School of Informatics computing review

1st May 2009

Contents

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 Information Systems Managed Linux desktop service, currently in use by several schools in the College of Science and Engineering, is based on our LCFG Linux technology - and beyond - our LCFG Linux technology underpinned the early CERN LHC DataGrid testbeds. A number of our COs are regularily invited to worldwide conferences to present talks on our development work on areas such as large scale machine configuration, authentication and directory services.

The majority of users use our commodity Unix computing environment known as DICE. This is currently based on the Scientific Linux 5.2 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 file services, 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

- 230 teaching and research staff (in 6 research institutes)
- 55 secretarial and administrative staff
- 22 computing staff

- 9 technical staff
- 20-30 visitors
- ~ 100 associates
- 500 undergraduates
- 280 research postgraduates
- 140 taught postgraduates

4. Internal Organizational structure

Informatics Computing has had the internal structure described below.

- A Head of Computing (Alastair Scobie), reporting to Head of School
- A 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 (Alison Downie)
- Managed Platforms (Alastair Scobie)
- Infrastructure services (George Ross)
- User Services (Craig Strachan)
- Research and teaching application/services (Tim Colles)

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
- Unit heads

It meets on a weekly basis to agree priorities and commit resources, both proactively and reactively. It also drafts policies for discussion at the Computing Strategy Group.

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.

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) - Prof. Steve Renals
- Head of Computing
- Senior Computing Staff
- Senior member from School's Research Advisory Committee
- Senior member from School's Teaching Committee

5.2. Computing Forum

A regular (termly) meeting, open to all staff, research students and representatives of taught students. The Head of Computing gives a presentation on recent work and developments in Informatics Computing.

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.

In 2008 there was one Innovation meeting which discussed Personal Computing.

5.4. Staff/student meetings

A computing staff representative is sent to many of the 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 Cosign and KX509 being used for services that can't directly use Kerberos. This provides a true single-signon for most 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.

We have our own locally developed system to automatically sign and distribute certificates for https web servers and other purposes (e.g. cosign clients). The way we deal with certificates will be changing because of industry developments.

Scope

School; SICSA (www.sicsa.ac.uk)

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. This would also introduce off-site dependencies on the EASE authentication servers, which could affect the reliability of our service.

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.

We added SPNEGO support to the upstream Cosign.

SICSA is making use of our authentication services to control access to its sites.

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.

We are currently bringing into service a new backup service which is more directly focused on our AFS based file service whilst offering greater scope for backup of other data such as that stored on laptops. The intention is that this service will have replaced the current backup regime by the start of the next academic year.

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).
- A Linux tape backup server (using Teradactyl's TiBS software).
- A Sun StorageTek L180 tape library with 4 LTO2 tape drives and 180 tape slots.

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 one beowulf cluster, totalling 58 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 making increasingly heavy use of the ECDF cluster and the current expectation is that the 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. The use of Condor must in the future be balanced against the opposing desire to save energy by turning off idle desktops.

Much of our cluster computing requires fast access to large amounts of data. We are currently running a test GPFS service on our compute clusters to evaluate performance. We are acting in partnership with the ECDF team so that this GPFS service can 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

60 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 :-

- automated mark returns to students
- reports for Boards of Examiners (calculated marks etc)
- staff teaching duties
- telephone and email lists
- research institute membership lists
- tutorial membership lists
- course membership lists

Some student data is sourced via a derived DACS 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.

Includes some local EUCLID enhancements. Building off an automated pull of data from EUCLID, our central database also manages the PAVD and PGT processes that are not covered by EUCLID.

