

System deployment using bittorrent

Steven Shiau, Thomas Tsai, Ceasar Sun

clonezilla.org

Q2, 2017



Outline

- Introduction to Clonezilla
 - Features
- System deployment using bittorrent
 - Challenge
 - Implementation
- Demo
- Q&A



Outline

- Introduction to Clonezilla
 - Features
- System deployment using bittorrent
 - Challenge
 - Implementation
- Demo
- Q&A



System imaging and cloning - backup



You want to crash!!!
I show you how to crash!!!


image source: maggiesfarm.anotherdotcom.com
www.compsults.com, and jervisdabreo.com

Massive system deployment



About us

- Developers of the free software DRBL, Clonezilla and more...
- Steven is also the maintainer of GParted live CD
- From Taiwan, working for the NPO NCHC (National Center for High-Performance Computing)



Taiwan image source: wikipedia.org

What is Clonezilla?

- A partition and disk imaging/cloning utility similar to True image® or Ghost®
- GPL license
- A bare metal recovery tool for



*1



*2



*3



*4

VMFS

VMware
ESX/ESXi

*5



MINIX

*6



*Logo source: (1) Larry Ewing, Simon Budig and Anja Gerwinski, (2) Apple, (3) Microsoft, (4) Marshall Kirk McKusick, (5) VMWare (6) Distrowatch.com

TAIWAN

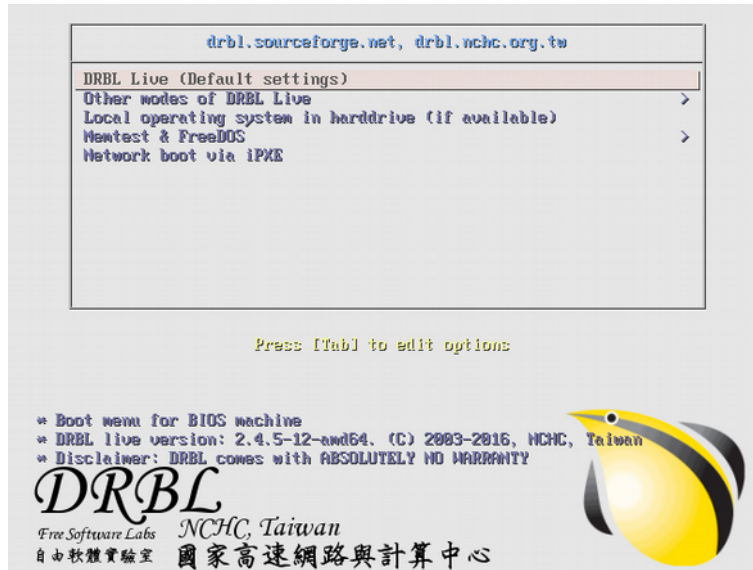
www.nchc.org.tw



Clonezilla Features

- Free ([GPL](#)) Software
- File systems supported:
 - [Ext2/3/4](#), [ReiserFS](#), [Reiser4](#), [XFS](#), [JFS](#), [HFS+](#), [BrfFS](#), [F2fs](#), [UFS](#), [Minix](#), [VMFS](#), [F2FS](#), [NILFS2](#), [FAT](#) and [NTFS](#)
 - Supports [LVM2](#)
 - Support some [hardware RAID](#) chips (by kernel)
- [Smart copying](#) for supported filesystem. For unsupported file systems sector-to-sector copying is done via [dd](#).
- Boot loader : [syslinux](#), [grub 1/2](#) ; [MBR](#) and hidden data (if exist)
- [Serial console](#)
- Unattended mode
- One image restoring to multiple local devices
- [Multicast](#) and [bittorrent](#) (in beta) deployment are supported in Clonezilla Server Edition (SE) and lite server
- The image format is transparent, open and flexible

DRBL live, i.e. Clonezilla Server Edition



Clonezilla Live




Press [F6] to edit options

Automatic boot in 25 seconds...

