



Dean Pemberton – NSRC  
PacNOG17

# Introduction to Virtualisation



These materials are licensed under the Creative Commons *Attribution-Noncommercial 3.0 Unported* license (<http://creativecommons.org/licenses/by-nc/3.0/>) as part of the ICANN, ISOC and NSRC Registry Operations Curriculum.

# Virtualization Technologies

**They are everywhere...**

By abstracting functionality away from hardware we gain

- Improved capacity usage
- Lower power consumption
- Reduced system administration overhead
- Better reliability (uptime, data loss)
- Possibilities that we are still thinking of...

# Two kinds of virtualization

## Consolidation


Run many services and servers onto fewer physical machines: increases ***efficiency***

## Aggregation

Distribute applications and resources across as many virtual servers as required, turning resources on or off as need: increases ***scalability***

# Some virtualization benefits

**Reduced power** use and better use of resources through **consolidation**





# More benefits

## Standardized platform

- Heterogeneous hardware platform hidden away behind virtualization
- Makes it easy to move hosts between platforms
- Not tied to a particular vendor – migrations are easy

**Open Source** offerings on par with commercial solutions and preferred by the “big boys”



# **AfNOG technical capacity trainings**

Instructors, staff, students and the institution have clearly benefitted...

# AfNOG Trainings 2000-2010...





# ...lots of setup



# Logistical benefits are obvious

## In the context of regional and local training:

- “Virtualizing” = less hardware (better use)
- Reduced shipping costs
- Reduced Logistics
  - Customs / import
    - Network equipment is often considered to be “telecommunications” – taxation issues, licensing
  - Small footprint – fits in a backpack or carry-on
  - Peripheral infrastructure (access points, desktop switches) are very small

# AfNOG 2013 we've virtualized

## Approximate numbers

- 140 PCs
- 40 routers
- 30 switches
- 3-5 full-sized tower PCs
- Keyboards, monitors, mice and network cabling down to...



# Virtual AfNOG 2013

- This represents ***significant overkill*** (2x or more).
- Much more could be virtualized...



# Motivations and benefits

Other benefits than logistics are well aligned with the needs of regional / decentralized training:

- Adaptability
- Educational

We will cover these in the next slides




# Benefits: Adaptability

- Single architecture multiple workshops
- Architecture and platform uncoupled
  - A **Mac** running **Linux** running **FreeBSD**
  - A **PC** running **Linux** running **Windows**
- ***Fast*** reconfiguration!
  - Can be done in minutes with templates.

## Let's see this now!

(Next slide for images)

# NMM Live in the room



# A smorgasbord of choices!

## Full virtualization

- KVM (Linux and Solaris only)
- Parallels (Mac OS X only)
- QEmu
- VirtualBox (Windows, Linux, Mac, FreeBSD)
- Virtual PC (Windows only)
- VMware (Workstation/Fusion, ESX)
- Windows Hyper-V
- Xen

## Lightweight/pseudo

- FreeBSD / Linux Jails/LXC/OpenVZ



## Network Simulation

- Marionnet
- Navy CORE

## Network Emulation

- Dynamips /  
Dynagen /  
GNS3
- Olive (Juniper)
- Cisco IOU (private)

# Some virtualization frameworks

## Manage/deploy virtualization in a controlled fashion

- VMWare vCenter (commercial – enterprise)
- Libvirt (for managing KVM, Xen, VirtualBox)
- Ganeti, Synnefo – clustering, small to med. size
- OpenStack, Eucalyptus – large scale (cloud provider)

# What do we use?

## Hardware

- MacMini Server, 16GB RAM, 2x256 SSD, i7 quad core
- (Optional) fanless, Gigabit, managed 8-port switch
- Ubiquiti UniFi Pro AP or Netgear AP

## Software

- Ubuntu Linux 12.04 LTS 64-bit
- KVM (Kernel based Virtual Machine)
- dynamips/dynagen
- Ansible




