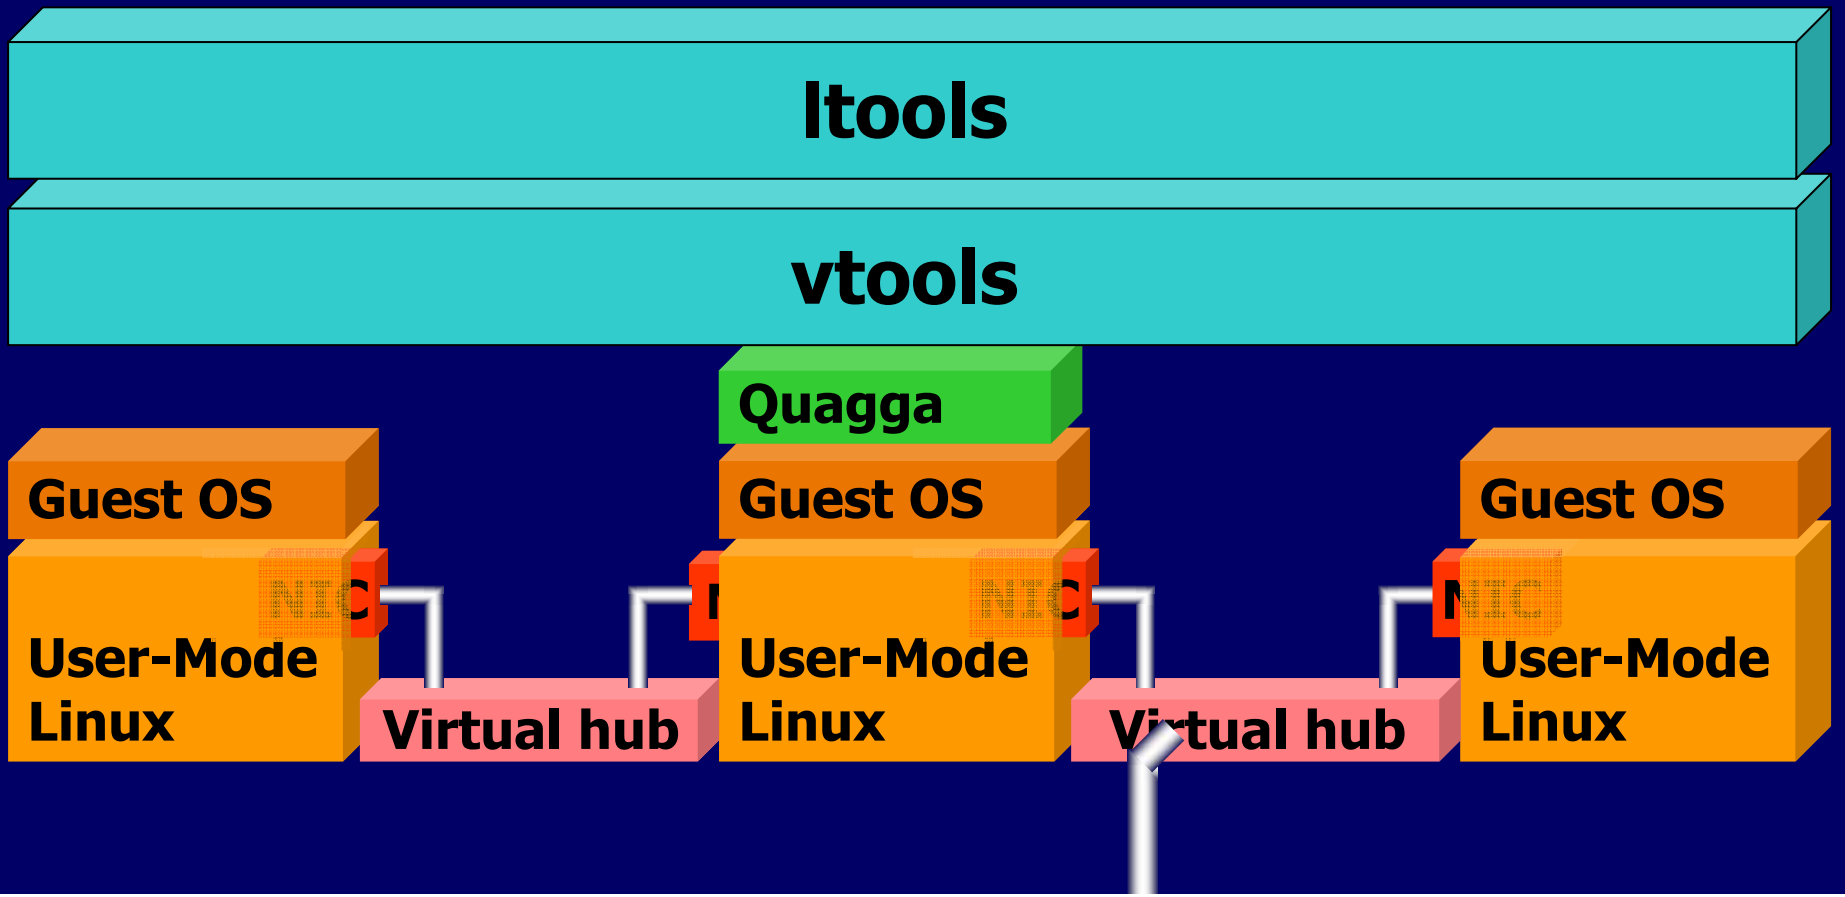

Architecture

- configure multiple virtual machines
- set up a virtual lab

ltools

vtools

Supported Technologies

virtual machine

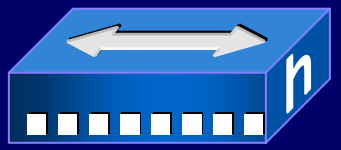

virtual hub software

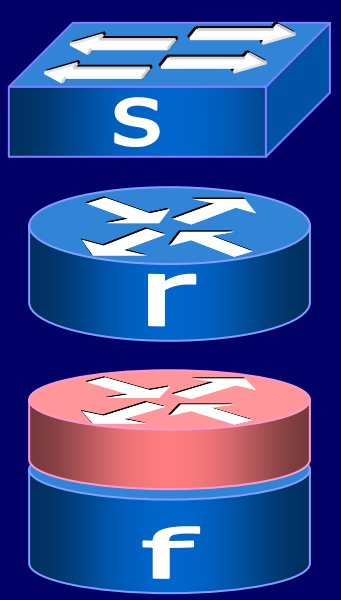

virtual machine running appropriate software

Supported Technologies

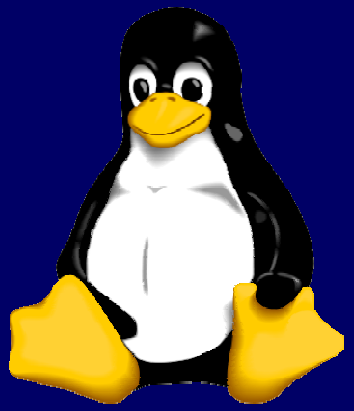

- Ethernet 802.3, 802.1d Bridging and Spanning Tree, 802.1Q VLAN tagging
- MPLS forwarding
- IPv4, IPv6, IP filtering and mangling (NAT, etc.), IPsec (transport and tunnel mode, ESP and AH), ARP, ICMP
- UDP, TCP, GRE tunnels, Equal cost multipath load balancing, PIM-SM
- ...

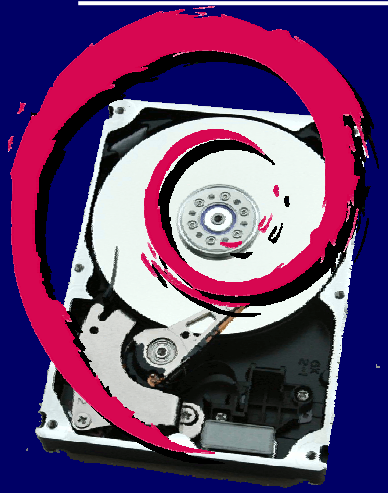

- DHCP, PPP, DNS, HTTP(S), Web proxy, MTA
- FTP, NFS, Samba
- Telnet, SSH
- RIP, OSPF, IS-IS, BGP, SNMP
- RADIUS, PAM, IKE, Snort, Traffic capturing and forging
- Scripting languages
- ...

Supported Technologies

MPLS forwarding

- ng and Spanning Tree, 802.1Q
- ngling (NAT, etc.), IPsec (transport and tunnel mode, ESP and AH), ARP, ICMP
- UDP, TCP, GRE tunnels, Equal cost multipath load balancing, PIM-SM
- ...

RIP, OSPF, IS-IS, BGP, SNMP

- DHCP, PPP, DNS, HTTP(S), Web proxy, MTA
- FTP, NFS, Samba
- rging
- Scripting languages
- ...

