

School of Informatics computing review

Alastair Scobie *

January 19, 2007

1 Introduction

Informatics Computing provides a large scale, reliable computing service to the staff and students of the School of Informatics.

The School of Informatics, as with the majority of Informatics academic departments worldwide, is very largely Unix based for both teaching and research. The use of Windows and MacOS PCs is largely confined to administrative staff, to personal laptops and for collaborative projects with non academic organisations.

However, the University computing infrastructure is very largely focused on Windows. For example, the University authentication and directory services is provided by Microsoft's Active Directory; as currently implemented, this is unsuitable for Unix clients. As a result, the school has been obliged to provide its own infrastructure.

Informatics, and particularly the previous Department of Computer Science, has historically played a leading role, across the University and beyond, in championing, deploying and developing new computing technologies. We are continuing to play this role both locally - the School of Engineering and Electronics are using our LCFG Linux technology to provide a Linux service for their staff and students - and beyond - our LCFG Linux technology underpinned the early CERN LHC DataGrid testbeds. A number of our COs are regularly invited to worldwide conferences to present talks on our development work on areas such as large scale machine configuration, authentication and directory services.

The vast majority of users use our commodity Unix computing environment known as DICE. This is currently based on Redhat's Fedora Linux distribution for the desktop and the majority of servers, and Solaris 9 for the remaining servers. Obviously in such a large school with a wide spread of computing research being carried out, this platform does not always meet every individual's research needs and a number of users self-manage their own machines; they continue, however, to use DICE services for email, printing, backups etc.

*with input from Paul Anderson, Ken Dawson and Steve Renals.

2 Mission statement

The objectives of Informatics Computing are :

- Provision of a high quality computing service, tailored to the needs of Informatics users;
- Responsiveness to teaching and research requirements;
- Innovation and development of new services;
- Contribution to strategic development of computing services at the School, College and University level.

3 Scope

- 190 teaching and research staff
- 45 secretarial and administrative staff
- 9 technical staff
- 20-30 visitors
- ~ 100 associates
- 800 undergraduates
- 250 research postgraduates
- 140 taught postgraduates

4 Internal Organizational structure

Until recently, computing staff were organized in a matrix structure with around 20 teams; these were managed by an executive group of 3 principal computing officers. This structure was appropriate while the school was merging the three previous computing teams from the former departments of Computer Science, Artificial Intelligence and Cognitive Science/HCRC. However, the structure was increasingly causing significant management problems, leading the school to reconsider the organization of the group.

Informatics Computing is currently in transition to the internal structure described below.

- A single Head of Computing, reporting to Deputy Head of School (Computing)
- Deputy Head of Computing reporting to Head of Computing
- Five units each managed by a Unit head, reporting to the Head of Computing.

The five units are :-

- User Support

- Managed Platforms
- Infrastructure services
- User Services
- Research and teaching application/services

Note that the Head, and Deputy Head, of Computing are also unit heads.

4.1 Computing Executive Group

The Computing Executive Group (CEG) is the main computing management group with the following membership :-

- Head of Computing
- Deputy Head of Computing
- Unit heads
- Convenor of Development meeting

It meets on a weekly basis to agree priorities and commit resources, both proactively and reactively.

4.2 Development meeting

Development and innovation is a core activity for Informatics Computing. The Development meeting monitors, assesses and prioritizes all significant development projects from initial proposal to operational signoff. All projects are subject to technical peer review. This meeting is held once a month, but perhaps more often when Informatics moves to the new Informatics Forum.

5 Interfaces to school

Simultaneously with reviewing the internal structure of the computing group, the school has also been considering the interfaces between the group and the rest of the school. Discussions are still ongoing, but the following describes the likely interfaces.

5.1 Computing Strategy Group

The Computing Strategy Group (CSG) is concerned with all operational and development issues, prioritization of work, budget decisions and resolution of computing issues in the school. It meets four times a year.

