

OpenAFS Status 2012

Nothing and a lot

- Derrick Brashear and Jeffrey Altman
- The OpenAFS Project
- 16 October 2012

History

- OpenAFS 1.6.0 was released on 1 September 2011.
 - Just 3 years late...
- You probably recall discussion of the “rx ping issue” last year.
- Data corruption issues as well.
- So we spent a lot of time making it right
- 1.6.1 on 3 April 2012

OpenAFS 1.6.1

- rx ping avalanche
 - leak of AFS connections (with one Rx connection each) during server probes
 - issues with which Rx connection would exist
 - tracked per server so exactly one exists
- corruption due to CopyOnWrite speedup
 - reverted

OpenAFS 1.6.1

- Lion Kerberos fixes. (Ugh)
- MacOS shutdown panic fixes.
- Positional I/O bugfix to avoid losing data due to writing to an obsoleted filehandle.
- Ubik log replay fixes
- Proper bosserv rxbind support.

OpenAFS 1.6.1a

- MacOS client issues (new packets)
 - Special listener
- MacOS Kerberos issues
 - Missing APIs
- MacOS Preferences Pane issues
 - admin powers

1.6 on the Road Map

openafs.org/roadmap.html

1.7

- The 1.7 series is the current Windows release branch

1.6

- The 1.6 series is the current “Stable”
- The 1.6 series includes significant improvements to source code quality.
 - Real issues are not overwhelmed in bogus compilation warnings. Some real issues fixed too.
- Last release without a Windows IFS
 - Sites wishing to test Windows IFS can use 1.7 series.

1.6 Platform summary

- AIX 5 and 6 (through 6.3)
- FreeBSD 7, 8, 9 and current
- HP-UX 11i v1 and v2
- Irix 6.5
- Linux 2.4, 2.6, 3.x(*) (ia32, ia64, amd64, ppc, ppc64, arm, sparc, sparc64)
- MacOS 10.3, 10.4, 10.5, 10.6, 10.7, 10.8 (ppc, i386, amd64).
- NetBSD 2.x, 3.x, 4.x, 5.x
- OpenBSD 4.4, 4.5, 4.6, 4.7, 4.8.
- Solaris 2.6, 7, 8, 9, 10, 11 (and OpenSolaris)
- Windows XP, 2003, Vista, 7

1.6 new features

- Demand-Attach Fileserver
- Disconnected Operation
- Linux improvements and page cache performance
- MacOS 10 64 bit client and GUI client prefs
- Rx performance, correctness
- FUSE cache manager
- Cache bypass
- Cache readahead
- Mountpoint-less volume access.

Development divergence

- Master
 - Where everything starts
- 1.7
 - Changes for Windows IFS builds
 - Still fast moving enough to not be “stable”
- 1.6
 - Stable branch for everything else
- 1.4
 - Legacy release branch
 - Rather long in the tooth

Changes to master

- Rx enhancements (atomic counters, correctness and performance fixes)
- Heimdal crypto (libdes killer)
- New token kernel interface
- Reading during ubik writes
- krb5.conf-style preferences file support
- Identity management layer for Rx
- Copious code reorganization and cleanup
- Libtool
- Restructured MTU and PMTU handling code in Rx
- * newer Linux kernel support
- Tests

Challenges

- Code cleanup makes patches not apply cleanly
 - Lots of porting work
 - Possible errors
- Libtool makes Makefile changes not apply
 - Errors less likely, still issues

Wither 1.6

- Plan for 1.6 to become 1.8 when Windows was ready
- Master has diverged heavily from 1.6
 - 1.7, less so from master
 - Additional testing of master needed

Forward plan

- 1.6.2 release
 - Linux kernel changes
 - Other substantial bug fixes
 - Mostly already pulled to 1.6.x
- Absent many additional volunteer hours, anyway

Status of Win7 Netbios Name Lookup Bug

- 2011 EuroAFS:
 - Microsoft has officially declared the bug “WONT_FIX”
 - The IFS is the only fix that OpenAFS can provide to the community
- Microsoft IFS PlugFest (Feb 2012) the root cause was identified
 - Not Netbios related!!!!