Supported Technologies

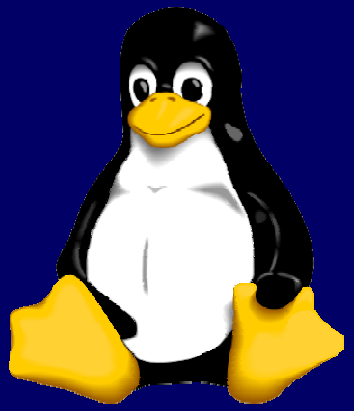

Compile time configurations provided

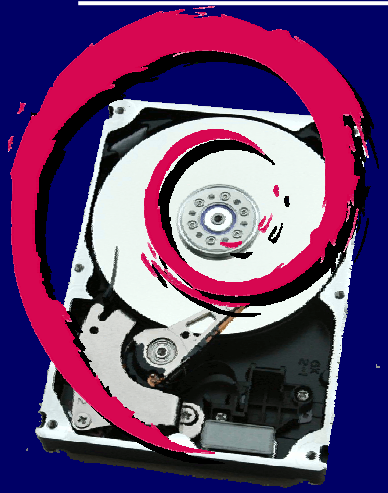

Debian package manager

More can be added

Complex Networks

Setting up a Lab

1. Define nodes

```
mkdir as20r1 as20r2 as200r1 as100r1
```

Setting up a Lab

1. Define nodes

```
mkdir as20r1 as20r2 as200r1 as100r1
```

2. Define topology

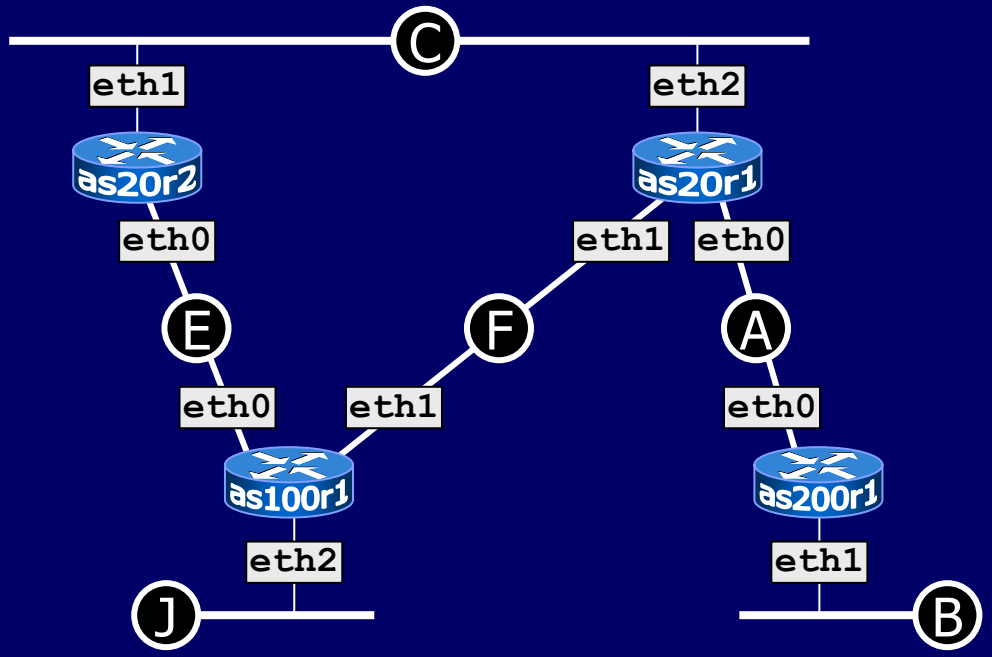


```
as20r1[0]="A"
as20r1[1]="F"
as20r1[2]="C"

as20r2[0]="E"
as20r2[1]="C"

as200r1[0]="A"
as200r1[1]="B"

as100r1[0]="E"
as100r1[1]="F"
as100r1[2]="J"
```

Setting up a Lab

1. Define nodes

```
mkdir as20r1 as20r2 as200r1 as100r1
```

2. Define topology

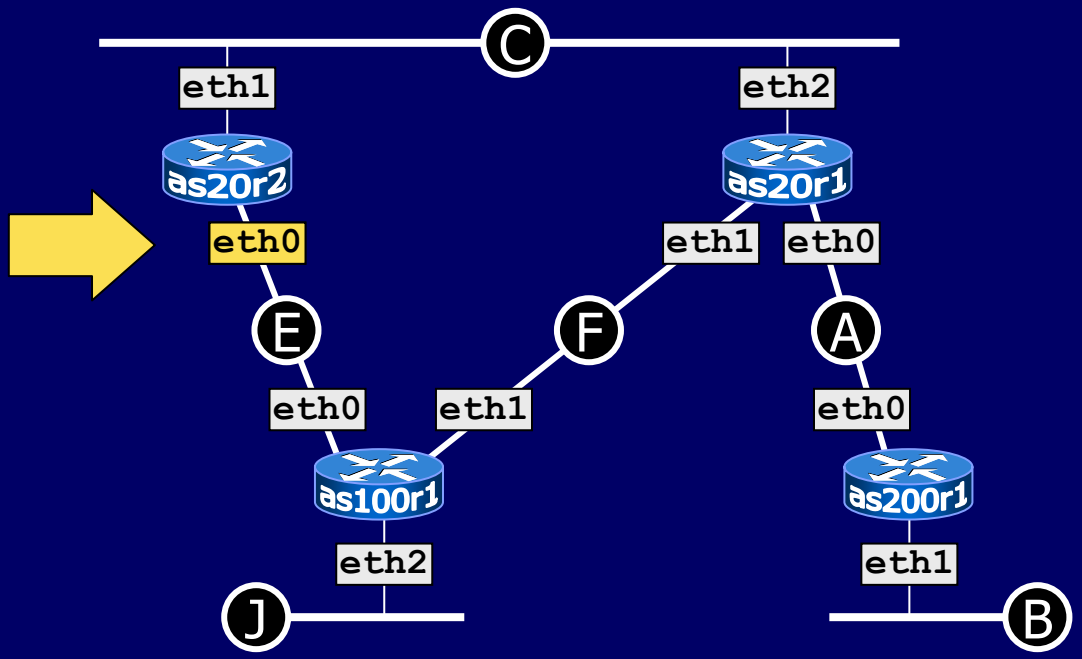


```
as20r1[0]="A"
as20r1[1]="F"
as20r1[2]="C"

as20r2[0]="E"
as20r2[1]="C"

as200r1[0]="A"
as200r1[1]="B"

as100r1[0]="E"
as100r1[1]="F"
as100r1[2]="J"
```

Setting up a Lab

1. Define nodes

```
mkdir as20r1 as20r2 as200r1 as100r1
```

2. Define topology

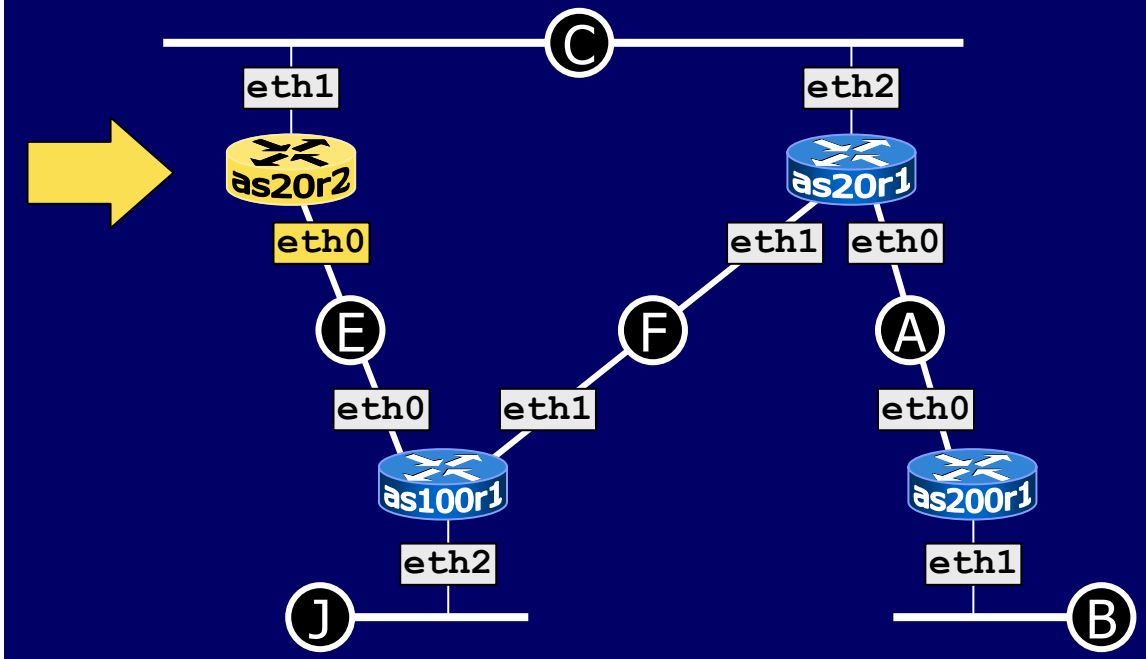


```
as20r1[0]="A"
as20r1[1]="F"
as20r1[2]="C"

as20r2[0]="E"
as20r2[1]="C"

as200r1[0]="A"
as200r1[1]="B"

as100r1[0]="E"
as100r1[1]="F"
as100r1[2]="J"
```

