dice NEWSLETTER

January 2006

Welcome to the first computing newsletter of 2006. This edition contains articles on the proposed new filesystem AFS, which will replace NFS across the school and also explains how users can find out about the new DIY DICE service. Web and network news detail the consequences for users of some recent security scares, and there is news from the Support team on arrangements for support at Forrest Hill. Users will also find information on proposed changes to support for laptops.

Morna Findlay <morna@inf.ed.ac.uk>

Why Do We Need a New Filesystem?

The current network filesystem (NFS) is very long in the tooth, and has serious security and performance flaws. It is not really suitable for the conditions and environment in which we now find ourselves.

In its DICE implementation, it has to rely on additional technologies, many requiring significant local development and effort for tasks such as internetworking, automounting and partition management. In addition, NFSv3 no longer provides the features we require for many new initiatives such as DIY DICE.

NFSv3's lack of security makes it difficult for us to allow access to home directories by machines that are not running a trusted operating system, and impossible to allow access by machines that are not on internal, "trusted", network wires. In deciding on a new network file system, we wanted to make sure it satisfied as many of the following criteria as possible:

- a secure filesystem (authentication required for use)
- flexible access permissions for users
- use of features such as local caching
- widely used and actively supported
- compatibility with existing disk file systems
- compatibility with backup infrastructure
- user learning curve
- support for new initiatives (such as DIY DICE)
- better support for Windows and MacOS (allowing us to phase out Samba)

Why Did We Choose AFS?

Six possibilities were considered (OpenAFS, NFSv4, CIFS, DCE/DFS, Coda, & Intermezzo), the two strongest contenders being OpenAFS and NFSv4. Both of these appeared to meet our basic needs well, both offered enhanced security via Kerberos/ACLs, and both were available on FC3 and Solaris. Production quality was a vital aspect in determining the filesystem choice - we did not want to risk user data on a filesystem which is still developing, and has yet to mature.

Both file systems have provision for the use of Access Control Lists (ACLs), and both systems have the option of using kerberos-based authentication and this feature was enabled when performing the benchmarks. However, OpenAFS was (and is) the maturer product, and is well supported. NFSv4 as currently shipped with both Solaris 9 and RedHat's FC3 just isn't in a production-ready state.

Whilst testing the two systems, it rapidly became clear that OpenAFS was much closer to being a production file system. In particular, we were unable to test important features of NFSv4 such as per-user Kerberos authentication (not yet supported), and key portions of the filesystem became unstable under heavy load.

Although NFSv4 is a rapidly emerging standard, and on paper is an admirable competitor to AFS, few (if any) vendors have implemented all of the standards in production systems. In contrast, OpenAFS has a long legacy of stability, and is in widespread use at academic and corporate institutions worldwide.

OpenAFS has a number of other benefits. Native support is available for both Windows and MacOS, allowing users of those platforms direct access to their files. It is designed for use over wide-area networks, allowing usable direct-filesystem access from authenticated users at home, at other universities, even whilst at conferences in other countries. It features a robust client-side caching mechanism, which dramatically improves performance for common tasks such as compilation when compared with NFS. It no longer requires the use of an automounter, meaning that system hangs and reboots when fileservers become unavailable will become a thing of the past.

What is AFS?

Each organisation (or part thereof) that runs and administers an AFS service is referred to as a Cell, and the Informatics AFS service is one such cell. These cells are linked together to form a global, navigable hierarchy which is implemented as a filesystem tree conventionally rooted at **/afs**, with cells listed below.

AFS behaves like a standard Unix file system in most respects; it groups files into collections called volumes, and it is usual practice for each users' home directory to be on a separate volume. The Informatics AFS filespace starts at **/afs/inf.ed.ac.uk.**

The Informatics AFS service places user volumes (home directories) below **/afs/inf.ed.ac.uk/user/**, with letter subdirectories for staff, and letter-digit-digit subdirectories for students. Projects filespace is also provided and will live under **/afs/inf.ed.ac.uk/project.**

Note that AFS supplements the standard Unix file protection mechanism in two ways - it associates an access control list (ACL) with each directory (and ACLs apply only to directories - all files within a directory share the same ACL), and it enables users to define a large number of their own groups, which can be placed on ACLs. AFS uses ACLs to protect files and directories, rather than relying exclusively on the standard Unix file protection mechanism.

For further technical information, see the URLs mentioned in the Acknowledgements section.

How will it affect ME?

The plan for introducing AFS is already underway, and any staff or postgrad user can request AFS filespace (via the Support Form). This can - if required - be used as home directory space (in which case it can also be made available to that user on self-managed and other networked machines, with the correct software installed). We are particularly interested in how AFS can meet the needs of research groups and anyone who would like to discuss this further should also use the Support Form to register an interest.

Early-adopters of AFS will - permissions permitting - be able to interact with other users who have NFS home directories. There should be no day-to-day difference between NFS home directory users and new, AFS home directory users. It is not intended to completely eliminate support for NFS, but we would want to reserve its use for atypical situations and we would be very keen to remove it as a front-line file-support mechanism. A deadline for the completion of the move from NFS to AFS for all user filespace has yet to be fixed, and details of the move from pilot to full service will be set out in a future newsletter article.

