


School of Informatics, University of Edinburgh

Issue 1

Contents

Farewell to Jeremy	1
Welcome to Our Newsletter	1
DICE Project Review	1
Support Policies and Guidelines	2
How To Get A New Account	2
Big Network Changes In Progress	3
The Future of Wireless	3
Informatics From Anywhere	4
New Servers, New Services	4
Solaris Services	5
Obsolete Services Replaced	6
Handling Microsoft Documents on DICE	8
Linux laptop support	8
Linux platform non upgrade	9
Redhat 9 on Laptops	9
System Configuration Research	10

► Farewell to Jeremy

Jeremy Olsen, a mainstay of Computing Support at Buccleuch Place for fourteen years, this week left for a new life in London. Jeremy joined Cognitive Science in August 1989, having previously studied for an MSc in the then AI Department in 1987/1988.



Then...

And now!

Jeremy and his partner Cate are getting married this month. We wish them both well and offer Jeremy our thanks and very best wishes.

► Welcome to Our Newsletter

This is first edition of a regular newsletter produced by Computing Staff for Informatics staff and students. We hope you will enjoy reading about our current work and future plans - and how these will affect you. Please tell us which articles you find useful, and tell us about any subjects you'd like us to cover in the next edition - due in late summer. Read this newsletter on-line at: <http://www.inf.ed.ac.uk/systems/newsletters/2004-05.pdf>

Morna Findlay and Chris Cooke
<morna@inf.ed.ac.uk, cc@inf.ed.ac.uk>

► DICE Project Review

Input from users required

The DICE Project¹ provides almost all of the computing infrastructure for the School of Informatics. This ranges from low-level network connections and underlying technology such as configuration management, to user-visible services like email and user-support.

It is now over four years since DICE was initially proposed, and we are currently performing a review of the project. A strategy document² is now available which is intended to reconsider the original aims, and to propose future priorities.

There are several issues which pose important strategic and policy questions that need to be addressed by the School in general, rather than by just the CO community, and input on all aspects of the document is very welcome, either directly to Paul Anderson (dcspaul@inf.ed.ac.uk), or to any member of the Computing Committee.

The document also presents a good overview of the aims of DICE, and the current state of the services, which should be of interest to those people who are unfamiliar with the origins of the project.

Paul Anderson <dcspaul@inf.ed.ac.uk>

¹Distributed Informatics Computing Environment

² www.dice.inf.ed.ac.uk/groups/technical_strategy/documents/-strategy04.pdf

Support Policies and Guidelines

Not only is there much technical development for DICE, there are many things happening on the support side. The main Systems pages have had a major revamp, to help users find information they need more easily. On the operational side, revised guidelines and policies relating to provision of computing to users in Informatics are now available. Current topics include the level and extent of support and how computing resources are provided. There will be additional information in due course.

Provision of policies and guidelines for such a diverse group of users and requirements as exists in Informatics is far from straightforward. Equally, the policies and guidelines are not rigid and unchanging. It is important that there is feedback about how they are working and what else is needed. Please let us have your comments.

Support is changing!

The advent of a cross-site computing service requires changes to the way that service is supported.

Traditionally legacy sites handled support in slightly different ways. The Frontline Support Team are currently standardising practices so that users can expect the same service wherever they are.

One example is the soon to be completed creation of dedicated computing support offices at each site. This provides a consistency of operation and allows us greater flexibility in providing support throughout the year.

However, we don't expect radical differences and we will be advising local users of any things that might change.

Why use a web form for support?

A notable change to provision of computing support in Informatics was the introduction of the support web form. It is clear from feedback received that this has not been entirely popular! So why did we do it?

It is generally reckoned that properly identifying a user's problem takes about 80% of the total time taken to resolve it. By having user and machine information clearly identified, that data is used during submission for database lookups of essential information about both user and/or the machine in question.

This can considerably speed the diagnosis of the problem by providing extra background/context.

It is worth noting that after the initial form is submit-

ted, all subsequent communication can be by email. So, while it may be a little less convenient for initial contact than email, it really does make it easier for us to help you.

