

Affa

From SME Server

 Languages:	English
--	----------------



Skill level: Advanced

The instructions on this page may require deviations from procedure, a good understanding of linux and SME is recommended.

Contents

- 1 Maintainer
- 2 Description
 - 2.1 Affa features at a glance
- 3 Installation or Update
- 4 Quick start example
- 5 Configuration
 - 5.1 Job configuration properties
 - 5.2 Default configuration properties
 - 5.3 Global disable
- 6 Usage and command line options
- 7 Example setups
 - 7.1 Dedicated backup server
 - 7.2 Backup of single ibays
 - 7.3 Two production servers backup each other
 - 7.4 Backing up a Windows computer
 - 7.4.1 Rsyncd setup on the Windows computer
 - 7.4.2 Affa Rsyncd mode setup (Quick start example)
 - 7.4.3 Affa Rsyncd mode manual setup
 - 7.5 Use Affa to backup to a NFS-mounted NAS or a local attached USB drive
 - 7.5.1 Setup NAS
 - 7.5.2 Alternatively setup a USB drive
 - 7.5.3 Setup Affa
 - 7.5.4 Limitations
 - 7.5.5 Automount
- 8 Restore
 - 8.1 Restore single files or directories
 - 8.2 Full restore
 - 8.3 Moving a SME 7 server installation to new hardware using the Affa rise feature
 - 8.4 Restore from USB drive on new server
- 9 FAQ

- 10 Uninstall
- 11 Bug report
- 12 Additional information
- 13 References

Maintainer

Michael Weinberger

Affa stands for **A**utomatische **F**estplatten **F**ernarchivierung

Description

The main purpose of this affa package is to make a SME 7 Server a dedicated backup box in a few minutes. Affa backs up as many SME servers as you like or any other servers which have sshd running and rsync installed. Once it is configured, Affa runs reliably unattended and sends warning messages in case of any errors.

All backup archives are full backups, as Affa make use of the hardlink technique. Therefore a new full backup only needs disk space for the differences plus the filesystem overhead for the hardlinks and directories (which is typically 2-3%).

Affa is based on the rsync program and supports the rsync --compress option. This allows you to run backups over the internet or VPN. A typical setup is one or more Affa backup servers placed in different locations, which backup the production server(s) over the VPN.

A special feature is the rise option, which allows you to rise the backup server to your production server from a backup archive in case of a dead loss of your production server. The rise is executed within a extremely short time, even with huge amount of data. The rise feature uses hardlinks and therefore does not use up additional disk space.

Affa is a command line tool for system administrators and is intentionally designed without a GUI. Therefore it can be efficiently managed on the console and over slow internet connections.

Note: This documents refers to the Affa Version 2 Release Candidate. Latest stable version 1 documentation can be found here (<http://wiki.contribs.org/index.php?title=Affa&oldid=11974>) .

Affa features at a glance

- Affa is secure: All transfers and connections are made by using the ssh protocol with public/private key authentication
- Periodically runs unattended full backups. Only differences to the previous full backup are transferred over the network
- Using rsync with optional bandwidth limit and compression allows backups over the internet
- Uses hardlink technique, i.e. physical disk space only needed for the differences between two full backups
- Keeps a configurable number of scheduled, daily, weekly, monthly and yearly full backup archives
- The archives are browseable. No need to unpack them first.
- Archives can be mapped to a Samba share.

- Backup jobs are started by the cron daemon
- Interrupted jobs continuing with already transferred data when restarted
- Backups the default e-smith directories and files, when property SMEServer is set to yes
- Additional directories and files can be included
- Directories and files can be excluded from the backup
- Non-SME server linux systems can be backed up by setting the SMEServer property to no and using an include list
- In ESXi mode, running virtual machines can be backed up. See Backup of ESXi Virtual Machines using Affa
- Configurable nice level for rsync processes on the backup and source server
- Optional run of custom programs before and after a job run (e.g. running tape backup)
- Checks the disk space left after a job run with warning levels strict, normal or risky
- Extensive checking of failure conditions
- Sends failure messages to a configurable list of email addresses
- Sends a warning message, if the backup server runs out of disk space
- Installs an optional watchdog on the source server in case the backupserver fails (SME Server only)
- Watchdog sends warning, if an expected backup did not run (SME Server only)
- Watchdog sends a daily reminder message, if the error continues unchecked (SME Server only)
- Option to display current status of all jobs showing times of last and next run, size and disk usage
- Status can be mailed on a daily, weekly or monthly schedule
- Option to display all existing archives of a job shown date, number of files, size and bytes transferred from the source
- Option to send (and revoke) the public key to the source server (SME Server and ESXi only)
- Option to rise the backup server to a production server from a backup (SME Server only)
- The rise feature does not physically move data and therefore is extremely fast and needs (almost) no extra disk space
- Rise option can be run remotely as the NIC driver configuration of the backup server are preserved
- Compares installed RPMs on source with backup server. Sends warning message, if not in sync
- Undo rise option to restore the backup server
- Configurable via a e-smith style db, with one record for each job and a default record for all jobs
- Logs to /var/log/affa/JOB.log and /var/log/affa/affa.log with optional debug switch for higher verbosity
- Log files are rotated weekly, with 5 logs kept

Installation or Update

```
wget http://mirror.contribs.org/smeserver/contribs/michaelw/sme7/Affa2/smeserver-affa-2.0.0-rc4.noarch.rpm
wget http://mirror.contribs.org/smeserver/contribs/michaelw/sme7/Affa2/perl-Compress-Bzip2-2.09-1.2.el4.rf.i386.rpm
/usr/bin/yum --enablerepo=smecontribs localinstall \
smeserver-affa-2.0.0-rc4.noarch.rpm \
perl-Compress-Bzip2-2.09-1.2.el4.rf.i386.rpm
```

When you have installed Affa for the first time run the following command to initialize the Affa database

```
affa --make-cronjobs
```

and logout and re-login to the console to take the bash auto-completion (TAB key) in effect.

