

# Backup with rdiff-backup and rsnapshot

*Christoph Mitasch & Thomas Niedermeier, Thomas-Krenn.AG  
@Cmitasch, @tk\_niedermeier*

*Open Source Backup Conference 2015  
September 29<sup>th</sup>-30<sup>th</sup> 2015, Cologne, Germany*



Open Source Backup  
Conference

**TH=MAS**  
**KRENN**<sup>®</sup>  
server.hosting.customized.

# Status quo

— Do you know **rsnapshot**?

| Yes | No |
|-----|----|
|     |    |

# Status quo

— Do you know **rdiff-backup**?

| Yes | No |
|-----|----|
|     |    |

# Status quo

— What do you use for backup now?

| Bareos / Bacula | rsync | other OSS | commercial |
|-----------------|-------|-----------|------------|
|                 |       |           |            |

# Agenda

- \_ intro (5')
- \_ rsnapshot (15')
- \_ rdiff-backup (15')
- \_ so what (should I use) – a comparison (5')

rsnapshot

rsnapshot?  
= rsync + hard links

# Basics

- A filesystem snapshot utility
- GNU General Public Licence
- Based on rsync
- Using hard links for increments
- Written in perl / no module dependencies
- Local and remote via rsync over ssh
- Triggered by cron or anacron





# Benefits

- Minimal disk space required
- Hourly, daily, weekly and monthly snapshots
- Direct access to every snapshot
- Easy to use
- Tested working on a lot of distributions

# How to get?

## — source-releases

- <https://rsnapshot.org/downloads/>

## — git

- `$ git clone git://github.com/rsnapshot/rsnapshot.git`



## — Packages

- Debian, Ubuntu
  - `apt-get install rsnapshot`
- Fedora
  - `yum install rsnapshot`
- ArchLinux
  - `pacman -S rsnapshot`
- Gentoo, FreeBSD, NetBSD, OpenBSD ...



FreeBSD

# Structure



- \_ A single configuration file
  - \_ /etc/rsnapshot.conf
  - \_ Key parameters
    - snapshot\_root
    - retain daily|weekly|monthly
    - verbose
    - logfile <file>
    - include <DIR>
    - exclude <DIR>
    - backup
    - backup\_script

# Key parameter 1/3



## — snapshot\_root

- By default: /var/cache/rsnapshot/

## — Retain policy

- retain hourly X
- retain daily X
- retain weekly X
- retain monthly X

# Key parameter 2/3



## — Verbosity

- Parameter: verbose {1,5}
- Levels 1 – 5
  - From errors only to debug mode



## — Logging

- Parameter: loglevel
- Parameter: logfile



# Key parameter 3/3



## — Include folders/files

- Parameter: include


## — Exclude folders/files

- Parameter: exclude

## — Exclude folders of remote backups


- Add „+rsync\_long\_args+=exclude=/some/folder/“

# Local backup example



```
_ /etc/rsnapshot.conf
_ snapshot_root /var/cache/rsnapshot/
_ retain hourly 4
  retain daily 7
  retain weekly 4
  retain monthly 3
_ verbose 3
_ loglevel 3
_ logfile /var/log/rsnapshot.log
_ backup /<DIR>/ localhost/
```

# Remote backup example




\_ /etc/rsnapshot.conf

- \_ rsync\_long\_args -ev -rsync-path=/home/rsnap/rsync-wrapper.sh
- \_ ssh\_args -i /home/rsnap/.ssh/id\_rsa
- \_ backup rsnap@<IP>:/<DIR>/ lesv2/<DIR>/



# Pull backup best practice



## — Own backup user

- e.g. rsnap
- Not root!
- Public key logins without password

## — MySQL Backups

- `backup_script /usr/bin/ssh -i /home/rsnap/.ssh/id_rsa rsnap@<IP_Address> 'mysqldump --all-databases | gzip --rsyncable > ~/alldb.sql.gz' unused1/`
- Placed above the backup entries of this server
- `backup_script` will be executed and the zipped file will be backed up with the backup `/home/` command

# Automation



always on

→ cronjobs



desktop and not always on

→ anacron

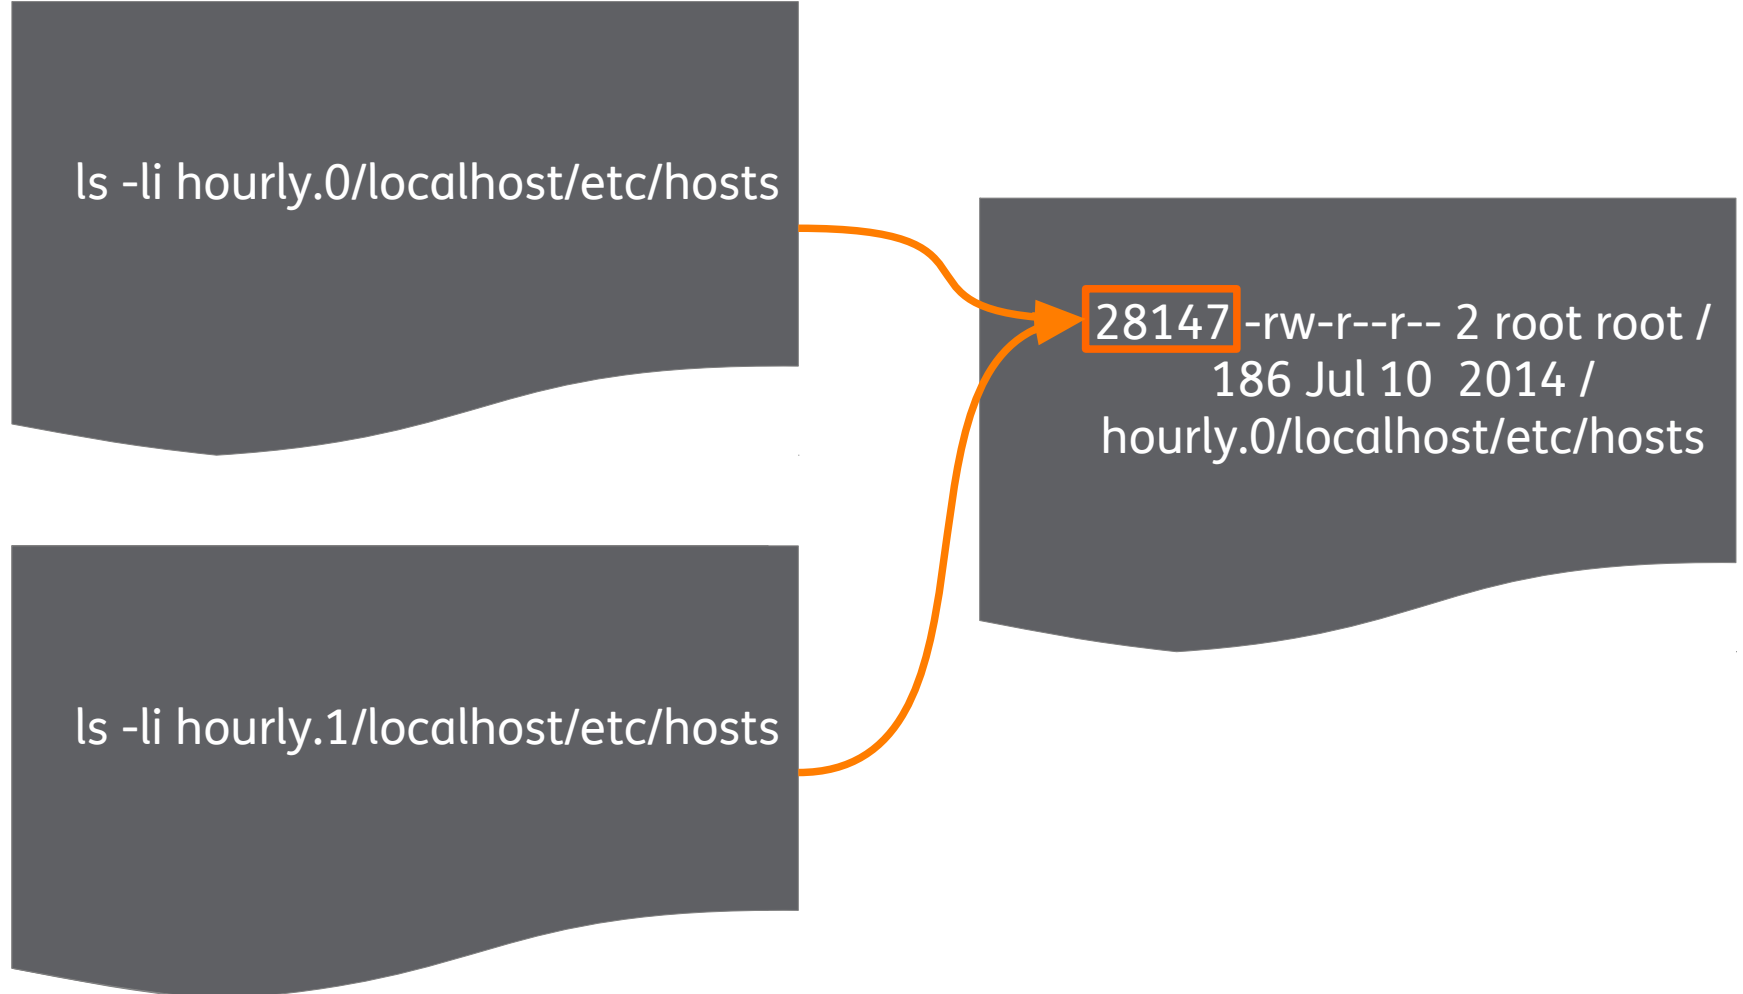
*for more details see:  
[tkwiki.cc/rsnapshot](http://tkwiki.cc/rsnapshot)*

# Hard link example 1/2

```
ls -li hourly.0/localhost/etc/hosts
```

```
ls -li hourly.1/localhost/etc/hosts
```

```
28147 -rw-r--r-- 2 root root /  
186 Jul 10 2014 /  
hourly.0/localhost/etc/hosts
```

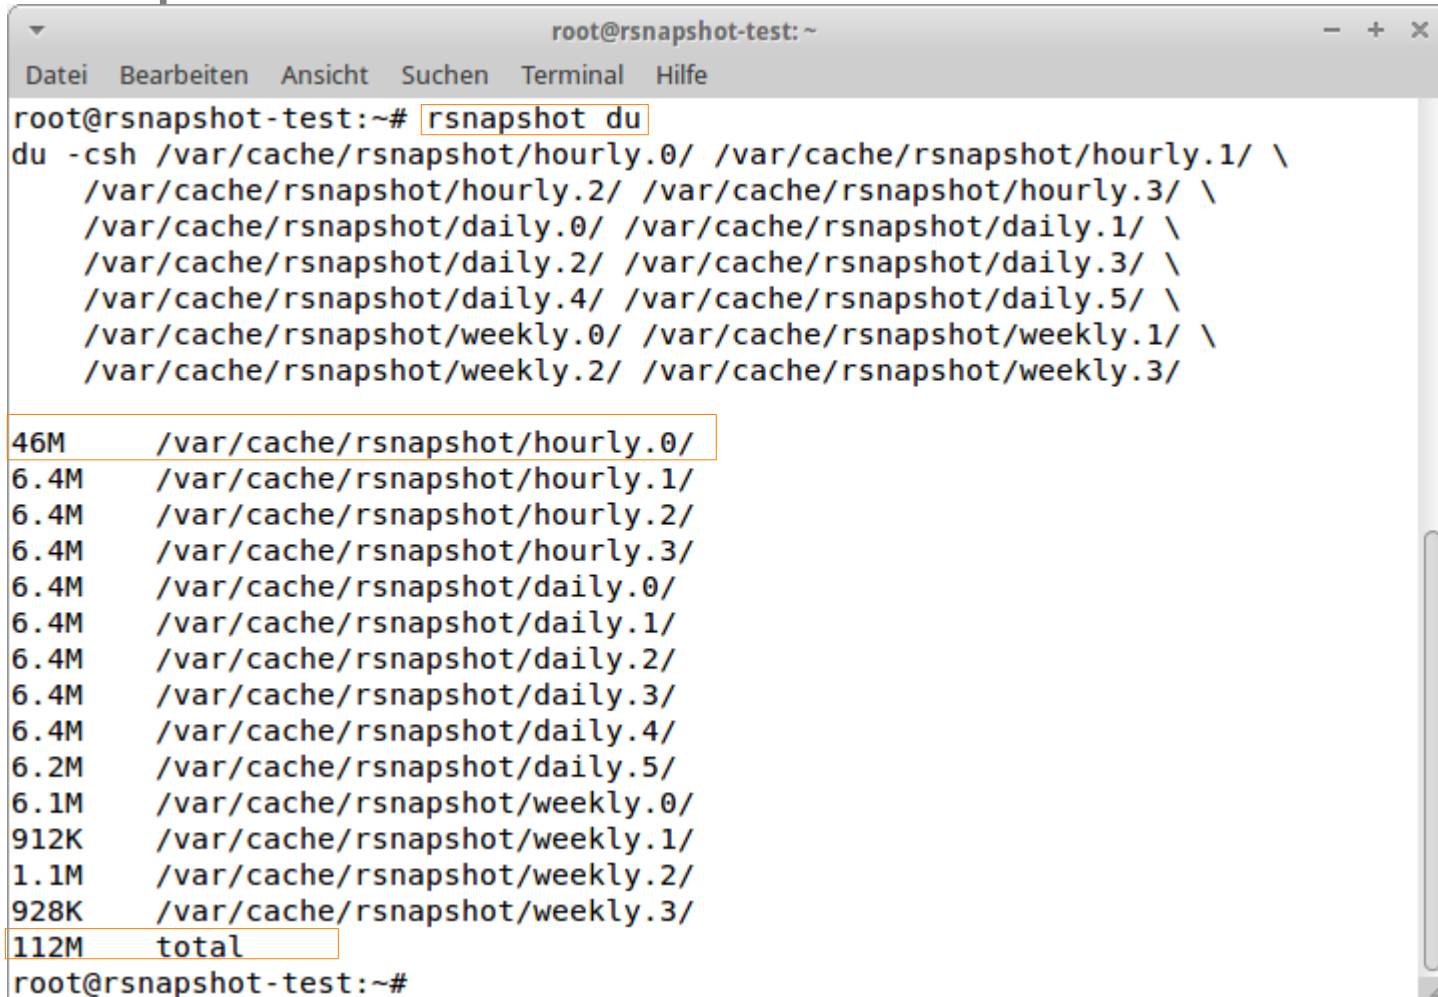


# Hard link example 2/2

```
— # ls
  _ hourly.0 hourly.1
— # ls -li hourly.*/localhost/etc/hosts
  _ 28147 -rw-r--r-- 2 root root 186 Jul 10 2015
    hourly.0/localhost/etc/hosts
  _ 28147 -rw-r--r-- 2 root root 186 Jul 10 2015
    hourly.1/localhost/etc/hosts
— # du -sh hourly.*
  _ 47M      hourly.0
  _ 908K     hourly.1
```