Setting up a Lab

1. Define nodes

```
mkdir as20r1 as20r2 as200r1 as100r1
```

2. Define topology

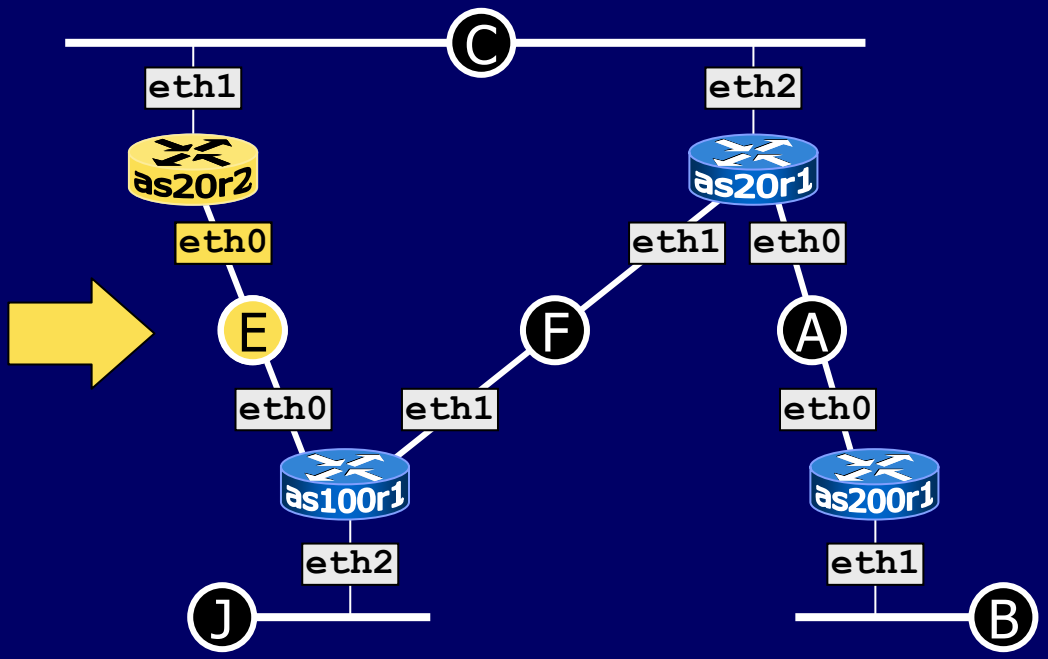


```
as20r1[0]="A"
as20r1[1]="F"
as20r1[2]="C"

as20r2[0]="E"
as20r2[1]="C"

as200r1[0]="A"
as200r1[1]="B"

as100r1[0]="E"
as100r1[1]="F"
as100r1[2]="J"
```

Setting up a Lab

1. Define nodes

   ```
   mkdir as20r1 as20r2 as200r1 as100r1
   ```

2. Define topology
3. Populate configuration files using native syntax

as100r1/etc/quagga/bgpd.conf

```
router bgp 100
network 100.1.0.0/16
neighbor 11.0.0.2 remote-as 20
neighbor 11.0.0.2 description Router as20r2 (primary)
neighbor 11.0.0.2 prefix-list defaultIn in
neighbor 11.0.0.2 prefix-list mineOutOnly out
!
ip prefix-list defaultIn seq 5 permit 0.0.0.0/0
ip prefix-list mineOutOnly seq 5 permit 100.1.0.0/16
```

Setting up a Lab

1. Define nodes

   ```
   mkdir as20r1 as20r2 as200r1 as100r1
   ```

2. Define topology
3. Populate configuration files using native syntax
4. Tell nodes to self configure

as100r1.startup

```
/sbin/ifconfig eth0 11.0.0.1 netmask 255.255.255.252 up
/sbin/ifconfig eth1 11.0.0.5 netmask 255.255.255.252 up
/sbin/ifconfig eth2 100.1.0.1 netmask 255.255.0.0 up
/etc/init.d/quagga start
```

Setting up a Lab


```
├── as100r1
│   └── etc
│       └── zebra
│           ├── bgpd.conf
│           └── daemons
├── as100r1.startup
├── as200r1
│   └── etc
│       └── zebra
│           ├── bgpd.conf
│           └── daemons
├── as200r1.startup
├── as20r1
│   └── etc
│       └── zebra
│           ├── bgpd.conf
│           └── daemons
├── as20r1.startup
├── as20r2
│   └── etc
│       └── zebra
│           ├── bgpd.conf
│           └── daemons
├── as20r2.startup
└── lab.conf
```

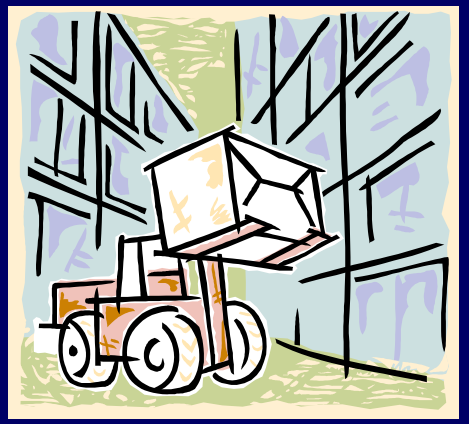

- Typical size: <200KB
- Email/Web friendly

Ready to Use Labs

- Basic topics
  - Routing with static routes and RIP
  - ARP
- Application level
  - DNS
  - Email
- Advanced
  - Bridging
  - STP
- Interdomain Routing
  - Prefix filtering
  - Stub & Multihomed AS
  - Transit AS

Applications

Prospective	Actual

Applications

	Prospective	Actual
Research	• Routing instabilities • Protocol development	• IGP–BGP interactions

Applications

	Prospective	Actual
Research	• Routing instabilities • Protocol development	• IGP–BGP interactions
Operation	• Configuration testing • Troubleshooting	• GARR emulation

Applications

	Prospective	Actual
Research	Routing instabilities Protocol development	IGP–BGP interactions
Operation	Configuration testing Troubleshooting	GARR emulation
Teaching	Basic/Advanced networking courses	University courses/exams

Other Network Emulators


UMLMON

VNUML


Einar

IMUNES

ModelNet


PlanetLab/VINI

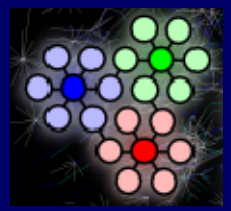
Emulab

Easy Emulation of Complex Networks on Inexpensive Hardware

Other Network Emulators

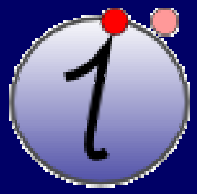

UMLMON

VNUML

- User-Mode Linux
- No kernel/filesystem image
- System wide configuration
- Targeted to sys admins

IMUNES

elNet

PlanetLab/VINI

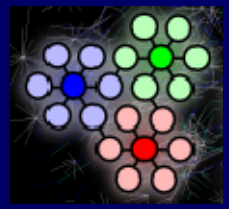

Emulab

Emulation of Complex Networks on Inexpensive Hardware

Other Network Emulators

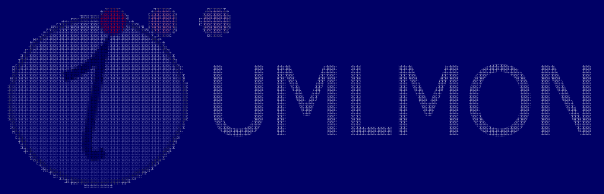

VNUML

Einar

- Live CD only
- Xen based

IMUNES

lodelNet

PlanetLab/VINI

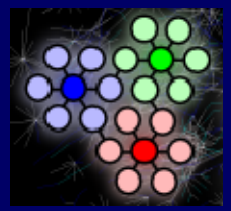

Emulab

Easy Emulation of

on

Other Network Emulators

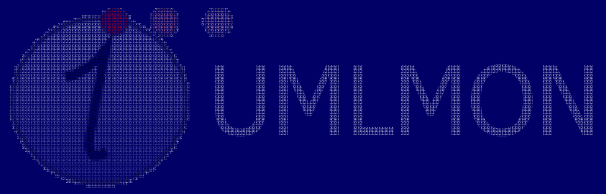

VNUML

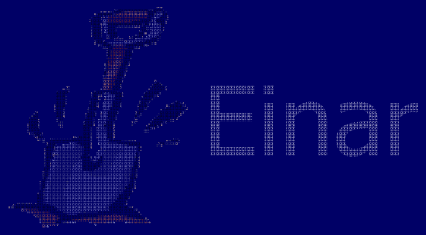

IMUNES

- Live CD only
- FreeBSD kernel

ModelN

University of Zagreb

- Marko Zec, M. Mikuc. Operating System Support for Integrated Network Emulation in IMUNES. Proc. ASPLOS-XI, Oct 2004.
- Marko Zec. Implementing a Clonable Network Stack in the FreeBSD Kernel. Proc. 2003 USENIX Annual Technical Conference, Jun 2003.

