

School of Informatics computing review

Alastair Scobie *

September 21, 2007

Contents

1	Introduction	4
2	Mission statement	4
3	Scope	5
4	Internal Organizational structure	5
4.1	Computing Executive Group	6
4.2	Development meeting	6
5	Interfaces to school	6
5.1	Computing Strategy Group	6
5.2	Innovation Meeting	6
5.3	Staff/student meetings	7
5.4	Computing Forum	7
6	Services	8
6.1	Authentication	8
6.2	Backups	9
6.3	Cluster computing	10
6.4	Database	11
6.5	Directory services	12
6.6	Email	13
6.7	File service	14
6.8	Front line support	15
6.9	LCFG	16
6.10	Linux platform	17
6.11	Miscellaneous infrastructure services	18
6.12	Miscellaneous user services	19
6.13	Network	20
6.14	Printing	22

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

6.15	Solaris platform	23
6.16	Storage Area Network	24
6.17	Teaching applications	25
6.18	Videoconferencing	26
6.19	Web service	27
6.20	Wiki	28
7	Current significant projects	29
7.1	Informatics Forum	29
7.2	Upgrade of DICE platform to Fedora Core 6	29
7.3	Move to AFS	29
7.4	Port of LCFG to SL5	29
8	Staff profile	30
9	Staff training/development	30
10	Allocation Policies	30
11	Resources	30
11.1	Figures	30
A	Active projects	32
A.1	LCFG buildtools rewrite	32
A.2	DICE client LDAP configuration	32
A.3	802.1X and RADIUS	32
A.4	Cluster Parallel filesystem	32
A.5	Revisit Account Management	32
A.6	AFS file system	32
A.7	Production Cosign Service	32
A.8	Review of Backup Technology	33
A.9	Production Condor Service	33
A.10	Improved LCFG Distribution	33
A.11	Web Technologies and Policy	33
A.12	Enhancing the Solaris LCFG Platform	33
A.13	Publishing and Discussion Media Survey Paper	33
A.14	Proof General Eclipse Improvements	33
A.15	Inventory System	33
A.16	OpenLDAP Replication and Server Configuration	34
A.17	System Monitoring	34
A.18	Infrastructure for New Informatics Forum	34
A.19	Infrastructure for Appleton Tower Refurbishment	34
A.20	Introduce a Production FC5 Desktop	34
A.21	Gridengine Configuration Review	34
A.22	Future Console Servers	34
A.23	Transition of Informatics staff mail services to Staffmail	34

A.24	LCFG SL5 port (inf level)	35
A.25	Improved RPM submission tool	35
A.26	Perimeter Filter Reporting	35
A.27	Network Management Review	35
B	Completed projects	35
B.1	FC5 services	35
B.2	"Friend" lightweight authentication system	35
B.3	Project Management System	35
B.4	64-bit FC5	36
B.5	Release Management	36
B.6	OpenLDAP Upgrade	36
B.7	Introduce a Web Services Server for Teaching	36
B.8	Development of Proof General for Eclipse	36
B.9	Development of HiGraph Application	36
B.10	IGS Database Work	36
B.11	Increase the number of nodes in the 64 node cluster.	37
B.12	New Administration Samba Server	37
B.13	Inf FC6 desktop	37
C	Proposed projects	37
C.1	Investigate power management options for DICE desktops/servers	37
C.2	Management Reports from the RT system	37
C.3	Review of user documentation	37
C.4	Replace kx509 on mailman	37
C.5	Replace Kx509 with CoSign for Web Services	38
C.6	Provide a General IMAP service	38
C.7	(re)introduce Kerberised Printing	38
C.8	Printing From Non-DICE Clients	38
C.9	Review of Repository Services	38
C.10	Bring Mail Forwarding on the Virtual Mail Relay Under the Control of RFE or Equi	38
C.11	Suexec for [groups—conferences—workshops].inf	39
C.12	Improve www.inf.ed.ac.uk search technology	39
C.13	Replacement of LPRng with CUPS	39
C.14	Printer Accounting and Quotas	39
C.15	LCFG boot component rewrite	39
C.16	DICE under emulators	39
C.17	Enhance LCFG fstab component	39
C.18	Enhance LCFG network component	40
C.19	Evaluate linux distributions	40
C.20	LCFG documentation technology	40
C.21	Review of LCFG header structure	40
C.22	LCFG log server	40
C.23	Solaris and Mac LCFG support mechanism	40
C.24	Installroot - investigate integration with PIE	40
C.25	Investigate benefits of moving to Solaris 10	41