Note that, during the interim period when not all data have moved to AFS, the visibility of home directories and projects filespace from self-managed machines will be limited - only AFS filespace will be visible. Users of self-managed machines could work around this by using shared AFS filespace, or mounting an AFS volume as a sub-directory of their home directory.

Acknowledgements

The above includes information from the AFS distributed filesystem FAQ at http://www.faqs.org/faqs/afs-faq/ and the AFS Administration Guide at http://www.openafs.org/pages/doc/AdminGuide.

The File Servers Team
<fileservers-team@inf>

Redhat 9 - End of Life

As we deploy new releases of the Linux platform, do de 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.

Computing Committee agreed, earlier this year, a policy on the duration of support for DICE platforms: http://www.inf.ed.ac.uk/admin/committees/computing/POLICIES/platformlength.html

Now that the vast majority of DICE services and clients have moved to Fedora Core 3 (FC3) and Redhat have withdrawn support for Redhat 9, it is time to implement that policy by deprecating the use of Redhat 9.

The Redhat 9 EOL will have two stages :-

Untrusted - 1st February 2006

- No bug fixes
- No security fixes
- Machines will be considered untrusted
- Machines will not be allowed access to NFS fileservers.
- Machines will not be allowed to be externally visible.
- Machine reinstalls still possible

From this date onwards, Redhat 9 machines will be unable to access the network home directories. It will, however, be possible to create local home directories on the machine's own disk (as currently done for laptops).

Users will still be able to access their network home directories from such machines using, for example, "scp".

It is expected that only a very few machines will be running Redhat 9 by this time - mostly servers.

Final - 1st March 2006

From this date, services to support the platform will be withdrawn. This will mean -

- No machine installs
- No RPM repository
- No profile generation

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

Alastair Scobie <ascobie@inf>

DICE Laptop Support

Fedora Core 3 (FC3) support for DICE laptops is now complete, with some caveats :-

- power management (ie suspend and resume) will only work on the most modern machines (eg IBM T41), and not on models such as HP 6000 and Dell C640s.
- FC3's disk footprint is significantly larger than that of RH9, so upgrading to FC3 will leave less space for your files.

As I reported in a newsletter article last year, we have been reviewing the cost effectiveness of supporting DICE on laptops.

The number of DICE laptops has been steadily declining over the years, for a variety of reasons :-

- the poor choice of laptops which would run DICE
- the availability of the Unix based MacOS and the improved support for Unix software under Windows (eg cygwin)
- the poor roaming capabilities of DICE

A recent survey of the remaining 12 academic DICE laptop users revealed that :-

- there is no common usage pattern amongst users
- most users could switch to Windows or Macs
- only a few users reported that they are using software that only runs under Linux or is awkward to run on anything other than Linux.

Our best estimate is that we spend around 0.25 FTE on DICE laptop support; our view is that this effort would be much better spent on improving support for self managed Windows and Mac laptops.

We will be proposing to Computing Committee that :-

- we continue support FC3 for existing laptops
- no new laptops should be installed with DICE
- support for laptops under the next DICE platform will be limited to that necessary to support the exam laptops.

Alastair Scobie <ascobie@inf>

DIY DICE

DIY DICE machines sit somewhere in between the centrally-managed (standard) DICE boxes and the self-managed machines. DIY DICE allows users to have full control over the configuration and administration of their desktops, while keeping many of the benefits of the centrally managed DICE service.

The centrally managed machines are completely administered and supported by Computing Support. However, the fact that users cannot directly intervene in its system administration, is a significant setback for some users. The alternative, self-managed machines, are fully administered by users. The disadvantage of self-managed systems is that being non-standard, they cannot be supported by computing staff. Self-managed machines do not benefit from the regular automatic software updates, and in case of failure there is no central support for rebuilding.

The aim of DIY DICE is to provide users with a system within which they can freely install software or services, whilst, if they choose, receiving the significant advantages of automatic software updates. Another key benefit of DIY DICE is that the underlying LCFG management system allows the configuration to be rapidly rebuilt in case of failure.

DIY DICE machines will be provided with some level of computing support, although not to the same extent as standard DICE boxes. DIY problems are likely to be of such complexity that they would consume disproportionate amounts of computing staff effort. A web forum has been set up to provide assistance for DIY DICE users (http://bb.inf.ed.ac.uk/boards/). Users can share their experiences, and can collaborate with solutions.Computing staff monitor this forum, however their involvement is not guaranteed.

Due to security constraints, NFS will not be available on DIY DICE machines. NFS will be substituted by AFS as the means to provide users with shared filespace. At the moment, AFS is a pilot service and is available on all DICE and DIY machines, allowing users from both systems to share their filespace.

Finally, there is a wiki section for DIY DICE (https://wiki.inf.ed.ac.uk/DiyDice) with more information on this topic.

How Do I Switch to DIY DICE?

Contact Computing Support. Users are provided with individual assistance to start them with DIY DICE. Individual requirements will be fed back into the DIY development and assessed by the developers.

