Grazer LINUXTAGE09
25. APRIL 2009 FH JOANNEUM GRAZ

Backup mit

The
Advanced
Maryland
Automated
Network
Disk
Archiver

DI Gernot Pirnbacher

www.it-to-fit.at

contents

- About me
- Features
- Architecture
- Config in Detail (and connected commands)
- Tapelist, Labels and Disklist {DLE}
- running it
- Restore / Recover
- Restore without Amanda Index or Tapeserver
- Advanced Usage – MySQL or VMWare Backup {LVM2}

About me

DI Gernot Pirnbacher

Bauingenieur and LinuxUser (LPI 1 done, LPI 2 underway)

Timberengineer & IT Fanatic

currently servicing:

- a 40 node mixed network (and some smaller ones)
- a micro-mpi-cluster with 16 CPUs
- several smaller numbercrunchers, vmware-servers and firewalls
- some webservers, my mailserver and some hosting {svn,darcs,web&mail}
- #might be able to provide vps-services to opensource projects in need
- ***pdf & scripts & contact: http://www.it-to-fit.at/doku.php/mywiki:amanda***

```
-----BEGIN GEEK CODE BLOCK-----
Version: 3.1
GE/IT d- s:- !a C++$ UL+++$ !P L+++$ E--- W++ N- !o? K--? w+ !O !M !V PS@
PE !Y? PGP+ t+ !5 !X !R tv b+++ DI++ D++ G e+++ h---- r+++ y?
------END GEEK CODE BLOCK------
```

Features – what can *amanda* do?

- LAN based Multi-Host Backups (Linux, BSD, Mac, Win)
- Using native tools to do the job (tar, dump, smbfs)
- Disk2Tape; Disk2Disk(2Tape); Disk2Cloud
- Automated planner for consistent space requirements
- Parallel dumps to holdingdisk then streaming to tape(s)
- client/server compression, priorities, dedicated interf.
- File by File recovery over Network
- ”battle tested” Software (16yrs)
- Authentication ”bsd” style or ssh-auth(+transfer)

Architecture

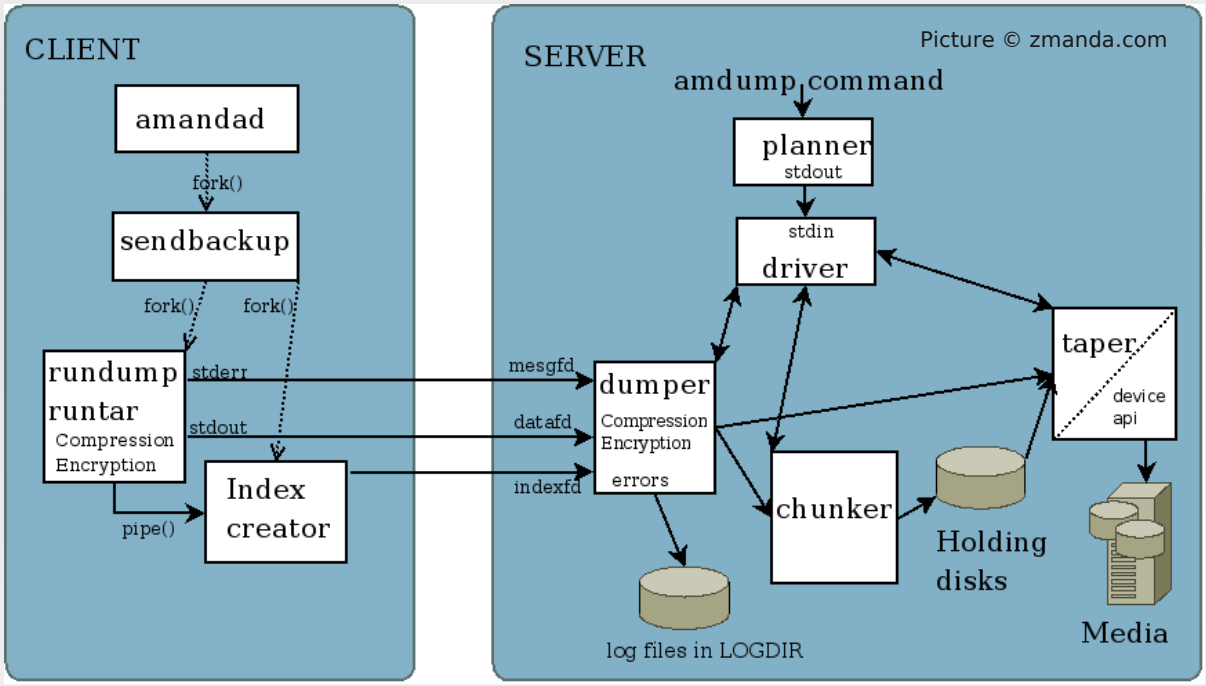

- Client – Server Architecture
  - Compression can be performed on both sides
  - Index is created on client and transferred to server
  - ”chunker”, queing, streaming to tape on server
  - Planning (full or incremental dump) on server

