secure endpoints

OpenAFS for Windows: One Year Later

Jeffrey Eric Altman

jaltman *at* secure-endpoints *dot* com

Why am I here at AFSig.se 2004?

- Introduce myself to the community
- Describe the state of OpenAFS on Windows today
- Describe the issues which must be solved
- Offer proposals for future directions
- Obtain your feedback

Who am I and what do I do?

- OpenAFS Gatekeeper for Windows
  - Audit code submissions
  - Manage Bug Requests
  - Build releases
  - Fix things
  - Plan for the future
- MIT Kerberos for Windows maintainer
- Internet Engineering Task Force (IETF)
- Project JXTA Board Member

secureendpoints

What else have I done?

- The Kermit Project
  - Cross platform (Unix, OS/2, Windows)
- Internet Access Methods
  - Java based Person to Person collaboration software
- Miscellaneous Network Security stuff
  - OpenSSL, Secure Remote Password, TELNET START_TLS, FTP AUTH TLS, SSH

secureendpoints

How bad things were in 2003 ...

- OpenAFS on Windows was under supported
- Other than the work added in 1.2.8 there were been close to zero changes since 1.0
- Submitted patches could not be applied as there was no one to audit them
- Bugs placed in RT could not be responded to.

secureendpoints

There is a new sheriff in town

- All items in RT queue have at least been responded to if not fixed
- Outstanding patches have been applied
- Code submissions obtained and integrated
- Resource leaks plugged
- “Stable” OpenAFS 1.3.75 announced
- Monthly Status Reports
- Active Development in Progress

1.2.11 vs. 1.3.75: Which definition of "stable" do we mean?

1. "Stable" meaning that the code does not change very much from release to release providing predictability: No
2. "Stable" meaning that the code performs reliably without crashing unexpectedly or adversely impacting the performance of the system: Yes

secureendpoints

Reasons 1.2.x is Not a Stable Release

- Un-initialized variables
- Memory leaks due to reference count management errors
- Kernel object leaks due to reference count and usage errors
- Thread deadlocks due to recursive use of single use lock implementation

More reasons 1.2.x is Not Stable

- Memory allocated in one DLL is de-allocated in another
- Operations which require both a pioctl and a RPC to send private data (ktc_GetToken and ktc_SetToken) are not atomic
- AFS RPC connection management errors can down your AFS file servers if crypt mode is used
- Integrated Logon DLL can prevent Windows XP SP2 from booting

secureendpoints

Even more reasons ...

- The number of NetBIOS control blocks used in protocol operations (100) exceeds the number of objects which Windows can wait on simultaneously (64).
- SMB messages with the “extended” bit set were not supported preventing file operations from being performed on a subset of files.

What is new in 1.3.75?

- Code Donations from:
  - Asanka Herath
  - Rob Murawski
  - Joe Beuhler
  - MIT
  - Morgan Stanley
  - Secure Endpoints
  - Sine Nomine
  - others
- New functionality
- Improved Performance
- Improved Reliability
- Two New Installers (NSIS and WiX MSI)
- Improved Developer experience