Quick start example

You have a SME 7 production server with hostname 'prodbox' and IP 10.200.48.1.
Set up a second SME 7 box as your backupserver with hostname 'affabox' and IP 10.200.48.2.

1. log into the 'affabox' and install the packages as described above.
2. copy the config helper script sample

```
cp /usr/lib/affa/jobconfig-sample.pl /root/prodbox-job.pl
```

3. edit /root/prodbox-job.pl and set

```
my $jobname='prodbox';
```

and

```
'remoteHostName'=>'10.200.48.1',
```

optionally, if the remote server port is configured to e.g. 2222 then set

```
'sshPort'=>2222,
```

4. write the configuration (this makes the database entries and sets up the cronjobs)

```
/root/prodbox-job.pl
```

5. generate the DSA keys and send the public key to the 'prodbox'



Note:

When initially doing this step, you will need to temporarily enable "Allow secure shell access using standard passwords" on the production server

```
affa --send-key prodbox
```

6. run the job manually

```
affa --run prodbox
```

Configuration

The configuration is stored in an e-smith style database. Use the db command to configure Affa. The jobname is the record key with the type 'job'.

To setup a new job with the name 'prodbox' enter:

```
db affa set prodbox job
```

then set the properties

```
db affa setprop prodbox remoteHostName 192.168.1.1
db affa setprop prodbox TimeSchedule '0030,0730,1130,1330,1730,2030'
db affa setprop prodbox Description 'My Production Server'
db affa setprop prodbox status enabled
```

and so on...

Alternatively you can use a script as described above in the 'Quick start' chapter.

To verify your work, type:

```
db affa show prodbox
```

Finally set up the cronjobs:

```
affa --make-cronjobs
```

Job configuration properties

Note: The default values shown in this table are the Affa program defaults and not to be confused with the preset values in the helper scripts, e.g. jobconfig-sample.pl.

Property	Value	Default	Description
remoteHostName	FQHN or IP		FQHN or IP of the source host (mandatory)
TimeSchedule	HHMM,HHMM,...		doesn't need to be ordered. At least one time is mandatory. Important: Using the proper format HHMM is essential. Affa does not check it. Badly formatted TimeSchedule will cause strange Perl errors.
Description	text string		
scheduledKeep	integer ≥ 1	1	how many of the scheduled backups should be kept
dailyKeep weeklyKeep monthlyKeep yearlyKeep	integer ≥ 0	7 4 12 2	how many of the daily, weekly, monthly or yearly backups should be kept
SMEsServer	yes or no	yes	when set to yes the default e-smith directories are automatically included and the property

			RPMCheck=yes can used
Include[0] Include[1] ...	full path		additional files or directories to include
Exclude[0] Exclude[1] ...	full path		additional files or directories exclude fi backup
RPMCheck	yes <i>or</i> no	no	<p>Only applicable to job that backups a SME server.</p> <p>Compares the package installation of the source host with this affa backup host. Sends a message with diff list not in sync. This check usefull, if you want have the option to rise the backup server to a production server from backup.</p>
DiskSpaceWarn	strict <i>or</i> normal <i>or</i> risky <i>or</i> none	strict	<p>run a disk space check after a job has been completed. With level 'strict' a warning message will be sent, if the available space is less than the size of the job completed backup. With level 'normal'/'risky' a message is sent, if less than 50%/10% of the backup size is still available.</p> <p>Note: When RootDir symbolic link to another filesystem the disk usage of the local filesystem rather than the linked filesystem is checked. Set the mountpoint as the RootDir to get correct checking.</p>
localNice	-19...+19	0	run rsync local process.

remoteNice	-19...+19	0	run rsync process on source niced.
Watchdog	yes <i>or</i> no	yes	Only applicable to jobs that backups a SME server. When a job is started affa installs a watchdog script on the source in /etc/cron.d/, which sends a warning message, if the next scheduled job (taken from the TimeSchedule property + 10 minutes) did not run. This guarantees, that you will be notified even in case of a affa server outage. The watchdog script sends daily reminder messages if the error continues. The next run job replaces the watchdog script with a new trigger time.
sshPort	service port	22	When sshd on the source host or your firewall listen on a non-standard port set the port here
ConnectionCheckTimeout	seconds	120	before the rsync process is started on the remote source host, affa checks the ssh connection and exits with an error after the configured time, if the host does not respond.
rsyncTimeout	seconds	900	Rsync exits, if no data transferred for the configured time. This avoids infinitely hanging in case of a network error.
BandwidthLimit	integer \geq 0 kilobytes per second	0	Limits the data transfer rate. A value of zero specifies no limit.
rsyncCompress	yes <i>or</i> no	no	compress the transferred data. May be useful on slow internet

