UPPSALA
UNIVERSITET

v1.0, April 14, 2016
Author: Jerker Nyberg von Below

SOP – Rsync backup to Btrfs snapshots

1 Purpose

This is documentation for the incremental file based backups via Rsync to Btrfs snapshots stored on disk. This is a faster and cheaper alternative than the university IBM Spectrum Protect, but not as reliable (unless backups are taken to multiple nodes). The backup is also only supported during office hours. This backup system using Rsync and Btrfs is replacing our current system using the same principles but ZFS instead of Btrfs.

2 Preparation

1. The backup server is running CentOS 7 with the built in Btrfs. Any Linux-distribution with a recent Btrfs works fine. Btrfs is mounted at a specific point at for example /data4.
2. Make it possible to login from the backup server to the backup client over SSH automatically by using keys.
3. Set up crontab for each client and the file systems to be backed up.
   ```
   1 21 * * * /usr/local/bin/backup.sh /data4 server.bmc.uu.se / /home
   ```
4. Set the email-address in backup-script to the one to get error messages.

3 First time backup

The first time the script is run the Btrfs subvolume **/data4/server.bmc.uu.se** is created. For every file system the individual subvolumes are created **/data4/server.bmc.uu.se/backup-:** and **/data4/server.bmc.uu.se/backup-:home.**

4 Completed backup

If all subvolumes has been created and Rsync complete the backup without errors (except for files disappearing or changing during backup) then a new read only Btrfs snapshot will be created named with the current time stamp. For example: **/data4/server.bmc.uu.se/snapshot-:-20160412-154147**

5 Performance

The backup server is connected to 1 Gbit/s Ethernet. Currently the performance of backing up large files is limited by CPU usage of Rsync and Btrfs together with the bandwidth. Current backup server hardware have CPU-performance to 1-2 Gbit/s Ethernet but that is only useful for the first time backup filling up the backup server. Performance of small files are limited by the disk IO-utilization traversing the directory tree. Modern hardware with more and faster CPU-cores should in theory scale up closer to 10 Gbit/s than 1 Gbit/s.

6 Errors

Errors are reported back via mail. This is mainly Rsync getting error messages or perhaps file system space running low above 95% usage.

UPPSALA
UNIVERSITET

7 Status webpage

The webserver on the backup server displays a status page with what backups are currently running or when they were last completed. Backups that are late or incomplete are marked such.

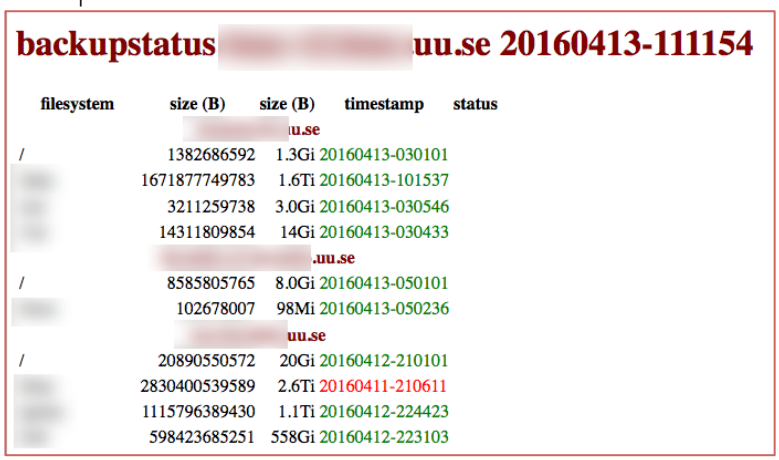

backupstatus [redacted] uu.se 20160413-111154

filesystem	size (B)	size (B)	timestamp	status
	[redacted] u.se			
/	1382686592	1.3Gi	20160413-030101	
[redacted]	1671877749783	1.6Ti	20160413-101537	
[redacted]	3211259738	3.0Gi	20160413-030546	
[redacted]	14311809854	14Gi	20160413-030433	
	[redacted] .uu.se			
/	8585805765	8.0Gi	20160413-050101	
[redacted]	102678007	98Mi	20160413-050236	
	[redacted] uu.se			
/	20890550572	20Gi	20160412-210101	
[redacted]	2830400539589	2.6Ti	20160411-210611	
[redacted]	1115796389430	1.1Ti	20160412-224423	
[redacted]	598423685251	558Gi	20160412-223103	

8 Restoring backups

All data for a particular file system is stored as they were in the read-only snapshot. This snaophot may for example be named **/data4/server.bmc.uu.se/snapshot-:-20160412-154147**. To restore from the backup Rsync the files back again.

9 Limitations and features

- The backup client **cannot destroy already backed up files** on the backup server. The only way already backed up content can be destroyed is when the read-only snapshots are removed on the backup server.
- Currently old snapshots have to be deleted manually - there is **no automatic expunge** of old snapshots.
- Since this system is using Rsync and not Btrfs snapshots for transferring of files, **all files are identified by file name**. This makes renaming large directories very heavy on disk usage. All files are transferred and backed up again. No deduplication of file content is being done.
- All files are stored using their uid and gid. But since the backup-server itself is not connected to any shared user-source (USER-AD) then this is just showing up as numbers.
- There is a locking file in **/var/lock/backup.server.bmc.uu.se.:data.lock.** If the lock is there that host and file system is skipped. This means that multiple backups for a particular host can be run in parallel. But if the whole backup is interrupted then the lock file may block further backups. This is shown in the status web page.
- Btrfs is compressing files via zlib.
- The computer to be backed up must run a SSH server with Rsync - in practice this is usually Mac OS X and Linux.
- The computer to be backed up must be contacted from the backup server using a host name. In practice this means the computer to be backed up must be a server that is always on.

- This is not a good way of backing up large files that are contantly changing - for example virtual machine images. Instead let the virtual machine be backed up.

10 References

Btrfs is an open source copy-on-write B-tree file system for Linux, originally designed at Oracle Corporation by Chris Mason in 2007 based on research at IBM by Ohad Rodeh in 2006.
https://btrfs.wiki.kernel.org/

Rsync is an open source utility that provides fast incremental file transfer and is currently being maintained by Wayne Davison. Rsync was created by Andrew Tridgell and Paul Mackerras in 1996.
https://rsync.samba.org/

ZFS is a combined file system and logical volume manager designed by Sun Microsystems and introduced in OpenSolaris in 2005. ZFS on Linux was created in 2008 by Lawrence Livermore National Laboratory (LLNL) and the POSIX layer released in 2011.
http://zfsonlinux.org/