amanda.conf

- **Configuration of Amanda's Features in**

  /etc/amanda/[*NameOfSet*]/amanda.conf

- **Configuration of DLE's in *..AsAbove..*/disklist**

  here the DiskListEntries – the volumes or partitions to be backed up - are configured

- **Configuration of used tapes in *..AsAbove..*/tapelist**

  managed by amanda or via `amrmtape`
  generally only edited because of defective tapes

amanda.conf in detail

Basic overview of absolutely necessary parameters:

- `org` Name of your BackupSet
- `mailto` list of operators for reports
- `dumpuser` user to run dumps
- `dumpcycle` days or weeks in planned cycle
- `runspercycle` amdumpruns inside the cycle
- `tapecycle` number of tapes in rotation
- `runtapes` number of tapes per run
- `tapedev` non-rewinding tape device (/dev/nst0)
- `labelstr` regex for allowed labels (enforced)

amanda.conf in detail

- `tapetype` used tape (size,speed,length,etc)

```
define tapetype HP-DAT-160GB {
    comment "just produced by tapetype prog (hardware compression off)"
    # fix blocksize problems i.e. of HP DAT 72 and 160
    readblocksize 64 kbytes # USB-Dat drives fail with default of 256kb {amlabel}
    length 77208 mbytes
    filemark 0 kbytes
    speed 6239 kps
}
define tapetype DVD47 {
    comment "A 4.7 Gbyte DVD-sized vtape"
    length 4482 mbytes
    filemark 4 kbytes
}

#special tapedev notation for B2D Usage:
#tapedev "file:/space/amanda-vtapes/tape"
#http://wiki.zmanda.com/index.php/How_To:Build_a_Test_Environment_With_Virtual_Tapes
#i.e. coupled with DVD-Sized vtapes and "chg-disk" as tapechanger – see link
#or with a tapetype containing all diskspace on that volume
#these directory are treated by amanda like tapes – including labeling and use of "mt"
#i.e. rewind and label a vtape
#$ ammt -t file:/space/amanda-vtapes/tape rewind
#$ amlabel test Test-01
```

amanda.conf in detail

- `holdingdisk` diskspace used to cache dumps

```
holdingdisk hd1 {
    comment "main holding disk"
    directory "/mnt/holdingdisk/hd1"    # where the holding disk is
    use -20000 Mb                       # how much space can we use on it
                                        # a non-positive value means:
                                        # use all space but that value

    chunksize 3Gb
}
```

- `dumptypes` parameters on how to dump a DLE

```
define dumptype root-tar {
    global
    program "GNUTAR"
    comment "root partitions dumped with tar"
    compress none
    index
#   exclude list "/etc/amanda/exclude.gtar"
    priority low
}
define dumptype comp-root-tar {
    root-tar
    comment "Root partitions with compression"
    compress client fast
}
```

tapetype, tapelist and labels

- each tape is labeled at the start with "*labelstr*"
  - Amanda tracks used tapes
  - defective tapes should be removed with `amrmtape`
  - Amanda does not overwrite active tapes (checked via the *tapelist* file)
  - `amtapetype`: generates the tapetype entry

```
#this command should be executed with hardware compression off
#can be set i.e. with mt-st
mt-st -f /dev/nst0 compression 0
#amtapetype call with estimate of 80GB and tee to store the stdout for amanda.conf
amtapetype -f /dev/nst0 -e 80G| tee ./amtapetype.out
```

disklist and DLEs {DiskListEntries}

- DLE's define the volumes / folder to be backed up

```
hostname diskname [diskdevice] dumptype [spindle [interface] ]
```