			connections. Increase CPU load on source ; backup host.
EmailAddresses	name@domain.com,name@domain.com,...	admin	comma separated list mail addresses, where the messages should be sent to Note: By default Affa only sends messages on errors, never on success (see property <code>chattyOnSuccess</code>).
RetryAttempts	integer ≥ 0	3	when set to a value > 0 Affa re-run a failed job <code>RetryAttempts</code> times with a delay of <code>RetryAfter</code> seconds.
RetryAfter	seconds ≥ 3	600	when set to a value > 0 wait <code>RetryAfter</code> seconds before re-running the job after an error. Only applicable with <code>RetryAttempts</code> > 0
RetryNotification	yes or no	yes	when set to no, Affa does not send an error message when a job failed and <code>RetryAttempts</code> is configured. An error message is only sent when the last attempt has failed.
chattyOnSuccess	integer ≥ 0	0	when set to a value > 0 Affa sends a message after a successfully completed job run and decrements the <code>chattyOnSuccess</code> value. When the value has reached zero, Affa falls back to the default and only sends messages on errors.
AutomountDevice AutomountPoint	full path		Device and mountpoint of backup device (e.g. USB disk). Device is automounted before job starts and unmounted after job completion. With both properties

			empty no automount done.
AutomountOptions	string		An option string pass to the mount comma
AutoUnmount	yes or no	yes	When set to 'no' the automounted device mounted after the Af run.
preJobCommand postJobCommand	full path		programs (local on th affa server) to be executed before/afte job run. The job nam and type (scheduled, daily etc.) are passed arguments to the program. The exit co is additionally passed the post job comman program. See /usr/lib /affa/ for sample perl scripts.
RootDir	full path	/var/affa	where to store the backup archives, Do use /home/e-smith or /root as these are included in the backu and therefore the rise option will not work! Recommended: /var/
SambaShare	yes or no	yes	Access to the job archives via SMB
Debug	yes or no	no	set to yes to increase verbosity
status	enabled or disabled	enabled	with set to disabled, 1 cron entries will mad You can still run a jo manually.
rsync--inplace	yes or no	yes	set to no, if the rsync version on the source does not support this option (like rsync on SME6)
rsync--modify-window	integer >= 0	0	When comparing two timestamps, rsync tre the timestamps as be equal if they differ by more than the modify window value. This i

			normally 0 for an exact match. A value ≥ 0 is useful if you can't get the clocks of the source; the Affa server in sync
rsyncOptions	string		additional option string to be passed to rsync
rsyncdMode	yes or no	no	set to yes to connect the rsync daemon on the remote host (instead of running rsync over ssh)
rsyncdModule	string	AFFA	the rsyncd module name (only applicable with rsyncdMode=yes)
rsyncdUser	string	affa	the username for authentication to the rsync daemon (only applicable with rsyncdMode=yes)
rsyncdPassword	string		the password for authentication to the rsync daemon (only applicable with rsyncdMode=yes)
remoteOS	cygwin		with remoteOS=cygwin the options --send-key and --revoke-key use the account 'Administrator' and the correct path for the public key on a Windows/Cygwin remote host.
ESXi	yes or no	no	enable VMware ESX virtual machine backup mode See Backup of ESXi Virtual Machines using Affa
ESXiVMName	string		The name of the VM displayed in the VI Client
ESXiUsername	string		The name of the VI client with permission to create and delete snapshots

ESXiPassword	string		The password of the ESXiUsername
chunkFiles	string		Filenames that Affa should chunk and compress. Multiple filenames are to be separated by the slash character, e.g. 'mysql.dump/pgsql.dump' chunks the two files mysql.dump and pgsql.dump. With ESXi=yes it is implicit set to '*.vmdk'. See a command line option --chunk-archive and --unchunk-archive

Default configuration properties

All properties can be set as defaults in the DefaultAffaConfig record. This is useful, when you set up many similar jobs.

Example: You want to set the property 'localNice' to 19 for all jobs. Then run

```
db affa setprop DefaultAffaConfig localNice 19
```

and don't set this property for the jobs. Properties set in the job record overrides the defaults.

The special property 'sendStatus' is only applicable to the DefaultAffaConfig record. It controls the status report sent by email and can be set to the values 'none', 'daily', 'weekly' or 'monthly'.

To setup a weekly status report run:

```
db affa setprop DefaultAffaConfig sendStatus weekly
```

then setup the cronjob:

```
affa --make-cronjobs
```

Global disable

All jobs can be disabled with setting the AffaGlobalDisable record type to 'yes'.

```
db affa set AffaGlobalDisable yes
affa --make-cronjobs
```

to re-enable run:

```
db affa set AffaGlobalDisable no  
affa --make-cronjobs
```

Usage and command line options



Note:

Options can be abbreviated to uniqueness, e.g. --mak is equal to --make-cronjobs

```
affa --run JOB
```

Starts a job run. Usually done by the cronjob.

```
affa --make-cronjobs
```

Configures the cronjobs as scheduled in the jobs records.

```
affa --send-key JOB  
affa --send-key --host=TARGETHOST [--port=PORT] [--remoteOS=cygwin]
```

This first generates the DSA key for the Affa Server, if not already done. Then it sends the public key to the host 'remoteHostName' as configured in the record of job JOB and generates the job specific ssh known host entry.



Note:

When initially doing this step, you will need to temporarily enable "Allow secure shell access using standard passwords" on the production server.



