

A QUICK GUIDE TO AFS

**SIMON WILKINSON
SCHOOL OF INFORMATICS
UNIVERSITY OF EDINBURGH
SIMON@SXW.ORG.UK**

AFS

- ✿ Originally developed by Carnegie Mellon University as part of Project Andrew
- ✿ Commercialised by Transarc, later acquired by IBM
- ✿ Became open source (and free!) in 2000
- ✿ Strong development community since then

AFS IS ...

“ ...a secure, global, federated, location independent open source storage system providing data access from a broad range of devices scaling from handsets to super computers. ”

A SECURE FILESYSTEM

- ✻ For today's purposes a "Secure Filesystem" is
 - ✻ One that does not rely on network trust
 - ✻ One that insulates users of multi-user machines from each other
- ✻ Generally not NFS before version 4

AFS SECURITY

- ✻ AFS uses Kerberos for user authentication
- ✻ Data transfer is (optionally) encrypted
- ✻ File access is policed on a per user, per access basis
- ✻ Per directory access control lists allow user control of all permissions

A GLOBAL FILESYSTEM

✱ How big is your world?

✱ A room

✱ A city,

✱ A country

✱ A planet?

☼ A solar system?



http://workshop.openafs.org/afsbpw06/talks/keynote_djbyrne.pdf

AFS ACROSS THE GLOBE

- ✻ AFS is designed for WAN use
- ✻ Works well over low bitrate, high latency links
- ✻ Operates across NAT devices

A FEDERATED FILESYSTEM

- ✱ Collaborative working requires file access by many different users across multiple organisations
- ✱ Difficult to handle with many filesystems
 - ✱ How do you manage cross-institution users
 - ✱ How do you control cross-site access?

AFS FEDERATION

- ✻ AFS leverages Kerberos's cross-realm capabilities
- ✻ If a user has access to a realm you trust, they can be given access to your files
- ✻ Lightweight authentication solutions can be used to give truly global access

A LOCATION INDEPENDENT SYSTEM

- ✱ Users don't *need* to know where their data lives
- ✱ Users don't **want** to know when you're moving their data between file servers

ACCESS YOUR FILES ANYWHERE

- ✻ At work
- ✻ At home
- ✻ At a conference
- ✻ On the train

... FROM A WEB CAFE ...


iFile: afs file management

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

school of informatics at the university of edinburgh monday february 11 2008

[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: 0.4.0

<input checked="" type="checkbox"/>	Type	Title ^	Size	Last Modified
<input type="checkbox"/>	PDF	AFS Workshop.pdf	55.1 KB	6/23/2005
<input type="checkbox"/>	Folder	Builds	2 KB	1/27/2008
<input type="checkbox"/>	Folder	Built	2 KB	1/14/2008
<input type="checkbox"/>	Folder	Desktop	2 KB	5/25/2007
<input type="checkbox"/>	Folder	Development	2 KB	2/5/2008
<input type="checkbox"/>	Folder	Favorites	2 KB	6/16/2007
<input type="checkbox"/>	Folder	Junk	2 KB	3/7/2007
<input type="checkbox"/>	Folder	OpenAFS	2 KB	12/17/2007
<input type="checkbox"/>	File	PHP-server-1.1.tar.gz	290 KB	3/24/2007
<input type="checkbox"/>	Folder	Public	6 KB	2/7/2008
<input type="checkbox"/>	Folder	Talks	2 KB	1/21/2008
<input type="checkbox"/>	File	UMICH-authz-ldap.txt	4.4 KB	11/2/2005
<input type="checkbox"/>	Folder	WeeklyNotes	2 KB	1/24/2008
<input type="checkbox"/>	Folder	WorkshopSlides	2 KB	3/19/2006
<input type="checkbox"/>	Folder	Yesterday	8 KB	2/10/2008
<input type="checkbox"/>	File	a.out	5.4 KB	2/1/2008
<input type="checkbox"/>	Folder	actest	2 KB	1/31/2008
<input type="checkbox"/>	File	afs.schema	0.3 KB	7/23/2006
<input type="checkbox"/>	File	afs.snoop	15 KB	12/13/2007

... FROM A PDA...