- Deputy Head of School (Computing)
- Head of Computing

- Deputy Head of Computing
- Convenor of Development Meeting
- Senior member from School's Research Advisory Committee
- Senior member from School's Teaching Committee
- Member of School's Policy and Resources Committee

5.2 Computing Forum

A regular (four per year) meeting, open to all staff, research students and representatives of taught students.

5.3 Innovation Meeting

The Innovation Meeting is a half-day meeting, open to all members of the Computing Forum and held once or twice a year. It consists of both formal and informal presentations, with plenty of discussion. The meeting reports to the Computing Strategy Group, which prioritises any ideas coming from the Innovation meeting against existing commitments.

5.4 Staff/student meetings

A computing staff representative is sent to all staff/student liaison meetings.

6 Services

6.1 Authentication

An authentication and authorization infrastructure, suitable for multiple platforms.

The underlying technology used for authentication is MIT's Kerberos, with KX509 and Cosign being used for services that can't directly use Kerberos. This provides a true single-signon for an increasing number of our services.

A powerful locally developed technology is used for authorization.

Scope

School

Resources

Several linux servers, with per-site replicated servers.

Local/Central

The university authentication service is based on Active Directory. As this is currently implemented, it is not suitable for DICE. It may be possible for us to use EASE, but there are a number of technical, security and support issues that make this unattractive at this time.

Export

Many of our enhancements and bugfixes to various software have been accepted upstream. Presented a number of papers on our deployment to conferences.

6.2 Backups

A data backup service for servers with master copies of data; eg file servers, software repository, WWW servers, self-managed machines.

This is achieved by a combination of nightly mirroring of data to off-site disk storage and weekly dumps to tape.

The off-site mirrors of users' home directories is accessible directly by the users via the network filesystem. This means that users can themselves restore files that they have accidentally deleted, without bothering support staff; this has obvious support savings.

The backup service is primarily for disaster recovery, not for archival purposes.

Scope

School

Resources

- A number of mirror servers using a combination of large local IDE disks or, increasingly, cheap SAN storage.
- A Solaris tape backup server (using Sun's Enterprise Backup software).

Local/Central

6.3 Cluster computing

The school currently has 5 beowulf clusters, totalling 150 nodes (some dual cpu). The individual nodes run the same DICE Linux platform as the commodity desktops; this is not only important from the view of avoiding duplication of effort, but it also means that users' experiments do not need porting to use the clusters.

Gridengine is used to manage the beowulf resources; eg job submission and monitoring.

The school's computational requirements are ever increasing and we expect this trend to continue. We are in the process of deploying Condor across all our desktop machines in order to make the most of our existing computing resource by harnessing any spare compute cycles available.

Scope

Principally research, but it is likely that teaching may require heavy computational power in the medium future.

Resources

150 Linux PCs, network switches for node communications, dedicated file server.

Local/Central

We are hoping that we can make effective use of the Edinburgh Compute and Data Facility when it is deployed. However, the extent to which this is possible will depend significantly on any effort required by our researchers to port their experiments to the facility.

6.4 Database

A database, developed within Informatics, dealing with all student and staff administration.

The database includes information on the following :-

- staff and visitors
- students
- taught courses
- student assessment results
- research publications
- research postgraduate applications

producing reports such as :-

- reports to Boards of Examiners (calculated marks etc)
- staff teaching duties
- telephone and email lists
- research institute membership lists
- tutorial membership lists

Some of the data is sourced via a feed from MIS databases. The database is also used to feed the DICE user account generation process and to create roles for the DICE authorization system.

Scope

School

Resources

Linux database server (running Ingres)

Local/Central

The deployment of EUCLID will necessitate a full review of the school database. However, EUCLID appears to only address student administration.

6.5 Directory services

A directory service infrastructure, suitable for multiple platforms. The directory service contains user data (such as home directory location, unix UID, group memberships, roles, email addresses) and information on physical devices such as printers.

The underlying technology used in the school for directory services is LDAP (specifically OpenLDAP).