Other Network Emulators

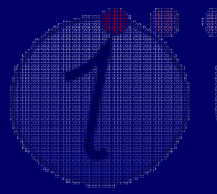

UMLMON

VNUML

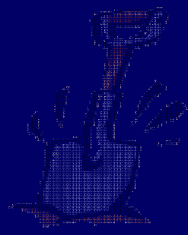

Einar

IMUNES

- Live CD only
- FreeBSD kernel

ModelN

PlanetLab/VINI

Emulab

Easy Emulation of on

Other Network Emulators

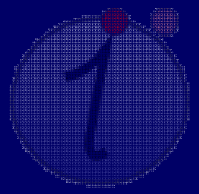

UMLMON

VNUML

Einar

IMUNES

ModelNet

PlanetLab/VINI

Emulab

Easy Emulation of Complex Networks on Inexpensive Hardware

Other Network Emulators

Communities started around 2002

- L. Peterson, A. Bavier, M. Fiuczynski, and S. Muir. Experiences Building PlanetLab. Proc. OSDI 2006, Nov 2006.
- L. Peterson and T. Roscoe. The Design Principles of PlanetLab. ACM SIGOPS Operating Systems Review, 40(1):11–16, 2006.
- A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In VINI Veritas: Realistic and Controlled Network Experimentation. Proc. SIGCOMM 2006, Sep 2006.
- P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. A Basis for Systematic Analysis of Network Topologies. Proc. SIGCOMM 2006, Sep 2006.
- R. Ricci, J. Duerig, P. Sanaga, D. Gebhardt, M. Hibler, K. Atkinson, J. Zhang, S. Kasera, and J. Lepreau. The Flexlab Approach to Realistic Evaluation of Networked Systems. Proc. USENIX NSDI 2007, Apr 2007.
- E. Eide, L. Stoller, and J. Lepreau. An Experimentation Workbench for Replayable Networking Research. Proc. USENIX NSDI 2007, Apr 2007.

Easy Emulation of Complex Networks on Inexpensive Hardware

Other Network Emulators

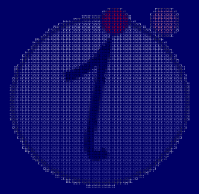

UMLMON

VNUML

Einar

IMUNES

Pl

- Large scale
- Server clusters
- Require affiliation & approval

Emulab

of Complex Networks

on

Other Network Emulators

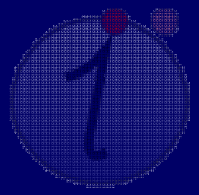

UMLMON

VNUML

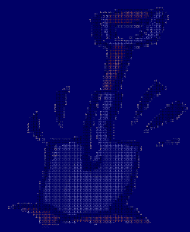

Einar

IMUNES

ModelNet

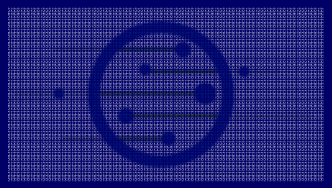

PlanetLab/VINI

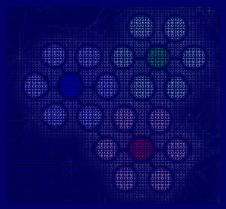

Emulab

Easy Emulation of Complex Networks on Inexpensive Hardware

Other Network Emulators

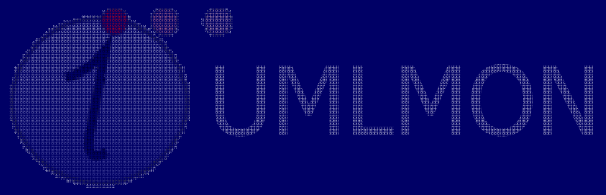

Universidad Politécnica de Madrid

- F. Galan, D. Fernández. Distributed Virtualization Scenarios Using VNUML. Proc. System and Virtualization Management Workshop (SVM'07). Oct 2007
- D. Fernández, F. Galán, T. de Miguel. Study and Emulation of IPv6 Internet Exchange (IX) based Addressing Models. IEEE Communications Magazine, 42(1):105–112, Jan 2004
- F. Galán, D. Fernández, J. Ruiz, O. Walid, and T. de Miguel. Use of Virtualization Tools in Computer Network Laboratories. Proc. International Conference on Information Technology Based Higher Education and Training (ITHET 2004), Jun 2004.

Other Network Emulators

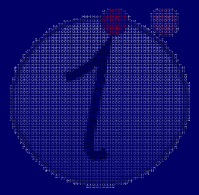

UMLMON

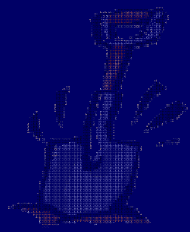

Einar

ModelNet

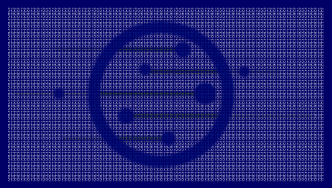

PlanetLab/VINI

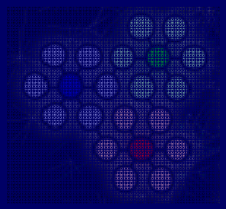

Emulab

VNUML

- User-Mode Linux
- XML
- Conceived to run as root

Easy Emulation of Complex Networks on Inexpensive Hardware

Inexpensive Hardware

Scalability

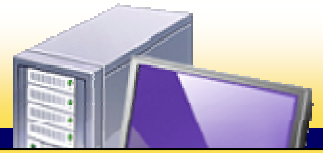

Pentium 4 3.2GHz, 2GB RAM (~350 €)

