

Precious precious... precious... data...

**by Mary Gardiner,
Sydney Linux User Group, Nov 28 2008**

**Notes and slides for this talk available at
<http://users.puzzling.org/users/mary/Presentations/SLUG2008/>**

**These slides are available under the Creative Commons Attribution-
Sharealike 2.5 Australia licence:**

<http://creativecommons.org/licenses/by-sa/2.5/au/>

See slide 47 for attribution of the images used in these slides.




Confession time

“I don't have backups”







Express backups

Run home right now...

(grab a big external hard drive on the way, format it with ext3)

Install rdiff-backup

Run the magical command

Magical command

```
sudo rdiff-backup --exclude-other-filesystems ::/ ::/media/disk
```

```
sudo rdiff-backup --remove-older-than 60D ::/media/disk
```


Recovery

```
sudo rdiff-backup --list-changed-  
since 1D ::/mnt/disk
```


Recovery

```
deleted Maildir/spam/cur/12276.flay:2,  
deleted Maildir/spam/cur/12277.flay:2,
```

Recovery

```
sudo rdiff-backup -r 1D  
/media/disk/Maildir/spam /tmp/spam
```







Disaster: Accidental deletion

```
# rm -rf /
```


Recovery from accidental deletion

Will work: At least one quite recent older copy of your data.

Emergency:

```
dd if=/dev/drive of=/mnt/otherdrive/drive.image
```

```
photorec /mnt/otherdrive/drive.image
```

Won't work: RAID

Media failure

ASM P/N 13N6704 4200RPM 40.0GB

FRU P/N 13N6705

Hitachi TM

www.hgst.com



D33082

02-4845

N155 E182115 T
LES

MODEL: IC25N040ATMR04-0 4200RPM
5V 1.0A DC 40GB ATA/IDE
MADE IN THAILAND BY Hitachi Global
Storage Technologies(Thailand).Ltd. TD
WARRANTY VOID IF ANY LABEL/
SCREW IS REMOVED OR BROKEN 06DEC03

IBM P/N: 92P6592 (HITACHI P/N: 13G1812)
MLC: H69555

(16383CYL, 16HEADS, 63SEC/T)
78.140.160 LBA'S



11S92P6592Z1Z8PZJRVZMH

Don't shock/
push cover



DO NOT COVER
THIS HOLE =>



Recovery from media failure

Will work: At least one *very* recent older copy of your data *on another drive*.

Will likely work: RAID

Emergency:

```
ddrescue /dev/old_drive /mnt/other/old_drive.image
```

```
fsck /mnt/other/old_drive.image
```

```
mount -o loop /mnt/other/old_drive.image  
/mnt/somewhere
```

Software failure

```
1120 # file kind...
1121 base_pair = contents_pair(self.base_tree)
1122 other_pair = contents_pair(self.other_tree)
1123 if base_pair == other_pair:
1124     # OTHER introduced no changes
1125     return "unmodified"
1126 this_pair = contents_pair(self.this_tree)
1127 if this_pair == other_pair:
1128     # THIS and OTHER introduced the same changes
1129     return "unmodified"
1130 else:
1131     trans_id = self.tt.trans_id_file_id(file_id)
1132     if this_pair == base_pair:
1133         # only OTHER introduced changes
1134         if file_id in self.this_tree:
1135             # Remove any existing contents
1136             self.tt.delete_contents(trans_id)
1137         if file_id in self.other_tree:
1138             # OTHER changed the file
1139             create_from_tree(self.tt, trans_id,
1140                             self.other_tree, file_id)
1141             if file_id not in self.this_tree:
1142                 self.tt.version_file(file_id, trans_id)
1143             return "modified"
1144         elif file_id in self.this_tree.inventory:
1145             # OTHER deleted the file
1146             self.tt.unversion_file(trans_id)
1147             return "deleted"
1148     #BOTH THIS and OTHER introduced changes; scalar conflict
```

Recovery from software failure

Will work: At least one quite recent older copy of your data *possibly on a separate operating system.*

Won't work: RAID

Massive local failure



Recovery from massive local failure

Will work: At least one quite recent copy of your data *as far away as possible*.

Won't work: Most other things.


Provider failure

Recovery from provider failure

Will work: A local copy of your data and meta-data, with normal backup procedures applied.

Won't work: Begging.

Won't work: Facebook



The important bits

1. Start making backups somehow, now
- 2... *on separate media.*
3. Easy and often (automatic is best, press of a button next best)
4. Backup everything
5. Test your media
6. Have local backups of your remote data
7. Have remote backups of your local data



Niggly bits

1. Dump your databases for backup
2. Don't (properly) backup your backups



Off-site backup

- Sneaker-net
- S3
- Dreamhost Personal Backup

Off-site backup: costs

- 10GB:
 - Sneaker-net: \$104 (actually, 160GB)
 - S3: US\$1 to put data in, US\$1.50 per month
 - Dreamhost: US\$5.95 a month (up to 50GB)
- 190GB:
 - Sneaker-net: still \$104
 - S3: US\$19 to put data in, US\$28.50 per month
 - Dreamhost: US\$19.95 per month

Off-site backup: risks





02

Useful tricks

- Automatic backups
 - Cron
 - Udev
 - Network manager

Advanced fancy stuff

```
$ du -sh ~/.thumbnails  
359M /home/mary/.thumbnails
```


Advanced fancy stuff

- You don't need to backup:
 - secondary, synced copies of photos, music, mail
 - caches (/var/cache/apt/archive, ~/.thumbnails)
 - re-installable packages (/usr/bin)
- ... unless you do.

