Free Software Infrastructure in Informatics

Simon Wilkinson <sxw@inf.ed.ac.uk>
Stephen Quinney <squinney@inf.ed.ac.uk>

Disclaimer

Nothing herein is warranted or guaranteed. This product is meant for educational purposes only. Any resemblance to real persons living or dead is purely coincidental. Any mention of commercial products is for information only; it does not imply recommendation or endorsement.This page is netfake enhanced. Void where prohibited. Some assembly required. List each check separately by bank number. Batteries not included. Contents may settle during shipment. Use only as directed. No other warranty expressed or implied. Do not use while operating a motor vehicle or heavy equipment. We do not warrant or assume any legal liability for the accuracy, completeness or usefullness of any information, apparatus, product or process discussed. Do not look into laser with remaining eye. Postage will be paid by addressee. Subject to CAB approval. This is not an offer to sell securities. Apply only to affected area. No postage necessary if mailed in the United States. Please remain seated until the ride has come to a complete stop. Breaking seal constitutes acceptance of agreement. For off-road use only. As seen on TV. One size fits all. Many suitcases look alike. Contains a substantial amount of non-tobacco ingredients. Colors may fade. We have sent the forms which seem right for you. Lots of other people were involved in the design and implementation of these systems. Slippery when wet. For office use only. Not affiliated with the American Red Cross. Drop in any mailbox. Edited for television. Keep cool; process promptly. Post office will not deliver without postage. List was current at time of printing. Return to sender, no forwarding order on file, unable to forward. Not responsible for direct, indirect, incidental or consequential damages resulting from any defect, error or failure to perform. At participating locations only. Not the Beatles. Penalty for private use. See label for sequence. Substantial penalty for early withdrawal. Do not write below this line. Falling rock. Lost ticket pays maximum rate. Your canceled check is your receipt. Add toner. Place stamp here. Avoid contact with skin. Sanitized for your protection. Be sure each item is properly endorsed. Sign here without admitting guilt. Slightly higher west of the Mississippi. You must be present to win. No passes accepted for this engagement. No purchase necessary. Processed at location stamped in code at top of carton. Shading within a garment may occur. Use only in a well-ventilated are. Keep away from fire or flames. Replace with same type. Approved for veterans. Booths for two or more. Check here if tax deductible. Some equipment shown is optional. Price does not include taxes. No Canadian coins. Not recommended for children. . List at least two alternate dates. First pull up, then pull down. Call toll free number before digging. Driver does not carry cash. Some of the trademarks mentioned in this product appear for identification purposes only. Objects in mirror may be closer than they appear. Record additional transactions on back of previous stub. Unix is a registered trademark of AT&T. Do not fold, spindle or mutilate. No transfers issued until the bus comes to a complete stop. Package sold by weight, not volume. Your mileage may vary. for any utility. Provided as is, with no express or implied warranty, except that provided by the law. If you don't like all this, parents can exercise your discretion. Simulated Picture. Photograph enlarged to show texture. These are performed by professionals. Do not try this at home. Not a toy, keep far away from children. If you are a child, ask your parents to keep you away from this .. and yes .. either parent will do, there no need to get both to do it, so you have no excuse for using this as a toy. Not a spermicide. No user serviceable parts inside. Conforms to FCC part B specifications for Spurious emissions. Fasten seat belts. Its not a good idea, its the law. Recycle and save the world. This mail written entirely with recycled electricity, 100%post-consumer, with vegetable inks. Void when printed on any printer. Whatever is not yet degraded is Bio-degradable. 99%cholesterol free. with 70%less fat than Land-o-lakes. You can't believe its not butter .. you don't have to. Non refundable, non transferable. The first washable waterproof mascara. The surgeon general has determined lots of things. This is not a joke .. if it was, you'd be laughing. This message may not be copied in any form of cover other than that originally sent in, and without such a condition being imposed on the subsequent copier. Including all color copiers too ... and remember, I did tell you that one of the guarantees above gets void when you print or fax. Not responsible for clothes left behind. The University doesn’t pay us enough to own our words - all opinions herein are our own. Don’t quote me on that. Don’t quote me on anything. Clothes left in dryer may be removed by next customer. Clothes right in the dryer may not be. Do not put clothes soaked in gasoline in the dryer. Gone for Lunch. No admission without permission. Good for your skin .. tested in a Swiss lab. No added salt or sugar. 0%sodium (acc. to the FDA, 0%sodium if less than so much sodium) .. uses potassium and radium instead. Void where prohibited. Offer expires tomorrow. This program was recorded live and edited for brevity. As seen on TV. Real Psychic - Don't sue, I'd know first. We don't care. Real bagel. real bagels are made fresh by hand. Machine made - untouched by hand. No preservatives. No artificial ingredients. All natural. the synthetic vanilla has been made from naturally occurring coal, methane, and organic compounds. Freshly reconstituted from concentrate. This unit not tagged for individual sale. Are you reading at all? This tag not to be removed under perjury of the law. Murder scene - Do not cross this line. Wheelchair accessible. Politically correct. This perm is guaranteed for life. For the life of the perm. Sale of cigarettes to persons below 18 prohibited, so if you are 18 don't push that special button that is low down and within your reach. We had to put it there to make this unattended vending machine wheelchair accessible, to comply with ADA. Offer not valid with any other offer. No expiry date. Coupon valid until actually presented. No more timeshare presentations. You aren't required to like this talk. Don't ask, don't tell, don't pursue. Mandatory 5day waiting period - waived for cash. Needles and rubbers perhaps, but Bullets will not be distributed in schools. Asking permission constitutes harassment. Between teacher and the taught is unethical and the board distances itself from the statements of the President. Contains some violent scenes. UL certified. Guarantee void if bar code removed. Stop - by opening this you are agreeing to everything. Valid only in continental US - not valid in Alaska, Puerto Rico Hawaii or Canada. Contains cryptographic code. Do not ftp outside the US. All responsibility that of the ftp-er. Do not recharge or swallow. Linux is a registered trademark of Linus Torvalds. Liability limited to replacement. Void if mixed with other types. This has been written entirely in ASCII. No EBCDIC or animal fat. Coconut Cookies - Two for the price of one Special - No tropical oils. Information on cholesterol is provide for those who are modifying their dietary intake under the advice of their doctor. Perky perky ... why did you read it? Actually none of this was designed to be read. It was meant to be read of course, don't mistake me. Only not designed to be read. Or rather designed to be not read. Directions: Use as desired. All wrongs re-served. All flames to /dev/null. No shirts no shoes no service. Affirmative Action/ Equal Opportunity. Specifications subject to change without notice. Do not puncture, incinerate or store above 120 degrees Farenheit.Disclaimer does not cover misuse, accident, lightning, flood, tornado, tsunami, volcanic eruption, earthquake, hurricanes and other Acts of God, neglect, damage from improper reading, incorrect line voltage, improper or unauthorized reading, broken antenna or marred cabinet, missing or altered serial numbers, electromagnetic radiation from nuclear blasts, sonic boom vibrations, customer adjustments that are not covered in this list, and incidents owing to an airplane crash, ship sinking or taking on water, motor vehicle crashing, dropping the item, falling rocks, leaky roof, broken glass, mud slides, forest fire, or projectile (which can include, but not be limited to, arrows, bullets, shot, BB's, shrapnel, lasers, napalm, torpedoes, or emissions of X-rays, Alpha, Beta and Gamma rays, knives, stones, etc.). Other restrictions may apply.

Overview

- Whistlestop tour of Informatics's computing environment
  - Contexts & philosophy
  - Infrastructure & Middleware
  - Authentication & Authorisation
  - Directory Services
  - Machine Configuration
  - File Services
  - Putting it all together
- Necessarily high level - ask for more detail if you are interested
- Questions taken at any point...

Contexts

- ~4600 user accounts (~4100 currently active)
  - ~800 staff
  - ~3200 students

- Approximately 2500 hosts (1400 managed)
  - ~ 800 desktops
  - ~ 400 public lab machines
  - ~ 160 servers

- 4 geographically distinct sites
  - Kings Buildings
  - Appleton Tower
  - Forest Hill
  - Buccleuch Place

Philosophical Issues

- Services should be accessible from any platform and from anywhere
- The network shouldn’t be trusted (both internally & externally)
- Machines should be resilient to both local and remote network disconnection
- There will always be non-web delivered services. Web-only solutions are not appropriate.
- All lab and desktop machines should have an identical configuration, unless there are specific requirements otherwise.
- Systems should be based around Open Standards

Infrastructure models

- Traditionally ....

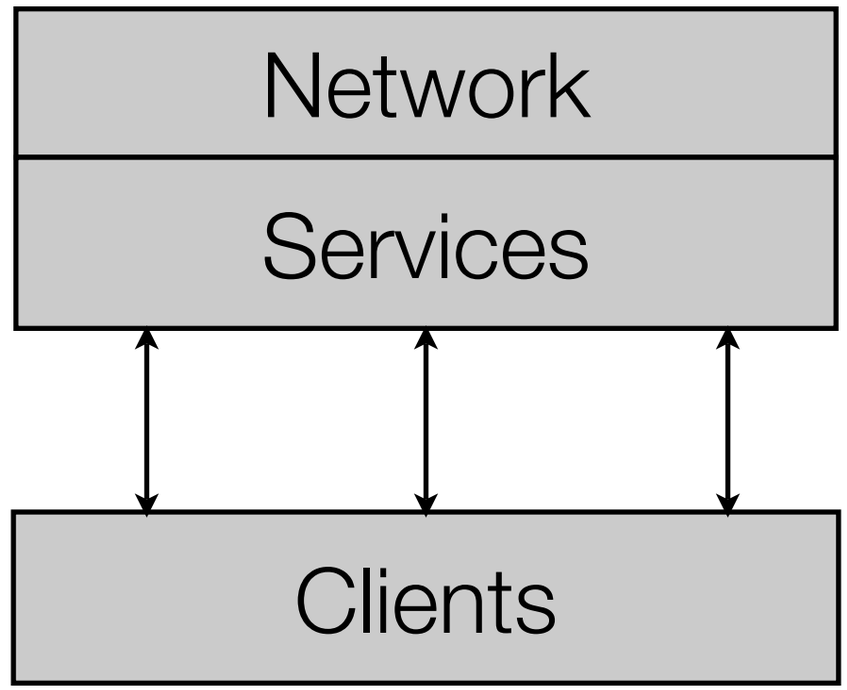

Infrastructure models

- Now ....

