




SINE NOMINE  
ASSOCIATES

# AFS Tools Review



EUROPEAN AFS & KERBEROS  
CONFERENCE 2010

Michael Meffie <mmeffie@sinenomine.net>

Sine Nomine Associates

European AFS & Kerberos Conference 2010



# Agenda



New tools in the OpenAFS project  
Community contributed tools



## What's new in OpenAFS (toolwise)

We have some new vos commands

- setaddr: manually set server address

- endtrans: manually end volume transaction

fs option to assist in restoring acls

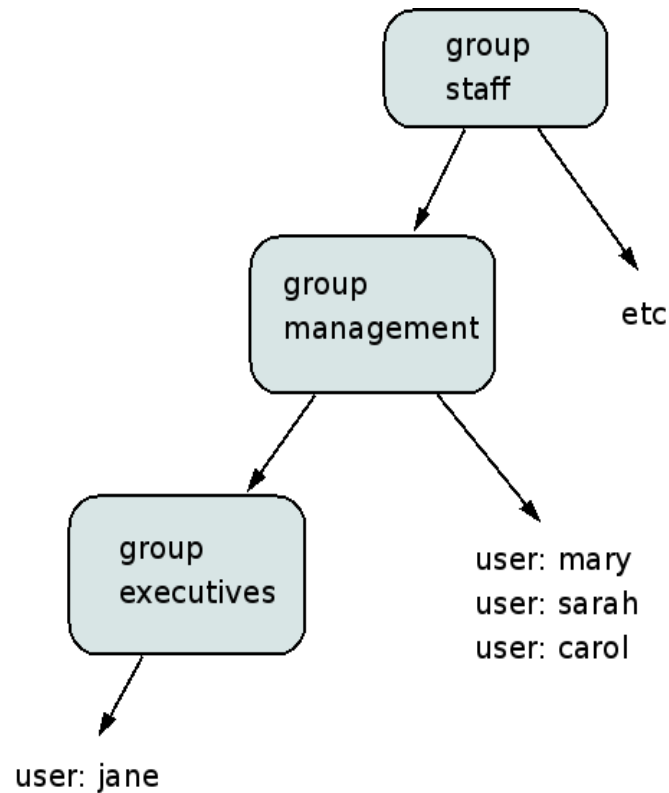
- fs la -cmd - output as 'fs setacl' command

afsmonitor - show more fs stats

Improved pts support for supergroups



# Example: -expandgroups



pts mem management  
-expandgroups

Expanded Members are:

jane  
mary  
sarah  
carol



## A new tool, afsio

- High throughput i/o of data to and from afs files.
- Requires a cache manager to create and stat files
- write, append, read data
- Author: Hartmut Reuter

## Some useful bits

- vldb\_check, prdb\_check
- volinfo, voldump, restorevol
- twiddle - adjust (some) cm parameters
- cacheout - flush fileserver cps cache entries
- afscp - basic afs file copy (testing)
- fsx - stress testing
- afsdump\_scan - CMU dumpscan



# Contributed Tools

- A few sites have made their tools available
- Some assembly may be required





# Stanford AFS Tool Suite

- Well documented, well maintained AFS and Kerberos tools used at Stanford University
- Authors: Russ Allbery, Neil Crellin, et al
- License: GPL and Artistic
- <http://www.eyrie.org/~eagle/software>



## Stanford AFS Tool Suite (cont)

- Volume Management
  - smart move, create, remove
  - delegated administration
  - balancing\*
- Monitoring and Reporting
  - meta-data and mount point database
  - nagios plugins, partition usage
  - read/write vs read-only tree diff

# Stanford AFS Tool Suite (cont)

And more!

- recursive fs command wrapper
- PAM modules
- k5start for daemons and cron jobs
- remctl for kerbertized remote control





# UNC Toolset

- University of North Carolina  
at Chapel Hill
- Author: Steven Joyce
- License: GPL\*
- <http://www.physics.unc.edu/~stephen/>




## UNC Tools (cont)

- AFS Disk Space Allocation
  - Visualize quota allocation of volumes for each partition, web based
  - Show volumes close to reaching quota limit
  - Show quota over commit (deficit)
- This is an in-house UNC tool, Steven offered to release



# AFS Disk Space Allocation

## **lithium.physics.unc.edu /vicepc**



## UNC Toolset (cont)

- phpafsfree
  - web-based AFS disk utilization
  - inspired by the tcl program afsfree
- BackupPC4AFS
  - use BackupPC to backup AFS volumes
- <http://www.physics.unc.edu/~stephen>



# BackupPC4AFS



## afs\_user\_s

- [afs\\_user\\_s Home](#)
- [Browse backups](#)
- [LOG file](#)
- [LOG files](#)
- [Edit Config](#)