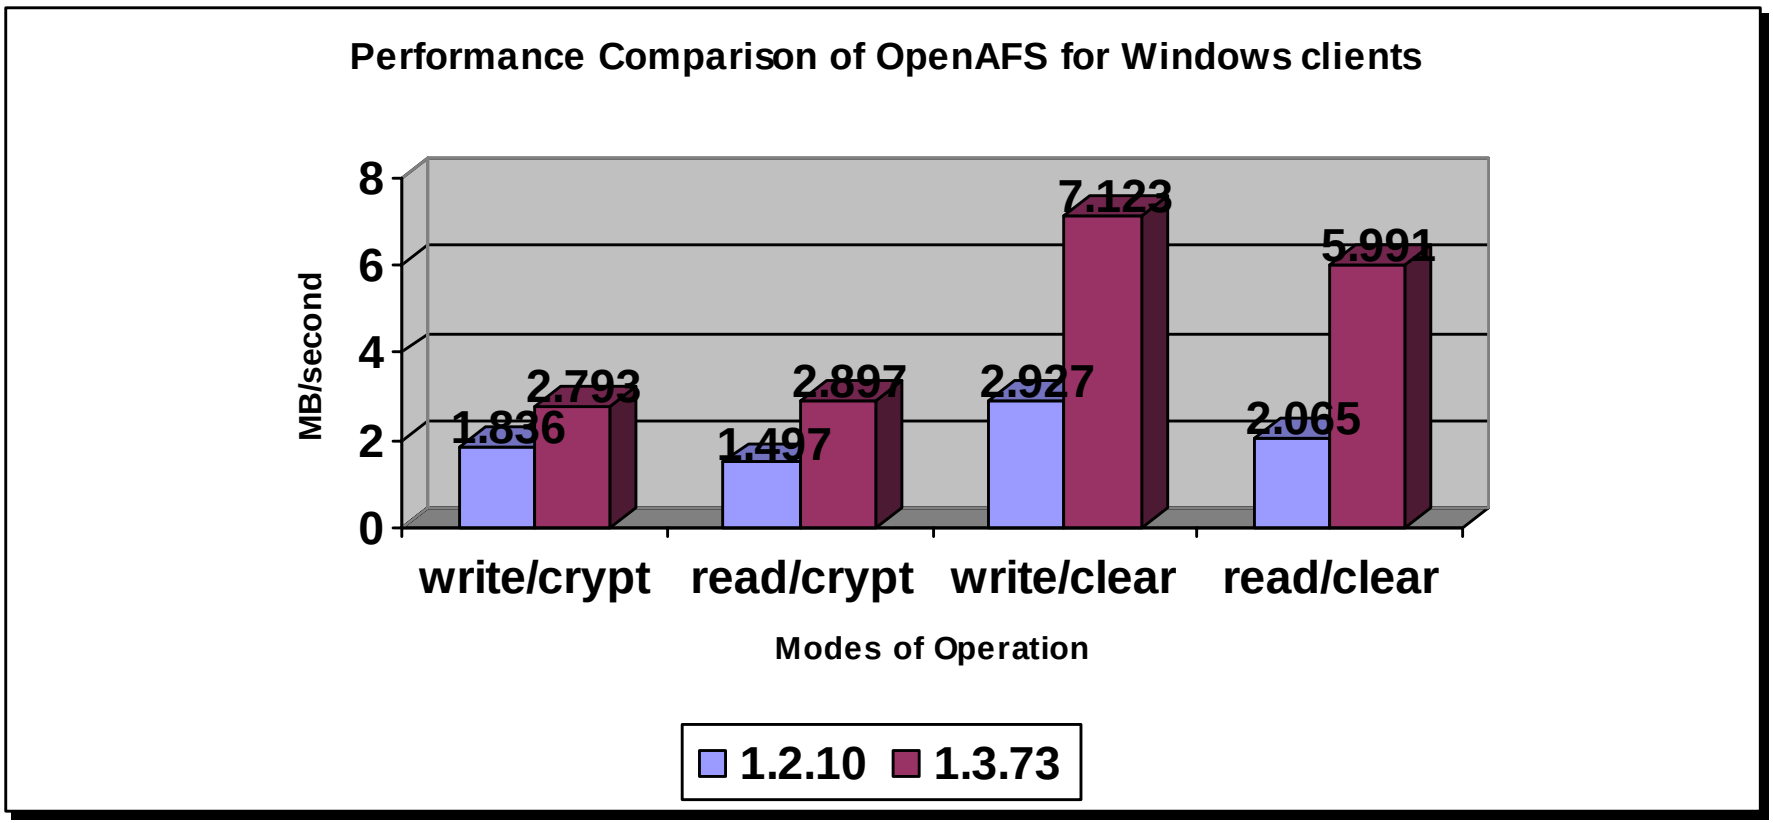

write/crypt mode: 52%; read/crypt mode: 93%;
write/clear mode: 143%; read/clear mode: 190%

SMB Server Improvements

- Extended CIFS messages are now supported
- SPNEGO authenticated connections
- Unique Virtual IDs assigned to each connection
- Garbage Collection of invalid or expired objects
- Partial support for SMB Share Browsing
- Soft symlinks properly followed
- Hard symlinks now supported
- Windows XP SP2 support (pioctl library)