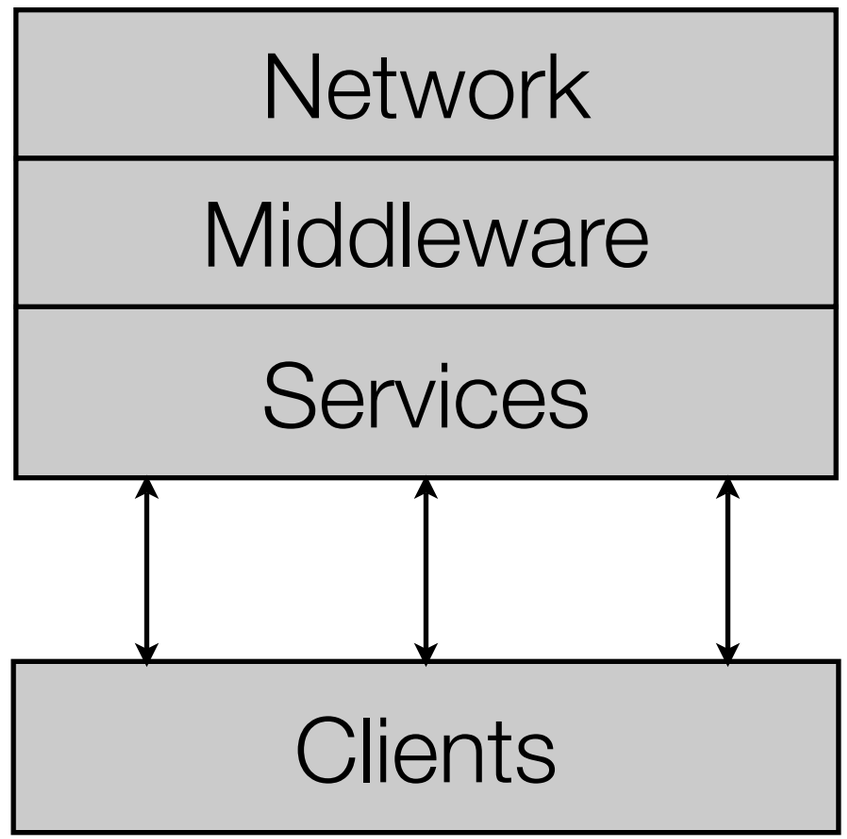

- Middleware is the glue that binds services together

Networking

- Information Services provide networking between our sites
- UKERNA/JANET provide wider network access
- We do all of our own internal networking
- Switches automatically configured from text based configuration data
- Firewalls automatically configured by our configuration management system
- Linux on all firewalls and routers

Managed Machines

- Provide an operating environment named DICE
- Fedora Core (some 5, now mainly 6) based
- Still have some Solaris servers, but they're a dying breed
- Looking at moving to Scientific Linux 5 for servers in the near future
- Predominantly use an 'off the shelf' distribution, so aren't going to talk much about it.
- However, we have a large number of local packages ...

