

Table of Contents

Operations Help Guide AFS_SERVERS.....	1
Contents.....	1
Service Description:.....	1
Logging in.....	1
Servers status.....	1
Startup and Shutdown Procedures:.....	2
afs_startup.....	2
afs_rundown.....	2
Old rundown procedures.....	2
Backups:.....	2
Status Checking.....	2
AFS status checking (all servers).....	2
Status checking (individual server).....	3
Machine blocked/hung out of hours.....	3
Basic trouble-shooting.....	3
Restarting the AFS file server process.....	3
Rebooting AFS servers.....	4
Soft reboot of AFS servers (all except afs7X).....	4
Rebooting AFS LINUX scratch servers (afs6X).....	4
Hardware problems - calling the manufacturer.....	4
IBM servers.....	4
SUN.....	5
Transtec disk servers.....	5
How to look at the error log on a SUN.....	5
Alarms.....	5
Offline volume(s) in partition(s): /vicepzz.....	5
False no_contact alarms following a "network" problem.....	6
Contacts and Support:.....	6
Keywords.....	6

Operations Help Guide

AFS_SERVERS

Contents

- Service Description
- Startup and Shutdown Procedures
- Backups
- Status Checking
- Machine blocked/hung out of hours
- Contacts

Service Description:

AFS is a network-distributed file system, supported by Transarc, which consists of AFS server machines located in the computer centre and AFS clients which are UNIX/Linux workstations or PCs running the AFS client, located all over CERN. If a user logs in to his account on LXPLUS, his files are located on one of the AFS servers. If this same user runs a batch job on LXBATCH, these same files will be located on the AFS server. AFS is a critical service at CERN and data should be available around the clock, throughout the year.

Logging in

To perform checks etc, please use from a UNIX session:
(ssh is a "secure telnet" - passwords are encrypted)

```
ssh -l ops afsxx
```

Servers status

Please note that not all monitored AFS servers are in production.

- ◆ **Production servers** can found by executing the script:
`/afs/cern.ch/project/afs/etc/afsconf.pl -s`
on the AFS console account. (Just type `afs_config` on afs console account).
- ◆ The location of the server can be found by looking to the column "Location Room" (from result of previous link)
- ◆ **Scratch Servers AFS6xx** are scratch servers running LINUX. They are less critical than the main AFS File Servers - see special instructions. The volumes on the scratch servers are NEVER backed up.
- ◆ **Backup servers** Many servers run backups overnight. **AFS15** controls backups. It does not contain any production volumes. Backups end up in TSM and in CASTOR.

Startup and Shutdown Procedures:

afs_startup

After a power cut, not all machines will come back automatically. The ideal sequence is as follows:

1. Turn on all SUN & IBM RAIDs (= disks)
2. Turn on afsdb1, afsdb2, afsdb3, in any order, in quick succession.
3. Wait 5 minutes or until one of afsdb1,2,3 is 'alive'.
4. Turn on afsmisc2 & afsmisc1.
5. Turn on afsdb4 & afsdb5 in building 613.
6. Turn on all other AFS servers, in any order.

afs_rundown

3-April-1998

The rundown procedures have now been modified

Login in to the AFS **console** account (e.g. on lxadm, lxplus, etc...) and use the **afs_rundown** command.

Please note the difference between the the AFS console account and the system console accessed by the **connect2console.sh** command.

The command **afs_rundown** is in `/afs/cern.ch/user/c/console`

```
usage : ~/afs_rundown [-reboot] [-all|-afs|-afsnfs|-dce|-host
```

Old rundown procedures

1. Login to the local account as **ops** (the AFS ops account for is CONSOLE)
2. Issue **rundown** or **rundown -r** for a reboot and wait 3 to 5 minutes
3. In case of reboot, check that the status command returns

```
fs is running normally, auxiliary status file server running
```

Backups:

AFS disaster/recovery :

Currently AFS restores can be requested by sending a mail to afs.support@cern.ch These requests are processed by SysAdmin Team.

Status Checking

AFS status checking (all servers)

On any terminal, run: `/afs/.cern.ch/project/afs/etc/afscnf.pl -c`

This will timeout on the failing server or at least show a drastically increased response time.

