

# **Informatics Computing Plan 2017**

School of Informatics

March 3, 2017

# 1 Long-term vision and strategic objectives

Informatics Computing serves around 330 staff (235 teaching and research), 320 research students, 460 taught postgraduate students, 1300 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** Providing added value over services offered by College and IS.
- S3** Maintaining an optimum level of interoperability of Informatics Computing with College and IS services.
- S4** Provision of expertise to support the teaching and research activities of the School.
- S5** 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

**R4** Development of prototype services from R&D projects

**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 2016

### Mandatory Goals

1. **Goal** Re-occupation of Appleton Tower from Forrest Hill and Wilkie (assuming no change to current plan)  
**Progress** The re-occupation was postponed to 2017 because of refurbishment contract over-run.
2. **Goal** Produce a register of medium and high risk data and a mechanism for users to self populate the register  
**Progress** We have been waiting on the College wide solution (being produced by School of Engineering). Meanwhile, we have been surveying our data to identify that which is medium and high risk.
3. **Goal** Take remaining steps to implement College security action plan  
**Progress** All managed Windows desktops are now encrypted and the College security policy included in the School's induction pack. We have a project currently looking at how to encrypt any user files stored locally on DICE Linux desktops.

### High priority goals

1. **Goal** Consideration of how best to make use of the new central RDM services  
**Progress** A project has recently started to cover this work. The School has created a new Research Data team.
2. **Goal** Continued consideration of appropriate use of central data storage facilities  
**Progress** Met with IS to discuss what is available and how we can use it. A current project is looking at what advice we should give our users.
3. **Goal** Engage in requirements capture for and design of proposed central archiving service  
**Progress** Still waiting for invitation from IS to engage.
4. **Goal** Engage with University Network Strategic Review  
**Progress** Have submitted, to the Review, a detailed requirements document describing the Informatics network. One of our computing team is representing the College of Science and Engineering on the Review board.
5. **Goal** Support the expansion of Data Science teaching facilities  
**Progress** New clusters installed and fully operational with hadoop/grid-engine/glusterfs.
6. **Goal** Complete the re-factoring of School Database back-end (database engine and client)  
**Progress** The re-factoring is now complete. Some snagging work and documentation remain to be completed.
7. **Goal** Continue School intranet review and revamp  
**Progress** The intranet has been revamped and is now hosted on a server running the EdWeb Drupal distribution.

8. **Goal** Migrate existing Institute web sites off Plone CMS service to School Drupal service (based on IS Drupal distribution)  
**Progress** No progress.
9. **Goal** Investigate whether the existing network file system and underlying storage infrastructure are still appropriate for the School's requirements  
**Progress** We have concluded that we no longer can afford nor justify a SAN and we will decommission this service over the next few years as we replace storage servers. We have a pending project to look at network file systems.
10. **Goal** Continue work on LCFG Scientific Linux 7 server platform  
**Progress** This work was completed early in 2016. It proved to be a much more substantial piece of work than anticipated, due to significant changes in Linux infrastructure between Scientific Linux 6 and 7.
11. **Goal** Start to migrate School services to Scientific Linux 7 platform  
**Progress** Projects to migrate services are well underway (around 50% complete)
12. **Goal** Perform a review of the future of the DICE desktop platform  
**Progress** A group has been formed to perform the review, but has made little progress so far.
13. **Goal** Full review of requirements and options for videoconferencing, particularly with external organisations in order to reduce travel. Including holding an innovation meeting  
**Progress** This was completed for the immediate ATI and CDT requirements.
14. **Goal** Review, and where required refresh, the AV facilities of the Forum  
**Progress** A review of the ground floor meeting rooms was performed and improvements identified. These improvements are currently being procured and scheduled.
15. **Goal** Review of self-managed servers (due to space, energy and security concerns)  
**Progress** The self-managed server room reached full occupancy. .... still need to review ... / expansion
16. **Goal** Continue work on migration to IPv6  
**Progress** IPv6 is now supported on all managed Linux systems (desktops and servers). We experienced problems with Managed Desktop Window systems which are being investigated by IS. Support for self-managed systems will require further development work and edge switch upgrades.
17. **Goal** Perform an audit of all research data within the School (depends on Mandatory Goal 2)  
**Progress** At least 50% completed. We have gone ahead with an audit without waiting on the College DB service.
18. **Goal** Produce a risk register, covering staff and equipment resources, financial processes and systems  
**Progress** No progress

