

Informatics Computing Plan 2016

School of Informatics

January 18, 2016

1 Long-term vision and strategic objectives

Informatics Computing serves over 300 staff (230 teaching and research), 330 research students, 350 taught postgraduate students, 950 undergraduates, and over 300 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

Commercialisation 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 commercialisation 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 2015

Mandatory Goals

1. **Goal** Appleton Tower decant to Forrest Hill and Wilkie
Progress Completed.
2. **Goal** Consider infrastructure requirements for refurbished Appleton Tower and plan for return from Forrest Hill and Wilkie
Progress Ongoing. Dependent on decision on whether to expand teaching provision in Appleton Tower.
3. **Goal** Complete DICE SL7 desktop platform
Progress Completed.
4. **Goal** Upgrade DICE desktops to Scientific Linux 7 (or other RHEL7 derivative)
Progress All student lab desktops were upgraded by start of academic year 15/16. In process of upgrading staff desktops.
5. **Goal** Deploy CDT dedicated computing clusters
Progress Completed.
6. **Goal** Web review and revamp (including migration of polopoly pages to IS Drupal service)
Progress Migration from polopoly to IS Drupal service completed. School intranet partially revamped.
7. **Goal** Produce a register of medium-high risk data and a mechanism for users to self populate the register
Progress Ongoing. Delayed largely due to desire to develop a register in a collegiate manner with other CSE schools.
8. **Goal** Identify and implement modifications to Theon required by the taught assessment regulation changes and the introduction of progression boards.
Progress Completed.
9. **Goal** Implement College security action plan
Progress Ongoing. MDP Windows desktops in process of being encrypted. New Scientific Linux 7 DICE installs now have certain key partitions encrypted. Where possible, new portable equipment is encrypted before handover to end-user.

Goals

1. **Goal** Continued consideration of appropriate use of central data storage facilities, specifically investigate AFS over DataStore.
Progress Investigation almost complete.
2. **Goal** Engage in requirements capture for and design of proposed central archiving service
Progress Still awaiting call for requirements capture from IS.
3. **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 Little progress on this. Considering using PURE for academic staff primary homepages.
4. **Goal** Continue investigations with using ECDF for Hadoop requirements

- Progress** Little progress as effort was diverted to work on deploying the CDT compute clusters. There are now doubts as to the suitability of ECDF for lab based Hadoop teaching. Considering a local dedicated Hadoop cluster.
5. **Goal** Engage in requirements capture for central media services, particularly with respect to similar in-house services
Progress Completed. Will now consider how to make use of Media Hopper.
 6. **Goal** Review impact of University activities wrt. teaching - timetabling, VLEs, Distance Learning (including MOOCS) , EUCLID developments.
Progress Tracking assessment EUCLID changes and PGR milestone tracking changes.
 7. **Goal** Further develop virtualised DICE for use on students' personal machines, if required. For example, more frequent updates?
Progress No progress this year. Concluded that we should assess take-up before developing further.
 8. **Goal** Consider how the School's computing staff could contribute to teaching activities
Progress Identified the third year system design exercise as most effective way to contribute.
 9. **Goal** Consider online exams in IS public labs (perhaps using virtualisation)
Progress Determined that no one IS public lab is sufficiently large to hold our largest class, so no benefit in investigating using IS public labs for online exams. We are now splitting large exams into two sessions, using our in-house online exam infrastructure.
 10. **Goal** Complete the re-factoring of School Database back-end (database engine and client)
Progress Significant progress has been made and work should complete in 2016.
 11. **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. Specifically - UG projects DB, Integration with RT, Reform.
Progress Have identified many requirements, but re-factoring work (see above goal) has been of higher priority.
 12. **Goal** Deploy an instantiation of IS Drupal distribution to replace the existing School Plone CMS
Progress Have several instantiations of EdWeb (eg ISS/IGS intranet). Further progress has been hampered by lack of clarity of EdWeb support and future.
 13. **Goal** Migrate existing content off Plone CMS service to School Drupal service (based on IS Drupal distribution)
Progress Ongoing. See preceding goal.
 14. **Goal** Revamp ISS/IGS web presence (for current students)
Progress Completed. Now hosted on a local EdWeb instantiation.
 15. **Goal** Start work on LCFG Scientific Linux 7 server platform
Progress Work has started and will continue into 2016.
 16. **Goal** Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation, DIY DICE, etc).
Progress Some limited progress has been made.
 17. **Goal** Improve access to School services from mobile devices (eg printing, AV, lab availability, school web accessibility, OpenVPN, VDI)
Progress OpenVPN in service. Ongoing project to permit printing from mobile devices. EdWeb provides an improved web experience for mobile devices.