# What do we use?

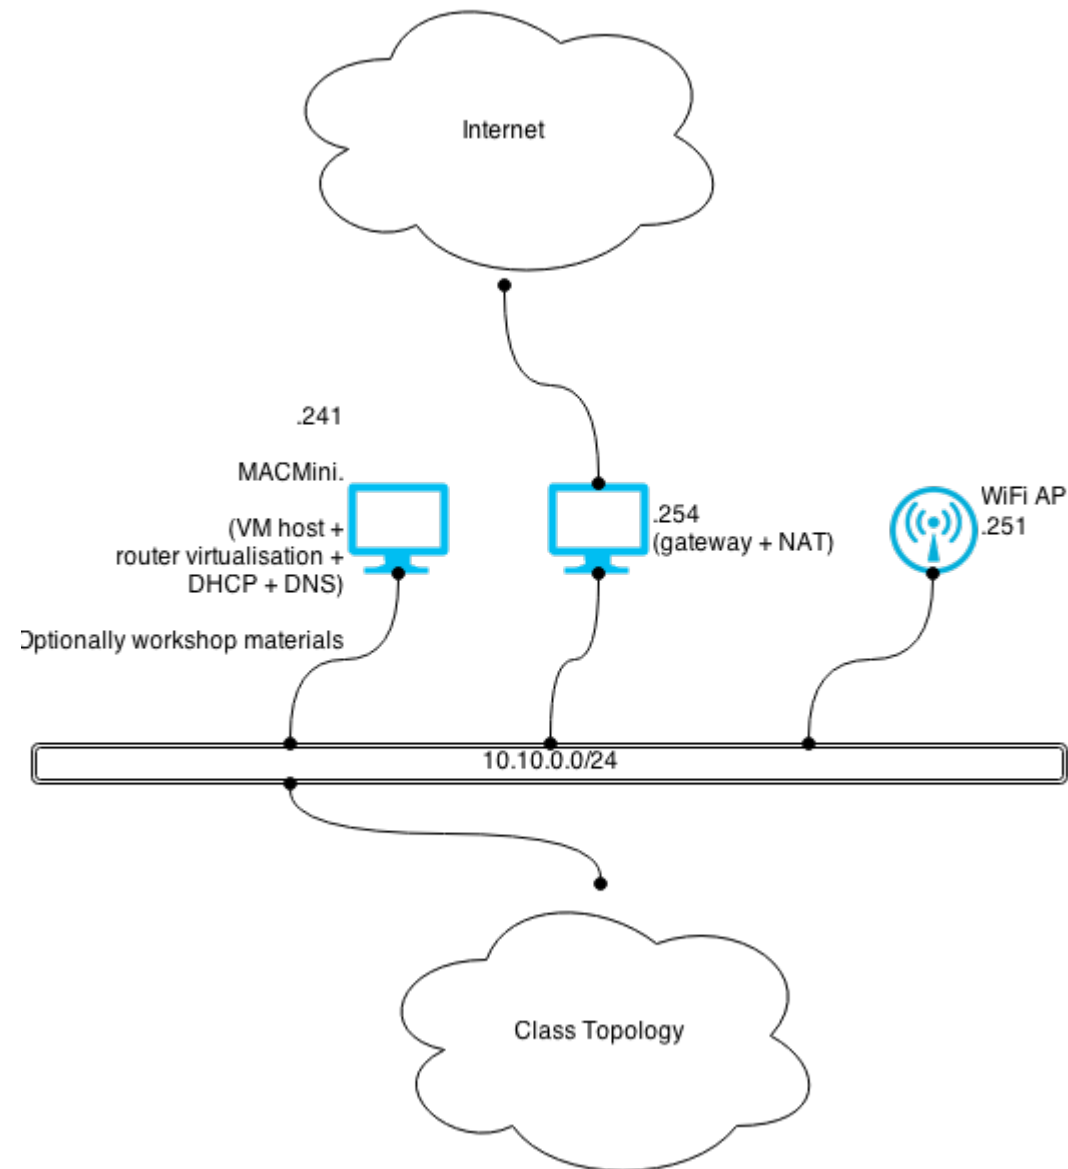
## Hardware

- We use the MacMini for convenience (power / size ratio = very good)
- Can use any reasonably modern machine with virtualization extensions (VT-x/VT-d) in the CPU
- Tower PC with Core i7, 7200 RPM disks
- Rackmount server, Xeon/Opteron, faster disks
- SSD disks are very nice, but not critical
- RAM is a big factor