C.26 Investigation of server virtualization	41
C.27 Design new LCFG package specification	41
C.28 Re-implement ISDD System	41
C.29 Personal Response Proxy	41
C.30 Virtual Open Day	41
C.31 Re-Implement Submit Practical Submission System	42
C.32 LCFG Linux install redevelopment	42
C.33 Improve Host key management	42
C.34 Centralise logging	42
C.35 PPTP endpoint	42
C.36 LCFG Core Refactoring	42

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. 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.

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

- 223 teaching and research staff (in 6 research institutes)
- 55 secretarial and administrative staff
- 27 computing staff
- 9 technical staff
- 20-30 visitors
- ~ 100 associates
- 500 undergraduates
- 277 research postgraduates
- 137 taught postgraduates

4 Internal Organizational structure

Since the reorganization of the group in the spring of 2006, Informatics Computing has had 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

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 considered interfaces between the group and the rest of the school. Discussions are still ongoing, but the following describes the current 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 monthly.

- Deputy Head of School (Computing)
- Head of Computing
- Senior Computing Staff
- Senior member from School's Research Advisory Committee
- Senior member from School's Teaching Committee

5.2 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.

In 2007 so far, there have been two Innovation meetings, focusing on web developments and distributed computing. There is a meeting proposed for early 2008 to focus on data backup and archival.

5.3 Staff/student meetings

A computing staff representative is sent to many of the staff/student liaison meetings.

5.4 Computing Forum

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

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.

We have recently developed and deployed a lightweight authentication system iFriend, similar to the Cosign Friend system, to allow users not affiliated to Informatics to authenticate to our systems. The advantage of iFriend over Cosign Friend is that it allows us to provide access to non web services, such as subversion.

Scope

School

Resources

Several linux servers, with per-site replicated servers.

Local/Central

The university authentication service is based on EASE.

We use kerberos in many innovative ways for a wide variety of services, such that we are recognised as being world leaders in kerberos deployment. Transitioning to the central EASE service, as it currently stands, would lose us the benefits of our innovative uses and affect the manageability of our managed platform.

Export

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

We participated heavily, contributing both language and example code, in the standardisation effort for SSH-GSSAPI and are acknowledged in the resulting RFC. Code derived from our SSH-GSS implementation ships with all Unix operating systems.

6.2 Backups

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

This is achieved by a combination of nightly mirroring of data to off-site disk storage and nightly 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. The issues of backups and archival are scheduled for review.

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

Our understanding is that the existing central backup service is too small in scale to meet our requirements. However we are investigating the possibility of making use of the forthcoming SAN based service.

We are very interested in the mooted central archival service and would welcome involvement in its specification.

6.3 Cluster computing

The school currently has 4 beowulf clusters, totalling 152 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 is already making heavy use of the ECDF cluster and the current expectation is that central provision will replace much of our local provision over time.

In order to make the most of our existing computing resources, we have deployed Condor across around 200 of our desktop machines to harness spare compute cycles when such machines are idle. We expect to increase the size of this Condor "flock" in the near future.

Much of our cluster computing requires fast access to large amounts of data. We are currently investigating, in partnership with the ECDF team, the deployment of GPFS on our compute clusters and SAN storage; this hopefully will provide high performance fileaccess of both ECDF storage and School storage on both ECDF cluster nodes and School cluster nodes.

The school's computational requirements are ever increasing and we expect this trend to continue.

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 already making use of ECDF and expect our local provision to decline over time. There will, however, be some residual requirement for our own compute clusters, particularly to meet the requirements of those researchers working on systems research where clusters need to be reconfigured at short notice as research requirements change.