Local packages

a52dec
a52dec-devel
aclisp
AcrobatReader-chsfont
AcrobatReader-chtfont
AcrobatReader-jpnfont
AcrobatReader-korfont
acroread
AdobeReader
AdobeReader_enu
aiai-projects
alice-complete
allbery-pam-krb5
allbery-pam-krb5-debuginfo
alpine
amrnb
am-utils
aopalliance
aopalliance-javadoc
argouml
argouml-menu
asm
asm-javadoc
aspectwerkz
aspectwerkz-demo
aspectwerkz-javadoc
aspectwerkz-manual
asrmanpages
atftp
atftp-debuginfo
atftp-server
atrium
bash
bash-completion
bash-debuginfo
bashdefenv
bashdefenv-debuginfo
bea-stax
bea-stax-api
bea-stax-javadoc
bib2html
bibtex2html
biojava
biojava-javadoc
bow
busybox
bytecode
bytecode-javadoc
c4.5
cglib
cglib-javadoc
claimform
cluster
cogent
condor
condor-6.7.20-linux-x86-
glibc23
conserver-client
conserver-debuginfo
conserver-server
coq
coqide
CorpusWorkbench
cosign
cosign-client-apache1
cosign-client-apache2
cosign-debuginfo
cosign-friend
cosign-html
coverlipse
crossover-pro
cross-pic30-elf-binutils
cross-pic30-elf-gcc
cryptyc
cryptyc-javadoc
ctat
cups
cups-debuginfo
cups-devel
cups-libs
cups-lpd
curt
cvs
cwb
cxbottle-managed_win2000
cyclone
darcs
defenv_bash_genrep
Dellmgr
detex
devproj
dhclient
dhcp
dhcp-debuginfo
dhcp-devel
dice-accntmgr
dice-accntutils
dice-afsutils
dice-apache
dice-apache-devel
dice-apache-manual
dice-apache-suexec
dice-apcpdu
dice-authorize
dice-authorize-debuginfo
dice-authorize-server
dice-bookmarks
dice-conference-kiosk
dice-conference-kiosk-
defaults-s1
dice-config
dice-contact
dice-cosign-html
dice-cosign-templates
dice-credchecker
dice-dcs-Xstartup
dice-dcs-Xstartup-debuginfo
dice-diy
dice-dns
dice-fvwm-startup
dice-groupsync
dice-ifriend
dice-iptables
dice-kx509-cert
dice-lcfgldap
dice-lcfgldap-defaults-s2
dice-ldapamdmaps
dice-ldapamdmaps-server
dice-logo
dice-LPRng-perms
dice-mailserver-utils
dice.mod_auth_kerb
dice-netman
dice-netman-defaults-s1
dice-nut
dice-openvpn
dice-orders
dice-partsync
dice-pcaputils
dice-quotas
dice-quotas-defaults-s1
dice-release
dice-rfemapsync
dice-slapd-tickle
dice-snmptrap
dice-snmptrap-debuginfo
dice-sslcerts
dice-styles
dice-submit
dice-TEC
dice-xdm
dice-xdm-debu…o
dice-x…
diskh…
dmi2xml
dmi2xml-debu…
dom4j
dom4j-demo
dom4j-javadoc
dom4j-manual
dresden-ocl
dresden-ocl-ar…
dresden-ocl-javadoc
dtdparser
dtdparser-javadoc
dvi2tty
eclipse
eclipse3.2
eclipse3.2-rcp
elfingrs
emacs-haskell-mode
emosml
enscript-conf
erg
esterel
etherape
etherape-debuginfo
eucs-sslcerts
eventd
examlaptop
exmh
exmh-debuginfo
faac
faac-devel
faad2
faad2-devel
fcgi
ffmpeg
ffmpeg-devel
ffmpeg-libpostproc
filesystem
finch
finch-devel
findlib
firefox
firefox-devel
firefox-jre-plugin
flash-plugin
foxwrapper
franzlisp
ftools
ftools-debuginfo
funnie
fvwm
fvwm-debuginfo
fvwm-icons
gafstoken
gafstoken-debuginfo
gaim
galax
gcc34
gecode
getupdates
ggobi
glimpse
glimpse-debuginfo
gl…
…trove
…u.trove-javado…
gocr
gocr-debuginfo
gocr-devel
gride…
g…
…-debuginfo
…-conf
gsm-devel
gstreamer-plugins-ugly
gturing
gurgle
gurgle-caingres
gurgle-mysql
gurgle-postgresql
gwt
happy
happy-k…
haskell-e…cs-hooks
…xml
HaXml-d…
HaXml-g…4.2
helium
hexcalc
hibernate…
hibernate…vadoc
hibernate…manual
hide
horde
html2ps
html2ps-debuginfo
hugs98
hugs98-alut
hugs98-demos
hugs98-glut
hugs98-hgl
hugs98-openal
hugs98-opengl
hugs98-x11
i18n-lib
i18n-lib-javadoc
i8kutils
i8kutils-debuginfo
icu4j
icu4j-javadoc
idzebra
ifhp
ifhp-debuginfo
imapsync
imp-h3
ingres2006
ingres2006-abf
ingres2006-dbms
ingres2006-documentation
ingres2006-ice
ingres2006-jdbc
ingres2006-license-gpl
ingres2006-net
ingres2006-odbc
ingres2006-rep
ingres2006-star
ingres-tools
intel_fc-lice…
intel-icc8
intel-iidb7
ip…
ipw2200…
ipw2200-fw-…uginfo
isabelle
isabelle-ccl
isabelle-c…
isabelle-ctt-src
isabelle-debuginfo
isabelle-fol
isabelle-folp
isabelle-folp-src
isabelle-fol-src
isabelle-holcf
isabelle-holcf-src
isabelle-lcf
isabelle-lcf-src
isabelle-pdf
isabell…
isab…sequents-s…
isa…-src
isa…
isa…e-zf-src
isor…
isore…vadoc
iv
jabberd
jabberd-debuginfo
JabRef
jape
java-1.5.0-javadoc
java-1.5.0-sun
java-1.5.0-sun-alsa
java-1.5.0-sun-demo
java-1.5.0-sun-devel
java-1.5.0-sun-fonts
java-1.5.0-sun-jdbc
java-1.5.0-sun-manual
java-1.5.0-sun-plugin
java-1.5.0-sun-src
java-1.6.0-sun
java-1.6.0-sun-alsa
java-1.6.0-sun-demo
java-1.6.0-sun-devel
java-1.6.0-sun-fonts
java-1.6.0-sun-jdbc
java-1.6.0-sun-plugin
java-1.6.0-sun-src
javabayes
javabayes-javadoc
javassist
javassist-demo
javassist-javadoc
javassist-manual
jaxen-bootstrap
jconv
jconv-debuginfo
jess
jess32
jflap
jflap-javadoc
jik…
jisp2…mo
jisp2…vadoc
jmf
jmf-…doc
jmf-…nual
jrex…vadoc
kad…evel
kcap…al
kdcregister
kdcregister-debuginfo
kernel
kernel-debuginfo
kernel-devel
kernel-doc
kernel-pxe
kernel-pxe-install-fc3
kernel-pxe-install-fc5
kernel-pxe-install-fc5-
debuginfo
…nel-pxe-ins…6_64
…nel-pxe-ins…l5
…nel-pxe-ins…l5_64
…nel-smp
…nel-smp-de…
…eperacomm…
…pkcs11
kpkcs11-debuginfo
kqemumod
krb5-auth-dialog
krb5-debuginfo
krb5-devel
krb5-libs
krb5-server
krb5-workstation
kros
kstart
kstart-debuginfo
kx509
kx509-debuginfo
kx509-server
lame
lame-devel
latexdiff
latex.local.informatics
latex.local.informatics.defs
lcfg-afs
lcfg-afs-defaults-s1
lcfg-alias
lcfg-alias-defaults-s1
lcfg-amd
lcfg-amd-defaults-s2
lcfg-apache
lcfg-apacheconf
lcfg-apacheconf-defaults-s1
lcfg-apacheconf-nagios
lcfg-apache-defaults-s1
lcfg-arpwatch
lcfg-arpwatch-defaults-s1
lcfg-auth
lcfg-auth-defaults-s1
lcfg-authorize
lcfg-authorize-defaults-s1
lcfg-axnet
lcfg-axnet-defaults-s1
…boot-defaults
…g-bugzilla-defau…s1
…g-buildhosts
…g-buildhosts-de…s-s1
…g-buildinstallro…
…g-buildtools
lcfg-client-defaults-s3
lcfg-condor
lcfg-condor-defaults-s3
lcfg-conserver
lcfg-conserver-defaults-s2
lcfg-conserver-defaults-s3
lcfg-consoles
lcfg-consoles-defaults-s1
lcfg-cosig…
lcfg-cosig…faults-s1
lcfg-cosig…faults-s2
lcfg-cosig…gios
…g-cron
…g-cron-d…ults-s2
…g-cups
…g-cups-d…ults-s1
…g-cvs
…g-cvs-d…ts-s1
…g-cvsw…
lcfg-cvsweb-defaults-s1
lcfg-cyrussasl
lcfg-cyrussasl-defaults-s1
lcfg-dbiproxy
lcfg-dbiproxy-defaults-s1
lcfg-defetc
lcfg-defetc-fc5
lcfg-dhclient
lcfg-dhclient-defaults-s1
lcfg-dhclient-defaults-s2
lcfg-dhcpd
lcfg-dhcpd-defaults-s2
lcfg-dialup
lcfg-divine
lcfg-divine-debuginfo
lcfg-divine-defaults-s1
lcfg-dns
lcfg-dns-debuginfo
lcfg-dns-defaults-s2
lcfg-errorlog
lcfg-errorlog-defaults-s1
lcfg-etcservices
lcfg-etcservices-defaults-s1
lcfg-eventd
lcfg-eventd-defaults-s1
lcfg-events-defaults-s1
lcfg-example
lcfg-example-defaults-s1
lcfg-ffox-defaults-s1
lcfg-file
lcfg-file-defaults-s1
lcfg-friend
lcfg-friend-defaults-s1
lcfg-fstab
…tab-defaults-s2
…g-ganglia-defaults-s1
…g-gbios
…g-gbios-defaults-s1
…g-gdm
lcfg-gdm-defaults-s1
lcfg-gdm-defaults-s2
lcfg-gridengine
lcfg-gridengine-defaults-s1
lcfg-grub
lcfg-grub-defaults-s2
lcfg-grub-defaults-s3
lcfg-hackparts
lcfg-hackparts-debuginfo
lcfg-hardware
lcfg-hardware-defaults-s2
lcfg-hlfs-defaults-s1
lcfg-ingres
lcfg-ingres-defaults-s1
lcfg-init
lcfg-init-defaults-s1
lcfg-install
lcfg-installcfg
lcfg-install-defaults-s1
lcfg-installfixups
lcfg-installroot
lcfg-inventory
lcfg-inventory-client
lcfg-inventory-defaults-s1
lcfg-ipfilter
lcfg-ipfilter-defaults-s1
lcfg-iptables
lcfg-iptables-debuginfo
lcfg-iptables-defaults-s1
lcfg-jabberd
lcfg-jabberd-defaults-s1
lcfg-jabberd-nagios
lcfg-jnotify
lcfg-jnotify-defaults-s1
lcfg-kerberos
lcfg-kerberos-defaults-s6
lcfg-kerberos-nagios
lcfg-kernel
lcfg-kernel-defaults-s1
lcfg-kx509
lcfg-kx509-defaults-s1
lcfg-laptopupdate
lcfg-lcfginit
lcfg-lcfgweb
lcfg-lcfgweb-defaults-s1
lcfg-ldap-defaults-s3
lcfg-localhome
lcfg-localhome-defaults-s3
lcfg-localhome-doc
lcfg-logserver
lcfg-logserver-defaults-s1
lcfg-LPRng
lcfg-LPRng-defaults-s1
lcfg-macboot-defaults-s1
lcfg-mailcap
lcfg-mailcap-defaults-s2
lcfg-mailng
lcfg-mailng-defaults-s2
lcfg-matlab
lcfg-matlab-defaults-s1
lcfg-minimal-client
lcfg-minimal-server
lcfg-mkinitramfs
lcfg-monitor
lcfg-monitor-none-compdefs-
s1
lcfg-mozilla
lcfg-mozilla-defaults-s1
lcfg-mysql
lcfg-mysql-defaults-s1
lcfg-nagios
lcfg-nagios-compdefs-s1
lcfg-nagios-defaults-s1
lcfg-network
lcfg-network-defaults-s2
lcfg-nfs
lcfg-nfs-defaults-s2
lcfg-ngeneric
lcfg-ngeneric-defaults-s1
lcfg-nscd
lcfg-nscd-defaults-s3
lcfg-nsswitch
lcfg-nsswitch-defaults-s2
lcfg-nsu
lcfg-nsu-debuginfo
lcfg-nsu-defaults-s1
lcfg-ntp
lcfg-ntp-debuginfo
lcfg-ntp-defaults-s2
lcfg-nut
lcfg-nut-defaults-s2
lcfg-om
lcfg-om-defaults-s1
lcfg-openldap
lcfg-openldap-defaults-s3
lcfg-openldap-nagios
lcfg-openssh
lcfg-openssh-defaults-s1
lcfg-openssh-nagios
lcfg-openvpn
lcfg-openvpn-defaults-s2
lcfg-pam
lcfg-pam-defaults-s2.0
lcfg-pcmcia
lcfg-pcmcia-defaults-s1
lcfg-perlex
lcfg-perlex-defaults-s1
lcfg-pkgtools
lcfg-pkgtools-debuginfo
lcfg-pkgtools-devel
lcfg-pkgtools-perl
lcfg-pkgtools-perl-debuginfo
lcfg-podim
lcfg-podim-defaults-s1
lcfg-postgresql
lcfg-postgresql-defaults-s1
lcfg-postgresql-defaults-s2
lcfg-powernowd
lcfg-powernowd-debuginfo
lcfg-powernowd-defaults-s1
lcfg-prelink
lcfg-pxelinux
lcfg-pxelinux-defaults-s1
lcfg-ramdisk
lcfg-ramdisk-defaults-s1
lcfg-reapusermounts
lcfg-reboot
lcfg-release-fc5
lcfg-release-fc5_64
lcfg-remctl
lcfg-remctl-defaults-s1
lcfg-rfe
lcfg-rfe-defaults-s3
lcfg-rmirror
lcfg-rmirror-defaults-s2
lcfg-routing
lcfg-routing-debuginfo
lcfg-routing-defaults-s2
lcfg-rpmaccel
lcfg-rpmaccel-defaults-s1
lcfg-rpmaccel-nagios
lcfg-rpmcache
lcfg-rpmcache-defaults-s2
lcfg-rsync
lcfg-rsync-defaults-s3
lcfg-rsync-defaults-s4
lcfg-rsync-nagios
lcfg-samba
lcfg-samba-defaults-s3
lcfg-sane-defaults-s1
lcfg-saslauthd
lcfg-saslauthd-defaults-s1
lcfg-schemes
lcfg-server
lcfg-server-apache-conf
lcfg-server-defaults-s3
lcfg-shezhu
lcfg-shezhu-defaults-s4
lcfg-skeleton
lcfg-smartd
lcfg-smartd-defaults-s1
lcfg-snmp
lcfg-snmp-defaults-s2
lcfg-snmptrap
lcfg-snmptrap-defaults-s1
lcfg-splus-defaults-s1
lcfg-stubweb
lcfg-stubweb-defaults-s1
lcfg-subversion
lcfg-subversion-defaults-s1

1234 of them!

Middleware

- Middleware is last year's buzz phrase
- Also called 'soft-infrastructure' and many other things by the marketing folks
- Layer that sits between the network and the services you run
- Typically includes
  - Authentication Services
  - Authorisation Services
  - Directory Services
  - Machine configuration systems
  - File service
- Most of this talk is going to be about these 5 areas

Buzzword bingo

- One of the perils of middleware is that everyone wants to sell you something
- Buzz words prevail
- No two vendors use the same word to mean the same thing!