# Tools 1/3

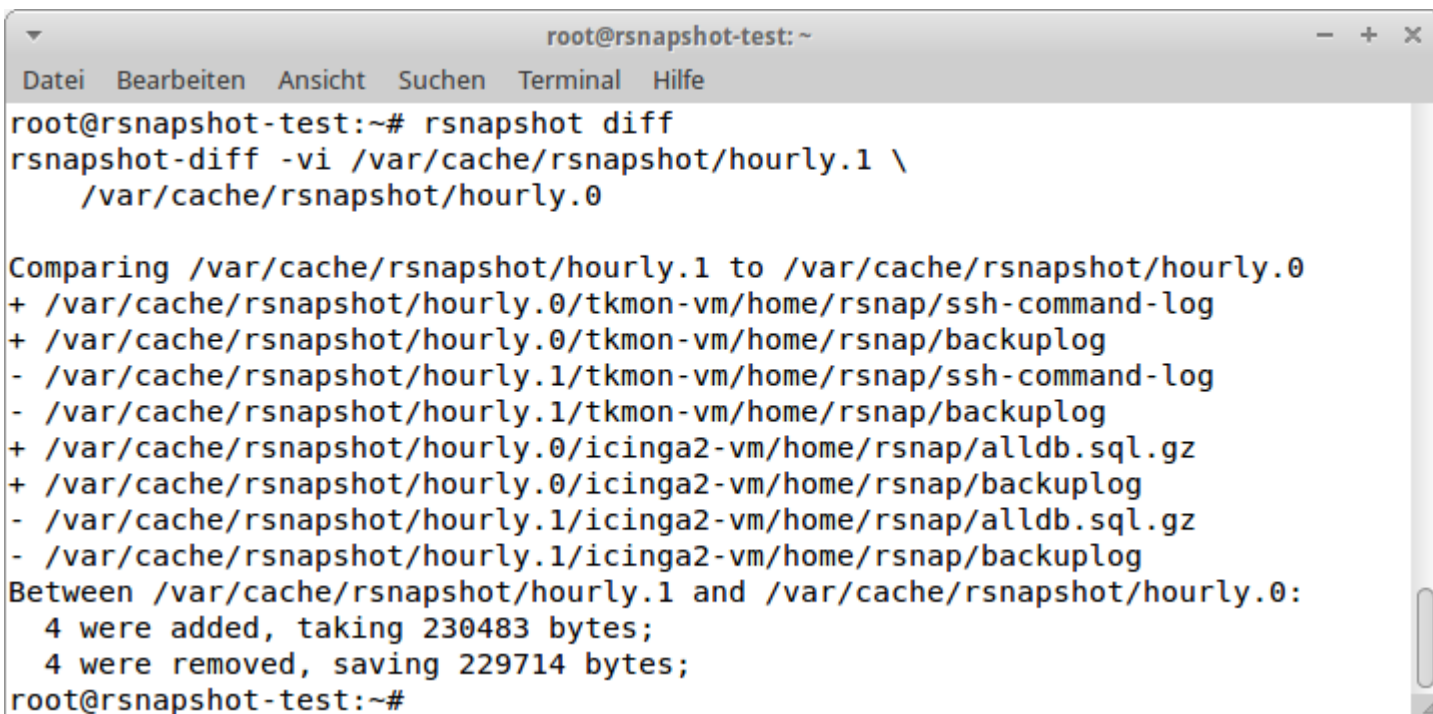
## — rsnapshot du



```
root@rsnapshot-test: ~  
Datei Bearbeiten Ansicht Suchen Terminal Hilfe  
root@rsnapshot-test:~# rsnapshot du  
du -csh /var/cache/rsnapshot/hourly.0/ /var/cache/rsnapshot/hourly.1/ \  
/var/cache/rsnapshot/hourly.2/ /var/cache/rsnapshot/hourly.3/ \  
/var/cache/rsnapshot/daily.0/ /var/cache/rsnapshot/daily.1/ \  
/var/cache/rsnapshot/daily.2/ /var/cache/rsnapshot/daily.3/ \  
/var/cache/rsnapshot/daily.4/ /var/cache/rsnapshot/daily.5/ \  
/var/cache/rsnapshot/weekly.0/ /var/cache/rsnapshot/weekly.1/ \  
/var/cache/rsnapshot/weekly.2/ /var/cache/rsnapshot/weekly.3/  
  
46M /var/cache/rsnapshot/hourly.0/  
6.4M /var/cache/rsnapshot/hourly.1/  
6.4M /var/cache/rsnapshot/hourly.2/  
6.4M /var/cache/rsnapshot/hourly.3/  
6.4M /var/cache/rsnapshot/daily.0/  
6.4M /var/cache/rsnapshot/daily.1/  
6.4M /var/cache/rsnapshot/daily.2/  
6.4M /var/cache/rsnapshot/daily.3/  
6.4M /var/cache/rsnapshot/daily.4/  
6.2M /var/cache/rsnapshot/daily.5/  
6.1M /var/cache/rsnapshot/weekly.0/  
912K /var/cache/rsnapshot/weekly.1/  
1.1M /var/cache/rsnapshot/weekly.2/  
928K /var/cache/rsnapshot/weekly.3/  
112M total  
root@rsnapshot-test:~#
```

# Tools 2/3

## \_ rsnapshot diff




```
root@rsnapshot-test: ~
Datei Bearbeiten Ansicht Suchen Terminal Hilfe
root@rsnapshot-test:~# rsnapshot diff
rsnapshot-diff -vi /var/cache/rsnapshot/hourly.1 \
  /var/cache/rsnapshot/hourly.0

Comparing /var/cache/rsnapshot/hourly.1 to /var/cache/rsnapshot/hourly.0
+ /var/cache/rsnapshot/hourly.0/tkmon-vm/home/rsnap/ssh-command-log
+ /var/cache/rsnapshot/hourly.0/tkmon-vm/home/rsnap/backuplog
- /var/cache/rsnapshot/hourly.1/tkmon-vm/home/rsnap/ssh-command-log
- /var/cache/rsnapshot/hourly.1/tkmon-vm/home/rsnap/backuplog
+ /var/cache/rsnapshot/hourly.0/icinga2-vm/home/rsnap/allldb.sql.gz
+ /var/cache/rsnapshot/hourly.0/icinga2-vm/home/rsnap/backuplog
- /var/cache/rsnapshot/hourly.1/icinga2-vm/home/rsnap/allldb.sql.gz
- /var/cache/rsnapshot/hourly.1/icinga2-vm/home/rsnap/backuplog
Between /var/cache/rsnapshot/hourly.1 and /var/cache/rsnapshot/hourly.0:
  4 were added, taking 230483 bytes;
  4 were removed, saving 229714 bytes;
root@rsnapshot-test:~#
```

# Tools 3/3

- \_ Check the configuration!
  - \_ `root@rsnapshot-test:~# rsnapshot configtest`
  - \_ Syntax OK