Current entry price: >400 €

- Startup time

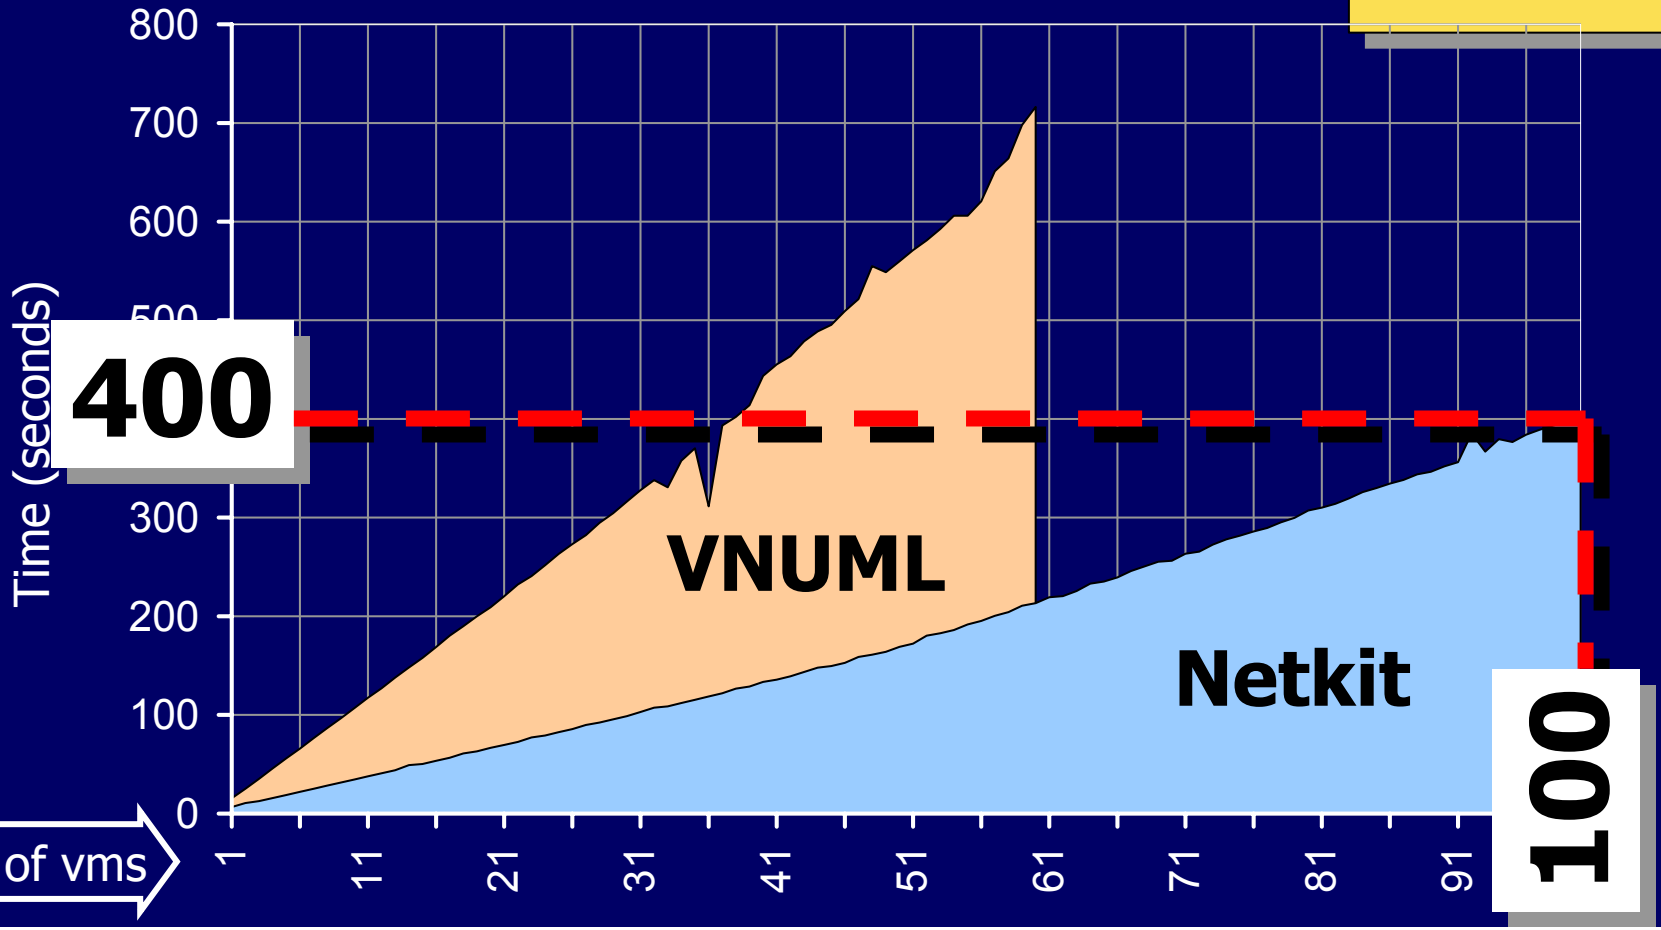

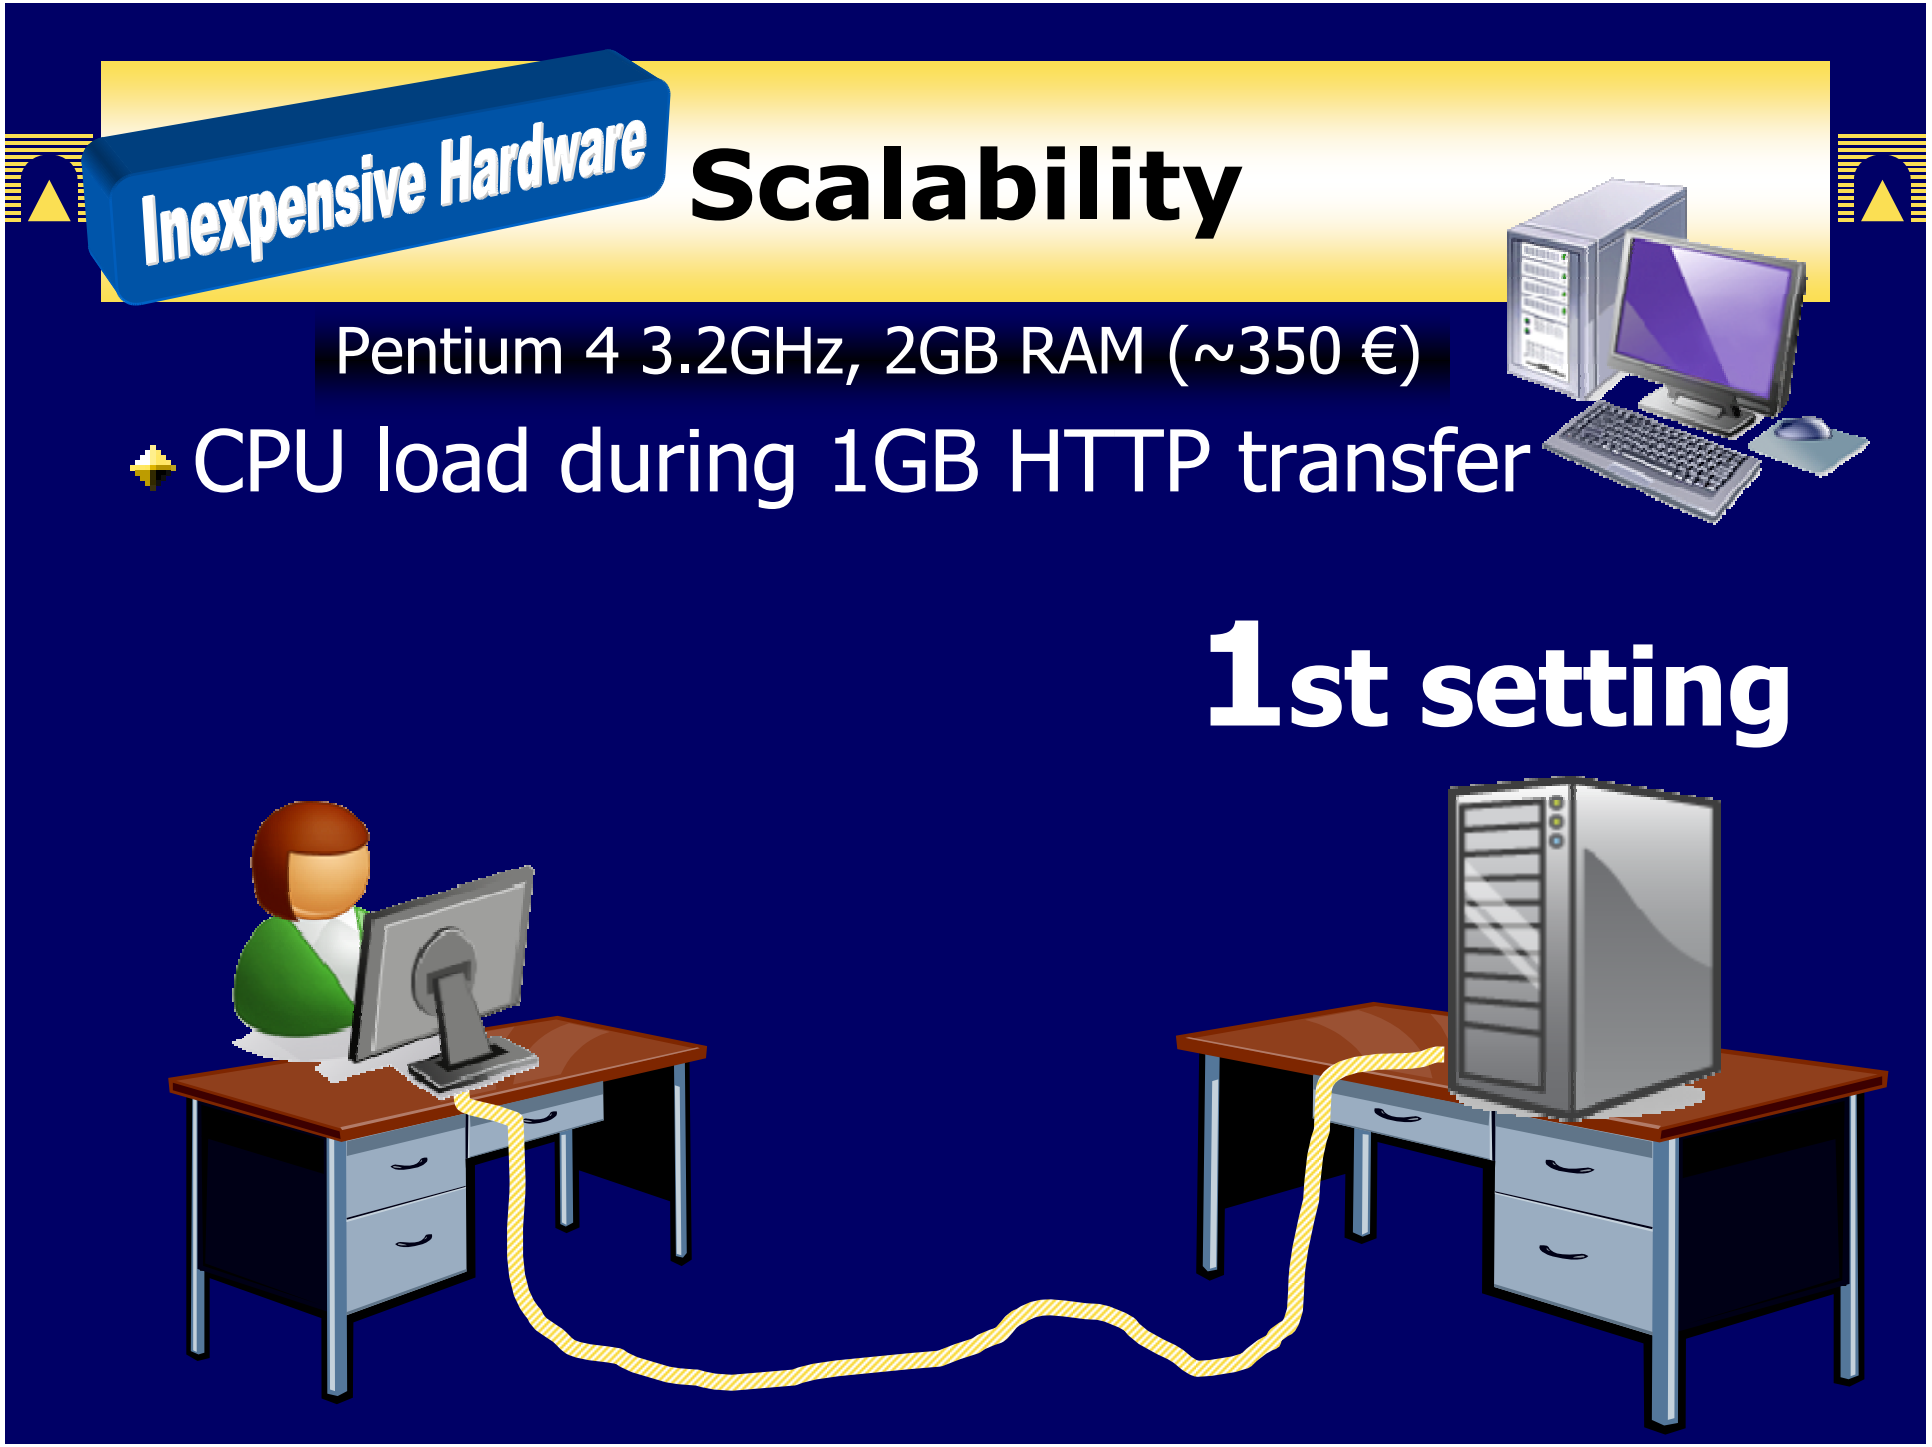
Inexpensive Hardware
Scalability
Pentium 4 3.2GHz, 2GB RAM (~350 €)
CPU load during 1GB HTTP transfer
1st setting