Other miscellaneous supporting services (generally separate from the database system itself although often taking data feeds) for the Teaching Organisation are:

- online course descriptors - web based system for creating and maintaining DPTs for courses - EUCLID will replace this
- online course proposals - web based system for proposing new courses - EUCLID will not provide this functionality
- UG4/MSc projects databases - web based system for students to select projects and for the projects organiser to manage allocations - EUCLID will not provide this functionality

- online course registration - web based system for students to select their courses - EUCLID will not provide this functionality
- online questionnaires - web based system for getting student feedback - EUCLID will not provide this functionality
- WebMarks - web based system for staff to enter project marks and comments for automatically producing paper forms for College - EUCLID will not provide this functionality
- multiple per-course student mailing lists, per-course lecturer/TA mailing lists, tutorial group mailing lists and other mailing lists - although there is a central mailing list service it does not support the automatic creation and maintenance of mailing lists from a data feed which we make significant use of locally
- PAVD - system for managing Post Application Visit Days, tied to EUCLID data feed
- PGT - system for managing letters and emails sent to taught postgraduates during the application process, tied to EUCLID data feed

Scope

School

Resources

Linux database server (running Ingres). Other LAMP-like services.

Local/Central

The deployment of EUCLID will necessitate a full review of the school database. The expectation was that EUCLID would replace most of the student related information in our local database. However, this is no longer the case. Our existing feeds will still be dropped as a result of EUCLID and an alternative feed from EUCLID will need to be provided from the centre, at the very least to support student account and authorization processes.

The new central mailing list service currently being investigated is intended to support the automated creation of lists from database information, in which case we hope to be able to use that in the future.

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

Much of the content of our directory service is automatically generated from the School database, our account management system and our host configuration database (lcfg). These systems are tightly coupled over authenticated channels to ensure the security of the data.

In any case, the use of central directory servers would introduce off-site dependencies, which could affect the reliability of our service.

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 mailing lists and forwarding of mail.

All user mail services are now devolved to the equivalent central services.

We are continuing to run a local mail service to handle system mail, mail sent to legacy domains and mail sent to former users.

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

Scope

School - system mail only.

Resources

Linux email server

Local/Central

All users use the central student email service.

We are in discussion with IS regarding their plans for a new central mailing list server. Should the new service meet our needs, we would plan on migrating to it in due course.

6.7. Examinations

The Laboratory Examination environment is a special locked down version of DICE we switch whole labs over to for running online exams. Used for two UG1 courses (200 students) and one UG4 course (30 students) at present. It includes all the standard teaching applications, the same base environment and the same submission system the students are familiar with.

We also provide specialist facilities for exam preparation - a cluster of special locked down DICE machines for academic and administration staff preparation of examination scripts using LaTeX.

Scope

Teaching

Resources

Local/Central

No equivalent central service.

6.8. File service

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

The migration of 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, is largely complete. The intention is that all users will have migrated by the start of the next academic year.

The majority of this service is now hosted on Linux servers. The remaining Solaris NFS/AFS servers are scheduled to be decommissioned by the start of the next academic year.

We are currently, in collaboration with IS, introducing a GPFS based file service which will provide our users with a high performance file system better suited for parallel processing applications.

Scope

School

Resources

- 5 Solaris NFS/OpenAFS file servers (to be decommissioned within the year)
- 8 Linux OpenAFS file servers.
- 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 current central IS service that meets all these requirements and would be very unlikely to migrate to a central service that did not meet them.

We have expressed interest in IS's proposed new general purpose central file service for research data and are currently drawing up a list of our requirements for such a service for their consideration.

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.

6.9. Front line support

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

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 unit.

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). The team also manage a significant number of 'compute' servers within the School. After the move to the Forum, the support team have also been involved with the day-to-day support of the AV equipment.

Scope

School

Resources

60 servers including RT and devproj servers, 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 progressed. IS have created a CMS queue called sch-Informatics for us which we will start using for 'cross-boundary' issues (e.g. staffmail) as soon as training has been arranged.

This will, of course, be a short-term solution as IS are planning on replacing their current CMS system. Up to date we have been actively involved in providing information to assist with creating a requirements document and plan to be actively involved with this project.

6.10. Teaching support

We provide a fault tracking system for the Teaching Organisation. This is using RT, a separate instance of the same system as used by Frontline support. We also provide a web form which students and applying students can use to contact the ITO via this system. The ITO use it for all student enquiries. It allows proper tracking of incoming requests and a complete history of all correspondence. The ITO also use it in order to triage requests to Directors of Studies and for handling Special Circumstances so all the information is there for the Board of Examiners.