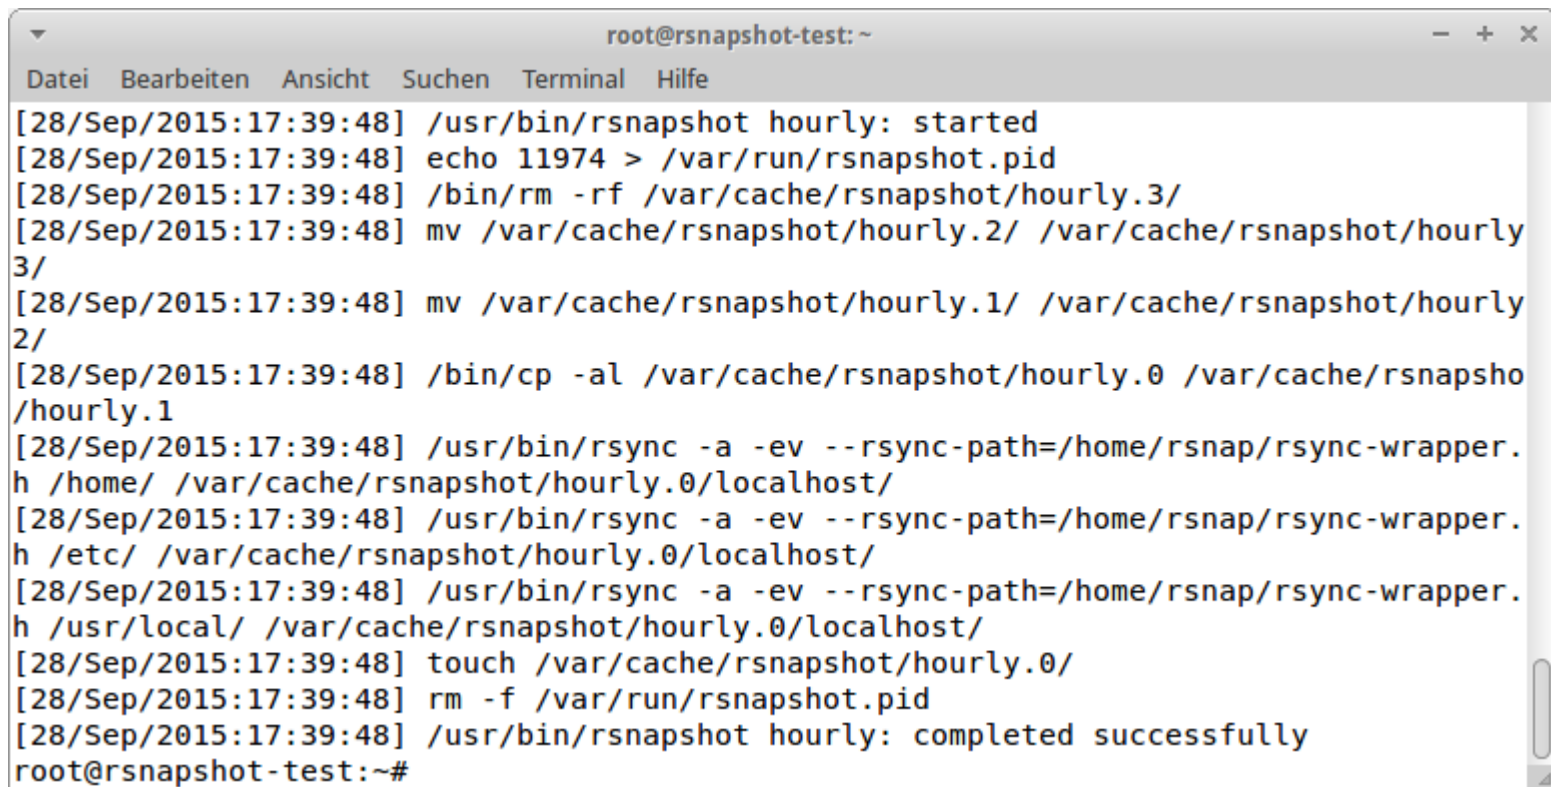


# how rsnapshot runs

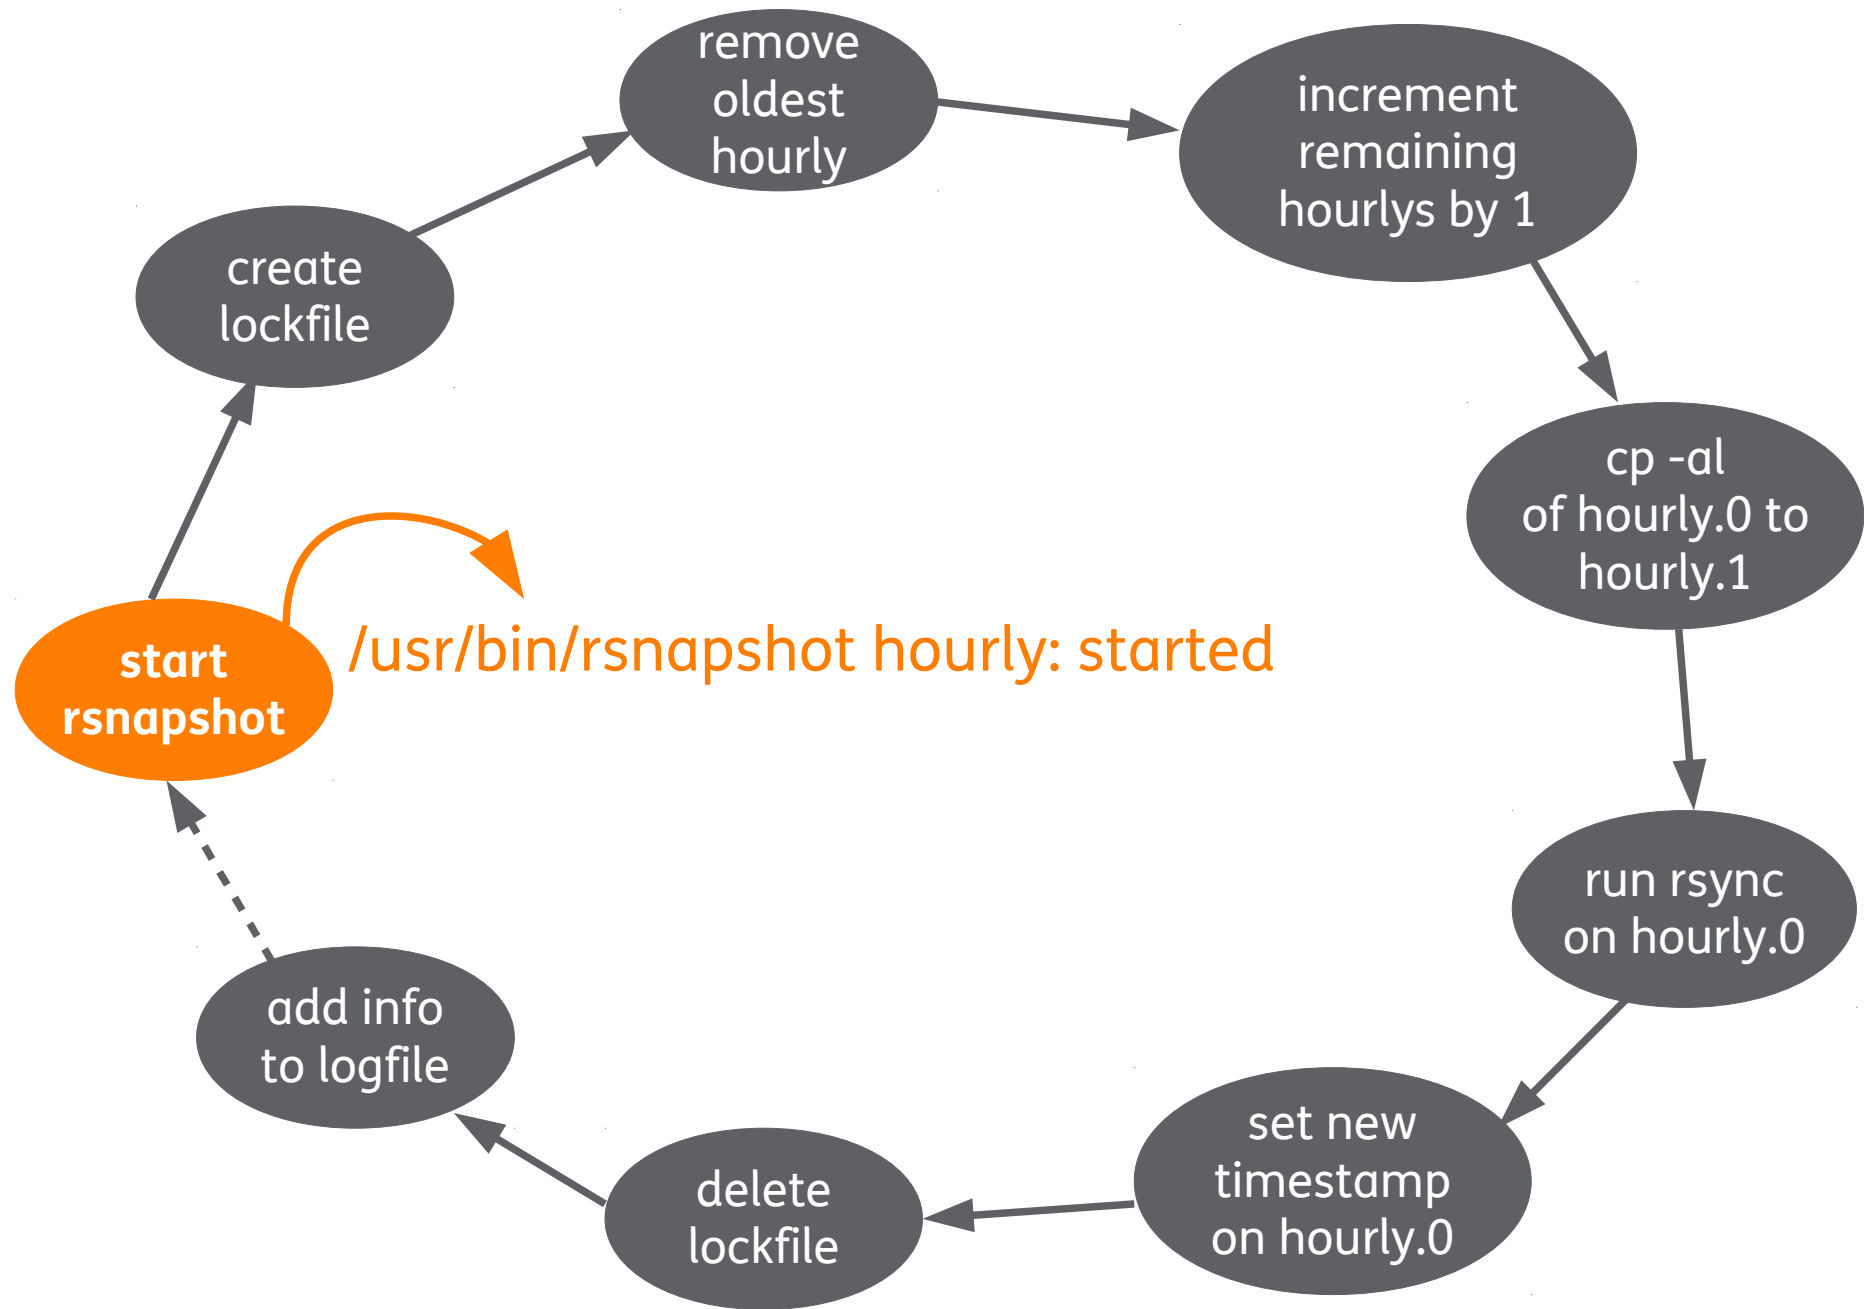


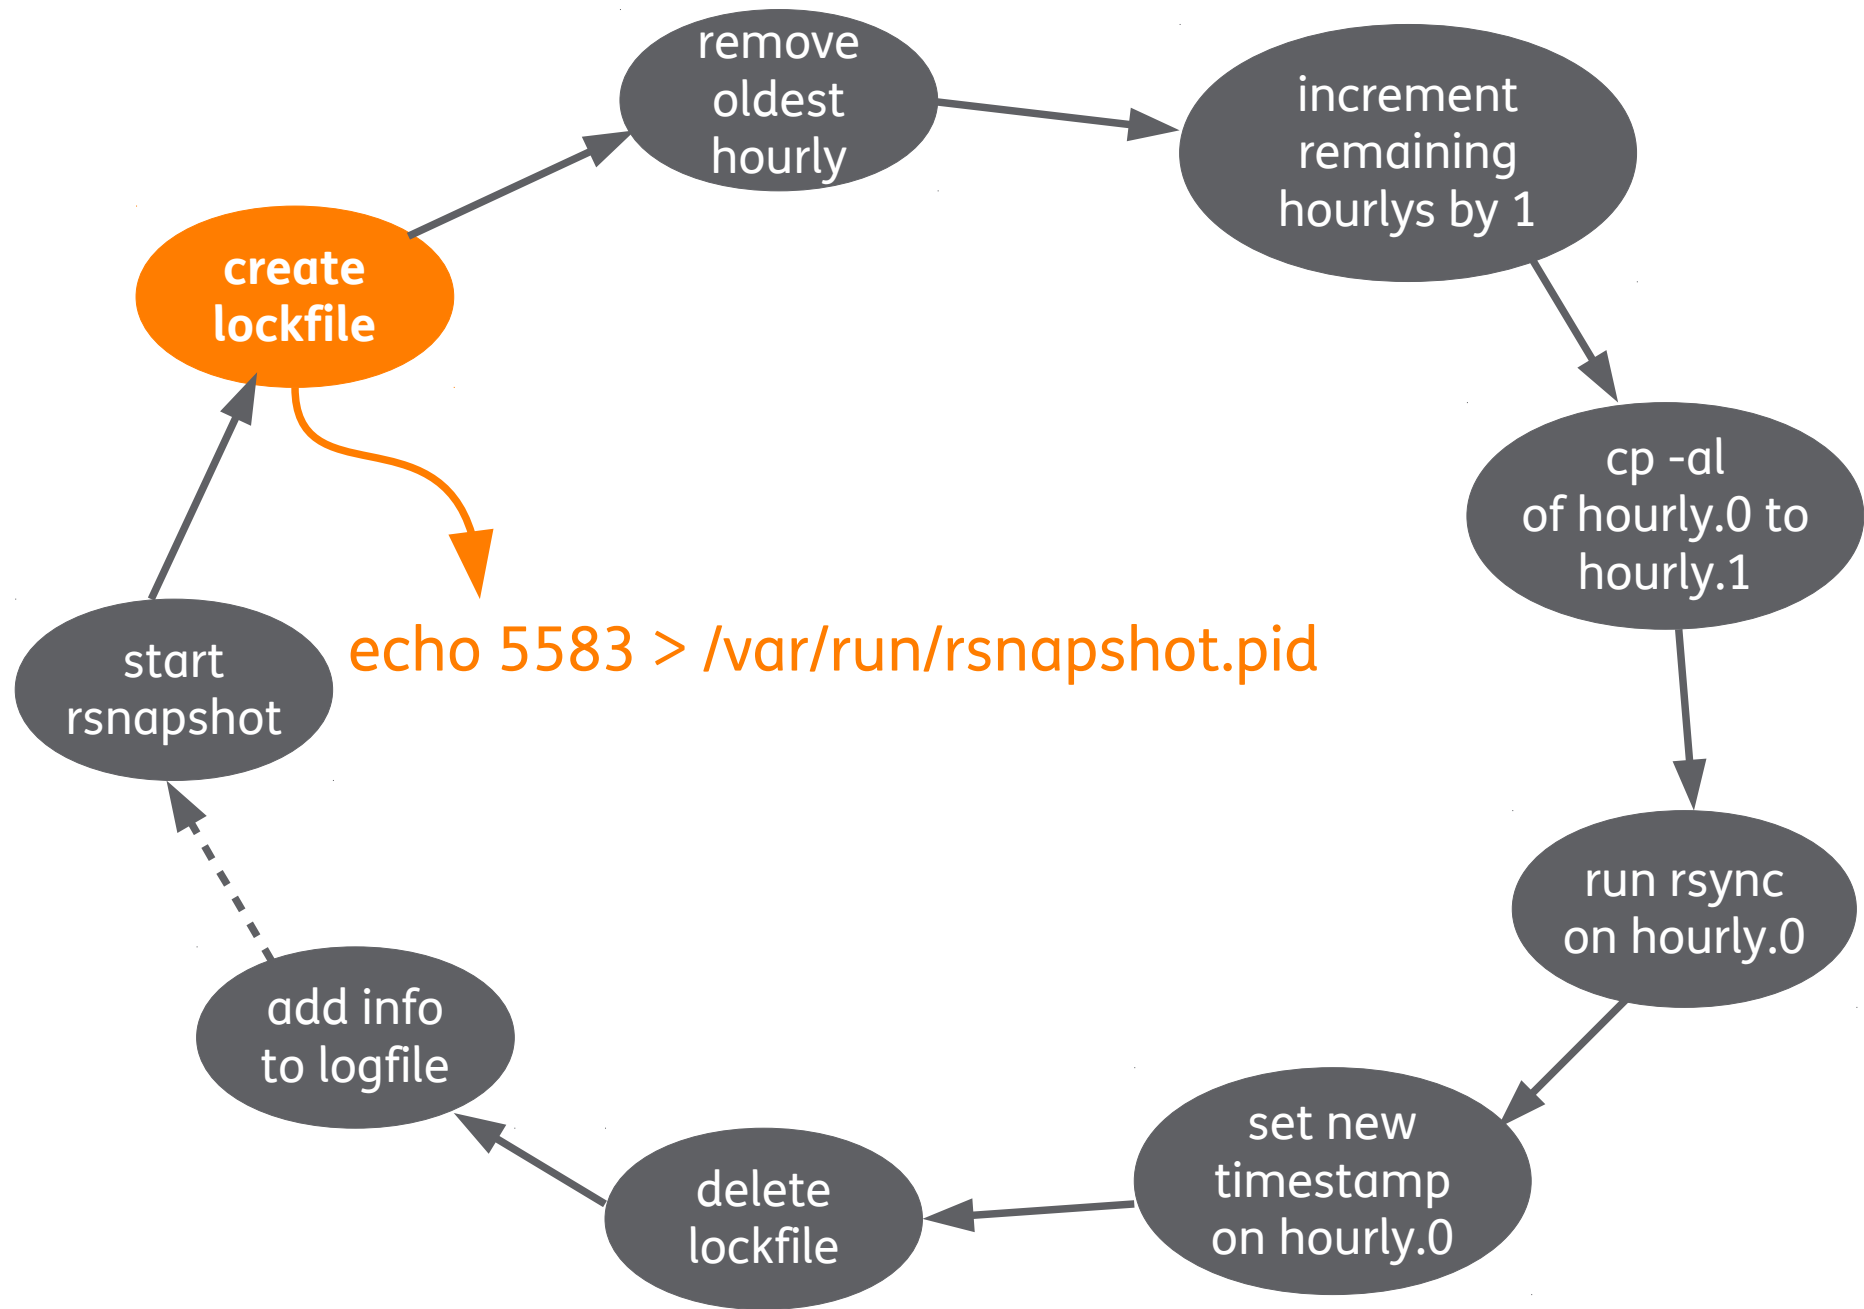
# How rsnapshots run

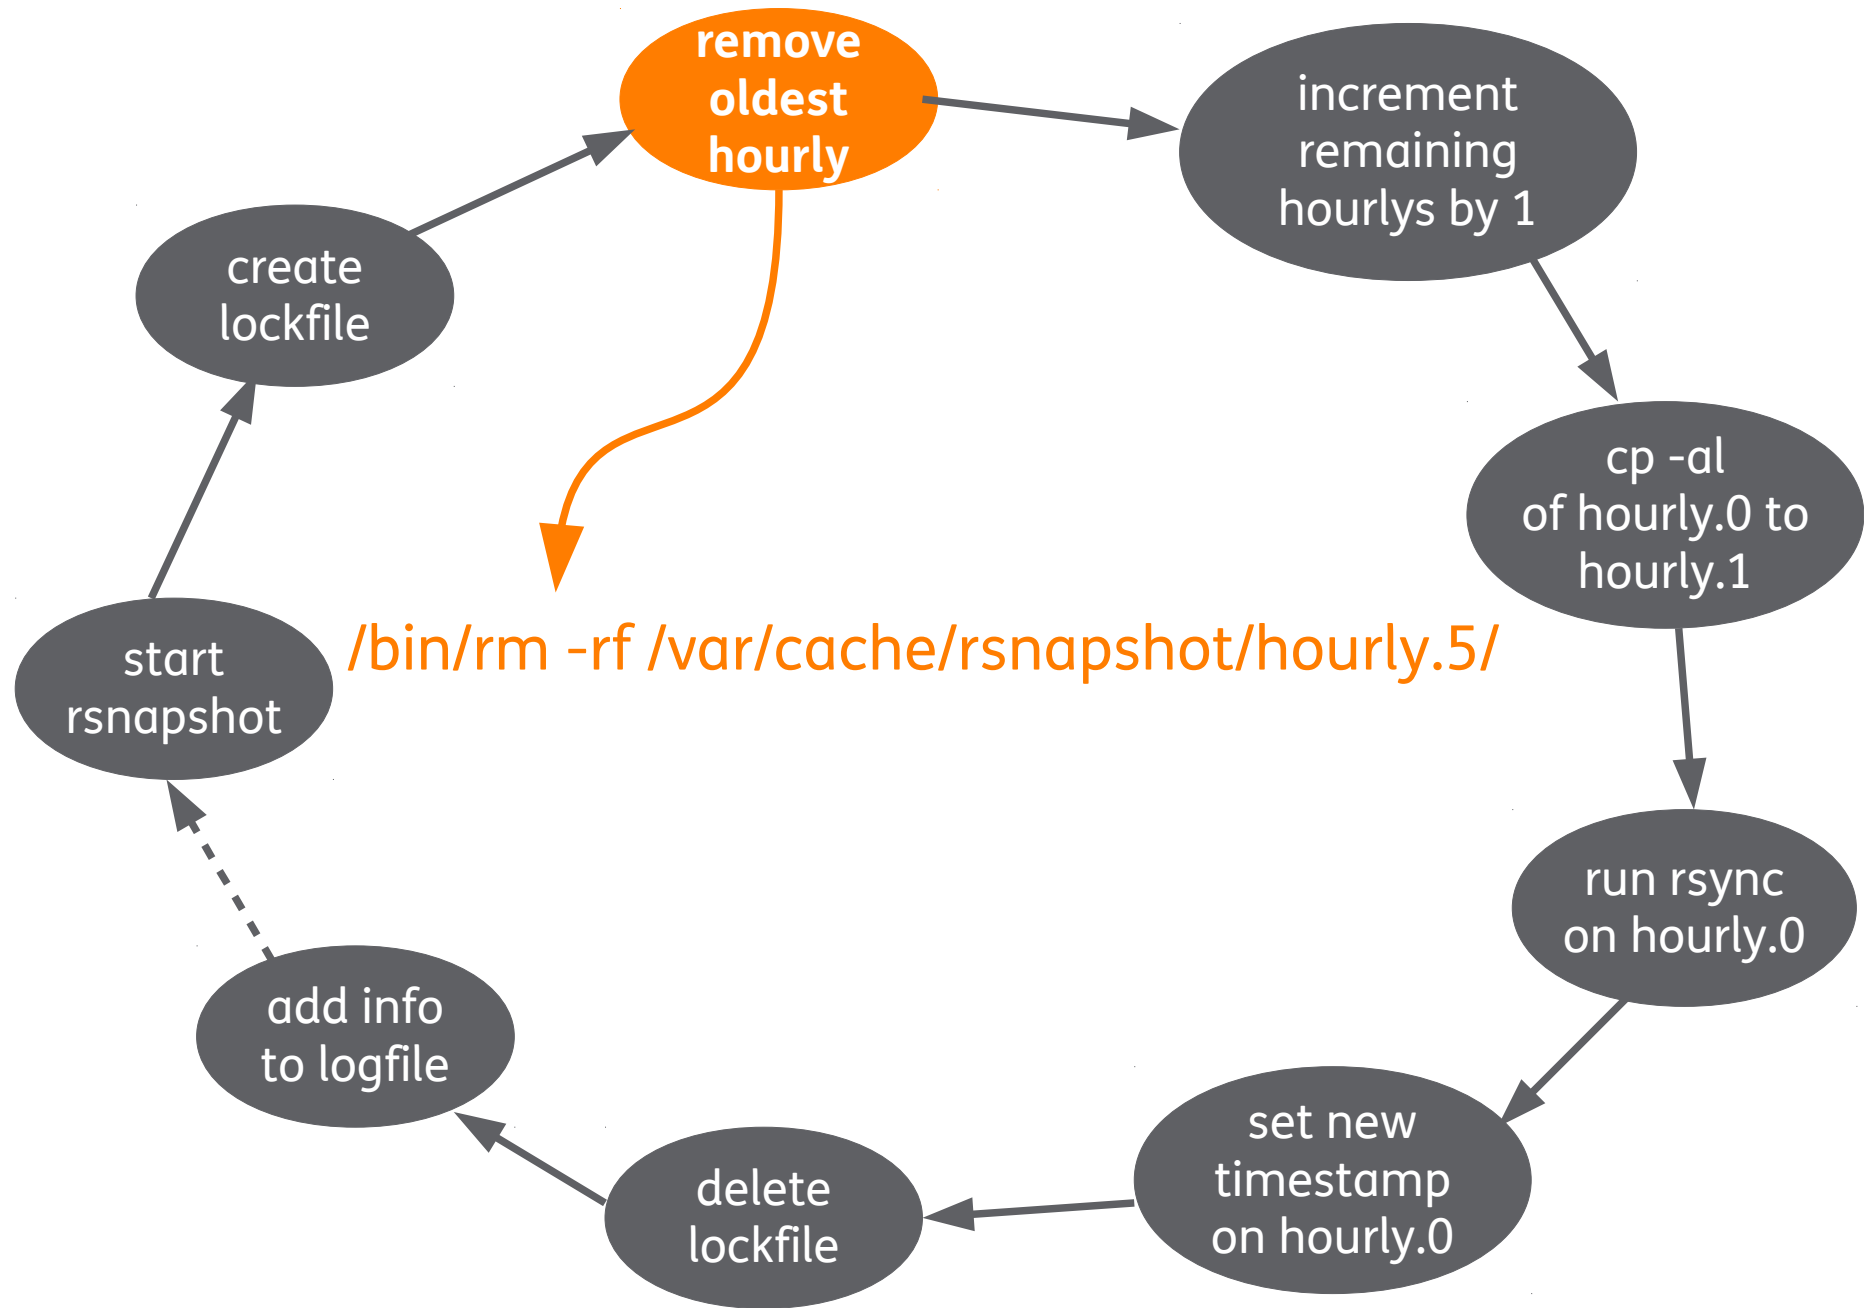
— `cat /var/log/rsnapshot.log`

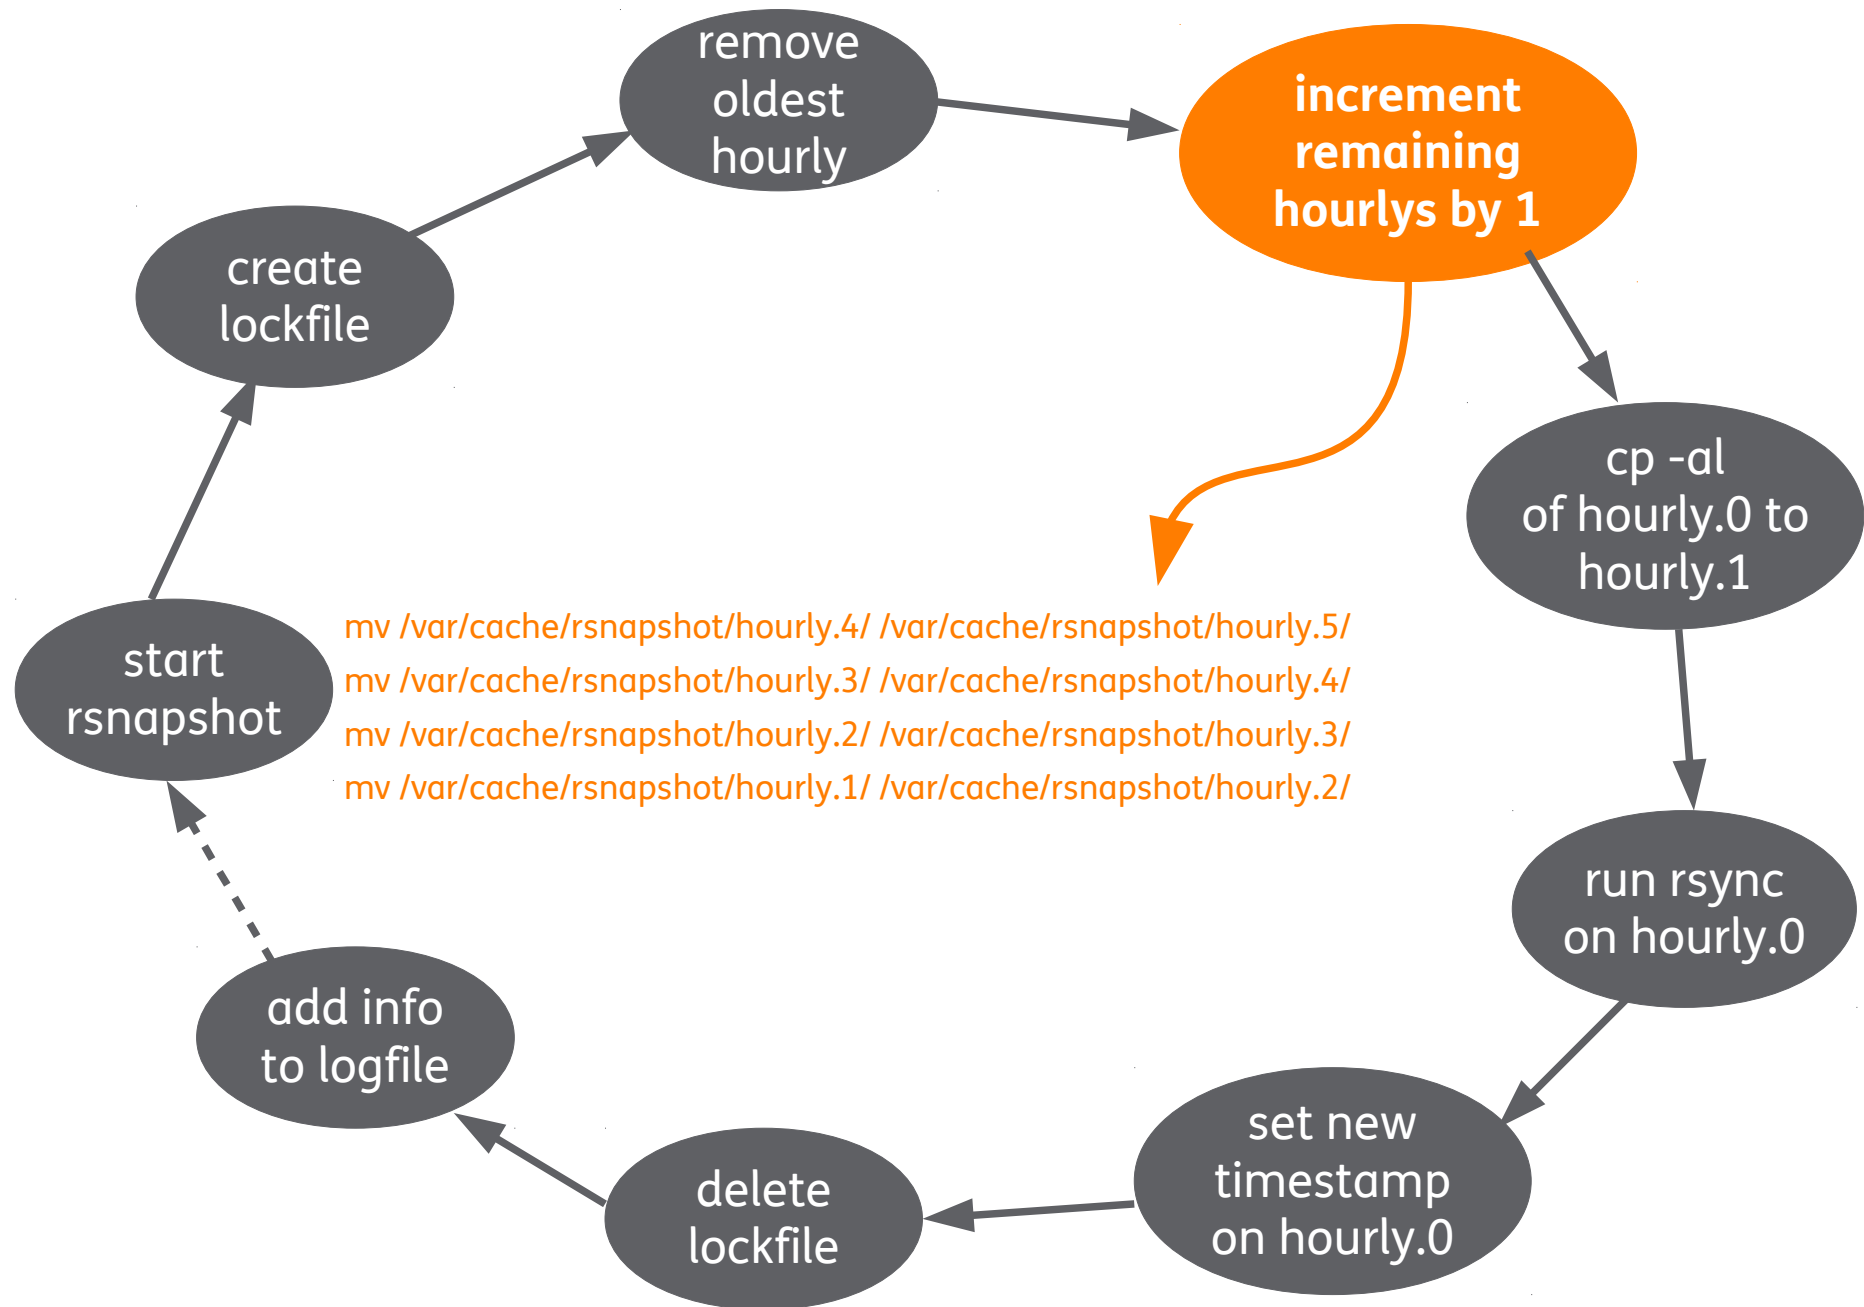


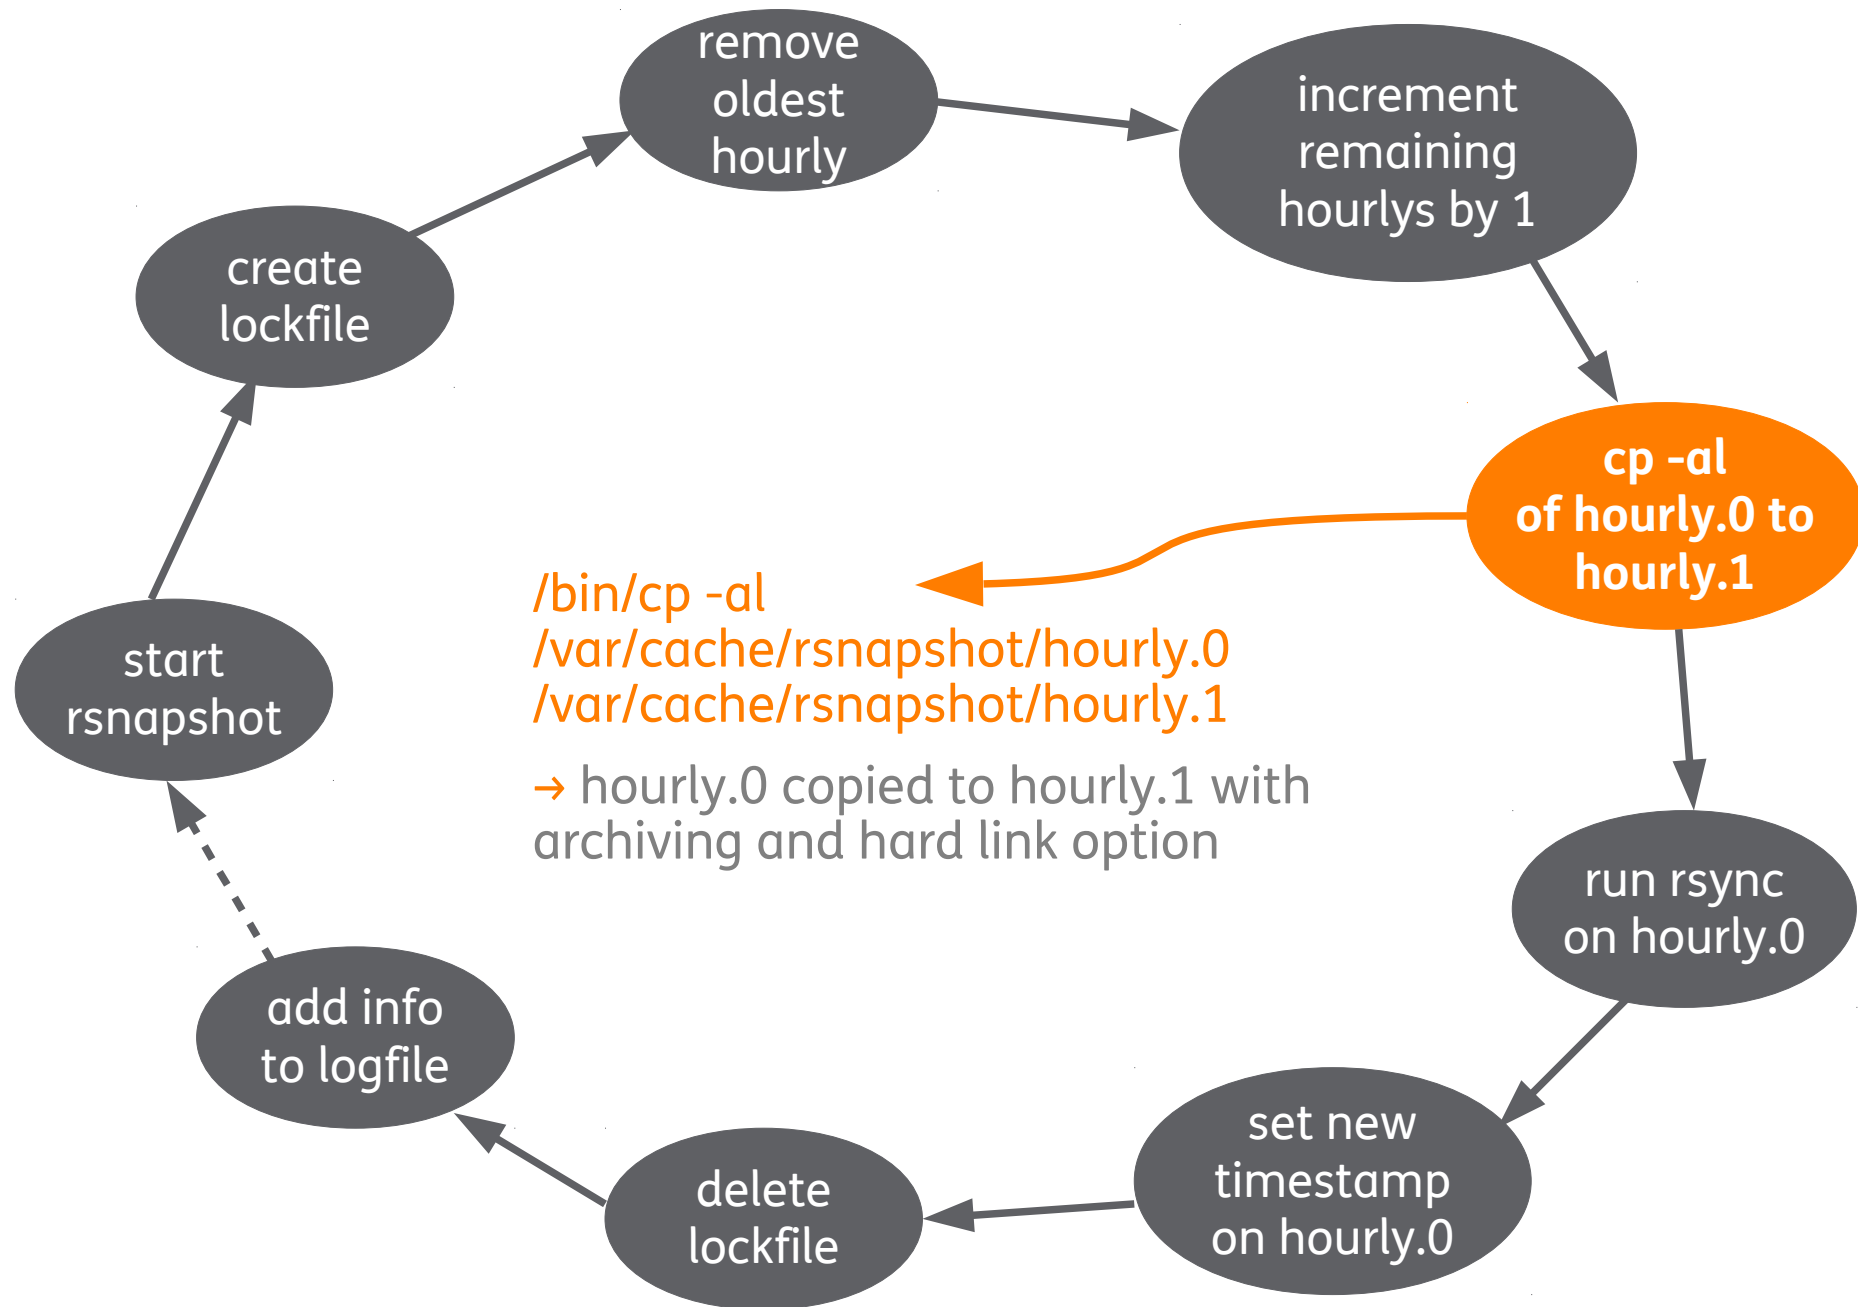
```
root@rsnapshot-test: ~
Datei Bearbeiten Ansicht Suchen Terminal Hilfe
[28/Sep/2015:17:39:48] /usr/bin/rsnapshot hourly: started
[28/Sep/2015:17:39:48] echo 11974 > /var/run/rsnapshot.pid