The top of the form also presents reminder links to topics available in the FAQ which may answer your question immediately.

The Frontline Support Team
<<http://www.inf.ed.ac.uk/systems/support/>>

How To Get A New Account

What to do and who to talk to.

When you need an account, contact:

David Dougal BP visitor or research accounts

✉ david.dougal@ed.ac.uk

☎ 504446

👉 2BP-1R17

Dyane Goodchild KB visitor or research accounts

✉ dyane@inf.ed.ac.uk

☎ 505116

👉 JCMB-3421

Jean Bunten All other staff and visitor accounts

✉ jeanb@inf.ed.ac.uk

☎ 502690

👉 AT-3.01

They will take the details from you and perform the administrative magic. They will then pass the details to the Support team, who will make the account.

Account details can usually be collected from your local Support Office³.

Making a new account can take up to three days, so it's important that *as much notice as possible* is given.

Staff have to be given a unique username by EUCS, and this has to be entered in the Informatics database before an account can be created.

Alison Downie <alisond@inf.ed.ac.uk>
Lindsey Brown <lmb@inf.ed.ac.uk>

³See www.inf.ed.ac.uk/systems/support for locations

► Big Network Changes In Progress

A large number of developments have taken place to the Informatics network over the last few months.

Some of this has been part of our ongoing upgrade and enhancement work. Much, however, has been as a result of the fire, which is still causing a significant workload, as well as building expansion and renovation.

This has taken a considerable amount of CO and Technician time to plan and implement, a major commitment which looks set to continue for some time yet.

At Appleton Tower we have gone from being a small slave-site reliant on KB, to a full free-standing operation fully equal to the other three sites. Along the way there have been several major network reorganisations required, but where possible we have used our built-in redundancy to minimise disruption to the building's users.

Although the bulk of the network reorganisation should now be behind us, there is still quite a bit of work required to bring the remaining Informatics areas fully on-line.

At Buccleuch Place we have been expanding along the street, into premises which have had a considerable amount of sub-standard wiring which we have had to replace.

Several of the switches serving the existing part of Informatics were also below-spec and have been upgraded.

At Forrest Hill a number of temporary network arrangements were required as a result of building work elsewhere.

We have also been integrating the ANC network, work which is still ongoing, which has improved connectivity and resilience.

At King's Buildings we have expanded into the new IPAB area, as well as upgrading the MSc workroom and 3311, an expansion of roughly 25%.

At all four sites we now have **Gigabit connections** to EdLAN and beyond, as well as using Gigabit links for all our primary internal connections.

A number of core servers have also now been given Gigabit connections to the network.

We have installed fast switch/routers at all four sites, so that there is essentially no difference in network latency between any two machines anywhere on the Informatics network.

We have **considerable redundancy** in place to protect against failure of our own and EdLAN equipment.

In the central area there are backup links connecting

together Appleton Tower, Buccleuch Place and Forrest Hill, while at King's Buildings we have an arrangement with the School of Engineering and Electronics to provide a mutual backup link.

Internally, almost all our switches continue to be part of a mesh of backup links; and our core switches have dual redundant power supplies, connected where possible through uninterruptible power supplies (UPSes) to separate mains phases or even completely separate mains feeds.

Wherever possible we have network infrastructure machines in place at all four sites, so that no site is dependent on any other for its network operation.

These arrangements have cut in automatically on quite a few occasions since they were first put in place, keeping the network running when there would in the past have been a break in service.

Our **perimeter firewall routers** are now blocking several tens of thousands of probes every day (not long ago a few dozen would have been unusual!).

Our network monitoring system has been extended to cover all four sites, providing us with comprehensive traffic graphs, as well as notification messages concerning network and end-system faults.

In order to satisfy packet auditability requirements, we can also associate network ports with addresses as necessary, though this still requires tedious manual correlation at present.

Developments to the "SRIF" network and the wireless service are described in separate articles.

George Ross <gdmr@inf.ed.ac.uk>

► The Future of Wireless

New Access Points for North and South

For the last few years Informatics has operated a prototype wireless network system, mainly in Buccleuch Place with some coverage at Appleton Tower and Kings Buildings.