Scope

Teaching

Resources

Local/Central

No equivalent central service. We could, as being done for frontline support, consider the use of the central CMS (or any replacement) instead of RT. However its not clear whether the central CMS is properly geared towards the use by administration staff for student correspondence, however we have fed this back as a requirement for any replacement to the existing CMS.

6.11. 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, with no manual intervention.

Scope

School, IS, various CSE schools, research grants, external

Resources

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

Local/Central

Central IS provision is based on the Informatics LCFG 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 and various CSE schools.

6.12. 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 Scientific Linux 5.2.

Scope

School (over 1000 machines), IS, various CSE schools

Resources

A number of Linux servers for delivery of software packages.

Local/Central

Central IS provision based on Informatics service

Export

IS Desktop services provide a managed Linux platform to a number of CSE schools, based on the Informatics Linux platform.

6.13. Miscellaneous infrastructure services

A number of infrastructure services, including:

- NTP time servers: a reliable NTP time service is critical, as synchronised time is fundamental to the correct operation of the Kerberos authentication protocol
- *nagios*-based service monitoring, which enhances the reliability of the whole of the School's computing service by providing immediate notification of system failures
- console servers, which allow remote access to managed and self-managed servers' serial and IPMI consoles, to simplify remote control and diagnosis
- power distribution bars under network control, again simplifying remote management of servers
- UPS monitoring of both small per-server UPSes and large whole-building UPSes, providing advance warning of faults as well as clean shutdowns as necessary during power failures

Scope

School, UK wide (NTP)

Resources

A number of Linux servers

Local/Central

No central provision, or central provision not appropriate

Export

NTP service UK Wide.

6.14. 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 (MRBS)
- 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

We are currently re-investigating the School's requirements for a calendar service to see if they are compatible with the new eDiary service. A decision on any future internal calendar service will be made when this is completed.

With the demise of the eDiary service in its former form, it seems likely that we will continue to run our own room booking services (now migrated to the open source MRBS rather than the internally created Shehzu) for the foreseeable future.

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

6.15. Network

A managed network on 5 sites, with the following features:

- around 8000 live network ports (including 1300 VoIP phone ports)
- 182 network switches, of which 8 provide hardware wirespeed routing
- five EdLAN connections and two SRIF connections
- private inter-site dark-fibre links, where appropriate
- 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 re-installation.
- soft patching — no manual patching once switches installed
- configuration rights delegated to technical, support and computing staff as appropriate to ensure responsiveness to user requests
- VPN service

There is a 10Gbps backbone in the Informatics Forum, with a 10Gbps link from there to our network in Appleton Tower.

Scope

School; non-School users in Forrest Hill and 1–5 Buccleuch Place

Resources

- 182 HP network switches
- 54 wireless access points
- a number of Linux servers acting as firewalls, routers and providers of network services

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 has 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 provides 1000baseT 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 50 access points; these are part of the centrally managed wireless service.

We are currently providing the network service for users in the School's former space in Forrest Hill and 1–5 Buccleuch Place. We are hoping to move our remaining machines out of the BP server room quite soon now, and at that point the intention is to hand over management of the network there to IS. We expect to be in FH until the Appleton Tower basement refurbishment is complete.

Export

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

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

6.16. Printing

Provision of a multi-platform networked printing service.

The underlying technology is CUPS, managed via LCFG. Windows clients connect via Samba.

At the urging of the procurement office, we have replaced colour printers with MFDs in the Informatics Forum. To date these have not proved entirely satisfactory and we are currently reviewing their future use within the school

Scope

School

Resources

- 32 network printers - 23 mono A4, 5 mono A3, 4 colour A4.

- 7 Canon multifunction devices (printers/photocopiers/fax machines)

Local/Central

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

6.17. Procurement

A procurement service for computing hardware and software. Assist academic staff specify computing equipment including PCs and PC components and ensure that purchases are in accordance with University Procurement guidelines.

Maintain hardware and software inventories, including warranty and license information.

Scope

School

local/central

Purchases are made using the University's E-financials package in accordance with University Procurement guidelines. Purchase details are also stored locally and are key to our local inventory system.

