

Venti Backups to Blu-Ray Discs

Geoff Collyer

Bell Laboratories
Murray Hill, New Jersey 07974

Overview

As a precaution against multiple disks in our Plan 9 *venti* server's RAID array failing at about the same time, or other catastrophic failure, we record *venti* arenas, after they are sealed, onto dual-layer Blu-Ray discs (BDs). One could use other large optical discs instead.

The advertised capacity of a dual-layer BD is 50GB, but those aren't even disk-manufacturer's (decimal) gigabytes, which would give a capacity of 50×10^9 bytes, which is roughly equivalent to 46.6 gigabytes, as the term is used by everyone but disk manufacturers. In the case of BDs, even that is an exaggeration, with the actual capacity being closer to 48.44×10^9 bytes, so the claimed capacity should be read as '50 BD-gigabytes', where a *BD-gigabyte* is 968,800,338 bytes. The default *venti* arena size is 512MB, but for some reason our server is configured with 1GB arenas, so we could fit 46 of them on a BD. To leave a little extra room for lead-in, lead-out, inter-track gaps, lossless-linking and the like, we record 45 arenas per BD.

The scripts used and records kept are in `/sys/lib/backup` and pertain to the file server on which they are stored. You will probably want to edit `funcs` to set default file server and Blu-ray device, at minimum.

Creating and Updating Backups

We start `cdfs` after inserting a disc (virgin or partially-written), then use `venti/rdarena` to copy the next sealed but not backed-up arena to the start of the unwritten portion of the BD by writing to `/mnt/cd/wd/x`. We don't fixate data BDs, as that seems to tickle a bug, perhaps in our Sony Blu-Ray burners, that only records the first track in the disc index upon fixation.

One can find out which arenas are sealed and which are open by viewing <http://venti-server/index>. We also periodically print the most recent *fossil* dump scores and save the paper with the discs.

All of this is packaged up as scripts in `/sys/lib/backup`, notably `backup`. A typical invocation would be just

```
cd /sys/lib/backup
backup
```

which will guide its invoker as to which discs to insert into a Blu-ray burner and when. This works for a first full backup and for subsequent incremental backups of just the newly-sealed arenas. There is provision for burning multiple backup sets, the default being `set1`. To burn a second set, we would

```
mkdir /sys/lib/backup/set2
backup -s set2
```

`backup` and related scripts keep track of which BD is current and how many tracks are

recorded, and which arenas have been dumped to BD.

We have recorded two sets of BD backups, one of which will go to Antwerp to seed their *venti* server and also serve as off-site backup for Murray Hill.

Restoring from Backups

After a disaster, or when setting up a new *venti* server from BD backups, the first step is to get Plan 9 running on the new *venti* server. This could be done by installing into a smallish (2GB) *fossil* partition from a Plan 9 installation CD, if necessary. One would then initialise the new disk partitions per *venti-fmt(8)* and read all the arenas on all the BDs into the new *venti* by running *venti/wrarena* once per BD track (arena).

To add the contents of a backup BD to a (possibly fresh) *venti* store, shut down the *venti* server, format the arenas partition, then run

```
restore first-arena-number
```

after inserting a BD into the Blu-ray drive. Repeat this for each BD in the backup set. When all the arenas have been restored, it will be necessary to build a new *venti* index, the usual steps being to run *checkarenas*, *fmtisect*, *fmtbloom*, *fmtindex*, and *buildindex -b*, all from *venti-fmt(8)*. Then the *venti* server may be restarted.

Once the *venti* storage has been restored, a new *fossil* partition (perhaps the existing one or another one) can be initialised from the last *fossil* dump score corresponding to the last arena on BD (see *fossil/flfmt* in *fossil(4)*):

```
fossil/flfmt -v c388...32b5 /dev/sdC0/fossil
```