

MSU Linux Cluster



**R. Moore,
Michigan State
University**

Security:

Unclustered machines not always up to date

Many different setups increases chance of a security hole

System administration

It s a lot easier to set up and run a cluster than 10 individual machines

Coordination

Access to centralized resources to avoid needless duplication

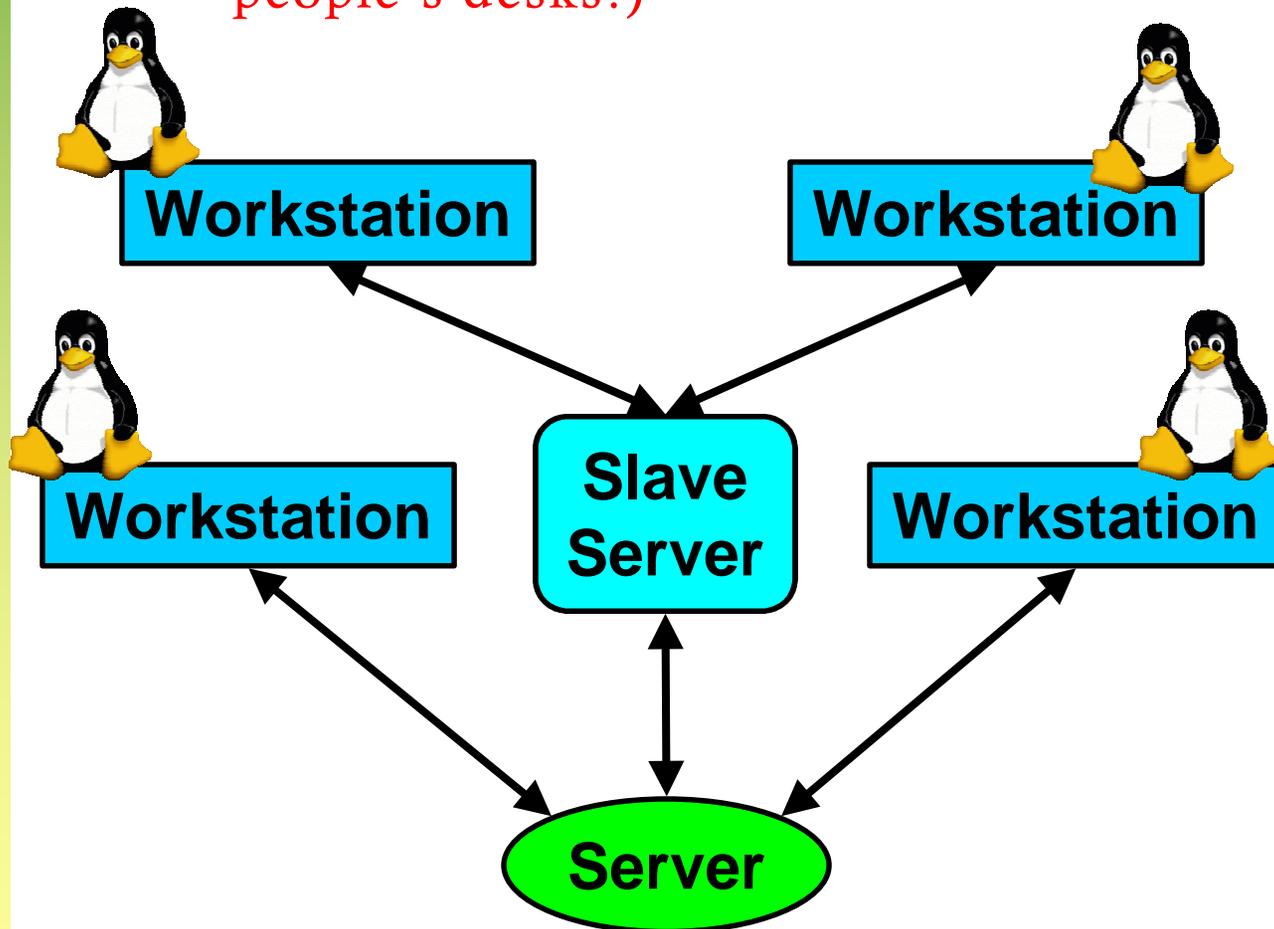
Backups, disks, batch queues etc.

Workstations

Sit on people's desks

Server(s)

Provide cluster services (and sits on people's desks!)



Home directory space on a server

Accessible from all cluster machines

Batch queues

Harness the power of the entire cluster to run your jobs

Backups

Centralized backup service available for any directory on any machine to any tape drive

Access to D0 software without having to install it yourself!

Must leave your machine running
Linux 24 hours a day:

No rebooting into Windows NT
(but see later)

Must not run externally visible
network services:

Your account is fully accessible
through the server

Your spare CPU cycles:

Used by batch jobs

You will get to use other people's
spare CPU cycles in return though!