Note:

By default, the --send-key option works for a SME Server as a remote server and for systems where the keys are stored in /root/.ssh/authorized_keys2 and the commands /bin/cat, /bin/touch, /bin/grep and /bin/mv are available. With remoteOS=cygwin it works for a Cygwin/Windows remote server.

```
affa --full-restore JOB [ARCHIVE]
```

Does a full restore of the standard and additional included files and directories from the backup ARCHIVE on the remote source server as defined in the JOB record. If ARCHIVE is not given, the

archive 'scheduled.0' is used as the default. The full restore reconstructs the server as it was at the time of the backup. After the restore the source host reboots.

```
affa --rise [--all] JOB [ARCHIVE]
```

Runs a full restore on the Affa server (!) of all standard files and directories from the backup ARCHIVE of job JOB. In other words: After completion, the Affa box reboots as a clone of the source server. Ensure, that the source server has been powered off before you reboot the Affa box, otherwise the network interface will not come up. This is important, when you run --rise remotely. The --rise feature only works with SME 7 servers und should only be used on dedicated backup servers.

With option --all, all files and directories of the archive as defined by the include[] properties are restored. Files or directories with the same name on the Affa server will be overwritten and cannot be restored by a undorise. This should not be an issue on a dedicated Affa server which does not hold any other data. After a possible undorise those additional restored dada must be removed manually.

Please note, that the rise process backs up the the Affa server itself before doing the restore from the archive. This backup is used by a possible undorise run to restore the Affa server. Only the standard files and directories are backed up. Data in non-standard loctions (like /opt) are untouched and will still exist after the rise run. See also: Backup_server_config#Standard_backup_.26_restore_inclusions

```
affa --undo-rise
```

This feature reverts a risen Affa box to a backup server. After a reboot, all configured jobs will work again.

```
affa --unchunk-archive JOB ARCHIVE
```

Concatenate and uncompress all chunked files in the archive ARCHIVE of job JOB to make the archive ready for a restore. Note that the most recent archive scheduled.0 is never chunked and therefore can be restored as it is without previous unchunking.

```
affa --chunk-archive JOB ARCHIVE
```

Chunk and compress all files that match the names in property chunkFiles to save disk space. The should only be used as the reversal of --unchunk-archive. Chunks are stored as bzip2 files in a directory named FILENAME.affa-chunks with 8 levels depth.

```
affa --create-backup-file JOB [ARCHIVE] [--outfile=FILE]
```

Creates a gzipped tar archive FILE from the ARCHIVE of job JOB. During creation the MD5 checksum is calculated against which the written tar is checked. The default ARCHIVE is scheduled.0 and default the FILE is ./smeserver.tgz. The checksum is written to FILE.md5sum.

A smeserver.tgz file stored on an USB disk can be used for restore during a SME Server fresh install.

```
affa --list-archives [--csv] JOB
```

Displays a table of all present archives of job JOB with date, number of files, size and and bytes received.

```
Affa version 2.0.0-rc4 on affa1.mydomain.de (10.204.48.2)
-----+-----
| Job: primmail |
| Description: Mailserver Albstadt |
| Directory: /var/affa/primmail/ |
| Hostname: 10.204.144.2 |
| Email: admin@mydomain.de |
-----+-----
| Run | Completion date | Exec Time | Files | Size | Recvd |
-----+-----+-----+-----+-----+-----
| M 9 | Sun 2008 Mar 02 06:11 | 0h29m47s | 679010 | 80GB | 60MB |
| M 8 | Sun 2008 Mar 30 06:12 | 0h26m59s | 701683 | 83GB | 48MB |
| M 7 | Sun 2008 Apr 27 06:16 | 0h28m01s | 731332 | 87GB | 47MB |
| M 6 | Sun 2008 Jun 01 06:20 | 0h27m37s | 755529 | 90GB | 50MB |
| M 5 | Sun 2008 Jun 29 06:24 | 0h28m39s | 800200 | 94GB | 52MB |
| M 4 | Sun 2008 Jul 27 06:28 | 0h29m10s | 835398 | 99GB | 59MB |
| M 3 | Sun 2008 Aug 31 06:34 | 0h34m53s | 869409 | 101GB | 65MB |
| M 2 | Sun 2008 Sep 28 06:38 | 0h33m19s | 910889 | 106GB | 59MB |
| M 1 | Sun 2008 Oct 26 06:42 | 0h38m41s | 932627 | 110GB | 61MB |
| M 0 | Sun 2008 Nov 30 06:48 | 0h51m44s | 954090 | 116GB | 75MB |
-----+-----+-----+-----+-----+-----
| W 3 | Sun 2008 Dec 14 06:50 | 0h53m34s | 962267 | 118GB | 68MB |
| W 2 | Sun 2008 Dec 21 06:52 | 0h53m11s | 974914 | 120GB | 69MB |
| W 1 | Wed 2008 Dec 24 06:52 | 0h48m58s | 974536 | 120GB | 64MB |
| W 0 | Sun 2008 Dec 28 06:58 | 0h39m01s | 977003 | 120GB | 61MB |
-----+-----+-----+-----+-----+-----
| D 6 | Tue 2008 Dec 30 06:52 | 0h54m45s | 976872 | 120GB | 77MB |
| D 5 | Wed 2008 Dec 31 06:52 | 0h49m10s | 977764 | 120GB | 63MB |
| D 4 | Thu 2009 Jan 01 06:28 | 0h28m09s | 977879 | 120GB | 60MB |
| D 3 | Fri 2009 Jan 02 06:53 | 0h53m21s | 978144 | 120GB | 68MB |
| D 2 | Sat 2009 Jan 03 06:53 | 0h53m15s | 978072 | 120GB | 60MB |
| D 1 | Sun 2009 Jan 04 06:53 | 0h53m16s | 976478 | 120GB | 70MB |
| D 0 | Mon 2009 Jan 05 06:52 | 0h52m42s | 977062 | 120GB | 61MB |
-----+-----+-----+-----+-----+-----
| S 1 | Mon 2009 Jan 05 17:57 | 0h27m48s | 976353 | 120GB | 98MB |
| S 0 | Tue 2009 Jan 06 06:52 | 0h52m55s | 976539 | 120GB | 70MB |
-----+-----+-----+-----+-----+-----
```