- * Boot menu for BIOS machine
- * Clonezilla live version: 2.4.5-20-amd64. (C) 2003-2016, NCHC, Taiwan
- * Disclaimer: Clonezilla comes with ABSOLUTELY NO WARRANTY

Clonezilla *Free Software Labs*
National Center for High-Performance Computing
Taiwan



Developers

- Steven Shiau
- K. L. Huang
- Ceasar Sun
- Jazz Wang
- Thomas Tsai
- Jean-Francois Nifenecker
- Louie Chen
- Nagappan Alagappan



Language file contributors

- English (en_US): Dylan Pack
- Catalan (ca_ES): René Mériou and Innocent De Marchi
- German (de_DE): Michael Vinzenz
- Hungarian (hu_HU): Greg Marki
- Spanish (es_ES): Juan Ramón Martínez and Alex Ibáñez López.
- French (fr_FR): Jean-Francois Nifenecker and Jean Francois Martinez.
- Italian (it_IT): Gianfranco Gentili.
- Japanese (ja_JP): Akira Yoshiyama and Annie Wei.
- Brazilian Portuguese (pt_BR): Marcos Pereira da Silva Cruz.
- Russian (ru_RU): Anton Pryadko and Igor Melnikov.
- Slovak (sk_SK): Ondrej Dzivy Balucha
- Turkish (tr_TR): Ömer YILDIZ
- Simplified Chinese (zh_CN): Zhiqiang Zhang and Liang Qi.
- Traditional Chinese (zh_TW): T. C. Lin.


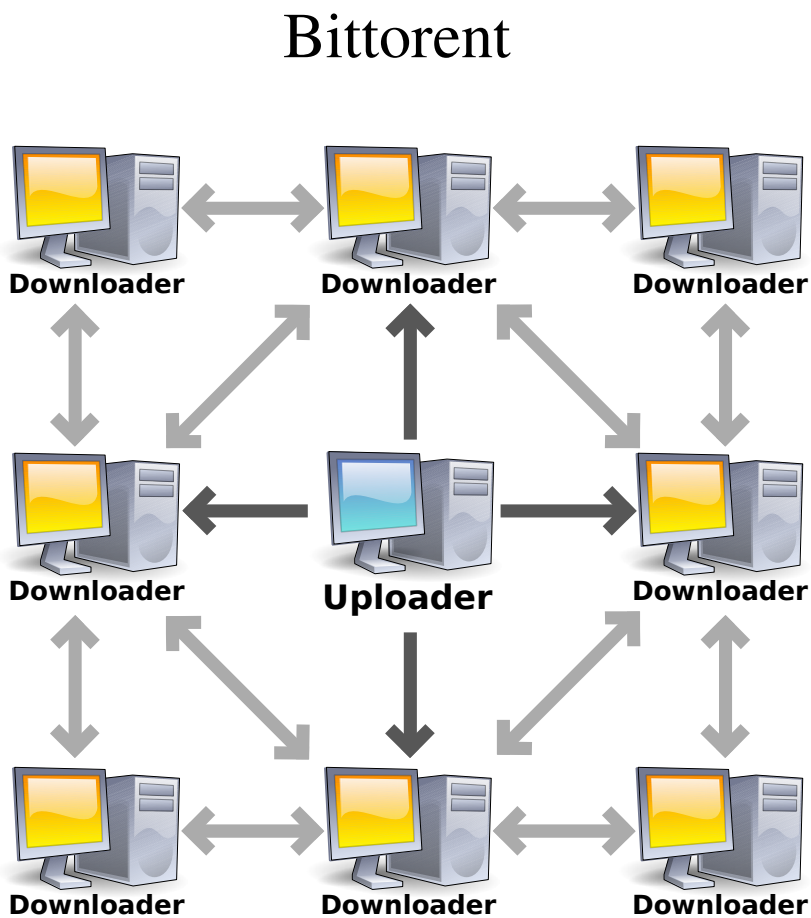


Outline

- Introduction to Clonezilla
 - Features
- System deployment using bittorrent
 - Challenge
 - Implementation
- Demo
- Q&A



Bittorrent vs. multicast



Scale

Image source: By Scott Martin. <https://commons.wikimedia.org/w/index.php?curid=31181606>



