Thousands of Linux Installations (and only one administrator)

A Linux cluster client
for the University of Manchester

A V Le Blanc
I T Services
University of Manchester

LeBlanc@man.ac.uk

Overview

- Environment
- Description of the system
- How to look after it
- Why is it this way?
- Evaluation

Satisfying the Demand?

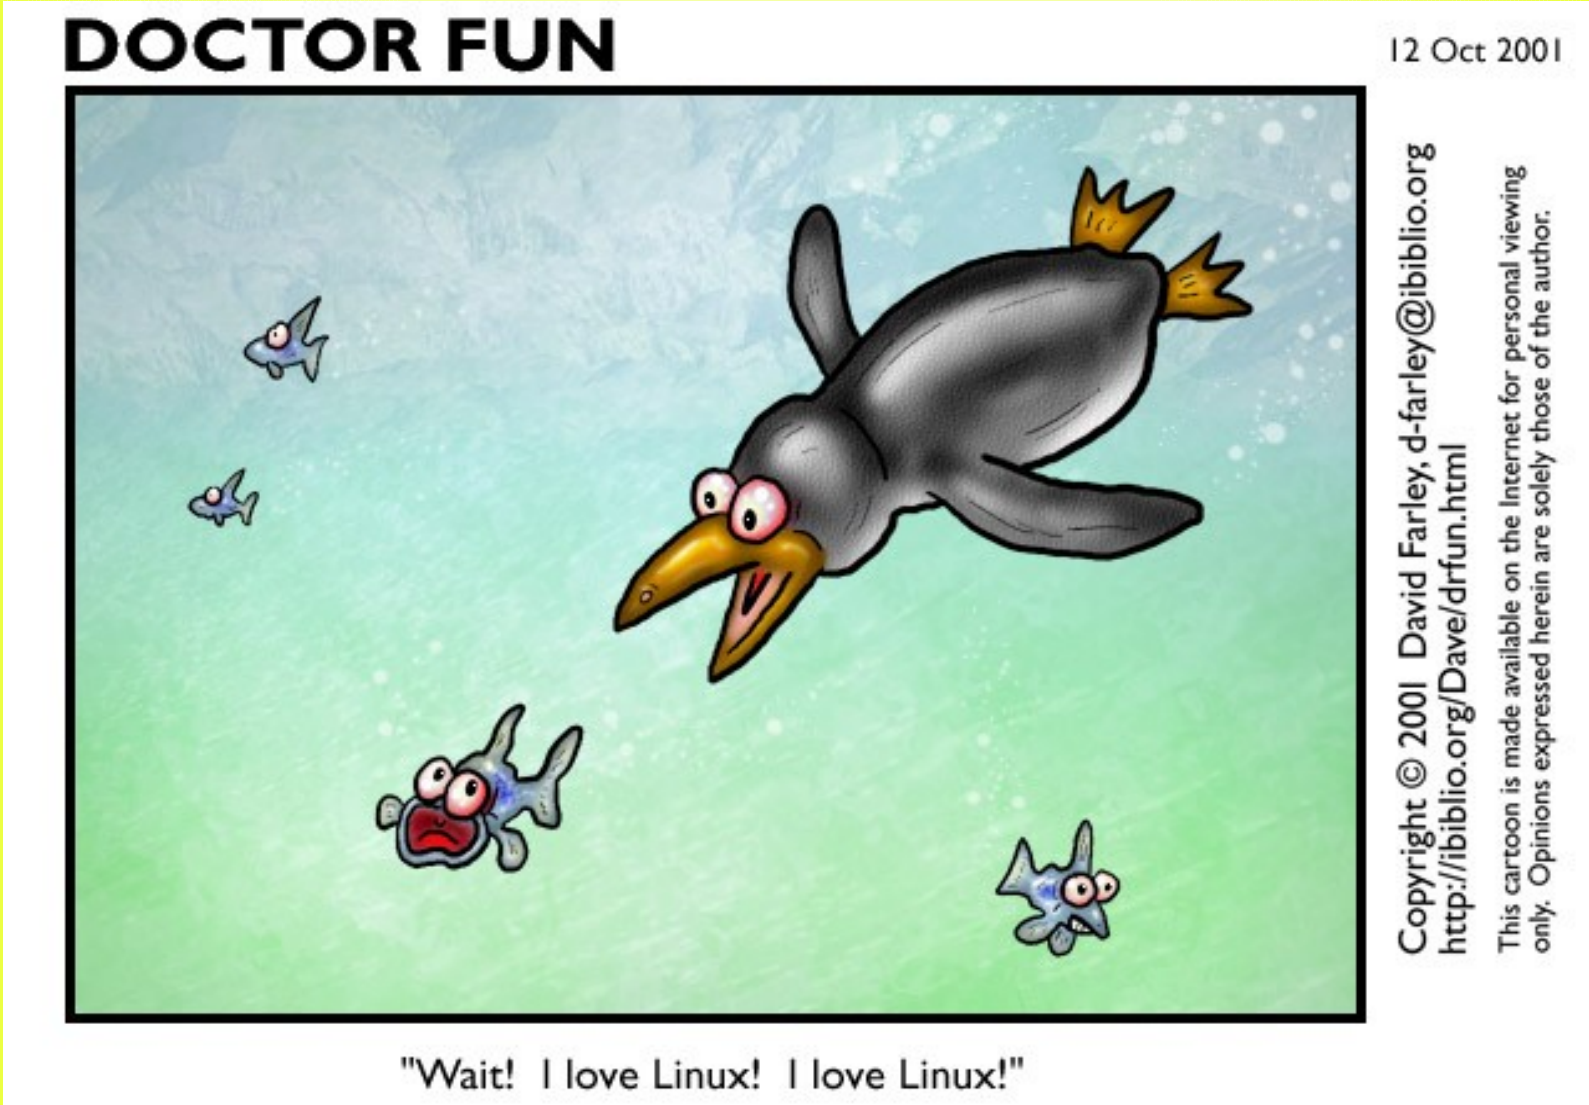

Where We Started

- Environment
- Overheads
- Underhanded, sneaky
- Sources of Ideas

Environment

- Existing clusters
- Windows XP support team
  - Including me for boot loader and Linux
- Supervisors who monitor hardware
- Thousands of users

Overheads

- By operating system
- By version
- By distribution
- By number of machines
- By type of user and usage

Underhand

- Be made available without fuss
- Installable without significant effort
- Maintenance must piggy back on existing arrangements
- All must interoperate with existing
  - user management system in LDAP
  - file systems

Sources of Ideas

- Filesystem Hierarchy Standard
  - read-only /usr
- Knoppix
- Debian supported some ability to configure the system at boot time
- Previous work on centrally managed HP systems
- Previous install system for teaching cluster