Middleware	Identity Theft	Two factor	Convergence	n-Tier
Assurance	Location Idenpendent	Entitlement	Policy	Schema
Distributed	Meta-directory	FREE SQUARE	Federation	SOA
Capability	Single Signon	Role based access control	Platform Independent	Identity Management
Directory Enabled Networks	Biometrics	TCO	Trust	Provisioning

Authentication

- Authentication is the act of proving your identity to the system.
- It’s not about determining what that identity can do.

Authentication systems

- Trivial systems store username/password lists for every service
  - Password syncronisation is hard
  - User has to remember many different passwords
  - User has to repeatedly enter passwords on many different occasions
  - Lots of machines have password lists lying around
  - Compromise of a machine allows the compromise of every other machine that uses the same passwords.

- Single signon systems solve these problems

Single signon

- A single password for all services (within the same trust domain)
- Password only required once at the start of a session
- Servers never obtain the user’s password
- The compromise of one service must not compromise the whole system

Kerberos

- Informatics use Kerberos for all internal authentication
- Kerberos is a cryptographically secure protocol linking clients and servers through a trusted third party - the Key Distribution Centre.
- Designed for use between trusted systems, across an untrusted network - the client’s password never crosses the network.
- Clients receive ‘tickets’ from the KDC which allow them to prove their identity to services, without allowing that service to impersonate them elsewhere
- Clients may ‘delegate’ their tickets to a service if they wish to permit impersonation (for example, when accessing a login service)

Kerberos services

- Whilst Kerberos can simply be used as a central password source, that doesn’t provide the full power

- To get the full benefit from Kerberos, you need support in your client and server applications

- Applications with support include:
  - **AdiumX**, amanda, Apache httpd, AFP, Apple CalendarServer, alpine, Balsa, Coda, CUPS, curl, cvs, Cyrus, DB2, Dovecot, evolution, Emacs, Eudora, Evolution, Fedora-DS, fetch, **firefox**, GridEngine, iCal, iChat, Internet Explorer, **jabberd2**, kmail, konqueror, Kermit, LPRng, Mail.app, Mutt, Mulberry, NFSv4, Oracle, OpenAFS, OpenLDAP, **OpenSSH**, qpopper, pine, **pidgin**, postfix, postgreSQL, Safari, Samba, sendmail, SQL Server, **SunSSH**, svn, **Thunderbird**, UW-IMAPD, Wildfire

Kerberos @ Informatics

- Use MIT Kerberos for servers - clients use whatever they ship with.
- Multiple deployed KDCs for redundancy, all running on Linux.
- DNS records to support zero configuration
- Locally written services deployed to automate key management as part of our machine configuration system
- Can authenticate to our machines from any compatible machine, anywhere on the Internet!

Kerberos demo

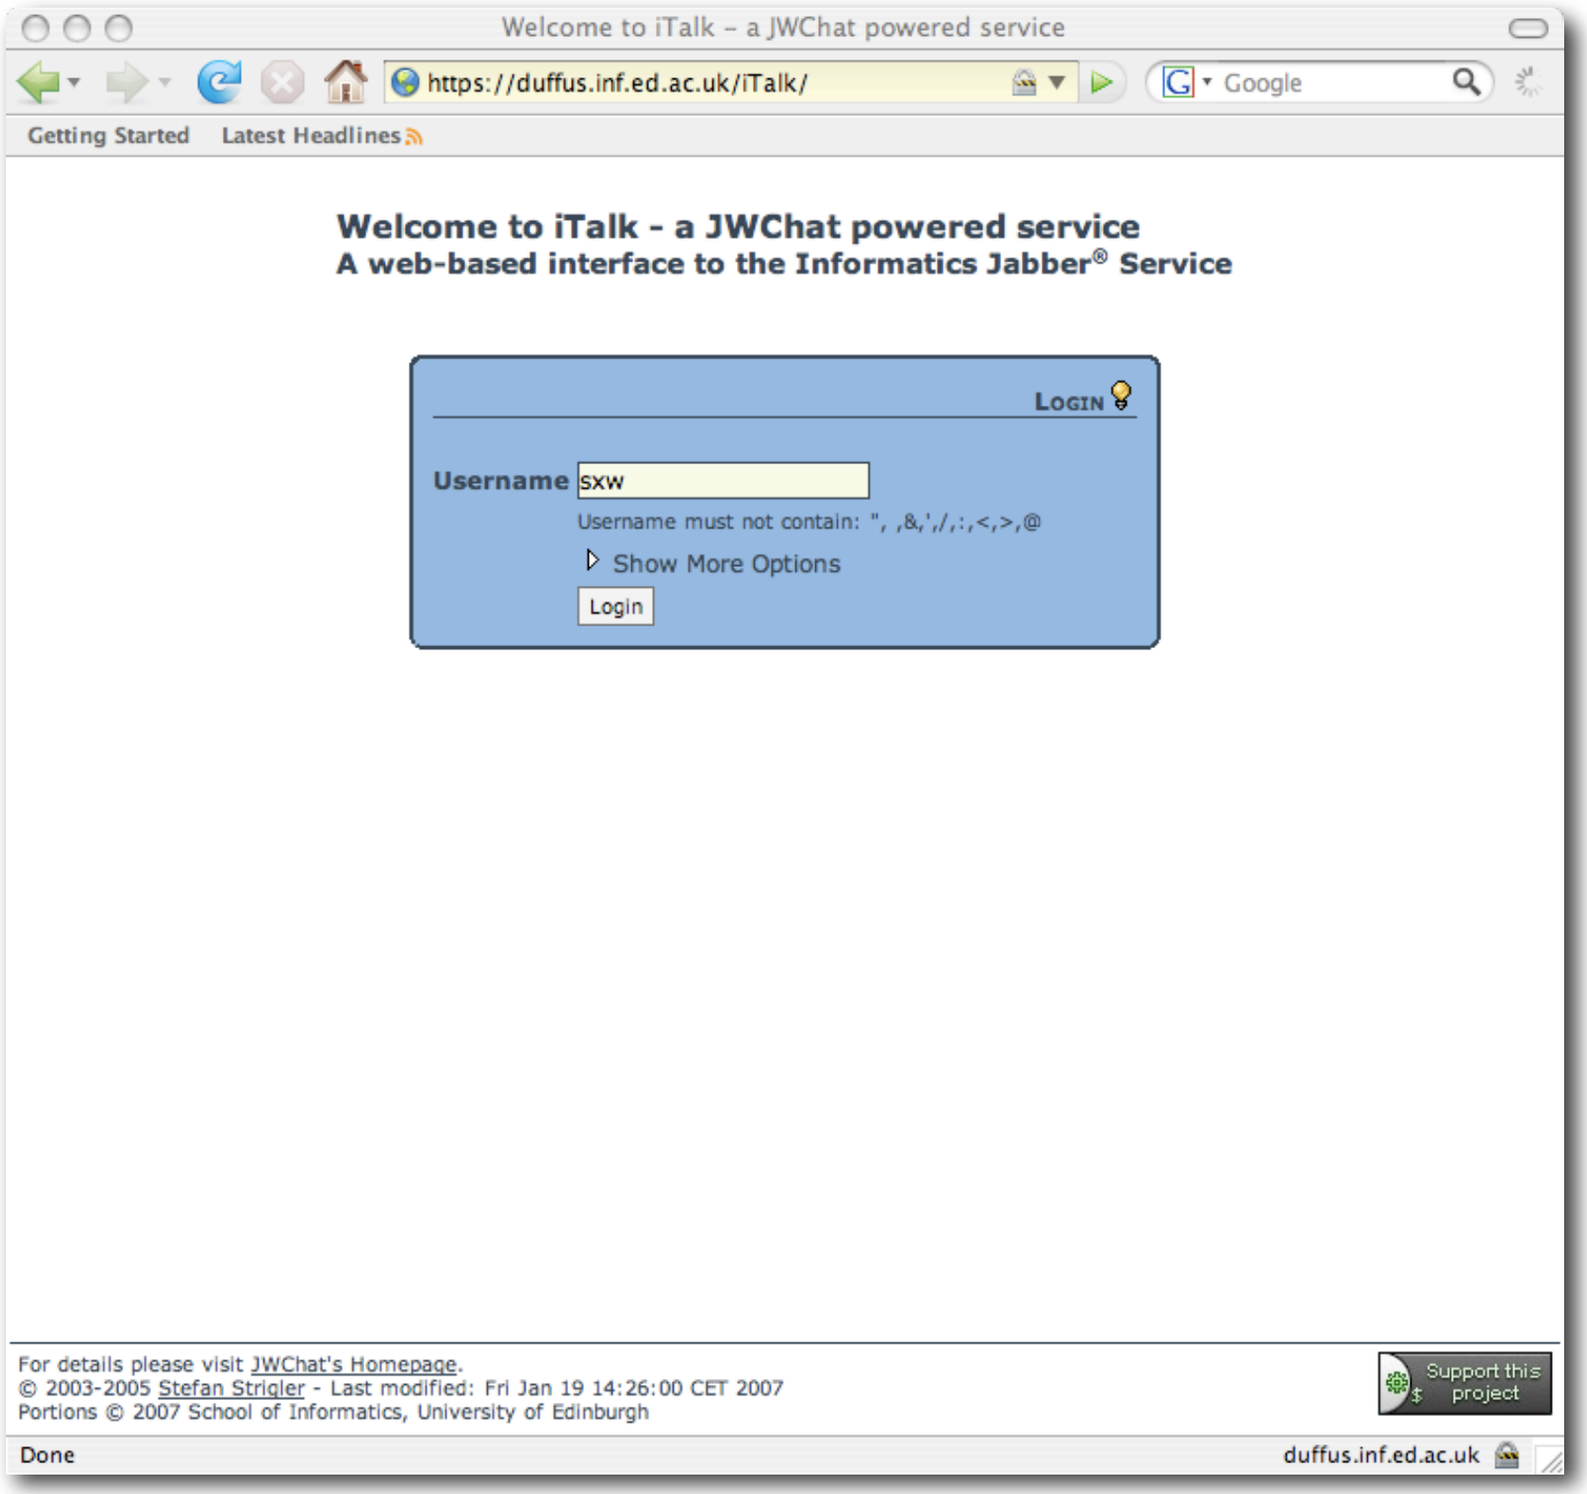

Kerberos key management

- We still use other key management schemes
  - X509 (for TLS/SSL protected services)
  - SSH public keys

- Have developed software to bootstrap all of these keys from a machine's Kerberos key.

