A New Computing Infrastructure for the Division of Informatics

by Paul Anderson <paul@dcs.ed.ac.uk>
Alastair Scobie <ajs@dcs.ed.ac.uk>

Division of Informatics
University of Edinburgh

1 Introduction

Previous proposals for improving the integration of Computing Facilities within the Division of Informatics have concentrated on merging the ex-departmental systems (See the paper *Computing Facilities in the Division of Informatics*[1]). Many of the suggested changes would be difficult and expensive to perform on a "live" network, and the inevitable compromises would make a poor foundation for a modern, sustainable computing facility.

It has been suggested that it might be preferable to construct a completely new "Informatics Computing Infrastructure" (*InfoLan*?), from scratch, and to gradually migrate existing users and services as the new facilities become available. This paper is an attempt to identify the main components that would be necessary for such an infrastructure, and to outline some of the choices to be made, and the possible resource requirements.

2 Summary

- Development should concentrate (at least initially) on a maintainable state-of-the-art infrastructure for commodity computing.
- Resource requirements are almost impossible to quantify because they depend on the extent to which appropriate staff can be made available, and on a number of design choices. However, it is unlikely that anything useful could be achieved with less than about two to three appropriate people for 18 months.
- User-level services, such as mail, news and printing should remain on the existing systems until the core is complete.
- In many cases, there is no perfect technology and difficult decisions will required which will inevitably be a compromise.
- An ongoing commitment will be required to maintain a state-of-the-art system.

3 Why a New Infrastructure?

Most of the existing systems within the Division are build around an infrastructure which was developed about ten years ago. Many of the fundamental assumptions on which this was based are no longer true; we can no longer assume that machines are always connected to the same network (or connected at all!), or that machines are managed by a small group of trusted professionals. The increase in connectivity has also drastically changed the way in which the facilities are used. A new infrastructure would allow us to support these new styles of working and provide a good foundation for the next ten years, which is not possible with the existing system.

Over the last ten years, the availability of hardware has increased enormously, but the availability of staff has decreased. Whereas it was once typical for one person to manage a single Unix machine, 50-100 machines per person is now normal. However, there is definitely an increase in the hidden costs incurred by other staff members managing their own personal machines. We have already made considerable progress in reducing this TCO (*Total Cost of Ownership*) but new infrastructure would allow us to reduce this even further.

Much innovative work has been done in the ex-departments on system management, but there has rarely been time to export this to other users. By taking advantage of this experience when building a new system, it should be possible to design the technology so that aspects of it can be exported to other installations within the University. If some of this technology is adopted by the wider "open source" community, then the ongoing maintenance and development costs can be considerably reduced.

[1] http://www.informatics.ed.ac.uk/admin/committees/computing/meetings/99-05-04/integration.html

4 Scope

We consider the following aims to be fundamental, and these have been used to guide the suggested developments:

1. The new infrastructure should provide the foundations for a maintainable, state-of-the-art computing facility for the Division of Informatics.
2. Maintainability and low TCO should have a very high priority.
3. Reliability, robustness and flexibility should continue to be important goals.
4. The increasing importance of security should be acknowledged by explicit policy statements and an infrastructure which provides an agreed compromise between useability and security.
5. Support for new working practices, such as mobile computing and tele-working should be an integral part of the design. This includes support for students (and staff) to run compatible systems on their own machines.
6. The infrastructure should support various levels of devolved system management in a secure way. This includes the ability for research groups to run their own specialised systems, while still taking advantage of the underlying infrastructure. It also recognises that distributed management of commodity computing within the Division is likely to continue in the foreseeable future.
7. Given limited resources, emphasis should initially be placed on the provision of high-quality, “commodity” computing facilities which benefit the majority of users.

4.1 Platforms

In keeping with aim [7] above, we are assuming that the new infrastructure should initially be targeted primarily at Linux on PC hardware as the commodity platform. However, since flexibility is a major consideration, the design must be capable of supporting other Unix platforms, and these would be included from the start. Actual implementations for some of these (for example, Linux on Alpha) may appear in parallel; others (such as Solaris) are likely to require more implementation effort and will probably not be available initially. Note that *Commodity* computing refers to non-specialist applications, such as text processing and email, where differences between the platforms are not usually significant to the end user.

Windows NT represents a significantly different, and difficult platform that resources are unlikely to be available to include specific support for this operating system. Management of NT machines is probably best handled in other ways (for example, using EUCS Technology[2]), although we would hope to provide integration with NT systems and consider their special requirements wherever possible.

We do not believe there is sufficient demand for other systems (such as Apple) to be worth significant consideration.