We also harness spare compute cycles of our Linux desktops using the Condor system, but given that this sits upon a School administered resource (our desktops), it is not clear how this could be a central provision.

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

There is currently no general purpose directory service provided by IS.

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.

We are planning to move staff to the central staff email service in the near future; all computing staff have already moved across in order to identify any issues.

Once we have moved staff to the central service, we shall 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 OpenAFS, NFS and Samba.

The school is part way through migrating all users from NFS to OpenAFS, which provides a more secure and flexible cross platform file system allowing self managed machines to access shared unix filesystems.

We are currently trialing Linux OpenAFS file servers.

Scope

School

Resources

7 Solaris NFS/OpenAFS file servers 2 Linux OpenAFS file servers (shortly) 2 Linux samba servers

Local/Central

The School's heterogeneous platform nature dictates the provision of a cross platform file service. We also require a file service to be performant and secure. We are unaware of a central IS service that meets all these requirements and would be very unlikely to migrate to a central service that did not meet them.

Export

We are active participants in the OpenAFS community, sending a delegate to the annual AFS and Kerberos Best Practices workshop. Paper presented on our OpenAFS deployment at UKUUG Spring 2007 conference.

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, 850 managed desktops

Local/Central

Our increasing shift to using more and more central services makes it more difficult to manage a single point of contact for our users' queries/fault reporting. Adopting the central IS CMS for all our queries/fault reporting would make it easier to pass tickets between ourselves, central services, and any other university service providers. A move to the central IS CMS is being investigated.

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, EPCC, IS, 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.

IS Desktop Services are using LCFG to configure aspects of their managed MacOS platform.

A LCFG Deployers group, consisting of those active in LCFG development, meets monthly; this currently has attendees from IS, SEE and EPCC.

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.

Currently supported platforms are both 32 bit and 64 bit versions of Redhat's Fedora Core 5 and 6.

The School and EPCC are currently collaborating on producing a Scientific Linux version of the managed Linux platform. Other CSE schools have indicated an interest in deploying this platform.

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, School of Mathematics, a number of IS 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

A reliable NTP time service is critical as synchronised time is fundamental to Kerberos.

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

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, eDiary, becomes available.

We are investigating the possibility of deprecating our room booking service in favour of the eDiary service.

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 4440 live network ports
- 75 network switches, of which 5 provide hardware routing.
- four EdLAN connections and four SRIF connections
- 70 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

We are planning on providing a 10Gb backbone when we colocate in the Forum.

Scope

School

Resources

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

Local/Central

Our wired network is our "virtual lab"; it underpins all our diverse research and teaching requirements as well as our commodity needs. As a result it has flexibility, performance and reliability requirements well in advance of what other Schools require.

The school has traditionally provided its users with a network in advance of what IS 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 IS 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 over 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. Windows clients connect via Samba.

We are currently investigating moving from LPRng to CUPS as the LPRng code base is no longer being actively maintained, and CUPS should allow us to meet our long held goal of authenticated printing. It is very likely that we will be collaborating with another School in this development.

We are also investigating the use of MFDs to replace some of our printer provision in the Informatics Forum.

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 Buccleuch 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 (SANs), providing a total of 58TB of fibre attached storage to around 20 servers.

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

The school also has an allocation of 35TB of SRIF storage. An Informatics server is directly connected to the SRIF SAN at the ACF to allow the school to access a portion of this storage from its systems. The remainder of the school's allocation is available, via GPFS, for the school's ECDF cluster users.

Scope

School

Resources

- 5 QLogic fibre channel switches.
- 7 Nexsan storage arrays, providing 58TB of storage.
- 35TB 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 Videoconferencing

The school has recently introduced a desktop videoconferencing service, based on Adobe Connect, for a research group. Should this service prove successful, it may be scaled up to support the whole school.

Scope

Research

Resources

A Windows 2003 server

Local/Central

No equivalent central service, but IS are being involved in the trialing.

6.19 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.

A review of the school's web service indicated that there was a clear and pressing requirement for a CMS service for, particularly, research institutes. After discussions with the IS web development team, we concluded that the central polopoly service would not be able to meet our requirements in the short to medium term. This service, based on Zope/Plone, should enter service in October 2007.

