

Re: Backups: rsync, software RAID, other strategies?

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2004-03/1064.html>

From: Dany Nativel (dany_list_at_natzo.com)

Date: 03/08/04

Date: Sun, 07 Mar 2004 23:22:28 -0800

To: Bob Johnson <bob89@bobj.org>

Hi Bob,

I use the following configuration on my file server.

It's a small cube based on a low power mini-ITX EPIA 5000 motherboard (fanless) and running two 120GB HDDs.

I looked at RAID but it doesn't help solving one of the potential issue ... me, user removing files that are not supposed to be removed so incremental backup is a plus. This baby is powered by FreeBSD 5.2.1 (used to be Debian).

DISK 0 (live)

128M ad0s1a /

512M ad0s1b swap

128M ad0s1d /var

200M ad0s1e /tmp

3200MB ad0s1f /usr

110GB ad0s2d /data

DISK 1 (Backup)

ad2s1d 128M /backup/os/root

ad2s0b 512M swap

ad2s1e 128M /backup/os/var

ad2s1f 200M /backup/os/tmp

ad2s1g 3199M /backup/os/usr

ad2s2d 108G /backup/data/backup

I use 3 different programs :

- Unison (<http://www.cis.upenn.edu/~bcpierce/unison/>): 2-way synchronization using rsync/ssh, multi-platform graphical interface. I can have the same files on my file server, laptop running win2k as well as my desktop running BSD. Very convenient especially with laptops when you can't be connected all the time. Very fast too (only transmit diffs)
- rsync (man rsync) : typical rsync that will mirror the source to the destination
- rdiff-backup (<http://rdiff-backup.stanford.edu/index.html>): it's based

freebsd-questions: Re: Backups: rsync, software RAID, other strategies?

on rsync but you get the advantage of incremental backups so you can restore from a specific date. You can also purge the backup by removing old stuff. No fancy file format, just .gz for the diffs.

Here is how I use those tools :

```
/data/current/user0_live (DISK 0) <-----> UNISON : 2-way  
synchronization with laptop/desktop  
/data/current/user0_incremental (DISK 0) <-----> RDIFF-BACKUP :  
incremental backup of user0_live using RDIFF system  
/backup/data/backup/user0_incremental (DISK1) <-----> RSYNC : quick  
mirror of the already incremental backup
```

```
/backup/data/backup/pictures (DISK1) <-----> RDIFF-BACKUP :  
in this case, rdiff-backup between drive0 and drive1 (no incremental on  
disk0)
```

```
/backup/os/root (DISK1) <-----> DUMP : 1:1 copy of the live root fs  
/backup/os/tmp (DISK1)  
/backup/os/var (DISK1)
```

```
/backup/os/usr (DISK1) <-----> RSYNC : (with -delete option) for  
a quick mirror of current /usr
```

PS: for user0, there are two copies of the data on disk0, 1 live synchronized with Unison and another one which is an incremental of the first one. For less critical data (like /data/current/pictures) I use rdiff-backup between disk0 and disk1. In that case I would lose incremental backups if disk1 fails.

I've simulated a crisis situation by removing the drive0 and swapping it with drive1. It worked (except for those entries in fstab referring to disk1).

I have a cron job taking of the different backups at night.

```
#!/bin/sh
```

```
# Duplicate /  
# erase slive before mirroring, any other way?  
umount /backup/os/root  
newfs /dev/ad2s1a  
mount /backup/os/root  
# dump with 'live filesystem' option  
dump -0 -L -f - /dev/ad0s1a | (cd /backup/os/root ; restore -r -v -f -)
```

```
# Duplicate /var  
umount /backup/os/var  
newfs /dev/ad2s1d  
mount /backup/os/var  
dump -0 -L -f - /dev/ad0s1d | (cd /backup/os/var ; restore -r -v -f -)
```

freebsd-questions: Re: Backups: rsync, software RAID, other strategies?

```
# Duplicate /tmp (probably a waste of time)
umount /backup/os/tmp
newfs /dev/ad2s1e
mount /backup/os/tmp
dump -0 -L -f - /dev/ad0s1e | (cd /backup/os/tmp ; restore -r -v -f -)
```

```
# incremental backup of the ./pictures directory on the second drive
rdiff-backup /data/current/pictures /backup/data/backup/pictures
```

```
# First, incremental of the user0_live dir on the same drive then rsync
on the second drive
rdiff-backup /data/current/user0_live /data/current/user0_incremental
rsync -a --delete /data/current/user0_incremental/
/backup/data/backup/user0_incremental
```

The only I don't like is the NEWFS command. Is there a cleaner way to do this dump ?

I use this configuration in a non-critical installation (my house) but it has been serving its purpose so far.

Dany

PS: On the rdiff-backup webpage there is a link to another tool called duplicity (<http://rdiff-backup.stanford.edu/duplicity.html>). You can do remote backup but in that case the image can be stored on a remote FTP server and encrypted with GPG... sweet if you're planning to use your ISP's disk space for backups!

Bob Johnson wrote:

>A bunch of related questions:

>

>I'm setting up a small mail and file server. The mail server part will
>be Courier, while the file server part will primarily be used via NFS
>and Samba to store backups of my desktop and laptop computers.

>

>The system has a pair of WD1600JB 160 GB ATA 100 drives in it, both on a
>single Promise PDC20268 UDMA100 controller, but each on a separate
>channel (i.e. both are masters with no slaves). My plan is to use one
>of the drives as a backup for the other. I want to use a backup method
>that creates a mirror of the working drive so that if it fails, I can
>simply mount the backup in place of the working drive, and get back in
>business. The operating system will (probably) not be on either of
>these drives, they will only host /home where mail and backup files
>will reside.

>

>I've tentatively decided to synchronize the mirror to the working drive
>with rsync run once a night (perhaps more often later). This risks
>losing up to a day of mail, but that's probably not the end of the
>world. The reasoning I used was that if I use software RAID, an
>unexpected power failure during a large write operation (yes, it will

Re: Backups: rsync, software RAID, other strategies?

freebsd-questions: Re: Backups: rsync, software RAID, other strategies?

>be on a UPS) could corrupt both drives. Running rsync once a night
>would reduce the risk of a failure that damaged both drives, since
>their write activity would not be so strongly correlated.
>
>Is my fear of losing both drives in a software RAID (mirrored drives)
>configuration a reasonable one? Or is that not going to happen?
>
>If I use rsync with -delete to maintain a mirror of the primary drive,
>what happens when the primary drive fails? Is there a scenario that
>causes rsync to duplicate all the missing data on the primary drive by
>deleting it from the mirror drive (I've heard of this happening to
>someone, but I believe he was using a homegrown perl script rather than
>rsync).
>
>Is Courier compatible with this scheme? Or does it care about inode
>numbers or some such thing that will make the backup copy useless?
>
>Is there any chance it would make sense to use the Coda file system for
>this?
>
>Do any of the answers change if the mail server ends up on a remote
>system, but I still want the maildirs backed up on the file server?
>
>Any other suggestions that don't involve buying new hardware? An IDE
>RAID controller would be nice, but buying one isn't on my list of
>things to do. But if I DID break down and buy a new controller card,
>what should it be?
>
>Thanks,
>
>- Bob
>

>freebsd-questions@freebsd.org mailing list
><http://lists.freebsd.org/mailman/listinfo/freebsd-questions>
>To unsubscribe, send any mail to "freebsd-questions-unsubscribe@freebsd.org"
>
>

freebsd-questions@freebsd.org mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-questions>
To unsubscribe, send any mail to "freebsd-questions-unsubscribe@freebsd.org"