Cache Manager Improvements

- Case-sensitive Directory Name lookups
- New filename pattern matching algorithm
- No more memory leaks
- Improved Callback Management
- Improved Call Timeout Management
- AFS RPC connections used more then once

New NSIS Installer

- Rob Murawski implemented a new installer using the Open Source Nullsoft Scriptable Installer Framework 2.0
- Supports new installs, uninstalls and upgrades from previous releases
- Designed for interactive installs (not an MSI)

NSIS Installer: Selecting Components

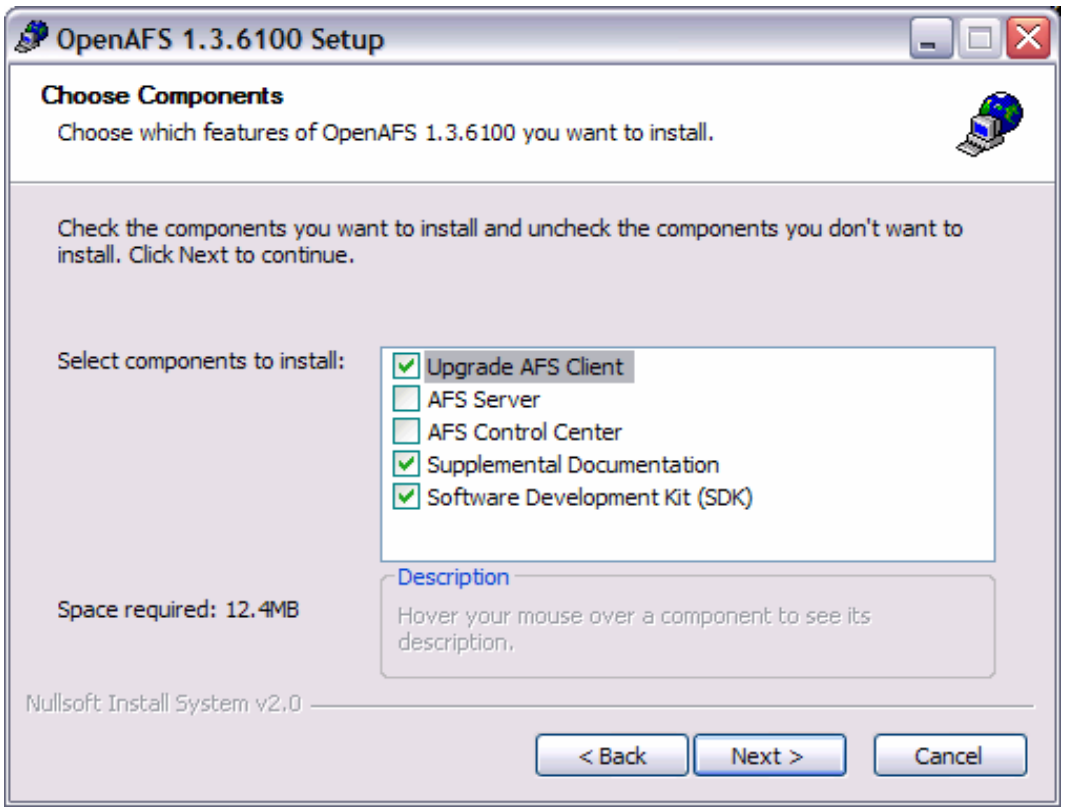

NSIS Installer: CellServDB

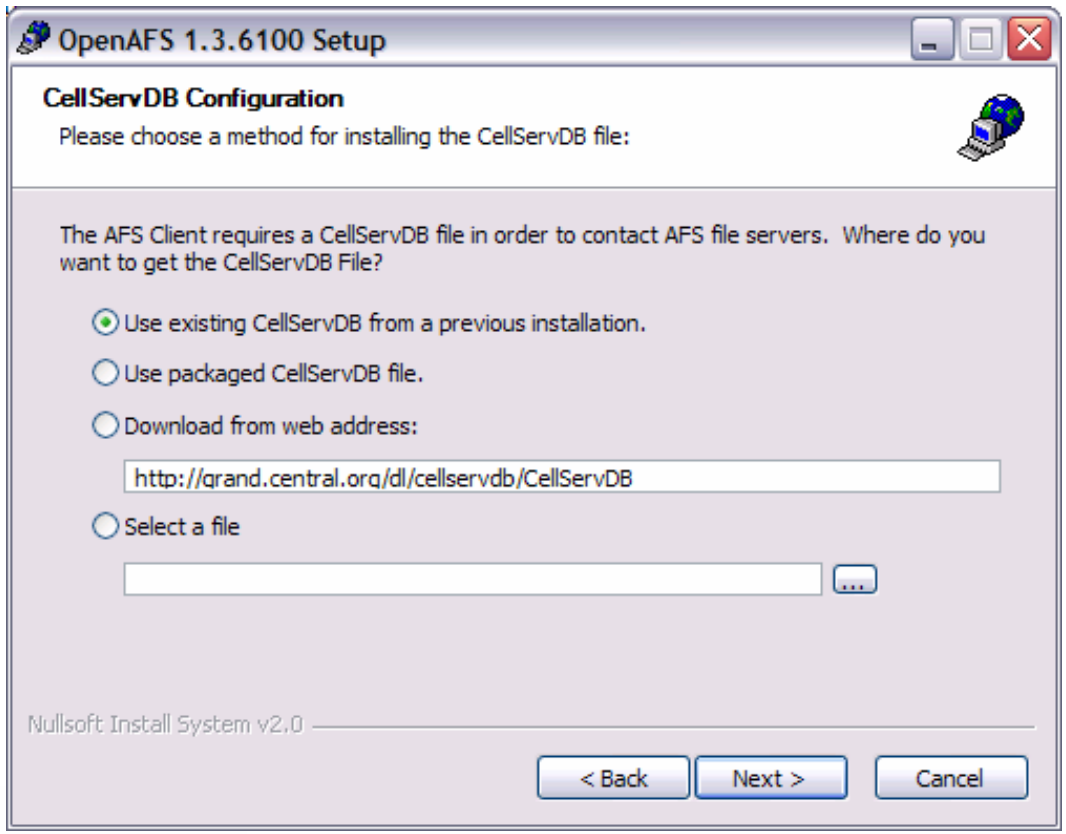

secureendpoints

NSIS Installer: Client Configuration

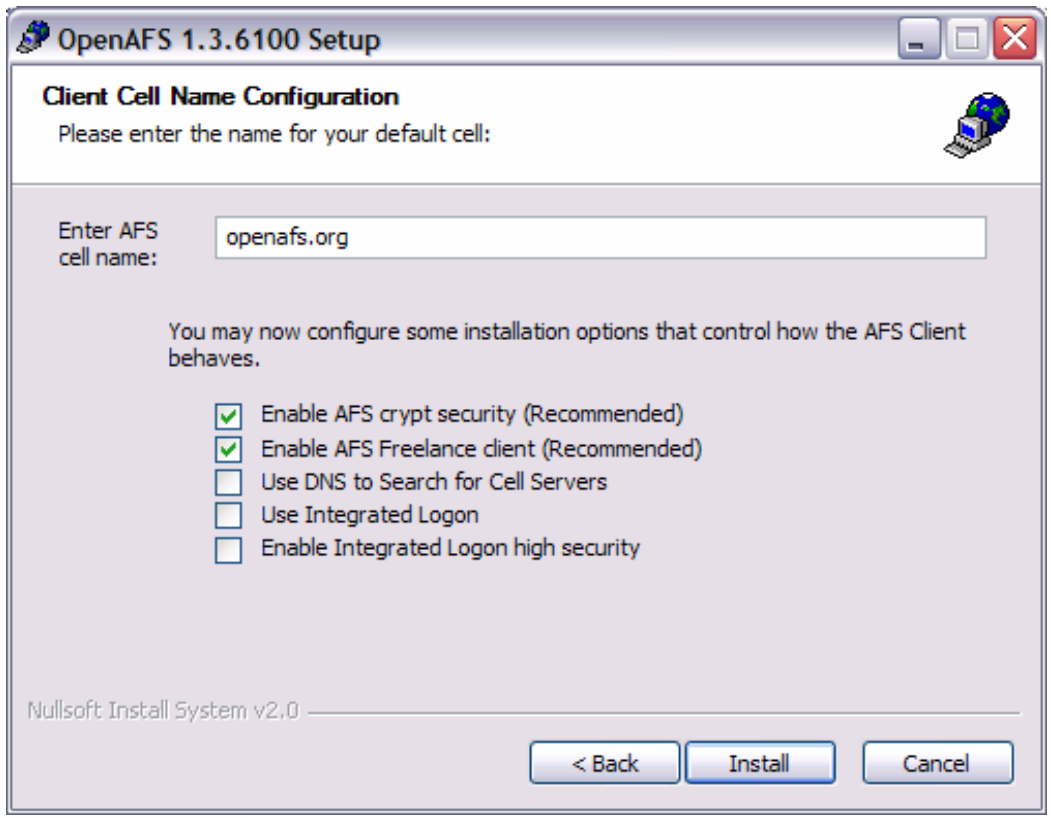

secureendpoints

Wix MSI Installer

- Asanka Herath implemented a new installer for use with Windows Group Policy deployements
- Easily customized via MSI Transforms to meet any organizations requirements
- Build a transform once, and apply to all OpenAFS.org MSI distributions

secureendpoints

Wix Installer Welcome Page

secureendpoints

Digital Signatures

- All binaries and installers are digitally signed by “Secure Endpoints Inc.” with a Thawte issued code signing certificate
- Signatures timestamped by Verisign
- Signatures used to ensure that installed files have not been replaced or modified