[28/Sep/2015:17:39:48] /bin/rm -rf /var/cache/rsnapshot/hourly.3/
[28/Sep/2015:17:39:48] mv /var/cache/rsnapshot/hourly.2/ /var/cache/rsnapshot/hourly
3/
[28/Sep/2015:17:39:48] mv /var/cache/rsnapshot/hourly.1/ /var/cache/rsnapshot/hourly
2/
[28/Sep/2015:17:39:48] /bin/cp -al /var/cache/rsnapshot/hourly.0 /var/cache/rsnapsho
/hourly.1
[28/Sep/2015:17:39:48] /usr/bin/rsync -a -ev --rsync-path=/home/rsnap/rsync-wrapper.
h /home/ /var/cache/rsnapshot/hourly.0/localhost/
[28/Sep/2015:17:39:48] /usr/bin/rsync -a -ev --rsync-path=/home/rsnap/rsync-wrapper.
h /etc/ /var/cache/rsnapshot/hourly.0/localhost/
[28/Sep/2015:17:39:48] /usr/bin/rsync -a -ev --rsync-path=/home/rsnap/rsync-wrapper.
h /usr/local/ /var/cache/rsnapshot/hourly.0/localhost/
[28/Sep/2015:17:39:48] touch /var/cache/rsnapshot/hourly.0/
[28/Sep/2015:17:39:48] rm -f /var/run/rsnapshot.pid
[28/Sep/2015:17:39:48] /usr/bin/rsnapshot hourly: completed successfully
root@rsnapshot-test:~#
```