- Only need to install one piece of key material on a given machine

Kerberos & the WWW

- Traditionally Kerberos support within web browsers has been poor
- Need to solve the internet cafe case - machine with no Kerberos support but still wants to access our services.
- We use **Cosign** to provide web based authentication that ties in with our internal authentication system
- A web user without Kerberos credentials is prompted for their username and password
- Cookie based authentication then allows the user to access local web sites from anywhere on the internet.

Kerberos & guest users

- We have many virtual visitors - people who need more than public access to our services, but have no physical affiliation to the School.

- Created ‘iFriend’, which allows anyone with an email address to create an account on our KDCs.

- Remember - authentication has nothing to do with access control!

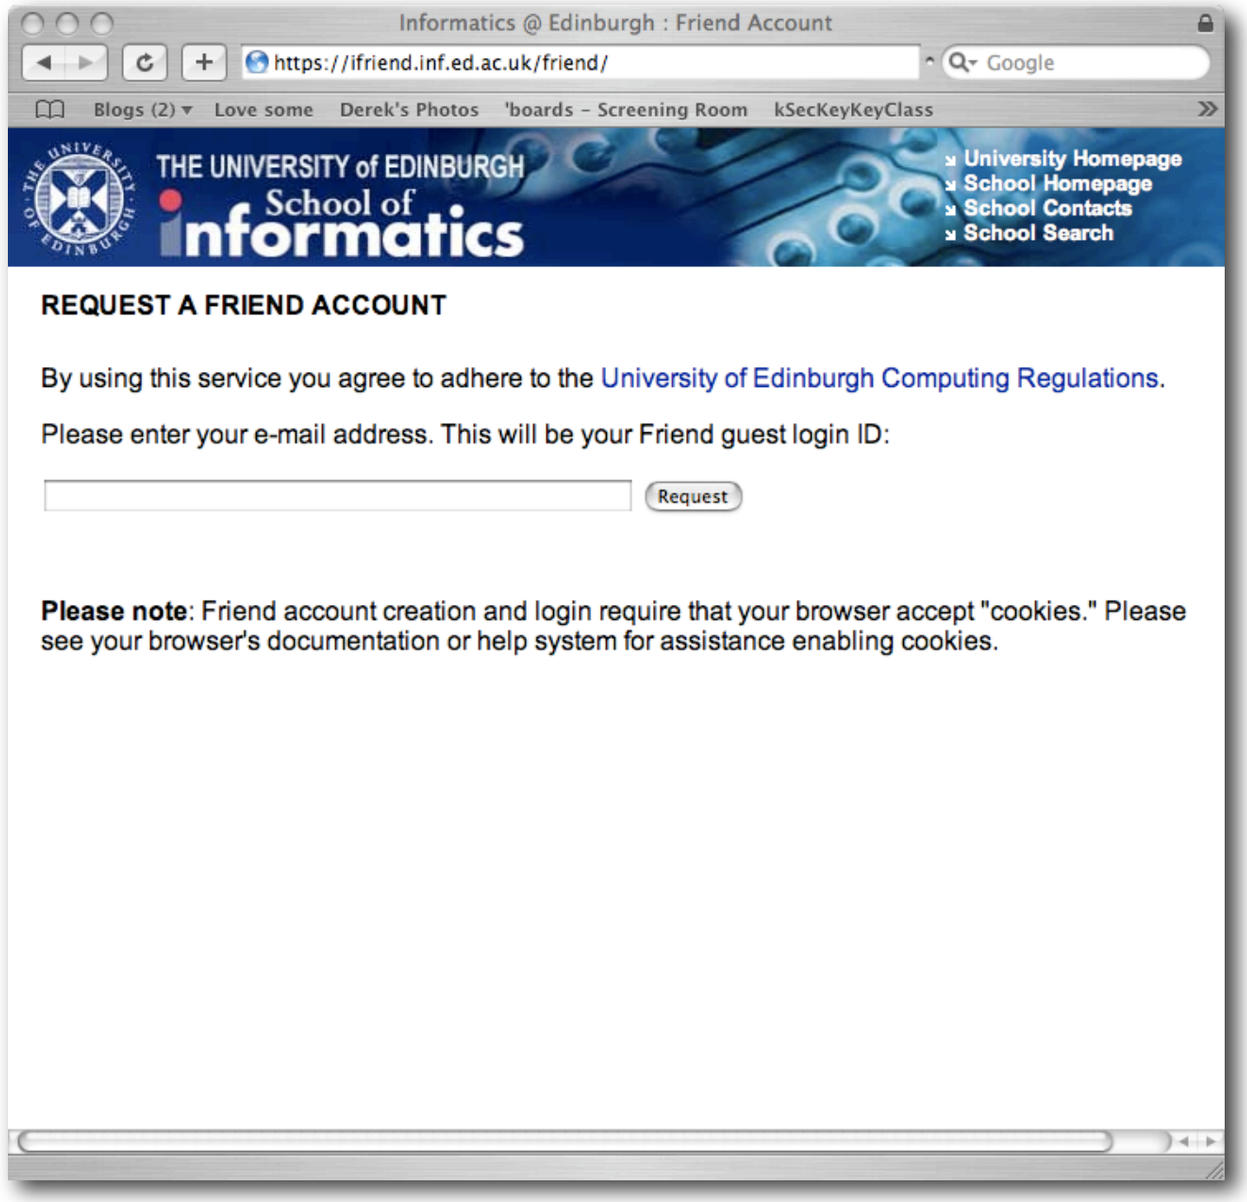

Authentication futures

- Passwords are becoming increasingly fragile
- Smartcards look like being the way to go
    - US Federal Government leading the way with PIV
- Smartcard support in Kerberos is under active development

Authorisation Services

- If authentication establishes “**Who you are**”, then authorisation controls “**What you can do**”

- General purpose authorisation services are hard
  - User X can print to printer Y
  - User X can print to printer Y, but only between 9am-5pm
  - User X can print to printer Y, but only documents less than 100 sheets
  - User X can print to printer Y, but only if they’re within their print quota

- Need to decide what kind be solved generally, and what is best left to the individual application

- No good, general purpose, open source solutions available in this space!

Authorisation @ Informatics

- System built around our directory service
- Each user has a set of ‘roles’. Some of these are determined from personnel data (‘staff’, ‘student’) - others are added by support staff (‘beowulfuser’)
- Each role gives a user a set of entitlements. Entitlements are service specific tags (‘print/laserprinter2’, ‘login/duffus’) which are checked to perform access control decisions
- Services then check whether a particular user has a particular entitlement in order to determine access
- Anything more complex is up to the individual service

Authorisation Implementation

- All source information stored in our directory service
  - Role => entitlement mapping information stored in custom maps
  - User role membership stored in their account record, and managed by account provisioning system

- Each entitlement has a LDAP group containing a list of users

- Entitlements also represented in netgroup format for use by legacy clients

- Trigger script updates entitlement groups and netgroups whenever source information is changed

- Code freely available (GPL) upon request

Directory Services

- Saying “we have a directory service” is a lot like saying “we have a database”!

- Informatics use our Directory Service for a number of distinct purposes
  - White pages
  - Yellow pages (NIS) information for Unix machines (passwd, group &c.)
  - Host information
  - Authorisation information
  - Application specific data
    - printcaps
    - KDC principal information
    - automounter maps

Directory Service Implementation

- OpenLDAP based service
- Single master, with public data replicated to every client (using locally written replication technology)
- Massive redundancy, but wouldn’t recommend this configuration!
- Currently looking to completely redesign our replication strategy
- All directory access is Kerberos authenticated, and controlled via authorisation entitlements