Description of the System

- Base layer
- Files on disk
- Cache
- Full system on the network

Description of the system

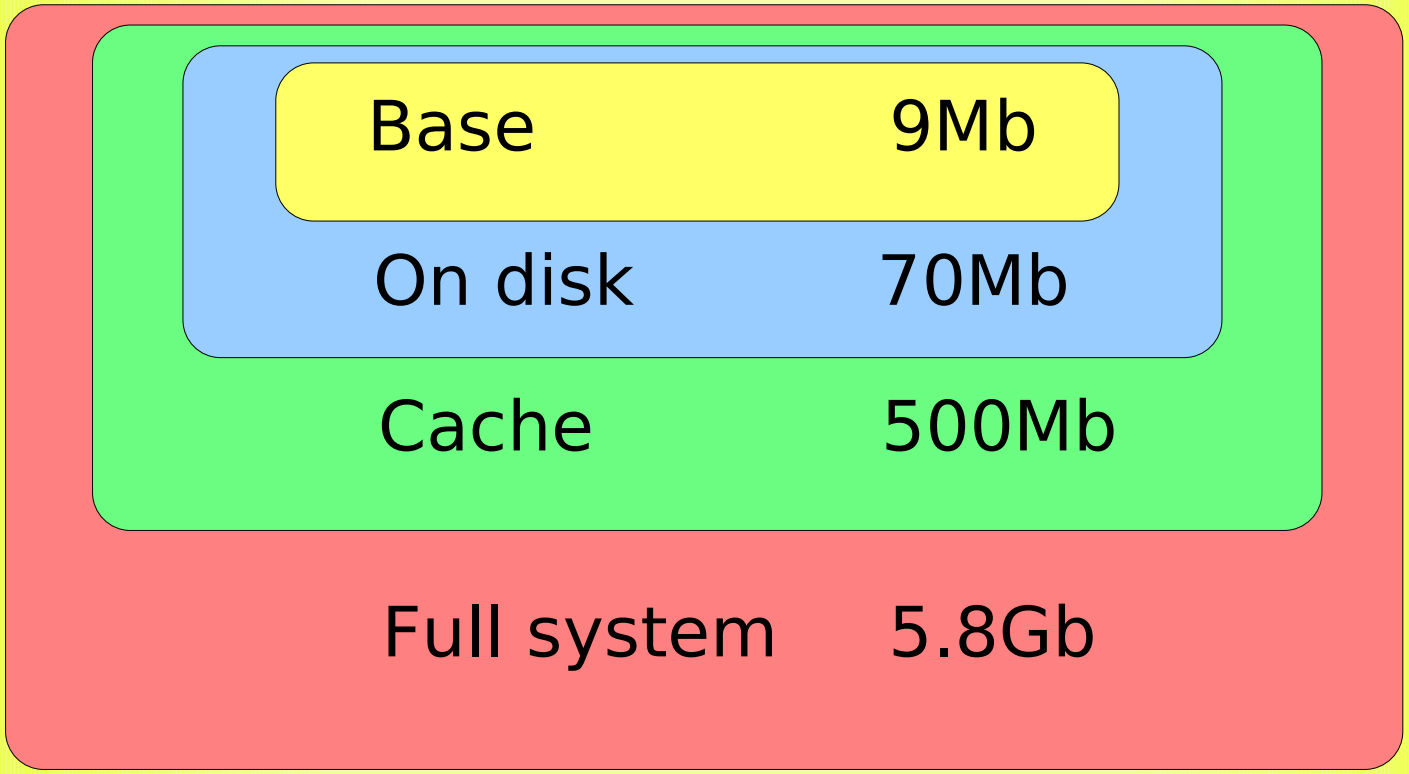

Base Layer

- In a separate partition, also containing:
  - GRUB
  - Novell's ZenWorks
- Kernel 1495 Kb
- Ramdisk 7506 Kb
- Config file 50 bytes

Base Layer (booting)

- Boot from this kernel with ramdisk as root.
- Locate config file, which describes other partitions on the disk(s).
- Swap on ($mkswap$ if necessary).
- Mount / ($mkfs$ if necessary).
- Copy/install minimum system from ramdisk.
- $pivot\_root$

Files on Disk

- System files (/bin, /etc, /lib, /sbin, /var)
- /local (/usr/local -> /local)
  - /local/home
- Symbolic links (/usr, /opt, many others)
- 521 files, 634 symbolic links, 70 Mb

Files on Disk (booting)

- AFS file system starts.
- The *package* command maintains directory structure.
- It also checks all resident files for size, date, permissions, and ownership.
- It also checks all directories, links, devices.
- It also deletes all local scratch and temporary files (not including cache files).

Adaptation of package

- Locally written utility (75 lines of C).
- Config file + hostname (FQ) + IP address.
- Package's own macro processor mpp.
- Package.
- Reboot option

Cache

- Feature of AFS file system.
- Set to 500 Mb maximum.
- Persistent across reboots.
- File access -> check -> use cached copy
  |-> get a new copy

Cache (booting)

- AFS file system already accessible.
- A special script (per cluster/machine) may run after the local system is checked.
- $/sbin/init$ starts after this.
- Changes must be made to ordinary startup and shutdown.

Mutation


und somit wirklich Geld. Mehr
Infos unter www.microsoft.com/
germany/windows2000

Full System

- Copy of normal distribution, modified to remove networking and some init scripts.
- On Network (in AFS) 3.2 Gb
- Read only – no access controls.
- AFS ACLs can be used for licensed software, if necessary.
- Users and authentication (various)
- User filestore (various)

Users and Authentication

- In central LDAP system: pam_ldap
- Special UNIX user names for courses created at boot time: pam_unix
- AFS users: pam_afs

Filestore

- Novell (mounted as NCPFS)
- SAN (mounted as SMBFS)
- Local UNIX filestore set up at boot time
- AFS for system files and a few users