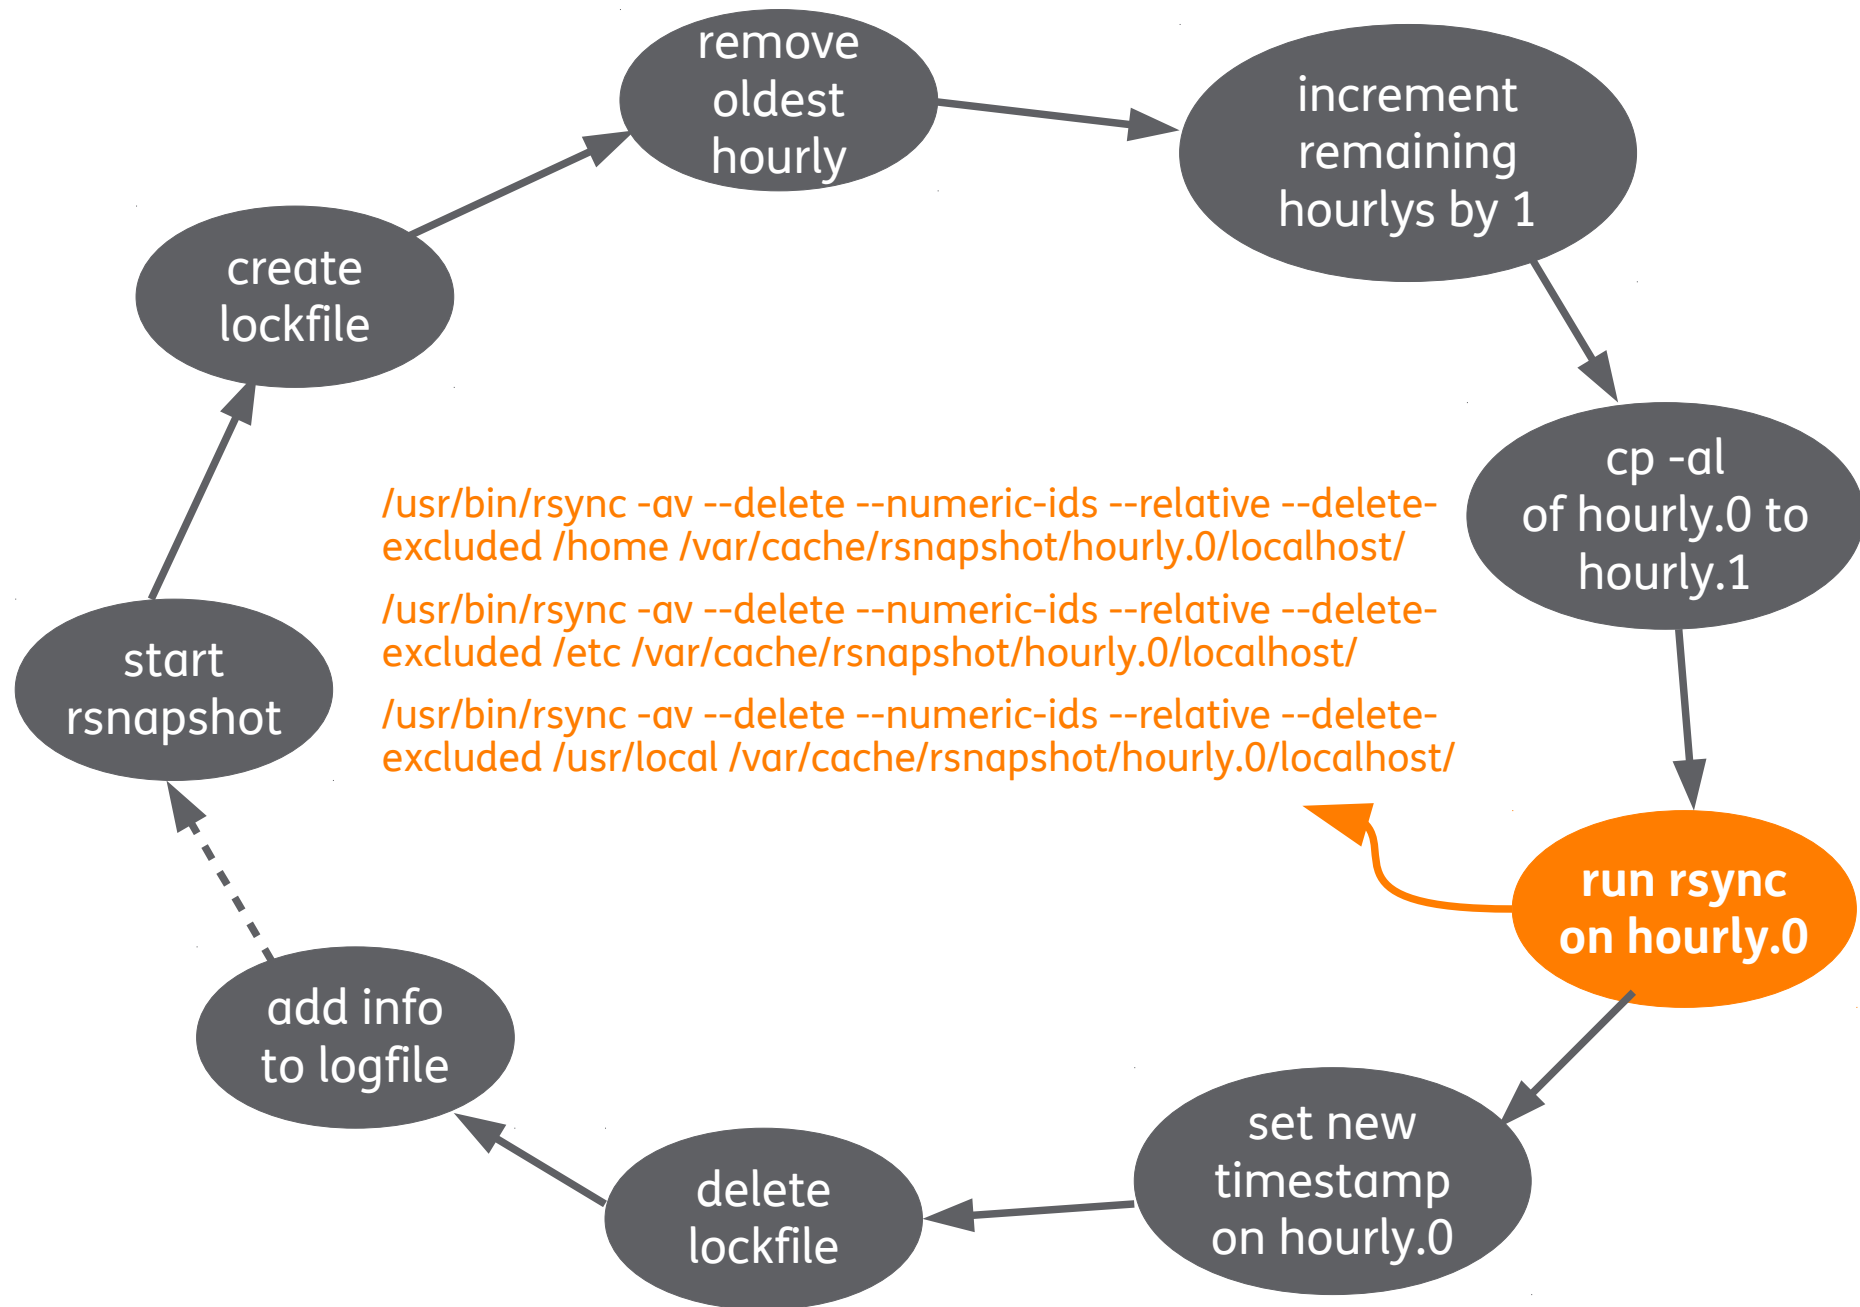


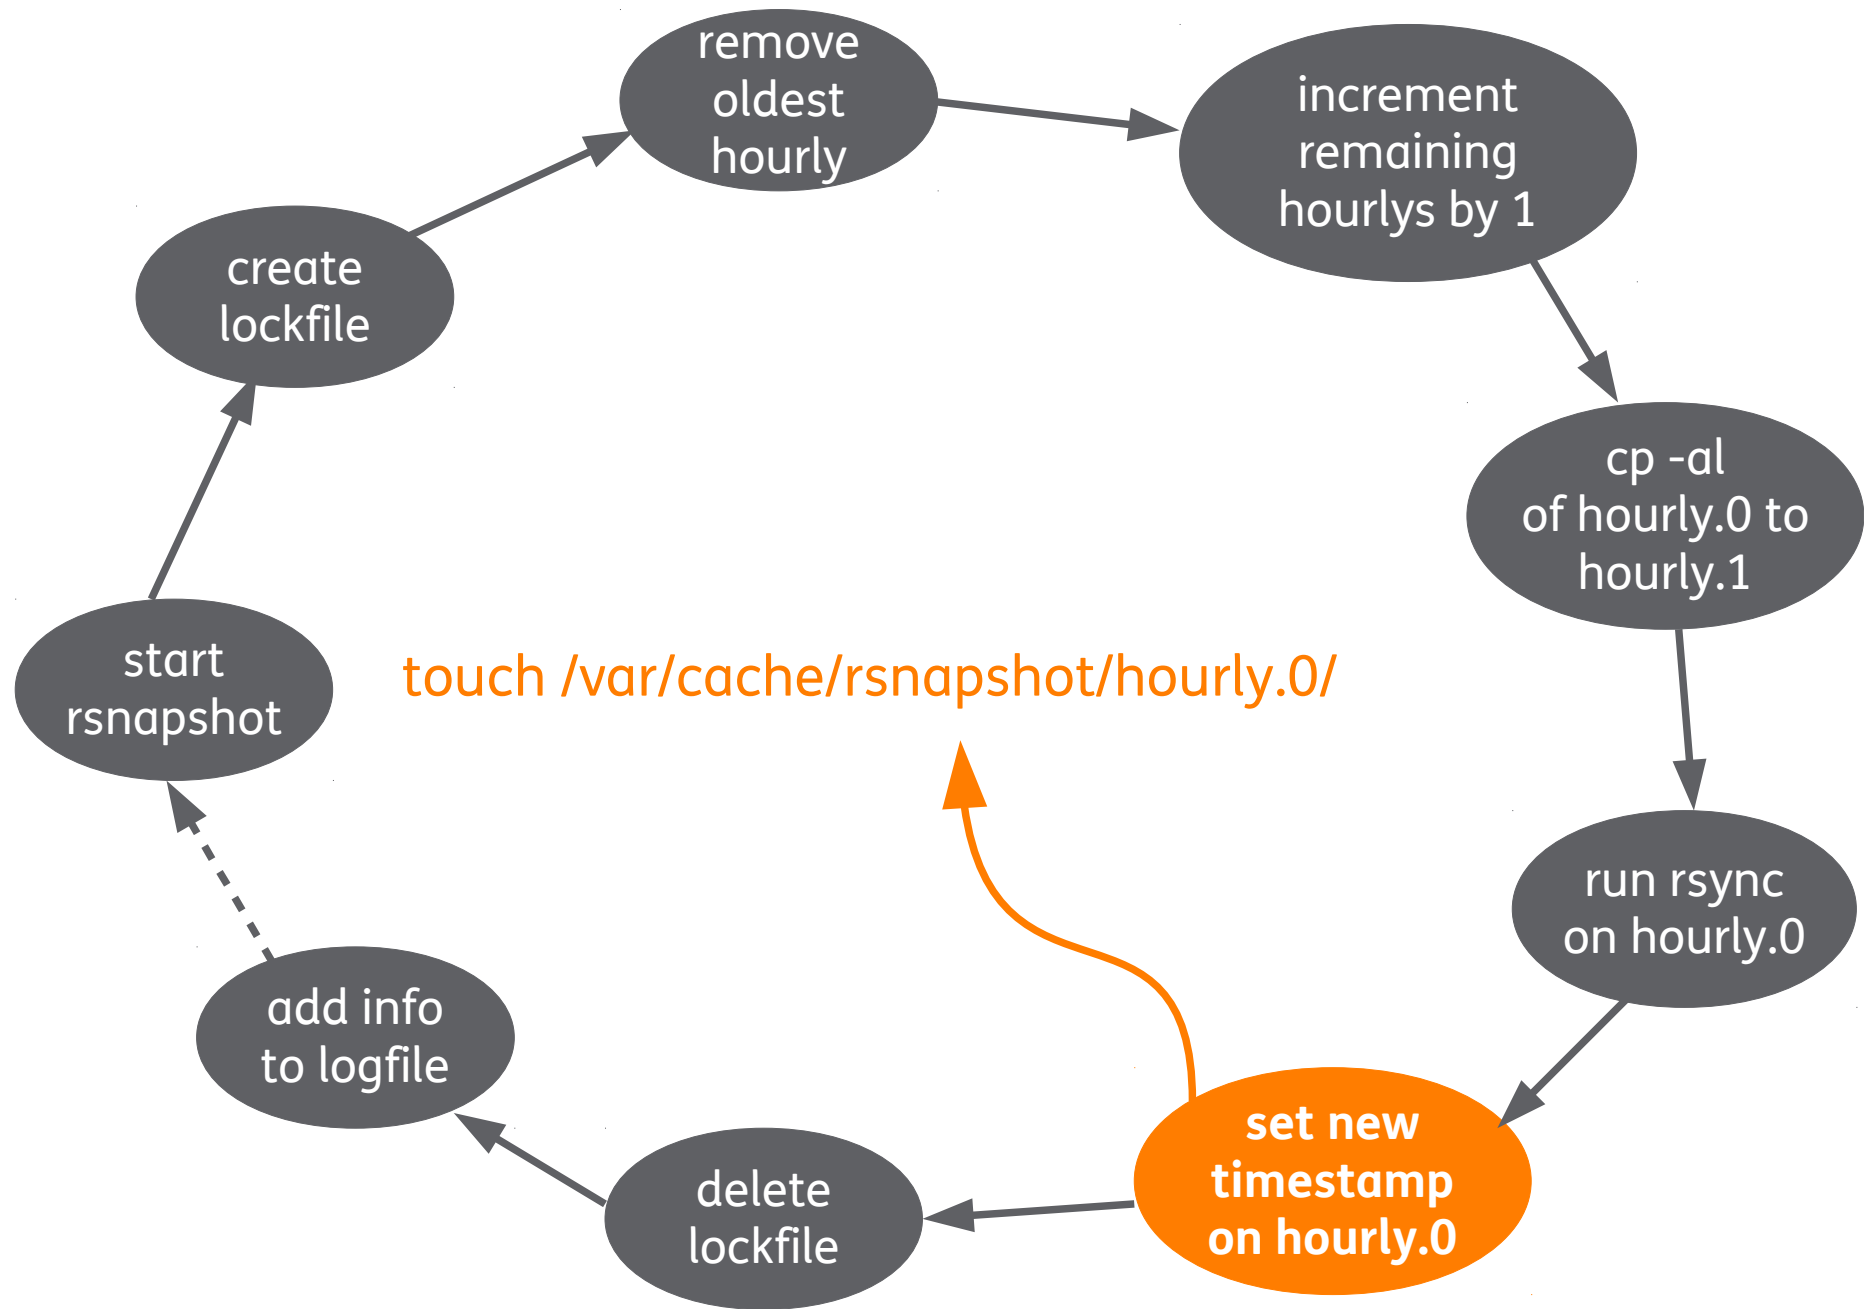




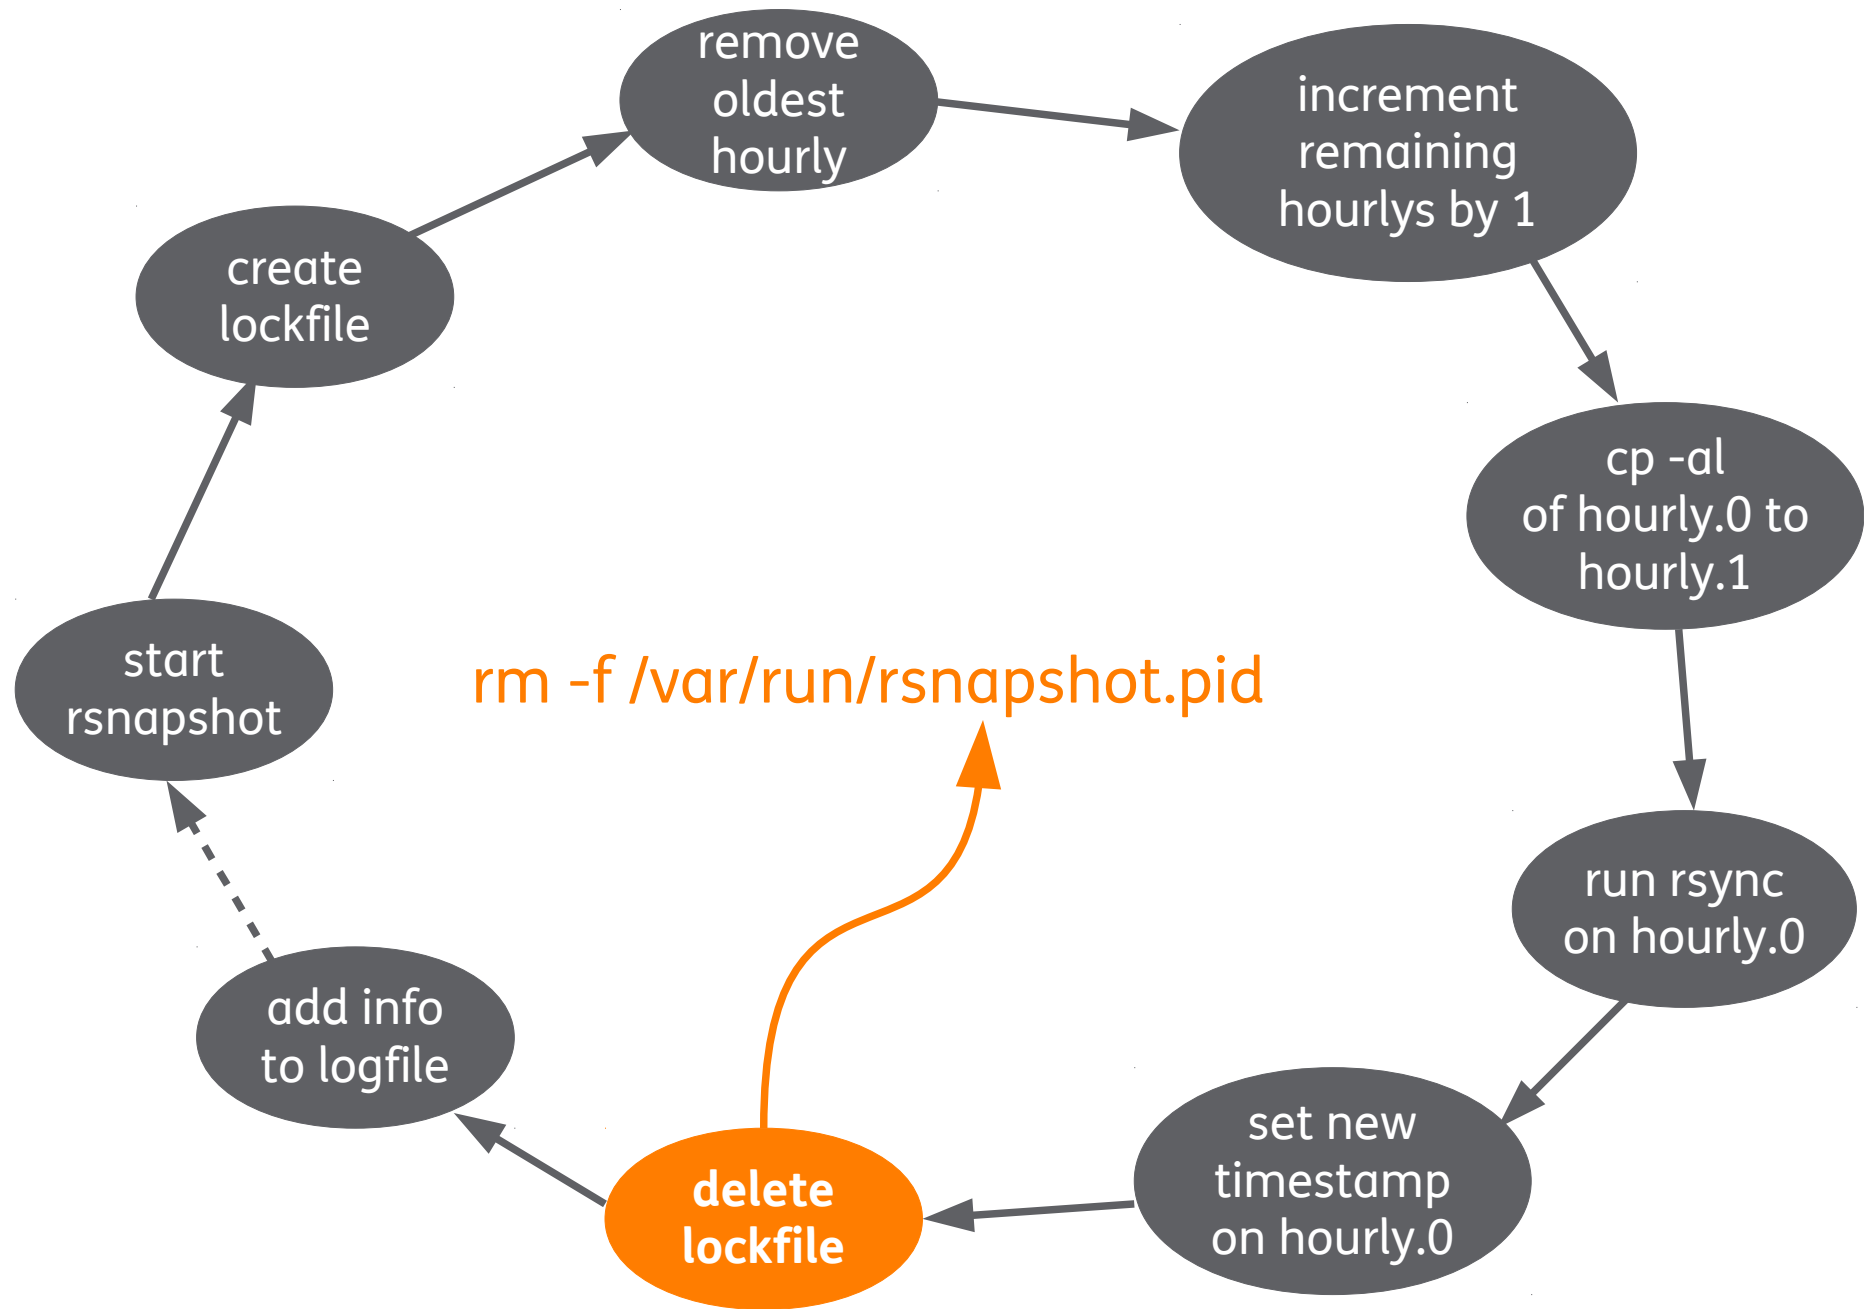


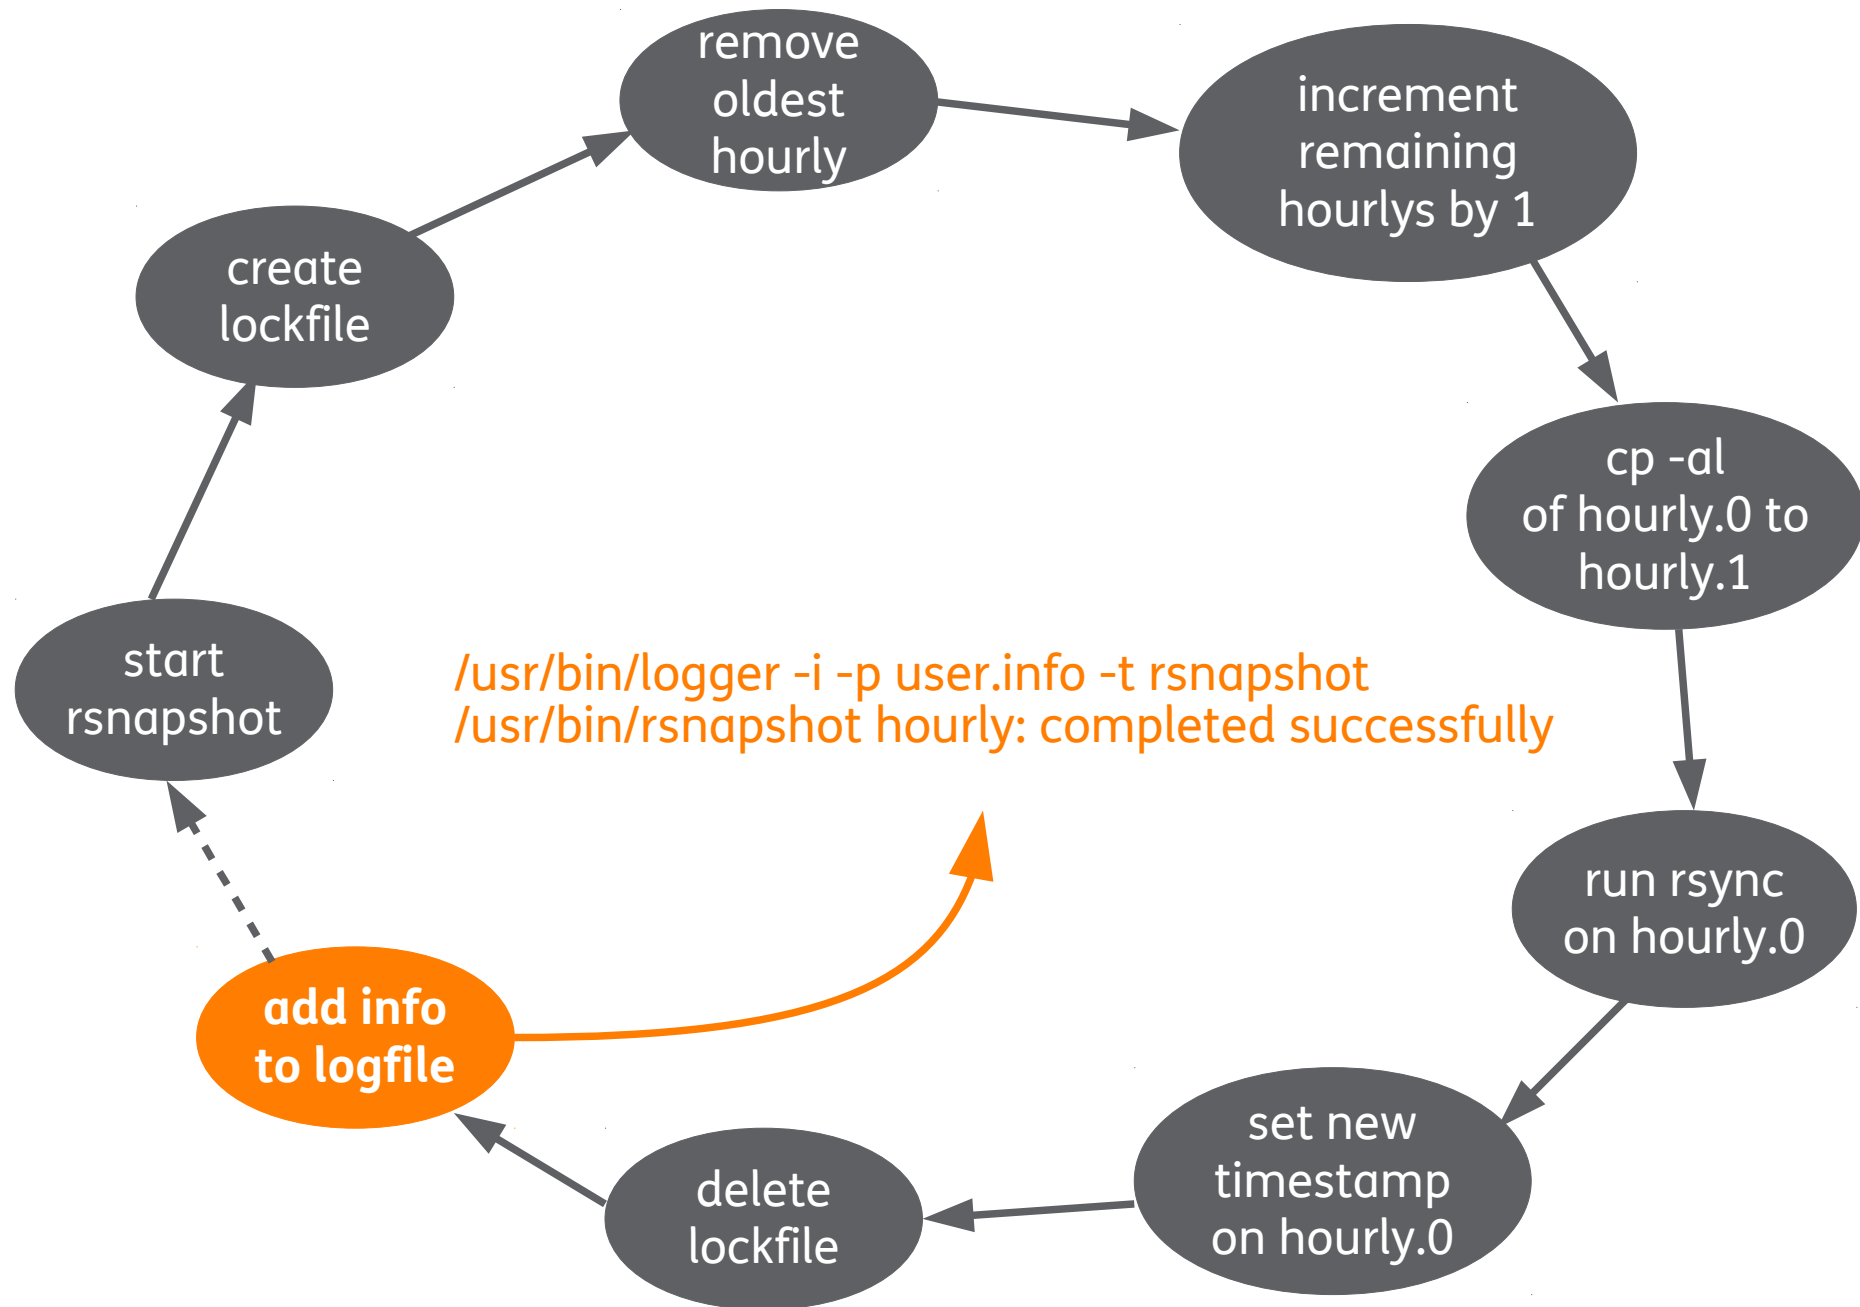












# Limitations and hints

- \_ rsnapshot.conf: elements separated by tabs
- \_ Trailing slash on the end of directory entries
- \_ Rsnapshot directory: chmod 700 & owned by root (if rsnapshot is run by root)
- \_ By default only pull backups
  - <http://linux.die.net/man/1/rsnapshot>
    - A lot more hints in the notes section

# Push backups really not possible?

## — rsnapshot server

- Attach a rsync daemon (with a configuration) to the public SSH Key in `authorized_keys`
- Configuration of the rsync daemon sets the backup directory
- Via `post-xfer exec` a script containing `rsnapshot` will be executed
- If more server should be managed, use a configuration independent of the `/etc/rsnapshot.conf` for the backup scheme

## — How it's executed by the client

- The backup client synchronizes via rsync and SSH the data to the backup server
- Via `authorized_keys` the rsync daemon and the `post-exec` script will be triggered → a snapshot will be created

## — Disadvantage

- Necessity intermediate step through rsync
- Data is stored double, not very space efficient

# Restore?

- `scp`

- Copy via scp the favored folders or single files from the server

- `rsync -avr`

- Sync the folder back to your machine

- `sshfs`

- Mount the folder from the rsnapshot server via ssh

# Monitoring?



— Plugins available


— Icinga/Nagios

— check\_rsnapshot on **github**

— Compatible with TKmon through NRPE

| Hostname                                  | Service    | Status | Last check          | Duration     | Attempt | Graphs |
|---|------------|--------|---------------------|--------------|---------|--------|
| rsnapshot Test-VM                         | net-ping   | OK     | 2015-09-22 14:59:40 | 0d 0h 3m 32s | 1/4     |        |
| PING OK - Packet loss = 0%, RTA = 0.60 ms |            |        |                     |              |         |        |
| rsnapshot Test-VM                         | nrpe-check | OK     | 2015-09-22 14:59:40 | 0d 0h 3m 32s | 1/4     |        |
| Backup OK - Multiple backups OK           |            |        |                     |              |         |        |





Thomas



Christoph



# rdiff-backup

A remote incremental backup  
of all your files could be as easy as

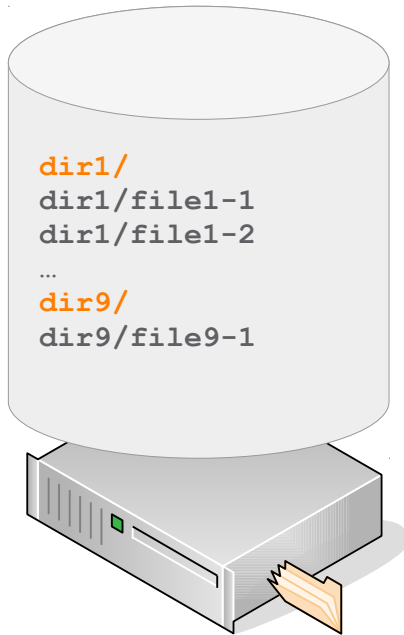
**"rdiff-backup / host.net::/target-dir"**

*source: <http://www.nongnu.org/rdiff-backup/>*

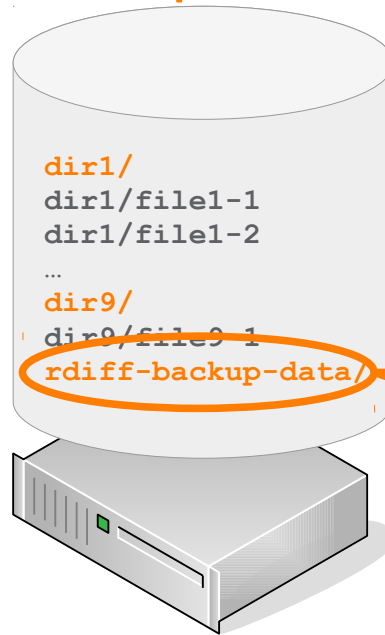