Scalability

Pentium 4 3.2GHz, 2GB RAM (~350 €)

- CPU load during 1GB HTTP transfer

2nd setting

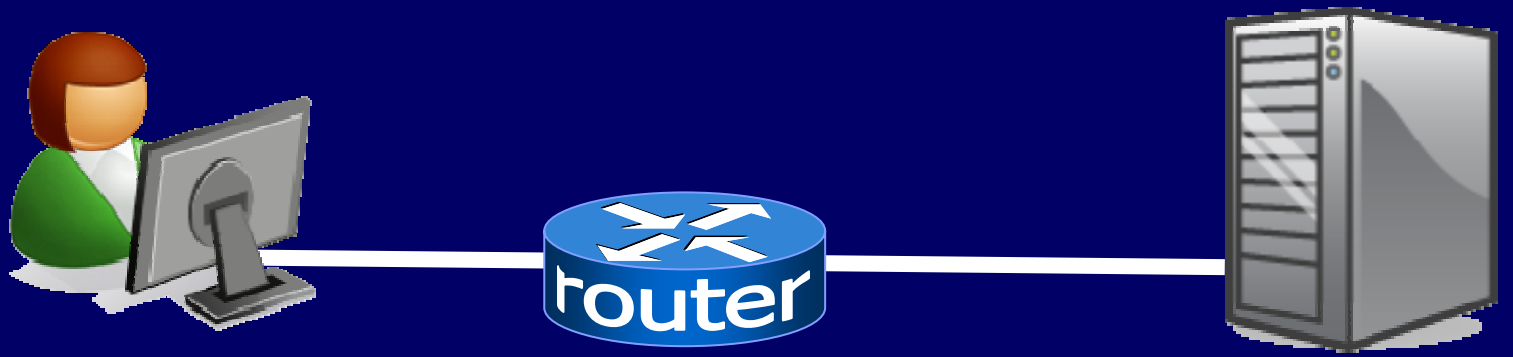

Inexpensive Hardware

Scalability

Pentium 4 3.2GHz, 2GB RAM (~350 €)

- CPU load during 1GB HTTP transfer

3rd setting

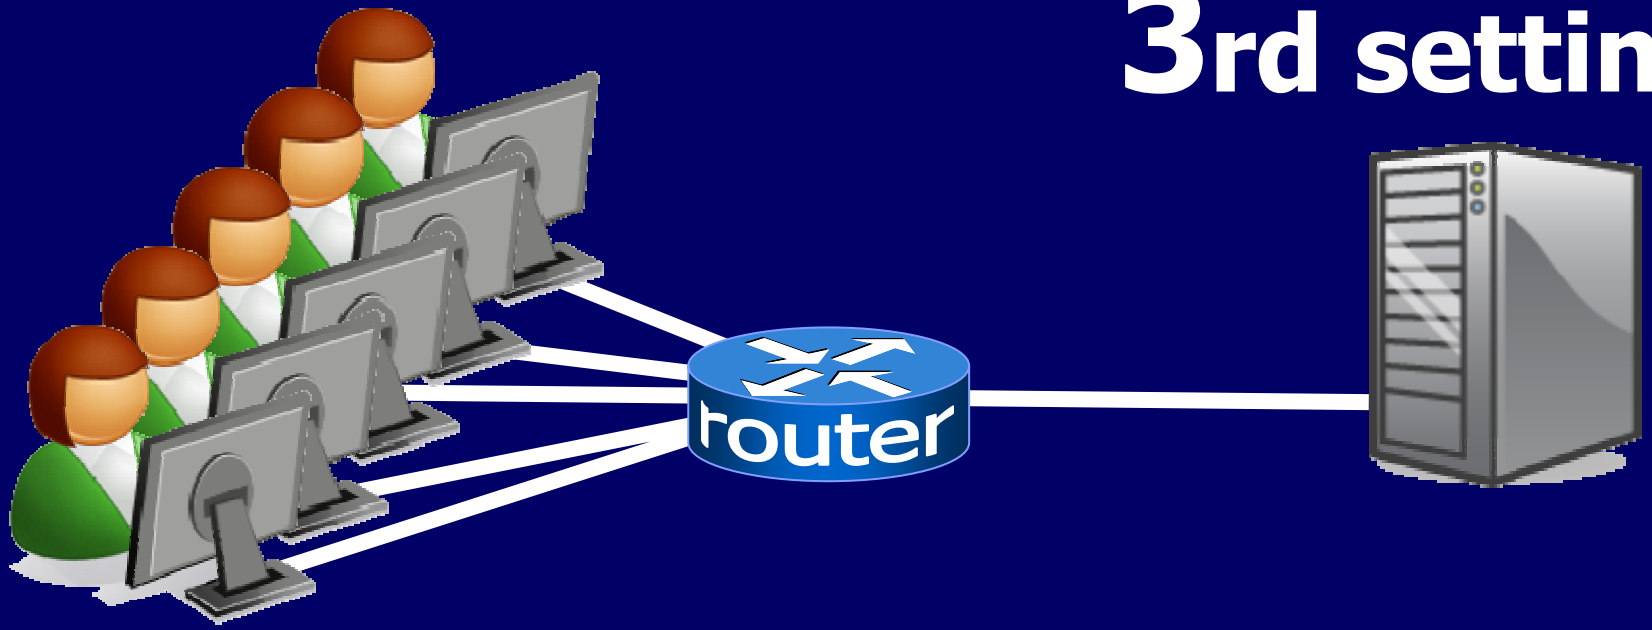

Inexpensive Hardware

Scalability

Pentium 4 3.2GHz, 2GB RAM (~350 €)