All machines see the same
apparent file system:

Home directories available on any
machine

Utilities/software available in same
location everywhere

e.g. StarOffice, D0 software etc.

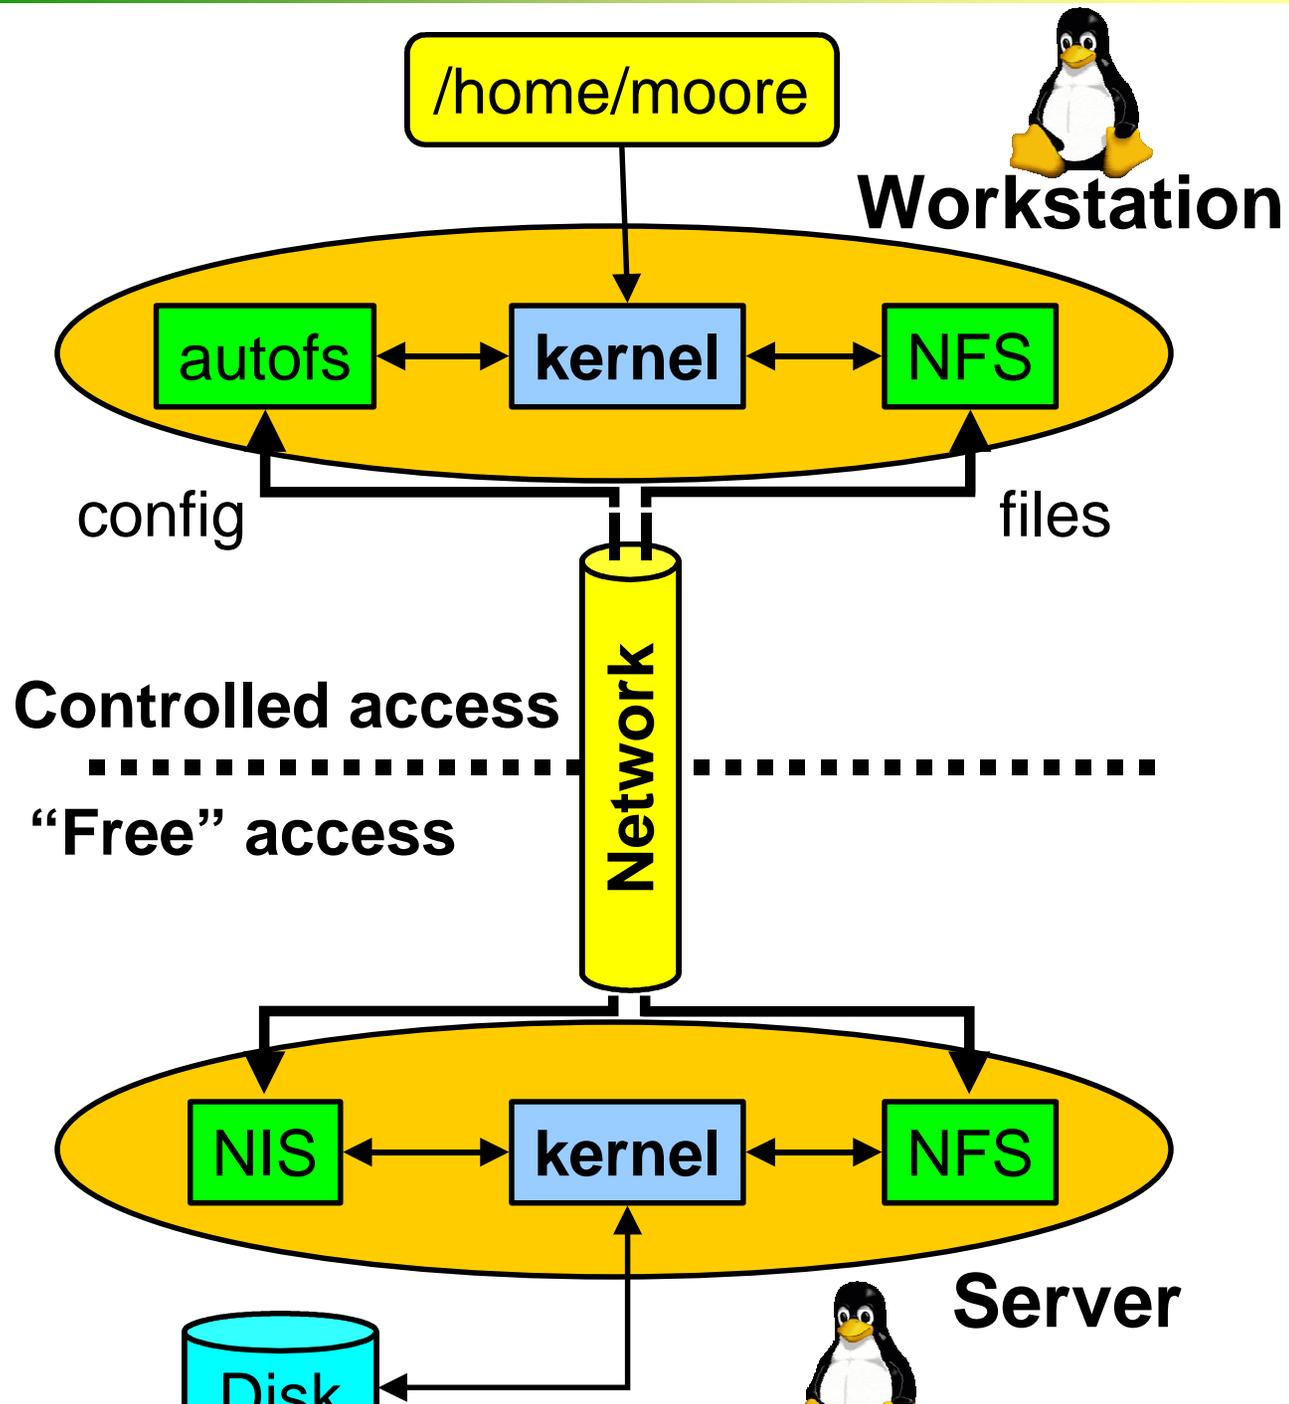
Achieved using autofs+NFS

Configured by server (NIS)

Tree-like mounting structure

Medium and long cluster-wide
batch queues available on all
machines

More complex queue configuration
possible



Configuration

**Two types of machine: server
and workstation**

Workstations (7)

200-750MHz Pentium II-III

80-256Mb memory

3-34Gb IDE/ATA disk

Server (1)

dual 600MHz Pentium III

512Mb memory

36Gb SCSI disks

DDS3 DAT drive

All run customized RedHat 6.1

Linux

**Minimal administration
required!**

Installation quick and easy

RedHat 6.1 customized CDROM

two install options available:

Workstation or Server

Supports >34GB IDE disks

**Configuration information
served using NIS and rsync**

Two levels of root access

Local root access (one PC)

**All system changes monitored by a
security program**

Cluster root access (all PCs)

Big concern!!

Workstations hidden

Only incoming “**intra-cluster**” connections allowed

Outgoing connections normal

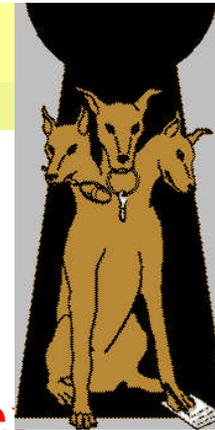
Automatic system updates

“**autorpm**” updates system with **RPMs** found on server

No unencrypted logins

Kerberos telnet/ftp + SSH

SSH web-page Java applet for when you are away at conferences etc.



Tripwire runs on all nodes

Detects changes in config files

Sends email to cluster sysadmins
once per day

Possible improvements:

Move Kerberos database to
inaccessible node

Have system logs sent to inaccessible
node & run swatch

Make a dedicated workstation act as
the cluster gateway

Hide workstations behind a hardware
firewall (CERN)



UPS database installed on server

Disk based database: no need to install on each workstation!

exported to workstations via autofs+NFS