Julieta Pineda <julieta@inf>

Mail News

Mail.inf FC3 upgrade

The previously announced upgrade of the main Informatics mail service from Redhat9 to Fedora Core 3, just prior to Christmas, did not happen. It will be happening before the end of January and may well have happened by the time you read this.

The main difference will be a new version of the web interface to mail.inf.ed.ac.uk, due to an upgrade from IMP3 to IMP4. This is to mirror the current interface running on the EUCS staffmail mail service. (https://www.staffmail.ed.ac.uk)

Mail Quotas

Another change with FC3, will be implementation of a quota for mail files. Again this is to try and keep us inline with the EUCS staffmail service. We are doing this, as we continually review the possibility of us ceasing to provide an Informatics mail service, and use the centrally provided mail service instead. The more like the staffmail service our mail service is, the easier it will be to migrate from one service to the other.

The default quota will be set at 200MB (just like staffmail), and should you reach this limit, mail will not be delivered to you until some space is freed up, eg deleting old mail, or moving it into your home directory. Mail messages will not be lost, just delayed, or returned to the sender (with an explanation why) if it has not been delivered after 5 days.

For those users (about 10% of you) who are already over (or near) the 200MB limit, your quota will be set to about 15% more than you are currently using to allow you to continue to receive mail. However, it will be expected that you'll try and reduce your usage to within the 200MB limit. The handful of individuals who are well over the 200MB limit, will be contacted individually.

Neil Brown <neilb@inf>

Web News - CGI Security

This is a reminder about how important it is for authors and users of CGI scripts to make them as secure as possible, and/or apply the latest security fixes for any software they may be using.

In a recent incident, a user owned Wiki on homepages.inf was exploited and used to send bulk SPAM via our mail relays. Due to the way we execute CGIs on homepages, the rogue process was running as the owner of the CGI, which somewhat

limits the damage, but the real danger here was that our domain "mail.inf.ed.ac.uk" gets blacklisted as a spammer friendly domain, and mail servers around the world stop accepting mail from @inf.ed.ac.uk addresses.

We are all governed by the Computing Regulations (http://www.ucs.ed.ac.uk/EUCS/regs.html), which says that private use of the computing facilities is allowed, but not a right, and that breaches of these regulations may lead to disciplinary action. Bringing the University into disrepute is a breach.

If you are running 3rd party CGI software, please make sure that you sign up to any "announcements" or "security" related mailing lists associated with that software. Hopefully you'll then be alerted to any known vulnerabilities with the software, and be able to take remedial action to fix it.

New Workshop URL

Similar to the existing http://groups.inf.ed.ac.uk/ and http://conferences.inf.ed.ac.uk/ services, a new URL http://workshops.inf.ed.ac.uk/ has been created. So if you are organising any workshops, and are looking for a place to put some web pages, hopefully this will be of some use. Requests should be submitted via the support form as normal.

Neil Brown <neilb@inf>

Support Team News

Staffing

As most of you will know, Shehzad Ali, one of our CSOs, has now left to take up employment elsewhere and Sarah Reed has not yet returned from her maternity leave. As a result, until further notice, the Forrest Hill Support Office (A05) will only be manned in the mornings as follows:

- Monday Lindsey
- Tuesday Carol
- Wednesday Ross
- Thursday Charlie
- Friday Richard

Consequently the Buccleuch Place Support Office will not be staffed on Wednesday mornings nor will the Appleton Tower Support Office be staffed on Thursday mornings.

Requests and queries can still, of course, be entered using the support form at any time and will be prioritised and dealt with by the Frontline Support team as soon as possible.

RT Statistics - 2005

The number of queries received (via the Support Form) in 2005 was as follows:

- queries received 5143
- resolved 4778
- open 337
- on hold 18

Visitor Accounts

In a recent newsletter, we explained the revised procedures for obtaining Visitor Accounts. This has now been published and can be found at:

www.inf.ed.ac.uk/admin/policy/visitors.html

Support for Workshops and Conferences

The level of computing support staff organising conferences and workshops can expect, and how this support should be requested can be found at:

http://www.inf.ed.ac.uk/admin/committees/computing/POLICIES/conferences.html

Alison Downie <alisond@inf>

Network News

A recent incident where a Professor in another School connected an unregistered and misconfigured WAP (Wireless Access Point) to EdLAN resulted in our School, and the School of Engineering and Electronics, being disconnected from EdLAN for nearly 24 hours.

Unfortunately, current network technology is such that the University cannot completely design EdLAN so that mis-configured network equipment can not impact other users on the network, whilst being flexible enough for the University's academic needs. Instead, it has to rely on users acting responsibly - in this case, understanding that EdLAN is a complicated network and that adding network equipment should be left to those computing staff who understand the issues.

EUCS are attempting, where possible, to ensure that, when incidents like this do occur, the problem is localised at the School level. Where possible, this will happen automatically; the EdLAN routers will drop a School's EdLAN connection(s) if they determine that some equipment downstream of that connection is compromising the EdLAN service.

Alastair Scobie <ascobie@inf>