- CPU load during 1GB HTTP transfer

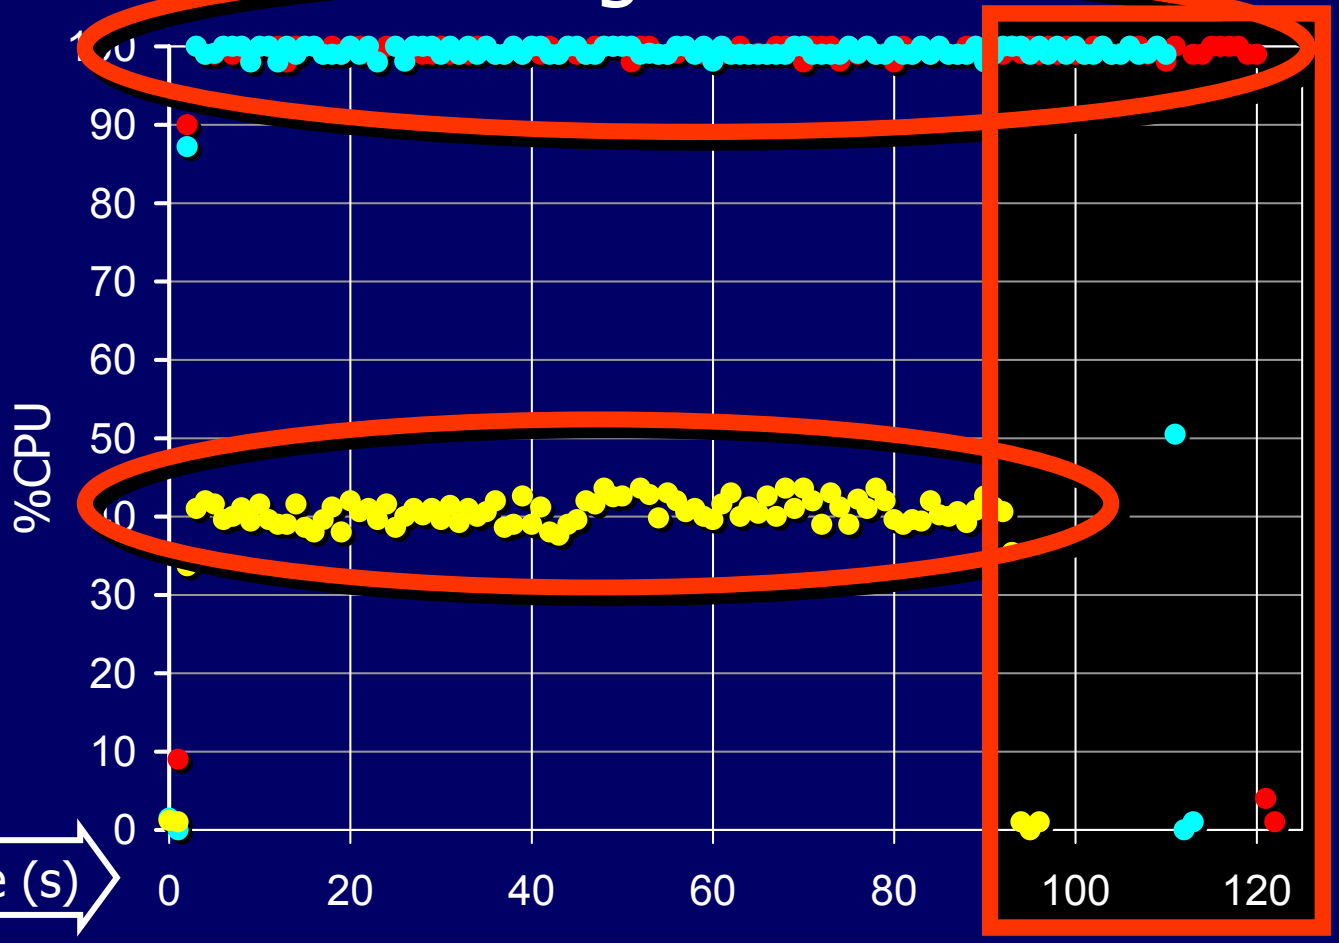

x1

Inexpensive Hardware Scalability

Pentium 4 3.2GHz, 2GB RAM (~350 €)

- CPU load during 1GB HTTP transfer

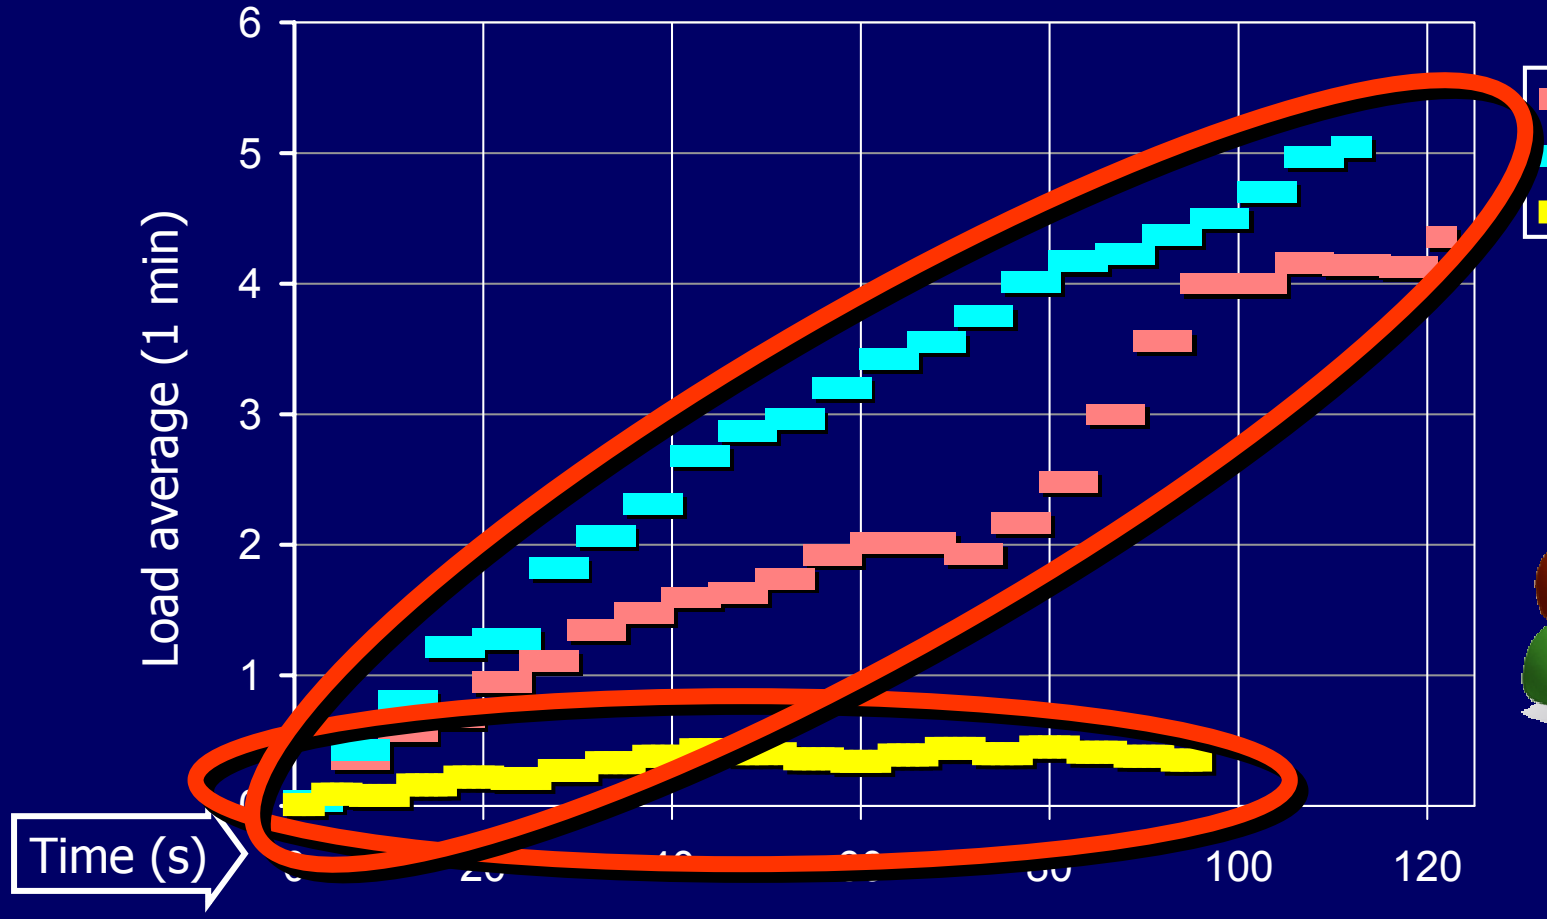

x1

Inexpensive Hardware

Scalability

Pentium 4 3.2GHz, 2GB RAM (~350 €)