AFS

- ✻ Units of file manipulation
 - ✻ file
 - ✻ directory
 - ✻ volume
 - ✻ partition
- ✻ The volume is the key unit for management

VOLUMES

- ✿ Basic volume is a read-write copy
- ✿ Multiple read-only replicas for redundancy and load sharing
- ✿ Single 'backup' volumes provide a snapshot for backup & recovery
- ✿ Volume replication is a manual process

BACKUP VOLUMES

- ✻ Backup volumes are maintained as deltas
- ✻ No protection against disk failure
- ✻ Provide means for users to access yesterday's data

MORE ON VOLUMES

- ✻ AFS allows volumes to be transparently migrated between servers
- ✻ Volumes are stitched together through mountpoints to produce the filesystem
- ✻ Filesystem is typically identical on **every** host running AFS

A GLOBAL FILESYSTEM

- ✻ Standard mountpoint on all clients - /afs
- ✻ Next level is 'cell' - your site - /afs/inf.ed.ac.uk
Derived from DNS or global config file
- ✻ Below this, is up to the individual site
- ✻ All AFS sites can access /afs/inf.ed.ac.uk
- ✻ Our clients can access all AFS cells

CELLS AND DATABASES

- ✻ A cell is the AFS organisational unit
- ✻ Each cell will have a number of database servers providing
 - ✻ Volume Location - which fileserver a given volume is on
 - ✻ Protection database - group membership and permissions for all users in a
- ✻ AFS has powerful multimaster replication for all databases - you want more than one!

FILESERVERS

- ✻ Each cell may contain any number of file servers
- ✻ File servers do not store data on disk in human readable form - all access must come through AFS client
- ✻ Possible to completely bypass the native filesystem and use the raw disk

CLIENTS

- ✻ Clients are comprised of a kernel module, plus a user space daemon - the cache manager
- ✻ Cache manager deals with fitting volumes together into the filesystem
- ✻ Also handles powerful local caching system

CALLBACKS

- ✻ AFS protects cache integrity using Callbacks
- ✻ When a client opens a file it registers a callback with the fileserver
- ✻ Any changes to that file will result in the fileserver notifying the client

AUTHENTICATION

- ✱ Originally AFS used Kerberos v4
- ✱ Can now use Kerberos v5 natively
- ✱ Only supports DES encryption
- ✱ On-the-wire encryption is even weaker
- ✱ Better encryption on the way...

ACLs AND GROUPS

- ✿ ACLs available to control access on a per directory basis
- ✿ ACLs can set permissions by user, system wide group, or by user defined group
- ✿ Permissions can be both positive and negative
- ✿ Special groups exist for 'any authenticated user' and 'any user'

PLATFORM SUPPORT

- ✻ AIX 5.1-5.3
- ✻ HPUX 11.0, 11i, 11i v2
- ✻ IRIX 6.5
- ✻ Linux 2.4 and 2.6 (through .24) kernels
 - ✻ Fedora Core 3-9, RHEL3-5
 - ✻ Debian, Ubuntu
 - ✻ Gentoo and others
- ✻ MacOS 10.3-10.5
- ✻ Microsoft Windows
 - ✻ 2000, XP, Server 2003,
 - ✻ Vista, Server 2008 (32-bit and 64-bit)
- ✻ NetBSD, FreeBSD and OpenBSD (server only)
- ✻ Solaris (Sparc and x86) 7-11 and OpenSolaris

FUTURES

- ✻ Disconnected Operation
- ✻ Better cryptography
- ✻ Object storage

USERS

Argonne National Laboratory, Carnegie Mellon University, CERN, DAPNIA, Duke University, Edinburgh University, ENEA, Fermi National Accelerator Laboratory, Iowa State University, Jet Propulsion Laboratory, KTH, Leiden University, Manchester Computing, MIT, Morgan Stanley, NASA, Naval Research Laboratory, New Jersey Institute of Technology, Intel, NCSA, Pictage, Pittsburgh Supercomputing Centre, Stanford Linear Accelerator Centre, Stanford University, Stockholm University, United Airways, University of Maryland, University of Michigan, University of North Carolina, University of Notre Dame, University of Pittsburgh, United States Geological Service, Rutherford Appleton Laboratory

180 public cells - and an increasing number of known private cells.

PICTAGE

- ✻ Wedding Photography Service
- ✻ Over 260 Terabytes of AFS storage (May 2007)
(planned growth to 420TB by now!)
- ✻ ... large number of Windows clients

UNITED STATES GEOLOGICAL SURVEY

- ✻ Provide “natural hazard information on the web ... even in the event of large-scale regional disasters”
- ✻ AFS replication used to mirror web data between multiple regional data centres

<http://www.openafs.org/success.html>

MORGAN STANLEY

- ✻ 25000+ hosts in 50+ sites on 6 continents
- ✻ “No AFS, No Unix”

UNITED AIRLINES

- ✿ All aircraft maintenance information stored in AFS
- ✿ Millions of pages of documentation
- ✿ Available at every gate, maintenance bay and vendor world wide

QUESTIONS?

This talk (and others on AFS) available from
<http://www.dice.inf.ed.ac.uk/publications/>

Feel free to email me
simon@sxw.org.uk