# like rsync but with increments based on reverse diffs

client



backup server



rdiff-backup-data/increments/

```
dir1  
dir1/file1.txt.2015-09-14T15:27:37+02:00.diff.gz  
dir1/file1.txt.2015-09-14T15:28:59+02:00.snapshot.gz  
dir1/file2.txt.2015-09-14T15:27:37+02:00.missing  
dir1.2015-09-14T15:27:37+02:00.dir  
dir1.2015-09-14T15:28:59+02:00.dir  
file1.txt.2015-09-14T15:27:37+02:00.snapshot.gz  
file2.txt.2015-09-14T15:27:37+02:00.missing  
file3.txt.2015-09-14T15:31:12+02:00.missing
```

# rdiff-backup Basics

- written in Python, a little C
- based on librsync
- version 0.1 released 07/2001
- latest version 1.2.8 released 03/2009
- well tested on UNIX
- for MacOS: extended attributes and ACLs
- Debian/Ubuntu: apt-get install rdiff-backup
- CentOS/RHEL: 1.2.8 in EPEL 6 & 7
- SUSE: SLES 11 SP3/4, SLES12, openSUSE



# backup examples



- local → local  
`rdiff-backup /dir1 /backupdir`
- local → remote (via SSH)  
`rdiff-backup /local-dir hostname.net:~/remote-dir`
- remote → local  
`rdiff-backup hostname.net:~/remote-dir /local-dir`
- remote → remote  
`rdiff-backup sourcehost.net:~/source-dir  
desthost.net:~/dest-dir`

# restore examples



— copy from last backup run's mirror

— restore 10 days ago:

```
rdiff-backup -r 10D host.net:~/remote-dir/file /local/file
```

— restore from increment file:

```
rdiff-backup host.net:~/remote-dir/rdiff-backup-  
data/increments/file.2015-09-30T08:21:41-07:00.diff.gz  
/local/file
```

— example MySQL dump

— **uncompressed Thomas-Krenn-Wiki dump: 207 MB**

— **daily diff ~300 – 600 KB**

```
wiki/rdiff-backup-data/increments/backup-mysql# ls -lh | tail  
-rwx----- 1 root root 368K Sep 12 03:25 wiki.sql.2015-09-12T03:28:52+02:00.diff.gz  
-rwx----- 1 root root 353K Sep 13 03:25 wiki.sql.2015-09-13T03:09:17+02:00.diff.gz  
-rwx----- 1 root root 621K Sep 14 03:25 wiki.sql.2015-09-14T03:16:53+02:00.diff.gz  
-rwx----- 1 root root 483K Sep 15 03:25 wiki.sql.2015-09-15T03:22:26+02:00.diff.gz  
-rwx----- 1 root root 557K Sep 16 03:25 wiki.sql.2015-09-16T03:17:51+02:00.diff.gz  
-rwx----- 1 root root 480K Sep 17 03:25 wiki.sql.2015-09-17T03:20:00+02:00.diff.gz
```

# maintenance examples 1/2



## list increments

```
wiki# rdiff-backup -l .
Found 35 increments:
...
  increments.2015-09-17T03:20:00+02:00.dir   Thu Sep 17 03:20:00 2015
  increments.2015-09-18T03:19:38+02:00.dir   Fri Sep 18 03:19:38 2015
  increments.2015-09-19T03:38:32+02:00.dir   Sat Sep 19 03:38:32 2015
  increments.2015-09-20T03:08:53+02:00.dir   Sun Sep 20 03:08:53 2015
  increments.2015-09-21T03:16:06+02:00.dir   Mon Sep 21 03:16:06 2015
Current mirror: Tue Sep 22 03:18:57 2015
```