- Crash reports sent to Microsoft are now delivered to Secure Endpoints Inc. for analysis

New Build System

- Supports
  - Visual Studio .NET;
  - Visual Studio .NET 2003 (release builds)
  - Visual Studio .NET 2005
- Windows XP SP2 SDK required
- Only Windows 2000 and above.

secureendpoints

New UNC Path format

- UNC paths of the form \\afs\cellname are now supported when using the MS Loopback adapter
- The “NetbiosName” registry value can be used to specify alternatives to “afs”
- No longer need to use \\afs\all\cellname

secureendpoints

Integrated Login

- Keberos 5 used to obtain tokens
- All uppercase usernames are retried as lowercase upon failure
- Closed Security hole which leaked plaintext passwords on the wire
- WinLogon Event Notification handler added to destroy AFS tokens on logout
- Domain specific configuration is supported

secureendpoints

Windows Explorer Shell

- AFS UNC paths now supported by Explorer
- Browsing of \\AFS is now supported (limited to 13 character names)
- AFS Context Sensitive Menu works on all files and directories located within the AFS namespace (including Freelance)

secureendpoints

Microsoft Loopback Adapter

- Installed by OpenAFS.org installers (if missing)
- Provides locally visible adapter to bind the SMB service name
- Allows AFS Client Service to start when network is disconnected
- Prevents AFS Client Service panic on network disconnect

secureendpoints

Select LAN Adapter By Name

- Display name of the LAN adapter used to select adapter for AFS Client Service
- Simply name the desired LAN adapter “AFS”
- This functionality may be disabled via a registry value: “NoFindLanaByName”

Power Management

- Receipt of Standby, Suspend or Shutdown notifications trigger an automatic flush of all AFS SMB shares
- Only works for single user machines
- Needs to be revisited

secureendpoints

Cisco IPSec VPN compatibility

- AFS RPC packets restricted to 1292 bytes to pass through IPSec VPN
- Exact value used determined by “RxMaxMTU” registry value
- Hurts performance but not by much

secureendpoints

Profile Data

- HKLM\Software\OpenAFS\Client key used to set system default values
- HKCU\Software\OpenAFS\Client key used to store user configuration data
- Used for:
  - Token Expiration Reminders
  - Use of MIT Kerberos for Windows for Kerberos 5
  - Use of the Kerberos 5 to Kerberos 4 translation service
  - Show Tray Icon (afscreds.exe auto start)
  - afscreds.exe shortcut parameters
  - Freelance data (mount points and symlinks)
  - Submount entries
  - Drive mappings
  - Default Authentication Cell
  - Windows' SMB Client Side Caching configuration