new updates installed on server will “appear” on every machine in the cluster

See Dylan’s talk for more information on this...

Batch queuing system installed on all machines

central queues exist on server to which all jobs are submitted

Job sent to least loaded, highest performance machine

“fair” queuing system

when a new job starts running all that user’s jobs are moved to the back of the queue!

Maximum time limits enforced

receive email when job starts and/or completes (optional)



Adding more machines is simple

custom CD install, ~30 minutes

Desktops place minimal load on server and network

OS installed and maintained locally

only home directories or cluster

software use network directories

As cluster grows two options:

Split into separate clusters

Add slave servers to reduce load on main server

Eventually may need high power CPU/disk servers for RunII data analysis

Possible blueprint for project servers...

NT under LINUX!

VMware have written a “virtual machine emulator”

emulates a standard PC

install entire operating system on the virtual machine. Supports:

Windows 95/98/NT/2000

Linux(!)

BeOS and other oddities

**disks either emulated using a file
OR via direct access to unused disk
partition**

**Access to UNIX home directory
via SAMBA**

Used to write this talk!

Performance:

95-98% performance of host CPU

**~60% of host disk performance
using file mode disk emulation**

~50% of host graphics performance

**None are a problem for standard
applications like**

Office 2000

Adobe Acrobat

Internet Explorer etc.

Cluster works very well

no major problems to date
very little administration needed
(printing needs an occasional tweak)

Fully integrated Linux/NT environment

users can decide on Linux or NT
for their desktops without penalty

Easy to implement at other locations

Planning to use something similar for online L2 Alpha Linux machines