Scope

School

Resources

A number of Linux servers

Local/Central

We are expecting the School site to transition to the central polopoly service. Transition of other sites (research institute/ collaborative sites) will depend on the ability of the central service to meet content representation and access rights requirements of individual research groups.

6.20 Wiki

A wiki service, based on twiki.

Scope

School

Resources

A linux server.

Local/Central

Ideally we would like to make use of the new central Wiki when it enters service. However, whilst we can almost certainly use the central Wiki for new requirements it is not clear whether it will be possible or practical to transfer existing content from our current school Wiki service to the central Wiki.

7 Current significant projects

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

7.1 Informatics Forum

Sometime in early 2008, 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 and the network infrastructure, and are shortly to procure the network equipment. During the last year, significant thought has been put into planning the relocation itself, and how the move will affect our working practices and the planning still continues as details of the move are confirmed.

7.2 Upgrade of DICE platform to Fedora Core 6

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 36 months.

7.4 Port of LCFG to SL5

Informatics and EPCC are currently engaged in a joint project to port the LCFG linux platform to Scientific Linux 5. This work will form the basis of a College of Science and Engineering SL5 platform for both desktops and servers.

The school policy is that the DICE desktop platform is upgraded to the latest underlying OS version on a yearly basis. Although there is no equivalent policy for servers, past practice has been that servers are upgraded on the same basis. Servers have been upgraded on this basis largely because of the short life of the Fedora platform and consequent concerns re availability of security patches.

However, this annual server upgrade has proven to be very costly in terms of effort in porting services, upgrading servers, and disruption to services. Moving to a more "stable", long-life, platform such as Scientific Linux, for servers, should allow the school to move to a more manageable biennial upgrade cycle.

8 Staff profile

The school employs 20 Computing Officers (19.0 FTE) and 7 Computing Support Officers (6.4 FTE).

- 6 FTE Grade 9
- 7 FTE Grade 8
- 7 FTE Grade 7
- 3.4 FTE Grade 6
- 3 FTE Grade 5

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
 - 250 student labs
- Managed Windows desktops

- 55 MDP managed desktops for administrative staff
- Managed DICE (Linux) servers
 - 120 infrastructure and user services servers
 - 40 research grant funded compute servers
 - 150 beowulf nodes (4 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.
 - 7 undergraduate teaching labs and 2 tutorial rooms, with a total of around 250 desktops

A Active projects

A.1 LCFG buildtools rewrite

The existing scripts for building LCFG packages (buildtools) are starting to creak seriously and are unsuitable for use outwith Informatics as they rely on DICE; this is a serious barrier to collaborating with other sites on LCFG development.

A.2 DICE client LDAP configuration

Investigation into DICE client LDAP configuration, to replace our existing home grown technology.

A.3 802.1X and RADIUS

802.1X is a standard mechanism for network authentication. Based on the credentials passed by the user, the port or wireless connection may be encrypted in various ways or configured onto a VLAN with the access rights considered appropriate.

A.4 Cluster Parallel filesystem

A parallel filesystem for use on the departments clusters which integrates with the university ECDF cluster.

A.5 Revisit Account Management

Rethink the way in which we perform account management, such that we can better manage multiple identities, and a large number of distributed services

A.6 AFS file system

This project will manage the implementation and deployment of the OpenAFS distributed file system within the Informatics network.

A.7 Production Cosign Service

Complete the introduction of the Cosign/WebLogin service to mostly-replace KX509.

A.8 Review of Backup Technology

The current backup system is not much changed from that in use when Informatics was created. The vast increase in the amount of data we need to back up and the forthcoming introduction of and AFS based file system means that we need to review the technology and methods we are using to carry out backups and also clearly define the purpose of backups. This project would carry out such a review.

A.9 Production Condor Service

Create a production quality Condor service for use on staff and lab machines allowing Research Staff and students to use spare CPU cycles for computationally intensive tasks.

A.10 Improved LCFG Distribution

The distribution of LCFG products (the headers and component packages) needs reworking to provide better means of access for external (to Informatics) users.