The EUCS is also introducing a University-wide wireless network to cover many lecture theatres, libraries and public areas.

There are clear advantages to users in being able to move seamlessly between wireless network areas. Moreover, due to the limited channel-space available we would not be able to offer duplicate coverage of those areas included in the central wireless network,

even if we thought it was a sensible use of resources to do so.

Our existing Access Points are of an old design which does not support the new higher-speed standards. They are also no longer obtainable and we would have to consider our options in any case if we wanted to expand our coverage.

Due to deficiencies in the original wireless standards, the “informatics” system as currently implemented is not secure and can not easily be made so. The necessary changes would result in something very similar to the “central” service in any case.

After due consideration, we have concluded that it makes much more sense overall to opt in to the University-wide “central” wireless network, rather than to attempt to duplicate parts of it ourselves. New coverage at Appleton Tower and Buccleuch Place is therefore being installed on that basis, using access points capable of supporting the faster 802.11g standard in addition to the original 802.11b. Existing “informatics” coverage at Kings Buildings and Buccleuch Place will transition to “central” over the next few months.

We will also shortly be introducing a pilot “central” service at Forrest Hill, though it may take a period of experimentation before we find suitable locations for the access points so as to provide acceptable coverage due to the nature of the building.

The [EdLAN Access web pages](#)⁴ contain instructions on registering for the central wireless service, along with helpful configuration advice, and a list of the public areas with existing and planned wireless coverage.

The [Informatics network team web pages](#)⁵ describe each of our wireless access points in detail.

George Ross <gdmr@inf.ed.ac.uk>

► Informatics From Anywhere

Virtual Private Network — Trial Users Wanted

A Virtual Private Network (“VPN”) is a mechanism whereby secure tunnels may be extruded from some network over insecure paths, such as wireless networks or the Internet, to remote machines, making them appear as though they were really inside the perimeter of the base network — in this case, the School of Informatics network.

⁴ www.ucs.ed.ac.uk/nsd/access/

⁵ www.dice.inf.ed.ac.uk/groups/infrastructure/network/docs/Wireless.html

This would typically allow services such as private web pages or mail relaying to be made available to those remote machines. Those who might benefit include users working at home or while travelling, and users of the new University-wide wireless network.

Most current VPN products are based around the IPsec (“secure IP”) protocol standards, developed as part of IPv6. There is still a little way to go, however, before these settle down enough to provide a full interoperable solution, particularly for users behind NAT (“network address translating”) gateways such as most home ADSL and cable modems. We will continue to track these developments, of course, and will introduce a service based on these standards in due course.

As an alternative we have been investigating OpenVPN⁶, which uses the TLS secure transport protocols to provide VPN functionality for UNIX and Windows.

We have a pilot service operating now for Linux laptops (RH7 and RH9), and although it is not yet fully implemented we would welcome some more trial users. Please contact us using the support form⁷ in the first instance if you’re interested.

Other mechanisms exist already, of course, such as ssh tunnels. The main advantage of a VPN-type solution over these is its transparency. Once the VPN is established, traffic is routed entirely automatically without the user or application being aware that the machine is not actually connected by a piece of real wire directly to the base network.

George Ross <gdmr@inf.ed.ac.uk>

► New Servers, New Services

Which machines to use for remote access, file transfer, heavy computing

These days reports of serious internet security breaches seem to come thick and fast, and malicious security probes on Informatics machines are a daily fact of life. With this very much in mind we’re in the midst of a revamp of our remote access and multi-user machines.

Remote Access Gateways

To minimise the number of vulnerable targets we present to the outside world, we provide only two machines to allow logins from outside Informatics: `staff.ssh.inf.ed.ac.uk` and `student.ssh.inf.ed.ac.uk`.

⁶ openvpn.sourceforge.net

⁷ See www.inf.ed.ac.uk/systems/support

Research students, visitors and ex-staff with accounts can all use the staff machine; anyone with a DICE account can use the student machine. Connect to these machines using `ssh`.