With --csv, the output is in machine readable colon separated format.

```
affa --status [--csv]
```

Displays a table of all configured jobs with enable status, time of last and next run, size of the most recent archive, execution time of the last run and the number of scheduled (S), daily (D), weekly (W), monthly (M) and yearly (Y) archives. Last time shows 'failed', if a job did not run in the last 24h. For

disabled jobs 'Last' always shows 'failed' after 24 h. To see the date and time of the last run of those jobs use the `--list-archives` option. Column 'Next' shows the time when the next run will be started, if 'Enabled' is 'yes'.

```
Affa version 2.0.0-rc4 on backup.mydomain.de (10.204.48.2)
```

Job	ENA	Last	Exec Time	Next	Size	N of S,D,W,M,Y
erp-alb-rsyncd	yes	00:20	0h00m14s	00:20	46MB	1, 7, 4, 10, 0
erp-fra-rsyncd	yes	02:46	0h01m03s	02:45	712MB	1, 7, 4, 10, 0
esxi-TS-W2K	yes	00:37	4h22m08s	20:15	60GB	1, 7, 1, 0, 0
esxi-W2KR2	yes	02:22	3h52m09s	22:30	40GB	1, 7, 1, 0, 0
helpdesk	yes	00:25	0h00m47s	00:25	117MB	1, 7, 4, 10, 0
imageserv	yes	running (pid 16477)				1, 7, 4, 10, 0
intraweb	yes	06:34	0h04m50s	18:00	2.1GB	2, 7, 4, 10, 0
pdcalb	yes	21:19	0h14m27s	21:05	89GB	1, 7, 4, 10, 0
pdcfra	yes	04:18	0h03m33s	04:15	33GB	1, 7, 4, 10, 0
primmail	yes	06:52	0h52m55s	17:30	120GB	2, 7, 4, 10, 0
rayofhope	yes	21:52	0h07m40s	21:45	26GB	1, 7, 4, 10, 0
smecrmpg	yes	22:35	0h00m21s	22:35	70MB	1, 7, 4, 10, 0
sozserv	yes	02:09	0h04m03s	02:05	5.9GB	1, 7, 4, 10, 0
wiki	yes	03:45	0h00m27s	03:45	278MB	1, 7, 4, 8, 0
az32share	no	-	0h00m22s	03:05	1.3GB	1, 7, 4, 8, 0

With `--csv`, the output is printed in a machine readable colon separated format.

```
affa --disk-usage
```

Shows the current disk usage of all root dir filesystems

```
Affa version 2.0.0-rc4 on backup2.mydomain.de (10.204.0.52)
```

Use%	Used	Avail	Root Dir
71%	938GB	365GB	/var/affa
20%	194GB	759GB	/mnt/affadev

With `--csv`, the output is printed in a machine readable colon separated format.

```
affa --show-schedule [--all]
```

Prints a 'graphical' timetable for all enabled jobs. The resolution is 30 minutes. An 'X' character marks the scheduled start times. The last performance duration is marked with '=' characters.

```

Affa version 2.0.0-rc4 on backup.mydomain.de (10.204.48.2)
      TIME 0:00   4:00   8:00  12:00  16:00  20:00
erp-alb-rsyncd X-----
      helpdesk X-----
      esxi-W2KR2 =====
      esxi-TS-W2K =====X==
      imageserv ---X---
      sozserv ---X---
erp-fra-rsyncd ----X--
      wiki -----X
      pdcfra -----X
      primmail -----X
      intraweb -----X
      pdcalb -----X
      rayofhope -----X
      smecrmpg -----X
1 disabled jobs not listed. Use --all to display.

```

```

affa --send-status

```

Sends the status table, the disk-usage and the archive list of all jobs to the email addresses configured in the 'DefaultAffaConfig' record. Used by the cronjob 'affa-status'.

```

affa --mailtest JOB

```

Sends a test email to the email addresses configured in the JOB record. With property Watchdog=yes, a test email is sent from the remote host, too. Use this to verify, that your mail processing is functional. **Note:** By default Affa only sends messages on errors, never on success (see property chattyOnSuccess).

```

affa --cleanup JOB

```

After you have lowered a keep value, e.g. scheduledKeep, then archives with a higher indices will no longer be shifted and will exist for ever. This option finds these archives and deletes them.

```

affa --rename-job JOB NEWNAME

```

Renames the job JOB to NEWNAME including all database records and archive directories.

```

affa --move-archive JOB NEWROOTDIR

```

Moves the archive directory of job JOB to the rootdir NEWROOTDIR and adjusts the property RootDir. NEWROOTDIR must be a full path starting with a slash. As moving across filesystems (e.g.

from an external USB drive to the local disk) is not possible, Affa uses a copy command in this case and deletes the source directory after that. Depending on the archive size, copying across filesystems can take a long time.

```
affa --delete-job [--revoke-key] JOB
```

Irreversibly deletes a job including all archives, configuration and report databases. With given --revoke-key option, the public key on the remote server will be deleted.

```
affa --revoke-key JOB
affa --revoke-key --host=TARGETHOST [--port=PORT] [--remoteOS=cygwin]
```

Deletes the public dsa key on the remote server.