Directory Service Layout

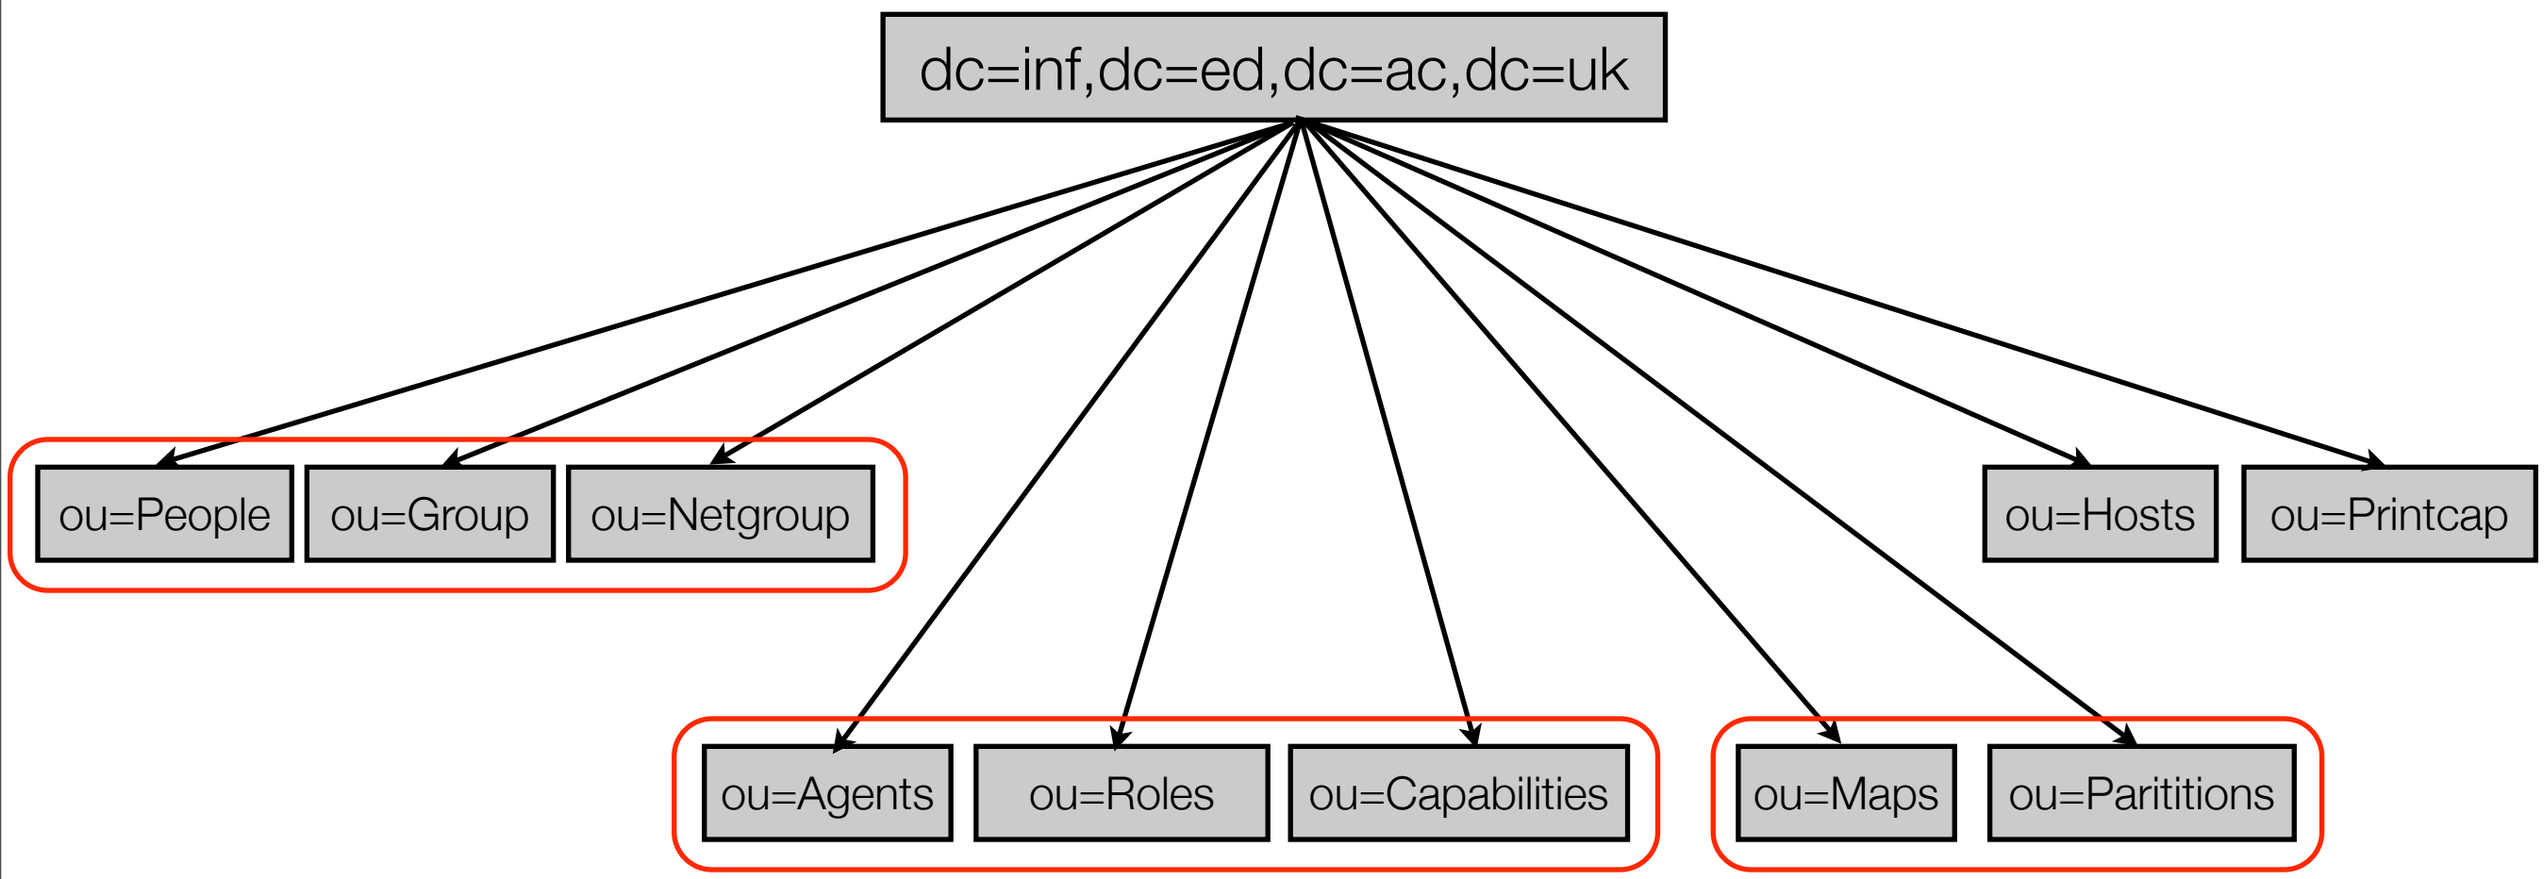

Directory Service Demo

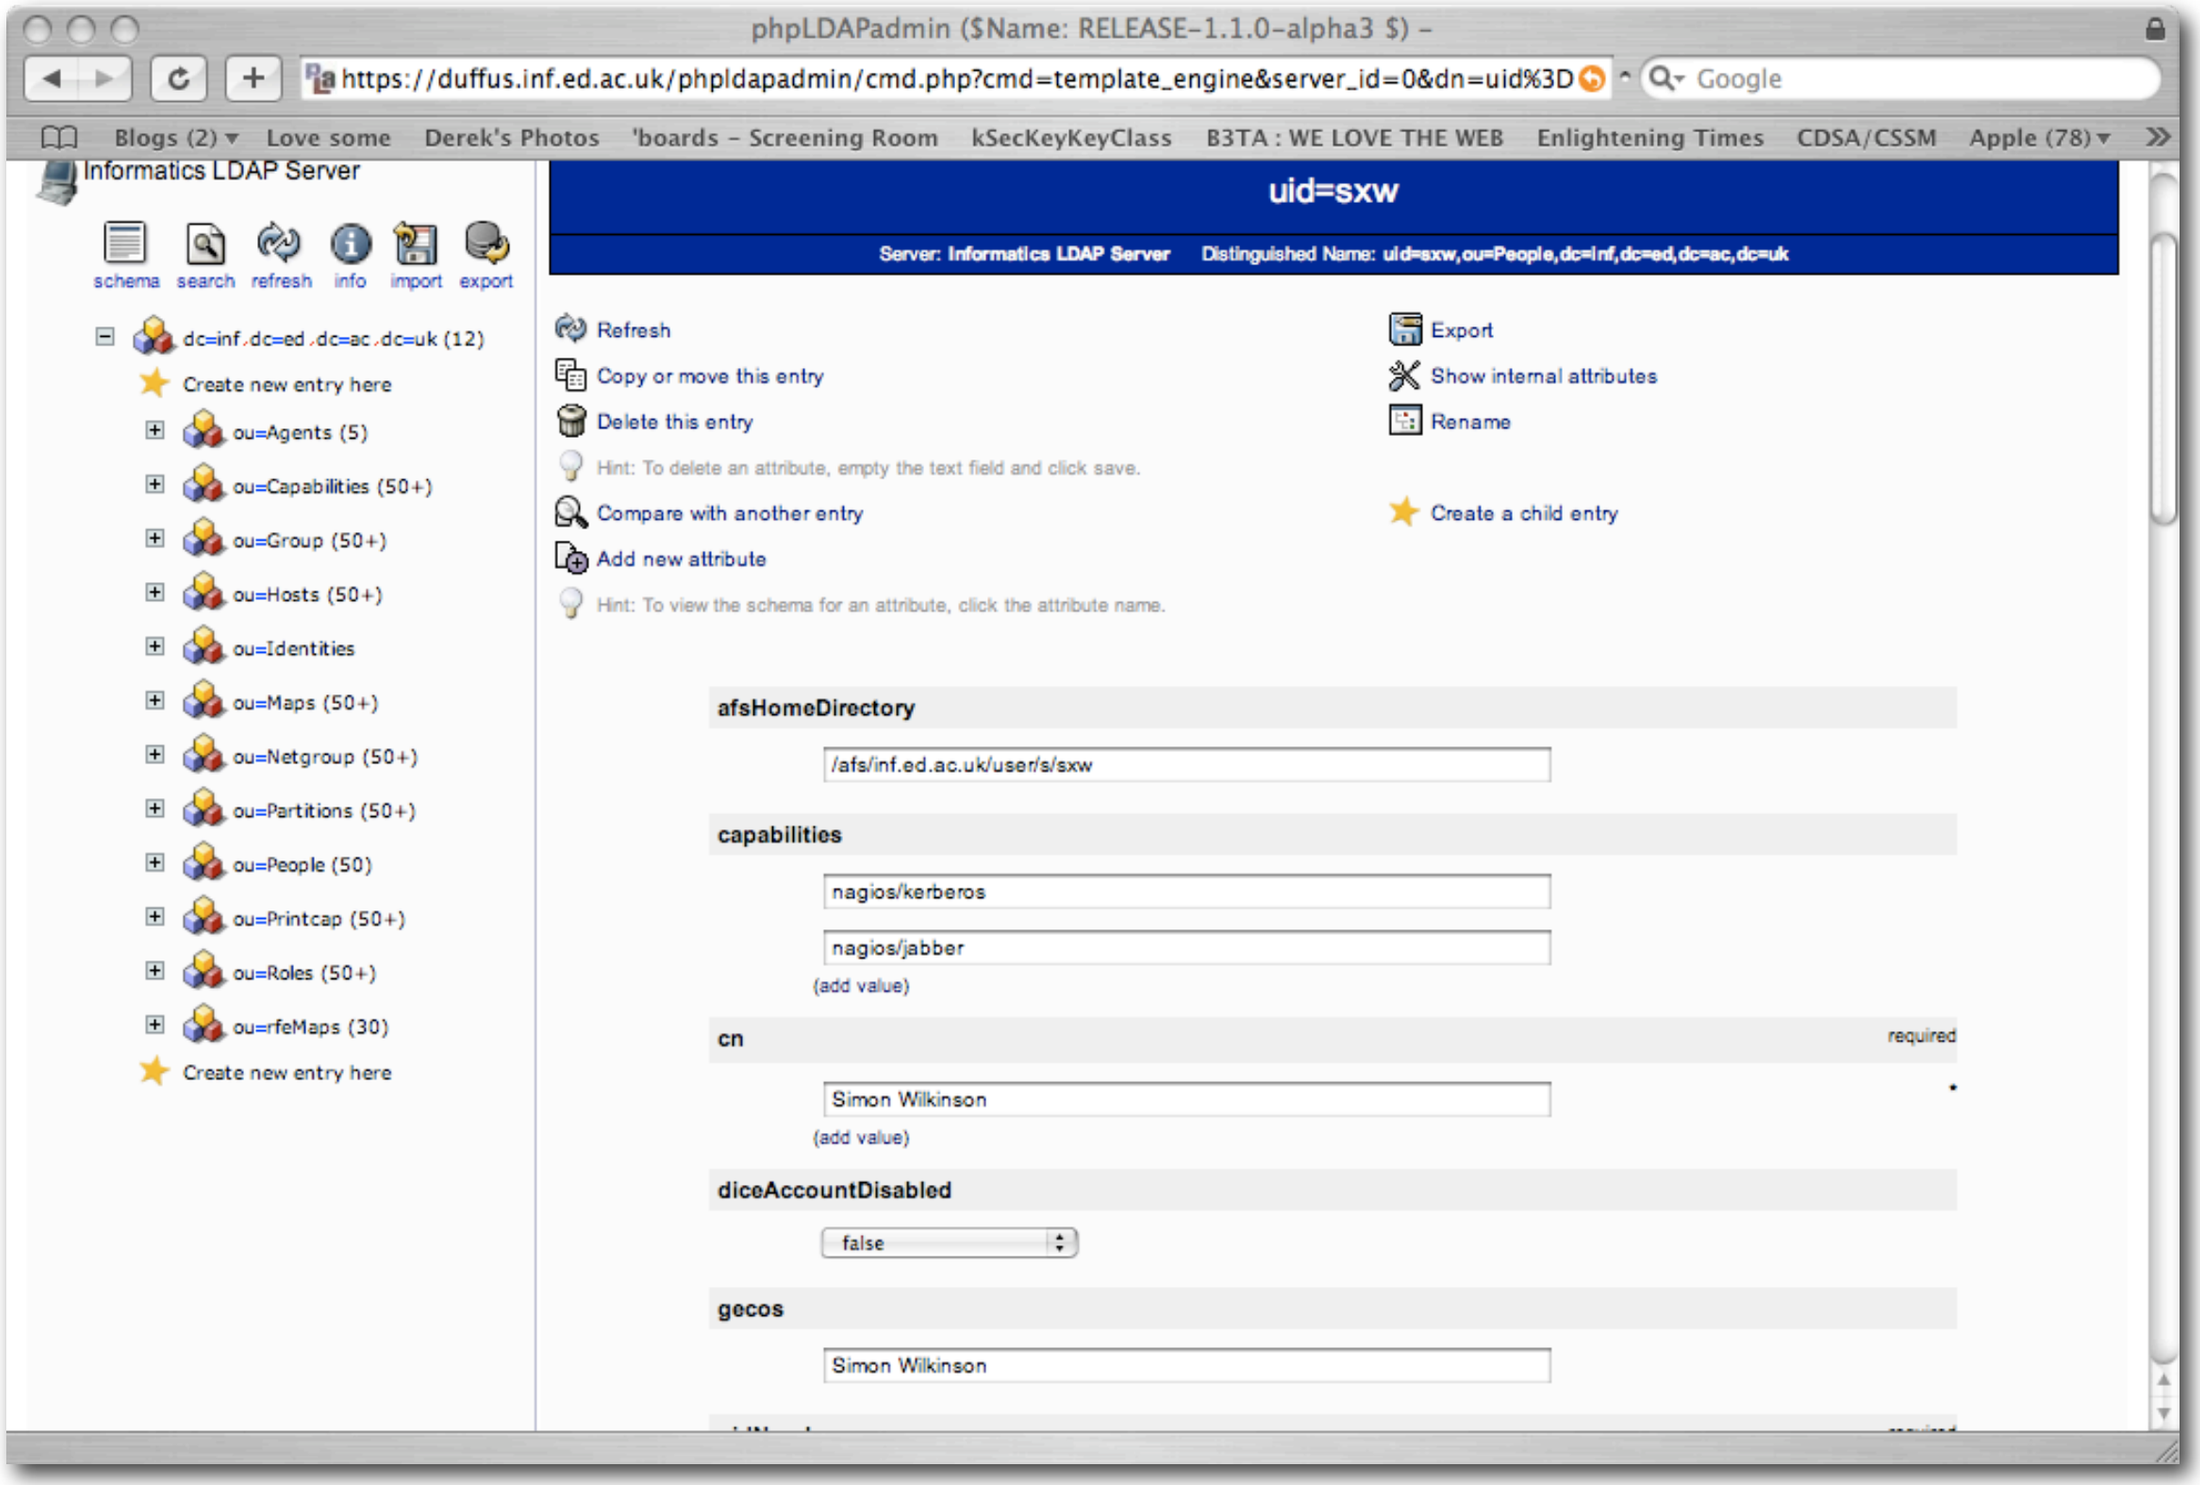

Identity

- Putting the last three topics together, you nearly have the makings of a complete Identity Management solution
- Haven’t really discussed account provisioning, as it is so site specific
- Our account management toolset pulls information from multiple local, and University-wide, data sources to create appropriate entries in LDAP and our KDC.

The view from here ...

- All of these services are designed to be usable from any client, and by any service.
- Clients can authenticate by just doing `kinit user@INF.ED.AC.UK`
- Any standards-compliant application can use our LDAP white pages system
- Any Unix client which supports LDAP access from the name service switch can use our LDAP yellow pages data (using nss_ldap, or similar)
- Any Unix service which supports PAM can accept our Kerberos passwords (using pam_krb5)
- Any Unix service which supports LDAP group authentication can use our authorisation system

Configuration Management Options

- System imaging
- Just hack the machines till they work!
- Store config files elsewhere and use a deployment tool.
- Generate templates for config files that can be filled in per-host and deployed.
- Put templates on the host and deliver some configuration data to build config files.

Requirements

- Reproducibility
- Scalability
- Ease/Efficiency of Management
- Validity
- Manage relationships
- Manage change

An example – Create a new web server

- **Infrastructure**
  - DNS entry
  - DHCP entry
  - Firewall hole
  - SSL certificates
  - Backups

- **Configure the machine:**
  - Configure disks, install software, etc..