For efficiency and security, each DICE client carries its own replicated copy of the LDAP directory.

Scope

School

Resources

Various Linux servers.

Local/Central

The university directory service is based on Active Directory. As for authentication, as this is currently implemented, it is not suitable for Unix clients.

Export

Informatics were one of the first large scale organisations to use LDAP as a replacement for the traditional unix directory service, NIS. Many of our enhancements and bugfixes to various software accepted upstream. Presented a number of papers on our deployment to conferences.

6.6 Email

An email service for staff only; with IMAP and IMP user access.

The service was deliberately designed to be as similar to the university's staffmail service as possible, to minimize the user disruption should the school choose to transition to the centrally provided service.

A mailing list service is also provided, using mailman. The mailing lists are automatically generated from the school database.

Mail forwarding for email addresses from the formative departments of the school is provided.

Scope

School - staff only

Resources

Linux email server

Local/Central

All students use the central student email service.

There is one perceived obstacle to the school using the central service for staff mail; a substantial number of our users have a significant amount of email stored on our email service; more than 100 have more than 200Mb, and nearly 20 have more than 1Gb. We understand that the standard staffmail quota is 200Mb.

In any case we would continue to require an internal email service for system mail, albeit with much reduced required functionality.

The majority of our mailing lists are generated from the school database so it is likely that we would need to continue to run our own mailing list service.

6.7 File service

Provision of a network file service to users using various technologies including NFS and Samba.

The school is in the process of deploying AFS to provide a more secure and flexible cross platform file system; this will allow self managed machines to access shared unix filesystems.

Scope

School

Resources

7 Solaris NFS/AFS file servers 2 Linux samba servers

Local/Central

We are not aware of any suitable central service.

6.8 Front line support

A team of six computing support officers providing front line user support.

All user fault reports and requests are entered into a fault tracking system (RT), either by users or CSOs responding to phone calls. CSOs process and deal with as appropriate, passing any that need more technical input to the relevant CO team.

It is worth noting that the school's user base, in general, is very computer literate; consequently the queries and requests made tend to be of a much higher technical nature than those of other schools.

This team also handles the allocation, deployment and operating system upgrades of all managed desktops (both Linux and Windows), along with the management of user accounts (eg account creation, home directory moves etc).

Scope

School

Resources

RT server, 900 managed desktops

Local/Central

6.9 LCFG

LCFG, developed in Informatics, is a system for managing the configuration of large numbers of Unix systems. It is particularly suitable for sites with very diverse and rapidly changing configurations, such as Informatics.

LCFG is used not only to configure machine configuration, but also most of the services described in this document. This means that (re)installing a server, eg a print server, often entails no more than booting the LCFG installation process off the network.

Scope

School, SEE, UCS, research grants, external

Resources

A number of Linux servers for machine configuration compilation and delivery.

Local/Central

No central provision

Export

Several papers presented to conferences worldwide.

The European DataGRID project originally used a version of LCFG to manage testbed Grid farms. The project has now developed a new configuration toolset (Quattor) based on the LCFG architecture.

EUCS are using LCFG to configure aspects of their managed MacOS platform.

6.10 Linux platform

A managed Linux platform, using LCFG to manage system configuration and locally developed tools for machine installation and software package management.

The current platform is based on Redhat's Fedora Core 3. This will be upgraded to Fedora Core 5 this summer.

Scope

School (over 1000 machines), SEE

Resources

A number of Linux servers for delivery of software packages.

Local/Central

No central provision

Export

LCFG/Linux is in use by the School of Engineering and Electronics, a number of EUCS public labs and is being considered by other schools for delivering a managed Linux platform.

6.11 Miscellaneous infrastructure services

A number of infrastructure services, including :-

- console servers
- NTP time servers
- UPS monitoring

Scope

School, UK wide (NTP)

Resources

A number of Linux servers

Local/Central

No central provision

Export

NTP service UK Wide.