Note:

By default, the --send-key option works for a SME Server as a remote server and for systems where the keys are stored in /root/.ssh/authorized_keys2. With remoteOS=cygwin it works for a Cygwin/Windows remote server.

```
affa --check-connections
```

Checks the ssh login for all configured jobs. For jobs where the public key was not yet sent, you are prompted for the password and the key will be sent then.

```
affa --kill JOB
```

Terminates the running job JOB and all its child processes (rsync processes).

```
affa --debug
```

Enables verbose logging. Overrides job and global configurations.

Example setups

Dedicated backup server

- Set up a job for every server you want to backup
- Send the public keys to every server

```
affa --send-key JOBNAME
```

- Check whether password-less logins are working

```
affa --check-connections
```

- Check whether the scheduled jobs are evenly distributed over the day

```
affa --show-schedule
```

- Create the cron jobs

```
affa --make-cronjobs
```

- Check the status after 24 hours

```
affa --status
```

Backup of single ibays

Suppose you want to backup the ibays 'staff1' and 'staff2' on your production server with WAN IP 82.123.1.1 to an Affa server in a different location over the internet every night at 2:30am.

- log into the Affa server and install the packages as described above.
- copy the config helper script sample

```
cp /usr/lib/affa/jobconfig-sample.pl /root/ibay-staff-job.pl
```

- edit /root/ibay-staff-job.pl and set

```
my $jobname='ibay-staff';
```

and

```
'remoteHostName'=>'82.123.1.1',
'TimeSchedule'=>'0230',
'SMEServer'=>'no',
'Include[0]'=>'/home/e-smith/files/ibays/staff1',
'Include[1]'=>'/home/e-smith/files/ibays/staff2',
```

- write the configuration

```
/root/ibay-staff-job.pl
```

- send the public key to the production server

```
affa --send-key ibay-staff
```

- check next morning

```
affa --list-archives ibay-staff
affa --status
ls /var/affa/ibay-staff
```

Two production servers backup each other

You have two sites connected via a VPN and a SME Server running on each site. In that case you don't need a dedicated Affa backup server. Both production servers can additionally act as Affa backup servers backing up the server of the other site. Simply install Affa and configure a job that backs up the other one. You can use all Affa features except of the rise feature. When using the rise feature the server become any of the backed up systems, which is less useful in this scenario as it would give you a running copy of the server of the other site while the server of this site is down.

To get redundancy and a faster restore you can configure a local backup to an external USB or NAS device.

Please pay attention, that you do not backup the archives back to the other site. Set the RootDir property to a path which is not included in the SME default backup list. When leaving the RootDir property to its default '/var/affa', this is guaranteed.

Backing up a Windows computer

Backing up data from a Windows system requires the Cygwin Rsyncd daemon installed and configured on Windows. The standard procedure *rsync over ssh* does not work, as the Cygwin rsync process always hangs after some files were transferred.

Rsyncd setup on the Windows computer

Install the Cygwin base, the rsync package and configure the Rsyncd service as described in this document: [Rsyncd setup on a windows computer for use with Affa backup](#)'. The installation of the sshd service is optional and not needed for the backup itself, but having a ssh login can be very helpful for administration or executing scripts on the Windows system. Affa supports sending the public key to a Windows Cygwin for password-less login.

Note: Affa does not backup the Windows Access Control List (ACL) information. You may need to correct the ACLs manually after a restore.

Affa Rsyncd mode setup (Quick start example)

You want to backup the *My Documents* folders of the users *ssorglos* and *bsimpson* from the Windows computer 'ws001' with IP 192.168.1.65 to your Affa server 'affabox' with IP 192.168.1.3.

- log into the 'affabox' and copy the Cygwin config helper script sample

```
cp /usr/lib/affa/jobconfig-cygwin-sample.pl /root/ws001-mydocs-job.pl
```

- edit `/root/ws001-mydocs-job.pl` and set

```
my $jobname='ws001-mydocs';
```

and

```
'remoteHostName'=>'192.168.1.65',
'rsyncdPassword'=>'secretword',
'Include[0]'=>'/c/Documents and Settings/ssorglos/My Documents/', # don't use backslashes in pathnames
'Include[1]'=>'/c/Documents and Settings/bsimpson/My Documents/',
```

where *secretword* must be replaced by the password you have chosen in the `rsyncd.secretsfile` on the Windows box.

- write the configuration (this makes the database entries and sets up the cronjobs)

```
/root/ws001-mydocs-job.pl
```

- run the job manually. After completion check the archive `/var/affa/ws001-mydocs/scheduled.0` and the logfile `/var/log/affa/ws001-mydocs.log`

```
affa --run ws001-mydocs
```

Affa Rsyncd mode manual setup

In case you want to do the setup manually using the `db` command, these are the mandatory settings for Cygwin Rsyncd mode

```
db affa setprop JOB rsyncdMode yes
db affa setprop JOB rsyncdModule AFFA
db affa setprop JOB rsyncdUser affa
db affa setprop JOB rsyncdPassword secretword
db affa setprop JOB SMEServer no
db affa setprop JOB RPMCheck no
db affa setprop JOB Watchdog no
```

Optional for ssh login

```
db affa setprop JOB remoteOS cygwin
```

Use Affa to backup to a NFS-mounted NAS or a local attached USB drive

You want to backup your SME 7 production server with hostname 'prodbox' and IP 10.200.48.1 on a mounted filesystem instead of setting up a dedicated Affa box.

