

Informatics Computing Plan 2014

School of Informatics

February 25, 2014

1 Long-term vision and strategic objectives

Informatics Computing serves over 300 staff (240 teaching and research), nearly 300 research students, over 150 taught postgraduate students, around 850 undergraduates, and over 100 visitors and associates.

The aim of the Informatics Computing staff 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, 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 five principal aims underpinning the Informatics Computing Strategy:

- S1** Maintenance, review and update of a computing environment fit for the purposes of 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.
- S4** Provision of expertise to support the teaching and research activities of the School.
- S5** Providing added value over services offered by College and IS.

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, Farr Institute, 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

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

T2 Support appropriate assessment of students (eg online examinations).

T3 Provision of expertise to support teaching activities

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.

Interaction with IS

We shall focus :-

- on being early adopters of services that may or may not become commodity
- on developing new services that are specific to, or inspired by, our environment

We shall use IS services wherever possible, unless there are sound academic reasons for not doing so. However, we shall take a careful approach when considering migration from a School service to the equivalent IS service.

2 Report on 2013

Goals

1. **Goal** Continue review of Computing team structure and implement any required changes
Progress Completed interviews with individual COs. Currently discussing, with likely changes and a reworking of the project management framework.
2. **Goal** Continued consideration of appropriate use of central data storage facilities, specifically investigate AFS over Eddie storage
Progress Have an active project to try AFS on top of GPFS, but blocked on availability of IS effort to collaborate on this. IS data storage project was delayed.
3. **Goal** Engage in requirements capture for and design of proposed central archiving service
Progress Nothing new from centre as they have been concentrating on storage for live data. Have a active project to capture our requirements.
4. **Goal** Continue engagement with shared timetabling project (with personal timetables)
Progress Engaged. Modified RBS system and have active project for displaying bookings in foyer.
5. **Goal** 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)
Progress A lot of activity on PURE related to REF. Also non-REF - staff profiles, creating landing pages, updating institute pages, closing down informatics research repository with research reports and publications migrated to PURE.
6. **Goal** Await invitation to discuss future business model for ECDF funding. Continue investigations with using ECDF for Hadoop requirements
Progress Still awaiting invitation. About to look at running MapReduce on ECDF.
7. **Goal** Review impact of University activities wrt. teaching - timetabling, VLEs, Distance Learning (including MOOCS) , EUCLID developments.
Progress Yes - timetabling, distance learning (AI Planning MOOC). No driver to make use of VLE at this time.
8. **Goal** Debate future of teaching platform
Progress Have debated. Decided that no one solution is best, so will develop a portfolio of solutions - existing labs, virtualised DICE, instructions for s/w installs on own machines.
9. **Goal** Consider how to cope with increasing student numbers, particularly with respect to on-line exams
Progress Ditto from above point. No solution yet for online exams. Would still like to investigate a solution making use of the IS central labs.
10. **Goal** Continue with work on providing a virtualised DICE for use on students' personal machines
Progress This has just shipped. Too soon to tell gauge success.
11. **Goal** 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).
Progress Active project to re-factor. Awaiting admin staff effort to determine requirements for moving other data to new system.
12. **Goal** Further develop data feed integration (from central services) - eg maximise use of

HR feed and take first feed from SAT.

Progress Use new HR feed, eg for internal contracts. Almost have a SAT feed - waiting on IS.

13. **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 Not this year. Consolidation is priority for this year.

14. **Goal** Engage with design of central IS drupal service, or produce own drupal service.

Progress Attempting engagement, but only recently with some success. Central service will be delivered much later than we had expected - at some point we may have to consider developing our own service.

15. **Goal** 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.

Progress Survey completed. Currently analysing responses.

16. **Goal** Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation etc).

Progress Done a bit as part of documentation project. Still need to work more on this.

17. **Goal** Consider how to improve access to School services from mobile devices

Progress Identified ways in which we can improve access - printing, AV, lab availability, school web accessibility, OpenVPN, VDI.

18. **Goal** Review DICE authentication technology, including possibility of outsourcing to EASE

Progress Active project to look at this. Only possible once Ed kerberos upgraded and facility exists to automate creation of host key credentials.

19. **Goal** Implement improvements to security of external facing services as identified in 2012.

Progress Active project to work on this.

20. **Goal** Code improvement on LCFG core client side

Progress Project has completed which has improved quality of code.

21. **Goal** Continue development work to take advantage of new account management framework (eg implement account life-cycle)

Progress Ongoing project working on this.

22. **Goal** Implement remaining improvements identified as result of review of resilience to disasters - off-site DR for School DB

Progress Done. Live and recovery process documented.

23. **Goal** Redesign and implement inventory system - current system does not cater well for self-managed machines

Progress Active project. Requirements captured. Design discussed and prototype written. Implementation will finish this year (hopefully)

24. **Goal** Consider infrastructure requirements for refurbished AT

Progress Pushed back as start of AT refurbishment has been delayed.

25. **Recurring Goal** Consideration of our existing commitments given reduced salary and non-salary budgets.

Progress None.

26. **Recurring Goal** Further promote School developed solutions to the rest of the University and beyond

Progress Theon, AFS in ECDF.

27. **Recurring Goal** Further improve communication between users and computing staff

Progress Going to attend CompSoc. COs to attend student demonstrator help-desk

28. **Recurring Goal** Ring-fencing 5% of individual computing staff's time for staff development

Progress Achieved 8% on average for computing officers, but very little for user support staff.