  - Configure dns, network, ntp, apache, etc..

LCFG

- Each managed host has a source profile
- Source profile pulls in headers
  - allows sharing of config data between machines
- Central server compiles this into XML.
- Client downloads XML profile and configures itself accordingly

The Profile

- Completely describes the required state of a machine
- Consists of a set of components and, optionally, a list of packages.

A Component

- Set of resources – basically key/value pairs.
- Optional templates for config files.
- Optional control code
  - Manage daemons (stop/start).
  - LCFG provides a standard framework for Perl and shell scripts.

Spanning Maps

- A component in one profile can ***publish*** resources to which a component in the profile of another machine ***subscribes***.

- Usage includes:
  - dhcp
  - ipfilter
  - nagios, system monitoring
  - inventory

Configuration Conclusions

- LCFG
  - Improves usability (common language, etc.).
  - Allows a complete description of the state.
  - Allows devolved management.
  - Makes relationship management easy.
  - Provides the ability to manage change.

File services

- Providing a reproducible environment requires every machine has the same view of the filesystem
- In particular user's homedirectories, but also research data
- Historically, used to use NFS
- Security was not good
- Manageability was even worse

AFS

- In the process of moving to OpenAFS
- AFS is a truly global filesystem - any AFS client, anywhere in the world can access files on any AFS server
- Data can be moved around between fileservers without the user noticing
- Security provided through Kerberos
- Many staff, and all new students now have AFS homedirectories. Roll out to continue over the next 4 years.

