

Informatics Computing Plan 2013

School of Informatics

August 6, 2013

1 Long-term vision and strategic objectives

Informatics Computing serves over 300 staff (215 teaching and research), nearly 300 research students, over 200 taught postgraduate students, around 700 undergraduates, and over 100 visitors and associates.

The aim of the Informatics Computing Service is to ensure that members of the School of Informatics (staff, students and visitors) receive computing services necessary for their research, teaching and knowledge transfer activities. These services should be efficient, state-of-the-art, fit to users' requirements, good value for money and use open standards. Appendix A outlines the evaluation processes that we have established to ensure that we are fulfilling this aim.

Strategic objectives

We have three principal aims underpinning the Informatics Computing Strategy:

- S1** Maintenance, review and update of state-of-the-art computing environment for all members of the School.
- S2** Maintaining an optimum level of interoperability of Informatics Computing with College and IS services.
- S3** Engagement with international best practice.

We have specific objectives relating to the computing infrastructure and to the activities of the School: research, teaching and knowledge transfer :-

Infrastructure We are committed to providing an infrastructure that ensures that members of the School get those services that they need. These services may be provided by the School, by IS or by external organisations.

- I1** Review and evaluate computing infrastructure change taking account of changing user needs and general computing trends.
- I2** Development of new services.
- I3** Provision of Informatics know-how and technologies to college and university level, and beyond.

Research In addition to providing a flexible, responsive environment for research in the School, we must meet the specific research requirements across our research institutes, and structure research computing support to be well-matched to the ways researchers propose and carry out research projects.

- R1** Continued development of lightweight, responsive support for research computing that is fully compatible with full economic costing of research
- R2** Ensuring that Informatics users get efficient, responsive access to high performance research computing and storage facilities
- R3** Provision of support for interdisciplinary and collaborative research projects (eg SICSA, InSpace, Digital Health).

R4 Development of prototype services from R&D projects (eg Data Intensive Research machine)

Teaching In addition to providing a stable environment for the School's teaching activities, we shall develop appropriate instrumentation in the teaching environment to support research/teaching synergy.

T1 Support research-led teaching by providing support for the transfer of research tools to our standard teaching platform.

T2 Ensuring that our infrastructure complies with open e-learning standards where possible.

T3 Support appropriate assessment of students (eg online examinations).

Commercialization and knowledge transfer Informatics Computing can support the School's knowledge transfer activities by providing a bridge between research and use.

C1 Using the School's commercialization infrastructure as a driver to develop prototype services from applied research in Informatics.

Management Information We shall support the ISS business processes. We also aim to support planning and decision making through the timely and effective maintenance and provision of Management Information.

2 Report on 2012

Goals

1. **Goal** Review of Computing team role and structure
Progress Head of Computing having lightweight discussions with each individual CO to identify any common concerns. A formal review will only be instigated should these interviews identify a common concern that could only be fixed by a major team restructure.
2. **Goal** Continued consideration of appropriate use of central data storage facilities
Progress Ongoing. A project to investigate providing AFS over Eddie file-storage has been resourced for 2013.
3. **Goal** Engage in requirements capture for and design of proposed central archiving service
Progress Ongoing. A project, to capture the School's requirements and engage with the design of the central service, has been resourced for 2013.
4. **Goal** Continue engagement with shared timetabling project
Progress Engaged, as required, with phase 1
5. **Goal** Engage with the PURE project to meet identified requirements for knowledge management functions other than those related to teaching administration (eg research grant management)
Progress Engaged.
6. **Goal** Consider the most effective usage of our share of ECDF
Progress We are investigating ECDF for our hadoop requirements. We believe that we have the best recovery rates in the College.
7. **Goal** Review impact of University activities wrt. teaching - timetabling, Personal Tutors, VLEs, Distance Learning, EUCLID developments
Progress Timetabling - yes, personaltutors - yes, VLE - various discussions, Distance Learning - ?, EUCLID developments - ? - none.
8. **Goal** Continue with work on providing a virtualised DICE for use on students' personal machines
Progress A project has been resourced, for 2013, to develop this.
9. **Goal** Complete re-factoring of School Database (database engine and client), and complete move of data and users from old system to new.
Progress Re-factoring of user-facing aspects of database engine and client now complete. Re-factoring of back-end in process of being completed. Still some data to move from old (2G) to new (3G) land.
10. **Goal** Further consider how best to maximise benefit of new School Database by reviewing which additional, often standalone, services can be brought into or better integrated with the School Database.
Progress Decided to defer until 2013/2014 in order to relieve pressure on development staff and consolidate on what we have done so far.
11. **Goal** Identify enhancements to the School database to better support tutors and demonstrators
Progress Identified and implemented (HTBN).
12. **Goal** Produce drupal service.
Progress Began looking at, but decided to defer pending IS deliberations over central