6.12 Miscellaneous user services

A number of user services, including :-

- CVS - change control system for storing software projects
- Subversion - replacement for CVS
- Software license service
- Room booking system (shezhu)
- Calendar service (not widely used)
- FTP service (for legacy domains)
- Corpora provision
- bugzilla - a bug tracking system used to track software bugs and work tasks
- PostgreSQL service for teaching and research databases
- Wiki

Scope

Resources

A number of Linux servers

Local/Central

It is expected that the school will drop its own Calendar service when the central University service is available.

We are not aware of equivalent central provision for the other services.

6.13 Network

A managed network on 4 sites, with the following features :-

- over 3500 live network ports
- 60 network switches, of which 8 provide hardware routing.
- four EdLAN connections and four SRIF connections
- 30 different VLANs - many different functions with different access rights.
- extensive firewalling, both externally and internally.
- traffic load monitoring
- DNS service
- MAC address tracking and port locking
- DHCP - configured via individual machine configuration
- powerful switch configuration system with change control for audit trail and rapid switch reinstallation.
- soft patching - no manual patching once switches installed
- policy that all wall ports are live
- VPN service

Scope

School

Resources

- 60 HP network switches
- 28 wireless access points
- a number of Linux servers acting as firewalls and secondary routers

Local/Central

The school has traditionally provided its users with a network in advance of what EUCS have been able to provide, both in terms of performance and functionality. For example, Computer Science deployed structured wiring in 1991, long before the rest of the University. The new Informatics Forum will present 1000T to the majority of desktops, with fibre ducting to offices in readiness for future technologies.

We will be using the port probing service from EUCS to probe us from out-with our firewalls and to probe self-managed machines through holes in our firewalls.

We have an extensive wireless service, with nearly 30 access points; these are part of the centrally managed wireless service.

Export

Switch configuration code is used in School of Engineering and Electronics.

2 of the 5 university external DNS servers are managed by Informatics

6.14 Printing

Provision of a multi-platform networked printing service.

The underlying technology is LPRng, with LDAP used for listing printers and their capabilities, and kerberos being used for authentication. Windows clients connect via Samba.

Scope

School

Resources

56 network printers - 41 mono A4, 6 mono A3, 9 colour A4.

The number of printers appears large, but this is mainly due to the physical layout of our Bucleuch Place accommodation. We expect to have significantly fewer printers once we move to the Informatics Forum.

Local/Central

We are not aware of a central multi-platform printer service.

6.15 Solaris platform

A managed Solaris platform, using LCFG to manage system configuration locally developed tools for machine installation and software package management.

The current platform is based on Solaris 9.

Scope

Whole school - underpins file and backup services and multi-user Solaris machines.

Resources

Local/Central

No central provision

6.16 Storage Area Network

Three separate Storage Area Networks, providing a total of 34TB of fibre attached storage to 12 servers.

We are planning to reduce to one SAN once the school is colocated in the central area.

An Informatics server will shortly be installed at the ECF directly connected to the SRIF SAN to allow the school to make best use of this resource. The school has a notional allocation of 50TB of SRIF storage.

Scope

School

Resources

- 5 QLogic fibre channel switches.
- 5 Nexsan storage arrays, providing 34TB of storage.
- 50TB notional allocation of SRIF storage

Local/Central

6.17 Teaching applications

Support for over 100 taught courses, most of which require software packages to be installed, sometimes ported, and tested on the DICE platform.

Well over 100 packages, some simple, some complicated, are involved.

Scope

Teaching

Resources

Local/Central

No equivalent central service. The large majority of software is specific to Informatics teaching.

6.18 Web service

A managed Web hosting service.

The main school Web site, www.inf.ed.ac.uk, uses technology to enforce publication of standards compliant HTML; this is to ensure the pages are as browser independent as possible. All content is change controlled to provide an audit trail. A substantial number of pages on this site are autogenerated as reports from the school database. CGIs are carefully controlled.