# Classroom Layout (physical)



# Classroom layout (logical pieces)





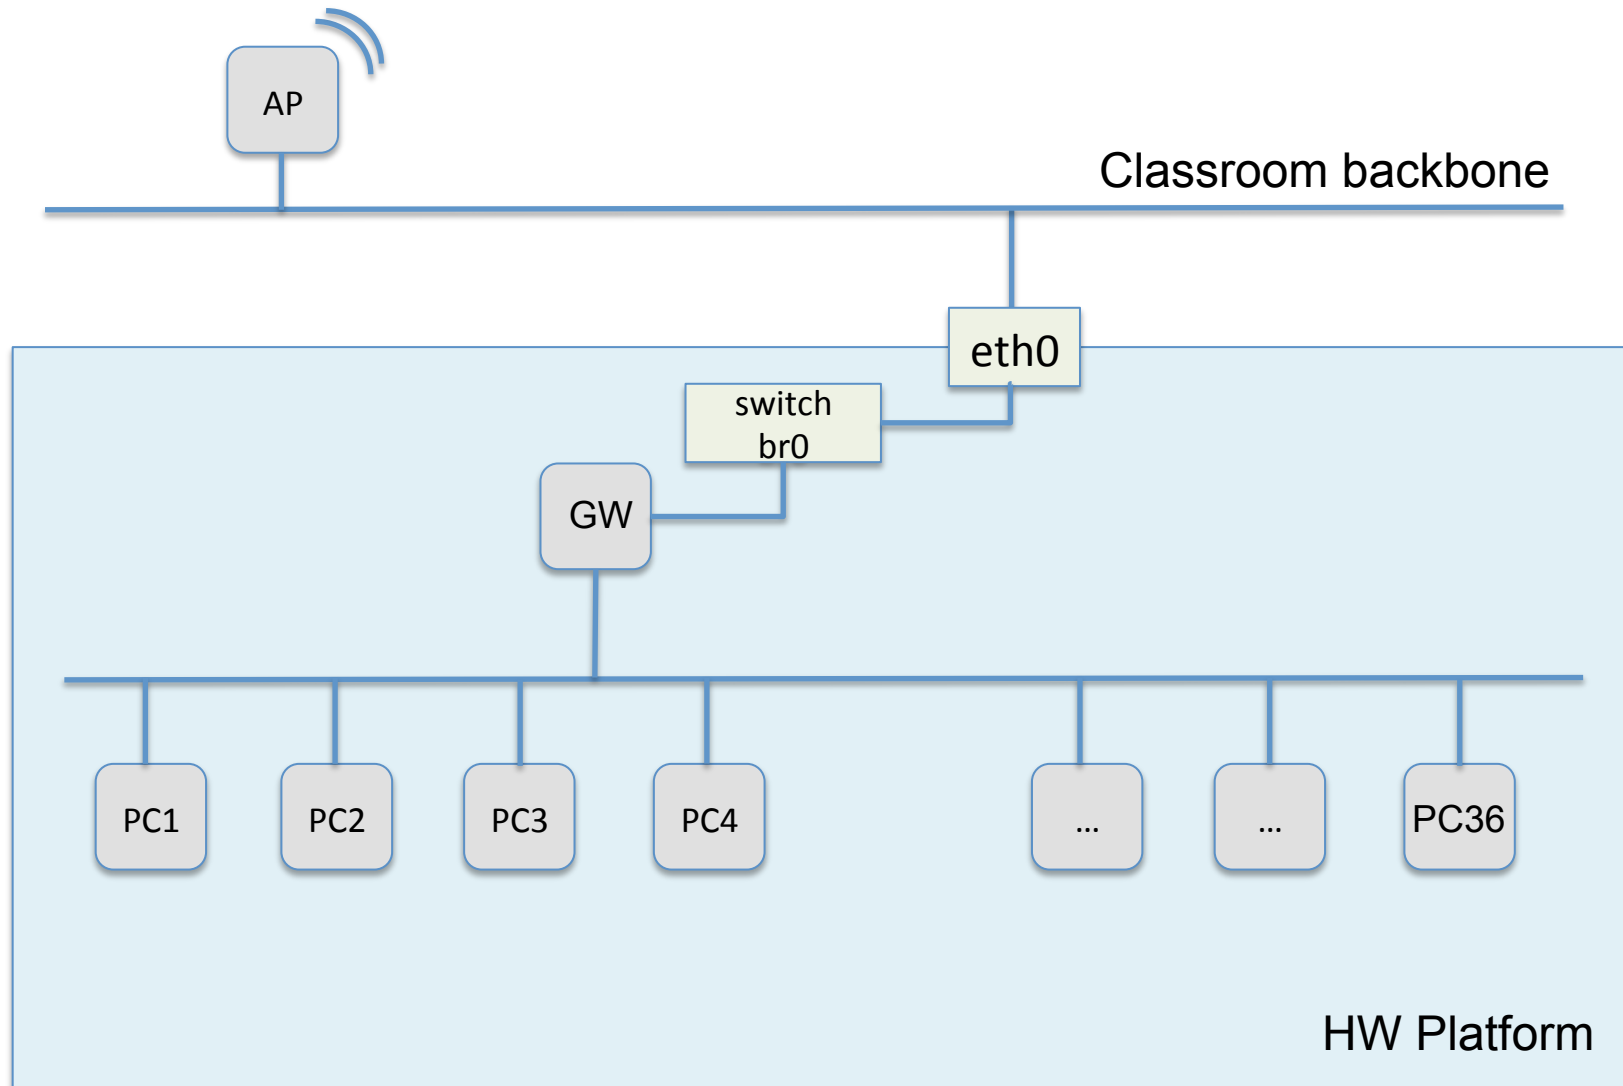
# Concerning materials

- The class materials need to be hosted somewhere so students can download them.