- `hostname` alias or fqdn of host
- `diskname` name of the disk to be dumped
- `diskdevice` default = diskname
  may be: full /dev/sda; only sda
  or a mountpoint like /home/user1
- `dumptype` dumptype to be used (can be inline)
- `spindle` am will only run 1 backup per ”spindle” - used for balancing
- `interface` network interface to use

disklist and DLEs {DiskListEntries}

- multiple DLEs per mountpoint can be defined by using distinct names and include/exclude statements

```
# I suggest to always use the diskdevice in the diskname.
# Don't forget to use the same spindle for all entries.
#hosta /diskA/all /diskA {
#       # all directories except the one that start with [a-u]
#       high-tar
#       exclude "./[a-u]*"
#       } 1
#hosta /diskA/ag /diskA {
#       # all directories that start with [a-g] except big1 and big2
#       high-tar
#       include "./[a-g]*"
#       exclude "./big1" "./big2"
#       } 1
#hosta /diskA/big /diskA {
#       # directories big1 and big2
#       high-tar
#       include "./big1" "./big2"
#       } 1
#hosta /diskA/gm /diskA {
#       # all directories that start with [h-u]
#       high-tar
#       include "./[h-u]*"
#       } 1
```

DLE for a samba share

- a DLE-Entry with special notation triggers the usage of samba to perform the dump for this DLE

```
# DLE syntax to perform smbclient based dumps
#[linux-client with smbclient] //host/share [dumtype using TAR]
xpc //webmaster/EigeneDateien user-tar
```

- Credentials have to be provided in `/etc/amandapass`
- a native Windows client is available too

/etc/amandahosts & ssh-auth

- `/etc/amandahosts` ist the bsd authentication file must be present on server an client

```
# on the server; enables amrestore access
localhost backup amdump amindexd amidxtaped
xpc.zuhause.local backup amdump amindexd amidxtaped
localhost root amdump amindexd amidxtaped
xpc.zuhause.local root amdump amindexd amidxtaped
KingKong.zuhause.local backup amdump amindexd amidxtaped
KingKong.zuhause.local root amdump amindexd amidxtaped
# on the client only dumps are allowed
xpc.zuhause.local backup amdump
xpc backup amdump
```

- ssh based authentication – keys without password

```
# amanda.conf entries:
auth "ssh"
ssh_keys "/home/amandabackup/.ssh/id_rsa_amdump"
# if not the default user (or other name)
#client_username "otherusername"
# maybe add in authorized_keys to increase sercurity:
from="amanda_server.your.domain.com",no-port-forwarding,no-X11-forwarding,\
      no-agent-forwarding,command="/absolute/path/to/amandad -auth=ssh amdump"
```

/etc/amandapass and /etc/hosts

- `/etc/amandapass` stores credentials for samba DLEs

```
# Syntax is samba-share user%pass
//webmaster/EigeneDateien  webdesigner%foobar
```

- `/etc/hosts`
  has to contain all nodes (client&server) participating
  dnslookups do not suffice in my experience

- usually via a cronjob; ie. like this

```
#!/bin/bash
# amanda cron script by Gernot Pirbacher
# including alternate B2D prestep via rsync; then backups to tape via amanda
....
    cut to display properly ...
....
echo "starte amcheck ..." >> /var/log/backup/daily.run
su backup -c "/usr/sbin/amcheck std >> /var/log/backup/daily.run"
echo "================================================================================"
echo "starte amdump ..." >> /var/log/backup/daily.run
su backup -c "/usr/sbin/amdump std >> /var/log/backup/daily.run"
echo "================================================================================"

echo "Backuplauf vollständig!" >> /var/log/backup/daily.run

....
    cut to display properly ...
....
echo "mail erinnerung an die backup admin(s) versenden ... ">> /var/log/backup/daily.run
mail  user@domain.com, admin@domain.com -s "TAPECHANGE REMINDER! Please ... < mail-text
echo "================================================================================"
mt -f /dev/nst0 offline
```

- `amcheck [`*`NameOfSet`*`]` tests config and all hosts

```
xpc:/etc/amanda/lt-demo# su backup -c '/usr/sbin/amcheck lt-demo'
Amanda Tape Server Host Check
-----------------------------
Holding disk /mnt/holdingdisk/hd1: 305110 MB disk space available, using 285110 MB
read label `lt-tape-001', date `X'
NOTE: skipping tape-writable test
Tape lt-tape-001 label ok
# there might be notes like this:
#NOTE: conf info dir /mnt/holdingdisk/amanda/lt-demo/curinfo does not exist
#NOTE: it will be created on the next run.
#NOTE: index dir /mnt/holdingdisk/amanda/lt-demo/index does not exist
#NOTE: it will be created on the next run.
Server check took 0.013 seconds

Amanda Backup Client Hosts Check
--------------------------------
Client check: 2 hosts checked in 0.033 seconds, 0 problems found

(brought to you by Amanda 2.5.2p1)
```