- drupal service
13. **Goal** Windows 7 upgrade for admin staff desktops
Progress Completed (just). Problems with AFS interaction with ancient Kaspersky anti-virus suite delayed progress.
 14. **Goal** Extend use of desktop virtualisation by encouraging users to adopt the technology - and facilitate this.
Progress Have been encouraging people by not giving them second physical machines.
 15. **Goal** Complete migration of user documentation to new site.
Progress Completed. New user documentation site entered service 25th February 2013.
 16. **Goal** Improve quality and validity of user documentation
Progress Completed. Editor appointed to keep up-to-date.
 17. **Goal** Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation etc).
Progress Little progress.
 18. **Goal** Consider how to improve access to School services from mobile devices
Progress No progress.
 19. **Goal** Improve security of services, particularly against external attacks
Progress Developed improved intrusion monitoring and detection. Further potential improvements have been identified.
 20. **Goal** Code improvement on LCFG core client side
Progress Project just started for this year.
 21. **Goal** Migrate from VMware service to KVM service
Progress Completed. VMware service now closed.
 22. **Goal** Continue development work to take advantage of new account management framework (eg implement account life-cycle)
Progress Deferred due to a critical member of staff having left.
 23. **Goal** Upgrade DICE servers to Scientific Linux 6 (or other RHEL6 derivative)
Progress Completed.
 24. **Goal** Accelerate replacement of physical servers with virtualised servers to reduce energy consumption and meet reduced capital budget
Progress Getting close to ideal mix of real/virtualised servers for infrastructure services. Now need to address research servers.
 25. **Goal** Implement remaining improvements identified as result of review of resilience to disasters - off-site DR for School DB
Progress Project in progress.
 26. **Recurring Goal** Consideration of our existing commitments given reduced salary and non-salary budgets.
Progress Computing Strategy Group working on identifying commitments that can be dropped or transferred to IS.
 27. **Recurring Goal** Further promote School developed solutions to the rest of the University and beyond
Progress Theon User group... Many talks at UKUUG/FLOSS conference and OpenAFS talk.
 28. **Recurring Goal** Further improve communication between users and computing staff
Progress No noticeable progress.
 29. **Recurring Goal** Ring-fencing 5% of individual computing staff's time for staff development

Progress This is being achieved by all units, and being exceeded by some.

30. **Recurring Goal** Consideration of ways to minimise our energy footprint

Progress Improved DICE desktop sleeping system. Looking at removing an unnecessary UPS system within the Forum.

Activities to be considered for de-prioritisation

Goal School Beowulf cluster

Progress Actively being considered. Gridengine service almost certain to close, but may need to continue a cut-down service for Hadoop provision.

Collaboration with others

We are very keen to collaborate with other CSE schools on development and even service delivery.

1. We continue to provide the base LCFG Linux platform to other schools (via IS).
2. A number of CSE schools are deploying aspects of Theon, with our assistance. Theon user group has been formed.
3. We co-hosted the FLOSS UKUUG 2012 conference.
4. We hosted the OpenAFS 2012 conference.