Status checking (individual server)

First method:

1. Login as **ops**, always log in from the server console or through the `connect2console.sh` command. Do not rely on telnet sessions.
2. Issue the `status` command. This
 - ◊ checks the fileserver processes on the AFS servers
 - ◊ checks the volume server processes on the AFS servers
 - ◊ checks the free disk space on the user (not project) home directories

Another way to detect if there has been/is a problem is to issue from any AFS account.

- ◆ `bos status afsxx -long`,
- e.g `bos status afs31 -long`

```
Instance fs, (type is fs) currently running normally.
Auxiliary status is: file server running.
Process last started at Wed Jul 31 22:25:32 1996 (5 proc starts)
Last exit at Wed Jul 31 22:25:32 1996
Last error exit at Wed Jul 31 22:19:21 1996, by file, due to signal 11
Command 1 is '/usr/afs/bin/fileserver -m 2 -L -nojumbo'
Command 2 is '/usr/afs/bin/volserver -nojumbo'
Command 3 is '/usr/afs/bin/salvager'
```

- ◆ Look at the last error exit time - this will tell you when the fs process last failed due to an error.
- If the output of a `bos status afsx -long` gives...
 - ◊ `has core file` - indicates process failed but this failure will not prevent the server being restarted.
 - ◊ `salvaging file system` - a file salvage operation is in progress and when finished, the fs process will be restarted.

Machine blocked/hung out of hours

Basic trouble-shooting

- run `/afs/.cern.ch/project/afs/etc/afscnf.pl -c` (it will block on the failing server)
- `bos status afsXX` (e.g. `afs42`)
 - this might indicate that the machine is actually salvaging in which case wait 30-40 minutes for the salvage to complete
- check for `no_contact` alarms; try to ping suspected server
- are there any network problems? AFS needs network to work properly
- check the console
- check the machine and its disks
- if no obvious solution in sight, problem looks serious (user complaints), then **reboot the machine** (see below)

Restarting the AFS file server process

Whenever an AFS file server seems to be blocked :

- ◆ users complain that a certain file is not accessible and that file is on `afsXX`
- ◆ `/afs/cern.ch/project/afs/etc/afs_checkservers` stops just before `afsXX` or

shows an access time for afsXX that is that is greater than 20'000'000 (unit is microseconds, therefore 20 seconds),

- ◆ no opportunity to inform the AFS team

you can restart the AFS fileserver process using the command:

```
/afs/cern.ch/project/afs/etc/fs_restart afsXX
```

where afsXX is afs32, afsdb1, or whatever AFS file server. You run this comand from the **console** account on a normal machine. You could run it from the file server console itself, but then you'd have to `klog.krb console` first. This procedure only restarts the file server process. It is pretty quick (much faster then a complete reboot), relatively non-intrusive in that applications would usually just block shortly but not fail, and collects major error logs. Always try it before a reboot - but in case of a serious OS problem a reboot might still help where a file server restart does not.

Rebooting AFS servers

- ◆ Always try a "soft" reboot first, i.e. use the `afs_rundown` command (see below)
- ◆ If this fails, reboot from the console.
 - ◇ On Sun servers, connect via the console manager and hit `RETURN, TILDE (~), CONTROL-B` wait for 'ok' then type `sync`.
 - ◇ On Linux, hit `alt-ctrl-f1`, then `alt-ctrl-del`.
- ◆ some machines have a reset button
- ◆ If this fails: power the machine off, wait 1 minute, power it on again.

Soft reboot of AFS servers (all except afs7X)

- ◆ First try to contact `afs.support` (see contact details below).
- ◆ If all the above checks have been run AND there are user complaints, try to reboot the system using the `afs_rundown` command from the afs console account:

```
~/afs_rundown -host afs44 -reboot
```

Rebooting AFS LINUX scratch servers (afs6X)

- ◆ **These instructions are for the AFS6x machines only !**
- ◆ They are real scratch servers - the difference is that the files/volumes are not backed up.
- ◆ In case of problems outside working hours (user complaints, alarms)
 - ◇ soft reboot the server i.e. `~/afs_rundown -host afs6x -reboot`
 - ◇ if this fails try a reset
 - ◇ if this fails, leave the machine down, mail `afs.support@cernNOSPAMPLEASE.ch` and call the `afs.support` team at a reasonable hour

Hardware problems - calling the manufacturer

IBM servers

Currently running: afs90-afs97 with an DS4300 RAID array controller. Equipment is covered by a 24x7, 4 hour intervention time contract! In case of H/W trouble, call IBM at 0800-55-5454. You need the machine model and serial number: both are on a small black label at the front of the machine.