18. **Goal** Videoconferencing (low hanging fruit)
Progress No progress.
19. **Goal** Continue development work to take advantage of new account management framework (eg multi-faceted identities)
Progress Lifecycle code completed. Account end/grace period/disabling now in service. Our new account management framework has enabled increased automation, particularly with respects to managing a user's account lifecycle. It has also helped us to greatly improve the accuracy and consistency of account data across our systems. Final stage (deletion) is a pending project.
20. **Goal** Complete redevelopment of new equipment inventory system
Progress Substantial progress has been made, but not yet complete.
21. **Goal** Review energy usage of research servers - perhaps sleeping idle servers and virtualising little used servers
Progress No progress
22. **Goal** Review of self-managed servers (due to space, energy and security concerns)
Progress Developing a space-lease plan.
23. **Goal** Start work on migration to IPV6 (initial investigations and planning)
Progress Work has started and significant progress has been made.
24. **Goal** Perform an audit of all research data within the School
Progress Ongoing. A framework for recording data has been developed.
25. **Goal** Continue to implement improvements to security of web services
Progress Completed.
26. **Recurring Goal** Aim for a minimum of 20% of development time to be dedicated to user submitted projects
Progress The large number of mandatory projects this year made this impossible to achieve.
27. **Recurring Goal** Further promote School developed solutions to the rest of the University and beyond
Progress No further progress.
28. **Recurring Goal** Further improve/maintain communication between users and computing staff
Progress No further progress.
29. **Recurring Goal** Ring-fencing 5% of individual computing staff's time for staff development, including user support staff.
Progress Largely achieved, but not for user support staff who have been particularly stretched this year because of the building decants.
30. **Recurring Goal** Consideration of ways to minimise our energy footprint, eg identifying under-used research servers
Progress No further progress.
31. **Recurring Goal** Further consideration of migration to central services (big ticket items only)
Progress No further progress.

Unplanned activities

1. The compromise of the IS certificate signing service. This resulted in some disruption and work -:

- Generating of a new Informatics CA certificate, having it signed and testing roll-out implications.
- Roll-out of CA chain for DICE.
- Roll-out new certs for services (e.g. https, ldap, openvpn).
- Dealing with MDP desktops.
- Advising users (e.g. self-managed).
- Writing documentation.

Activities to be considered for de-prioritisation

- **Goal** Legacy web sites
Progress No progress.
- **Goal** Legacy email domains
Progress No progress.
- **Goal** Legacy filespace
Progress Some limited progress has been made.

Collaboration with others

We would like to register our continued appreciation of the assistance of Angus Rae, who has acted as a very responsive and effective interface with IS.

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. We co-hosted LDAPCon 2015, the international Conference on LDAP, Directory Services and Identity Management
3. Staff attended Floss Spring 2015 and LDAPCon 2015 conferences.

3 Revised plan for 2016

Each project has a cost effort estimate, where small is 1 to 3 FTE weeks, medium is 4 to 7 weeks and large is 8+ weeks.