Bittorrent

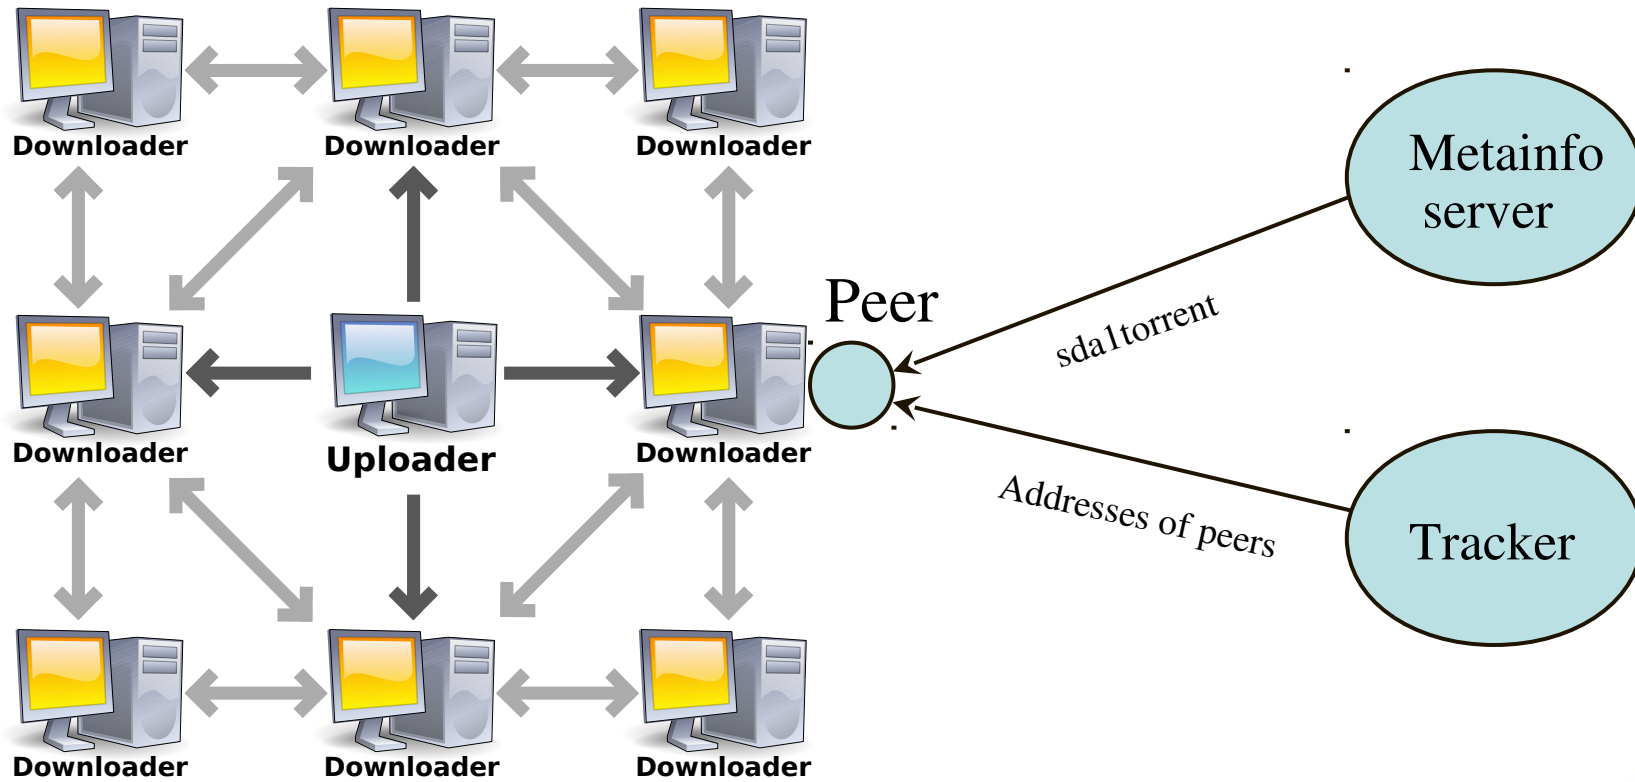


Image source: By Scott Martin. <https://commons.wikimedia.org/w/index.php?curid=31181606>