A.11 Web Technologies and Policy

Assess the future requirements of the stakeholders in the Informatics Web Service and form a strategy to meet those requirements.

A.12 Enhancing the Solaris LCFG Platform

The core LCFG platform is incomplete on Solaris and would benefit from being enhanced in a number of ways to improve reliability and manageability.

A.13 Publishing and Discussion Media Survey Paper

Investigate what publishing and discussion facilities staff require and what facilities are already available.

A.14 Proof General Eclipse Improvements

Further improvements to the Proof General Eclipse system.

A.15 Inventory System

An inventory system to replace the existing mysql inventory and orders database with web/cgi interface.

A.16 OpenLDAP Replication and Server Configuration

Investigation into OpenLDAP replication and proxy-caching server and client configuration.

A.17 System Monitoring

Deploy a monitoring system configured via LCFG resources. In stage 1, this system will purely monitor the AFS service, stage 2 will expand this to being usable by all service component authors.

A.18 Infrastructure for New Informatics Forum

This project encompasses all the Infrastructure Unit activities relating to the new Informatics Forum.

A.19 Infrastructure for Appleton Tower Refurbishment

Ensure that the network infrastructure is in place for the ongoing Appleton Tower refurbishment.

A.20 Introduce a Production FC5 Desktop

Create a production quality FC5 Desktop for staff and lab machines. This is a subproject of the FC5Upgrade

A.21 Gridengine Configuration Review

Review the gridengine configuration used on each of the gridengine clusters and develop a new policy for scheduling jobs.

A.22 Future Console Servers

Review our options for "remote" console servers for the new building.

A.23 Transition of Informatics staff mail services to Staffmail

This project will decide whether the time is right to delegate Informatics staff mail services to the Staffmail service provided by EUCS and, if the decision is taken to proceed, plan and carry out the move.

A.24 LCFG SL5 port (inf level)

Develop Inf level support for Scientific Linux 5 (SL5) targetting the desktop and server environments on i386 and x86_64 cpu architectures. This work will form the basis of a College of Science and Engineering SL5 platform for both desktops and servers. ECDF runs Scientific Linux. It is highly likely that it will form the basis for a DICE server platform.

A.25 Improved RPM submission tool

Convert the rpmsubmit tool to use AFS instead of NFS.

A.26 Perimeter Filter Reporting

The aim of this project is to revamp, substantially, the mechanisms for reporting on filter holes, probably to include more targetted automatic mail and some generated web information and status pages.

A.27 Network Management Review

Review our existing network management technology, to see how it will fit in with the new building and potential 802.1X deployment, and whether it could usefully be enhanced.

B Completed projects

B.1 FC5 services

Coordinate and monitor the transition of all DICE FC3 based services to FC5. The actual responsibility of transitioning the services will lie with the Unit responsible for that service.

B.2 "Friend" lightweight authentication system

Develop a lightweight authentication system, similar to (or based on) the cosign "Friend" system, to allow users not affiliated with Informatics to authenticate to our systems.

B.3 Project Management System

A tool to be used for managing all stages of the project process.

B.4 64-bit FC5

Develop a 64-bit version of the DICE FC5 platform.

B.5 Release Management

This project breaks the direct connection between the LCFG include and package files and the individual LCFG profiles of DICE machines, interposing a regime of weekly DICE releases. This allows us to test each DICE release for integrity and stability before it is widely deployed.

B.6 OpenLDAP Upgrade

Upgrade OpenLDAP clients and servers to latest stable version.

B.7 Introduce a Web Services Server for Teaching

Create a student web services server, like homepages.inf.ed.ac.uk, for course work. To use Tomcat as a server engine with JSP (Java Servlets) and a backend database and (for example) Hibernate for persistent object storage.

B.8 Development of Proof General for Eclipse

Proof General is a generic interface for proof assistants (such as Isabelle, developed by LFCS and with a worldwide user base. Its interface is currently based on, and runs in XEmacs or GNU Emacs, but development has begun on porting this system to the IDE / development 'Platform' Eclipse.

B.9 Development of HiGraph Application

HiGraph is a system for presenting graphs, for displaying and manipulating hierarchical proofs generated by the proof planner IsaPlanner.