6.18. Storage Area Network

Three separate Storage Area Networks (SANs), providing a total of 106TB of fibre attached storage to around 30 servers. The Forum SAN is configured to allow us to make use of multi-pathing for additional stability, and will be linked to the Appleton Tower SAN over private fibre for flexibility and redundancy.

Scope

School

Resources

- 9 QLogic fibre channel switches.
- 7 Nexsan storage arrays, providing 58TB of storage.
- 2 R/Evolution 2730 storage arrays providing 48TB of storage.

Local/Central

No central switching capability to allow servers to access remote devices.

6.19. 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.

It is planned to phase this platform out by Summer 2009.

Scope

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

Resources

Local/Central

No central provision

6.20. Online Practical Submission

We provide a simple command line tool for online practical submission system. This is used by students for submitting practical work and by markers to access practical work for marking. The same system is used for approximately half of our taught courses.

Scope

Teaching

Resources

Local/Central

The central equivalent is WebCT and none, or very few, of our academic staff currently make use of that for teaching materials. The reasons cited for not using WebCT are that it restricts access to course materials and none of our staff and students want that, there is a general lack of awareness and no idea what might be involved in using it, fallout from staff that used a previous version and then were told that all their course materials would have to be redone when moving to the next version and a lack of any compelling reason to use it. We believe that systems in general (and e-learning systems in particular) should be compliant to open standards and should be open source if possible so there is resistance to the use of WebCT/Blackboard. We also believe that open courseware is becoming an important recruitment (and visibility) channel and would like to see a VLE that supported this. We plan to investigate the use of some aspects of WebCT/Blackboard over the summer period since we may have need of coursework submission and gradbook functionality. We don't anticipate widespread adoption of WebCT amongst staff in the short term.

6.21. 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.

Approximately 200 specific software packages, some simple, some complicated, are involved. Bigger packages include Matlab (150 classroom licenses), SAS, Xilinx, Maple, Sicstus, Allegro Common Lisp, Haskell, NLTK, Ocaml, Webots.

We also provide some course specific teaching services. A web development environment for Tomcat/Java (UG4), Web-CAT Java (UG1), PostgreSQL database providing automatic per-student accounts and databases, also with associated web space for Web 2.0 work (UG3, UG4 and MSc), TFTP boot service for robotic robotic brain brick hardware (UG4).

Scope

Teaching

Resources

Local/Central

No equivalent central service. The large majority of software is specific to Informatics teaching. The central linux/LCFG platform is also SL5 based, so there is potential to slowly add some of our local teaching software to that environment and hence make more use of the central lab linux provision.

6.22. Videoconferencing

The school runs a browser based desktop videoconferencing service using Adobe Connect for a research group. While available to anyone in the School it is not as yet being widely used.

Scope

Research

Resources

A Windows 2003 server

Local/Central

No equivalent central service, but IS were involved in trialing.

6.23. Virtualisation service

A simple server virtualisation service running on VMware server 2 (on top of DICE servers). Virtual hosts are stored on the school's SAN to make it easy to migrate virtual guests between virtual hosts. Prime motivation is to reduce power and capital equipment costs.

This service is an interim solution pending university wide discussions on virtualisation technologies.

Scope

School

Resources

Currently three servers. Expected to expand to ten servers in summer 2009.

Local/Central

There is no equivalent service provided centrally for school use.

Export

6.24. Web service

A managed Web hosting service.

The main school Web site, www.inf.ed.ac.uk, currently uses internally developed 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 auto-generated as reports from the school database. CGIs are carefully controlled.

There are around 20 virtual web servers for School institutes, specific projects and interdisciplinary groups based on this technology.

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

After a wide reaching consultation with our user base on their future web needs, we have started deploying a new web services based on the open source Zope application server and Plone content management system. The initial intention was that this would provide a better mechanism for research groups within the School to create and manage their web presence. This technology is currently being used to deploy the websites for SICSA (the Scottish Informatics and Computer Science Alliance) and the commercial development arm of the School of Informatics.