Setup NAS

You have a FreeNAS (<http://www.freenas.org>) box with IP 10.200.48.2 up and running with NFS

service enabled. The disk is mounted to /mnt/affashare. You have authorized the IP address of your prodbox server to access share /mnt/affashare.

- log into the 'prodbox' and install the NFS packages

```
/usr/bin/yum install --enablerepo=smecontribs smeserver-nfs
```

- now enable and start the portmapper service

```
config setprop portmap status enabled
service portmap start
```

- mount the NFS share

```
mkdir -p /mnt/affadevice
mount 10.200.48.2:/mnt/affashare /mnt/affadevice
```

Alternatively setup a USB drive

- log into the 'prodbox'
- connect a USB hard disk to the USB Bus. Now you must determine what device the kernel has assigned to the drive. View the /var/log/message and search for *Initializing USB Mass Storage driver*. A few lines below you'll find the name of the device. In this example it is *sdh*. Replace /dev/sdh by *your device* in following instructions.
- use the fdisk program to create a linux partition. **Verify that this is really the attached USB drive before you continue!**

```
fdisk /dev/sdh
```

You'll most likely find an existing vfat dos partition, which you have to delete first. In the following we assume, that you have created a single partition /dev/sdh1.

- now format the drive with an ext3 filesystem

```
mkfs.ext3 /dev/sdh1
```

- make the mount point

```
mkdir -p /mnt/affadevice
```

- add the following line to the /etc/fstab

```
/dev/sdh1 /mnt/affadevice ext3 defaults
```

- mount the drive

```
mount /mnt/affadevice
```

- crosscheck your work using the df command

```
df
```

Setup Affa

You want to run backups on 11:30 h, 15:30 h and 19:30 h and you want to keep the last 3 scheduled backups, 7 daily, 5 weekly, 12 monthly and 1 yearly backups.

- log into the 'prodbox' and install the Affa packages as described above.
- copy the config helper script sample

```
cp /usr/lib/affa/jobconfig-sample.pl /root/prodbox-job.pl
```

- edit /root/prodbox-job.pl and set

```
my $jobname='prodbox';
```

and

```
'remoteHostName'=>'localhost',
'TimeSchedule'=>'1130,1530,1930',
'scheduledKeep'=>3,
'dailyKeep'=>7,
'weeklyKeep'=>5,
'monthlyKeep'=>12,
'yearlyKeep'=>1,
'RootDir'=>'/mnt/affadevice',
```

Review the other properties and change them to your needs.

- write the configuration

```
/root/prodbox-job.pl
```

- run the job manually

```
affa --run prodbox
```

Limitations

With this kind of setup you cannot use the affa rise feature, as it requires the backup archive to be located on the **same filesystem** as the server installation. The rise option uses hardlinks, which are not

working across filesystems.

Automount

Having the backup archives in the same filesystem is always a risk, which can be minimized by using the automount feature. Then the external filesystem is only mounted during a job run.

In the NAS example set

```
'AutomountDevice=>'10.200.48.2:/mnt/affashare',
'AutomountPoint =>'mnt/affadevice',
```

and skip the step 2.

In the USB drive example set

```
'AutomountDevice=>'/dev/sdc1',
'AutomountPoint =>'mnt/affadevice',
```

and skip the steps 5 to 8.

The mount point will be automatically created, if it does not exist. To access the archive directory, you need to mount it manually.

Restore

Restore single files or directories

Example 1: It's Tuesday January 6th 2009, when user 'briedlin' asks you to restore the messages of his mailbox 'orders' he has accidentally deleted on Monday.

- You first must check what backup archives are available. The jobname of this server backup is 'primmail'. To get a listing of all archives run

```
affa --list-archives primmail
```

(see the example listing in chapter Affa#Usage_and_command_line_options)

- Choose the daily.0 archive, which was created Monday night. Now restore the mailbox 'orders' using the rsync command.
- Now run the rsync command *(note the trailing slash!)* on the Affa backup server:

```
export RDIR=/home/e-smith/files/users/briedlin/Maildir/.orders/ # this variable is used to shorten
rsync -av /var/affa/primmail/daily.0/$RDIR 10.204.48.1:$RDIR
```

If the servers are configured to use a different ssh port eg 2222, then instead do:

```
export RDIR=/home/e-smith/files/users/briedlin/Maildir/.orders/ # this variable is used to shorten
rsync -av -e 'ssh -p 2222' /var/affa/primmail/daily.0/$RDIR 10.204.48.1:$RDIR
```

Example 2: A user has deleted the file `orderform.pdf` from ibay 'docs' on the server 10.204.48.1 and asks you to restore it.

- You have searched and found the latest version of this file in weekly archive of job 'prodserv'.
- To copy it back to the server 10.204.48.1 run on the Affa server

```
export RFILE=/home/e-smith/files/ibays/docs/files/orderform.pdf # this variable is used to shorten
rsync -av /var/affa/prodserv/weekly.1/$RFILE 10.204.48.1:$RFILE
```

- If the servers are configured to use a different ssh port eg 2222, then instead do:

```
export RFILE=/home/e-smith/files/ibays/docs/files/orderform.pdf # this variable is used to shorten
rsync -av -e 'ssh -p 2222' /var/affa/prodserv/weekly.1/$RFILE 10.204.48.1:$RFILE
```

Example 3: Restoring a file with special characters in the path

Do not use shell variables to shorten the command. It would complicate things more than it would help. Quote the source and destination path. In the destination path escape blank characters with a double backslash and brackets with a single backslash. On the Affa backup server do:

```
rsync -av "/var/affa/fileshare/daily.4/home/e-smith/files/ibays/mechfiles/files/Valve Control (Des
```

If the servers are configured to use a different ssh port eg 2222, then instead do:

```
rsync -av -e 'ssh -p 2222' "/var/affa/fileshare/daily.4/home/e-smith/files/ibays/mechfiles/files/Va
```

Full restore

To run a full restore of user and configuration data run on the Affa server

```
affa --full-restore <JOB> [<ARCHIVE>]
```

This rsyncs the data from the backup ARCHIVE back to the 'remoteHostname' defined in the configuration of the job JOB.