running it

- `amdump` does not print to stdout
- `amstatus [`*`NameOfSet`*`]` is used to track progress

```
xpc:/mnt/holdingdisk/amanda# !amsta
amstatus lt-demo
Using /mnt/holdingdisk/amanda/lt-demo/amdump from Sat Apr 18 01:08:16 CEST 2009

xpc:xpc-homes        0       209m finished (1:09:17)
xw8200:xw8200-homes  0      2512m writing to tape (1:10:55)

SUMMARY          part      real  estimated
                           size       size
partition       :   2
estimated       :   2                2722m
flush           :   0         0m
failed          :   0                   0m           (  0.00%)
wait for dumping:   0                   0m           (  0.00%)
dumping to tape :   0                   0m           (  0.00%)
dumping         :   0         0m        0m (  0.00%) (  0.00%)
dumped          :   2      2722m     2722m (100.00%) (100.00%)
wait for writing:   0         0m        0m (  0.00%) (  0.00%)
wait to flush   :   0         0m        0m (100.00%) (  0.00%)
writing to tape :   1      2512m     2512m (100.00%) ( 92.29%)
failed to tape  :   0         0m        0m (  0.00%) (  0.00%)
taped           :   1       209m      209m (100.00%) (  7.71%)
  tape 1        :   1       209m      209m (  0.27%) lt-tape-001
```

reports generated by amdump

- `amdump` sends a mail report after finishing the current dump
- this report includes:
  tapes (used&next)
  statistics
  tape usage
  planner infos
- if no or a wrong tape is inserted all dumps are kept in the *hd* - with `amflush` these can be flushed to tape

```
REPORT
...cut...

Hostname: xpc
Org     : kell.hound's lt demo amanda setup
Config  : lt-demo
Date    : April 18, 2009

These dumps were to tape lt-tape-001.
The next tape Amanda expects to use is: a new tape.

STATISTICS:
                          Total       Full      Incr.
                        --------   --------   --------
Estimate Time (hrs:min)     0:00
Run Time (hrs:min)          0:09
Dump Time (hrs:min)         0:03       0:03       0:00
Output Size (meg)         2722.8     2722.8        0.0
...cut...
USAGE BY TAPE:
  Label                Time      Size      %    Nb    Nc
  lt-tape-001          0:07     2723M    3.5     2     0

NOTES:
  planner: Adding new disk xpc:xpc-homes.
  planner: Adding new disk xw8200:xw8200-homes.
  taper: tape lt-tape-001 kb 2788224 fm 2 [OK]

...cut...
```