B.10 IGS Database Work

The following combines multiple separate requests for improvements to the Informatics Graduate School Database into one project for efficiency. The requests covered are Tracking Applications, Recording Disability, Recording Institute Actions, Refining Status and User Reports and are outlined in detail in the Proposal section.

B.11 Increase the number of nodes in the 64 node cluster.

There is currently space in the beowulf racking for additional nodes for this cluster. With the addition of a shelf in rack 11 and utilising existing spare network kit it should be possible to add up to 16 desktop machines to lion.

B.12 New Administration Samba Server

This project covers the introduction of a new Samba server for administration staff. All data provided by the server will be locally mounted thereby eliminating unreliability in the Samba service due to automounter related problems

B.13 Inf FC6 desktop

Develop Inf level support for Fedora Core 6 targetting the desktop environment on i386 and, if possible, x86_64 cpu architectures. This work will form the basis of the DICE FC6 managed platform and allow other Units/COs to target FC6.

C Proposed projects

C.1 Investigate power management options for DICE desktops/servers

This project will investigate what power management techniques we could apply to DICE desktops, and perhaps servers, to reduce the school's energy consumption.

C.2 Management Reports from the RT system

The RT response tracker system holds a lot of data about tickets and statistics distilled from this data would be valuable management information. The aim of this project is to create a system that can deliver accurate, timely and relevant management information from the data held in the RT database.

C.3 Review of user documentation

Review of the Informatics documentation for users and computing staff that is currently spread across various Web sites with a view to developing a coherent, searchable structure.

C.4 Replace kx509 on mailman

The mailman program on lists.inf currently uses kx509 for authentication. We should replace this with CoSign

C.5 Replace Kx509 with CoSign for Web Services

Currently we use kx509 to control access to restricted web pages and services such as www.inf and wiki.inf. This project would replace the kx509 controls with CoSign in line with other Informatics' services.

C.6 Provide a General IMAP service

Even if the proposed move to Staffmail takes place, it may still be worthwhile to run a general Informatics IMAP service allowing storage and retrieval of files on a remote server via IMAP.

C.7 (re)introduce Kerberised Printing

Originally printing under DICE was controlled via Kerberos. Following an OS upgrade, we saw continual problems with runaway lpd processes when under heavy load. The solution was to turn off kerberised printing. Since then, printing executables and kerberos libraries have both changed. We should investigate whether this problem still exists and if not re-enable kerberised printing. Note that this may have an impact on the 'Printing from Non DICE Clients' project.

C.8 Printing From Non-DICE Clients

This project would reconsider how we manage printing from non-DICE clients such as Windows PCs and Macs.

C.9 Review of Repository Services

At present we have two separate repository systems, CVS and Subversion, This project would carry out a review of repository services with the aim of unifying the existing services as much as possible and in addition consider how best access to the repositories might be managed (including access by non-DICE users).

C.10 Bring Mail Forwarding on the Virtual Mail Relay Under the Control of RFE or Equi

Currently, the config files controlling how mail is forwarded from the virtual mail relay can only be edited by computing staff. This project would make these maps editable by rfe giving non-computing users control over where their mail is forwarded to.

C.11 Suexec for [groups—conferences—workshops].inf

Currently all CGIs on the above services run as apache. This project would implement a suexec on the above services similar to that used on homepages.inf except that the CGIs would run with an effective GID.

C.12 Improve www.inf.ed.ac.uk search technology

Currently we are using Google as a stopgap solution to providing a search facility on www.inf. We should look into how this could be improved.

C.13 Replacement of LPRng with CUPS

Some work has already been done on writing a CUPS component with a view to possibly replacing LPRng with CUPS. This project should consider whether such a move is desirable and if so complete the work.

C.14 Printer Accounting and Quotas

At present we do not impose printing quotas or do any accounting of printer usage. Some work has been done on this in the past but not completed. This project should decide whether we wish to implement printer accounting and quotas and if it is agreed that we should do so complete the work.