The Management

How to Look After It

- Installation
- Administration
- Routine maintenance
- Bugs and problems

Installation Timeline

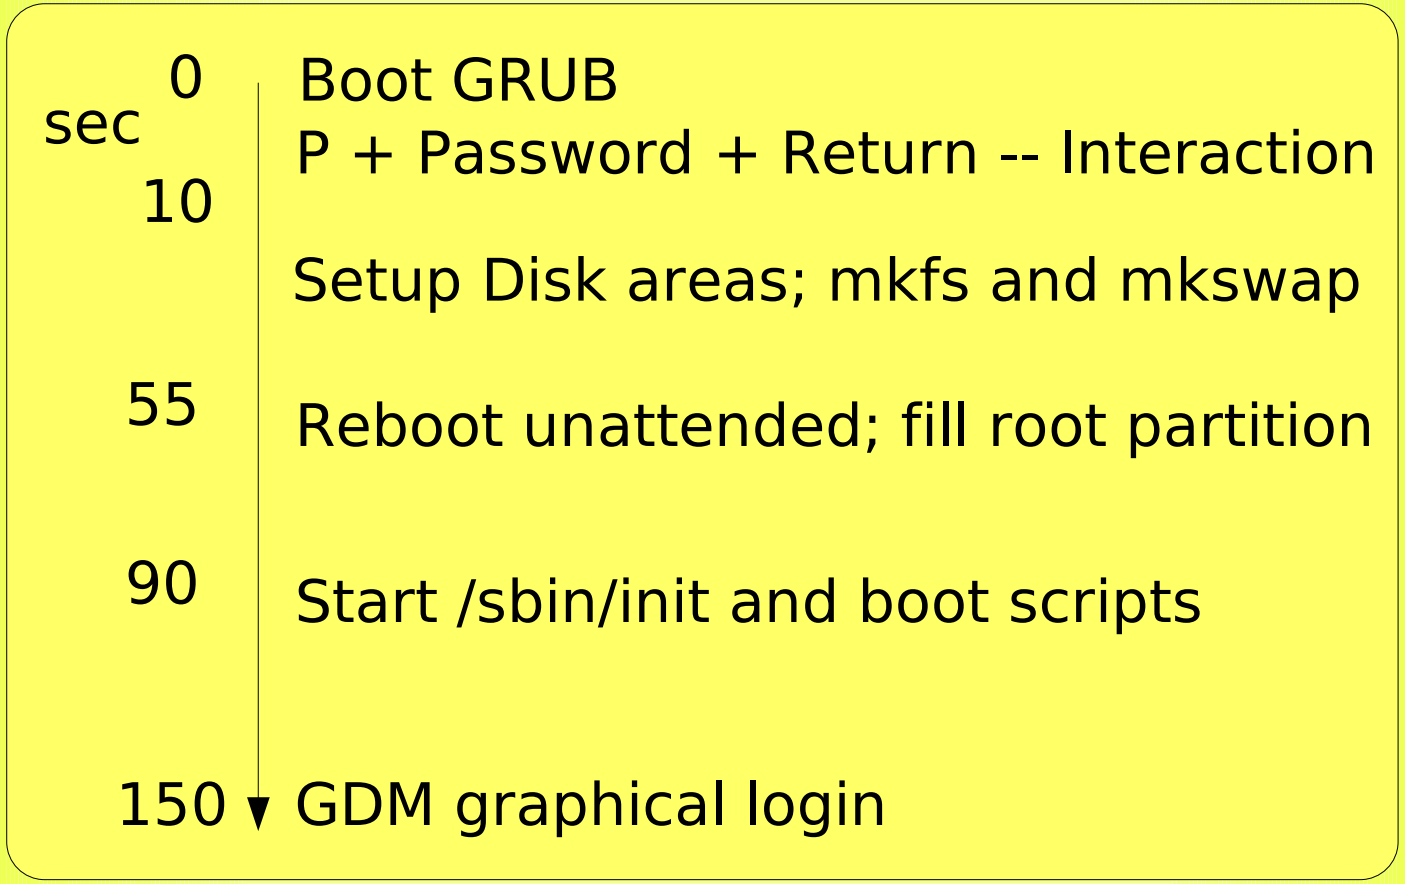

NB: This can run in parallel.

Installation with X configuration

- 2 min to line mode login
- login
  - configure X
  - save XF86Config-4
  - define new hardware type
- reboot

Administration

- rsync AFS read-only image to directory
- chroot
- apt-get update; apt-get dist-upgrade
- rsync directory to AFS
- test
- release read-write copy to read-only copy

Maintenance

- Config files for *package*
- Changes in init scripts
- Preventing services from running:
  - no exim
  - no sshd
  - no portmap
- Bugs in debian

Problems and Bugs

- Currently based on Debian
- Politics
- Changes in our local user infrastructure
- Using the system and applications in a way not foreseen by the developers

Why is it This Way?

- Constraints of politics
- Constraints of resources
- Constraints of management
- This is the only way we can do this without:
  - money
  - people
  - hardware

Evaluation

- Not rocket science
- Delays
  - at boot time
  - when using any program for the first time on that machine
- It works

Thousands of Linux Installations (and only one administrator)

A Linux cluster client
for the University of Manchester

A V Le Blanc
I T Services
University of Manchester

LeBlanc@man.ac.uk