## Discretionary goals

1. **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** None.
2. **Goal** Provide significant development support to the revised System Design exercise  
**Progress** This was delayed pending a change in the academic staff running the exercise.
3. **Goal** Perform a review of Tardis's role and sustainability  
**Progress** A review has been carried out and a Memorandum of Understanding is being produced. Tardis are now connected to EdLAN via the Informatics network, rather than being connected directly. This puts them under the control of the Informatics firewall systems.
4. **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 Request Tracker (RT), Reform.  
**Progress** There is a new MSc/UG4 projects database.
5. **Goal** Move administrative home and group filesystem from AFS to Datastore  
**Progress** A project is underway to carry out this work. We are taking the opportunity to rationalise group space (some of which is still based along formative Departments' structures)
6. **Goal** Investigate options for a more reliable and sustainable virtual Linux desktop service  
**Progress** No progress.
7. **Goal** Carry out feasibility study and cost/benefit analysis for deploying Cloud based printing within the School  
**Progress** A pilot service in the Forrest Hill teaching labs was introduced. The feasibility study for introducing to the rest of the School has still to be produced.
8. **Goal** Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation, DIY DICE, etc).  
**Progress** The structure has been identified, but content has yet to be produced.
9. **Goal** Continue development work to take advantage of new account management framework (eg continue work on multi-faceted identities)  
**Progress** No progress.
10. **Goal** Complete redevelopment of new equipment inventory system  
**Progress** Considerable progress has been made on development of this system, but unfortunately it is still not ready for deployment.
11. **Goal** Review energy usage of research servers - perhaps sleeping idle servers and virtualising little used servers  
**Progress** No progress.
12. **Goal** Investigate 802.1X for some or all of School network ports  
**Progress** No progress.

13. **Goal** Implement separate backup streams for Medium to High Risk (MHR) and non-MHR data to meet differing retention policies (dependent on data audit)

**Progress** A project to carry out this work has started, but has been delayed pending an upgrade of the TiBS backup software.

## Recurring goals

1. **Goal** Aim for a minimum of 20% of development time to be dedicated to user submitted projects

**Progress** Unfortunately, only 3% of development time was spent on user submitted projects.

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

**Progress**

3. **Goal** Ring-fencing 5% of individual computing staff's time for staff development, including user support staff.

**Progress** Not all units achieved this figure this year, which may be accounted for by fewer people attending conferences.

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

**Progress** No progress other than replacing elderly kit with virtualised servers or newer, more energy efficient, hardware.

5. **Goal** Assess system security and identify potential improvements

**Progress** We held a security review and a number of improvements were identified and resulting work prioritised.

6. **Goal** Further consideration of migration to central services (big ticket items only)

**Progress** No progress.

7. **Goal** Review impact of University activities wrt. teaching - VLEs, Distance Learning (including MOOCS), EUCLID developments (assessment).

**Progress**

## Unplanned activities

1. The expected return to AT did not occur in 2016 and has been deferred to 2017. The School has taken the advantage of this delay to reconfigure AT office and lab space to cope with increased PGT numbers. This reconfiguration required some detailed planning.
2. The rapid increase in GPU demand was unexpected. This has potential implications for our energy reduction drive.

## Activities to be considered for de-commissioning

1. **Goal** Legacy web sites

**Progress** Captured usage data, but still analysing.

2. **Goal** Legacy email domains

**Progress** Captured usage and analysed. Made a proposal (to do what) to CSG for consideration.

3. **Goal** Legacy filespace

**Progress** Mostly just admin data remaining - and this is being tackled during the migration to DataStore.

## 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 Schools and IS on development and even service delivery.

1. We continue to provide the base LCFG Linux platform to other schools (via IS) and host the monthly LCFG deployers meeting.
2. One of our team was involved in the EdWeb code sprints.
3. We maintain a close working relationship with ECDF, have advised on large scale hardware procurement, and have been particularly involved in the recent GPU service expansion

### 3 Revised plan for 2017

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.

Information on computing projects is available at <http://computing.projects.inf.ed.ac.uk>

#### Mandatory goals

1. Re-occupation of Appleton Tower from Forrest Hill and Wilkie  
*Who: School, Cost: large, Project:*
2. Produce a register of medium and high risk data and a mechanism for users to self populate the register  
*Who: School, Cost: small, Project: 307*
3. Take remaining steps to implement College security action plan  
*Who: School, Cost: small, Project: 372*
4. Continue to engage with University Network Strategic Review  
*Who: School, Cost: small, Project:*
5. Adapt processes and systems for migration of assessment from Theon to APT  
*Who: Teaching, Cost: small, Project:*