The remote access gateways are there purely to provide access to the Informatics network, rather than as machines on which a lot of work can be done. A few of the more demanding software packages have been deleted from the gateway machines in an attempt to maximise their availability and ensure that they don't get too slow. Once logged in to either gateway you can then login to a more suitable machine for working on.

The older departmental remote access machines are being withdrawn - See page 7

General Use

When you're logged in to a gateway machine, or any Informatics machine come to that, and you need access to a DICE machine to get some work done, you can use `ssh` to login to one of these:

```
student.login.inf.ed.ac.uk
bpstaff.login.inf.ed.ac.uk
kbstaff.login.inf.ed.ac.uk
fhstaff.login.inf.ed.ac.uk
atstaff.login.inf.ed.ac.uk
```

Once again the student machine is accessible to everyone, and the staff machines are accessible to staff, visitors and research students. There's a staff machine at each site, in case of network failure, but there are no special site-based login restrictions — just use whichever machine is closest to where you normally work.

Extra Compute Power

The multi-user login servers are all very well for a small amount of work, but they don't have enough power to support compute-intensive work. For these we've provided two compute servers. `student.compute.inf.ed.ac.uk` is for taught students, and `staff.compute.inf.ed.ac.uk` is for staff, visitors and research students. Feel free to use these for long-running or demanding computing jobs.

File Transfer

Files can be transferred between DICE home directories and remote machines via the two remote access gateways. The `scp` and `sftp` utilities can both be used. The *Downloads and File Transfers* section of the

Support FAQ⁸ may help here, or refer to the man page for each command.

Older file transfer facilities such as dedicated ftp servers will be withdrawn, after consultation with users. We're currently considering how to go about this and what extra replacement facilities might be needed.

Chris Cooke <cc@inf.ed.ac.uk>

► Solaris Services

Old servers decommissioned, New servers installed

The recent introduction into service of the new file-servers *sphinx* and *wyvern* at KB (see next issue!) means that within the next couple of months the last few Informatics machines running Solaris 2.6, 7 and 8 will be taken out of service.

Crusty hackers who still remember the Sun 2 fondly will be pleased to hear that this doesn't signal the final demise of Solaris within Informatics. We are running Solaris 9, configured via LCFG, on the new file-servers and certain other infrastructure machines.

This does however mean the end of support for day-to-day Solaris use and in particular for Solaris on the desktop. Users still using applications running under Solaris should make strenuous efforts to port these applications to DICE as soon as possible. It's appreciated however that there are circumstances where this will simply not be possible, either because the application cannot be made available under Linux for some reason or because the application is one being supported by members of Informatics on Solaris sites outwith the School.

To support these applications, we will be continuing to provide one or more (depending on demand) Solaris 9 multiuser machines which can be used for developing and running Solaris applications. The first of these machines, a SunBlade 100 named *troll* is already available; most of the standard applications such as `emacs` and `gmake` can be accessed by adding `/usr/sfw/bin` and `/opt/sfw/bin` to the `$PATH` shell variable and Informatics users can log in using their normal username and password.

These multiuser machines are NOT intended for day-to-day use. If you need help with porting Solaris applications to Linux, please contact Support⁹.

Craig Strachan <cms@inf.ed.ac.uk>

⁸www.inf.ed.ac.uk/systems/support/FAQ/

⁹See www.inf.ed.ac.uk/systems/support

► Obsolete Services Replaced

Redhat 7 and Legacy Services to be replaced

Maintaining outdated and obsolete machines and services is always costly in time, effort and money, so all hardware and services reach a natural end of life (*EOL*). This article covers some of those services which will be replaced over the next few months.

Legacy (cogsci, dai, dcs) End Of Life

There are still a considerable number of legacy (pre-DICE) services and clients.

We propose an EOL of 1st September for these, with only certain agreed and documented services being allowed to pass this EOL.

Some of these services will simply die, while others will be replaced by DICE-based equivalents. We have started a review to determine how to deal with each service and will be consulting with Computing Committee on our proposals.