Terminal Server Compatibility

- All configuration files have been removed from the %WINDIR% directory
- All user configuration data is now stored in the per-user registry allowing for multiple users and user instances
- SMB/CIFS sessions are authenticated to the logon session thereby removing the need for the random SMB Names utilized by “High Security” mode to enforce token ownership separation
- Corrected detection of the current logon user name

Windows XP SP2 Support

- Sets a magic registry value to permit the use of SMB/CIFS Service Names which do not match the local hostname
- Sets a magic registry value to allow SPNEGO authentication over a loopback connection
- Communicates with the Windows Integrated Firewall to dynamically open the ports used for callback messages
- DLLs no longer initialize the AFS RPC library in DllMain entry-point which caused the Integrated Logon support code to block the successful startup of the Windows operation system.

secureendpoints

Kerberos 5 Support

- Integrates with MIT Kerberos for Windows 2.6.5
- Obtain tokens using Kerberos 5 and optionally krb524d
- Imports credentials from both the MSLSA and CCAPI credential caches
- Automatically renews tokens and tickets as they approach expiration
- Architecture supports obtaining tokens for multiple cells from a single krb5 tgt (no UI)
- Not yet supported by Integrated Logon
- Can be disabled on a per user basis (no UI)
- Maximum token size increased to 12,000 bytes for Windows 2003 Active Directory support

secureendpoints

DNS AFSDB Support

- Cells not specified in the CellServDB may be discovered via DNS
- DNS AFSDB lookups enabled by default
- Windows DNS Query API now used instead of home grown implementation
- Controlled by “UseDNS” registry value

secureendpoints

Dynamic Root Volume (Freelance mode)

- On by default
- Supports r/w mount points
- Supports symlinks to arbitrary AFS paths
- Dynamically creates mountpoints for cells upon first access
- Stores local mount points and symlinks in Registry
- “fs mkmount” and “fs rmmount” may be used to configure mount lists
- Symlink command
- Controlled by “FreelanceClient” registry value
- Provides for better disconnected user experience

secureendpoints

Hidden Dot Files

- Following Unix tradition, files/directories whose names begin with a period are given the Hidden attribute when the “HideDotFiles” registry value is set

Logging Changes

- afsd_init.log and afsd.log moved to the %TEMP% directory (usually %WINDIR%\TEMP for the SYSTEM account)
- Stack Trace data logged to afsd_init.log during assertion failure or unhandled exception

secureendpoints

AFS System Tray Tool improvements (afscreds.exe)

- No longer requires Administrator Account
- Uses Kerberos 5 instead of kauth to obtain tokens
- Automatically renews Kerberos 5 tickets and acquired tokens
- Supports multiple Kerberos 5 principals
- One Kerberos 5 principal may obtain tokens for multiple cells
- Monitors network connectivity for IP Address changes, obtains tokens when cell becomes accessible
- Drive mappings restored at startup
- Automatic PTS registration of new user to foreign cells

secureendpoints

AFS Control Panel Tool improvements (afs_config.exe)

- Fields which modify the AFS Client Configuration now require logon account membership in the Windows “AFS Client Admin” group

Command Line Tools

- aklog.exe added
- afsshare.exe modified to support registry
- fs.exe and vos.exe no longer do stupid things
- UNC paths supported in fs commands

secureendpoints

Known Issues: New User Interface required

- Poor separation of functionality between tools
- AFS Client Configuration options displayed to non-administrative users
- Unable to choose between Kerberos 5 and Kerberos 4 on per-cell basis
- No ability to map multiple cells to a given Kerberos principal for single sign-on
- No ability to configure SysTray startup options
- No notion of default Kerberos principal
- Default authentication cell must match the root.afs cell
- No configuration for integrated logon options
- No configuration for all of the new functionality since 1.2
- AFS SysTray and MIT KFW Leash duplicate functionality

secureendpoints

Known Issues: AFS Client Service

- AFSD Client Service unable to handle dynamic changes to network configuration when MS Loopback Adapter is not installed
- SMB redirector overhead imposes performance restrictions
- Large File (> 2GB) support not yet implemented
- Memory Cache is not persistent