#### High priority goals

1. Consideration of how best to make use of the new central RDM services  
*Who: Research, Cost: small, Project: 386*
2. Engage in requirements capture for and design of proposed central archiving service  
*Who: Research, Cost: small, Project:*
3. Complete the documentation for the School Database back-end (database engine and client)  
*Who: Admin, Cost: small, Project: 207*
4. Migrate existing Institute web sites off Plone CMS service to School Drupal service (based on IS Drupal distribution)  
*Who: School, Cost: medium, Project: 388*
5. Investigate whether the existing network file system is still appropriate for the School's requirements and identify possible alternatives  
*Who: School, Cost: medium, Project:*
6. Complete the move of administrative home and group filespace from AFS to Datastore  
*Who: Admin, Cost: small, Project: 369*
7. Consideration of email provision given that fewer Schools are now using Staffmail  
*Who: School, Cost: small, Project:*
8. Complete the migration of School services to Scientific Linux 7 platform  
*Who: School, Cost: medium, Project: 362,357,358,359,360,361*
9. Perform a review of the future of the DICE desktop platform  
*Who: School, Cost: small, Project: 379*
10. Investigate options for a more reliable and sustainable virtual Linux desktop service  
*Who: School, Cost: small, Project: 389*

11. Continue 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, Project:*
12. Implement upgrade of AV facilities in Forum G.07/G.07A, in collaboration with IS LTS  
*Who: University, Cost: small, Project:*
13. Establish a target of maintaining energy consumption (by computing equipment) at 2016 level, or lower. Possible approaches include increased use of cloud services for computation work  
*Who: University, Cost: small, Project:*
14. Review policies with respect to self-managed servers (due to space, energy and security concerns)  
*Who: Research, Cost: small, Project:*
15. Improved management of Commercial tenants, with respect to network provision and user support  
*Who: Research, Cost: small, Project:*
16. Complete the audit of all research data within the School (depends on Mandatory Goal 2)  
*Who: School, Cost: medium, Project:*
17. Implement separate backup streams for MHR and non-MHR data to meet differing retention policies (dependent on data audit)  
*Who: School, Cost: medium, Project: 346*
18. Investigate options for additional disk encryption on DICE desktops  
*Who: School, Cost: small, Project: 382*
19. Second factor authentication for system administrator accounts  
*Who: School, Cost: medium, Project: 399*
20. Produce user security training materials  
*Who: School, Cost: small?, Project: 403*

## Discretionary goals

1. Continued consideration of appropriate use of central data storage facilities  
*Who: Research, Cost: small, Project: 387*
2. Continue engagement with the PURE project to meet identified requirements for knowledge management functions other than those related to teaching administration (eg use of PURE to capture impact and improve public engagement).  
*Who: Admin, Cost: small, Project:*
3. Provide significant development support to the revised System Design exercise  
*Who: Teaching, Cost: small/medium?, Project:*
4. Investigate streaming lecture slides to teaching lab desktops  
*Who: Teaching, Cost: small?, Project:*
5. Complete feasibility study and cost/benefit analysis for deploying Cloud based printing within the School.  
*Who: School, Cost: small, Project:*
6. Update existing CVS/SVN based web content to current University style  
*Who: School, Cost: small, Project:*
7. Produce guidance on resources available for research projects (eg software repositories, wikis, VMs for software preservation, DIY DICE, etc).  
*Who: Research, Cost: small, Project:*

8. Review, and where required refresh, the AV facilities of the Forum other than G.07/G.07A  
*Who: University, Cost: small, Project:*
9. Improve end-of-life account management.  
*Who: Infrastructure, Cost: medium, Project:*
10. Complete redevelopment of new equipment inventory system  
*Who: Admin, Cost: medium, Project: 269*
11. Investigate 802.1X for some or all of School network ports  
*Who: School, Cost: medium, Project:*
12. Investigate improved storage platform for virtualisation service  
*Who: Infrastructure, Cost: medium, Project:*
13. Continue work on migration to IPv6 - support for self-managed machines  
*Who: School, Cost: medium, Project:*
14. Review energy usage of research servers - perhaps sleeping idle servers and virtualising little used servers  
*Who: Research, Cost: small, Project:*
15. Develop and document policies and procedures for physical security  
*Who: School, Cost: small, Project: 394*
16. Investigate whole disk encryption of DICE machines  
*Who: School, Cost: small, Project: 396*
17. Firewall self-managed machines (from rest of School)  
*Who: School, Cost: small, Project: 397*
18. Improve security of package distribution system  
*Who: School, Cost: small, Project: 401*
19. Improve security of LCFG profile access  
*Who: School, Cost: medium, Project: 402*
20. Further improve standard web configuration (LCFG headers)  
*Who: School, Cost: medium, Project: 400*
21. Produce a risk register, covering staff and equipment resources, financial processes and systems  
*Who: School, Cost: small, Project:*
22. Consolidate internal computing documentation  
*Who: School, Cost: medium, Project: 391*
23. Appoint a Learning Technologist  
*Who: Teaching, Cost: small, Project:*