Mandatory goals

1. Re-occupation of Appleton Tower from Forrest Hill and Wilkie (assuming no change to current plan)
Who: School, Cost: large
2. Produce a register of medium-high risk data and a mechanism for users to self populate the register
Who: School, Cost: small
3. Take remaining steps to implement College security action plan
Who: School, Cost: small

High priority goals

1. Consideration of how best to make use of the new central RDM services
Who: Research, Cost: small
2. Continued consideration of appropriate use of central data storage facilities
Who: Research, Cost: small
3. Engage in requirements capture for and design of proposed central archiving service
Who: Research, Cost: small
4. Engage with University Network Strategic Review
Who: School, Cost: small
5. Support the expansion of Data Science teaching facilities
Who: Teaching, Cost: small
6. Complete the re-factoring of School Database back-end (database engine and client)
Who: Admin, Cost: medium
7. Continue School intranet review and revamp
Who: School, Cost: small
8. Migrate existing Institute web sites off Plone CMS service to School Drupal service (based on IS Drupal distribution)
Who: School, Cost: medium
9. Investigate whether the existing network file system and underlying storage infrastructure are still appropriate for the School's requirements
Who: School, Cost: small
10. Continue work on LCFG Scientific Linux 7 server platform
Who: University, Cost: medium
11. Start to migrate School services to Scientific Linux 7 platform
Who: School, Cost: large
12. Perform a review of the future of the DICE desktop platform
Who: School, Cost: small
13. Full review of requirements and options for videoconferencing, particularly with external organisations in order to reduce travel. Including holding an innovation meeting.
Who: Research, Cost: medium

14. Review, and where required refresh, the AV facilities of the Forum.
Who: University, Cost: small
15. Review of self-managed servers (due to space, energy and security concerns)
Who: Research, Cost: small
16. Continue work on migration to IPV6
Who: School, Cost: medium
17. Perform an audit of all research data within the School (depends on Mandatory Goal 2)
Who: School, Cost: medium
18. Produce a risk register, covering staff and equipment resources, financial processes and systems
Who: School, Cost: small

Discretionary goals

1. 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: small
2. Provide significant development support to the revised System Design exercise
Who: Teaching, Cost: small/medium?
3. Perform a review of Tardis's role and sustainability
Who: Teaching, Cost: small
4. 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. Specifically - UG projects DB, Integration with RT, Reform.
Who: Admin, Cost: medium
5. Move administrative home and group filespace from AFS to Datastore
Who: Admin, Cost: small
6. Investigate options for a more reliable and sustainable virtual Linux desktop service
Who: School, Cost: small
7. Carry out feasibility study and cost/benefit analysis for deploying Cloud based printing within the School.
Who: School, Cost: small
8. Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation, DIY DICE, etc).
Who: Research, Cost: small
9. Continue development work to take advantage of new account management framework (eg continue work on multi-faceted identities)
Who: Infrastructure, Cost: medium
10. Complete redevelopment of new equipment inventory system
Who: Admin, Cost: medium
11. Review energy usage of research servers - perhaps sleeping idle servers and virtualising little used servers
Who: Research, Cost: small
12. Investigate 802.1X for some or all of School network ports
Who: School, Cost: medium
13. Implement separate backup streams for MHR and non-MHR data to meet differing reten-

tion policies (dependent on data audit)

Who: School, Cost: medium

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/maintain 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
6. Assess system security and identify potential improvements
7. Further consideration of migration to central services (big ticket items only)
8. Review impact of University activities wrt. teaching - VLEs, Distance Learning (including MOOCs) , EUCLID developments (assessment).

Activities to be considered for de-prioritisation

- Legacy web sites
- Legacy email domains
- Legacy filespace

IS services critical to Informatics

As far as we are aware, these are the IS services that are critical to the School. It is possible that there are other IS services that are widely used by our users; presumably IS maintains usage statistics that could be used to identify these. The following are in a rough priority order.