- NSRC has been using wikis – and we use trac for it.
- The optimal place should be linked to from wherever the website describing the training is.
- In either case, a local copy of materials is sometimes useful to save bandwidth.

# Some classroom virtual environments

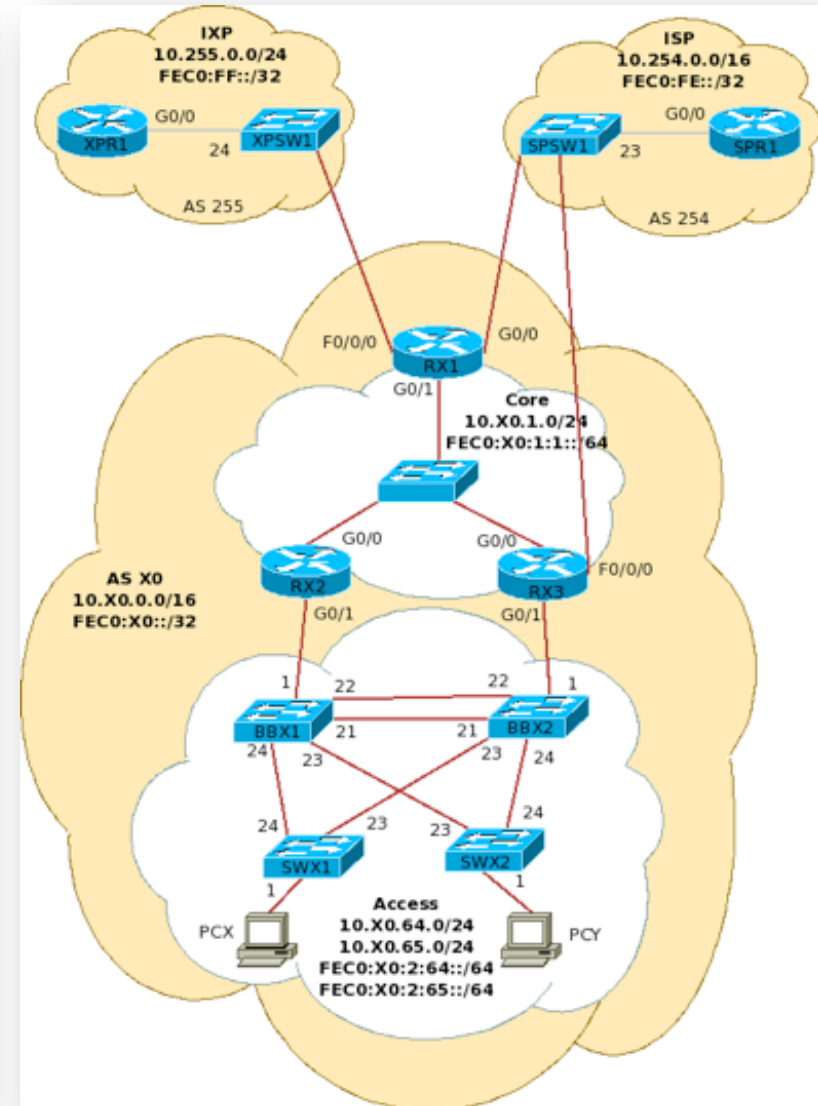
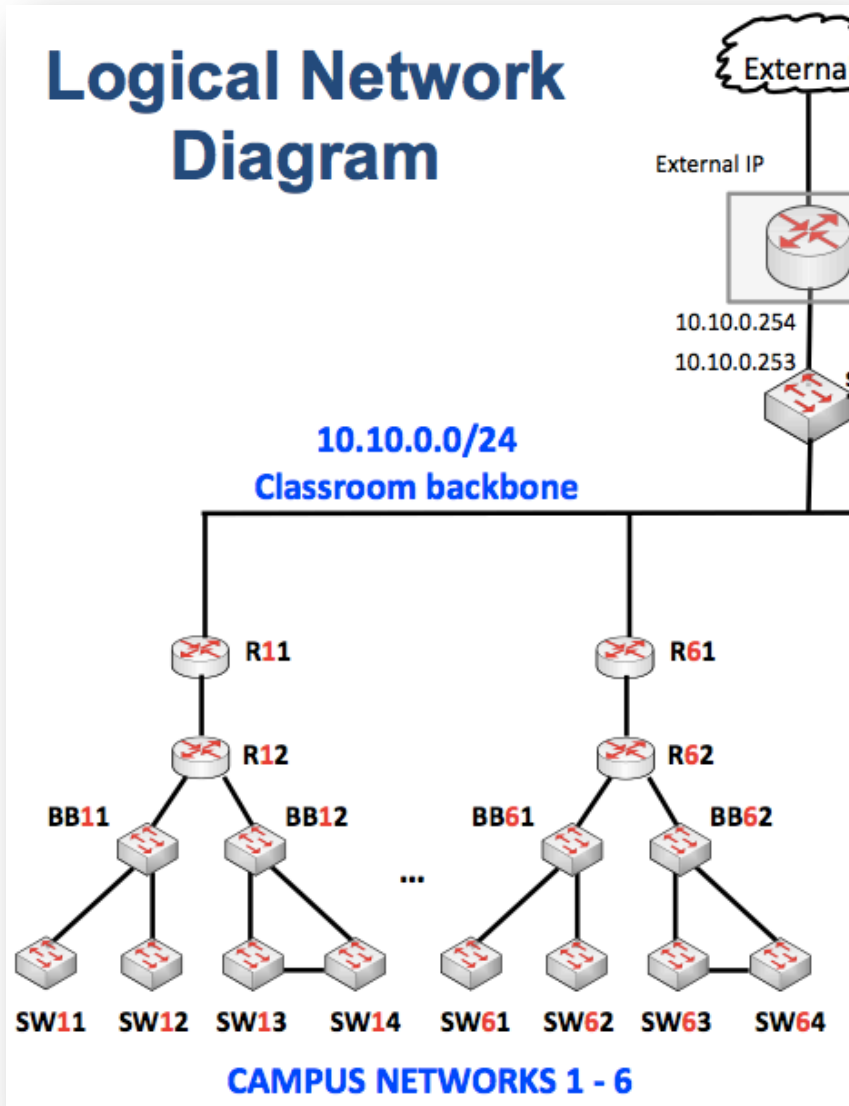
- Flat, simple network (UNIX / Linux intro)
- Campus Network Design (L2 and L3)
- Network Monitoring and Management
- Hybrid of virtual networks
- DNS/DNSSEC workshop