In discussion with the central web project, it was agreed that a lack of certain features in the current Polopoly service, most notably a fine grained authorisation system allowing access from within and outwith the University, made it impractical to move the majority of the School's web content to Polopoly in the short to medium term. Given the need to move away from our old internally developed system, we have agreed to move 10 or so of our most visible and visited pages to Polopoly and begin the migration of our remaining content to the Zope/Plone service. initially, we are creating an 'institute in a box' which will allow the School's institutes to easily update and refresh the contents of their individual sites whilst presenting a more uniform appearance across the sites. Once this is complete, the issue of whether the rest of the School's web content should be moved to Zope/Plone or Polopoly will be reconsidered,

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. There are no plans to make any changes to this service at the moment.

Scope

School

Resources

A number of Linux servers

Local/Central

A small number of outward facing web pages will be served from the central Polopoly server. The remainder of the School's web presence will continue to be served from internal servers

in the short to medium term. The intention is that more and more pages will be moved to the Polopoly service as the introduction of new technology permits.

6.25. Wiki

A wiki service, based on twiki.

Scope

School

Resources

A linux server.

Local/Central

Ideally, all Informatics wiki content would be stored on the central wiki service. This is not feasible at the moment since the Informatics' wiki service has been modified and extended to provide finer grained authorisation, managed via the School's roles and capabilities mechanism, and access to external collaborators via the School's iFriend system which many of the School's wiki pages depend on and which, for the moment at least, the central service cannot provide. We are therefore required to continue running our own service for those users who require these features. Users who do not fall into this category are directed to the central wiki service.

7. Staff profile

The school employs 15 Computing Officers (14.6 FTE) and 7 Computing Support Officers (6.6 FTE).

- 4 FTE Grade 9
- 7.6 FTE Grade 8
- 3 FTE Grade 7
- 3.6 FTE Grade 6
- 3 FTE Grade 5

For comparison, the 2008 figures were 20 Computing officers (19.0 FTE) and 7 Computing support officers (6.4 FTE). This very significant reduction in effort is having a seriously detrimental effect on our ability to provide the wide spectrum of services that our users expect.

8. 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.

9. Allocation Policies

Academic teaching staff are generally entitled to a DICE desktop and a laptop. Research staff are provided for by their associated research grant. Research postgraduates are entitled to a DICE desktop (a 3 year old desktop in their first year, and a new desktop in their second year).

The school policy is for a staff desktop writedown of 3 years and student desktop writedown of 4 years.

10. Resources

10.1. Figures

- Managed DICE (Linux) desktops
  - 480 staff and research postgraduates

  - 260 student labs

- Managed Windows desktops

  - 75 MDP managed desktops for administrative staff

- Self-managed

  - Around 200 Linux/Windows/MacOS desktops/laptops

- Managed DICE (Linux) servers

  - 120 infrastructure and user services servers
  - 40 research grant funded compute servers
  - 58 beowulf nodes (2 clusters)

- Managed Solaris servers

  - 7 NFS/AFS file servers
  - 2 multiuser login servers
  - 1 backup server

- Accommodation

  - 3 air-conditioned machine rooms, with a total of around 160 square meters.
  - 7 undergraduate teaching labs and 2 tutorial rooms, with a total of around 260 desktops

A. Active projects for (February - September 2009)

A.1. Infrastructure for Appleton Tower Refurbishment

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

A.2. Decommission of FH and BP

This project will decommission the FH and BP server rooms.

A.3. Rehousing at KB

This project is concerned with relocating servers from the temporary KB server room to the new KB college server room.

A.4. Moving kit from old sites to Forum

This project is concerned with the relocation of all active servers from the machine rooms at KB, BP and FH to the server rooms in the Forum, Appleton tower amd Kings Buildings.

A.5. Develop strategy for laptop support

Develop strategy for addressing the increasing move to portable and/or personal machines, and how this impacts our commodity computing provision.

A.6. Improve communication with users

Improve communication with users, including awareness of available services, communication of service availability, restructured and updated documentation, and increased exposure of COs to end users.

A.7. Informatics Commercialisation

Zope/Plone development work for the Informatics Commercialization and Ventures web sites. Create/extend a number of Plone modules and page templates to support these web sites.

