

The logo for the 'dice' newsletter. It features a red circle above a blue square containing the word 'dice' in white lowercase letters. To the right of this is the word 'NEWSLETTER' in large, bold, blue, outlined uppercase letters.

December 2009

Contents

- 1 Perchance to Dream**
- 2 Informatics Server Rooms**
- 3 OpenVPN**
- 4 FlexLM Licence Status**
- 5 SSHFS**
- 6 VirtualBox on DICE**
- 7 Virtualised DIY DICE**
- 8 Collaborative LaTeX**
- 9 School Database**
- 10 Strategy for self-managed machines**
- 11 Progress with move to AFS**
- 12 The Future of the School Web Site**
- 13 Accessing Your School File Space From Outside the School**
- 14 Results of the Multi Function Device Survey**

1 Perchance to Dream

The newest DICE computers in the student labs now sleep (and wake up again) automatically.

The User Support folks are often asked about sleep: not the restorative mammalian variety of course, but the energy-saving computer state. For a computer a sleep state is a kind of suspended animation where most large components are powered off to save energy. Most of us are probably familiar with laptops sleeping when they're not being used.

For a while we've recognised that we should also be trying to use sleep on our DICE desktop machines: the University is committed to improving its energy efficiency and environmental sustainability, and Informatics should play its part. More prosaically Informatics will very soon have to pay its own electricity bills.

Sleep has been a problem area on Linux: it's not always easy to get it to work reliably. Windows and Mac OS have benefited hugely from close cooperation between hardware and software designers, and this shows in their sleep capabilities; Linux developers on the other hand haven't always had such privileged contact with hardware vendors. Development effort also seems to have focused largely on laptops.

Sleep on desktop Linux machines therefore doesn't always work very well. A common problem is that X windows displays can fail to resume properly after sleep: the display can become corrupted and unusable until reboot. Another is that machines can occasionally hang when entering sleep. The latter problem unfortunately plagues our Dell computers, at least

with the version of Scientific Linux we use.

Even when it works, sleep isn't always appropriate. A DICE Linux machine can often be doing something important even when nobody is logged in: it might be running a Condor job, automatically adjusting its installed software or updating its user database; a large variety of remotely managed automatic configuration tasks are performed at various times of the day and night. So, when is it safe for a DICE machine to sleep?

After some development we came up with a solution. The new LCFG sleep component monitors HP dc7900 computers in the student labs. The sleep component performs two functions: it sends a machine to sleep, and it tells it when to wake up again.

Each machine is regularly checked for idleness, measured in several ways, and the appropriateness of sleep is judged using a number of factors. When sleep looks appropriate the machine is sent to sleep. Before this happens the component also examines all upcoming "cron" (automatically timed) jobs, decides which of them are too important to ignore, and sets an alarm which will wake the machine in time to run the first of these jobs.

So far this seems to be working well. At the moment we have only deployed the sleep component in the student labs, but if anyone with a desktop HP dc7900 DICE machine would like it to sleep automatically, please mention it to User Support.

Sadly we haven't yet persuaded sleep to work reliably enough on our DICE Dells, despite much effort: older models can suffer from video corruption and newer models have a frustrating tendency to hang occasionally (perhaps once or twice a week) when entering sleep. We can only hope that ongoing Linux development will at some point make sleep more reliable on these models. If that happens we do now have the technology to make use of any improvements in Linux. There is some hope of considerable improvement in OS support for power management: in particular the DeviceKit software included in Fedora 12 at first sight seems to reimplement most of this area in a fresh new way.

You can also read about automatic sleep on

DICE machines at this web address¹

Chris Cooke.

2 Informatics Server Rooms

At the time of writing, Informatics still had server rooms at five sites. By the time this is published we hope to be down to four, and to three not long after that.

2.1 Forum

There are two server rooms in the Forum:

- The main server room (B.02) houses the bulk of our managed servers and core network equipment, particularly those relating to Forum users.
- The "self-managed" server room (B.Z14) is intended as a more accessible area, where researchers can install their own self-managed servers. There are some shelves for legacy deskside equipment, though new servers should be bought as rack-mounted.

Access arrangements for this room are far from ideal at the moment. However, once the maglocks are finally sorted out we hope to be able to set up a more flexible regime.

If you would like to discuss access to this server room, or think you have kit which might be housed there, please send in a support ticket in the usual way.

- (There are also two or three "IT closets" on each floor. These contain from three to eight network "edge switches" each, to which the floorboxes in offices and labs are connected. These switches are connected in turn by fibre to the core switches in the basement.)