## 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. Ring-fencing 5% of individual computing staff's time for staff development, including user support staff.
4. Consideration of ways to minimise our energy footprint, eg identifying under-used research servers
5. Assess system security and identify potential improvements
6. Further consideration of migration to central services (big ticket items only)
7. Review impact of University activities wrt. teaching - eg. VLEs, Distance Learning

(including MOOCS)

## Activities to be considered for de-commissioning

- Legacy web sites
- Legacy school database
- Plone CMS

## 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. A standards compliant IMAP mail service (eg Staffmail)
3. Managed Windows Desktop
4. ECDF GPU cluster (based on expected use)
5. Phones and Access Control
6. Central authentication and directory services
7. Central administrative services (and feeds from) such as EUCLID/HR/BIS/PURE/EUGEX etc.
8. DataStore
9. Office365 mail (for students)
10. ECDF subversion service
11. Public PC labs - we would like to investigate more effective use of these labs for our 1st and 2nd year students.
12. ECDF Eddie cluster
13. EdWeb distribution
14. EdWeb hosted service
15. Learn
16. WIKI
17. IDM (Identity management system)
18. Lecture capture service (currently using HSS panopto service)
19. Pcounter printing, as linked to our School printing service
20. Cloud printing
21. SSL certificate signing service
22. MOOC
23. Software purchasing
24. 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. Support for secure, open standards based, email.
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" assessment data to central systems instead of manual copy/paste processes.
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
9. Replacement for ESISS scanning service
10. A more pervasive attitude, particularly with respect to increasing provision and access, to the use of video conferencing across the University

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 2018

### Goals

1. Migrate Robotics teaching and research and Commercialisation to the Bayes building
2. Any required actions resulting from review of DICE desktop platform
3. Any required actions resulting from video-conferencing requirements review
4. Any required actions resulting from research server energy review
5. Any required actions resulting from network file system review
6. Establish a target of reducing 2018 energy consumption by computing devices by 10%
7. Any additional work required to refresh the Forum AV
8. Start work on porting LCFG and DICE to Scientific Linux 8
9. Produce an Information Architecture plan

### 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, 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 computing team 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 Research, 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 643 managed DICE (Linux) desktops; 312 personal machines for staff and research students, and 331 in student labs. 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 374 managed DICE (Linux) servers (207 physical, 167 virtual) and a further 36 cluster nodes. There are an additional 97 physical self-managed servers. Many servers are used to host GPUs - we now have around 200 GPUs with nearly 400,000 cores.

Trends :-

- There has been a reduction in the number of managed DICE desktops used by staff and research postgraduates. This is probably mainly due to desktop Linux becoming much more easy for end users to install and configure, and also a significant shift to Apple Mac equipment by academic staff.
- There has been an increase in the number of managed DICE desktops in student labs (with the move to Forrest Hill). This number is expected to further increase by around 100 later in 2017 with the return to Appleton Tower.
- There has been a marked increase in the number of both managed and self-managed servers, largely caused by the School being awarded three Centres for Doctoral Training

and an increased demand for GPU provision.

Our servers are housed in 3 air-conditioned machine rooms, with a total area of around 160 m<sup>2</sup>.

(Figures as of 10/01/2017).



## C Categories and activities

<b>Support</b>	Front Line	Documentation	General Web	Online Exams	Procurement
	User accounts	AV and VC	Inventory	Exam prep	
<b>File</b>	<b>Web</b>	<b>Admin</b>	<b>Misc</b>	<b>Desktop</b>	<b>Compute</b>
AFS	Web (CVS)	Theon	ssh servers	DICE desktop	DICE compute
Samba (admin)	EdWeb	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	Gridengine
SVN	Homepages	ISS and KM RT	Postgresql	DICE virtual desktop (NX)	CDT clusters
GIT	Plone	Room Booking	Coltex	Virtual DICE image	MSc cluster
Gluster	Legacy		Forums		
	Wordpress				
<b>Infrastructure</b>	Linux platform	LCFG	Kerberos	Directory	Network
	S/W repository	pkgforge	Cosign + Certs	Authorisation	Network services
	Server rooms	Consoles	Monitoring	Prometheus	System backups
	Virtualisation				