3 Revised plan for 2013

Goals

1. Continue review of Computing team structure and implement any required changes
Who: School, Cost: 2w
2. Continued consideration of appropriate use of central data storage facilities, specifically investigate AFS over Eddie storage.
Who: Research, Cost: 3w
3. Engage in requirements capture for and design of proposed central archiving service
Who: Research, Cost: 2w
4. Continue engagement with shared timetabling project (with personal timetables)
Who: Admin, Cost: 4w
5. Continue engagement with the PURE project to meet identified requirements for knowledge management functions other than those related to teaching administration (eg research grant management)
Who: Admin, Cost: 1w
6. Await invitation to discuss future business model for ECDF funding. Continue investigations with using ECDF for Hadoop requirements
Who: Research/Teaching/Energy, Cost: 1w
7. Review impact of University activities wrt. teaching - timetabling, VLEs, Distance Learning (including MOOCS) , EUCLID developments. *Who: Teaching, Cost: ?*
8. Debate future of teaching platform
Who: Teaching, Cost: ?
9. Consider how to cope with increasing student numbers, particularly with respect to on-line exams *Who: Teaching, Cost:*
10. Continue with work on providing a virtualised DICE for use on students' personal machines *Who: Teaching, Cost:*
11. Complete the re-factoring of School Database back-end (database engine and client), and complete the move of data and users from old system to new (particularly HR).
Who: Admin, Cost: ??
12. Further develop data feed integration (from central services) - eg maximise use of HR feed and take first feed from SAT. *Who: Admin, Cost: ??*
13. Further consider how best to maximise benefit of new School Database by reviewing which additional, often standalone, services can be brought into or better integrated with the School Database.
Who: Admin, Cost: 1w
14. Engage with design of central IS drupal service, or produce own drupal service.
Who: School, Cost: 10w??
15. Survey user requirements for virtualisation/cloud provision - including consideration of replacement of research servers with virtualised servers to reduce energy consumption and encourage resource sharing.
Who: Research, Cost: ?
16. Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation etc).
Who: Research, Cost: 2w
17. Consider how to improve access to School services from mobile devices

- Who: ?, Cost: ?*
18. Review DICE authentication technology, including possibility of outsourcing to EASE
Who: ?, Cost: ?
 19. Implement improvements to security of external facing services as identified in 2012.
Who: School, Cost: 4w?
 20. Code improvement on LCFG core client side
Who: Infrastructure, Cost: 4w
 21. Continue development work to take advantage of new account management framework (eg implement account life-cycle)
Who: Infrastructure, Cost: ??
 22. Implement remaining improvements identified as result of review of resilience to disasters - off-site DR for School DB
Who: Admin, Cost: ?
 23. Redesign and implement inventory system - current system does not cater well for self-managed machines
Who: Admin, Cost: ?
 24. Consider infrastructure requirements for refurbished AT

Recurring goals

1. Consideration of our existing commitments given reduced salary and non-salary budgets.
2. Further promote School developed solutions to the rest of the University and beyond
3. Further improve communication between users and computing staff
4. Ring-fencing 5% of individual computing staff's time for staff development
5. Consideration of ways to minimise our energy footprint, eg identifying under-used research servers

Activities to be considered for de-prioritisation

- School Beowulf cluster
- Authentication

Collaboration with others

We are very keen to collaborate with other CSE schools on development and even service delivery.

1. We shall continue to collaborate closely with other Schools deploying our LCFG technology.
2. We shall assist other schools in deploying our Theon database technology.
3. We shall work with IS on investigating the feasibility of providing AFS on top of Eddie storage.
4. We shall work with IS on investigating the feasibility of providing Hadoop on Eddie.

What we would like of IS

- Assistance with investigating the feasibility of hosting our online exams in the central IS labs.
- Provision for data archiving and, perhaps, curation. Note that this archiving should not be limited to research data.
- Audio Visual support from LTSTS for non Informatics events in the Forum
- We are keen to make effective use of our share of the ECDF service with a view to deprecating our remaining Beowulf cluster. We are particularly interested in the possibility of running Hadoop on the ECDF cluster. We are also interested in the proposed ECDF Cloud service.
- A barrier to further migration of services to "equivalent" central IS services is the tight integration of our services with our authorisation roles service which is fed from our School Database - allowing us, for example, to create mailing lists and subversion repositories for individual tutorial groups. The ability to feed into Grouper, from our School Database, might reduce this barrier.
- Our user documentation pages refer to various pages on the IS documentation pages. The URLs for the IS pages appear to change on a frequent basis - it is very frustrating to have to continually fix-up references to these pages.
- We are keeping a close eye on the upgrade of the central wiki service with respect to the reported limitation of WYSIWYG editing only in the latest version of the Confluence software.