There are around 20 virtual web servers for specific projects and interdisciplinary groups.

All users, including undergraduates, can publish their own content on the homepages.inf.ed.ac.uk web service, though use of this service for official material is strongly discouraged. CGIs are permitted, except to first and second year undergraduates.

A small number of virtual servers support the web sites of the formative departments of the school; the content of these sites is frozen.

Scope

School

Resources

A number of Linux servers

Local/Central

The University Web redevelopment project will require the school to review its web services.

7 Current significant projects

The following are current projects which require many man months of effort :-

7.1 Informatics Forum

Sometime during the middle of 2007, the school will be relocating from Kings Buildings, Buccleuch Place and Forrest Hill to the new Informatics Forum and Appleton Tower. Computing staff have been heavily involved in the design of the Forum, are currently designing the network infrastructure and will shortly be considering the procurement of the network equipment. Over the next year, significant thought will be required as to planning the relocation itself, and how the move will affect our working practices.

7.2 Upgrade of DICE platform to Fedora Core 5

The school has a policy of upgrading the DICE platform on an annual (occasionally biennial) basis. This is a resource intensive task, but the school considers that it is essential to keep the platform up-to-date in order to discourage individuals from duplicating effort by managing their own machines.

7.3 Move to AFS

A longstanding significant weakness in our infrastructure has been the reliance on the inherently insecure network file access protocol NFS. We have recently been working on deployment of the Andrew File System (AFS) which is inherently far more secure, provides multi-platform support and will allow remote file access to user files from out-with the University. AFS is widely used in similar academic organizations, such as MIT and CMU. The intention is that AFS will completely replace NFS over the next 12 months.

7.4 Restructure

As explained earlier, the Computing group is being restructured. This will inevitably result in some shifting of responsibilities between individuals; this will cost effort as knowledge is transferred, but it has the beneficial side effect of sharing knowledge over a larger number of individuals.

7.5 Port of DICE platform to 64 bit

Increasingly, our research is hitting the limits of a 32-bit address space. Fortunately our base Linux platform is available in a 64 bit flavour, but we will need to expend effort ourselves in porting our LCFG Linux platform to this variant.

8 Staff profile

The school employs 21 Computing Officers (19.3 FTE) and 7 Computing Support Officers (5.9 FTE).

- 3 FTE AD5
- 2 FTE AD4
- 9.3 FTE AD3
- 5 FTE AD2
- 2.4 FTE CD3
- 3.5 FTE CD2

9 Staff training/development

The school's intention is that 10% of all computing staff's time is ring-fenced for personal technical development - keeping abreast of developments, learning new technologies etc.

Technical training is largely achieved by a combination of training on the job and technical presentations. We consider most commonly available technical courses to be insufficiently technical, though tutorials at conferences tend to be worthwhile attending.

Staff are encouraged to attend general development courses, eg time-management etc, and technical conferences, eg those organised by UKUUG and USENIX.

10 Allocation Policies

Academic teaching staff are generally entitled to a DICE desktop and a laptop. Research postgraduates are entitled to a DICE desktop.

The school policy is for a desktop writedown of 3 years. Although this policy continues for staff desktops, recently environmental concerns have led us to redeploy 3 year old machines in less demanding roles, such as for visitors and the less commonly used student labs.

11 Resources

11.1 Figures

- Managed DICE (Linux) desktops
 - 520 staff and research postgraduates
 - 310 student labs

- Managed Windows desktops
 - 55 MDP managed desktops for administrative staff
- Managed DICE (Linux) servers
 - 60 infrastructure and user services servers
 - 40 research grant funded compute servers
 - 150 beowulf nodes (5 clusters)
- Managed Solaris servers
 - 7 NFS/AFS file servers
 - 2 multiuser login servers
 - 1 backup server
- Accommodation
 - 4 air-conditioned machine rooms, with a total of around 120 square meters.
 - 10 undergraduate teaching labs, with a total of 310 desktops