29. **Recurring Goal** Consideration of ways to minimise our energy footprint, eg identifying under-used research servers

Progress Sleep research servers? Research grant closure process + annual review will identify little used research servers and perhaps virtualise them.

Activities to be considered for de-prioritisation

- **Goal** School Beowulf cluster

Progress Decided to stop Gridengine service - just leaving Hadoop service. But we have concerns over Eddie's future.

- **Goal** Authentication

Progress Active project - see item above.

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 looking at deploying aspects of Theon, with our assistance.
3. We have been trying to collaborate with IS on investigating AFS on top of the new Research Data platform, but effort availability (on both sides) has hampered this.

3 Revised plan for 2014

Goals

1. Implement any required changes as a result of review of Computing team structure
Who: School, Cost: ?
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 (But see "What we would like of IS", below)
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. Further develop virtualised DICE for use on students' personal machines, if required.
Who: Teaching, Cost: ?
9. Consider how the School's computing staff could contribute to teaching activities
Who: Teaching, Cost: ?
10. 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: ??
11. Further develop data feed integration (from central services) - eg maximise use of HR feed and take first feed from SAT.
Who: Admin, Cost: 2w min
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.
Who: Admin, Cost: 1w
13. Capture requirements for CDTs (Centres for Doctoral Training) and implement.
Who: Admin, Cost:
14. Adopt central space management system and drop our internal room booking system
Who: Admin, Cost:
15. Engage with design of central IS drupal service, or produce own drupal service.
Who: School, Cost: 10w??
16. Port of LCFG to Scientific Linux 7 (or other RHEL7 derivative)
Who: Teaching/Research, Cost: 10w
17. Upgrade DICE desktops and servers to Scientific Linux 7 (or other RHEL7 derivative)

- Who: Teaching/Research, Cost: 10w*
18. Act on results of virtualisation/cloud provision survey, as appropriate.
Who: Research, Cost: ?
 19. Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation etc).
Who: Research, Cost: 2w
 20. Improve access to School services from mobile devices (eg printing, AV, lab availability, school web accessibility, OpenVPN, VDI)
Who: Teaching/Research, Cost: ?
 21. Decide on whether to outsource DICE authentication to EASE, and produce a migration plan if necessary
Who: ?, Cost: ?
 22. Continue development work to take advantage of new account management framework (eg implement account life-cycle)
Who: Infrastructure, Cost: ??
 23. Complete redevelopment of new equipment inventory system
Who: Admin, Cost: ?
 24. Consider infrastructure requirements for refurbished Appleton Tower and plan for 2015 decant
Who: , Cost: ?
 25. Revamp ISS/IGS web presence (for current students) *Who: Admin, Cost: ?*
 26. Review energy usage of research servers - perhaps sleeping idle servers and virtualising little used servers. *Who: , Cost: ?*
 27. Perform an audit of all research data within the School
Who: School, Cost: ?
 28. Produce a register of medium-high risk data and a mechanism for users to self populate the register
Who: School, Cost: ?
 29. Implement improvements to security of web services
Who: School, Cost: 4w?
 30. Planning to improve awareness of medium-high risk data at induction of new staff
Who: School, Cost: ?

Recurring goals

1. Aim for a minimum of 20% of development time to be dedicated to user submitted projects
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, including user support staff.
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
- Room booking system

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

- Provision for data archiving and, perhaps, curation. Note that this archiving should not be limited to research data.
- 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.
- With respect to the SAT (Shared Academic Timetabling) project, Informatics' position remains that our full engagement with this project is dependent on Student Allocator functionality. We gather that consideration is being given to how this missing functionality can be provided - meanwhile it is likely that our use of the minimal ECP software will continue as is.
- Provide support for S/MIME and PGP encryption and signing of email, including institutional key- and certificate-signing
- Longer term (2015), renewed assistance with investigating the feasibility of hosting our online exams in the central IS labs.
- 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. We welcome the invitation to be involved in the Grouper enhancement project.

v

4 Plan for 2015

Goals

1. Appleton Tower decant
2. Consider use of students' own laptops in exams
3. Consider online exams in IS public labs (perhaps using virtualisation)
4. Implement any ideas arising from consideration of computing staff involvement in teaching
5. Migrate existing web content off Plone CMS service to IS (or School) Drupal service
6. Further consideration of migration to central services (big ticket items only)
7. Further promote School developed solutions to the rest of the University and beyond
8. Further improve communication between users and computing staff
9. Ring-fencing 5% of individual computing staff's time for staff development.
10. 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 (19.6 FTE).

There are 685 managed DICE (Linux) desktops; 375 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 255 managed DICE (Linux) servers (150 physical, 105 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 User accounts	Documentation AV	General Web Inventory	Online Exams Exam prep	Procurement Corpora
File	AFS Samba (admin) Backups (user) iFile SVN	Admin	Misc	Desktop	Compute
	Web (CVS) CMS plone Wiki Group web Homepages I in a Box Legacy Blog	Theon ISS desktop IGS desktop HR desktop ISS RT RBS	ssh servers Printing Mailing lists Licence servers Postgresql Coltex DICE labs	DICE desktop Teaching s/w Research s/w MDP Virtual DICE desktop	DICE compute Grant servers Beowulf Gridengine Hadoop
Infrastructure	Linux platform S/W repository Server rooms	LCFG pkgforge Consoles	Kerberos Cosign Monitoring	Directory Authorisation Prometheus	Network Network services System backups
	Virtualisation				