SUN

Equipment is covered by a 24x7, 4 hour intervention time contract! In case of trouble call Sun, you need the serial number of the device.

Transtec disk servers

issue a ITCM ticket

How to look at the error log on a SUN

- Login as ops on the AFS server which has problems.
 - Type: `more /var/adm/messages`
-

Alarms

Offline volume(s) in partition(s): /vicepzz

During working hours, please log the problem and send a mail to **afs.support**.
Outside normal hours apply the following procedure.

The above message indicates that one or more volumes has gone offline, i.e unavailable. Depending on the volume(s), this could be a minor or major problem.

From the afs console account, type:

```
vos listvol afsxx zz | grep -v "On-line"
```

(where *xx* is afs server concerned and *zz* is the partition)

One or more volumes will be listed offline in the listvol. If the problem occurs during the day, mail AFS support

Salvage the volume if these conditions are met:

1. there is an "offline" alarm for this volume **and**
2. has a prefix of (p. user. s. or q.) - see types of AFS volumes below **and**
3. there are complaints from the user(s) **and/or**
4. other alarms indicate that they are AFS-related service problems.

To salvage a volume:

1. make sure the volume name does not end in `.backup` or `.readonly`, those cannot be salvaged!
2. log in as **console** on an AFS machine
3. type `/afs/cern.ch/project/afs/etc/salvage 'volume-name'`
4. salvaging a volume blocks all accesses for the time of the salvage, which can take several minutes so please be patient (**do not hit ctrl-c !**)
5. issue: `vos exam 'volume-name'`, it will indicate whether the volume is offline/online

Types of AFS volumes:

1. user volumes (e.g `user.tim`) normally belong to one user and affect only one user.

OpmOpProcAfsServers < Sandbox < TWiki

2. project volumes (e.g. `p.lsf`) which belong to experiments, projects and service and are shared by multiple users/services.
3. scratch volumes (e.g. `q.alsoft.scratch.0`) - general scratch volumes
4. personal scratch volumes (e.g. `s.cms.slehti.0`) scratch volumes belonging to a specific user.
5. system volumes (e.g. `sys.sun4x_56.34a.12`) - architecture specific.
6. other volumes (`root.afs`) used by the AFS cell - do NOT salvage these volumes.

Example

```
[lxplus055] ~ > vos partinfo afs20
Free space on partition /vicepa: 154903706 K blocks out of total 1717438
Free space on partition /vicepb: 49699179 K blocks out of total 88606368
Free space on partition /vicepda: 40807435 K blocks out of total 4990677
[lxplus055] vos listvol afs20 b | grep -v "On-line"
Total number of volumes on server afs20 partition /vicepb: 2642
**** Could not attach volume 537229091 ****
Total volumes onLine 2641 ; Total volumes offLine 1 ; Total busy 0
[lxplus055] ~ > vos exam 537229091
**** Could not attach volume 537229091 ****
RWrite: 537229091      Backup: 537229093
number of sites -> 1
server afs20.cern.ch partition /vicepb RW Site
[lxplus055] ~ > vos listvldb 537229091
user.johnh
RWrite: 537229091      Backup: 537229093
number of sites -> 1
server afs20.cern.ch partition /vicepb RW Site
[lxplus055] ~ > vos exam user.johnh
**** Could not attach volume 537229091 ****
RWrite: 537229091      Backup: 537229093
number of sites -> 1
server afs20.cern.ch partition /vicepb RW Site
```

False no_contact alarms following a "network" problem

If no_contact alarm persists even if the server is reachable, then on that server, type:

```
◇ ps -ef | grep monitor
◇ kill the pid for the process /usr/local/bin/perl ./monitor and the
  monitor will restart by itself.
```

Example:

```
ps -ef | grep monitor
ops 27392 1 0 16:56:36 ? 0:00 /usr/local/bin/perl ./monitor
kill 27392
```

Contacts and Support:

Please consult SDB for Cluster afs service.

Keywords

AFS servers salvage

This topic: Sandbox > OpmOpProcAfsServers

Topic revision: r2 - 2009-06-30 - FabioTrevisani



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)