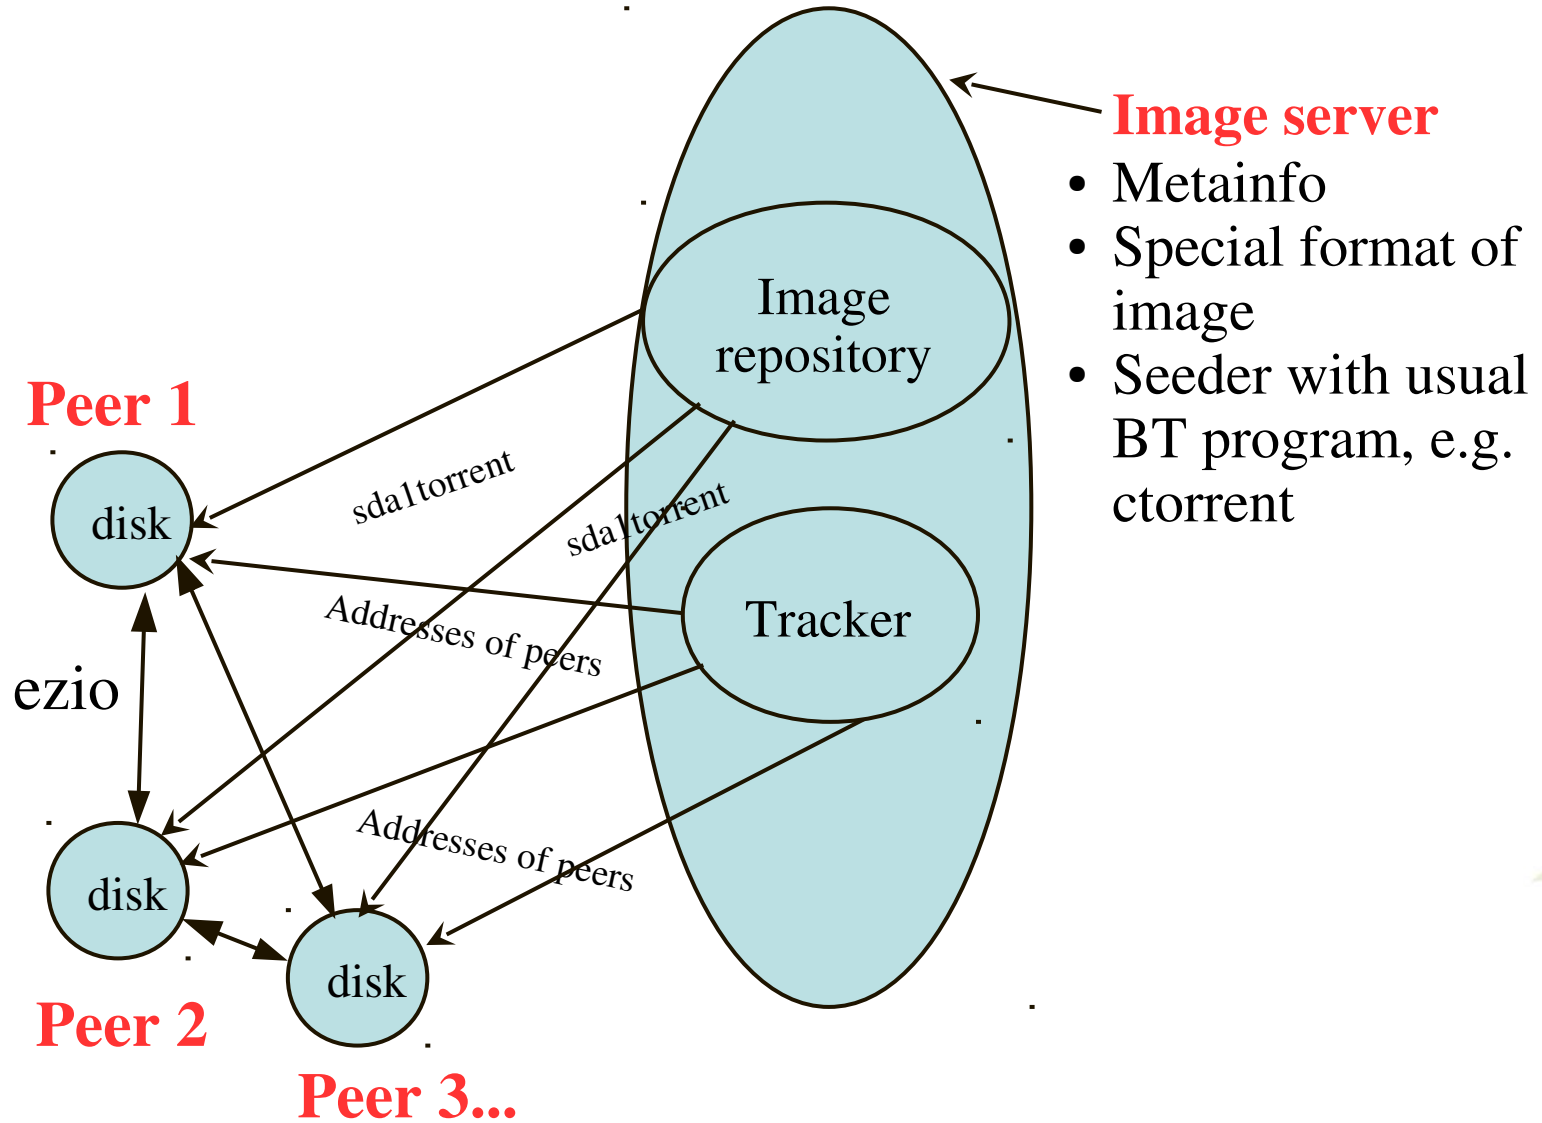


System Deployment via Bittorrent- Challenge and Solution

- Challenge
 - For peer to peer file sharing, a temp storage space is required
 - For bare metal recovery, the available storage space is RAM
 - Because the existing disk is the destination disk to be overwritten
 - However, RAM disk is normally not big enough for system deployment
 - Disk image: ~ GB to 100 GB or more
- Solution
 - By Date Huang, etc.
 - Separate the seeder and downloader
 - Seeder
 - Special format image file
 - Downloader
 - Directly write image to disk and share the blocks from disk



Bittorrent in Clonezilla





System Deployment via Bittorrent

- The image has to be in **special format for bittorrent** purpose. Hence extra space on hard drive is required.

```
- root@debian:~# ls /home/partimag/btzone/stretch-x64-20170626/sda1  
00000000000000000000000000000000 00000000000000000000000001f0fa000  
0000000000000000000000000440c0000 00000000000000000000000000100000  
00000000000000000000000001f120000 000000000000000000000000044100000  
00000000000000000000000000200000 00000000000000000000000001f123000  
000000000000000000000000044120000 00000000000000000000000000300000  
00000000000000000000000001f125000 000000000000000000000000044140000  
0000000000000000000000000400000 00000000000000000000000001f175000  
000000000000000000000000044180000 00000000000000000000000000500000  
00000000000000000000000001f179000 0000000000000000000000000441af000  
00000000000000000000000000600000 00000000000000000000000001f193000  
0000000000000000000000000441e0000 00000000000000000000000000700000  
00000000000000000000000001f196000 000000000000000000000000044200000
```

...

- Each file stores a section of used continuous file system blocks. The file name denotes its offset on the partition.



System Deployment via Bittorrent

- Extra packages are required:
 - **Ezio, ocs-bttrack, ctorrent, mktorrent**