For example, the following are still running on, or rely on, legacy servers :-

- www.dcs.ed.ac.uk, www.cogsci.ed.ac.uk (See page 7)
- the Informatics wireless service (See page 3)
- many multi-user servers (eg ssh.dcs, ssh.cogsci etc) (See page 7)
- mail to dcs.ed.ac.uk, cogsci.ed.ac.uk (See page 6)
- NFS file servers at KB and BP
- self managed network wires

There are [more details on the web](#)¹⁰.

Let us know your concerns: We plan to seamlessly replace many of these services before September: for example, content under www.dcs.ed.ac.uk will not be deleted, and email to legacy domains will be forwarded. Users who have concerns should contact the Computing Committee to ensure that these are addressed.

Alastair Scobie <ascobie@inf.ed.ac.uk>

¹⁰see <http://www.dice.inf.ed.ac.uk/groups/infrastructure/news/-legacyeol.html>

Legacy Mail — EOL Approaching

The old legacy mail servers are due to be replaced by a new virtual mail server. The old services are already being phased out, and the project is planned to be complete by June 1st.

Temporary Mail Domains Now Possible: The new server can also handle virtual mail domains - for example the Edinburgh-Stanford Link domain: @eslink.org or the LCFG domain: @lcfg.org.

Mail for users within these domains is not delivered locally, but is forwarded to the users usual mail addresses: however this feature allows working groups, conferences and so on to publish their own mail addresses if they so choose.

Users should contact support if they are interested in having a virtual mail domain.

No change for current users: This new DICE mail server will forward legacy mail addresses for current users onto the user's equivalent @inf mail box: ie there will be no local mail storage on this service. As this is already our practice, the vast majority of DICE users will not notice the switch-over.

However, anyone exploiting weaknesses in the existing forwarding mechanism, to continue receiving mail into their legacy mail boxes, will find that this no longer works; all mail will be delivered to their @inf mail box, or forwarded to a non-DICE email address if they wish.

Former users must forward legacy mail onwards: From the end of April, any mail addressed to old cogsci legacy users (for whom we don't have a forwarding address) will be returned to the sender as "user unknown". Likewise for any legacy pseudo-user ("nsu") accounts.

If you do have a cogsci mail address and have not yet provided forwarding information for it, contact **support**, who will make the appropriate arrangements. Making changes to the .forward file in your legacy cogsci home directory will have no effect.

Once the new service is in operation, the next step will be to kill off any remaining legacy IMAP and POP services. We'll provide more details of that, and of changes to old legacy "dcs" accounts nearer the time, and we'll contact those few remaining users individually.

Neil Brown <neilb@inf.ed.ac.uk>

Redhat 7 — EOL October

Now that the vast majority of DICE services and clients have moved to Redhat 9 and Redhat have withdrawn support for Redhat 7, the Computing Committee have agreed that we should EOL Redhat 7 by 1st October 2004. This date may be exceptionally deferred if serious problems are discovered.

As we deploy new releases of the Linux platform, we eventually withdraw support for the previous releases. We do this for two main reasons :

- ❑ Maintaining multiple versions of a platform is effort costly; new features have to be back ported, compatibility with new services has to be ensured etc. This cost becomes less easy to justify as we reach the stage where only a few services or clients continue to run the previous release.
- ❑ Support for the Linux distribution on which the previous Linux release was based is withdrawn by the distribution source; currently Redhat.

Support to be withdrawn in three phases: The Redhat 7 EOL will have three stages :

- ❑ Development Freeze — Immediate

This will have little noticeable effect to users of Redhat 7 services or clients. If your machine breaks, it will still be possible (though perhaps not desirable) to reinstall its replacement as a Redhat 7 machine.
- ❑ Untrusted — 1st July 2004

From this date onwards, Redhat 7 machines will be unable to access the network home directories. It will, however, be possible to create local home directories on the machines's own disk (as currently done for laptops). Users will still be able to access their network home directories from such machines using e.g. `scp`.
- ❑ Final — 1st October 2004

Although machines may continue to run after this date, they will be completely unsupported.