## Hosts

afs\_user\_s

Go

## Server

- [Status](#)
- [Host Summary](#)
- [Edit Config](#)
- [Edit Hosts](#)
- [Admin Options](#)
- [LOG file](#)
- [Old LOGs](#)
- [Email summary](#)
- [Current queues](#)
- [Documentation](#)
- [FAQ](#)
- [SourceForge](#)

## Restore Options for afs\_user\_s

You have selected the following files/directories from share /, backup number #135:

- /user. /user. -2007\_08\_19\_23\_10\_29-L0.vdmp
- /user. /user. -2007\_09\_24\_23\_08\_10-L1.vdmp
- /user. /user. -2007\_09\_26\_23\_07\_42-L2.vdmp
- /user. /user. -2007\_09\_28\_23\_07\_18-L4.vdmp
- /user. /user. -2007\_09\_30\_23\_05\_24-L6.vdmp
- /user. /user. -2007\_10\_02\_23\_37\_23-L8.vdmp

You have three choices for restoring these files/directories. Please select one of the following options.

### Option 1: Direct Restore

#### Restoration of afs\_user\_s via vos

You can start a restore that will restore these volume(s) directly into AFS.

**Warning:** If you do not specify an extension, the VOLUME YOU'RE RESTORING WILL BE OVERWRITTEN!

Restore the files to host (don't change)

afs\_user\_s

Pass the following options to "vos restore"

-server SERVERNAME -partition PARTITIONNAME -e

Unused for AFS restores (don't change)

/user.

Start Restore

### Option 2: Download Zip archive

Archive::Zip is not installed so you will not be able to download a zip archive. Please ask your system administrator

### Option 3: Download Tar archive

You can download a Tar archive containing all the files/directories you have selected. You can then use a local appli the files.

**Warning:** depending upon which files/directories you have selected, this archive might be very very large. It might t and you will need enough local disk space to store it.

Make archive relative to /user.stephen (otherwise archive will contain full paths).

Download Tar File



# CMU AFS Tools

- Authors: Many and various
- License: BSD-type
- [ftp.andrew.cmu.edu/pub/](ftp://andrew.cmu.edu/pub/)
- [/afs/grand.central.org/dl/software/dumpscan](http://afs/grand.central.org/dl/software/dumpscan)







# CMU Tools

- dumpscan
  - tools and libs to examine and construct AFS volume dumps
- balance
  - auto-balance by size, volumes, access
  - needs a maintainer
- ws – walk subtrees
- venusmon - show cache manager fetches

# University of Michigan Tools

- Filedrawers – web-based AFS access
  - <http://www.filedrawers.org>
- Apache modules
  - mod\_waklog, mod\_cosign
- webafsmen – web-based admin
  - <http://rsug.itd.umich.edu/cgi-bin/cvsweb.cgi>



# Mitre AFS Tool Suite

- A set of python tools for AFS admin
- Author: Jeff Blaine
- License: GPL v2
- <http://ats.sourceforge.net>





## Mitre AFS Tool Suite (cont)

- smart remove
- volume sanity check
- replication policy check
- assist in placing volumes by space
- report quota commitments by partition



# kladmin - Account Admin

- Account provisioning for MIT Kerberos, OpenLDAP, OpenAFS
- Author: J.J.E. Winius
- License: Perl
- Wrapper to keep ids in sync
- Mass account creation

<http://www.rjsystems.nl/en/2100-kladmin.php>


# AFS Manager





- Web-based AFS administration tool
- Authors:
  - Fabrizio Manfredi, Peppe Guarino
- License: TBD
- Contact:
  - Fabrizio Manfredi (manfred)
  - [fabrizio.manfredi at gmail.com](mailto:fabrizio.manfredi@gmail.com)









# AFS Manager (cont)



AFS Manager 


Cells    

- Cell
- Pts
- Server

Volume  Partition  Disk Space 

Add Volume | Release | Remove Volume | Rename Volume | Clone Volume | Search  

id	Vol Name	Server	Partition	Type	Status	Need Check	Dim	%
	example	afs01	a	RW	0			
	user.mrbig	afs01	a	RW	0			

Page 1 of 1 

# AFS Space Statistics (afsss)

- Perl script to generate reports about AFS fileserver partitions and volumes
- Author: Stephan Wiesand
- <http://www-zeuthen.desy.de/~wiesand/AFS/>



# AFS Tools for Mac OS Clients

- AFS Commander
- afsinstall
- Contextual Menu for finder
- afstokens

<http://www.ibiblio.org/macsupport/afs/afstools.html>



# Sine Nomine AFS Tools

- Contributed tools and helper scripts for AFS
- <http://git.sinenomine.net/git>
- project `afs-tools.git`





## Sine Nomine afs-tools (cont)

- afs-audit-volumes
- Perl script to report volume inconsistencies or errors, vos listvol vs vos listvldb
- Useful for detecting orphaned read-onlies