 - Ezio is a blocks deployment program with bittorrent, and is developed by Date Huang (tjjh89017), Ching-Hsuan Yen (mangokingTW), and Pu Lee (leepupu):
<https://github.com/tjjh89017/ezio>
 - Ocs-bttrack is a bittorrent tracker, and is developed by Clonezilla team based on BitTornado
<https://github.com/stevenshiau/ocs-bttrack>
- **Partclone >= v0.3.5**
- **More tests and improvements** are required before official release



Screenshots

```

Created a new partition 1 of type 'Linux' and of size 6 GiB.
/dev/sda2: Created a new partition 2 of type 'Extended' and of size 2 GiB.
/dev/sda3: Created a new partition 5 of type 'Linux swap / Solaris' and of size 2 GiB.
/dev/sda6:
New situation:

Device Boot      Start         End      Sectors  Size Id Type
/dev/sda1 *            2048    12584959    12582912     6G 83 Linux
/dev/sda2            12587006    16775167     4188162     2G  5 Extended
/dev/sda5            12587008    16775167     4188160     2G 82 Linux swap / Solaris

The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
This was done by: LC_ALL=C sfdisk --force /dev/sda < /home/partimag/xenial-x64-desktop-20161210/sda-
pt.sf 2>&1 | tee -a /var/log/clonezilla.log
Checking the integrity of partition table in the disk /dev/sda...
*****
Informing the OS of partition table changes.... done!
*****
The first partition of disk /dev/sda starts at 2048.
Restoring the hidden data between MBR (1st sector, i.e. 512 bytes) and 1st partition, which might be
useful for some recovery tool, by:
dd if=/home/partimag/xenial-x64-desktop-20161210/sda-hidden-data-after-mbr of=/dev/sda seek=1 bs=512
count=2047
2047+0 records in
2047+0 records out
1048064 bytes (1.0 MB, 1.0 MiB) copied, 0.0342036 s, 30.6 MB/s
*****
Restoring partition /dev/sda1...
*****
Clean filesystem header in device /dev/sda1...
Running: ezio-static /home/partimag/btzone/xenial-x64-desktop-20161210/sda1.torrent /dev/sda1
Start downloading
[P: 27%] [D: 7.4 GB/min] [DT: 9 secs] [U: 0 GB/min] [UT: 0 secs] _