There are [more details on the web](#)¹¹.

If Applications won't port: One possibility for applications that will not port successfully to Redhat 9 is

¹¹see <http://www.dice.inf.ed.ac.uk/groups/infrastructure/linux/news/rh7eol.html>

to run Redhat 7 under VMware on top of a Redhat 9 machine.

Alastair Scobie <ascobie@inf.ed.ac.uk>

DCS Web Cache Transferred to EUCS Service: `squid.dcs.ed.ac.uk`, the legacy DCS web cache has been turned off. All the DICE services have been configured to use the EUCS [proxy web cache](#). Anyone with manually configured web browsers should update their settings to use the auto configuration file at <http://wwwcache.ed.ac.uk/config/proxy-config.pac>

DCS Web Service — EOL June: The legacy DCS web site, `www.dcs.ed.ac.uk`, is to be moved to the Informatics domain, served by a DICE managed machine. It will continue to respond to `www.dcs.ed.ac.uk` requests, however, due to the technical differences between the old DCS world and the DICE world an identical service will not be possible. Particular difficulties will be changes to PHP (ie from version 3 to version 4) and user run CGI scripts.

On a date yet to be decided, but hopefully before the end of June 2004, the content of `www.dcs` will be frozen and users will no longer be able to update `www.dcs.ed.ac.uk` pages. After that point any further updates will have to go via the web team, and they will only implement redirects to the new location of the content, at a `.inf.ed.ac.uk` style URL. More technical details of the move process will be available soon at <http://www.inf.ed.ac.uk/systems/web/olddcs.html>

Neil Brown <neilb@inf.ed.ac.uk>

legacy ssh services withdrawn

The older departmental remote access machines are being replaced (See page 4).

`ssh.dai.ed.ac.uk` has already been withdrawn, `ssh.dcs.ed.ac.uk` is no longer directly accessible from the internet and will soon be switched off. `ssh.cogsci.ed.ac.uk` and `ssh.cstr.ed.ac.uk` will be withdrawn before the end of May; anyone still dependent on them should apply for a DICE account so that they can use the new Informatics services. (Apply on the web¹². You should either be a staff member or get a staff member to sponsor you.)

Toby Blake <toby@inf.ed.ac.uk>

¹² www.inf.ed.ac.uk/cgi-bin/support-nodice.cgi

▶ Handling Microsoft Documents on DICE

What you use to handle Microsoft documents very much depends on their complexity and what you want to do with them.

If all you need to do is view and/or edit simple Microsoft Office documents, then **OpenOffice** should meet your requirements. It's installed on all DICE machines and can load and print Microsoft Office documents. The currently installed release of **OpenOffice** is v1.0, the new release, v1.1, is being tested and should be available on all DICE machines soon. The new release is much better at handling Microsoft Office documents.

We would expect OpenOffice to suit the majority of cases. It can be launched from the command line by typing `ooffice` or `ooffice <filename>`. The first time you run it some setup is done for you.

There are limitations with OpenOffice when it comes to handling complex Microsoft Office documents such as Excel spreadsheets with heavy macro usage or charting and some PowerPoint presentations and when "sharing" a document with a real Microsoft Office user.

In these cases there is currently not a DICE option available. However, we intend to make **Crossover Office** available in due course for this purpose.

The **Crossover Office** application allows the real Microsoft Office applications to be run natively under Linux so there are no problems with interoperability. However, to run this application a licence will be required by individuals for both **Crossover Office** and **Microsoft Office**.

MS Office compatibility	Open Office v1.0	Open Office v1.1	Crossover Office
View (simple)	YES	YES	YES
View (complex)	Not always	Much better	YES
Edit	YES	YES	YES
Print	YES	YES	YES
Availability	NOW	VERY SOON	SOON

Should you also need to run non-Office Windows-based applications, then you need to use **VMWare** which emulates the whole Windows environment under DICE. A licence will be required for this as well as for the Windows OS you wish to run and for Windows-based applications that require it. It is under testing now for the current version of DICE.