## remove increments older than 2 months

```
rdiff-backup --remove-older-than 2M host.net::/remote-dir
```

## keep only last 30 backup sessions

```
rdiff-backup --remove-older-than 30B host.net::/remote-dir
```

## exclude / include

```
rdiff-backup --exclude /tmp --exclude /proc --exclude /sys
/ host.net::/remote-dir
```

# maintenance examples 2/2



## — list changed files

```
wiki# rdiff-backup --list-changed-since 5D .
changed backupmysql/wiki.sql
...
changed www/de/wikiDE/images/thumb/1/15/Foto_Werner_Fischer.jpg
new      www/de/wikiDE/images/thumb/1/15/Foto_Werner_Fischer.jpg/100px-Foto_Werner_Fischer.jpg
new      www/de/wikiDE/images/thumb/1/15/Foto_Werner_Fischer.jpg/150px-Foto_Werner_Fischer.jpg
new      www/de/wikiDE/images/thumb/1/15/Foto_Werner_Fischer.jpg/200px-Foto_Werner_Fischer.jpg
...
```

## — when rdiff-backup gets interrupted (killed, network, ...)

### — normally next run automatically fixes that

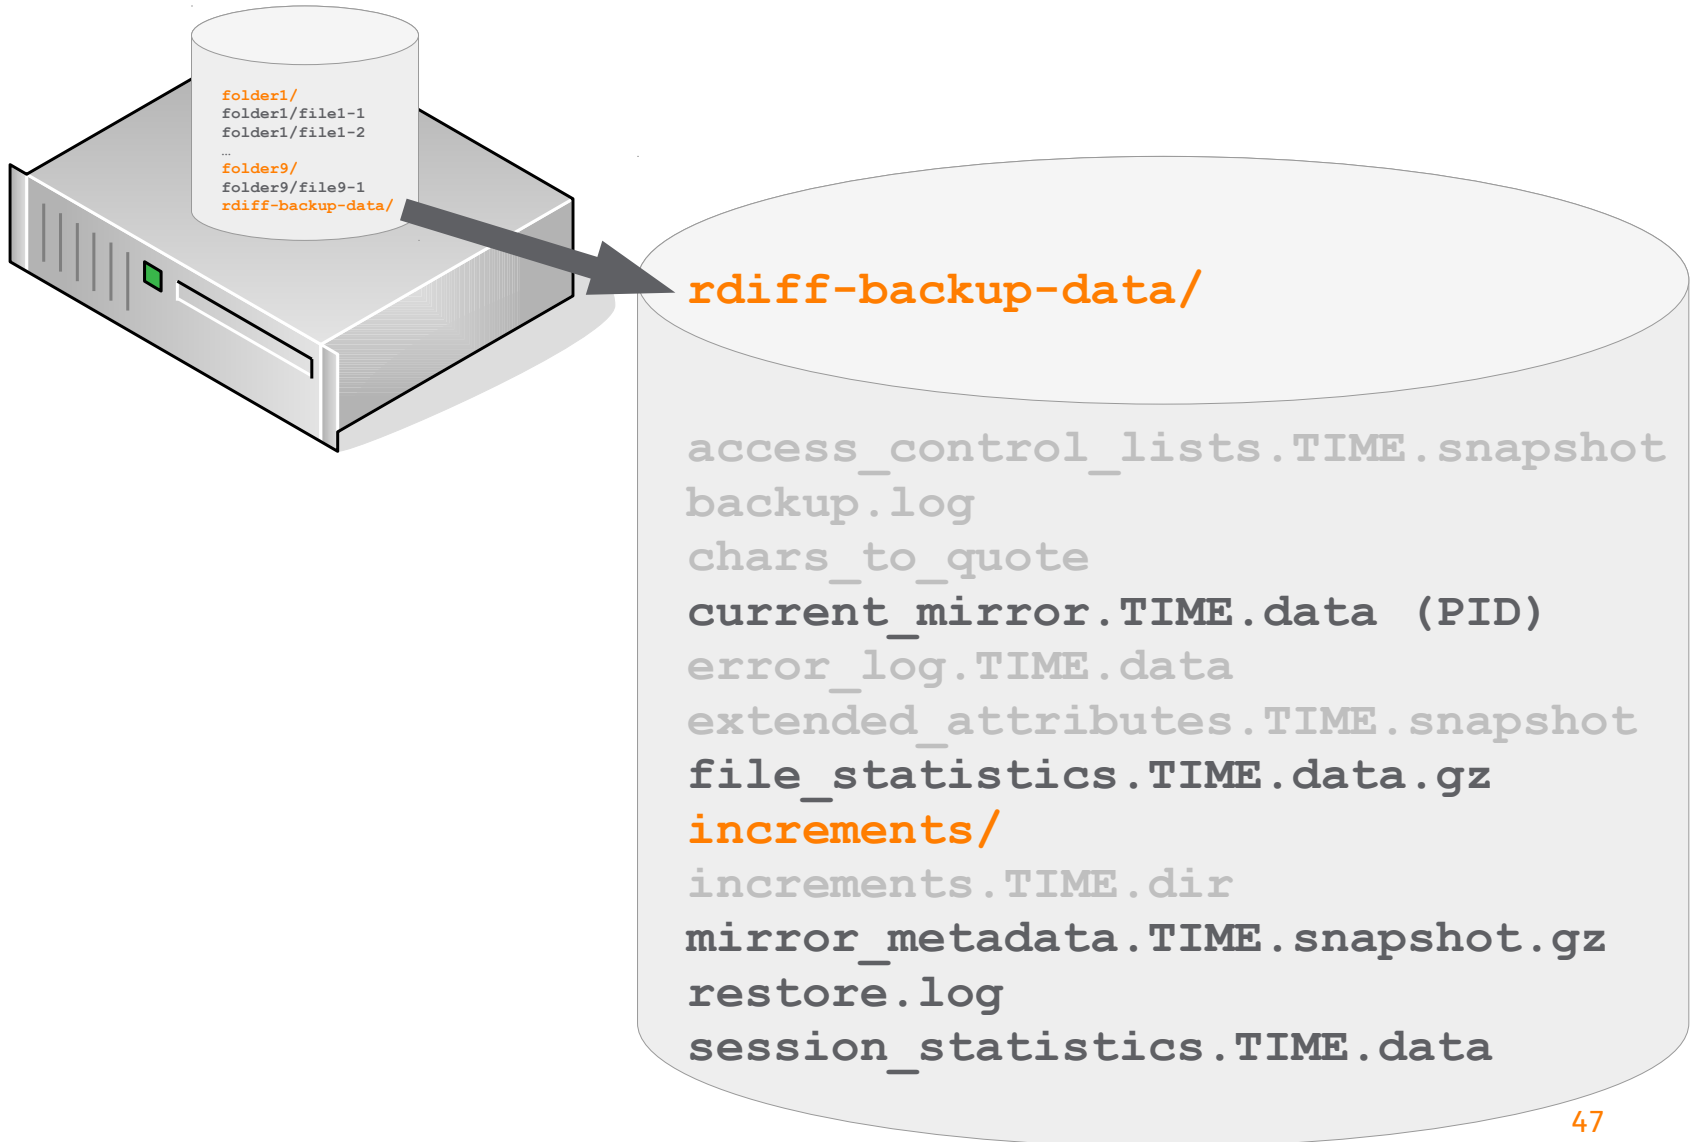
```
Previous backup seems to have failed, regressing
destination now.
```

```
Regressing to Tue Sep 22 12:26:17 2015
```

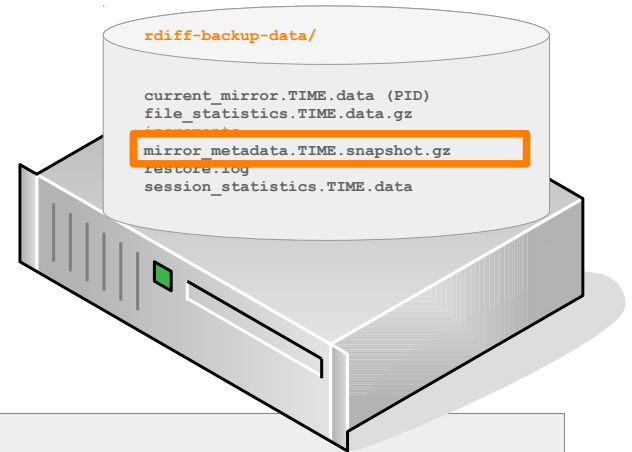
### — if that does not work

- restore from backup's backup ;-)
- try <http://www.timedicer.co.uk/programs/help/rdiff-backup-regress.sh.php>

# Metadata and Logfiles



# Metadata and Logfiles



```
mirror_metadata.TIME.snapshot.gz
```

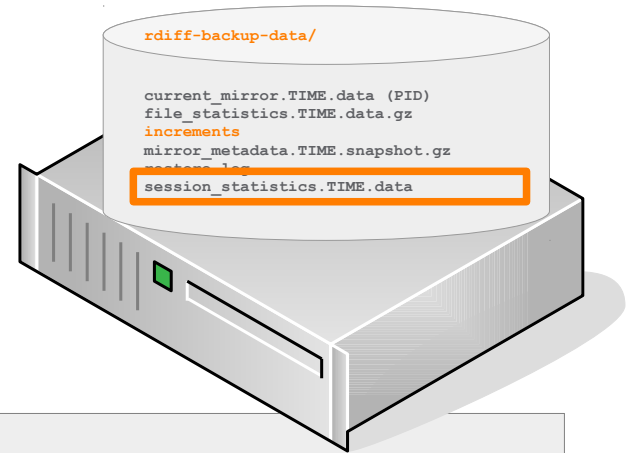
```
File .  
  Type dir  
  ModTime 1425809355  
  Uid 0  
  Uname root  
  Gid 0  
  Gname root  
  Permissions 493
```

```
File bin  
  Type dir  
  ModTime 1418727096  
  Uid 0  
  [...]
```

```
echo 'obase=8;493' | bc  
755
```



# Metadata and Logfiles



```
session_statistics.TIME.data
```

```
StartTime 1441144877.00 (Wed Sep  2 00:01:17 2015)
EndTime 1441147057.65 (Wed Sep  2 00:37:37 2015)
ElapsedTime 2180.65 (36 minutes 20.65 seconds)
SourceFiles 225971
SourceFileSize 21649997461 (20.2 GB)
[...]
NewFiles 1625
NewFileSize 266819836 (254 MB)
[...]
ChangedFiles 3335
[...]
TotalDestinationSizeChange 336682009 (321 MB)
Errors 0
```

*for more details see:  
[tkwiki.cc/rdiff\\_meta](http://tkwiki.cc/rdiff_meta)*

# Do's and Dont's



- separate backup repos for different time frames (e.g. daily, weekly, yearly...)
- exclude files that change often and/or are not essential
  - .deb, .rpm, /var/lib/php5/, ...
- use uncompressed files (e.g. MySQL-Dump) for efficient diff or use "gzip --rsyncable"
- use timestamps for logfiles
  - e.g. dateext for logrotate
- monitor data growth rate



- backup compressed or encrypted files (no useful rdiff)
- very short backup intervals (leads to long restore times)
- fixing metadata by hand without having a backup ;-)
- no regular backup of rdiff-backup archives

# Add Ons



## — FUSE Mount

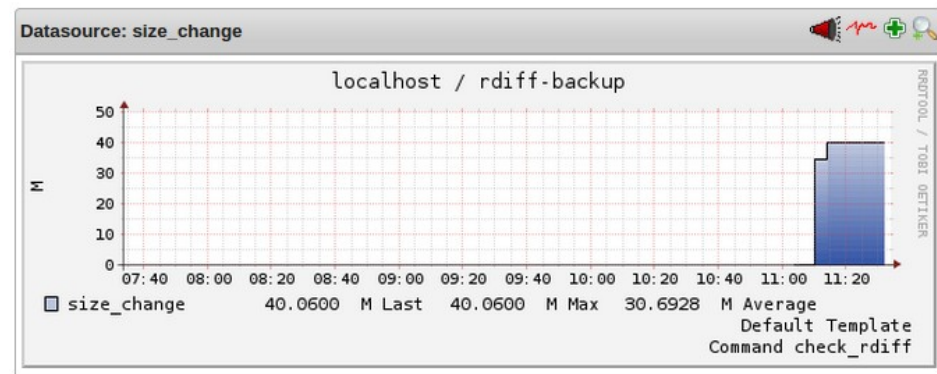
- <http://code.google.com/p/rdiff-backup-fs/>

## — GUI (not official)

- <http://rdiffbackupweb.sourceforge.net/>
- <http://www.patrikdufresne.com/en/rdiffweb/>
- <http://www.nongnu.org/jbackpack/index.html>

## — Icinga Plugin with perfdata patch

- [tkwiki.cc/rdiff\\_monitoring](http://tkwiki.cc/rdiff_monitoring)



so what?

so **what** should I use?  
a comparison!

# Comparison

|                           | <b>rdiff-backup</b>                        | <b>rsnapshot</b>                                     |
|---------------------------|--|--|
| <b>Data transfer</b>      | rsync via librsync                         | directly via rsync                                   |
| <b>Data storing</b>       | older versions stored as increments/deltas | unchanged files saved as hard-links across snapshots |
| <b>Data access</b>        | only newest version directly accessible    | all data directly accessible                         |
| <b>Removal of backups</b> | --remove-older-than                        | retain rule  |
| <b>Efficiency</b>         | highly efficient through compressed deltas | more space needed when files change often            |
| <b>Speed</b>              | slower (CPU intensive)                     | fast through hard links                              |
| <b>Push backups</b>       | yes  | not by default                                       |

# Questions?



Open Source **Backup**  
Conference

**THOMAS**  
**KRENN**<sup>®</sup>

server.hosting.customized.





Questions?  
... no more ???



Open Source **Backup**  
Conference

**TH=MAS**  
**KRENN**<sup>®</sup>

server.hosting.customized.





Thank you and have  
fun with backups and restores ;)



Open Source **Backup**  
Conference

**TH=MAS**  
**KRENN**<sup>®</sup>

server.hosting.customized.



# Images

- <https://www.flickr.com/photos/shardayyy/4793995249/>
- <https://www.flickr.com/photos/jakerust/16846257921>
- <https://www.flickr.com/photos/polsifter/4047982682>
- <https://www.flickr.com/photos/kevandotorg/6229660191>
- <https://www.flickr.com/photos/borispumps/5192678983>
- <https://www.flickr.com/photos/intelfreepress/6722295999>
- [http://www.hdwallpapers.in/beach\\_horizon-wallpapers.html](http://www.hdwallpapers.in/beach_horizon-wallpapers.html)
- photo of Weissensee on slide 39 taken by Christoph Mitasch