2.2 Appleton Tower

We decanted out of our Appleton Tower basement server room in February for a refurbishment which should have been completed by

¹<http://www.inf.ed.ac.uk/systems/sleep/>

late summer. At the time of writing (mid-November) the room still isn't ready for us to move back in. In the meantime, network provision for the upper floors of Appleton Tower is through switches in the Forum, while our AT servers have been temporarily rehoused in the Forrest Hill machine room.

The new room is an open plan area shared with IS. It's rather larger than both our old areas combined, as it incorporates the old post room too, and it's considerably better provided for in terms of power and cooling. When we do eventually move back, we'll use our space there mainly for servers directly related to our Appleton Tower user base (students, AIAI and commercialisation).

2.3 JCMB

Following our move out of JCMB, our old server room there was refurbished for use as a College-wide facility (and in the meantime we were able to use 2501, as Physics didn't immediately need it). That project is now complete, and we were the first occupants of the new space.

We intend to use this site primarily as a "disaster recovery" operation. We'll keep mirror filesystems there, and sufficient other infrastructure so as to be able to pick up the core Informatics business should something untoward happen to our primary sites (Forum and Appleton Tower).

2.4 Forrest Hill

We had hoped to be out of the Forrest Hill server room long ago, but with the Appleton Tower basement project overrunning we have been forced to keep decanted equipment there much longer than we would have liked. The power and air-conditioning arrangements there are far from ideal, and there is some impact on us from the other (themselves decanted) occupants of the building. The sooner AT basement is finally ready and we can move out of FH the better!

2.5 Buccleuch Place

By the time this newsletter is published we should finally have moved out of the Buccleuch Place server room (apart, probably, from some tidying and the disposal of old machines to the recyclers). Hooray!

George Ross.

3 OpenVPN

OpenVPN (<http://www.openvpn.net>) has been mentioned in these newsletters in the past as a pilot not-quite-a-service. We have recently, however, put our implementation on a more solid foundation, and now offer it as a supported service for those who might benefit from it.

What is it? There are a couple of problems which VPN (Virtual Private Network) systems can help solve. The first is where you're working at a remote site but you need to appear as though you were a local network user in order to access some resources. The second is where there are restrictions on your network access, often for audit-trail reasons. A "VPN tunnel" is a way to make your machine appear as though it's attached to the network somewhere other than where it really is; and OpenVPN is the system we have adopted to achieve this.

Using the Informatics OpenVPN means that you appear *inside* the Informatics network. This is in contrast with the central University VPN service, which will tunnel you to inside EdLAN but outside Informatics. This distinction may be important when accessing internal Informatics resources.

There are instructions for setting up OpenVPN on a variety of platforms²

These have been tested against MacOS and Windows XP and Linux, and are reported to work for Windows Vista too.

If you're happy with the University's central VPN provision then there's probably no reason to switch. Our OpenVPN service is offered as

²<http://www.inf.ed.ac.uk/systems/network/OpenVPN>

a lightweight alternative in the hope that some users find it useful. (I use it all the time when connecting in from my Linux laptop at home.)

George Ross.

4 FlexLM Licence Status

A majority of the population are fortunate enough to have never heard of FlexLM, the self-proclaimed “flexible licence manager” utilised by several licensed software packages used in Informatics. Many of those who have might consider its sole purpose to be to hamper efforts to run their favourite mathematical processing packages.

Several months ago the Research and Teaching unit refreshed the LCFG FlexLM component which manages all such software, attempting to make FlexLM less of an obstacle by allowing consistent management and tying in important improvements ranging from service monitoring to block licence reservation. The new component runs on all DICE FlexLM servers and, though invisible to users, one feature — the licence usage page — should not be.

Each server now produces a regularly-updated summary of the licences and features it provides, along with a list of checked-out licences and their users. This is accessible at <http://licence-server:1881/> (that’s port 1881) from any Informatics PC. Currently-licensed software includes Maple, Matlab and SIMICS; contact support if you don’t know the name of your licence server.

Graham Dutton

5 SSHFS

This semester sees a majority of Informatics using AFS, the globally accessible filesystem which gives users direct access to their filesystem from a variety of operating systems and anywhere in the world. However, AFS does not provide all the answers, and sometimes you need access to a data that isn’t or can’t be stored in your AFS file space, for example off-site or on a removable disk. Shuffling files back

and forth between devices is rarely convenient. This is where *SSHFS*, the secure-shell remote filesystem, comes in.

SSHFS utilises FUSE (Filesystem in Userspace), a comparatively new addition to Linux and other kernels which allows userspace code, such as the ssh utility, to be used as part of a filesystem implementation.

Its function is relatively easy to guess, given the name. It allows a user to mount any path on a remote host (to which they have access), given SSH access to that path. It can be mounted on any local path on your DICE machine.