- amrecover is able to restore single files using the index

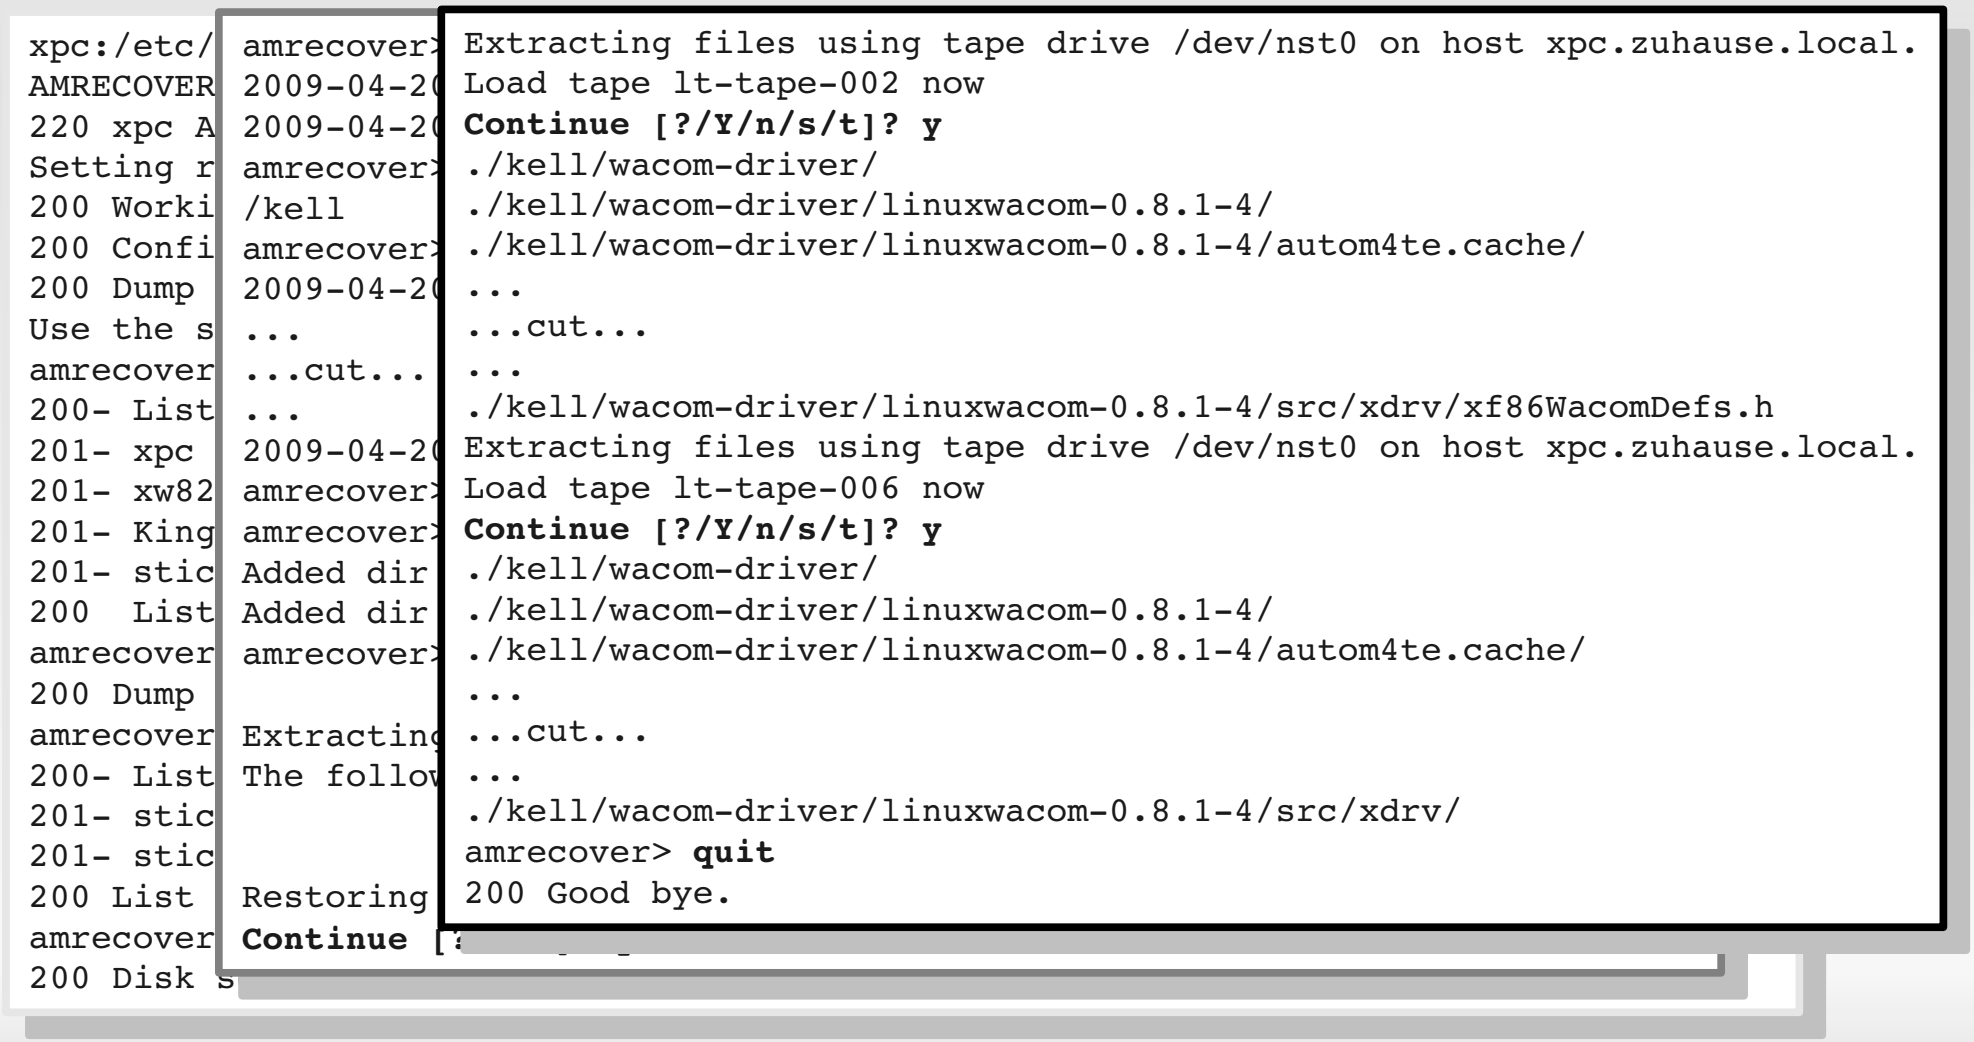

amrestore

- use `amrestore` to recover complete dumps

```
amrestore -p -f XX /dev/nst0 | pipe through tar, restore, nc , etc. ...
```