Known Issues: SMB Server deficiencies

- Restricted to 8-bit Code Pages; no Unicode
- Path names restricted to 256 characters
- Share names restricted to 13 characters
- Authentication restricted to 8-bit OEM Code Pages; no Unicode
- Remote Administration Protocol is incomplete
- Per message integrity protection and authentation (digital signatures) is not supported
- File writes are not cached and then written to AFS upon file close

Known Issues: Miscellaneous

- Need to synchronize with Unix sources
  - Multi-homed Server support
- AFS Server does not work reliably
- AFS Admin Tools do not work reliably
- Documentation
- AFS File System Support for:
  - DOS attributes (hidden, readonly, system)
  - Extended attributes
  - Streams

secureendpoints

3 year development roadmap

- Synchronize threads at shutdown (1st 2005)
- Implement memory based persistent cache (1st 2005)
- Design new UIs
  - Combined KFW and AFS System Tray Tool
  - AFS control panel for end-user configuration
  - AFS Administration tool for Client Service configuration
- 64-bit system support (Itanium and AMD64)
- SMB Unicode Support
- Large file support (> 2GB)
- Finish SMB Remote Administration Protocol
- Disconnected Mode
- Implement support for DOS attributes, EAs, and stream in AFS file system
- Implement SMB Digital Signatures
- Implement mini-port Installable File System

Using a Network IFS as the basis for an alternative architecture

- An IFS increases throughput by removing the delays imposed by the Ack/Nak SMB protocol and the overhead of protocol translation
- Once an IFS has been implemented, cache management can be performed by manipulating the contents of the Windows File System Cache instead of storing our own. Thereby reducing resource overhead.
- Use of an IFS would allow tighter integration with the Windows Security model. Tokens would be associated with user SIDs instead of SMB names.

secureendpoints

All IFS Kits are not created equal

- There are several issues to be concerned with when selecting an IFS Kit upon which to base development:
  - Licensing requirements of the kit. The license must be compatible with the OpenAFS.ORG license used for the existing code base.
  - Compatibility with existing and future releases of Microsoft Windows operating systems.
  - Cost of the kit. This higher the cost reduces the number of OpenAFS.ORG users who are capable of experimenting and producing binaries on their own. Fewer developers with access to the kit mean fewer developers who can contribute to the project simply by donating their time.

Comparison of Available IFS Kits

IFS Kit	License Requirements	Windows Compatibility	Retail Pricing
Microsoft IFS Kit	Royalty free. 60 day notice to Microsoft before shipping binaries. All binaries must be digitally signed by Microsoft.	API known to be compatible with shipping operating systems. MS reserves the right to alter the API in future operating systems.	$995 per developer per version (no support)
GNU IFS Kit	GNU Public License Version 2	Unknown	Free (no support)
OSR File System Framework	Royalty free.	OSR will upgrade their product to maintain compatibility with future releases.	$95,000 per commercial product (includes support)

- the Microsoft IFS Kit and the GNU IFS Kit contain licensing terms which are incompatible with the existing OpenAFS.ORG license
- Binaries produced with the OSR FSF are redistributable without royalty payments. OSR has in the past provided open source projects, CODA, with a modified version of their libraries at reduced pricing and may be willing to do so for OpenAFS.ORG.

Cross Platform Roadmap

- New Security Class to provide confidentiality and integrity protection based on GSS-API and Kerberos Crypto
- AFS RPC support for TCP connections
- Support for IPv6

secureendpoints

The Future of OpenAFS for Windows Rests in the Hands of the Community

- The work to be done is significant in scope
- How much can be done and in what time frame will be determined by the resources the Community can donate to the project.
- Donations can be in the form of person hours (programming, UI design, documentation) or money (support and work for hire contracts or general use grants).

secureendpoints

Q & A

You ask, I answer

Contact Information

Jeffrey Eric Altman

jaltman *at* secure-endpoints *dot* com