Let’s say you keep a large library of temporary files in scratch space on your DICE desktop. You’d like to use this same cache on a remote server, but copying the whole directory would be inefficient since most of it remains unused. Instead, you simply create a directory to act as a mount point, and mount your remote disk there. On the remote server:

```
$ mkdir /tmp/largelib
$ sshfs me@mypc:/disk/large /tmp/largelib
```

That’s it. Your files are accessible as if in /tmp; they’re accessible only to you and only on the DICE host which did the mounting. When you’re done, you should unmount it like any other filesystem or medium. In this case it’s a FUSE filesystem, u[n]mounted with the FUSEmount utility. On your DICE desktop:

```
$ fusermount -u /tmp/largelib
```

SSHFS makes no performance guarantees (consider the work it is doing to bring you your file...) and is not designed to replace local disk access for demanding tasks. And don’t forget that your removable disks and remove storage does not necessarily come with the comprehensive backup mechanisms which protect your DICE home directory.

Graham Dutton

6 VirtualBox on DICE

All DICE machines now have *VirtualBox* installed by default. You can use this to manage and run a virtual machine on your desktop

(such as another Linux distribution or Windows). For more local details on using this application see

<https://wiki.inf.ed.ac.uk/DICE/VirtualBox>

Also see <http://www.virtualbox.org/> for instructions on setting up and running a virtual guest operating system.

Tim Colles

7 Virtualised DIY DICE

DIY DICE, where users can manage their own DICE like machine using the same technology (LCFG) that COs use for normal DICE has been around for a number of years. It is useful for people who want to change some aspects of the configuration of their machine (eg software upgrades) but don't want the hassle of fully self-managing their machine.

We have recently added support for DIY DICE under the virtualisation tools VMware and VirtualBox. This allows those of us who use MacOS or Windows as our principal platform to have our own DICE like machine without the expense (and footprint) of another physical machine. For example, I run a DIY DICE virtual machine on my Windows laptop, giving me a DICE like environment when I'm working at home or travelling.

For further details see:

<https://wiki.inf.ed.ac.uk/DiyDice>

Alastair Scobie

8 Collaborative LaTeX

The Collaborative LaTeX system (at <https://coltex.inf.ed.ac.uk/>) is a Subversion repository with a backend LaTeX build system. This allows users in the School and outside external users (via iFriend) to collaborate on the production of documents such as research proposals, reports and books.

All members of staff in Informatics have an account on the system by default and you can use it straight away to create a

repository for your document and control who can edit/view that document. See the online documentation and introduction page at the above URL for further details. A number of members of staff have been using the system successfully for over a year and you can see the current list of repositories at <https://coltex.inf.ed.ac.uk/repos.html>. If you are not a member of staff you can request an account via the support form.

Behind the scenes whenever a change is made to the document the backend will automatically rebuild the PDF from the LaTeX source (using `latexmk` and `pdflatex`) and commit that into the repository. It will also build a change document (using `latexdiff`) highlighting changes between the most recent version and the previous version (or explicitly set version). You can if you choose also enable LaTeX syntax checking on pre-commit to avoid trivial typos being version controlled.

As a trial feature, each newly created document repository also gets an attached instance of `trac` providing a web based issue tracking, collaboration and content management system and which also has direct read access to the base subversion repository for a document (see <http://trac.edgewall.org/> for more information about this system). There are some gotchas with this at present and if you want to use it please contact support in the first instance.

Tim Colles.

9 School Database

The School of Informatics Database Service (pretty much as it is now in terms of technology) was introduced in the Department of Artificial Intelligence by Tim Colles in late 1995 (although some of the technology used actually dates from 1992). Over the next few years it gradually replaced each of the pre-existing standalone database and spreadsheet systems used for various administrative tasks. It finally incorporated mark recording, processing and reporting in 1998. In late 1998 it was adopted by the Informatics Teaching Organisation at the formation of the School of In-

formatics. In so doing it replaced private local PC-based databases and spreadsheets used in the Departments of Computer Science and Cognitive Science and enabled staff on each of the four sites of the School to work together on a single central database. It can be argued, and has been officially recognized by management and academia, that having this central multi-user networked database for the School of Informatics has been a very strong unifying force, in particular for the ITO.

A year later in 1999 the Informatics Database (as it was now known) was adopted for handling postgraduate admission records for the Informatics Graduate School as well as Staff/Visitor records for the School Office. It was later adopted to hold the new Informatics Report series, and has also been used for the RAE exercise, holding Inventory records and many other miscellaneous kinds of information. Over and above the core functionality provided to the ITO such as managing students, courses and assessment processing for the Board of Examiners meetings the database service drives the automated generation of many web pages, mailing lists, practical submission system data and computing data used for account generation and authorization control. The current systems have now run from 1995 until 2009 and will need to continue to run until at least the end of the 2009/10 academic session.