C.15 LCFG boot component rewrite

The current LCFG boot component uses a traditional rc script style numbering scheme to determine the order in which it starts and stops individual LCFG components; this is difficult and error-prone to configure correctly. This project would involve enhancing/rewriting the boot component to use explicit component dependencies to determine component start order. It may also consider any parallelisation that may be possible

C.16 DICE under emulators

Investigate running DICE under VMware and Parallels.

C.17 Enhance LCFG fstab component

The current LCFG fstab component does not support LVM and RAID - these need manually configured. This project would add support for these to the component.

C.18 Enhance LCFG network component

The current LCFG network component does not allow for fetching network configuration via DHCP - all configuration must be via static LCFG resources. This project would enhance the network component to support DHCP. This might be needed for the move to the new building, but has been requested, at least, by SEE.

C.19 Evaluate linux distributions

We have been running Redhat since 1997! This project would evaluate other linux distributions to see which distribution, or distributions, would best meet our needs.

C.20 LCFG documentation technology

The existing LCFG implementation is documented reasonably well in the LCFG guide. However, this is becoming less useful and needs updating.

C.21 Review of LCFG header structure

The current LCFG header structure has now been in place for a year. This project would perform a review of the structure, particularly considering multi-platform support and minimal core LCFG support.

C.22 LCFG log server

The current log server does a lot more than just server LCFG logs from the client machines. It also serves documentation, resource and status information. It would be good to replace the log server with a new info server which would feature improved internals and preferably a set of templates for the interface to extract the html interface from the Perl code.

C.23 Solaris and Mac LCFG support mechanism

The current LCFG core is in production use for both Solaris (Informatics) and OsX (EUCS). There is currently no standard mechanism for distributing these versions, or for ensuring that new releases are tested on these platforms.

C.24 Installroot - investigate integration with PIE

EUCS have developed an installation technology, PIE, for the Managed Desktop Project. They also use this, in a somewhat hacky way, to install LCFG machines. This project, which would be joint with EUCS, would investigate whether there is any mileage in integrating the PIE technology with our LCFG installroot.

C.25 Investigate benefits of moving to Solaris 10

This project would investigate the benefits of moving to Solaris 10.

C.26 Investigation of server virtualization

This project would take a look at the various server virtualisation implementations and report on whether this technology could be deployed and whether it would produce savings in practice.

C.27 Design new LCFG package specification

We should change the line format of the package lists to be more generic. Currently it is very oriented towards specifying RPM packages and is not totally suitable for other systems, such as Solaris or MacOSX. The idea would be to create a single library (probably in C) with a perl higher level interface which would do the line parsing and return a struct (or similar) of the useful information about the package (name, version, release, architecture, flags).

C.28 Re-implement ISDD System

The ISDD system was developed with not fully defined design requirements. As such some features now seen as important and not incorporated and the implementation language chosen is likely to be different.

C.29 Personal Response Proxy

Provide a Tracker/Response System for ITO. For personalization of recruitment for upcoming students (Undergraduate) to get questions answered by an academic member of staff but proxied through the ITO. So the students use a web form to submit a question which feeds into an RT system for the ITO. The ITO staff perform the triage and respond to 99

C.30 Virtual Open Day

Provide a Virtual Open Day framework for the Web. Set up a Virtual Open Day web site for streaming mini videos providing a virtual tour of informatics.

C.31 Re-Implement Submit Practical Submission System

The current Submit submission system is written in C. This project is to rewrite it in Perl and to convert the current back-end which uses NFS storage space to a version control system such as Subversion.

C.32 LCFG Linux install redevelopment

This project will rework the LCFG linux install technology so that it uses an additional step to load the kernel modules required for mounting the installroot.

C.33 Improve Host key management

Improve the host key management system (kdcregister and associated utilities) so it no longer has dependencies on non-public MIT Kerberos header files

C.34 Centralise logging

The evaluation phase of this project will look at the options to decide whether or not it's sensible to run a centralised logging service.

C.35 PPTP endpoint

Investigate and perhaps implement a PPTP endpoint.

C.36 LCFG Core Refactoring

A project to rewrite the LCFG core (server/client) to be cleaner and easier to maintain and develop, without significant change in functionality.