1. EdLAN / eduroam - delegated, fine-grained control would make IS management of the School's internal network more attractive
2. ECDF Eddie cluster
3. Staffmail
4. Managed Windows Desktop
5. Phones and Access Control
6. Central authentication and directory services
7. Central administrative services such as EUCLID/HR/BIS/PURE etc.
8. Public PC labs - we would like to investigate more effective use of these labs for our 1st and 2nd year students.
9. EdWeb distribution - we would like clarity on the commitment of IS to this.
10. EdWeb hosted service
11. WIKI - this service continues to prove unreliable
12. CapturED - this service has become increasingly unreliable
13. Pcounter printing, as linked to our School printing service
14. ESISS scanner - the limited number of login accounts available is overly restrictive
15. SSL certificate signing service

16. MOOC
17. Software purchasing
18. Visitor Registration Service

Additional services we would like

1. Provision for data archiving and, perhaps, curation. Note that this archiving should not be limited to research data.
2. We are interested in the proposed ECDF cloud virtualisation service.
3. Provide support for S/MIME and PGP encryption and signing of email, including institutional key and certificate signing
4. Additional programmatic interfaces to central administrative systems, as documented in Colin Higg's note on "Arguing for Authorised APIs". eg SOAP to EUCLID, SAT and VRS.
5. An API to "upload" marks data to EUCLID, instead of the current manual copy/paste process.
6. An API to apply for Janet SSL certificates
7. The ability to feed into Grouper, from our School Database, would reduce the barrier to the School making more use of central IS services
8. The ability to make more use of centrally provided group data but this remains dependent on the quality and accuracy of the data and suitable APIs

In order to achieve the aforementioned improvements and additional services, we are very keen to collaborate with both IS and other CSE schools on development and even service delivery.

4 Plan for 2017

Goals

1. Any re-occupation of Appleton Tower from Forrest Hill and Wilkie not completed in 2016
2. Plan for occupation of the new DTI building
3. Any required actions resulting from review of DICE desktop platform
4. Complete migration of School services to Scientific Linux 7
5. Continue work on migration to IPV6
6. Any required actions resulting from video-conferencing requirements review

De-prioritised areas

To be identified

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. Any member of the School can submit a project proposal via a web form. Each project is categorised into one of three prioritised categories - Mandatory, Strategic (meets one of the goals in Section 2 of this document) or Objective (meets one of the Strategic Objectives in Section 1). Projects are resourced in priority order when effort becomes available. There is a target of 20% of development time to be dedicated to user submitted projects. Teaching software requirements are met through a stable and well established system. 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 reports on activities in the past four months and future plans for work in the forthcoming four months. It includes figures on effort spent on development activities.

The School's Computing Strategy Committee (Head of School, Director of Computing, Director of Teaching, Head of Computing, Deputy Head of Computing, Director of Professional Services) provides over-site of the above processes.

B Staffing and Resources

The school employs 20 computing staff (19.8 FTE).

There are 675 managed DICE (Linux) desktops; 365 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 310 managed DICE (Linux) servers (165 physical, 135 virtual) and a further 24 beowulf compute nodes. There are an additional 50 physical self-managed servers. Our servers are housed in 3 air-conditioned machine rooms, with a total area of around 160 m².

(Figures as of 09/01/15).

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. Exter-

nally, our computing staff interact with organizations such as Usenix and FLOSS (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	Hadoop
iFile	Group web	HR desktop	Licence servers	MDP	CDT clusters
SVN	Homepages	ISS RT	Postgresql	DICE virtual desktop (NX)	
GIT	I in a Box	RBS	Coltex	Virtual DICE image	
	Legacy		Forums		
	Blog				
Infrastructure	Linux platform	LCFG	Kerberos	Directory	Network
	S/W repository	pkgforge	Cosign	Authorisation	Network services
	Server rooms	Consoles	Monitoring	Prometheus	System backups
	Virtualisation				