# UNIX / Linux Introduction



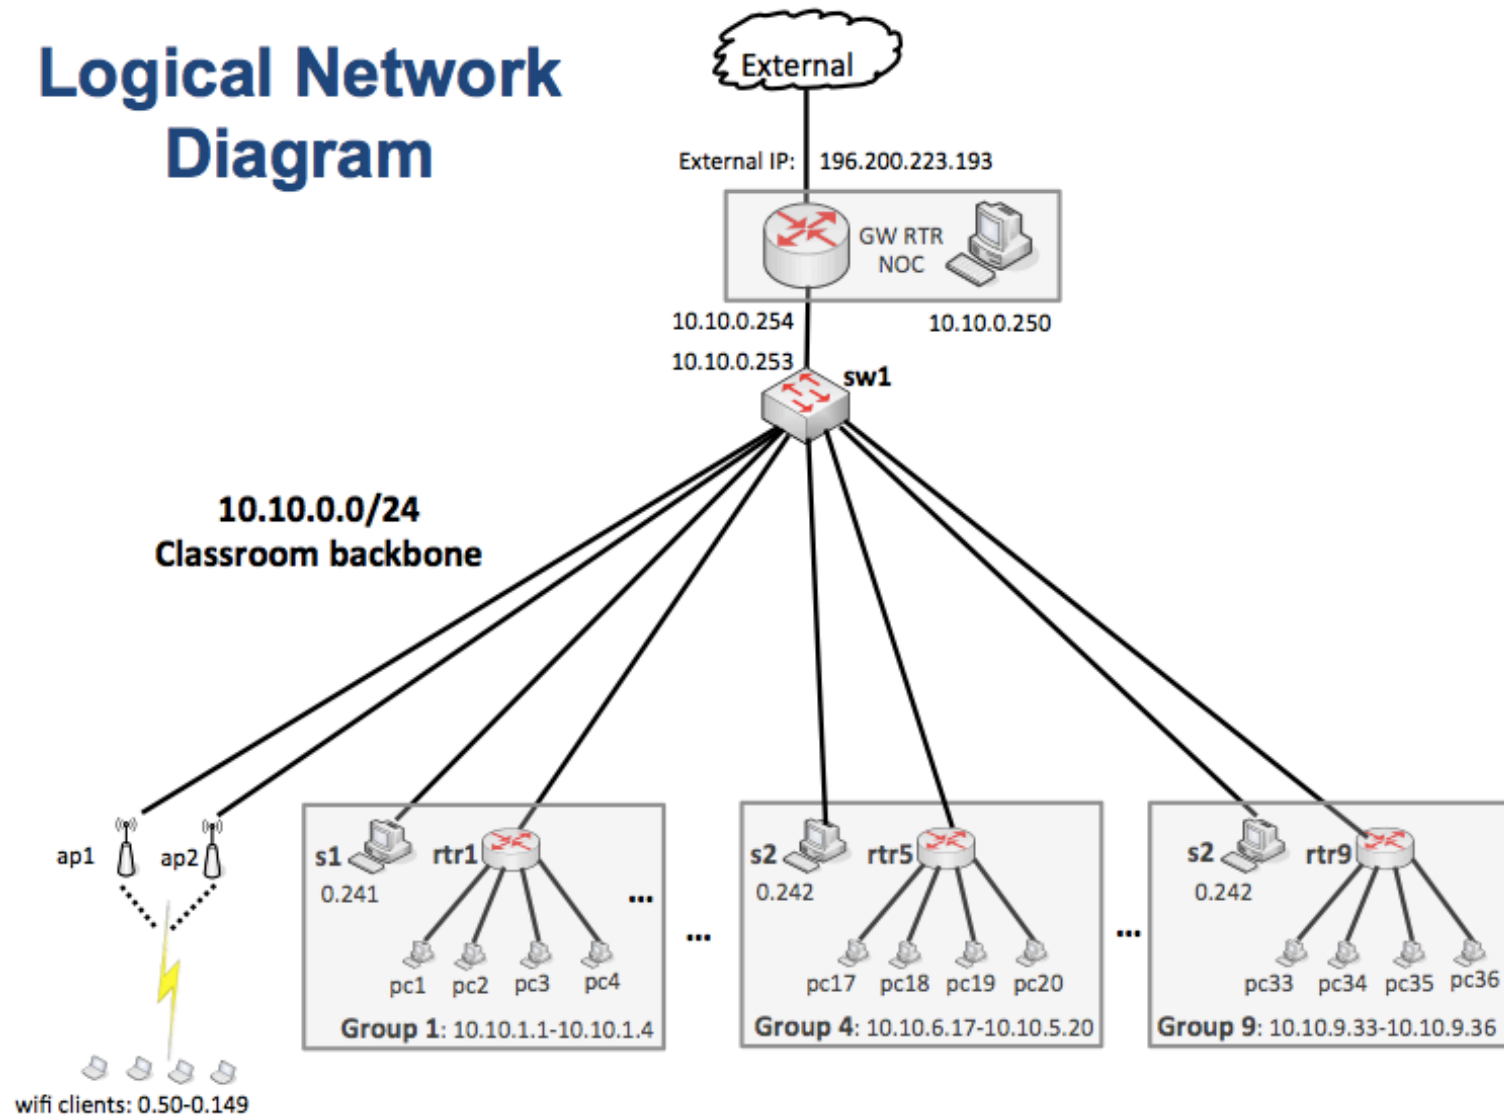
36+ virtual machines on one server

# Campus Network Design (CND)




# Network Management (NMM)

## Logical Network Diagram



# CND and NMM over 2-3 Machines

## Logical Network Diagram

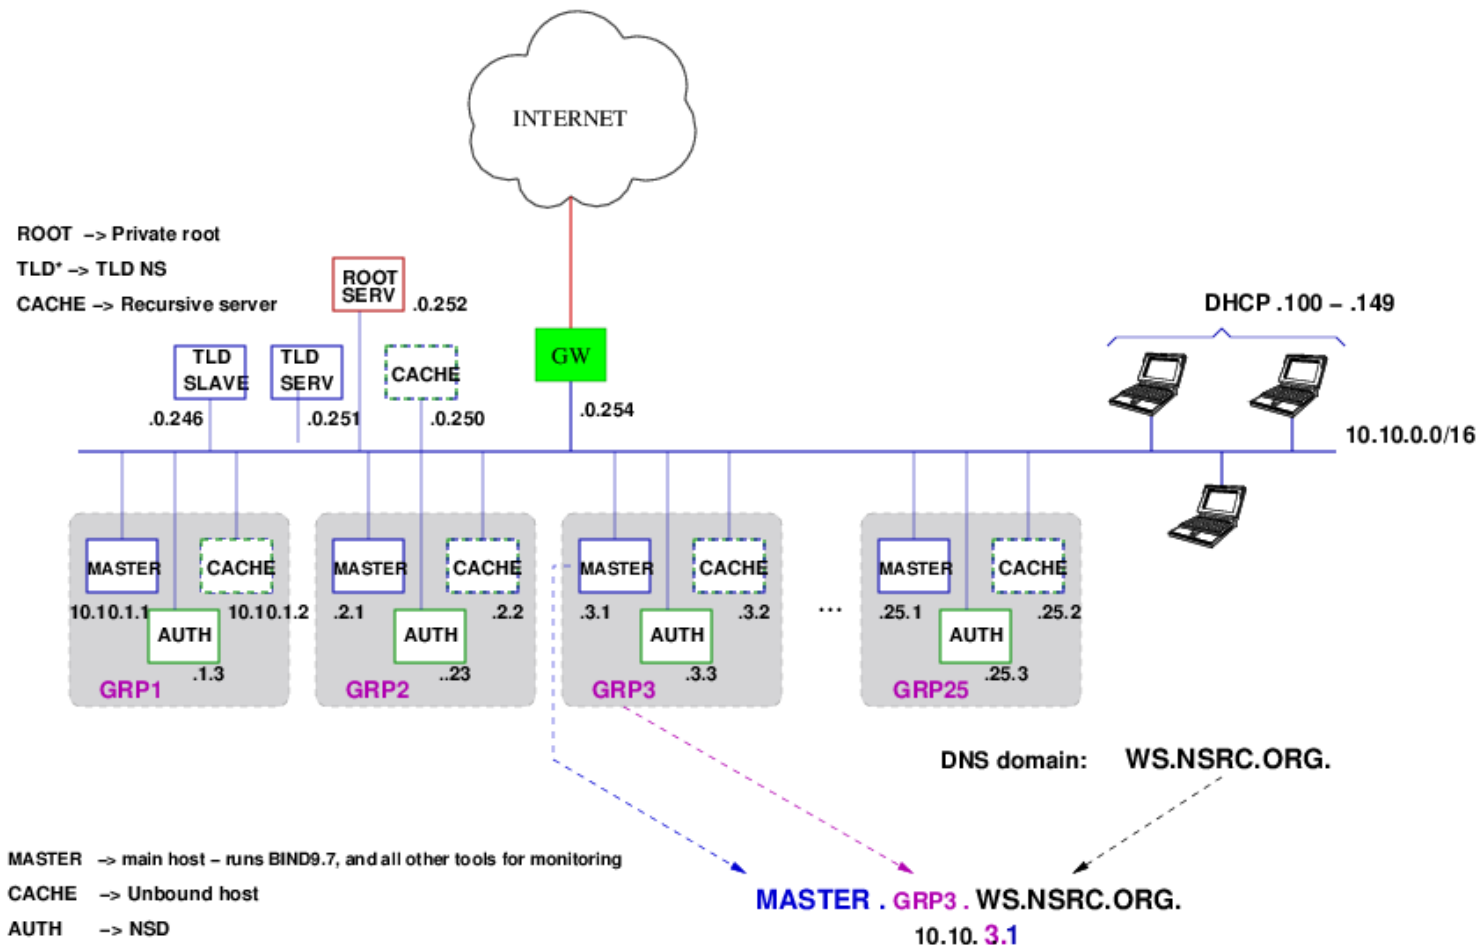


# DNS/DNSSEC workshop

## NETWORK LAYOUT

login: adm  
pass: \*given in class\*

WiFi SSID: DNS or DNS2  
WiFi pass: 8888888888



# Benefits: Educational

- Shipping small CPUs much less expensive.
- Other benefits not tied to simple cost benefit or logistics
- Virtualization technologies part of modern IT infrastructure
  - Not just for training
    - Virtualized OS (“hypervisors”)
    - Virtualized network (VLANs, virtual switches, routers, SDN)
    - Virtualized storage (iSCSI, disk images)



## **Benefits: Educational (con't)**

- Clear benefits for institutions offering applications and services to staff and faculty
- Professionally relevant for students and instructors
- Participants can easily re-create lab environments on laptops
- Much simpler to provide network and systems training.