- CPU load during 1GB HTTP transfer

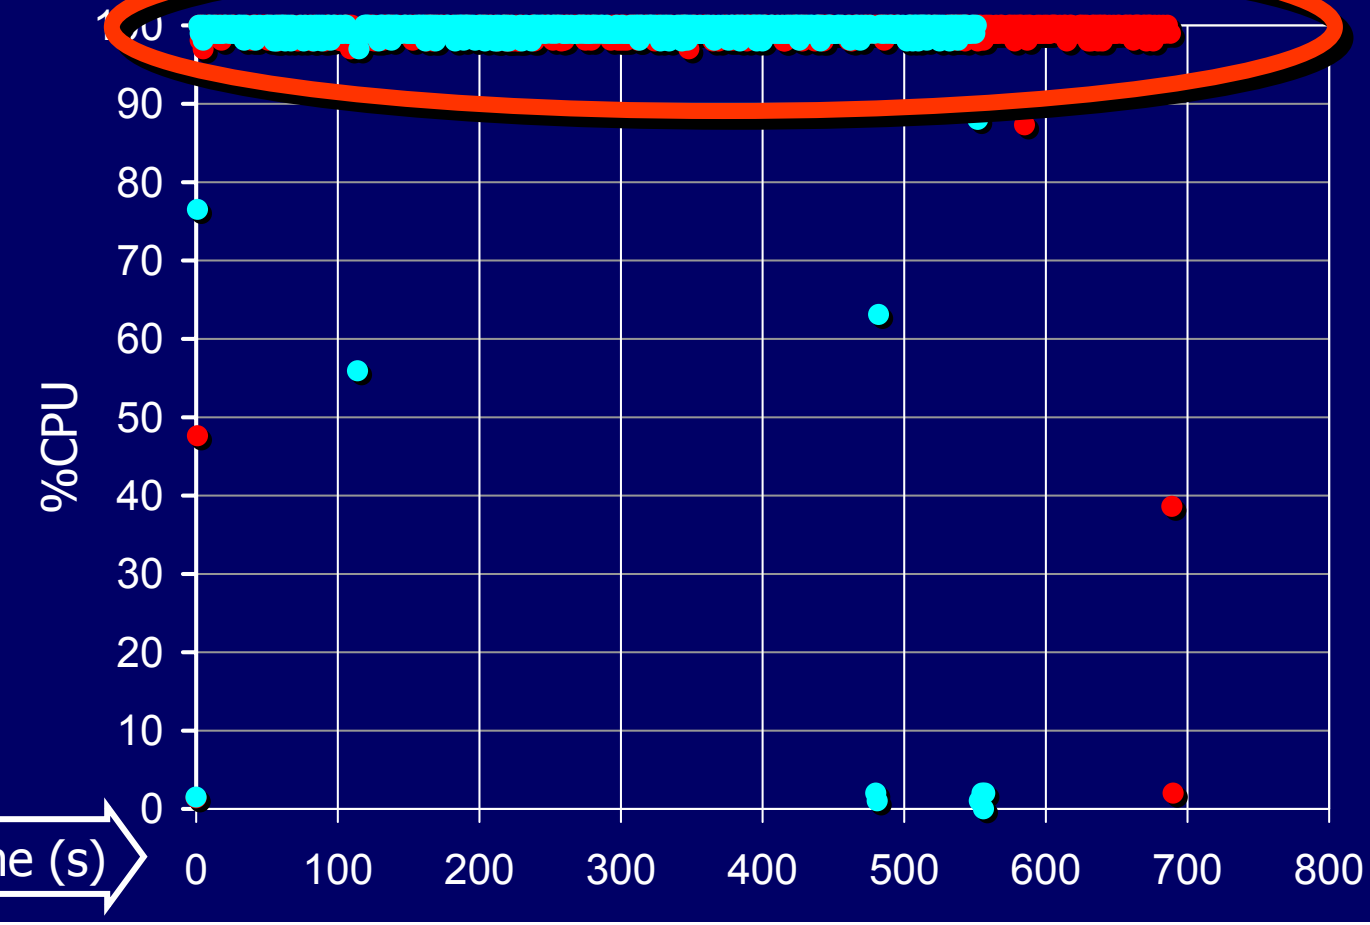

x5

Inexpensive Hardware

Scalability

Pentium 4 3.2GHz, 2GB RAM (~350 €)

- CPU load during 1GB HTTP transfer

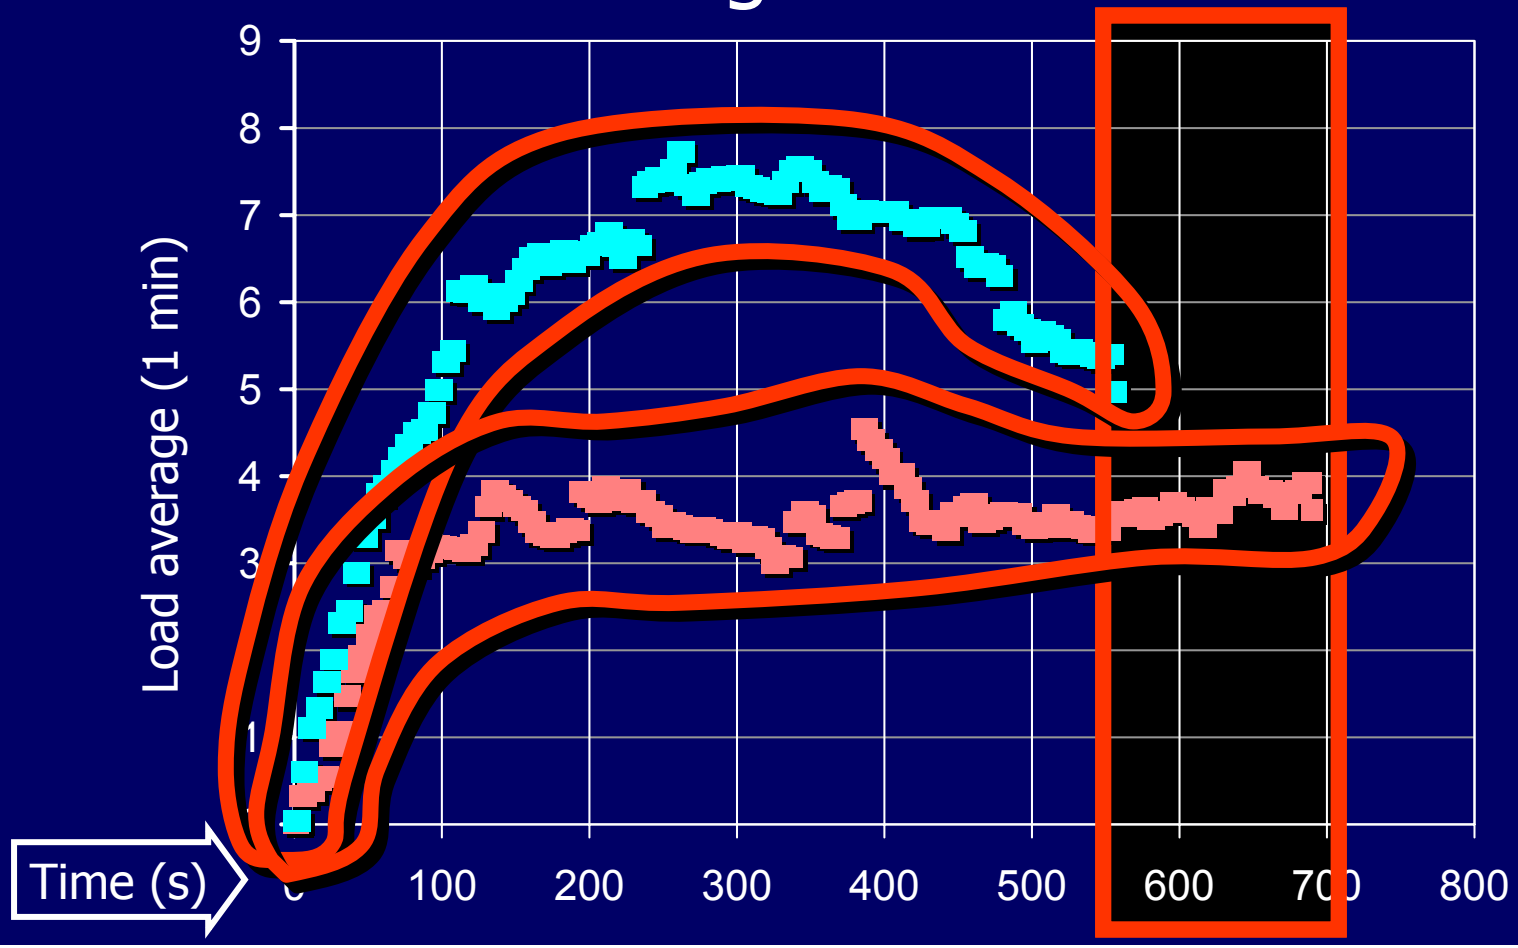

x5

So What?

- Take home
  - Lightweight
  - Easy
  - Turn key (labs)
  - Applications
  - Good scalability

- What next?
  - UI improvements (e.g., test procedure)
  - More labs (e.g., ISP topologies)
  - Better scalability by distribution (VDE?)

Where to Go Next

http://www.netkit.org/

- Releases & Documentation
  - Live CD
- Ready to use Labs
- Mailing list (138 subscribers) & FAQ
- Publications

Thanks to: You, Giuseppe Di Battista, Maurizio Patrignani, Stefano Pettini, …