4.2 Layers

For the purposes of design and development, it is useful to consider three main *layers*:

- *Low level network components* (such as cabling, network topology and management of the traditional network services) are an essential foundation for a reliable computing facility, but they can be considered largely independently of the other layers. A prototype *InfoLan* could probably be built on top of the existing low-level facilities, although this would not be suitable for a production environment.
- A set of *Essential Services* form the core of the infrastructure architecture which is necessary before any hosts at all can be supported. These include names services, user account management, machine configuration management and software distribution. A distributed file service is also necessary to support both these services, and real users. This layer would form the bulk of the design and development work, since most of these technologies in the ex-departmental systems are unsuitable as the basis for a new infrastructure.
- *Application Services* are important user services which are largely independent of the core infrastructure. In most cases, these could probably be constructed in parallel with work on the core services. However, in practice, the availability of resources is likely to mean that many of these facilities would remain on the existing systems until the core of the new infrastructure is in place. This includes mail, news and printing, for example.

5 Low Level Network Components

A number of important decisions will have to be made about cabling, topology, and technology for the new

[2] http://celia.ucs.ed.ac.uk/presentations/edwin/default.htm

network. However, as mentioned above, many of these decisions are comparatively independent and are not likely to require significant development work. Some areas which will require investigation, and possible development, include:

5.1 Network Management:

At present, we have very little traditional network management. However, network technology is moving from “passive” to “active” devices, where the topology of virtual networks is established by software, rather than by physical connections. It is crucial that we can configure and maintain these devices as effectively (if not more so) than the hosts themselves.

5.2 Firewalls

Firewalls on the existing networks provide some degree of security using various ad-hoc techniques for controlling traffic. The goal of supporting a large, secure network with various levels of devolved management, demands a more coordinated security policy, implemented with dedicated firewalls.

5.3 Mobile Computing

We believe that mobile computing is going to become very important and we intend support for this to be an integral part of the new system. This requires investigation of services such as dial-up access, mobile-IP and DHCP.

6 Essential Services

All of the following services are necessary before any new infrastructure can support client machines. Some of these require difficult decisions on the most appropriate technology; others require significant in-house design and development work. Direct re-use of any technology from the existing systems is unlikely to be appropriate.

6.1 Name Services

We use term *Name Services* to refer to all the technologies which support the various distributed databases containing essential system information. Two technologies are currently in use throughout the Division: DNS and NIS (NIS+ at BP).

It is likely that we will want to consider replacing at least some of these with a more modern technology such as LDAP.

6.1.1 DNS: is used largely for hostname lookup, but it also supports Hesiod which is required for printing and NFS automounter maps. DNS would be crucial to any new implementation and the basic technology is well supported, requiring little additional development work. However, two areas will require implementation effort:

1. Some technology is required to provide distributed editing and access control for the DNS source files. It is unlikely that the technology in use at any of the existing sites will scale. It may also be appropriate to take DNS information from a central configuration database (See [6.4]) instead.

2. If DNS continues to be crucial for printing and filesystem access, then some method is needed to maintain and update DNS servers on disconnected machines, such as portables.

6.1.2 NIS: is currently in use for a number of different purposes, including:

1. User account information.
2. Machine configuration data.
3. Netgroups (mostly for access control).
4. Information for booting machines (ethers and bootparams).
5. Host information for the local domain (duplicating DNS).

NIS is not nearly so widely accepted as DNS and we would want to consider replacing it with some other technology. There are performance and security issues with the existing mechanisms, and it is not well suited to disconnected operation or self-managed machines. Alternative mechanisms would need to be found for some, or all, of the above cases. It would also be necessary to implement some way of maintaining the source data, but it is likely that this could use the mechanism developed for the DNS (See above).

6.2 User Accounts

User account information is currently stored in NIS maps. Different sites have different mechanisms for managing these accounts, including procedures for adding and deleting users, either individually, or in bulk (from MIS data). Several issues need to be addressed:

1. A mechanism is needed to replace NIS for distribution of account information securely between machines.
2. A mechanism is needed for secure, distributed account management, both for individual accounts and for bulk-management of student accounts. It may be, for example, that this is linked to the Divisional database in some way.
3. Some mechanism is need to replace the NIS netgroups for secure access control information.
4. It is likely that we will want to separate the issues of account management and user authentication which are currently tied together by the concept of a Unix password file.

There also needs to be a clear policy for allocation/de-allocation and naming of accounts. This requires development of appropriate administrative procedures.

6.3 File Service

Highly-developed distributed filesystems form an integral part of the ex-departmental systems (See CS-TN-21[3], for example). The current implementations have a number of fundamental problems which make them weak points of the existing systems, and unsuitable as a basis for any new infrastructure. These filesystems were created at a time when most machines had small disks (or none at all), were centrally-managed, and were permanently connected to the same network. These are no longer good assumptions, and deciding on an alternative to the current file service is probably the single most difficult design issue.

Current remote filesystem usage falls largely into three categories:

- User home directories.
- Shared directories (for example, package sources).
- Program binaries.

Using modern Linux technology, and clients with large disks, the third of these requirements is considerably reduced, and the structure of any virtual filesystem could be greatly simplified. However, some form of distributed filesystem will still be required.

Security developments in NFS have not been sufficient for it to form the basis of a Division-wide virtual filesystem. If NFS is used, it is likely to be within smaller, trusted “islands” (See George Ross’ security paper[4] for more details).