SMB 1.x GSS SPNEGO authentication error

- The SMB specification permits the server to save a round trip in the GSS SPNEGO negotiation by sending an initial security blob.
- Windows 7 / Server 2008 R2 SMB 1.x redirector ignores the blob after initial connection.
- SMB 1.x reuses the original authentication context.
- Workaround:
 - The SMB 1.x server sends no security blob in the SMB_COM_NEGOTIATE response.
 - Force the client to send an initial GSS init_sec_context blob.
- The Deadlock:
 - After a SMB disconnect, reconnections appear to fail due to SMB connection resets.
 - The SMB 1.x redirector will retry indefinitely
 - All threads with outstanding requests to \\AFS will block
 - Reboot required

1.6.1

- Workaround for Win7 SMB 1.x Reconnect Bug
 - GSS SPNEGO optimization error
- Microsoft is working on a patch
 - Does anyone care?

1.6.1 – other changes

- VBUSY failover
- Improved idle dead time handling
- NAT ping constraints (one rx conn)
- Restrict processor affinity to 2
- Microsoft Advanced Firewall support

1.6.2

- VNOSERVICE processing
 - Indicates that file server did not process the RPC request
 - Triggered by file server idle dead timeout
 - Safe for client to retry

1.7 News

- 1.7.17 is current
 - 16 releases since DESY conference
- All 1.6.x improvements
- Windows 8 and Server 2012 support
- Explorer Shell integration
- Short Name generation
- Integrated Logon changes

Windows Short Names

- Short names are optional as of Windows 7
- 1.7 does not generate short names on Windows 8 and above
- Anti-virus vendors are thrilled
 - Reduced memory and CPU utilization
- Faster path evaluation
- Short names can be disabled on Windows 7 in 1.7
 - “ShortNames” TransarcAFSDaemon Parameter

1.6 -> 1.7 Upgrades

- 1.7 and beyond will no longer provide:
 - Windows 2000 support
 - afscreds.exe
 - afs_config.exe
 - SMB Submount functionality
 - NSIS (EXE) installers for 32-bit Windows
- Drive letter mappings to “Microsoft Network” must be deleted
- Integrated Logon changes for LOCALHOST
 - Long delays when mis-configured

Integrated Logon: Four Logon Domain Types

- Local Machine Account
 - (LOCALHOST domain)
- Domain or Forest Account
- Domain or Forest Account NETBIOS-compatible name
- Kerberos Principal mapped to a local or domain or forest account

Integrated Logon: Per Domain configuration

- Obtain AFS Tokens?
- Alternate Kerberos realm?
 - Required for LOCALHOST
- Tokens for additional cells?
- Error handling?
- Per user configuration
 - Name mapping?
 - All other options

Integrated Logon: Registry Hierarchy

- HKLM\SYSTEM\CurrentControlSet\Services\TransarcAFSDaemon\NetworkProvider\Domain key.
- For example:
 - ... \Domain\LOCALHOST\
 - ... \Domain\LOCALHOST\Administrator\
 - ... \Domain\AD\
 - ... \Domain\AD.EXAMPLE.ORG\
- Full domain name and the NETBIOS-name are separate entities.

Known Issues

- 10 second Extent processing stalls
 - Race between kernel and service
- Object Information / File Control Block dependency race
- Kernel memory pressure when large numbers of directory entries are evaluated

Blue Screens of Death

- BSOD reports are almost always triggered by Anti-virus or other filter driver interactions
- Some sites experience none
- Others experience weekly crashes

Changes to governance

- New gatekeeper
 - Ken Dreyer (USGS)
- New approvers
 - Plan not yet fully implemented
- Branch managers
 - Once we have volunteers
 - ...
- Moving on the foundation again
 - foundation-discuss@openafs.org

Money, money, money

- Total cost so far for 1.7 is approaching \$1.6 million
- End user organizations are asked to spend \$20 per in use copy

Credit where due

- Lots of people, lots of companies contributed to getting OpenAFS here.
- <http://www.openafs.org/credits.html>

Have a bug, send a report

- Do not assume that someone else has reported your bug
- BSODs are frequently triggered by environmental factors
- Ability to reproduce locally is limited
 - openafs-bugs@openafs.org
 - <http://www.openafs.org/support.html>

Talk back to us

- Mailing lists:
 - Openafs-info <http://lists.openafs.org/mailman/listinfo/openafs-info>
 - Openafs-devel <http://lists.openafs.org/mailman/listinfo/openafs-devel>
- IRC chat room: #openafs on freenode
- Jabber developer MUC:
openafs@conference.openafs.org