```

Clonezilla - Opensource Clone System (OCS) | Mode: restore-disk
Choose the mode to restore client disk

```

multicast multicast restore
broadcast broadcast restore
bittorrent bittorrent restore
unicast unicast restore

```

<Ok>

<Cancel>

```

/dev/sda3: Created a new partition 5 of type 'Linux swap / Solaris' and of size 2 GiB.
/dev/sda6:
New situation:

Device Boot      Start         End      Sectors  Size Id Type
/dev/sda1 *            2048    12584959    12582912     6G 83 Linux
/dev/sda2            12587006    16775167     4188162     2G  5 Extended
/dev/sda5            12587008    16775167     4188160     2G 82 Linux swap / Solaris

The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
This was done by: LC_ALL=C sfdisk --force /dev/sda < /home/partimag/xenial-x64-desktop-20161210/sda-
pt.sf 2>&1 | tee -a /var/log/clonezilla.log
Checking the integrity of partition table in the disk /dev/sda...
*****
Informing the OS of partition table changes.... done!
*****
The first partition of disk /dev/sda starts at 2048.
Restoring the hidden data between MBR (1st sector, i.e. 512 bytes) and 1st partition, which might be
useful for some recovery tool, by:
dd if=/home/partimag/xenial-x64-desktop-20161210/sda-hidden-data-after-mbr of=/dev/sda seek=1 bs=512
count=2047
2047+0 records in
2047+0 records out
1048064 bytes (1.0 MB, 1.0 MiB) copied, 0.0342036 s, 30.6 MB/s
*****
Restoring partition /dev/sda1...
*****
Clean filesystem header in device /dev/sda1...
Running: ezio-static /home/partimag/btzone/xenial-x64-desktop-20161210/sda1.torrent /dev/sda1
Start downloading
[P: 100%] [D: 11 GB/min] [DT: 24 secs] [U: 0 GB/min] [UT: 0 secs]
Start high-performance seeding
[U: 1.1 MB/min] [T: 68 secs]

```


Tests

- Bittorrent (BT) deployment testing results at NCHC computer classroom
 - 1 PC as the server, 38 clients. Image size is 6 GB (uncompressed)
 - To deploy the image to 38 clients
 - 360 secs with multicast mechanism
 - 750 secs with BT mechanism
 - 250 secs (180+70) if deploying 3 clients first (180 secs), once it's done, then deploying the rest of 35 (70 secs).




```
Checking the integrity of partition table in the disk /dev/sda...
[Informing the OS of partition table changes...] 39.302001] sda: sda1 sda2 * s
das
[ 39.534444] sda: sda1 sda2 * sda5
..... done!
.....
The first partition of disk /dev/sda starts at 2048.
Restoring the hidden data between MBR (1st sector, i.e. 512 bytes) and 1st parti
tion, which might be useful for some recovery tool. by:
dd if=/home/partimag/xenial-x64-desktop-20161210/sda-hidden-data-after-mbr of=/d
ev/sda seek=1 bs=512 count=2047
2047+0 records in
2047+0 records out
1048064 bytes (1.0 MB, 1.0 MiB) copied, 0.0668129 s, 15.7 MB/s
.....
Restoring partition /dev/sda1...
.....
Clean filesystem header in device /dev/sda1...
Running: ezio-static /home/partimag/btzone/xenial-x64-desktop-20161210/sda1.torr
ent /dev/sda1
Start downloading
[P: 74%] [O: 6.3 GB/min] [OT: 40 secs] [U: 0.00022 GB/min] [UT: 0 secs]
```



Demo

- Clonezilla live experimental 2.5.3-1exp
 - Not ready for public testing yet until early Aug
- Deploy Debian Stretch by Bittorrent to 2 clients



More info on site

- Workshop
 - Clonezilla workshop
 - Wednesday 5/Jul, 09:40-10:40 (Workshop, Room A 013 Côté Serveur)
- Booth
 - Clonezilla booth at the Village of associations



Reference

- Clonezilla: <http://clonezilla.org>
- DRBL: <http://drbl.org>



Questions ?

Great!



??????