Alison Downie <alisond@inf.ed.ac.uk>

▶ Linux laptop support

One of the goals of DICE is to ensure that mobile computers are considered first class citizens in the DICE world. The most important aspect of this is the centrally managed DICE Linux laptop, of which we have around 65.

We are one of the very few sites in the world to attempt to centrally manage Linux laptops. Our aim is that a laptop running DICE should look identical to a desktop running DICE, whether running on an Informatics network or some foreign network, eg at home. This aim is a difficult one to achieve, for the following reasons :-

Linux on laptops Linux support for new hardware always lags significantly behind that for Windows. This is because most Linux hardware support is written by open-source developers rather than the hardware manufacturers. This is a serious problem for laptops as it isn't possible to swap out system components (eg graphics cards) for ones that are supported by Linux.

Another significant problem at the moment is that almost all currently available laptops implement the ACPI spec for power management. Unfortunately ACPI is only partially supported by the current Linux kernel (2.4); in particular, this means that suspend and resume are not possible with ACPI laptops.

Different contexts While desktops are installed in a fixed location with a fixed network connection, laptops are mobile devices and have to cope with a number of possible situations ("contexts"), eg. informatics network, EUCS wireless, home ADSL network, dialup, VPN, arbitrary conference wireless network etc. Various subsystems have to behave in different ways depending on the current "context", eg routing, mail, authentication, LCFG profile delivery, filesystem exports etc.

We can do little about the first set of reasons; one possibility we are considering is running DICE managed Linux under VMware on top of a managed or self-managed Windows laptop. The second set of reasons can largely be solved by applying CO effort, but perhaps we need to reconsider the priority of this work? Feedback very welcome.

Alastair Scobie <ascobie@inf.ed.ac.uk>

Linux platform non upgrade

We have decided, with Computing Committee agreement, to continue with Redhat9 as our base linux platform for 2004/2005. This will allow us to concentrate our efforts on improving the stability and usability of DICE.

Most of you will be aware that we aim to upgrade to the latest Redhat release each summer. We do this for a number of reasons :-

1. to continue to be able to receive security and bug fixes (Redhat drops support for old releases after a period of time).
2. to obtain support for new hardware (in particular, graphics cards and laptops).
3. to update software (particularly development related) required for teaching and research.

Until recently, Redhat produced a new release roughly annually. This release would be supported (in terms of security and bug fixes) for around 15 to 18 months. This fitted in well with our annual "summer" release cycle.

Redhat, however, have recently decided to split their Linux offering into two branches :-

- a. Redhat Enterprise (RHEL) — a subscription product designed with stability as its core attribute with a new release every 12 to 18 months and a support window for 5 years.
- b. Fedora — a free, community developed branch coordinated by Redhat with a new release every 4 to 6 months and a support window of 7 to 9 months. This branch will tend to be more "bleeding edge", often using developmental versions of software. It is primarily intended for the hobbyist.

The linux team recently met to discuss which of these branches we should choose as our linux platform for 2004/2005; we had already dismissed moving to a different vendor (eg Suse, Mandrake, Debian) this summer as we have insufficient time to develop support for a different vendor's platform.

The conclusion we reached from these discussions was that we should stick with Redhat 9 for 2004/2005, for the following reasons :-

1. RHEL, although attractive due to stability and support, ships without a significant amount of

software, particularly development related, which ships as part of Redhat 9 and Fedora. We would have to spend considerable effort in porting this additional software from either Redhat 9 or Fedora on top of the normal effort required to support a new release.

2. The current version of Fedora will cease to be supported in August, not long after Redhat 9 support has ceased. The next version of Fedora, shipping in April, is based on the latest 2.6 kernel; this has only just been released and is not yet considered sufficiently proven for serious service use.
3. The software version differences between Redhat9, Fedora and RHEL appear to be sufficiently small that the small advantages that might be gained from upgrading would be outweighed by the effort required.
4. The effort that would otherwise have been used to upgrade could be much better employed in improving the stability and usability of DICE.