A.8. EUCLID interoperability

Once EUCLID goes online it will replace some of our local database information pertaining to students. However we drive a lot of local processes from this information - such as account management. Some mechanism is needed to preserve the functionality of our local processes once EUCLID is in place.

A.9. EUCLID Informatics enhancement

EUCLID won't do everything we need. Some critical local processes will not be supported. This project encapsulates all of them although some could potentially be forked off into separate projects in the future.

A.10. SICSA Database

A SICSA (Scottish Informatics and Computer Science Alliance) administration support database. The IGS are managing the student data for current SICSA PhD students (even those not at Edinburgh meaning that such information will not be held in EUCLID).

A.11. Future direction for the main Informatics web service

Determine what changes/improvements, if any, should be made to the main Informatics web service.

A.12. Transition of Informatics staff mail services to Staffmail

This project will complete the migration of Informatics staff mail services to Staffmail.

A.13. Final stages of commodity CMS work

Though we now have a working commodity CMS based on Zope/Plone in service, there are still some aspects of the service to be finalised.

A.14. 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.

A.15. AV facilities

The installation of the AV facilities within the Forum needs to be completed, including providing appropriate documentation and training.

A.16. Forum Information Display Screens

Make operational the eight 'information display screens' currently installed at various locations within the Forum.

A.17. Revisit Account Management (Prometheus)

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

A.18. Develop power management solution for DICE desktops

The project will implement a component to manage sleep on DICE machines. It will coexist with Condor. This should save energy.

A.19. Improved RPM submission tool

Redevelop the existing RPM submission technology to use AFS instead of NFS (for improved system integrity).

A.20. DICE under emulators

Support DICE under VMware for desktops (particularly Windows and MacOS).

A.21. Simple server virtualisation

We currently minimize the number of services running on any one server to reduce the dependencies when we perform the annual DICE upgrade. The downside to this is that we have a larger number of servers to maintain with resulting financial and environmental costs. This project will provide a simple server virtualisation solution.

A.22. 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.

A.23. Desktop Virtualization

Make an easily user managed VirtualBox configuration, to be deployed on all desktops so that staff and students can easily experiment with alternative operating systems.

A.24. Collaborative LaTeX System

The Collaborative LaTeX system allows end users to create and manage their own subversion repositories for LaTeX documents (primarily ACLs) with support for external collaborators (via iFriend). The backend manages the LaTeX build process, also creating change documents and can mark up merge conflicts. Currently a prototype service this project is to (re)develop the existing prototype service into a full production service.

A.25. 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 computationaly intensive tasks.

A.26. Infrastructure for New Informatics Forum

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

A.27. Cluster Parallel filesystem

A parallel filesystem (GPFS) for use on the school's clusters which integrates with the university ECDF cluster.

A.28. AFS file system

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

A.29. Progress TiBS backup system to full service

The project covers the full integration of the TiBS backup system into DICE.

A.30. Replacement of LPRng with CUPS

This project will complete the move from LPRng to CUPS.

A.31. Upgrade Twiki

This project concerns the upgrading of the School's Twiki installation to a more modern version.

A.32. Logging policy and centralised logging

The evaluation phase of this project will look at logging policy, and the options available for the provision of one or more loghosts.

A.33. DICE client LDAP configuration

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

A.34. Proof General Eclipse Improvements

Further improvements to the Proof General Eclipse system.

A.35. Publishing and Discussion Media Survey Paper

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

A.36. OpenVPN service

Convert existing prototpye OpenVPN service into a full service.

B. Completed projects (October 2007 to March 2009)

B.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.

B.2. Production Cosign Service

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

B.3. 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. Also produced a new www.lcfg.org site with wiki and bugzilla service.

B.4. Web Technologies and Policy

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

B.5. 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.

B.6. Inventory System

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

B.7. OpenLDAP Replication and Server Configuration

Development and deployment of replicated openldap servers.

B.8. 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.

B.9. Future Console Servers

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

B.10. LCFG SL5 port (inf level)

Develop, with EPCC, 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.

B.11. 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.

B.12. DICE SL5 Services

Coordinate the transition to Scientific Linux 5 of all DICE services currently running on FC5/FC6 machines.