- and find correct tapes with `amadmin CurrentSet info`:

```
#!/bin/bash
# mailing nicely formatted amadmin info (toc of tapes)
# alernatively use amadmin DailySet export for full dump (can be imported)
su backup -c "/usr/sbin/amadmin std info" > /etc/scripts/amadmin.info.last
mail -s "current amadmin tape-infos (worst case recovery) for `date "+%\
Y-%m-%d_%H:%M"`" user@domain.com < /etc/scripts/amadmin.info.last

this gives a toc per tape like this:
Current info for lignum-server.lignum.tugraz.local /mnt/samba/labor:
  Stats: dump rates (kps), Full:  8819.6, 9163.4, 8278.6
                     Incremental:  1673.0, 564.0, 2085.5
  Dumps: lev datestmp  tape              file   origK   compK secs
          0  20090417  std-16            12 14428890 14428890 1636
          1  20090424  std-21            16 107070 107070 64

Current info for
lignum-server.lignum.tugraz.local /mnt/samba/temp/Alle_hbf:
  Stats: dump rates (kps), Full:  7335.0, 11736.0, 8382.9
                     Incremental:    5.0,  10.0,  10.0
  Dumps: lev datestmp  tape              file   origK   compK secs
          0  20090421  std-18            12 58680 58680 8
          1  20090424  std-21            53 10 10 2
```

worst case recover

- Amanda saves tar archives directly after the label
- by skipping over the label at the start of the tape you can restore simply by untaring ond/or unzipping

```
dd if=<tape> bs=32k skip=1 | /bin/tar -xpGf - ...
```

changing files (dbms & vmware)

- by using LVM-snapshots rapidly changing data can be saved to (slow)tapes
- classic examples are dbms-storage files or virtual machines (memory and disk files)
- but intensively used /home-partitions may also suffer from ”*file changed while we read it*” errors

→ these areas can be backed up using LVM snapshots

LVM-Snapshot

- is a ”frozen copy” of a volumes state at a specific time
- changes to the original fs are recorded in the snapshot
- if the amount of changed blocks exceed the snapshot size the snapshot gets deactivated (and unusable)
- LVM2 enables read/write snapshots – it becomes possible to test changes inside snapshots by writing to the snapshot
- this may be used i.e. for shared XEN Volumes – i.e. one base volume and each Xen instance operates using a writeable snapshot recording only changes (i.e. even for VMWare Machines)

create snapshots & back these up

- recreate the snapshot before the next backup run to enable incremental backups

```
#!/bin/bash
# snapshot rotation script by Gernot Pirnbacher
DEVPATH="/dev/ml115/snap"
# 2 lines needed because create via /dev/mapper/vg-lv fails - /dev/vg/lv needed
BASEPATH="/dev/ml115/data"
GREPPATH="/dev/mapper/ml115-data"
echo "Checking for a snapshot-device mounted (ml115-specific) device = $DEVPATH"
echo "------------------------------------------------------------------------"
MOUNTPATH= mount | grep $GREPPATH
if [ $? = 0 ]; then
    echo "Found device mounted: $MOUNTPATH"
    # unmounten des Devices
    umount /mnt/snapshot
    # löschen des alten snapshots
    lvremove -f $DEVPATH
    # neuen Snapshot generieren
    lvcreate --name snap --snapshot $BASEPATH -L 35G
    # mounten des Snapshots
    mount $DEVPATH /mnt/snapshot -o ro
    # mounten erfolgreich?
    if [ $? = 0 ]; then
        echo "Snapshot $DEVPATH erfolgreich gemountet!"
    fi
elif  [ $? = 1 ]; then
    # Meldung über nicht gefundenes Device ....
    echo "Did not find configured device ($DEVPATH) mounted!"
```

links for amanda & the end

Thanks for your attention!

e-mail: office@it-to-fit.at

main links for amanda:

http://www.amanda.org

http://wiki.zmanda.com