Alastair Scobie <ascobie@inf.ed.ac.uk>

Redhat 9 on Laptops

The vast majority of managed laptops are currently running Redhat 7, with a few still running old DCS Redhat 6.2. We have recently, somewhat belatedly, completed DICE laptop support for Redhat 9. We now need to begin the process of upgrading all laptops to Redhat 9, particularly given the EOLs described in this newsletter.

As with all operating system upgrades, Redhat 9 is somewhat larger than previous releases. The disk footprint of a standard DICE machine is around 7Gb without teaching software (9Gb with), so a realistic minimum disk size is 10Gb. Memory and CPU requirements have not grown significantly; 256Mb and a 700Mhz PIII processor should be sufficient for most text processing requirements. Almost all laptops bought since Jan 2001 should be fine.

The following laptop models *will* run Redhat 9 :-

Dell Latitude C640

HP Omnibook 6000

Toshiba Portege 2100

Toshiba Satellite Pro 4360

Toshiba Satellite Pro 4600

Laptops which don't meet the minimum hardware requirements will be considered "self-managed" from the appropriate EOL date (1st October for Redhat 7 machines and 1st September for Redhat 6.2 machines).

The support team will shortly be contacting laptop owners to schedule upgrades.

Alastair Scobie <ascobie@inf.ed.ac.uk>

► System Configuration Research

The demand for Grid computing clusters is increasing the interest our expertise with large-scale system configuration.

For many years, the Edinburgh Computer Science Department managed a large network of Unix workstations and servers, using a locally-developed configuration management tool called LCFG¹³. A much enhanced version of this tool is now a fundamental part of the DICE technology.

Initially developed as a practical tool for managing the departmental network, it became clear that LCFG was based on a number of important principles which gave it several advantages for dealing with large and complex collections of machines.

These principles have contributed to a growing interest in the fundamentals of "large scale system configuration", which include a surprising number of difficult and interesting theoretical problems.

The development of very large Grid fabrics (~5,000 machines) and the spread of Linux on the desktop, has also increased the demand for tools and techniques to manage such configurations efficiently and correctly.

The practical experience, and the understanding gained from several years of live deployment (on ~1000 machines), means that Informatics has been well-placed to participate in system configuration research.

Apart from an informal research community¹⁴, this now includes several funded research projects, such as GridWeaver¹⁵ and OGSAConfig¹⁶, as well as RA/Phd funding for collaboration with the European DataGrid.

So, what is "System Configuration"?

The question is most easily illustrated by describing the problem facing a system administrator attempting to configure a new site (this may be a Grid fabric, or

a departmental facility with staff desktops and student laboratories). The raw material consists of:

- ❑ A large number of hardware boxes with various different properties, and (effectively) empty disks.
- ❑ Some repository of software, typically accessible over the network.
- ❑ A specification of the required behaviour of the final complete system — this is usually very informal, and comparatively high-level — for example, every desktop machine must "have email".

The task is to load the appropriate software onto the various machines, and set the individual configuration parameters so that the machines interact to produce an overall behaviour which meets the specification.

In addition, the system must evolve so that the specification is maintained, even if the specification, or the physical reality changes (things break).

To give an idea of the scale, the DICE installation involves about 1000 machines, each with around 5000 configuration parameters and 1500 software packages.

Of course, in practice, it is the complexity, and the relationships between systems which are more difficult to manage than the sheer scale.

The design of specification languages is a current major area of interest. These are basically data-description languages (not programming languages), but the application has particular requirements which lead to problems with composition of independent aspects, and constraint resolution, for example.

The technology for the deployment and maintenance of the configuration specifications is another important area, involving problems with distributed systems and peer-to-peer protocols.

So — System Configuration is a practical application with many unsolved problems, and we would certainly like to talk with anyone who may have relevant theory, or who would be interested in learning more.

Paul Anderson <dcspaul@inf.ed.ac.uk>

¹³see <http://www.lcfg.org>

¹⁴see <http://homepages.inf.ed.ac.uk/group/lsscconf/>

¹⁵see <http://www.gridweaver.org>

¹⁶see <http://groups.inf.ed.ac.uk/ogsaconfig/>