4 Plan for 2014

Goals

1. Act on deliberations on dealing with increasing student numbers
2. Act, if appropriate, on review of teaching platform
3. Act on survey of user requirements for virtualisation/cloud provision
4. Act, if appropriate, on review of authentication
5. Capture requirements for CDTs (Centres for Doctoral Training) administration
6. Plan for likely AT decant and various other estate developments
7. Improve availability and uptake of video capture
8. Further consideration of migration to central services (big ticket items only)
9. Further promote School developed solutions to the rest of the University and beyond
10. Further improve communication between users and computing staff
11. Ring-fencing 5% of individual computing staff's time for staff development.
12. Further consider how best to maximise benefit of new School Database by reviewing which additional, often standalone, services can be brought into or better integrated with the School Database.

De-prioritised areas

Plone (CMS) service?

A Evaluation

We have established a number of evaluation processes, to ensure that we are delivering a service in line with our strategic objectives.

- **Fit to requirements** User requirements are captured using various mechanisms. Teaching requirements are met through a stable and well established system for the collection, negotiation and delivery of computing requirements. We have implemented a newer mechanism to capture research computing needs, based on a basic level of recharge per researcher, in return for which certain services (eg disk space, network connectivity, cluster computing usage) are provided. Specific requirements are also captured in depth via focused innovation meetings, to which all members of the school may attend.
- **Value for money** This is a criterion for the annual review document, and is related to transparent support for research computing, centralised procurement that remains close to academic needs, and official audits of various research project expenditure.
- **Objective evaluation** Each unit provides a triannual report, which includes proportions of staff time spent on various activities, projects undertaken, etc. This data is used to inform strategy, and management: for example, consistently lower proportions of time spent on development activities (due to operational demands) than planned can be identified, and emphases changed.

B Staffing and Resources

The school employs 20 computing staff (19.4 FTE).

There are 655 managed DICE (Linux) desktops; 345 personal machines for staff and research students, and 310 in student labs (7 undergraduate teaching labs and 2 tutorial rooms). There are a further 70 managed Windows desktops for administrative staff.

In addition there are several hundred self-managed Linux, Mac OS and Windows desktops and laptops.

There are 270 managed DICE (Linux) servers (175 physical, 95 virtual) and a further 24 beowulf compute nodes. Our servers are housed in 3 air-conditioned machine rooms, with a total area of around 160 m².

C College, University, External Relationships

The School has a high degree of interaction and engagement at the College and University level, arising in particular from the expertise within the School. We are engaged with university committees concerned with authentication, security, and information architecture, for example, and play a leading role in envisioning the development of computing at a university level. Externally, our computing staff interact with organizations such as Usenix and UKUUG through workshops, conferences and tutorials.

D Categories and activities

Support	Front Line	Documentation	General Web	Online Exams	Procurement
	User accounts	AV	Inventory	Exam prep	Corpora
File	Web	Admin	Misc	Desktop	Compute
AFS	Web (CVS)	Theon	ssh servers	DICE desktop	DICE compute
Samba (admin)	CMS plone	ISS desktop	Printing	Teaching s/w	Grant servers
Backups (user)	Wiki	IGS desktop	Mailing lists	Research s/w	Beowulf
iFile	Group web	HR desktop	Licence servers	MDP	Gridengine
SVN	Homepages	ISS RT	Postgresql		Hadoop
	I in a Box	RBS	Coltex		
	Legacy		Inspace		
	Tomcat		DICE labs		
Infrastructure	Linux platform	LCFG	Kerberos	Directory	Network
	S/W repository	pkgforge	Cosign	Authorisation	Network services
	Server rooms	Consoles	Monitoring	Prometheus	System backups
	Virtualisation				