## Sine Nomine afs-tools (cont)

- afsmon - SNIPS plug-in
  - perl script plus SNIPS config files
- AFS fileserver monitoring
  - fileserver up/down (bos status)
  - fileserver busy (calls waiting)
  - partition space (vos listpart)

<http://netplex-tech.com/snips>

<http://git.sinenomine.net/git>, [afs-tools.git](http://git.sinenomine.net/git)



# Sine Nomine afs-tools (cont)

## afsmon for SNIPS

```
SNIPS TextView                               Fri May
Device      +-Variable--+ -Value--+ -Units--+ State
fs+bm1vsrv04      up          0          Info
fs+bm1vsrv04      calls_wait  0          Info
fs+va            up          0          Info
fs+va            calls_wait  0          Info
fs+tx2srv03       up          2          Critical
fs+tx2srv03       calls_wait  0          Info
fs+tx2srv04       up          1          Critical
fs+tx2srv04       calls_wait  0          Info
--X--
```



## Sine Nomine (cont)

- `cw_graphify.pl`
  - create gnuplot graphs of calls waiting for thread data for all the file servers in a cell
  - perl script
  - author: Andrew Deason
  - <http://git.sinenomine.net/git>, `afs-tools`, `git`

## Sine Nomine afs-tools (cont)

### *volscan (work in progress)*

- Scan partitions to list information about vnodes
- Show fids, mount point targets, directory acls, paths (relative to volume root), filter by partition, volume, vnode type
- Uses fs-sync to checkout volumes, so can be run while the fileserver is running
- Very much like Kim Kimball's vol-details program
- Written by Tom Keiser, with updates from Andrew Deason and Michael Meffie.



## Sine Nomine afs-tools (cont)

- ideas for volscan
  - nightly generate mount point index
  - audit acls to find system:anyuser abuse
  - nightly generate fid to file index for fileserver audit log analysis (see Kim Kimball's vol-details presentation from AFS BPW 2010)





## Sine Nomine afs-tools (cont)

Usage: volscan

[-online]

[-partition <eg. /vicepa (default: all)>]

[-volumeid <volume id>]

[-voltypes <all, rw, ro, bk>+]

[-vnodetypes <all, file, dir, mountpt,  
symlink>+]

[-format <output format>]

[-delim <field delimiter>]

Where: -online Synchronize with running  
fileserver



## Sine Nomine afs-tools (cont)

```
# volscan
```

```
dir = {  
  voltype=RW, volname=root.afs, parentid=536870912,  
  fid=536870912.1.1, dv=3,  
  path=/  
  acls = { size=36, version=1,  
           total=2, positive=2, negative=0,  
           acl = { { id=-204, rights=rlidwka },  
                  { id=-101, rights=rl } } } }  
  
mountpt = {  
  voltype=RW, volname=root.afs, parentid=536870912,  
  fid=536870912.2.2, dv=1,  
  path=/example.com,  
  target=#root.cell }
```



## Sine Nomine afs-tools (cont)

```
# volscan -partition /vicepa -voltypes rw  
-vnodetypes mountpt -format dsv -delim :  
  
mountpt:RW:root.afs:536870912:536870912.2.2  
:1:/example.com:#root.cell  
  
mountpt:RW:root.afs:536870912:536870912.4.3  
:1/.example.com:%root.cell  
  
mountpt:RW:root.cell:536870915:536870915.2.  
2:1:/user:#user
```



## Sine Nomine afs-tools (cont)

Recursively walk up the directory to find the the relative file paths.

```
# volscan /vicepa 536870921 -f dsv -d :
```

```
...
```

```
file:RW:user.mike:536870921:536870921.2.20  
5:1:/testdir/undertestdir:
```



## Sine Nomine - afsdb updates

- Work in progress, updates for the Stanford AFS Reporting Database (afsdb)
- Configurability for site-specific data
- Replaced Standford specific DBI wrappers.
- Added support for Mysql
  - published afsdb only supports Oracle



## Sine Nomine - afsdb updates (cont)

- More columns and tables
  - added vnode weekly access counts
  - new vice partition tables
  - optional history table for last N days
- More queries for afsdb-query
  - ranking reports for most accessed volumes
- New afsdb-cgi script for web-based info on servers and volumes
  - (work in progress, very barebones today)



# Sine Nomine - afsdb-cgi

## OpenAFS Reporting Database

[servers](#) | [active volumes](#)

Volume: **user.mmeffie**

Volume	Id	Type	Used	Status	Server	Part	Parent	Readonly	Backup	Quota	Created	Updated	Accesses	W
<a href="#">user.mmeffie</a>	536871188	RW	19746	OK	<a href="#">afs01</a>	a	536871188	0	536871190	0	2006-05-23 12:24:35	2010-08-10 08:21:07	10	14

History

Recorded	Accesses	Weekly
1282157614	10	146
1282157569	10	146



# Thank you



Michael Meffie <[mmeffie@sinenomine.net](mailto:mmeffie@sinenomine.net)>