Example: You have backedup your production server 'prodsrv' as job 'prodbox'. To restore from the latest backup run

```
affa --full-restore prodbox
```

To restore from the older archive daily.3 run

```
affa --full-restore prodbox daily.3
```


**Warning:**

A full restore reconstructs the server as it was at the time of the backup. That means, that all files created or server configuration changes made after the backup will be lost. After the restore is done, the restored server reboots automatically.

Moving a SME 7 server installation to new hardware using the Affa rise feature

Please see this Howto: [Moving SME to new Hardware](#)

Restore from USB drive on new server

This tip comes from <http://forums.contribs.org/index.php?topic=42412.0>

Q) I have complete backups using affa stored on a usb hard drive connected to our affa backup server. I need to restore an earlier monthly backup of our job "mailbackup" to a test server rather than back to the original system. If I did it from the backup server I see the instructions of how to rise that server to the current backup on that server but I want to restore a point about a month ago before some strange things happened. And I want to do it on a machine that is not our backup server or our production server. I tried to figure out how but am lost in the options. My goal is to do some testing.

A) On your testserver setup a job mailbackup identical to that on your backup server but set property RootDir to /var/affa and property status to disabled. Connect the USB drive and copy the archive of the job mailbackup to /var/affa. Then run `affa --rise mailbackup ARCHIVE`

Remember that Affa places job setup scripts into the archive directories. In your case it is mailbackup-setup.pl. Copy it to your testserver, change the RootDir and status properties and execute it.

FAQ**What files and directories are included by default?**

With SMEServer=no nothing at all.

With SMEServer=yes the SME default backup list. The following command displays the list:

```
perl -e 'use esmith::Backup;$b=new esmith::Backup;print join("\n",$b->restore_list)."\n"'
```

Can I exclusively backup image files from a specific directory?

Yes. Assuming you want to backup all gif and jpg files from directory /home/e-smith/files/ibays/pictures/files use this configuration

```
db affa setprop JOBNAME SMEServer no # do not include SME Server default directories and
db affa setprop JOBNAME Include[0] '/home/e-smith/files/ibays/pictures/files' # start searching here
db affa setprop JOBNAME Include[1] '*' # search the tree recursively
db affa setprop JOBNAME Include[2] '*.gif' # copy files that match these patterns
db affa setprop JOBNAME Include[3] '*.jpg'
db affa setprop JOBNAME Exclude[0] '*' # exclude all others
db affa setprop JOBNAME rsyncOptions '-m' # do not create empty folders
```

How can I move a backup archive from my production server in a remote branch office to the Affa server using an USB disk? I want to avoid to download the first full backup over the internet. Install Affa on your production server, connect and mount the USB disk. Setup a job to write the archive to the USB disk. After the job run, connect the USB disk to your Affa server and setup the job in the same way. Use the setup script from the archive directory. You only need to change the remoteHostName property. Now use the --move-archive option to move the archive to the local disk. If not already done, send the public key to your production server. You're done.

How do I backup two SME servers behind a firewall? First you need to configure port forwardings for the ssh service on your firewall. Use a non-standard port, e.g 2200->22, for your second server. Setup a job on your Affa server for each of your production servers. Set up the job property sshPort=2200 for second server.

Uninstall

This removes the installed Affa package, all configuration data and all backup archives.

Run the following commands for each job:

```
affa --revoke-key JOBNAME
affa --delete-job JOBNAME
```

Verify that all jobs have been deleted

```
affa --status
```

Remove the Affa package

```
rpm -e smeserver-affa
```

and, if these packages are not needed by any other packages:

```
rpm -e perl-Filesys-DiskFree perl-Compress-Bzip2
```

Cleaning up

```
rm -f /etc/cron.d/affa /etc/cron.d/affa-status
rm -rf /home/e-smith/db/affa /home/e-smith/db/affa-report
rm -rf /var/affa /var/log/affa
```

Bug report

Affa is listed in the bugtracker contribs section (http://bugs.contribs.org/enter_bug.cgi?product=SME%20Contribs&component=smeserver-ffa) . Please report all bugs, new feature requests and documentation issues there.

Additional information

For details of performance, changelog etc see [Affa:Additional information](#)

References

- Wikipedia article <http://en.wikipedia.org/wiki/Rsync>
- Rsync homepage <http://rsync.samba.org/>
- Wiki article [Moving SME to new Hardware](#)
- Wiki article [Backup with Affa and FreeDup](#)
- Wiki article [Rsyncd setup on a windows computer for use with Affa backup](#)
- Wiki article [Backup Windows Computers Using Affa](#)
- Wiki article [Backup of ESXi Virtual Machines using Affa](#)

Retrieved from "<http://wiki.contribs.org/Affa>"

Categories: [Advanced](#) | [Contrib](#) | [Backup](#)

- This page was last modified on 25 January 2009, at 15:43.
- Content is available under [GNU Free Documentation License 1.2](#).