Putting it all together

- A quick demo of ‘iFile’ - web based access to AFS
- This shows off nearly everything we’ve talked about ...
  - We’re authenticating with cosign (which gets Kerberos tickets for us)
  - The application checks us against the authorisation service
  - The web interface uses our Kerberos ticket to authenticate us to AFS
  - And the whole thing is configured and managed through LCFG

iFile Demo

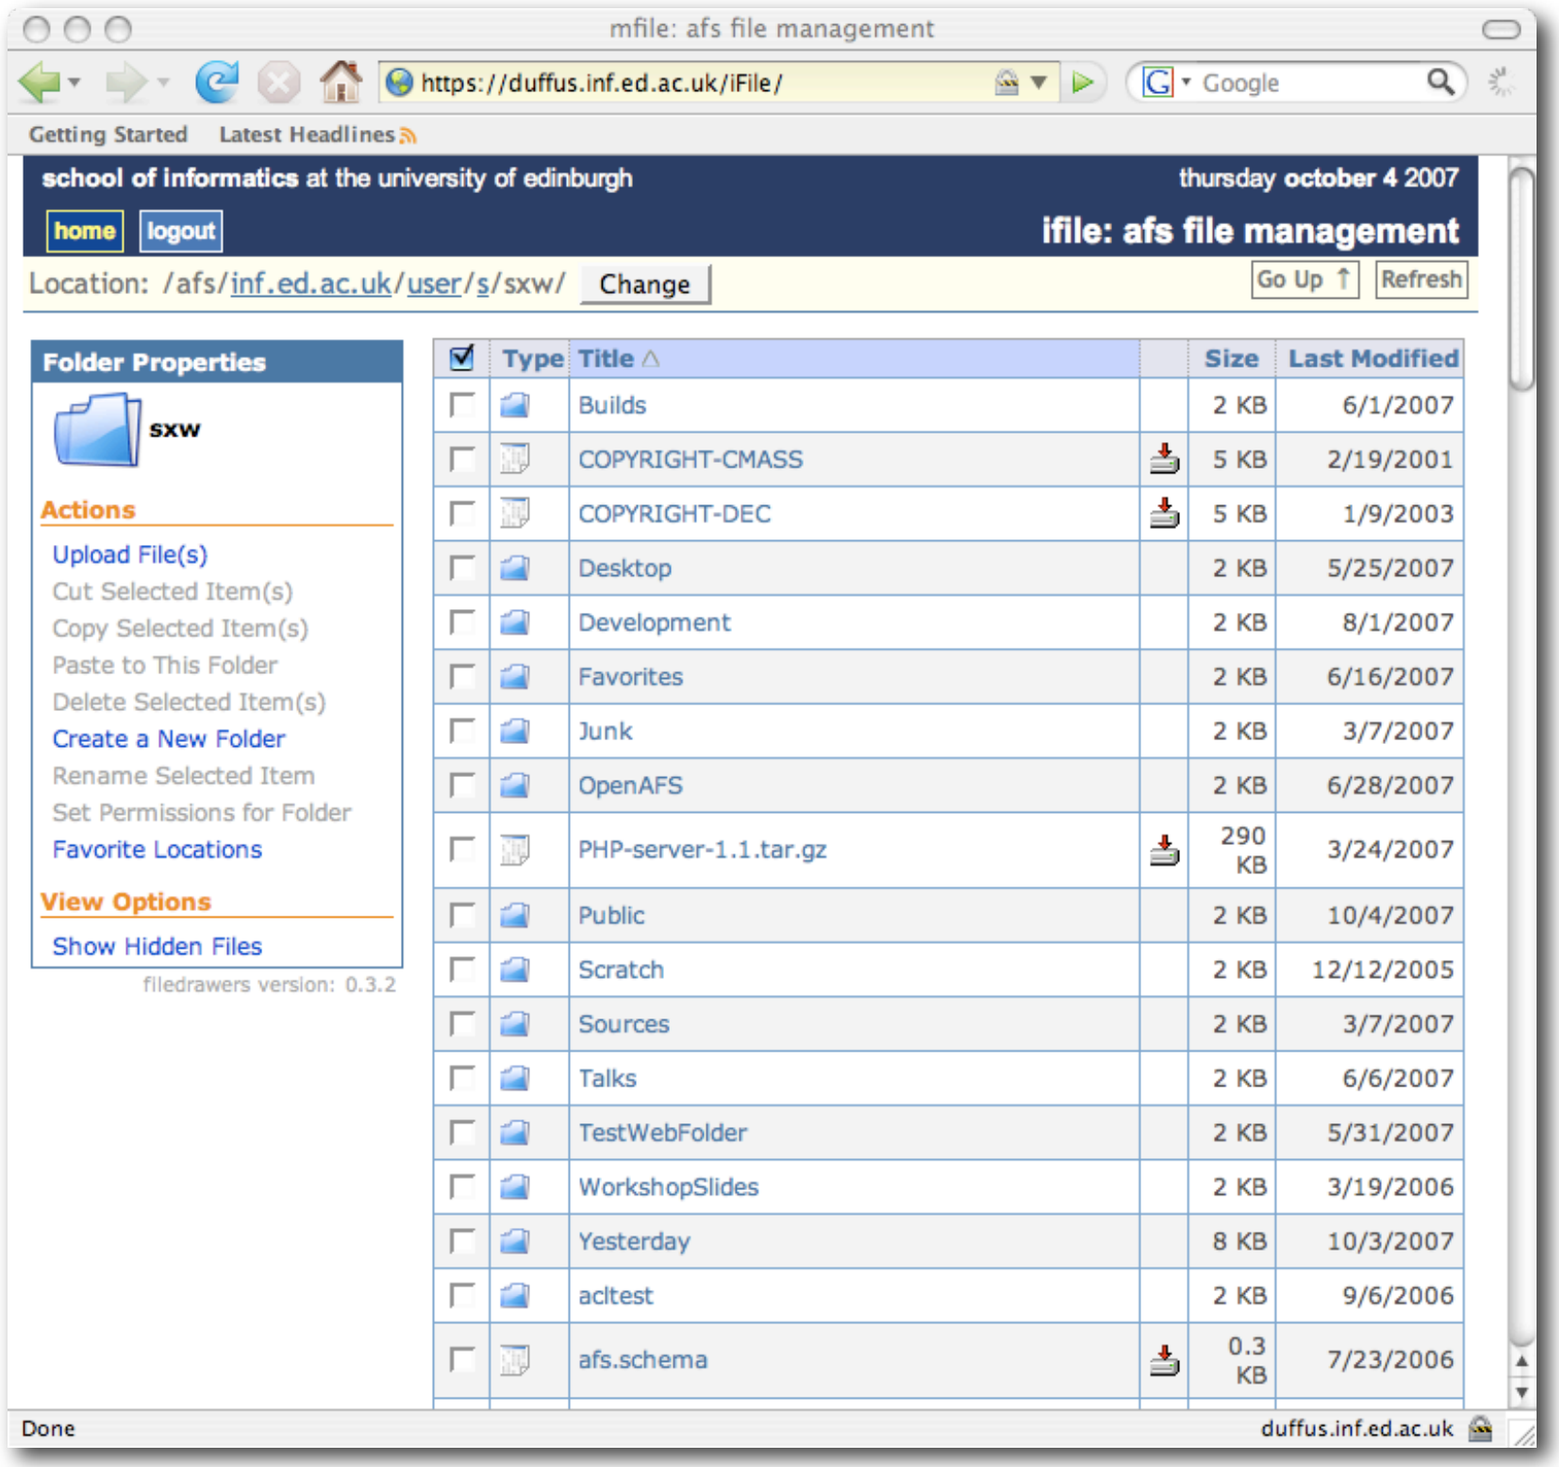
mfile: afs file management
https://duffus.inf.ed.ac.uk/iFile/
Google
Getting Started
Latest Headlines
school of informatics at the university of edinburgh
thursday october 4 2007
home
logout
ifile: afs file management
Location: /afs/inf.ed.ac.uk/user/s/sxw/
Change
Go Up
Refresh
Folder Properties
sxw
Actions
Upload File(s)
Cut Selected Item(s)
Copy Selected Item(s)
Paste to This Folder
Delete Selected Item(s)
Create a New Folder
Rename Selected Item
Set Permissions for Folder
Favorite Locations
View Options
Show Hidden Files
filedrawers version: 0.3.2
Type
Title
Size
Last Modified
Builds 2 KB 6/1/2007
COPYRIGHT-CMASS 5 KB 2/19/2001
COPYRIGHT-DEC 5 KB 1/9/2003
Desktop 2 KB 5/25/2007
Development 2 KB 8/1/2007
Favorites 2 KB 6/16/2007
Junk 2 KB 3/7/2007
OpenAFS 2 KB 6/28/2007
PHP-server-1.1.tar.gz 290 KB 3/24/2007
Public 2 KB 10/4/2007
Scratch 2 KB 12/12/2005
Sources 2 KB 3/7/2007
Talks 2 KB 6/6/2007
TestWebFolder 2 KB 5/31/2007
WorkshopSlides 2 KB 3/19/2006
Yesterday 8 KB 10/3/2007
acltest 2 KB 9/6/2006
afs.schema 0.3 KB 7/23/2006
Done
duffus.inf.ed.ac.uk

Further Information

- Informatics Infrastructure
  - http://www.dice.inf.ed.ac.uk/publications/
    Papers and talks about our Kerberos, OpenAFS and LCFG deployments

- Kerberos
  - http://web.mit.edu/kerberos/ - the MIT Kerberos site
  - http://www.pdc.kth.se/heimdal/ - Heimdal Kerberos
  - http://www.kerberos.org/ - the Kerberos consortium
  - http://www.eyrie.org/~eagle/software/pam-krb5 - PAM krb5 module

- Directory Services
  - http://www.openldap.org/
  - http://www.padl.com/OSS/nss_ldap - LDAP Name Service integration

Further Information (II)

- LCFG
  - http://www.lcfg.org/
  - http://wiki.lcfg.org.uk/
  - mailto:lcfg-discuss@inf.ed.ac.uk

- OpenAFS
  - http://www.openafs.org
  - http://www.filedrawers.org

Questions?

Talk available online at http://www.sxw.org.uk/computing/talks/edlug1007.pdf