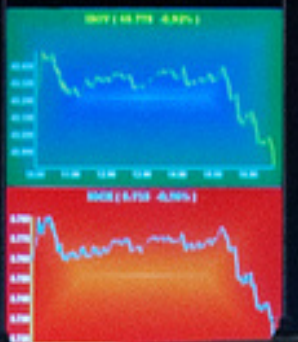
Ill-advised fancy stuff

If you're really short of room, /home, /etc and /var is a good opt-in.

If you're really *really* short of room, try:

- Your documents (if reliably located)
- E-mail (IM logs etc)
- Photos

PREV	OSTR3	TTM34	CRE3B	ESTR4	CRE3D	ESTR5	CRE3E	THLP3	CRE3F	THLP4	CRE3G	THLP5	GOLL4	GOLL4	US188	CRE3D	POST3	US188	US188	US188	US188	US188	US188	US188	US188	US188	US188	US188	US188	US188	US188		
11.36	17.50	62.80	17.50	1.32	17.50	23.83	17.50	23.83	17.50	66.25	17.50	56.49	75.50	113.17	96.31	17.50	31.06	17.50	96.31	0.24	1.49	25.83	0.22										
NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	NEG ATIU0	ULTIMO	OSC	
23 I800I3	849.98	0.8	5 I800P3	282.04	0.3	235 I800D4	24.89	0.2	48																								
17 COPV0	24.10	1.8	582 SCAR3	18.79	0.8	89 SZP04	4.86	1.2	116																								
39 OHLE0	45.79	0.8	24 SHT03	28.40	0.7	290 ITRN04	52.80	0.8	748																								
30 PDCR3	31.15	1.7	82 SBRB3	73.20	2.3	272 UOLL4	18.25	0.0	61																								
14 POSI3	14.10	2.7	33 TALE3	17.58	1.4	80	NIVEL - I																										
61 IPRG0	30.70	0.8	20 TCRB3	11.25	1.3	114 ATIU0	ULTIMO	OSC	NEG ATIU0	149.00	0.0	12																					
203 PSS00	30.81	0.0	728	NIVEL - II																													
4 RCH13	65.65	0.3	58 ATIU0	ULTIMO	OSC	NEG ATIU0	48.80	1.5	240																								
685 RENT3	19.80	0.5	74 ICLSC0	38.71	1.3	107 I800C4	11.21	1.4	322																								
74 RWR3	20.82	0.3	305 HELPL6	105.00	0.0	327 I800P4	43.38	1.1	855																								
75 ROM10	0.88	1.1	13 I800L4	59.39	1.0	582 I800H5	17.00	1.1	996																								
36 RSI03	16.97	0.1	107 I800E4	28.88	0.0	693 I800T04	12.47	1.2	1117																								
184.00 RTE	27.80	2.3	303 PON04	59.30	1.0	82 I800T4	21.75	0.0	337																								



Local backups

- Consider:
 - your (g|e)mail
 - your calendars
 - your address book
 - your photos
 - your blog
 - your social network
 - ... all your meta-data.

Remember to incrementally back these up.

my site's global
navigation is
pulled in / rendered
using
javascript
includes
:-)

favorite
in Flash
accessible.

We worked with a non-profit
that specifically works with
the disabled and they talked
us into accessible frames...

When I'm feeling really
lazy, I'll put in

 tags to put in
extra space between
elements

My first website was
made in MS Word
... Save As HTML

Confession time

- When did I start backing up?
- What made me start?
- ... what didn't make me start?
- What am I still not backing up?
- Do I have remote backups?

Credits

- Slide 2: “[The Lure of Gold](#)” by [bogenfraud@flickr](#), CC-BY-SA
- Slide 4: “[fright night couple](#)” by [lizjones112@flickr](#), CC-BY
- Slides 5: “[The Fright](#)” by [piddy77@flickr](#), CC-BY
- Slide 14, 25, 34: “[House fire from across the lake](#)” by [wickedchimp@flickr](#), CC-BY
- Slide 19: “[Hard Drive 016](#)” by [jon_a_ross@flickr](#), CC-BY
- Slide 22: Screenshot of [Bazaar](#) code by Mary Gardiner, code copyright Canonical Ltd and available under the GPL
- Slide 28, 37, 43: “[That was supposed to be going up, wasn't it?](#)” by [rednuht@flickr](#), CC-BY
- Slide 37: “[Pirate Riley. Aaarrhh Me Hearties!](#)” by [peasap@flickr](#), CC-BY
- Slide 38: “[Tricks](#)” by [cole24_@flickr](#), CC-BY-SA
- Slide 45: “[Web Design Confession Booth \(viii\)](#)” by [allaboutgeorge@flickr](#), CC-BY
- All clip art from <http://openclipart.org/>, public domain

The important bits

1. Start making backups somehow, now
- 2... *on separate media.*
3. Easy and often (automatic is best, press of a button next best)
4. Backup everything
5. Test your media
6. Have local backups of your remote data
7. Have remote backups of your local data