Any improvements and significant additions to the database service have largely been on hold for the last five years on the assumption that EUCLID would replace the vast majority of the existing functionality and a newer light weight approach could be adopted for the remaining parts. However, given the fact that EUCLID will not now replace any of the core ITO functionality, the existing systems now very fragile support status due to lack of any serious investment over the last five years and the various new things its ageing technology is being pushed into doing by advances in surrounding technology it was time to formally review its future.

This review took place a little while ago. The result being that a large project was initiated a number of months ago to completely revamp

the database in time for the start of the 2010 academic session. This decision was taken in the light of the still ongoing de-scoping of the EUCLID project and based on the expectation that we will need to continue to provide a comparable and similarly integrated local service for at least the next five years. The project consists of two phases. The first is directed at moving the database back end to newer technology and improving the management framework surrounding this in order to drastically reduce support and maintenance load and provide scope for more cost effective future enhancements. This phase is now nearing completion. The second phase will be directed at re-factoring the existing user front-end and reporting tools used primarily by the ITO, but also by the other administration groups in the School. This phase will be guided in part by technical considerations (given the extinct development environment that the current code base is built on) but also by business requirements that have been largely set out by the ISS, ITO, IGS and HR administrative groups.

For further information the full review document is available at this link³

Tim Colles

10 Strategy for self-managed machines

In April 2008, we held an iMeeting to discuss support for personal computing. The iMeeting is minuted at:

<https://wiki.inf.ed.ac.uk/twiki/pub/ComputingStrategy/PersonalComputingMeeting>

The move to the Forum, and subsequent related work, distracted us from acting on the discussion at that meeting, but we have recently been able to revisit this issue and have developed a strategy for supporting self-managed machines.

The strategy, in full, is available at this link⁴ under the Computing Guidelines and Policies Section.

³<https://coltex.inf.ed.ac.uk/live/SchoolDatabaseReview/>

⁴<http://www.inf.ed.ac.uk/systems/>

but briefly the key points are :-

- We shall introduce a web based discussion forum to create a community of users who mutually help each other with technical problems.
- Our user documentation will be restructured, with user input, and with increased focus on using DICE services from self-managed machines.
- We shall develop more MacOS and Windows skills, particularly amongst the front-line support team and with using DICE and IS services from these platforms.
- We shall introduce frequent open surgeries where users can discuss technical issues with the computing staff responsible for running DICE services.
- We shall provide an option for users to request, on delivery of a new self-managed machine, that the machine be pre-configured for DICE services.

Alastair Scobie

11 Progress with move to AFS

Work on moving our remaining NFS file space to AFS is progressing well and we expect the process to be complete by the end of the year. This will allow us to turn off the last of our Solaris based file servers, a sad moment perhaps for some veteran members of the School.

In the past few months, we have introduced several new features to allow this move to take place, one particularly noteworthy example being the ability to serve web pages out of AFS file space. In addition, we are responding to the demand by some users to be able to run a job accessing AFS file space for longer than a month by introducing a facility for users to be able to create keytabs on individual machines, allowing jobs to run indefinitely. Announcements will be made when this facility becomes generally available.

This move would not have been possible without the efforts of Lindsey Brown and Ross

Armstrong of the User Support Unit who have laboured tirelessly to move the terabytes of data involved to the new servers.

Craig Strchan

12 The Future of the School Web Site

As users are hopefully aware, there has been a debate taking place for some time now on the future direction of the School's various web-sites, particularly www.inf.ed.ac.uk. There has long been a desire for the site's current home-grown content management system to be replaced by something more up to date. At the same time, the University's central web service, based on Polopoly, has been introduced and Schools are being strongly encouraged to move their web presence to this new service. There have been concerns however over the ability of Polopoly to manage some of our content.

How best to resolve these issues? After much discussion with the major stakeholders in the School web sites and the University web project team, we believe that we have come up with a solution which meets the needs of the School best while still committing to the central web service. A relatively small number of the most outward facing pages in the School web site will be migrated to the Polopoly service and we have been working with the University web project to ensure that the School's look and feel will be retained for these pages. The remainder of the School's content will be migrated to a new web service run by the School and based on the Zope/Plone CMS introduced into service this year.