# Limitations...

**There are, of course, some tradeoffs:**

Hands-on is limited

- No manipulation of “real hardware”
  - Some people grasp concepts better
  - Cables vs VLANs
  - Reality for present-day networks and systems
- Not always possible to virtualize all hardware

What tool(s) to use?

## **... and new possibilities**

**All this simplifies capacity building...**


- Lower cost, better outreach
- Increased adaptability
- More can be done
- Easy to adjust to new topics / themes


# Ganeti terminology

Cluster - group of nodes

Node - physical host

Instance - virtual machine, aka guest





# Primary node failure



```
# gnt-node list
```

Node		DTotal	DFree	MTotal	MNode	MFree	Pinst	Sinst
nsrc1.nsrc.org	465.5G	337.4G	7.7G	4.2G	5.9G	4	0	
nsrc2.nsrc.org	465.5G	392.8G	3.8G	1.4G	2.9G	2	1	
nsrc3.nsrc.org	465.5G	369.6G	3.8G	2.0G	2.8G	1	1	
nsrc4.nsrc.org	465.5G	375.6G	3.8G	2.2G	2.3G	1	1	
nsrc5.nsrc.org	465.5G	339.6G	3.8G	2.2G	2.3G	1	1	
nsrc6.nsrc.org	465.5G	365.5G	3.8G	489M	3.4G	0	3	

```
# gnt-node list-storage
```

Node	Type	Name	Size	Used	Free
------	------	------	------	------	------

Allocatable

nsr6.nsrc.org	lvm-vg	ganeti	465.5G	100.0G	365.5G	Y
---------------	--------	--------	--------	--------	--------	---

nsr5.nsrc.org	lvm-vg	ganeti	465.5G	125.9G	339.6G	Y
---------------	--------	--------	--------	--------	--------	---

nsr4.nsrc.org	lvm-vg	ganeti	465.5G	89.9G	375.6G	Y
---------------	--------	--------	--------	-------	--------	---

nsr3.nsrc.org	lvm-vg	ganeti	465.5G	95.9G	369.6G	Y
---------------	--------	--------	--------	-------	--------	---

nsr2.nsrc.org	lvm-vg	ganeti	465.5G	72.8G	392.8G	Y
---------------	--------	--------	--------	-------	--------	---

nsr1.nsrc.org	lvm-vg	ganeti	465.5G	128.1G	337.4G	Y
---------------	--------	--------	--------	--------	--------	---



```
# gnt-instance list -o +snodes
```

Instance	Hypervisor	OS	Primary_node	Status	Memory	Secondary_Nodes
Example1.nsrc.org	kvm	noop	nsrc3.nsrc.org	running	512M	nsrc5.nsrc.org
Example2.nsrc.org	kvm	noop	nsrc1.nsrc.org	ADMIN_down	-	nsrc6.nsrc.org
Example3.nsrc.org	kvm	snf-image+ubuntu	nsrc1.nsrc.org	running	1.0G	nsrc3.nsrc.org
Example4.nsrc.org	kvm	snf-image+ubuntu	nsrc2.nsrc.org	ADMIN_down	-	
Example5.nsrc.org	kvm	snf-image+ubuntu	nsrc5.nsrc.org	running	1.0G	nsrc6.nsrc.org
Example6.nsrc.org	kvm	snf-image+ubuntu	nsrc4.nsrc.org	running	1.0G	nsrc2.nsrc.org
Example7.nsrc.org	kvm	noop	nsrc2.nsrc.org	running	128M	
Example8.nsrc.org	kvm	noop	nsrc1.nsrc.org	ADMIN_down	-	nsrc4.nsrc.org
Example9.nsrc.org	kvm	noop	nsrc1.nsrc.org	running	256M	nsrc6.nsrc.org

```
# hbal -C -m nsrccgnt
```

```
...
```

```
Loaded 6 nodes, 9 instances
```

```
Group size 6 nodes, 9 instances
```

```
Cluster score improved from 5.13458892 to 2.26351813
```

```
Solution length=2
```

Commands to run to reach the above solution:

```
echo jobset 1, 1 jobs
```

```
echo job 1/1
```

```
gnt-instance failover -f example1.nsrc.org
```

```
gnt-instance replace-disks -n nsrc4.nsrc.org example1.nsrc.org
```

```
echo jobset 2, 1 jobs
```

```
echo job 2/1
```

```
gnt-instance replace-disks -n nsrc3.nsrc.org example2.nsrc.org
```

```
gnt-instance migrate -f example2.nsrc.org
```

# Questions?