Some new developments, such as Coda[5] look promising. This builds on concepts of the Andrew filesystem and provides support for disconnected operation as well as many other useful features.

Some work is required to evaluate the options and any solution will probably be a difficult compromise. Some decisions on hardware are also required, such as the number and location of servers, reliability and replication (RAID?) issues, and performance.

6.4 Machine Installation & Configuration

The existing ex-DCS installation and configuration technology[6] has been well-proven and the *principles* should scale well to a Divisional level. However, many aspects of the existing *implementation* were only ever intended to be temporary, and these will require re-implementation to be suitable for wider use. We would also want want any new implementation to provide better support for self-managed, and disconnected machines, which is missing from the current implementation.

6.5 Software Distribution

Software distribution under Linux is handled well by the existing ex-DCS technology (*updaterpm*), and this should scale to Division-level without too much effort. A completely different mechanism[7] is used under Solaris; this is outdated and unsuitable for a new infrastructure. Support for Solaris would require significant work, perhaps involving a port of the Linux technology.

6.6 Backups

Backups are currently handled using a wide variety of tools, and there are a number of problems that we would like to address; reliability of backup hardware, easier location of files on backups and (self?) restore. We would also like to provide support for portables and self-managed machines, including other platforms.

6.7 Authentication

Security can not really be considered as a separate issue; all services need to be aware of security implications. However, *authentication* is the process by which

[3] http://www.dcs.ed.ac.uk/doc/Users/21/

[4] http://www.dcs.ed.ac.uk/~gdmr/MergeSec.dvi

[5] http://www.coda.cs.cmu.edu/ljpaper/lj.html

[6] http://www.dcs.ed.ac.uk/~paul/Publications/LISA8_Paper.pdf

[7] http://www.dcs.ed.ac.uk/~paul/Publications/LISA5_Paper.pdf

a user identifies themselves to the system, and this can be thought of as a independent service. We need to consider whether it is worthwhile implementing some pervasive authentication infrastructure, such as Kerberos, or whether more ad-hoc individual technologies will be used.

7 Application Services

The following application services are largely independent of the core services. These should probably remain on the existing systems until the core of the new infrastructure is in place, and they have therefore not been considered in as much detail:

7.1 Printing

Would probably be based on LprNG technology currently being developed in ex-DCS. This may require some work to move onto a new infrastructure, depending on the adopted name services.

7.2 Mail

Is likely to continue to be based on Sendmail. A number of peripheral issues will require more attention than the core mail service itself; for example, user names and aliases, and mailing list maintenance and archiving. Student (undergraduate) email would be likely to transfer to SMS once EUCS provide POP or IMAP access.

7.3 News

A local news service may or may not be required. Information dissemination in general is a major requirement though and this is likely to involve some work on News, Mail and/or Web technology.

8 Policies

In addition to the technical developments necessary for a new infrastructure, there is a definite need for the Division to develop explicit policies on several aspects of its' use. The lack of such policies in the past has often defeated the technical efforts to provide an effective service. Some of these include:

- A Security Policy is required to determine exactly who should be permitted to do what. Without this, any technical solutions are worthless and all data on the network should be considered public.
- An Ethical Policy is required to determine the extent of system manager's "power". With a large and distributed management group, many people will have the ability, for example, to read any user's mail.
- System Management may need to be more controlled, with more explicit documentation and, for example, more coordination over releases of new software versions.

9 Resourcing

We are extremely reluctant to attempt to quantify the resources required for the development of a complete new infrastructure, for several reasons, including:

- Much of the work is highly specialised and relies on the availability of staff with the appropriate skills and experience. We believe it is unlikely, in practice, that such staff can be released to work completely free from other commitments.
- Decisions need to made about the point at which quality is traded for speed of development. In some cases, it may be possible to provide visible results at an earlier stage by using "temporary" measures, but this should only be done where there is a genuine commitment to replace these as soon as possible.
- One of the reasons why such a large amount of work is now required, is that insufficient resources have been allocated for continuous development and updating of the existing systems. This type of evolution is essential to maintain a state-of-the-art infrastructure in an area where technology changes so rapidly.
- This is notoriously difficult, and any estimate will probably be wildly wrong!

However, the following is a very rough guess at the absolute minimum resources required (in person-months) to develop the core services to the point where they could support a small client community. There would probably be as much work again in bringing these services to a sufficient standard to support a full Informatics-wide infrastructure, and an ongoing commitment would be required to maintain a state-of-the-art system. This also assumes real commitment from suitable staff, and does not cover the additional "user services".

Switch configuration and network management	4
Security policies and firewalls	4
Name Service (LDAP?)	4
DNS configuration and management	3
NIS passwords and access control	4
Account creation and management	4
Machine configuration	6
Software distribution	3
Distributed file service	4
Backups	2
Security Infrastructure	4
General Integration	6

The above table deliberately includes no total, since it would be misleading to interpret this as a timescale by which a production infrastructure could be operational; even given the pre-requisites necessary to complete the individual components within the given timescales, it is not clear that the development effort could be sufficiently sustained to complete all of them within the sum of these times.