Moving the Schools web content to a new server will be a considerable task. It should not be viewed as merely an exercise in copying the existing web sites to the new server. There is much in the School's existing web space that is outdated, redundant or badly presented and this move should be grasped as a chance to reorganise and reinvigorate the School's web presence. We hope to begin work on this great task early in the new year. Look out for further announcements!

Craig Strachan

13 Accessing Your School File Space From Outside the School

The recent concern over a possible Swine Flu pandemic and in particular the possibility of the University being closed for an extended period may have led readers to worry about how they would access data stored on the School's file systems in the event that they were not able to gain access to the University's buildings as normal.

Fortunately, accessing the School's file systems from outwith the University should be an extremely simple task thanks to the adoption of the OpenAFS file system by the School. Users can either install the AFS client on their home machine or laptop and enjoy fully integrated access to their AFS file space or they can use the web based front end to the AFS file system if it's not practical to install the AFS client on the machine they are using. Users should be aware that depending on the operating system they are using, traffic between the AFS server and client may not be encrypted by default (for example encryption is turned on by default for Windows and Debian Linux but not for Mac OS and Fedora Linux). If this is a concern, the current encryption settings can be displayed by running the command 'fs getcrypt' and encryption turned on, if necessary, with the command 'sudo fs setcrypt on'.

To install the AFS client on a machine, follow the instructions at

<http://www.inf.ed.ac.uk/systems/AFS/>

for your operating system.

If installing the AFS client isn't appropriate, you can access the file system via a web browser by going to

<https://ifile.inf.ed.ac.uk>.

You will be sent to the School's authentication page where you should enter your DICE user

name and password. After successfully doing so, you will be taken to the iFile page. You can then view files, upload files to AFS file space and download files from AFS to your local hard disk provided you are using a web browser other than Internet Explorer (we recommend Firefox) which has a bug which prevents file downloads working at the moment. There is a workaround for this problem but it involves editing the Windows registry so we do not recommend this approach unless it is completely unavoidable. If this is the case, contact support for details. A new version of iFile should be available shortly which will avoid this problem completely.

Craig Strachan

14 Results of the Multi Function Device Survey

A big thank you to all of you who took the time to fill in the multi function device survey during the summer. There were 105 respondents in total, a far higher number than we had expected.

There isn't room here to cover the responses to every question in the survey but several issues emerged which will be addressed in the near future

- There isn't enough clear simple information about how to use the MFDs as scanners, photocopiers and fax machines available to users. Indeed several respondents were not aware of the full range of facilities the devices offer. We will be making more user orientated documentation available close by the devices.
- This lack of information was particularly felt when trying to use the devices as scanners. It was felt that the process of accessing the scanned material was over-complicated and not clearly outlined. To address this, we will be adding the option of having scanned documents emailed directly to the user's mail box.
- The reliability of the MFDs was called into question with many respondents recalling

a period when it seemed that every week brought at least one email about a broken MFD. A questionable design feature of the MFDs allowed toner to drop directly onto a sensor accounting for many of these breakdowns but our MFDs have now been fitted with a redesigned part which should prevent this happening and our sense is that the reliability of the MFDs has improved since then. If you disagree, please let us know!

People also complained that on many occasions, their printouts failed to appear at all. We believe that this was a problem with all printers, not just the MFDs and that it was related to the introduction of the CUPS printing service. Once again we believe that changes to the CUPS configuration should have improved this problem.

- Users expressed frustration that when scanning or photocopying, their jobs could be interrupted by people sending print jobs to the devices. It is possible to prioritize scanning and photocopying over printing and we will set the MFDs up this way for a trial period.
- A few users expressed frustration over the speed of the MFDs when photocopying and the lack of sophisticated stapling and collating facilities. Part of this may be due to the lack of documentation referred to above about the MFD's features but we will be looking into the possibility of having at least one 'heavy duty' photocopier in the Forum.
- Several respondents asked why we couldn't go back to having separate devices for photocopying, printing, scanning and faxing. It's perhaps worthwhile reiterating that the University has a strict policy that wherever possible MFDs such as these should be used in place of individual devices unless there are compelling reasons for not doing so. Given that a large proportion of the responses to the survey were positive (92% found the ease of use of the MFDs to be good or acceptable and 79% found the reliability to be good or acceptable) we don't really have a compelling

case for replacing these devices at the moment. We have taken the decision that for the moment at least, we will not be replacing monochrome printers with MFDs and to this end have ordered up further HP P4015s to replace the ancient HP 8150s in the South West printer rooms.

Once again, we would like to thank all those people who completed the survey with special thanks going to those who took the time to add their own comments. An extra special thank you goes to the anonymous author of the limerick in dubious taste which shocked us to our cores. Sadly, it cannot be printed in a genteel publication such as this one.

Craig Strachan